Cinderella:
an
Online Filter
for the
COMPASS Experiment

Diploma Thesis
by
Thiemo Nagel

TECHNISCHE
UNIVERSITÄT
MÜNCHEN

19th January 2005

Contents

List of Figures

List of Tables

Chapter 1

Introduction

"The good into the pot,
the bad into the crop."
— Cinderella, in [Gri84]

With the knowledge of physics of the smallest particles continuously expanding, new experiments set off to explore increasingly smaller effects, requiring higher statistics to achieve significant measurements. Consequently, higher beam intensities are necessary, interactions need to be detected with higher frequency and total data volume increases accordingly. This description certainly holds true for the COMPASS experiment at CERN.

One of its major challenges is to satisfy its large demands in bandwidth and storage. Thus already in the COMPASS proposal [Bau96] it was foreseen to establish an online filter to reduce the enormous data rate *prior to tape writing*. Yet, in the years of build-up of the experiment, 2001 with technical run and 2002 with the first physics run, the main focus was to achieve stable production data taking. Not all detectors had been commissioned at the beginning, so that bandwidth and storage limitations were not highly pressing at first.

Consequently, it was not until November of 2002 that the work on an online filter had begun. Delivering the functionality of a 2^{nd}-level trigger, the online filter quickly picked up the name Cinderella.[1] In 2003 the development of Cinderella was carried out partly by Roland Kuhn [Ku04] and partly by the author, so that in January of 2004, when the work was continued in the form of a diploma thesis, Cinderella already had been tested during parts of the data taking of 2003.

With the volume of recorded data having risen from 2002 to 2003 and having been projected to rise again sharply in 2004, it seemed to be well indicated to push the online filter

[1] The name is borrowed from the homonymous fairy tale in which the girl Cinderella is forced to pick lentils from ashes by her evil step-mother and is being magically helped in this tedious task by *all the birds beneath the sky* (cf. [Gri15] for a German version or [Gri84] for its English translation). Yet it has to be acknowledged that this allegory had already been used in [Boh92] where the acronym "**C**ondition **IN**terpreter for **D**ata and **E**vent **RE**duction by **L**ow **L**evel **A**nalysis" is given, which however was chosen *not* to be adopted for the online filter of the COMPASS experiment.

to production use for data taking of 2004. (A survey of recorded data volume is shown in Table 1.1.)

This diploma thesis describes the development and the first productive use of the online filter Cinderella. In chapter 2 first the COMPASS experiment is introduced to the reader, giving a summary of the experimental apparatus and its physics goals.

Due to hard constraints on available processing time it has been necessary to avoid existing decoding libraries and create high-performance decoding modules custom-tailored to the job. A great part of this effort was spent on the silicon micro-strip detectors and is detailed in chapter 3.

General concepts of data reduction as well as general facets of their implementation at COMPASS are explained in chapter 4. Aspects dedicated to either the muon or the hadron programme are elucidated in separate chapters 5 and 6, respectively.

Combined for both types of beam, chapter 7 illustrates the interfaces between Cinderella and the experiment. Details of design and implementation of the software are illuminated in 8. And finally, the conclusion is drawn in chapter 9.

As the members of the COMPASS collaboration probably constitute the largest audience of this text, some remarks here are enclosed for them:

For a general and thorough explanation of Cinderella as a whole, chapters 4–7 and 9 may be referred to. In contrast to that, chapters 3 and 8 rather cater to specific interests.

Almost every member of the collaboration will eventually get in touch with Cinderella somehow.[2] Especially for them chapter 7 was written, explaining the vital facts about Cinderella from a practical point of view and in a condensed manner. Therefore: *If you only have time to read one chapter, then read chapter 7!*

Year	Raw Data	Data on Tape	Ratio
2002	195 TByte	142 TByte	73%
2003	243 TByte	189 TByte	78%
2004	618 TByte	406 TByte	66%

Table 1.1: The volume of data acquired by the DAQ is compared with the storage space effectively allocated on tape, which is lower since the tape drives of the storage facility perform data compression. In 2004, additional data reduction by Cinderella came into effect, saving over 70 TByte of tape memory. (Without Cinderella, the ratio of tape storage over acquired data would have been 77% in 2004.)

[2]Though some may feel rather that Cinderella is getting in touch with *them*.

Chapter 2

The COMPASS Experiment

The **CO**mmon **M**uon and **P**roton **A**pparatus for **S**tructure and **S**pectroscopy COMPASS is a high-rate, fixed target experiment at the Super Proton Synchrotron SPS of the European Organisation for Nuclear Research CERN[1]. It is utilising longitudinally polarised muon and unpolarised hadron beams of energies well above 100 GeV/c, to pursue a variety of physics goals. At the dawn of the LHC[2]-era, COMPASS (numbered NA-58) currently is the largest experiment at CERN and a precursor to LHC experiments in a sense that many technologies originally developed for them already are employed at COMPASS.

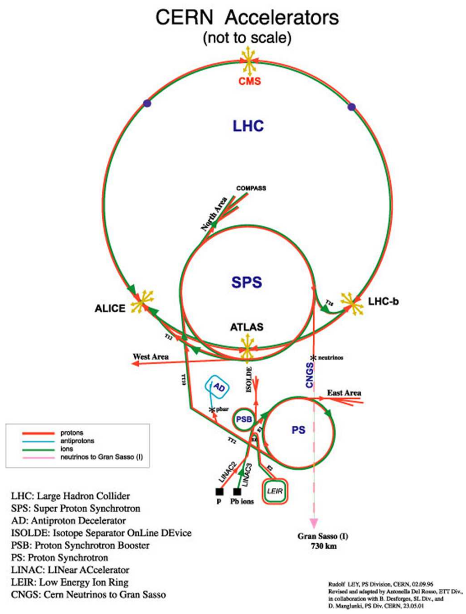

Figure 2.1: Overview of CERN accelerators (not to scale)

Initiated with the proposal of 1996, the COMPASS collaboration now[3] encompasses 273 members from 25 institutes in 11 countries on 2 continents. Physics data have been taken in the years 2002–2004. After a technical break in 2005 the experiment will continue with an upgraded spectrometer in 2006.

Before describing details of the setup of the experiment, the targeted physics goals are outlined.

[1]**C**onseil **E**uropéen pour la **R**echerche **N**ucléaire, founded 1954 by 12 member countries, see [CER] for more information

[2]The **L**arge **H**adron **C**ollider, currently under construction, after its start-up in 2007 will advance to previously unreached energies and is expected to have a lasting effect on the comprehension of fundamental physical theory.

[3]as of 9th of December 2004

2.1 Physics Goals

As COMPASS has emerged from combination of two competing collaborations, CHEOPS[4] and HMC[5], that share similar demands to the experimental apparatus, diverse physics topics are addressed, focussing around structure and spectroscopy of hadrons.

2.1.1 Physics with Muon Beam

The main objective of the muon programme is to gain insight in the spin-dependant structure of the nucleon. After three years of data taking (2002–2004), the analysis chain for data taken with muon beam is well established and for many areas of interest first results have been released or are expected to be released very soon.

Gluon polarisation $\Delta G/G$: While the spin of the nucleon has been known to be $\frac{1}{2}$ for a long time, its composition still is a mystery. A calculations from a parton model yields:

$$\frac{1}{2} = \frac{1}{2}\Delta\Sigma + \Delta G + L_z^q + L_z^g \tag{2.1}$$

From these four possible contributions, quark spin, gluon spin, and quark and gluon orbital angular momentum (in order of eq. 2.1), up to now only $\Delta\Sigma$ has been measured with reasonable accuracy. To the surprising result of $\Delta\Sigma = 0.27 \pm 0.13$ (from EMC, SMC and SLAC in [Lam00]), existing models were a bad fit.

A major goal for the muon programme is to shed some light on this issue by contributing a measurement of $\Delta G/G$. This is achieved by analysing the Photon-Gluon Fusion (PGF) process in Deeply Inelastic Scattering (DIS) through the open charm, high p_T pair or high p_T single channels. First results already have been presented or will be presented within the near future in [Ber04, Sch04a, Ku05].

Spin dependent parton distributions $\Delta q(x)$: The flavour decomposition of the quark contribution $\Delta\Sigma$ to the nucleon spin, is subject to this field of research, accessing the spin distribution functions $\Delta q(x)$ of the individual flavours. For strange quarks, the distribution may be extracted from studies of longitudinal polarisation of Λ and $\bar{\Lambda}$ [Ale04, Pes04].

Transverse spin dependant parton distributions $\Delta_T q(x)$: The transverse spin distribution $\Delta_T q(x)$ is the least well-known of the three independent parton distribution functions that together in leading order fully describe the structure of the nucleon in polarised deeply-inelastic scattering. Due to its odd chirality, $\Delta_T q(x)$ may only be measured in conjunction with another chiral-odd distribution. A path of analysis currently pursued is the measurement of $\Delta_T q(x)$ together with the Collins fragmentation function $\Delta D_a^h(z, p_T^h)$ [Pag04].

[4] **CH**arm **E**xperiment with **O**mni **P**urpose **S**etup, cf. letter of intent [Ale95]
[5] **H**adron **M**uon **C**ollaboration, cf. letter of intent [Nap95]

Transverse Λ^0 polarisation: The spontaneous polarisation of the Λ hyperon perpendicular to the production plane in hadronic collisions yet needs to be fully understood. COMPASS is well suited to the study of this process though the topic was not included in the original proposal. First results of transverse polarisation of Λ^0 generated from quasi-real photo-production are given in [Wie04, Fri04b].

Generalised Parton Distributions (GPD): A topic not included in the initial proposal, the determination of GPDs is planned for a later phase of COMPASS. Processes conceivable for this measurement are Hard Exclusive Meson Production (HEMP) or Deeply Virtual Compton Scattering (DVCS). While the former may accessed through the ρ^0 channel already with the current experimental setup, the latter requires extensive upgrades to the spectrometer [dHo02]. As a precursor to GPD extraction, several elements of the ρ^0 spin density matrices already have been determined [Ney04], also yielding indication for weak violation of the S-Channel Helicity Conservation (SCHC).

2.1.2 Physics with Hadron Beam

Primakoff scattering: The Primakoff reaction [Pri51] (radiative pion scattering) may be utilised for measurement of the electric and magnetic polarisabilities $\bar{\alpha}_\pi$ and $\bar{\beta}_\pi$ of the pion, permitting a test of predictions of chiral pertubation theory. First results from COMPASS soon should arise from analysis of the data sample taken in the hadron pilot run of 2004 [Col05], yielding statistics ten times higher than that of the existing data from the Serpukhov experiment [Ant85].

Using the $\sim$4.5% fraction of kaons of the hadron beam together with beam flavour tagging, the same measurement may be performed for the kaon polarisabilities, which currently is unprecedented.

Exotic States: Hadronic states outside the well-known meson $q\bar{q}$ and baryon qqq systems are labelled *exotic*. While the only restriction imposed on bound states by QCD[6] is that of colour neutrality, current evidence as to the existence of exotics is contradictory.

There may exist different classes of exotics, which may be characterised by their constituents. Bound states of more than three quarks may be referred to as *unconventional hadrons*. Several experiments[7] have reported indications for states that could be *tetraquarks* with quark content $qq\bar{q}\bar{q}$ and *pentaquarks* with quark content $qqqq\bar{q}$. *Hexaquarks* $qqqqqq$ or $qqq\bar{q}\bar{q}\bar{q}$ also have been predicted. In an analysis motivated by recent observation of the Ξ^{--} pentaquark by the NA-49 collaboration [Alt04], COMPASS could not establish the generation of the Ξ^{--} pentaquark from

[6] **Q**uantum **C**hromo**D**ynamics, a gauge field theory describing the strong interaction of quarks with a concept of colour charge, establishing gluons as exchange vector bosons, is a constituting part of the Standard Model.

[7] The LEPS collaboration in [Nak03] first published detection of a pentaquark and the BELLE collaboration in [Cho03] was the first to publish discovery of a tetraquark.

photo-production. Rather than that, from the data taken in 2002 and 2003 an upper limit to the production cross section was determined in [Bro04] and [DeM04].

The second class of exotics—called *hybrids*—is characterised by a combination of quarks and gluons in a bound state. Due to their composition of quarks (fermions) and gluons (bosons), hybrids may carry quantum numbers that are not possible in particles consisting purely of quarks, a fact that is of tremendous help for separating hybrids from conventional hadrons. The most simple hybrid may be imagined as resulting from excitation of the gluonic string binding quark and anti-quark in a meson.

Utilising diffractive production, the wave function of the $J^{PC} = 1^{-+}$ state which is forbidden for conventional hadrons and thus a candidate for an exotic state—first described in [Ada98]—may be explored in the $\eta\pi^-$ system [Dor02]. A small diffractive data sample could be recorded in 2004.

The third class of exotics is those only consisting of gluons, dubbed *glue-balls*. As with hybrids, a promising approach to their detection is selection of quantum numbers that are forbidden for conventional hadrons to avoid mixing of conventional and exotic states.

Doubly charmed baryons: Doubly charmed baryons Ξ_{cc}^+, Ξ_{cc}^{++} and Ω_{cc}^+ provide a venue to the understanding of baryonic structure that current experiments just have begun to reach out to. Indications for the former two particles have been observed by SELEX [Moi03], with a lifetime of Ξ_{cc}^+ well below theoretical estimates.

From the very beginning it has been envisioned to tackle this topic in COMPASS, too. However as the production cross section for doubly charmed baryons in COMPASS is expected to be very small, such an endeavour requires extremely high rates and very efficient particle identification. As a consequence, measurements in this area require upgrades of various parts of the COMPASS spectrometer (DAQ and target) and are planned for after 2006.

2.2 The Experimental Setup

The COMPASS spectrometer is built as a succession of two similar spectrometer stages, featuring two Spectrometer Magnets (SM1 and SM2) of increasing strength (cf. Figure 2.2). Each of the two stages, called Large Angle Spectrometer (LAS) and Small Angle Spectrometer (SAS), comprises detectors for tracking, particle ID and calorimetry. The *absorbing* detectors of the LAS (calorimeters and muon filter) have a central hole matching the acceptance of the following second stage, which with its stronger magnet and larger lever arm provides measurement of better accuracy for particles with high momentum. The angular coverage of the LAS is ±180 mrad behind SM1 with up to 1 Tm integrated field strength (with large gap) whereas the SAS provides ±25 mrad of coverage behind up to 5.2 Tm integrated field strength delivered by SM2. The gap of SM1 may be narrowed from 1.72 m to 0.82 m by removing some of the modular yoke pieces, increasing integrated field strength while at the same time reducing fringe fields. The benefit of the

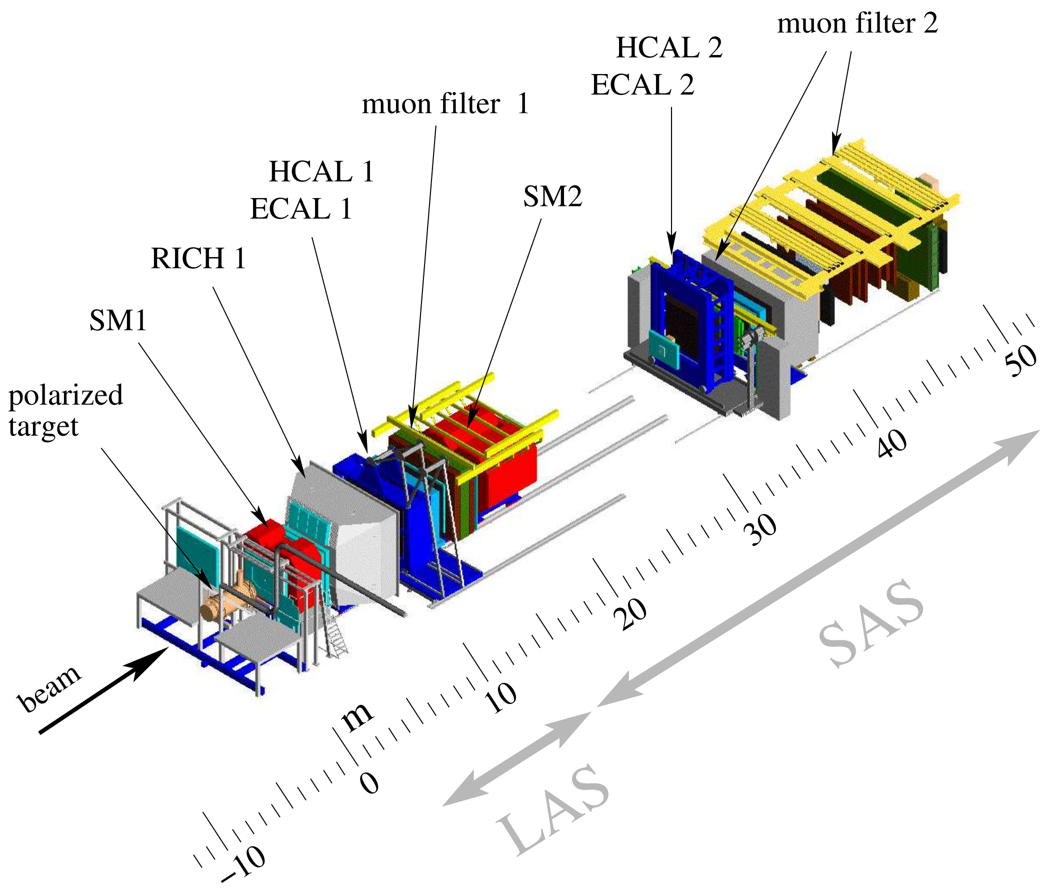

Figure 2.2: Overview of the COMPASS Spectrometer

staggered setup is large acceptance in x_{Bj} and Q^2, covering nearly six orders of magnitude in either.

2.2.1 The Beam

COMPASS uses secondary or tertiary beams which are generated from the primary beam of approx. $1.2 \cdot 10^{13}$ protons per cycle delivered to the T6 production target[8] by SPS with energies up to 450 GeV/c. The SPS follows an injection-acceleration-extraction cycle of currently 16.8 s duration, divided into 12 s for acceleration and 4.8 s for extraction, the latter period being called *spill* or *burst*. De-bunching is provided by turning off the acceleration structures a short time before extraction.

After the T6 target, which constitutes the beginning of the M2 beam line [M2] depicted in Figure 2.3, remaining beam protons are partly removed by momentum selection via magnets B1–B3 and collimators. The secondary beam of pions, protons and kaons then traverses a decay line of ∼600 m, after which momentum is selected again at bending

[8]May be chosen among *beryllium* of varying lengths and widths, *silicon* and *none*.

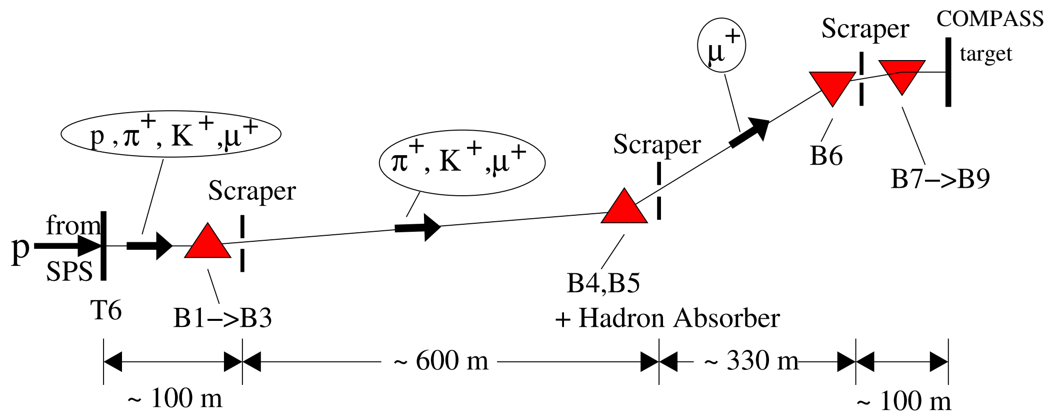

Figure 2.3: M2 beam line in *muon* configuration (from [vH02])

magnets B4–B5. At the same position, a hadron absorber of 9.9 m beryllium may be moved into the beam. The remaining $\sim$400 m of the beam line are used for transport to the experimental hall.

This setup is flexible enough to bring beams of positive and negative polarised muons, positive and negative hadrons, and electrons to the experiment.

Muon beam: By selecting positive hadrons of $\sim$177 GeV/c momentum at B1–B3 and inserting the hadron absorber at B4–B5, positive muons are transported to the hall, stemming from pion decay during flight through the decay channel are transported to the hall. The maximal parity violation of the weak $\pi^+ \rightarrow \mu^+ + \nu_\mu$ decay of the pion is exploited by selecting muons of $\sim$160 GeV/c at B4–B5, leading to natural longitudinal polarisation of -0.75$\pm$0.04.[9] The resulting muon intensity is $\sim 2.8 \cdot 10^8$ per spill with a remaining hadronic impurity of $\sim$1%; the momentum spread is $\sim$3% RMS. Due to its nature as a tertiary beam, the beam size at the target with $8 \times 8\,\mathrm{mm}^2$ RMS is rather large with divergence of $\pm$0.5 mrad horizontally and $\pm$1 mrad vertically. A strong halo is present. (This description refers to the setting considered best for physics data taking, variations are conceivable.)

Hadron beam: A hadron beam may be extracted simply by removing the absorber at B4–B5. However as the resulting beam still contains large amounts of (scattered) beam protons, an additional twist is added: Polarities of all magnets are reversed so that the bending magnets B1–B3 select *negative* hadrons. For momentum 190 GeV/c (with spread of 0.7% RMS), the composition of the hadron beam is 95% π^-, 4.5% K^- and 0.5% $\bar{p}$. The beam size at the target with $3 \times 4\,\mathrm{mm}^2$ RMS is five times smaller than that of the muon beam. The intensity initially was foreseen to be larger than 10^7 pions per spill [COM04, Fer05], yet in the interest of stable data taking it was not pushed higher than $\sim 4 \cdot 10^6$ pions per spill during the pilot run of 2004.

[9]This value has been obtained exclusively by simulation. Measurement was not deemed necessary as that of the SMC experiment, precursor to COMPASS in the same hall, had shown good agreement with the simulated value.

Electron beam: The electron beam may be gained by selecting a negative secondary beam of momentum approx. 100–120 GeV/c at B1–B3 after the T6 production target. In the decay channel a 5 mm lead converter is inserted, decelerating the electrons by bremsstrahlung. From the now tertiary beam of decelerated electrons, a momentum may be selected in the range of 30–60 GeV/c at B4–B5 (the absorber being retracted, of course) and transported to the spectrometer. The electron beam mainly is used for calibration of the calorimeters.

2.2.2 The Targets

For the two different physics programmes, two completely different types of targets are used.

Muon physics: For physics with polarised muon beam, a target consisting of two cylindrical cells of ^{6}LiD is utilised. The cells of 3 cm diameter and 60 cm length are polarised in opposite direction with two separate microwave systems, employing the method of Dynamic Nuclear Polarisation (DNP). The polarisation then is sustained by the extremely homogeneous 2.5 T field of a trim-able superconducting solenoid at temperatures of 80–100 mK maintained by a dilution refrigerator[10]. The polarisation is measured by Nuclear Magnetic Resonance (NMR), routinely exceeding its design value of 50%.

In addition to the longitudinal field of the solenoid, a transverse field of up to 0.5 T may be generated by a dipole magnet in vertical direction. This setup allows for reversal of longitudinally oriented nucleon spins by field rotation (ramping both magnets in a way that slowly rotates the total field vector, taking spins with it) or for holding the spins in transverse orientation.

Further information on the polarised target may be found in [Tak02].

Hadron physics: For hadron beam physics, unpolarised targets of lead, copper, carbon and CH_2 (plastic) are used, ranging from 12% radiation length (CH_2) to 53% radiation length (lead) and from 2.9% interaction length (lead) to 9.1% interaction length (CH_2). For Primakoff data taking, of course the lead target with its high radiation length is used while diffractive measurements call for the high interaction length of the CH_2 target. The targets with intermediate Z (copper and carbon) are employed for systematic studies of Z-dependence.

The targets are placed on a foam holder inside the **Veto Box**, a barrel of 12 scintillator plates and 96 lead glass blocks used to identify recoil nucleons from the target. The diffractive processes cause recoil energy less than $\sim$700 MeV/c and those of Primakoff reactions are lower, still, so that recoil energies above a certain threshold may be used as veto for the hadron triggers [Fer05].

After 2005 a liquid hydrogen target will we used for production of exotic hadrons through diffractive processes and central production.

[10] A cooling method that is taking advantage of phase-transition effects inside a mixture of ^{3}He and ^{4}He, distinguished by its high cooling power at very low temperatures.

2.2.3 The Detectors

Tracking: The tracking detectors of COMPASS may be differentiated by the size of their active area and their rate capability. Most detectors contain de-activated centres or beam holes to avoid discharges and to reduce the amount of matter in the beam. Throughout the $\sim$60 m of the spectrometer, several tracking stations are situated, consisting of a staggered setup of different trackers with increasing sizes, each taking advantage of a smaller detector with higher rate capability to fill its blind spot close to the beam.

Very Small Area Tracker (VSAT): These trackers are the only ones to stand the full rate and radiation dose of the beam and do not have a blind spot in the centre. Active areas are ranging from 16 cm^2 to 150 cm^2.

The Beam Momentum Station (**BMS**) is positioned *around* bending magnet B6, 60–140 m upstream of the target, consisting of four planes of horizontally mounted hodoscope slabs that were reinforced by two planes of scintillating fibres oriented likewise. They provide tracking in the vertical plane, from which the momentum of incoming beam particles may be determined in conjunction with knowledge of the field strength of B6. Each scintillator plane consists of slabs 5 mm high and 20 mm wide (in direction of the beam). The resulting time resolution is $\sim$260 ps at an efficiency of 80–90% per plane [vH02]. The BMS detectors have been removed during hadron beam data taking due to their high hadronic interaction length.

During muon beam data taking, eight stations of scintillating fibres (**SciFi**) are installed, of which two (FI01–02) serve for beam definition upstream of the target, constituting a part of the *beam telescope*. With their excellent time resolution of 350–500 ps and spatial resolution of 130–250 μm they serve for time-tagging of the events [Web04]. Two slightly different designs have been implemented by groups from Nagoya, Japan (FI01–04, "SciFiJ") and from Bonn, Germany (FI05–08, "SciFiG"). The size of the active area is increasing in beam direction, ranging from 4×4 cm^2 to 12×12 cm^2. During hadron beam data taking, some of the fibre stations were removed because of their high interactions lengths.

Silicon micro-strip detectors serve complementary purposes with regard to the fibre trackers: They excel with a spatial resolution of $\sim$10 μm while their time resolution of 2–4 ns is worse. Their active area is 5×7 cm^2. For muon beam data taking of 2004, three stations have been used as the other constituent of the beam telescope, while during the hadron beam, the three stations were used as vertex detectors directly after the target and two new stations took over the task of beam definition, so that a total of five stations now are operational. A more detailed description of the silicon trackers and their decoding is given in chapter 3.

Small Area Tracker (SAT): Detectors of this class cover an area of 0.1–0.2 m^2. Their centres are deactivated (at least in regular data taking conditions). They utilise a very small gas volume and different types of gas amplification to achieve high resolution in space and time.

Three stations of **MicroMeGas** (Micro-Mesh Gas) detectors with four projections each, mounted between target and SM1 record tracks that have missed the acceptance of the silicon trackers. Employing gas amplification through a metallic micro-mesh, they provide an active area of $40\times40\,\mathrm{cm}^2$ with spatial resolution of $\sim$70 μm and time resolution of $\sim$10 μs [Mag02].

The backbone of the SAT system after SM1 is constituted by 11 **GEM** (Gas Electron Multiplier) stations that use perforated, copper-clad kaption foils for gas amplification. They provide coverage of $32\times32\,\mathrm{cm}^2$ with spatial resolution of $\sim$50 μm and time resolution of $\sim$12 μs [Wei03, Ket03].

Large Area Tracker (LAT): The LAT consist entirely of wire chambers of different types, covering an area of 1.4–13.5 m^2. The centres of all LATs are deactivated.

Three stations of "Saclay" drift chambers (**DC**) are mounted behind the target and behind SM1, covering an area of $1.2\times1.2\,\mathrm{m}^2$. Their spatial resolution is approx. 0.25 mm.

To avoid excessive dead time due to space charge in the gas, the active volume of **Straw** drift chambers is segregated into a large number of cylindrical compartments of 6mm respectively 10mm diameter, so-called "straws". Five modules, each containing three double layers, have been mounted in the LAS and SAS. They provide a spatial resolution of $\sim$0.2 mm inside an active area of up to $3.25\times2.77\,\mathrm{m}^2$. For more information on the Straws, refer [Ilg03, in German].

The 15 chambers of **MWPC** (Multi-Wire Proportional Chamber) inherited from the OMEGA[11] spectrometer provide coverage of up to $1.78\times1.2\,\mathrm{m}^2$ with 2 or 3 projections per chamber [Bar98]. The accuracy of measurement is 0.5 mm.

An heritage from the EMC experiment, the **W45** chambers with $5.2\times2.6\,\mathrm{m}^2$ cover the largest area by far. Their spatial resolution is $\sim$1.9 mm.

Particle Identification: COMPASS possesses various systems for particle identification:

Two CERN-supplied **CEDAR**[12] detectors are mounted in the beam line between B6 and B8–B9 during hadron beam, serving to tag pions and kaons, respectively.

In the first stage of the spectrometer (LAS), a custom-built **RICH**[13] is situated for identification of secondary particles. Cherenkov photons created in a 83 m^3 volume of the radiator gas C_4F_{10} are reflected by focussing mirrors onto CsI photo-cathodes behind quartz windows. The resulting photo electrons receive gas amplification and induce charges on the photo-cathodes also serving as pad readout. The design should allow separation of pions, kaons and protons of momenta up to 55 GeV/c while detection thresholds of the gas are 2.5 GeV/c for pions, 8.9 GeV/c for kaons and 17.0 GeV/c for protons. The average number of photons per ring was 19 in 2002 and 20 in 2003 while 36 have been proposed initially [Fau04].

[11] a former spectrometer facility at CERN West Area

[12] **CE**renkov **D**ifferential counter with **A**chromatic **R**ing focus, cf. [Bov82]

[13] **R**ing **I**maging **CH**erenkov Counter

A similar RICH-2 detector for the second spectrometer stage (SAS) is planned for after 2005.

The large penetration power of muons is exploited for their identification. Two **Muon Filters** (often sloppily referred to as *muon walls*) are positioned at the ends of LAS and SAS, respectively. They consist of a series of tracking detectors (Iarocci tubes in MF1 and scintillator, aluminium drift tubes, and MWPC in MF2) before and after a massive absorber (60 cm iron and 240 cm concrete, respectively). Particles that may be detected before *and* after the absorbers are identified as muons, being the only particles with sufficient penetration power to traverse the absorber (and sufficient life time to reach the detector).

Calorimetry: To determine the energies of photons and electrons, two electromagnetic calorimeters **ECAL1** and **ECAL2** have been foreseen after the tracking detectors of the two spectrometer stages. Utilising well-known lead glass blocks from the GAMS calorimeter, the energy resolution may be specified as $5.5\%/\sqrt{E/\mathrm{GeV}} \oplus 1.5\%$ and the spatial resolution is quoted as $6\,\mathrm{mm}\,/\sqrt{E/\mathrm{GeV}} \oplus 0.5\,\mathrm{mm}$. While ECAL1 up to now only exists as mechanical structure (no detector material or readout installed, yet), ECAL2 in 2003 was operated with 2000 lead glass blocks with 12 bit ADC read-out and in 2004 additional 1000 lead glass blocks blocks have been installed with newly developed 10 bit sampling ADC read-out. The new read-out was tested during muon beam data taking and work was completed in time for production use in the hadron beam 2004 [COM04, Kon05].

Measurement of the energies of hadrons is handled by the two hadronic calorimeters **HCAL1** and **HCAL2**. They consist of iron-scintillator sandwich and provide energy resolutions of $80\%/\sqrt{E/\mathrm{GeV}} \oplus 8\%$ and $60\%/\sqrt{E/\mathrm{GeV}} \oplus 6\%$ for pions on active areas of $4.2 \times 3\,\mathrm{m}^2$ and $4.4 \times 2\,\mathrm{m}^2$, respectively. The spatial resolution of HCAL1 varies between 4 mm and 14 mm, depending on the position of the hit relative to block boundaries.

2.2.4 The Triggers

Muon triggers: The COMPASS muon triggers may be differentiated by the reaction they trigger on. Figure 2.4 shows the location and the shapes of all muon trigger hodoscopes. All triggers share a veto system designed to reduce contamination by the large muon halo. The four veto hodoscopes VBL, VI1, VO1 and VI2 are placed around the beam line upstream of the target and inhibit triggering on beam particles that are missing their openings (halo).

The **Inner Trigger** (IT) and **Ladder Trigger** (LT) serve to identify the PGF process by exploiting the fact that the scattered beam muon suffers an energy loss, yet receives only small transverse momentum. To separate those muons from muons scattered at larger angles *and* larger energies it is necessary to setup hodoscopes at *two* points of the trajectory. The two hodoscopes (HI04 and HI05 for IT, HL04 and HL05 for LT) for every trigger are connected to a coincidence matrix. Only those combinations

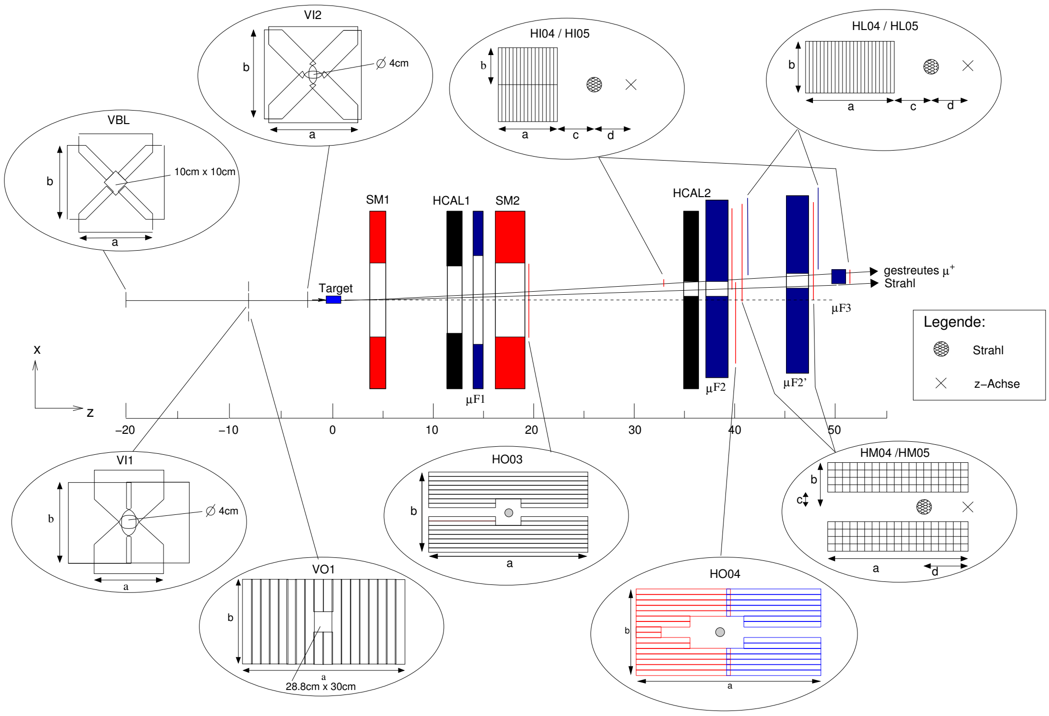

Figure 2.4: Locations and shapes of the muon trigger hodoscopes (from [vH02], modified in [Leb02])

are configured to initiate a trigger signal, that correspond to the desired trajectory of small transverse momentum, high energy loss muons. For further suppression of background, a certain minimal energy in one of the hadronic calorimeters is required additionally.

The **Middle Trigger** (MT) and **Outer Trigger** (OT) select the DIS process by requiring the four-momentum Q^2 of the virtual photon to be larger than $\sim 0.5\,\mathrm{GeV}^2/c^2$. Similarly to the PGF triggers, this is accomplished by a pair of hodoscopes for each trigger. The hodoscopes HO03 and HO04 of OT have horizontally orientated slabs mounted symmetrically around the beam while HM04 and HM05 of MT, each consisting of two planes of perpendicularly oriented slabs, even provide two-dimensional information. Again the desired interactions are selected by coincidence matrices, disentangling scattering angle and magnetic deflection and providing basic target-pointing. The Middle Trigger also requires a minimal energy in one of the HCALs, yet there also is the version of the **inclusive Middle Trigger** (inclMT) that has no calorimeter threshold. This, of course, results in a higher trigger rate for inclMT which in turn has required to pre-scale it with a factor of two in 2004.

Hadron triggers: The following elements are shared by all hadron triggers: Beam particles are defined by requiring hits to two scintillators position in the beam line. To avoid triggering on beam particles that do not hit the target, upstream of FI01 a scintillator with a hole of $\sim$4 cm diameter is situated as Beam Veto. Additional vetoes

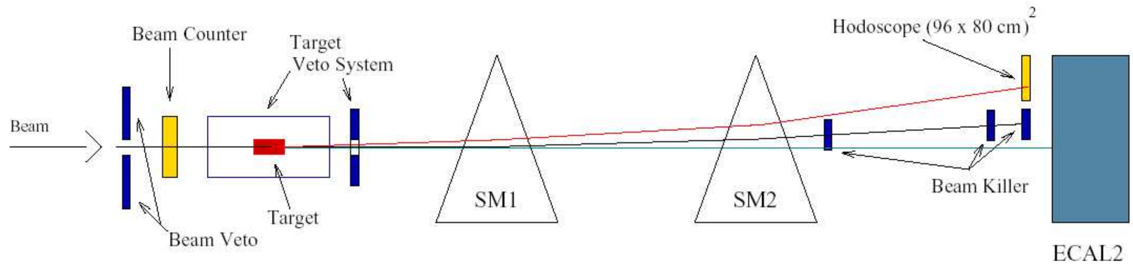

Figure 2.5: Setup of the Prim_1 trigger, from [Fer05]

are provided by the Veto Box barrel around the target and by an Aperture Sandwich veto after the target that is activated by secondary tracks with large angles.

The **Prim_1** additionally requires a hit to the Primakoff Hodoscope, geometrically imposing a scattered beam particle with momentum in the range 20–110 GeV/c. For ECAL2 a threshold of 40 GeV is set as minimal energy of the Primakoff photon.

The **Prim_2** trigger serves to catch Primakoff events leaving less than 20 GeV/c of momentum to the beam particle, which consequently by the magnets is bent too far to hit the Primakoff Hodoscope. To increase selectivity, an ECAL2 threshold of 100 GeV is required.

For the **Diff_1** trigger a multiplicity larger or equal to two is required in the Multiplicity Counter (MC), a scintillator placed after the target from which rough multiplicity information may be deduced from the signal level of the attached photomultiplier tube. In HCAL2, an energy threshold of 6 GeV needs to be exceeded. Additionally a **Beam Killer** veto is introduced. It consists of three small scintillators (BK1, BK2 and BK3) mounted in the SAS, tracing the beam. A hit to at least two of the three is taken as indication that a beam particle exists that has not been scattered and causes a veto on that event.

The **Charge Exchange** trigger requires zero multiplicity in the MC and uses an ECAL2 threshold of 40 GeV in addition to the common elements of all hadron triggers. BK3 is used as veto.

The planned **Diff_2** trigger was not realised in the pilot run of 2004.

2.2.5 The Data Acquisition System (DAQ)

The COMPASS DAQ system has the task to read-out over 200,000 detector channels stretched over approx. 190 m (from BM01 to HI05, implying over 600 ns time shift between the first and the last detector due to time-of-flight), at a trigger rate of 10–20 kHz in 2004 (with a design goal of 100 kHz after upgrades).

This is accomplished by a layered system of continuing data concentration depicted in Figure 2.6. Situated at the top of the read-out chain are the detector front-ends and their read-out modules, which upon reception of the trigger signal from the Trigger Control

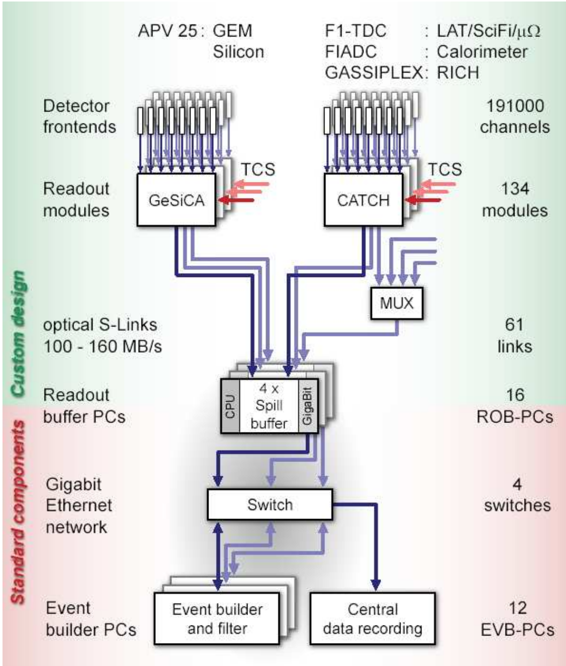

Figure 2.6: Architecture of the DAQ system (from [Sch04c])

System (TCS) take a snapshot of their detector channels and after zero-suppression transmit the data to the spill buffers via optical link. The spill buffers perform "data derandomisation": They absorb the high on-spill data rates into their buffer while emitting data continuously at a lower rate towards the event builders, where the data from different detectors are assembled into events and online filtering takes place and accepted events are written to disk. From disk the data are transferred to its long-term storage on CASTOR[14] by CDR[15] over a fibre-optic link of 1 GBit/s bandwidth.

Trigger Control System (TCS): The TCS is the component that is steering detector readout. Its 38.8 MHz clock provides stable and synchronous time to all read-out modules (who pass it on to the front-ends) through a unidirectional laser-driven fibre-optic network with passive fan-out. The same network is used to transfer synchronous trigger signals along with event number and type, and meta-data such as configuration information. The TCS controller provides for configuration of one fixed and two flexible dead times and also may be used to broadcast calibration triggers [Gru01, Kon01, Sch04c].

Read-out Modules: With a TCS receiver module, the two types of read-out modules, CATCH[16] and GeSiCA[17], pick up the signals of the TCS controller. Their task is to

[14]**C**ERN **A**dvanced **Stor**age Manager, [CAS]

[15]**C**entral **D**ata **R**ecording, cf. [CDR]

[16]COMPASS **A**ccumulate, **T**ransfer and **C**ontrol **H**ardware, documented in [Fis02], is the front-end driver used for readout and data concentration of most detectors.

[17]**Ge**m **Si**licon **C**ontrol **A**cquisition, the front-end driver for Silicon and GEM, is specialised in read-out of

Beam	Ev. Size	Rate	Bef. Filter	Aft. Filter	Bef. Filter	Aft. Filter
Muon	39 kB	9 kHz	44 kTrigger	34 kTrigger	102 MB/s	79 MB/s
Hadron	28 kB	17 kHz	82 kTrigger	74 kTrigger	133 MB/s	120 MB/s

Table 2.1: Typical DAQ performance parameters 2004 for muon beam and hadron beam. The columns contain in order: beam type, event size, on-spill trigger rate, trigger per spill before online filter, trigger per spill after online filter, data rate before online filter, and data rate after online filter.

drive the front-end modules and to serve as a first layer of data aggregation. Connected to multiple front-ends via short-range interconnects such as HotLink, they combine incoming data, perform partial header suppression and data formatting and transfer the output by a single S-Link [Boy97] fibre connection to a spill buffer in a ROB.

Read-out Buffer (ROB): The read-out buffers are PCs made from widely available standard components, yet including up to four of the custom-design spill buffers with 512 MBytes derandomisation FIFO as PCI[18] cards. Data are fetched from the spill buffers using a custom-built driver in the Linux kernel and taken over by the ALICE DATE software, which distributes them through a switched TCP/IP/Ethernet[19] network to the event builders. To increase the bandwidth *inside* the ROB, the spill buffers and the Ethernet controller reside on two different PCI busses.

Event Builder (EVB): The sub-event information is sent from the read-out buffers to the 13 event builder PCs round-robin in a way that all parts of the same event are received by the same event builder, respectively. At the event builder complete events are assembled from the received fragments passed on to the online filter which decides to keep or to reject them, only writing them to disk in the former case. Each event builder provides storage of up to 640 GByte on the local RAID array[20], so that all 13 event builders together provide 8.1 TB of temporary storage, which is used as a buffer to uncouple the experiment from fluctuations of tape writing. (This capacity was equivalent to two days of data taking in 2003 and $\sim$1.3 days in 2004.)

Several factors are limiting the capacity of the DAQ at various points in the read-out chain, yet for practical purposes it is interesting which of the limits is lowest and effectively imposes the upper bound to the amount of data that can be recorded. Table 2.1

APV chips. It is using the same data format as the CATCHes, commonly referred to as CATCH blocks, and has been described in [Gru01].

[18]**P**eripheral **C**omponent **I**nterconnect, a bus used on the motherboards of almost all current PCs

[19]The de facto standard for local-area networks. Described according to the Open Systems Interconnection (OSI) Reference Model it consists of the **T**ransmission **C**ontrol **P**rotocol (transport layer) built upon the **I**nternet **P**rotocol (network layer), which in turn relies on the Ethernet protocol (data link layer) to transfer data over optical fibres (physical layer). Network topology is that of several interconnected "stars" with computers centred around switches operating on the network layer.

[20]The **R**edunant **A**rray of **I**ndependent **D**isks denotes a method of pooling several hard disks into one logical drive to increase reliability and/or performance.

gives an overview of crucial parameters of DAQ performance in 2004. Possible limits that are to be considered include:

- The trigger rate is limited by detector/front-end capabilities, the bandwidth between CATCH/GeSiCA and ROB, and transfer capacity inside the ROB. This complex has been tested successfully up to 100–120 kTrigger per spill at event size of $\sim$39 kByte during DAQ tests early in 2004, albeit only with a subset of all detectors.

- Another limitation is that of I/O bandwidth on the RAID arrays on the event builders. While the RAIDs are expected to perform very well unidirectionally ($\sim$125 MByte/s for reading and $\sim$30 MByte/s for writing per event builder), performance degrades drastically during simultaneous input and output as it is conducted during regular physics data taking. However the throughput could be increased significantly by aggregation of output into large blocks by Cinderella before handing it over to the operating system, instead of calling for many segmented writes as is done by ALICE DATE.[21] Still, presently it is unclear which is the maximal transfer rate that the RAIDs may handle in this fashion.

- The third boundary is the capacity of 1 GBit/s[22] (119 MByte/s) of the fibre-optic link to CASTOR. The theoretic upper limit to usable bandwidth (after overhead of protocol stack) is approximated to be 110 MByte/s [Ku04]. There are plans to upgrade the link to tenfold capacity.

- Tape writing is limited by the bandwidth that is provided by CASTOR. While this posed a severe limitation in 2002, following the continuing extension of CASTOR it has been less of a problem in the following years. For the future it is expected to be of further decreasing concern. The cost of storage of currently 0.80 SFr/GByte (2 SFr/GByte in 2002) however will have to be considered also in the next years.

- The capacity of the switches connecting read-out buffers and event builders need to be taken into consideration, too. However in the case that it is needed, a hardware upgrade would be feasible.

In this environment, the role of Cinderella is to ease the load on RAID, link, and CASTOR. While this would not have been absolutely critical for the data taking of 2004, it certainly helped to increase the stability of the DAQ in general. However it has to be remembered that the operation of Cinderella in 2004 was its first use in production, future versions are likely to make a greater contribution.

[21] This is one of the reason for increased DAQ stability when running with Cinderella.

[22] Throughout this document, for specification of data volumes the convention has been utilised to let the quantifiers k, M, G and T denote 2^{10}, 2^{20}, 2^{30} and 2^{40}, respectively, as is customary in computing applications. Unfortunately, this custom does not extend to the specification of Ethernet bandwidth, so that the 1 GBit/s quoted actually equal 10^9 Bit/s.

2.2.6 Data Reconstruction

The extraction of physics results from the recorded events is an elaborate undertaking. First, in an operation commonly referred to as *production*, the raw data are processed with the CORAL[23] software that performs decoding, clustering, tracking, particle ID and vertexing, and stores its results in mDST[24] files on CASTOR. The size of a typical 200-spill run, representing 56 minutes of muon beam data taking, is reduced by a factor of $\sim$60 from 260 GByte of Cinderella-filtered raw data to $\sim$4.1 GByte mDST data. The production is handled entirely by the `lxbatch` cluster of CERN.

In the next step, performed within the framework of PHAST[25], the (partial) knowledge of tracks, vertices, momenta, charges, and particle types is exploited to reconstruct the observables of physical processes according to user-specified C++ functions. To speed up analysis, PHAST may generate sub-samples of mDSTs (so-called micro-DSTs), again utilising user-specified functions, which later may be re-processed with PHAST, reducing the amount of data that has to be iterated over in further steps of the analysis.

[23]**CO**MPASS **R**econstruction and **A**nalysis **L**ibrary, cf. [COR]
[24]mini **D**ata **S**ummary **T**ape
[25]**PH**ysics **A**nalysis **S**oftware **T**ools, for more information see [PHA]

Chapter 3

The COMPASS Silicon Micro-Strip Detectors

Among other detectors, Cinderella is making heavy use of time information provided by the COMPASS silicon detectors. In this chapter their decoding is described.

3.1 General Principles of Operation of Silicon Micro-Strip Detectors

A silicon particle detector consists of a p-n junction operated at reverse bias. Usually the bulk of the junction is wide and only slightly doped while the other side is a narrow implant of high density of charge carriers of the other sort. This design allows full depletion of the bulk with relatively low bias voltage. Upon passage of charged particles through the silicon bulk, clouds of charge carriers form along their track via two processes: Either electrons of the valence band are excited directly to the conduction band, creating electron-hole pairs with a deposited energy of 3.62 eV per pair[1] at room temperature[2], or electrons receive a higher energy transfer from the charged particle and become knock-on electrons, which again in their turn loose their energy to the bulk by one of the two mechanisms. Before recombination of electrons and holes can take place they are separated by the bias voltage and drift towards opposite contacts where they finally are read out.

To make the detector sensitive to track position, the readout contacts are divided into narrow strips. Assuming that all of the charge created by one particle is deposited on the same readout strip, the position of the track can be taken to be the strip position. Then it is easy to derive the spatial resolution from the pitch of the readout strips δx to be $\sigma_x = \frac{\delta x}{\sqrt{12}}$ [Wag01]. For a typical pitch of 50 μm, spatial resolution thus would be $\sim$14 μm.

[1]Si being an indirect semiconductor, the ionisation energy is larger than the band gap because excitation of a phonon is necessary for momentum conversation.

[2]Deposited energy rises to 3.81 eV per pair at 77 K [Leo94].

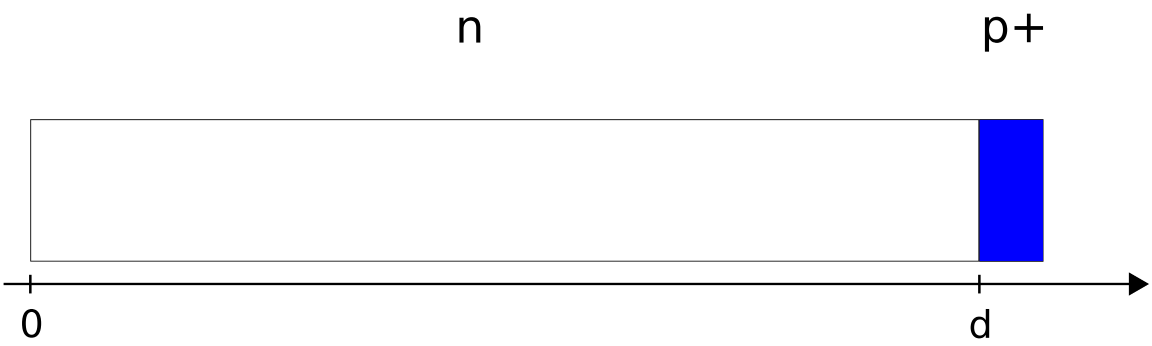

Figure 3.1: Coordinate system used for the calculations

However when the signal of a track is spread over multiple readout strips, the track position can be determined using more sophisticated approaches like centre-of-gravity calculation or fitting of a Gaussian[3] to the distribution of signal height over the strips. This can yield higher resolution by a factor of two [Fuc03] or three [Pei92]. This consideration would suggest designing a detector with readout strips as narrow as possible. Yet, with growing number of strips sharing the charge of a track, the signal-to-noise ratio of the individual strips and also that of a reconstructed cluster decreases, degrading resolution both in space and in time.

A good solution is the introduction of an uncontacted (floating) intermediate strip between every pair of adjacent readout strips. The charge collected by an intermediate strip is distributed to both neighbouring readout strips equally via capacitive coupling (AC coupling), thus widening the effective size of the charge cloud, increasing cluster size.

3.2 The Biased p-n Junction

The silicon detectors of COMPASS consist of a bulk with light negative (phosphorus) doping and a thin region of high positive (boron) doping. The following calculations assume a coordinate system with the x-coordinate extending through the bulk of the material perpendicular to the surface of the wafer, as depicted in figure 3.1. The interface of the n-type bulk to the external contact is taken as zero position while the edge towards the p-type implant is set as $x = +d$, the thickness of the bulk being denoted by d.

Silicon detectors always are operated fully depleted (active area). In that condition, an external voltage is applied that is sufficient to fully remove (deplete) the majority carriers (electrons) from the bulk.

The starting point for calculation of electric field inside the bulk is the one-dimensional Poisson equation

$$\Delta\,\phi(x) = -\frac{\rho(x)}{\epsilon}, \tag{3.1}$$

[3] or another function which fits the shape of the charge distribution

with the electric potential $\phi(x)$, electric charge density $\rho(x)$ and permittivity ϵ (cf. [Jac99]).

Assuming an external bias voltage V_{dep} which is just high enough to fully deplete the bulk, the space charge density is given by the doping density to be $\rho(x) = N_D \cdot e$. Then the Poisson equation is

$$\Delta\, \phi(x) = -\frac{N_D \cdot e}{\epsilon}. \tag{3.2}$$

In the absence of external fields eq. 3.2 is the complete description of the electrostatic problem. It may be solved by double integration. As the potential is defined with sign opposite to the field in $E(x) = -\frac{d}{dx}\phi(x)$, the electric field $E(x)$ is yielded by the first integration step as:

$$E(x) = -\int_0^x \frac{\rho(x)}{\epsilon}\, dx = \int_0^x \frac{N_D \cdot e}{\epsilon}\, dx = x \cdot \frac{N_D \cdot e}{\epsilon}. \tag{3.3}$$

Due to electric neutrality of the wafer as a whole, the electric field is zero outside of the detector. The potential $\phi(x)$ is gained by the second step of integration (choosing $\phi(0)$ to be zero potential):

$$\phi(x) = \int_0^x -E(x)\, dx = -x^2 \cdot \frac{N_D \cdot e}{2 \cdot \epsilon}. \tag{3.4}$$

The voltage drop over the width d of the bulk now can be obtained by evaluating the potential

$$V_{dep} = \phi(0) - \phi(d) = \frac{d^2 \cdot N_D \cdot e}{2 \cdot \epsilon}, \tag{3.5}$$

which may be transformed to yield the depth d of the depletion zone depending on the applied external voltage as

$$d = \sqrt{\frac{2 \cdot V_{dep} \cdot \epsilon}{N_D \cdot e}}. \tag{3.6}$$

For the p-n junction with depletion depth of the bulk d_n, to assure electrical neutrality of the detector as a whole, the p^+ side needs to be depleted to a depth of $d_p = \alpha \cdot d_n$, with $\alpha = \frac{N_D}{N_A}$. The voltage drop over the p^+ side then follows to be $V_p = \alpha \cdot V_n$, so that with $N_D \ll N_A$, its contribution to the bias voltage may be neglected. This consideration leads to an expression of the depletion voltage[4] dependent on the properties of the detector material:

$$V_{dep} = \frac{d_n^2 \cdot N_D \cdot e}{2 \cdot \epsilon}. \tag{3.7}$$

In the next section the utility of this equation is demonstrated by determining the concentration of the dopants in the bulk of the wafer, which is not well known from other sources.

[4]unfortunately misprinted by a factor of 1/2 in [Wag01]

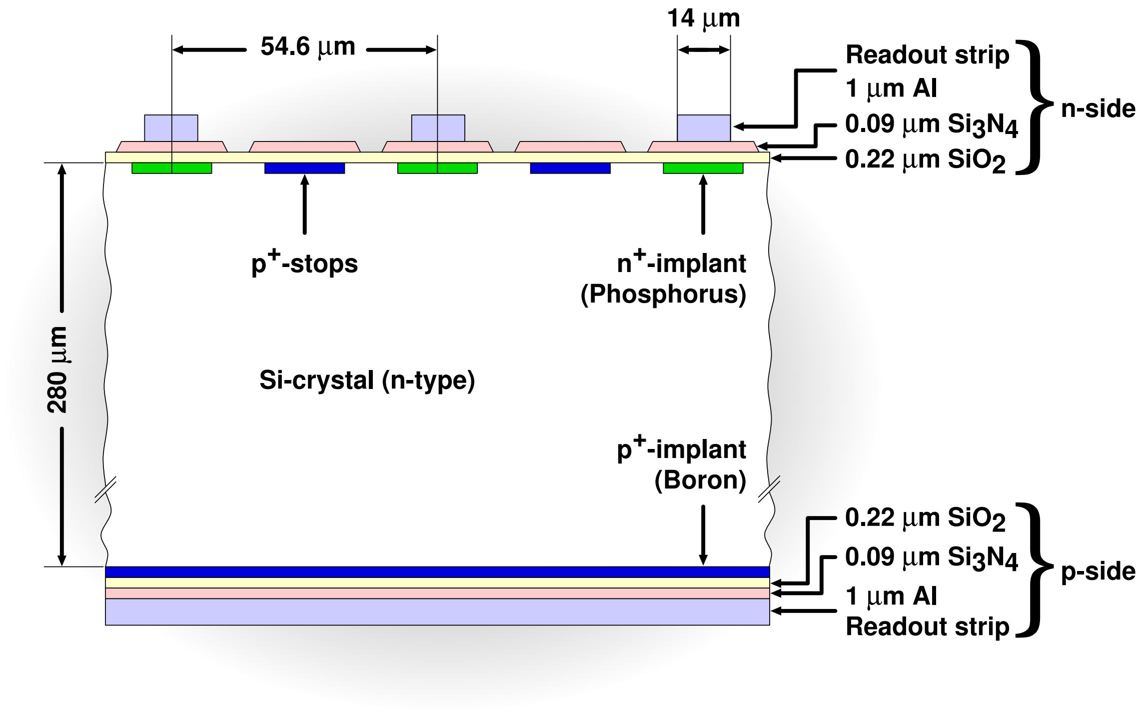

Figure 3.2: Schematic cut view of the silicon detector used in COMPASS [Gru01]

3.3 COMPASS Detector Design

The COMPASS experiment uses SI wafers that were designed by Max-Planck-Institut für Physik, München for the HERA-B experiment and produced by SINTEF[5]. The wafers of 300 μm thickness are n-doped with phosphorus. SINTEF quotes depletion voltage of 84–92 V for one batch ("Pasing p-spray") and 26–48 V for another ("SINTEF p-stop"). Using eq. 3.7 the donator concentration—accessible only indirectly and inaccurately from the data sheet [SINb]—can be calculated to be $3.1-3.4\cdot10^{12}\,\mathrm{cm}^{-3}$ for the first batch and $1.0-1.8\cdot10^{12}\,\mathrm{cm}^{-3}$ for the second batch.

The active area is of size $5\times7\,\mathrm{cm}^2$ and has orthogonal readout strips on both sides of the wafer that are coupled capacitively to the bulk to reduce leak current. The pitch of the readout strips is approx. 50 μm. Owing to details of the manufacture process, only the p-side[6] is equipped with intermediate strips, so that on this side spatial resolution is better and temporal resolution is worse than on the n-side. Bearing the fact that the spectrometer magnets deflect particles in the horizontal plane, the wafers detectors are mounted in such a way as to use the side with better spatial resolution in determining the horizontal coordinate of the tracks.

[5] The Foundation for Scientific and Industrial Research at the Norwegian Institute of Technology, a non-profit organisation for applied science [SINa].

[6] unfortunately misprinted as n-side in [Wag01]

The readout strips are connected to APV 25 front-end chips (see [Gru01, chap. 5.1] for details) where the signals first pass a shaper/pre-amplifier circuit which transforms the irregularly shaped incoming charge pulses[7] to voltage pulses of well-defined uniform shape (only differing in amplitude). This signal then is continuously sampled into an analogue pipeline at the TCS[8] frequency of 38.8 MHz. Upon trigger, three consecutive samples[9] are read out from a position in the pipeline which is earlier by a configurable amount of time to compensate for trigger delay.

The three samples then are digitised using a 10 bit differential ADC[10] and transferred to an FPGA[11] named 'zero-chip' where baseline correction, common-mode noise correction and zero suppression are performed (for details, again see [Gru01]).

More information on the COMPASS silicon detectors may be found in [Wag01, Gru01, Wie04, Fuc03].

3.4 The Silicon Signal

3.4.1 Equation of Motion for Electrons and Holes inside the Bulk

Using a linear relation between electric field and velocity

$$v_e = -\mu_e \cdot E \qquad \text{resp.} \qquad v_h = \mu_h \cdot E \tag{3.8}$$

for electrons and holes generated by ionisation at position $x(t=0) = x_0$ with $v(t=0) = 0$, the equations of motion follow from the electric field (eq. 3.5) as

$$\dot{x}_e(t) = -x_e(t) \cdot \frac{\mu_e \cdot N_D \cdot e}{\epsilon} \qquad \text{resp.} \qquad \dot{x}_h(t) = x_h(t) \cdot \frac{\mu_h \cdot N_D \cdot e}{\epsilon} \tag{3.9}$$

with solutions

$$x_e(t) = x_e(0) \cdot e^{-t/\tau_e} \qquad \text{resp.} \qquad x_h(t) = x_h(0) \cdot e^{t/\tau_h} \tag{3.10}$$

with

$$\tau_e = \frac{\epsilon}{\mu_e \cdot N_D \cdot e} \qquad \text{and} \qquad \tau_h = \frac{\epsilon}{\mu_h \cdot N_D \cdot e}. \tag{3.11}$$

Owing to the triangle shape of the electric field, electrons follow an exponentially decelerated path with $x_e(t) \to 0$ and $v_e(t) \to 0$ for $t \to \infty$ whereas holes are accelerated exponentially until they hit the p^+ section at $x_h(t_{max}) = d$ with $t_{max} = \tau_h \cdot \ln \frac{d}{x_h(0)}$.

[7]Because the capacitance of semiconductors varies with temperature an amplifier sensitive to *charge* is preferred over one sensitive to voltage.

[8]**T**rigger **C**ontrol **S**ystem, see [Gru01]

[9]This refers to the 'multi' mode of the APV readout. There are other modes, none of which are of great importance in the COMPASS experiment.

[10]**A**nalogue-to-**D**igital **C**onverter

[11]**F**ield **P**rogrammable **G**ate **A**rray

3.4.2 Charge Induction by movement of Electrons and Holes inside the Bulk

Following [Leo94, chap. 10.4] the charge signal on the readout strips is not created by electrons and holes actually hitting the strips but by induction created by the *movements* of the charges.

Considering a system consisting of a capacitor formed by the contacts of the bulk and a charge q within, for purpose of calculation it may be assumed that the external voltage supply is too slow to follow the changes in detector capacitance that are induced by the moving charge inside, so that the system in effect is a closed system and energy conservation law may be applied.

$$E = E_{cap} + E_{pot} = \frac{1}{2} \cdot C \cdot V^2 + q \cdot \phi(x) = const \tag{3.12}$$

Now movement of the charge q by a distance dx changes potential energy by

$$dE_{pot} = q \cdot E(x) \cdot dx, \tag{3.13}$$

and induction of a charge dQ on the electrodes leads to change in capacitative energy of $dE_{cap} = V \cdot dQ$. Due to energy conservation $dQ = -dE$, so that the induced charge is given by

$$Q(x) = \frac{q \cdot N_D \cdot e}{V \cdot \epsilon} \cdot \int_{x_0}^{x} x' \, dx' \tag{3.14}$$

and with the equation of motion of the charge $x_q(t)$ may be written as

$$Q(x) = \frac{q \cdot N_D \cdot e}{V \cdot \epsilon} \cdot \int_{0}^{t} x_q(t) \cdot \frac{dx_q(t)}{dt} \, dt. \tag{3.15}$$

For electrons and holes (eq. 3.10) results:

$$Q_e(t) = \frac{x_0^2 \cdot e^2 \cdot N_D}{2 \cdot V \cdot \epsilon} \cdot \left(1 - e^{-2t/\tau_e}\right) \qquad \text{resp.} \qquad Q_h(t) = \frac{x_0^2 \cdot e^2 \cdot N_D}{2 \cdot V \cdot \epsilon} \cdot \left(e^{2t/\tau_h} - 1\right) \tag{3.16}$$

With eq. 3.7 follows

$$Q_e(t) = -e \cdot \frac{x_0^2}{d^2} \cdot \left(1 - e^{-2t/\tau_e}\right) \qquad \text{resp.} \qquad Q_h(t) = -e \cdot \frac{x_0^2}{d^2} \cdot \left(e^{2t/\tau_h} - 1\right) \tag{3.17}$$

so that the pulse shape of the total induced charge $Q_{tot}(t) = Q_e(t) + Q_h(t)$ of one electron-hole pair is given as:

$$\begin{aligned} Q_{tot}(t) &= -e \cdot \tfrac{x_0^2}{d^2} \cdot \left(e^{2t/\tau_h} - e^{-2t/\tau_e}\right) && \text{for} \quad t \in [0, t_{max}] \\ Q_{tot}(t) &= -e + e \cdot \tfrac{x_0^2}{d^2} \cdot e^{-2t/\tau_e} && \text{for} \quad t \in [t_{max}, \infty] \end{aligned} \tag{3.18}$$

with $t_{max} = \tau_h \cdot \ln \frac{d}{x_h(0)}$. In the limit $t \to \infty$ the induced charge therefore is $-e$.

To yield the induced charge from passage of a particle, eq. 3.18 is to be integrated over all charges generated along the particle track.

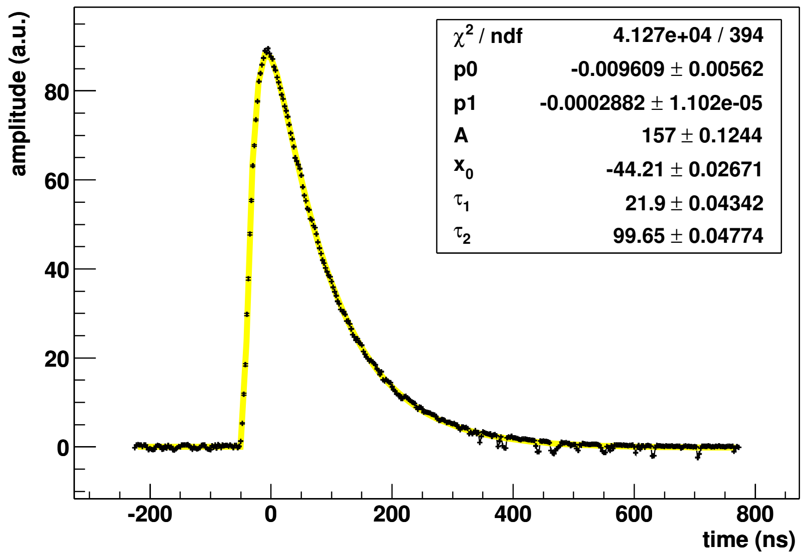

Figure 3.3: Response of APV shaper to δ-peak-like input fitted with the parametrisation of eq. 3.19 plus first order polynomial (from [Wie04]).

3.4.3 The Shaper

By fitting to experimental data, M. Wiesmann has shown in [Wie04], that the response of the APV shaper to the δ-function-like input from a pulse generator can be parametrised as

$$A(t) = A_0 \cdot (1 - e^{-(t-t_0)/\tau_{rise}}) \cdot e^{-(t-t_0)/\tau_{fall}}, \tag{3.19}$$

with $\tau_{rise} \approx 22\,\text{ns}$ and $\tau_{fall} \approx 100\,\text{ns}$.

The signal that is finally digitised at the ADC may be described as convolution of the signal as it is induced by the charge cloud on the readout strip and the characteristics of the shaping circuit.

3.5 Time Reconstruction

Upon trigger, the signal of the silicon detector is sampled three times at 25.8 ns intervals and stored in the "digits" a_0, a_1 and a_2. Due to varying pulse amplitude it is useful to begin time reconstruction with normalisation of the digits. This is achieved by calculation of ratios:

$$r_0 = \frac{a_0}{a_2} \qquad \text{and} \qquad r_1 = \frac{a_1}{a_2} \tag{3.20}$$

It is customary to plot the ratios against each other in a two-dimensional histogram, dubbed the *banana plot* due to its shape. A banana plot as calculated by Cinderella is

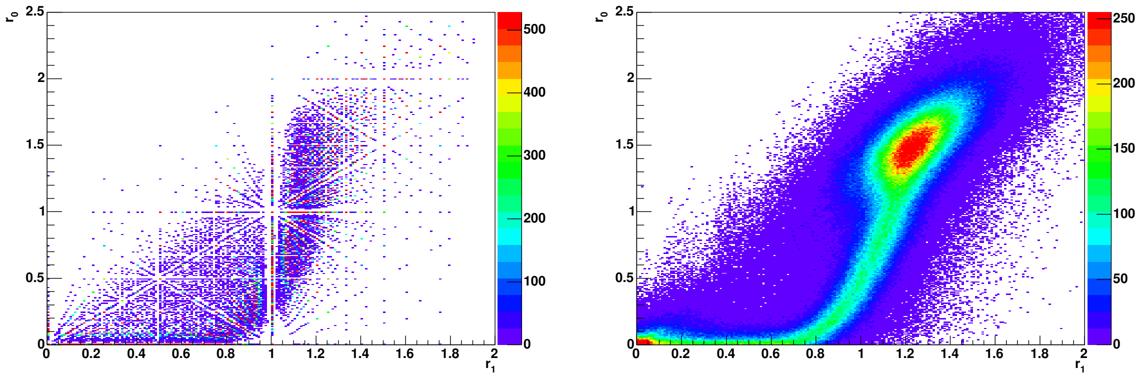

(a) The prominent line structure of this plot is an artifact resulting from the fact that r_0 and r_1 are fractions of small integers plotted against each other.

(b) The artifacts may be avoided by blurring the the digits a_0, a_1 and a_2 through adding of a random number between -0.5 and 0.5. This allows for easier analysis (especially fits are more likely to converge), however at the cost of slightly increased noise.

Figure 3.4: The banana plot illustrates the occurrence of combinations of r_0 and r_1.

shown in Figure 3.4 in two variations. The accumulation in the lower left corner around $r_0 = r_1 = 0$ results from occasions when a strip is sampled too early, i.e. a_0 and a_1 are measured *before* the signal has begun (reading only noise with low amplitude) and it is only a_2 that with the rising edge of the signal is sampling a larger value. For signals sampled somewhat later, a_1 is next to pick up parts of it which results in r_1 rising, while r_0 still stays close to zero. Only when the signal is measured still later, r_0 begins to rise, forming the bend of the "banana". At this point, time resolution is best. For signals that are measured later again—all sampling points now on the falling edge of the signal—the time resolution degrades again, as in the exponential decay the ratios approach a constant value, forming the large accumulation at $r_0 \approx 1.5$ and $r_1 \approx 1.25$.

The ratios forming the part of the "banana" between the two accumulations may be used to retrieve time information. This is illustrated in Fig. 3.5, showing the relation of ratios and particle time[12]. The underlying data were generated by A.-M. Dinkelbach [Din04] using tracking of COOOL[13] The plots demonstrate that the particle time may be reconstructed inside a time window of 40–60 ns for every ratio. Combining the partly overlapping periods of sensitivity of both ratios, if is safe to say that time may be reconstructed inside a sensitive period of 60–80 ns. To access the time information from the ratios it was suggested by [Fri04a] to employ the following parametrisation of the ratio as a function

[12] The polarity of the time is reversed with respect to the viewpoint taken in this text: The convention used for this plot sets the time of the signal as its fixed point, relating the time of the measurement to the time of the signal. This means that small ratios are seen to result from *early sampling* of a signal whereas large ratios are obtained from *late sampling* of the a signal. The point of view of the text is to employ the sampling time (which is related to the trigger time by an offset) as a fixed point, observing the time of the signal with regards to the time of the measurement. Large ratios in this terminology are yielded by sampling of an *early signal* (looking at its tail) whereas small ratios result from measuring a *late signal*, viewing its rising edge.

[13] COMPASS **O**bject **O**riented **O**n**L**ine, the monitoring software of COMPASS, see [COO]

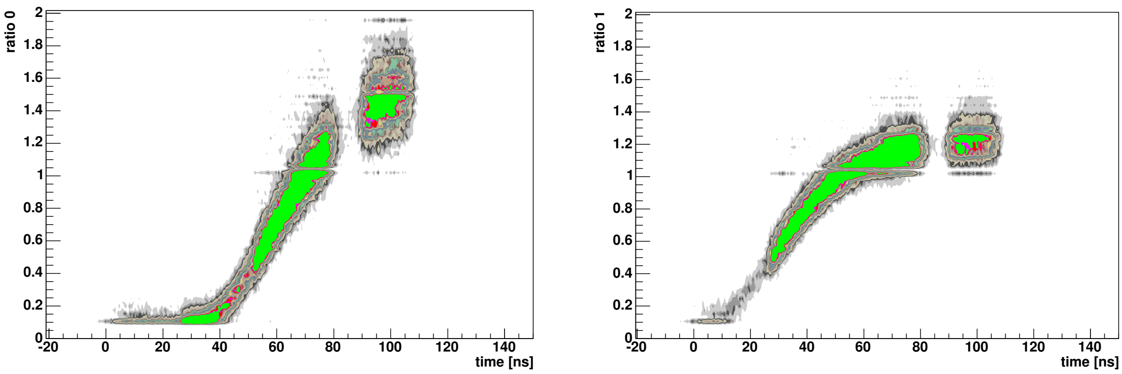

Figure 3.5: Relation of particle time and digit ratio r_0 (left) and r_1 (right), respectively. The zero point on the time axis is arbitrary. Note that the polarity of the time is opposite to the convention used in the text. Data are courtesy of [Din04].

function of time:

$$r(t) = r_0 \cdot \exp\left(-\exp\left(-s(t-t_0)\right)\right) \quad \text{with} \quad s(t') = \frac{a+c}{2} + \frac{a-c}{2}\left(\sqrt{t'^2+b^2} - b\right) + d \tag{3.21}$$

This parametrisation has the beneficial property of being invertible, so that the time may be reconstructed from a ratio as follows:

$$\begin{aligned} t(r) &= t_0 + s^{-1}\left(-\log\left(-\log\left(\frac{r}{r_0}\right)\right)\right) \\ \text{with} \qquad s^{-1}(r') &= \frac{1}{2ac}\left[(a+c)\cdot f - (a-c)\cdot\sqrt{f^2+acb^2}\right] \\ \text{and} \qquad f &= x - d + \frac{b}{2}\cdot(a-c) \end{aligned} \tag{3.22}$$

The error of reconstructing time in that way may be estimated starting from the uncertainty in the individual digits which is constituted by *noise* (and, to a lesser extent, quantisation error in the ADC). The noise level of the individual channels is easily accessible to measurement.[14] Thus an approximation for the error σ_a of the measured digit a is well-known, only needing to be read from file. By Gaussian error propagation, the *relative* error being the square root of the sum of the squares of the individual relative errors, the *absolute* error σ_r for the ratio $r_x = a_x/a_2$ follows to be

$$\sigma_r = \frac{a_x}{a_2}\cdot\sqrt{\left(\frac{\sigma_a}{a_x}\right)^2 + \left(\frac{\sigma_a}{a_2}\right)^2} = \frac{\sigma_a}{a_2^2}\cdot\sqrt{a_x^2+a_2^2}. \tag{3.23}$$

For calculation of time error σ_t, due to high curve of $t(r)$ it is more accurate to perform two evaluations of $t(r)$, determining an downward error σ_t^- and an upward error σ_t^+,

[14]Such measurements are conducted routinely to calibrate zero-suppression in the front-ends.

rather than utilising the Gaussian method, which approximates the slope of $t(r)$ with a constant.

$$\sigma_t^+ = t(r) - t(r + \sigma_r) \qquad \text{and} \qquad \sigma_t^- = t(r - \sigma_r) - t(r) \tag{3.24}$$

So much for determination of time and error from a ratio. However as there are two ratios to every hit, two sets of time and errors are reconstructed from every hit. Ratio r_0 yields t_0, σ_{t0}^+ and σ_{t0}^- whereas t_1, σ_{t1}^+ and σ_{t1}^- are calculated from r_1.

The overall strip time t—the final result—is calculated as average of t_0 and t_1, each weighted by its uncertainty in the direction of the other:

$$t = \frac{\frac{t_0}{\sigma_0^2} + \frac{t_1}{\sigma_1^2}}{\frac{1}{\sigma_0^2} + \frac{1}{\sigma_1^2}} = \frac{t_1\sigma_0^2 + t_0\sigma_1^2}{\sigma_0^2 + \sigma_1^2} \tag{3.25}$$

The uncertainty of t_0 *in direction of* t_1 is indicated by σ_0, and σ_1 denotes that of t_1 towards t_0. The asymmetric errors σ_t^+ and σ_t^- of the combined time are calculated by Gaussian error propagation separated by direction: the overall upward error as combination of the two individual upward errors, and likewise for the total downward error.

$$\sigma_t^+ = \frac{1}{\sqrt{\frac{1}{(\sigma_{t0}^+)^2} + \frac{1}{(\sigma_{t1}^+)^2}}} \qquad \text{and} \qquad \sigma_t^- = \frac{1}{\sqrt{\frac{1}{(\sigma_{t0}^-)^2} + \frac{1}{(\sigma_{t1}^-)^2}}}, \tag{3.26}$$

For judgement of the results of reconstruction, the calculated time is drawn relative to the time of the trigger. The results of time reconstruction may be viewed in Figure 3.6 which compares the timing peaks of two sides of the same wafer. The time period around trigger time in which the silicon detectors are sensitive may be induced from the relation of reconstructed time and reconstructed error as exhibited in Figure 3.7. From the plot may be deduced that with good accuracy, time may be measured in the interval between approximately -10 and +40 (up to 40 ns after trigger and up to -10 ns before trigger). For times much earlier than trigger time, the errors may be observed to increase due to decay of the signal. Times much later than the trigger may not be recorded at all since at the moment of measurement they have not yet traversed the detector.

3.6 Clustering

As already mentioned in earlier sections, the design of the COMPASS silicon detectors is optimised for spatial resolution. The charge induced by the passage of a particle usually is distributed over more than one readout strip, resulting in signals in several adjacent channels. Using *clustering*, the information of all strips associated with one particle may be combined, leading to a single cluster which is defined in space and time more precisely than the individual hits.

While it is beneficial for general operation of Cinderella to improve the time resolution of the silicon detectors by clustering, for cuts on track multiplicity—as employed for

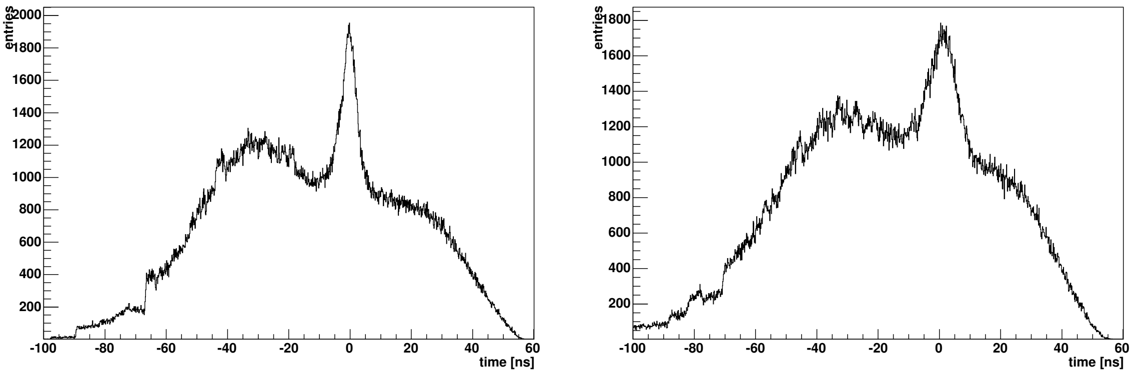

Figure 3.6: Time reconstructed from strips of both sides of the U/V wafer of SI01 with respect to trigger time. Clearly visible is the time resolution of the n-side plane SI01U1 (left plot) being superior to that of the p-side plane SI01V1 (right plot) due to the intermediate strips of the latter. (As the underlying data were not subject to any cuts, this plot should not be compared directly with plots from other sources that have been refined in some way.)

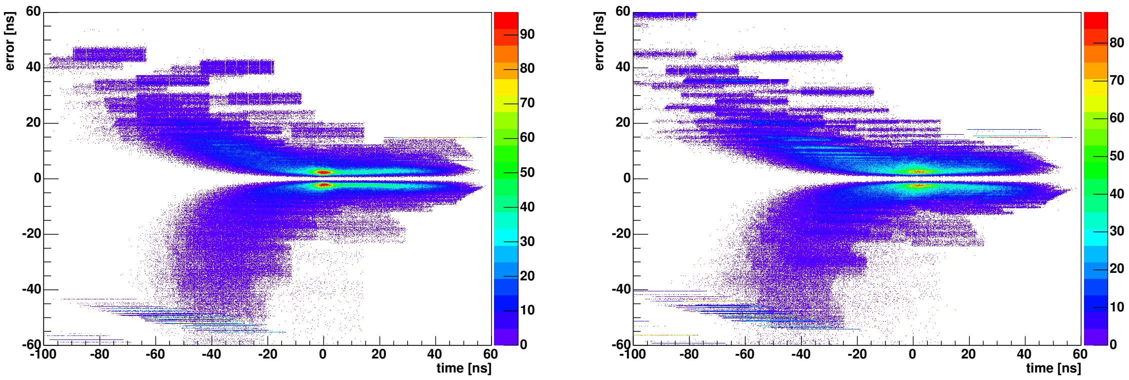

Figure 3.7: Reconstructed strip time plotted against its error (left: n-side, right: p-side). Time is plotted on the abscissa, its errors on the ordinate. The reconstructed error being asymmetric, for every reconstructed strip time, two errors are plotted: the error towards later times in positive direction and that towards earlier times in negative direction of the ordinate. Consequently the errors in direction of the trigger time are contained in 2nd and 4th quadrant. Just as in Figure 3.4(a) the line artefacts result from division of small integers.

filtering of hadron beam data—it is imperative to merge multiple hits originating from the same particle into one cluster so that the number of clusters in a plane equals the number of tracks of charged particles through it.

Due to the operation of the silicon detectors in high-rate conditions in the centre of the beam, the greatest difficulty for a clustering algorithm is to separate "legitimate" clusters from cases in which hits to adjacent channels stem from two different particles, which may or may not have the same timing.

In general, clustering is a very complex optimisation problem. Depending on application of the output, the two possible errors of wrongly separating or wrongly combining adjacent hits have an impact of different gravity. For the online filter, additionally a compromise between speed and correctness needs to be made. As a consequence, only comparatively simple heuristic may be employed.

The first part of any clustering algorithm is to find sets of neighbouring activated channels, which are identified as *cluster candidates*. (There are subtleties regarding as to the treatment of dead strips, which cannot be addressed in Cinderella.)

As the clustering algorithm has been designed for use in muon beam data taking (where all silicon stations are positioned upstream of the target), the initial assumption has been made that the hits that are to be separated result from two different beam particles, being improbable to coincide (the average time difference of two succeeding beam muons being 17 ns at $\sim 2.8 \cdot 10^8$ muons per spill).

Then three cases need to be differentiated: The first—most likely—is that both tracks are too far apart to even form a linked neighbourhood. In this case no action is taken by the clustering algorithm. The second possibility is the two tracks approaching close enough to impact adjacent strips: In that case, the cluster candidate should be separated at the point in the neighbourhood with time inconsistency between two adjacent strips. The third case is that of the tracks affecting overlapping strips. At that moment, two consecutive signals are induced on the same strip, piling up in the truest sense of the word. As time reconstruction relies on signal shape, the two ratios are expected to yield two inconsistent times. Thus under that circumstance, the cluster candidate should be separated at the position of time inconsistency between the two ratios of the same strip.

A consideration of legal cluster sizes leads little way, unfortunately. There is no upper limit to cluster size that could be introduced sensibly because it is not uncommon for knock-on electrons created by particle passage to travel many strips inside the silicon wafer, creating large clusters; an incidence that is promoted by the kinematics of the reaction, demanding an angle between muon and electron track that is close to perpendicular. On the other side, clusters only one strip wide are perfectly normal, too.

A fast algorithm is constructed following these considerations by evaluating every *boundary* between adjoining strips for *inconsistency* and divide the cluster at the point(s) where the inconsistency exceeds a certain threshold. The inconsistency ζ is taken to be the weighted average of the intra-strip inconsistencies of the two adjoining strips (weight 1 each) and the inter-strip inconsistency between the two strips (weight 2). For purpose of this algorithm, the inconsistency $\Upsilon(t_1, t_2)$ of two times t_1 and t_2 with errors σ_1 and σ_2 is

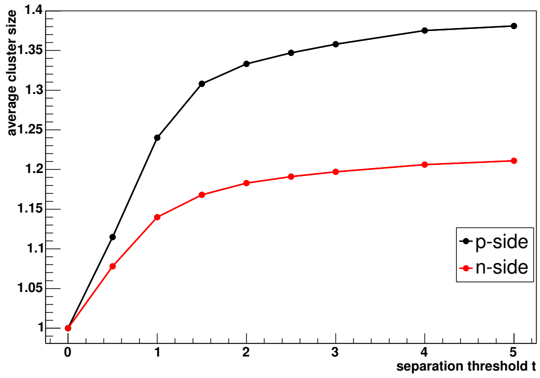

Figure 3.8: Average cluster sizes in dependency of separation threshold for n and p sides.

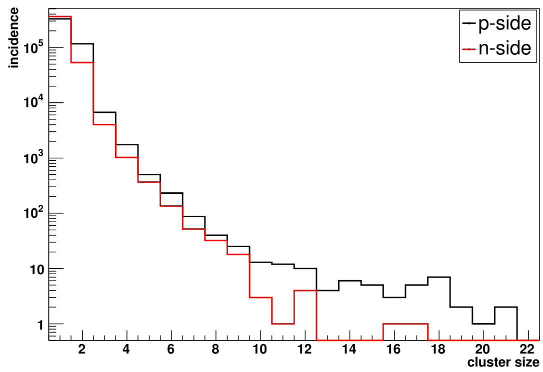

Figure 3.9: Cluster size distribution for the default separation threshold of 1.5 for both sides of the detector plane SI01U/V.

defined as

$$\Upsilon(t_1, t_2, \sigma_1, \sigma_2) = \left| \frac{t_2 - t_1}{\sqrt{\sigma_1^2 + \sigma_2^2}} \right|, \tag{3.27}$$

which in light of eq. 5.1 is to be interpreted as the smallest confidence level on which the two times agree. The inconsistency criterion $\zeta(a, b)$ between adjacent strips a and b from that follows to be

$$\zeta(a, b) = \frac{1}{4}\Upsilon(t_0^a, t_1^a, \sigma_0^a, \sigma_0^a) + \frac{1}{4}\Upsilon(t_0^b, t_1^b, \sigma_0^b, \sigma_1^b) + \frac{1}{2}\Upsilon(t^a, t^b, \sigma^a, \sigma^b), \tag{3.28}$$

where t_0^a, t_1^a, σ_0^a and σ_1^a denote the time and error information gained from ratios r_0 and r_1 of strip a and the same symbols indexed with b denote times and errors of strip b in an analogue way. The *combined* time end error information of the two ratios of strip a is indicated as t^a and σ^a whereas the same of strip b is referred to as t^b and σ^b.

Cluster candidates then are separated at the positions where the inconsistency of two neighbouring strips exceeds the configured threshold t. The dependency of average cluster size on the threshold t is pictured in Figure 3.8. The average cluster sizes yielded by this algorithm are smaller than that determined through other means since strips that have an amplitude too low for time reconstruction are not incorporated into clusters.

Having been designed initially for discrimination of beam particles and pile-up in the beam telescope, this algorithm works reasonably well for separation of tracks of different timings. Yet for two tracks with identical timing, as it is custom after the target, the algorithm has a tendency to combine them into one cluster. An estimation of the minimal distance for two *coincident* tracks not to be combined into one cluster is given in section 6.4.

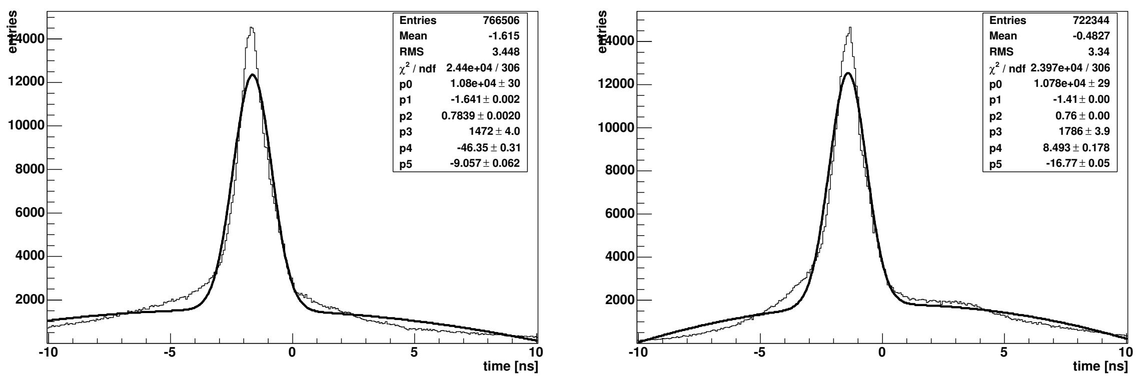

Figure 3.10: The shapes of the pulls (n-side: left plot, p-side: right plot) of the reconstructed cluster time (without cut) resemble Gaussian distributions with quadratic background and were fitted as such. Their sigmas are denoted by parameter *p2* and approximately equal 0.8. The displacements of the peaks, approx. -1.5, are specified by parameter *p1*.

3.7 Results of Silicon Time Decoding

To verify the agreement of calculated cluster time and error, the quantity t/σ_t is plotted in Figure 3.10, usually referred to as *pull*. The uncertainty σ_t is taken as average of the asymmetric uncertainties σ_t^+ and σ_t^-. Ideally a normalised Gaussian distribution should be the result. While the plots certainly do not show a perfect Gaussian distribution, they are not too far away, either. Although there is some background and a displacement of the peak of approx. -1.5 ns, the fitted sigmas with approx. 0.8 are close to the expected value of 1. To a certain degree this result may be interpreted as a confirmation, while on the other hand further improvements are encouraged.

After time reconstruction and clustering, a cut is applied on the inconsistency (in the meaning of eq. 3.27) of the reconstructed time t with trigger time (which is taken to be zero and with zero uncertainty). Thus the cut is expressed as

$$\Upsilon(t, 0, \sigma_t, 0) \leq c, \tag{3.29}$$

where σ_t denotes the error of t in the direction of trigger time and c is the confidence level on which a reconstructed time cluster needs to be consistent with trigger time in order not to be rejected. For the muon beam data taking of 2004 this confidence level was identical to that specified as c in section 5.3.2. (However since then, Cinderella has been flexibilised to allow two different values as there is to specific reason to tie together these two “confidence levels”.) The evolution of the timing peaks from reconstructed *strip time* over *cluster time* towards *cluster time after cut* is exhibited in Figure 3.11.

Figure 3.12 shows the time resolution of the silicon time after decoding and cuts. It has been fitted with a Gaussian distribution plus quadratic background. The uncertainty of time measurement is determined to be in the range of 4–5 ns.

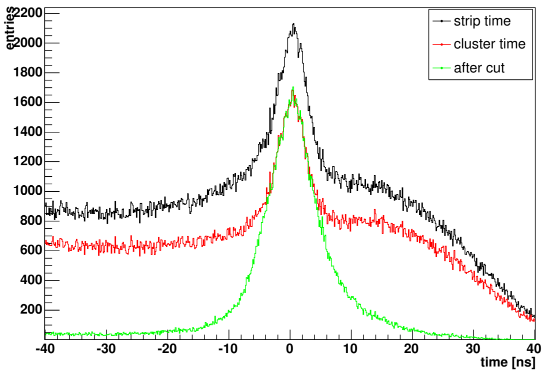

Figure 3.11: Evolution of the silicon time peak through the stages of processing for n-side (left plot) and p-side (right plot): The black curve shows the timing peak as decoded from the individual strips, the red curve displays the same peak after clustering, and the green line shows the timing peak after clustering and cut.

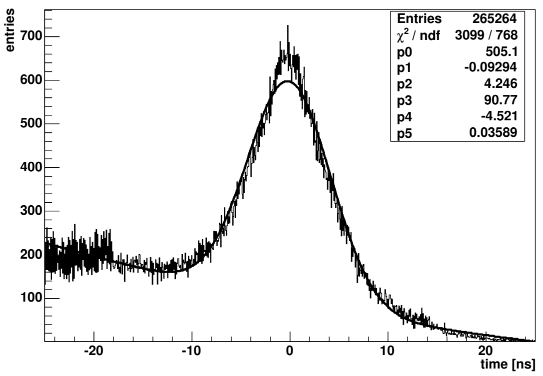

Figure 3.12: The time resolution (after cut) of silicon decoding for n-side (to the left) and p-side (to the right) fitted with Gaussian distribution plus quadratic background; its sigma is denoted by parameter *p2*.

Chapter 4

General Concepts of Online Filtering at COMPASS

4.1 Concepts of Data Reduction

Generally, the goal of data reduction can be accomplished by two fundamentally different approaches: By employing lossless compression, redundancies in a data stream are removed while completely retaining the inherent information, so that at any time the original data stream may be reconstructed from the compressed stream. Methods most commonly applied are exploiting sequences of identical characters, recurring patterns, or inequalites between frequency of occurence of different characters (eg. Huffman entropy coding, described in [Huf52]). Lossless compression today is in common use to reduce size of texts of all kinds (including textual representations stored in databases and even executable machine code) and there exist many generic, *all-purpose* algorithms. Well known representatives of lossless compression software are gzip, bzip2 and ZIP. CASTOR, too, makes use of lossless compression *transparently* in the firmware of the tape drives.

The opposed mechanism is lossy compression, where an algorithm tries to differentiate between "important" and "dispensable" aspects of the data, only including the former in the compressed stream while *unrecoverably* ommitting the latter. To achieve this, the algorithm needs a certain degree of "understanding" or "comprehension" of the data that it is handling. Inherent to every mechanism of lossy compression consequently is a *model* of the data it is working on. The higher the data reduction that is desired, the higher sophistication is required for the model. (The last two decades have seen extensive research in the field of lossy audio and video compression, leading to development of advanced *psycho-acustic* and *psycho-visual* models. The existing knowledge of the properties and especially the limitations of human sense organs and the processing of their signals in the brain is exploited to reduce data rates *without being noticed by the audience*.)

The advantage of that approach is the achievement of greatly increased compression rates compared to lossless compression methods. The drawbacks however are that development of the *highly specialised* models usually is a elaborate task where the obtained compression rate needs to be weighted carefully against the amount of information rejected. Lossy compression algorithms are in widespread use for sound, image and video compression, often referred to as *codecs*[1]. Familiar examples are JPEG image compression, MPEG video (".mpeg") compression, and MPEG Audio Layer-3 compression[2] (".mp3").

Given the characteristics of the two concepts, the choice of *lossy compression* for use in Cinderella is evident. The demanded reduction rates cannot be reached with lossless compression alone, which already is being used at CASTOR anyway. However a combination of lossy and lossless methods is conceivable for future developments.

4.2 Event-Based Filtering

The key feature of lossy compression is to distinguish "important" from "dispensible" data, plainly spoken, to separate the wheat from the chaff. Yet, this is a highly abstract view. For an actual implementation first it is necessary to constitute the definition of the smallest unit of data which then may be filed in one of the two categories. While other choices are conceivable,[3] initially it seemed most rewarding to define the *event* as the smallest unit of data upon whose fate the online filter has to decide.

With that approach, the online filter may very well be characterized as 2^{nd}-level trigger that serves to clean up the impurities of the 1^{st}-level triggers. As for any trigger, the notions of purity and efficiency as fundamental parameters of performance apply identically to the online filter. Consequently the *efficiency* of the online filter is the ratio of *recognised* useful events to all useful events. The efficiency of the online filter usually being close to 1, it is generally more convenient to quote the complementary inefficiency I instead:

$$I = \frac{N_{rej}^{good}}{N^{good}} \tag{4.1}$$

[1] **co**der-**dec**oder

[2] Developed by Fraunhofer Institut Integrierte Schaltungen [FII] and standardised by ISO.

[3] Another possible choice would be a *digit*. Online-filtering of digits is being used widely across the whole experiment, commonly referred to as *zero suppression*. The topological argument strongly speaks in favour of implementing this functionality close to the individual detectors. That way, data reduction happens at the earliest possible position in the readout chain, easing the load on the following links and processing stages.

Being, in its simplest form, nothing more than enforcement of a threshold for detector channels, zero suppression is comparatively easy to implement in electronics. But the greatest advantage of this mechanism also turns out to be its greatest deficiency: The lack of knowledge of even so much as the digits of the neighbouring channels restricts the utility of the procedure. In the name of separating detector hits from underlying noise, often the valuable tails of the collected charge distributions are cut off, worsening resolution substantially.

The online filter has the benefit of being capable of access to all digits of all detectors without requiring hardware modifications and thus could mitigate the aforementioned problem. Consequently it was being discussed in earnest to make use of Cinderella for zero suppression of calorimetry digits. By considering neighbouring digits (even across two different types of readout) for zero suppression Cinderella would be able to increase the ECAL2 energy resolution in parallel to its main task, the event-based filtering. The plan was shelved due to more urgent tasks at hand, but it never was abandoned.

In this connexion, the issue of *definition* of what constitutes an useful event—and more delicately, of *what does not*—is deferred to chapters 5 and 6 which deal with the subjects that are specific to a certain beam type.

The other fundamental parameter of performance is the *reduction ratio*, which either may be defined as the ratio of rejected bad events to total events ("total reduction ratio") or as the ratio of rejected bad events to all bad events ("bad event reduction ratio"):[4]

$$R_{tot}^{bad} = \frac{N_{rej}^{bad}}{N} \quad \text{and} \quad R_{bad}^{bad} = \frac{N_{rej}^{bad}}{N^{bad}} \quad \text{with} \quad R_{bad}^{bad} = R_{tot}^{bad} \cdot \frac{N}{N^{bad}} \tag{4.2}$$

Both definitions exist in their own right. The first is easier to communicate and may be used for rate calculations to determine overall trigger setup, whereas the second, being independent of (sub-)trigger impurity, has a distinct advantage when comparing Cinderella performance of different sub-triggers. An effectively equal formulation is that of acceptance ratio:

$$A_{tot}^{bad} = \frac{N_{acc}^{bad}}{N}, \quad A_{bad}^{bad} = \frac{N_{acc}^{bad}}{N^{bad}} \qquad \text{with} \quad A_{bad} + R_{bad} = 1 \quad \text{and} \quad N_{rej} + N_{acc} = N \tag{4.3}$$

Naturally, the aim is for the online filter to attain high efficiency and high reduction ratio *at the same time*. Unfortunately these two parameters are complementary in a sense that an effort undertaken to increase any one of them usually will decrease the other. As a consequence, apart from finding a good selection criterion to differentiate between useful and undesired events, it is also necessary to tune the free parameters of that criterion in such a way as to balance trigger efficiency and veto efficiency to suit the requirements of the data taking situation.

4.3 Considerations of Implementation

The decision to discard a certain event in the online filter always is final and cannot be undone in a later stage of data processing. This is a fundamental difference with regards to "offline" reconstruction or analysis software, where errors can be (and in fact often are) corrected by fixing the problem in source code and then re-running the program. As a consequence, Cinderella is *designed* and *implemented* with the goal of highest reliability.

To ease studies of the properties of Cinderella, every n^{th} event of the input stream is always included with the output, regardless of the filter decision.[5] This *clean sample* or *monitoring sample* is indispensable for many types of analysis, not only to search for biases that might be introduced by the online filter but also to determine the inefficiency of Cinderella. (Physics analysis with the clean sample is explained in chap. 7.5.)

[4] It may be argued that in the view of the online filter as a 2$^{\text{nd}}$-level trigger, the parameter corresponding to efficiency would be *purity*. However it has proven to be more handy and intuitive not to strain the analogy too much but rather use a definition in this fashion.

[5] In all of 2004 n was configured to be 30.

The online filter is adapted to the COMPASS trigger system, consisting of partly overlapping sub-triggers ("triggers" for short), which trigger on different physical reactions or different kinematic regions of the same reactions. It also is possible, that multiple sub-triggers are activated for the same event. Thus in Cinderella it is possible to specify per sub-trigger configuration, allowing to run two completely different algorithms for two different sub-triggers. In the case of multiple sub-triggers being activated, a well defined order of precedence takes effect and the event is dealt with using the settings for the sub-trigger with the highest precedence.

The environment that Cinderella is running in currently is constituted by the 13 dual-processor event builder computers[6] whose 26 CPUs have to be shared with other processes. Depending on trigger rate and load caused by other components of the DAQ system, the time budget that is available to Cinderella is in the range of only 4–8 ms of CPU[7] time *per event*. Consequently, Cinderella is restricted to perform only fast, basic operations like decoding and correlating time information of *some* detectors. Track or even vertex reconstructions are out of scope for the time being.

[6]The 13$^{\text{th}}$ computer used to be a hot spare but now is used in production.
[7]AMD Athlon MP at 1.6 GHz

Chapter 5

Filtering of the Muon Beam

Since the muon programme takes a predominant part of COMPASS beam time and a large share of the overall data volume accordingly, the online filter was initially designed for filtering of muon beam data.

Before concepts of an algorithm for online event categorisation are developed, it seems appropriate to give some thought to the topic of the evaluation of the algorithm. Doubtless, a thorough understanding of its properties is imperative before deployment may be considered. The decisions of the online filter need to be verified. Yet an universal gauge does not exist. The utility of any given event is subject to opinion, after all.

For the muon programme, this situation is resolved by resorting to some kind of minimal consensus, that already has been established. All physics analysis is being done using the mDST output files of the CORAL reconstruction software, which has its very own criteria for inclusion of events. Consequently, for purpose of evaluation of the online filter, the definition of a "good" event is based on the requirement of it being included in the mDST by CORAL. The only additional prerequisite is that a beam track must have been reconstructed for that event, a condition that may be applied without concern because all physics analysis of muon beam data at COMPASS crucially depends on the knowledge of track and momentum of the incoming muon.

Yet the outcome of CORAL needs to be taken with a grain of salt. The past has shown that with ongoing development the efficiency of CORAL has improved, and events that formerly could not be reconstructed now may be used for physics analysis very well.[1] Following these considerations, the result of CORAL is not trusted blindly for evaluation of Cinderella. It is taken for what it is: a very useful indicator of the quality of the filtering algorithm, but not the definitive benchmark for its judgement.

Now regarding the fact, that the amount of reconstructed events in CORAL is lower than the total amount of events recorded by a factor of approx. 3, the potential for data reduction can be estimated to be large, the reduction ratio bounded only by $\sim 2/3$ of

[1] The work on beam track reconstruction by M. v. Hodenberg [vH04], that directly impacted the muon filtering scheme, may serve as a prominent example.

overall data. However since physics analysis of COMPASS muon data to large parts depends on the study of asymmetries in distributions of different physical properties, it is of paramount importance to avoid biasing the distribution of *almost any* physical property; a requirement which limits the data reduction ratio that may be reached effectively.

5.1 Criterion: Beam Track Reconstructabilty

The constitution of a filtering criterion represents the establishment of a model of COMPASS event data that may be interpreted by an algorithm (cf. section 4.1).

One idea for reduction of muon data is to reject events which do not contain enough information to reconstruct the beam track. Since all physics analysis of muon beam data depend on reconstruction of the primary vertex (point of interaction between muon and polarized target), it can be safely assumed that events from which no beam track can be reconstructed are useless for that purpose.

As a second thought, any physics analysis also depends critically on knowledge of the momentum of the beam muon (due to momentum spread of the beam of $\sim$5% this value has to be measured individually for every beam particle), so that events that do not contain enough information to reconstruct the beam momentum may be filtered out, too, without losing physics data.

Third, the time information of the beam particle has to match the trigger time to ensure that the trigger really was initiated by a beam muon (and/or its reaction).

Selected elements of the spectrometer setup[2] before the target that play a role for setup of Cinderella are shown in Table 5.1. For beam track definition there are 12 planes of silicon micro-strip detectors in 4 different projections ((U,V,X,Y) $\times$ 3) and 4 planes of scintillating fibres in 2 different projections ((X,Y) $\times$ 2). (The (U,V) coordinate system, with respect to the (X,Y) system, is rotated by 5° around the beam axis.) The beam momentum is determined in the beam momentum station (BMS) by measurement of y-coordinate in 6 planes of scintillator, of which 3 are before and 3 are after the vertical bending magnet B6, close to the end of the M2 beam line.

Following the initial considerations, the geometric minimum for beam track definition is four *coincident* hits in four different detector planes of which not more than any two may be of the same projection (U, V, X or Y). Presuming the beam track before the target is defined, the geometric minimum for momentum reconstruction would be one *coincident* hit in the BMS upstream of the bending magnet B6. However to reach the desired accuracy in the beam momentum, M. v. Hodenberg [vH04] has shown that at least one additional hit anywhere in the BMS is necessary.[3]

With an efficiency of 99% for silicon detectors and an efficiency of 83-92% for BMS hodoscopes [vH02], for an event with beam track the *expected* amount of hits are $\sim$15 out

[2] taken from detectors.34930.minus.dat

[3] This is the minimal requirement for which his "rescue-algorithm" is able to reconstruct a beam particle in CORAL.

Position [m]	TBName	Type	Function
-137.2	BM01P1	sc. slabs	beam momentum
-130.6	BM05P1	sc. fibre	beam momentum
-123.8	BM02P1	sc. slabs	beam momentum
-104.6 ... -85.7	B6	magnet	beam guide
-73.7	BM03P1	sc. slabs	beam momentum
-70.8	BM06P1	sc. fibre	beam momentum
-61.3	BM04P1	sc. slabs	beam momentum
-7.60	FI01X1	sc. fibre	beam position
-7.58	FI01Y1	sc. fibre	beam position
-5.69	SI01U1	silicon	beam position
-5.69	SI01V1	silicon	beam position
-5.68	SI01Y1	silicon	beam position
-5.68	SI01X1	silicon	beam position
-5.01	SI02U1	silicon	beam position
-5.01	SI02V1	silicon	beam position
-5.00	SI02Y1	silicon	beam position
-5.00	SI02X1	silicon	beam position
-3.53	SI05U1	silicon	beam position
-3.53	SI05V1	silicon	beam position
-3.52	SI05Y1	silicon	beam position
-3.52	SI05X1	silicon	beam position
-2.88	FI01X1	sc. fibre	beam position
-2.88	FI01Y1	sc. fibre	beam position
-1.00 ... -0.40	upstream cell	^{6}LiD	polarised target
-0.30 ... +0.30	downstream cell	^{6}LiD	polarised target

Table 5.1: Selected elements of the spectrometer setup upstream of the target.

of 16 planes in the beam telescope and $\sim$5 out of 6 in the BMS. Yet for an event without beam track, due to noise and pile-up a hit count substantially larger than zero can be assumed for BMS and beam telescope. These considerations suggest that the criterion in plane multiplicity for the optimal separation of events with and without beam track resides somewhere in between the absolute geometric minimum and the expectation of close to full efficiency.

5.2 The Conditional Coincidence Algorithm

In the last section it was concluded, that for reconstruction of beam track and beam momentum a certain amount of hits would be necessary to several planes of BMS and beam telescope. Furtheron the timing of these hits would need to be consistent among themselves as well as with the time of the trigger. An event failing to fulfil this requirement cannot be reconstructed in CORAL an thus may be discarded by Cinderella.

These considerations constitute the demands to a *conditional coincidence algorithm*: A flexible condition comprised of many planes need to be evaluated with the constraint of pairwise coincidence of all partaking hits. Talking plainly, the algorithm needs to be able to generate an answer to a freely configurable question along the lines of: "Was there a point in time when from the group of planes X at least k showed a signal and at the same time from the set of planes Y at least m were active and also at the same time out of the series of planes Z at least n saw a particle... ?" The structure of an algorithm that is capable of that is described in the following.

The concept of coincidence is closely related to the uncertainty of the measured time, i.e. to the time resolution of the participating detectors. Two times t_1 and t_2 with errors σ_1 and σ_2 are defined to be consistent on a confidence level c if their difference is consistent with zero, i.e. its absolute value is smaller than c times the Gaussian error of the computed difference:

$$|t_2 - t_1| \leq c \cdot \sqrt{\sigma_1^2 + \sigma_2^2} \tag{5.1}$$

Unfortunately, due to the nonlinearity of this condition, to determine coincidences for n hits, $\mathcal{O}(n^2)$ evaluations of the inequality 5.1 need to be performed (every hit compared to every other hit).[4] A rough calculation shows that for 22 planes (6 from BMS, 4 tracking fibres and 12 silicons), depending on previous cuts, the number of hits easily may exceed 100 in regular conditions and occasionally (beam instability, noise spikes) may even be much higher. Thus for the sake of stability, quadratic complexity *in the number of hits* cannot be tolerated.

[4]The condition of inequality in this form lacks *transitivity*. That is, for three hits t_1, t_2 and t_3, if t_1 is consistent with t_2, and t_2 is consistent with t_3, nothing may be learned from that regarding to the consistency of t_1 and t_3. It especially may not be reasoned that t_1 then is consistent with t_3. As there is no shortcut through transitivity, a complete evaluation of consistency cannot avoid calculating inequality 5.1 once for every pair of time points, leading to a complexity of $\mathcal{O}(n^2)$.

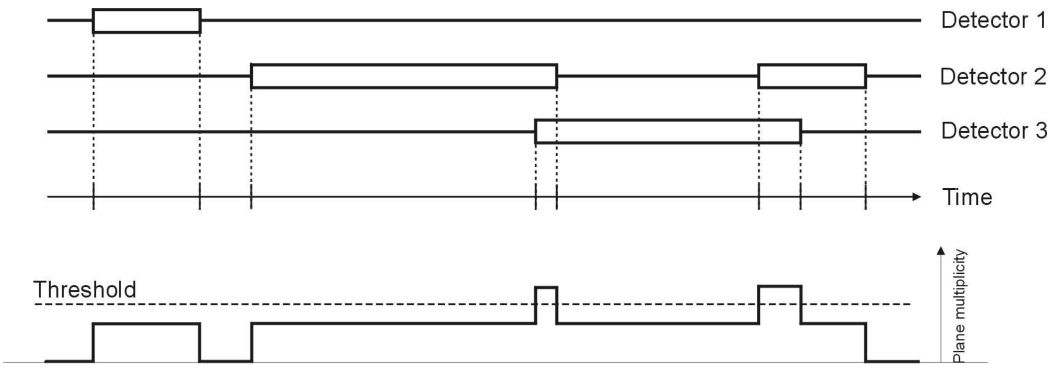

Figure 5.1: For illustration of the fast algorithm for conditional coincidence a simple example is chosen, consisting of three detector planes for which the criterion "two out of three" shall lead to accepting of the event. A width corresponding to the uncertainty of their time determination is assigned to the individual hits (symbolised as boxes) to these planes. The boundaries of these time intervals then are projected on the time axis, which represents the sorted list. When the sorted list then is gone through from left to right, at every boundary marking the start or the end of a plane hit, the multiplicity counter is increased or decreased, respectively. As soon as the counter exceeds the configured threshold, the criterion has been satisfied, the event is accepted and the processing stops immediately. However if the end of the list is reached without exceeding the threshold, the event is rejected.

However, using the triangle inequality, it is possible to find an upper limit to the right side which linearises the relation:

$$t_2 - t_1 \ \leq c \cdot (\sigma_1 + \sigma_2) \tag{5.2}$$

Assuming w.l.o.g. that $t_1 \leq t_2$, the simple condition for coincidence

$$t_2 - c \cdot \sigma_2 \leq t_1 + c \cdot \sigma_1 \tag{5.3}$$

is yielded.

Now that separation of variables was successful, a quicker algorithm of $\mathcal{O}(n \log n)$ may be employed. The individual hits may be seen as having an extension or *width* of $2{\cdot}c{\cdot}\sigma$ with *boundaries* $t - c{\cdot}\sigma$ and $t + c{\cdot}\sigma$. Calculation of boundaries for all hits and sorting of the results (using one of many well known algorithms of $\mathcal{O}(n \log n)$, see [Bra88]) yields a list from which determination of coincidences is possible in linear time[5] so that the overall complexity of the algorithm is $\mathcal{O}(n \log n)$, which is a huge improvement over $\mathcal{O}(n^2)$ of the initial approach. This algorithm is illustrated in Figure 5.1.

[5]When traversing the sorted list of boundaries, a kind of *transitivity* may be used to speed up calculation. If the boundaries of t_1 are smaller than those of t_2 which in turn are smaller than these of t_3 it may be implied that t_1 is inconsistent with t_3 without further calculation:
$t_1 - c\sigma_1 < t_1 + c\sigma_1 < t_2 - c\sigma_2 < t_2 + c\sigma_2 < t_3 - c\sigma_3 < t_3 + c\sigma_3 \quad \Longrightarrow \quad t_1$ and t_3 are not consistent

The drawback, though, is some sacrifice in accuracy. While the linearisation is legal in a mathematical sense, it has the effect, that hits may be mistaken as coincident on the chosen confidence level while in fact, they are not. While such an effect is undesired, it may be tolerated because it "errs on the right side". It may lead to an event being wrongly accepted but not to one being wrongly discarded. The maximal inaccuracy occurs for $\sigma_1 = \sigma_2$. Then the minimal confidence level at which two hits could be seen as coincident is underestimated by a factor of $\sqrt{2}$. In the limit of $\sigma_1 \ll \sigma_2$, the error becomes negligible.

With a high-performance mechanism for coincidence established, it is easy to add the conditional part: While iterating over the sorted list of hit boundaries, a logic structure is updated at every step. The overall condition then is evaluated as true, if there is any point in time that has enough overlap of hits to fulfil the condition.

This conditional coincidence algorithm is the work-horse of the online filter. Its input is taken from decoding modules for scintillators (fibres and slabs) and for silicon (cf. section 3.5). For scintillators, the uncertainty of time measurement is determined by automatic calibration (cf. section 3.5) while for silicon detectors for every single hit the time error is calculated individually by the decoding module.

The free parameters of this algorithm, constituted by the logic condition and the confidence level c, need to be carefully tuned to achieve the desired functionality, as is described in the next section.

From the three considerations for the filtering criterion of section 5.1, the first two—ability to reconstruct beam track and beam momentum—are covered by the algorithm introduced in this chapter. The third requirement, coincidence of detector time and trigger time, easily could be integrated with it. However currently the third requirement is handled by a separate cut in an earlier processing stage, allowing for some more specialised functionality.

5.3 Configuration

5.3.1 Method of Comparison to Reconstructed Data

With the conditional coincidence module described in the last section, the tool is at hand for filtering muon beam data according to the criterion defined in section 5.1. The only thing that is left to do is assembling the parts and finding a useful set of configuration parameters.

In spring 2004, the COMPASS experiment already could look back on two beam times of physics data taking, of which 2002 data already was processed in full with CORAL and 2003 data was processed partially. After the first physics runs had been taken in 2004, a fast pre-production was started, so that even some processed runs of 2004 were available quickly.

A promising approach to finding a good parameter set is to first devise a sensible scheme, fixing some "obvious" parameters and setting rough boundaries for others. Then the

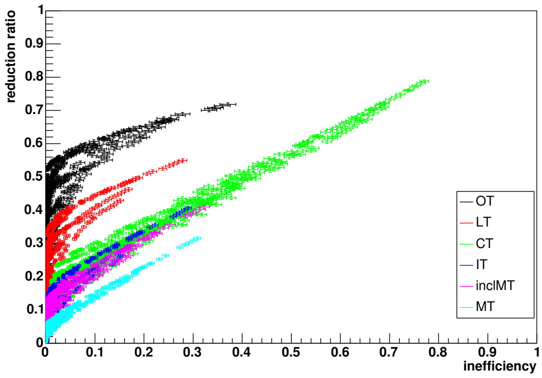

(a) All tested variations of configuration.

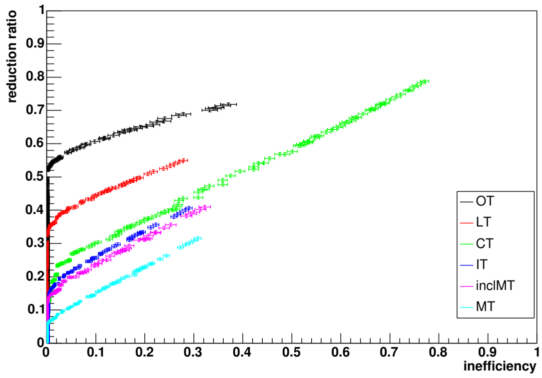

(b) After exclusion of all settings which unambiguously can be categorised as sub-optimal.

Figure 5.2: For all variations of parameters the yielded reduction ratio is drawn against the resulting inefficiency. The quality of a parameter setting increases from the lower right (low reduction ratio, high inefficiency) to the upper left (high reduction ratio, small inefficiency).

filter algorithm is tested on already recorded raw data many times with many different variations of parameters. With CORAL results serving as a benchmark being readily available, the performance of Cinderella (reduction ratio and efficiency) is computed for every of these test runs. From the resulting variety of results, that parameter setting finally is chosen which has produced the most beneficial reduction ratio *and* efficiency.

To select the optimal parameter variation among the many that have been tried, it is useful to plot these two variables in relation, as shown in Figure 5.2(a) for every sub-trigger. Depending on the type of analysis that is envisioned for a certain sub-trigger, a compromise between reduction ratio and efficiency needs to be made.

As the choice of the desired parameter set contains a trade-off between reduction ratio on the one side and efficiency on the other side, generally one setting cannot be preferred to another *for all purposes*. However for *some* of the settings it holds true that they are less optimal than others in both aspects (i.e. reduction ratio *and* efficiency). As a consequence, they are not considered any further. Having removed all those unequivocally undesired parameter sets, as exhibited in Figure 5.2(b), the plot clears up considerably.

The parameters to the Conditional Coincidence Module that have been varied to generate the data points of these plots are listed in Table 5.2. The table is to be interpreted in the following fashion: For every row, the group listed in column *Elements* needs to coincidently satisfy the condition of *Criterion*. Yet for the whole event to be considered useful by Cinderella, the conditions of all rows *coincidently* need to be satisfied. Beginning at the top of the table, this means: For an event to be accepted by Cinderella, it need to show a hit to at least 1 out of the 3 planes BM1, BM5 and BM2. Yet at the same time, it needs to show a hit in at least 2 out of the 6 planes BM1–BM6. Again coincidently, it needs to show a hit in at least L out of the 3 planes SI1U, SI2U and SI3U, and so on... However there is a speciality: The four “Tracking” criteria that serve to ensure proper beam defi-

Group	Elements	Criterion	Type
BMS (before B6)	BM1, BM5, BM2	1 out of 3	fixed
BMS (complete)	BM1–BM6	2 out of 6	fixed
Tracking: U proj.	SI1U, SI2U, SI3U	L out of 3	var.
Tracking: X proj.	FI1X, SI1X, SI2X, SI3X, FI2X	M out of 5	var.
Tracking: V proj.	SI1V, SI2V, SI3V	L out of 3	var.
Tracking: Y proj.	FI1Y, SI1Y, SI2Y, SI3Y, FI2Y	M out of 5	var.
Meta-Group	U-proj., X-proj., V-proj., Y-proj.	N out of 4	var.

Table 5.2: Overview of configuration options that have been tested for the *Conditional Coincidence* module. While some have been tried in variations, others have been fixed beforehand to reduce the amount of calculations. (Please refer the main text for further explanation.)

nition are relaxed by a meta-criterion. Differing from the requirement that the condition of every row of the table needs to satisfied, the “Meta-Group” defines the criterion that out of the 4 “Tracking” groups only N need to be satisfied concurrently for the event to be accepted.[6]

Additionally to the free parameters L, M and N of the previous description, the confidence level c plays an important role for the overall filtering criterion. As a constituting part of the *definition of coincidence* in eq. 5.1 and its simplification in eq. 5.3 it defines the confidence level for every single coincidence computation, effectively specifying the *hardness of every single time cut* in multiples of σ_t. (Though this analogy may look like an invitation to do it, the confidence level c may not be taken as the “single all-important” configuration parameter. To the contrary, it only gains relevance in conjunction with the combinatoric parameters[7], *and vice versa*.)

The values denoted by symbols L, M, N and c (and some others which have not been presented individually due their subordinate importance) have been varied through all combinations that seemed sensible, the outcome of which already has been discussed in this section and presented in Figure 5.2. The criterion for momentum reconstruction in the BMS was fixed as quoted by M. v. Hodenberg [vH04], reducing the amount of evaluations that had to be performed.[8]

[6]It should be noted that this “exception” is *generically* supported by the implementation of the Conditional Coincidence Module and is *not* hard-coded as it may appear from this description. To the contrary, the current implementation supports three layers of nesting up to “meta-meta-groups”.

[7]fixed as well as variable

[8]Still Figure 5.2(a) includes over 10,000 data points, and in total over 1,000,000 data points have been calculated to determine the configuration of Cinderella for muon beam data taking. As the computation of a single data point is a time-consuming process, involving filtering of a representative data sample by a Cinderella instance started with a specific parameter set, the necessary computations were performed on the E18 cluster and took about a week to complete.

	L	M	N	c_1	c_2
IT	1	4	3	3.5	2.75
MT	2	4	3	3.5	2.75
LT	2	4	3	3.5	2.75
inclMT	1	4	3	3.5	2.75
OT	2	4	3	3.5	2.75
CT	1	4	3	3.5	3.75

Table 5.3: Configuration of Cinderella as it has been used during muon beam data taking of 2004. The value of c was decreased by 0.75 on 22$^{\text{nd}}$ of July 2004 because better time calibration got available for the silicon detectors at this date. Thus c_1 denotes the setting for c before that date and c_2 for thereafter. (For elucidation of the meanings of the parameters, refer to the previous section.)

Order	Abbreviation	Name	Number	Bit Mask
1	IT	Inner Trigger	0	0x0001
2	MT	Middle Trigger	1	0x0002
3	LT	Ladder Trigger	2	0x0004
4	inclMT	inclusive Middle Trigger	8	0x0100
5	OT	Outer Trigger	3	0x0008
6	CT	Calorimeter Trigger	4	0x0010

Table 5.4: Order of precedence of physics triggers. The *number* column refers to the trigger number that is used in the DAQ whereas the *bit mask* column visualises the equivalent position in the DATE *trigger mask*.

5.3.2 Filtering Criteria for the Individual Sub-Triggers

After the collaboration had agreed upon use of Cinderella, asserting utilisation of *conservative* settings, the values of Table 5.3 have been extracted from Figure 5.2(b), for each trigger selecting the set of options belonging to the data point highest in y-coordinate (reduction ratio) without notable differing from zero in the x-coordinate (inefficiency). The filtering inefficiency resulting from these settings has been estimated to be smaller than 0.5%.

The last item of configuration that is mentioned here is the *order of precedence* for sub-triggers, providing an answer for the question which setting to be used for filtering of an event that has been triggered on by multiple sub-triggers. (However this matter is not of extremely high importance as the settings for the individual sub-triggers were chosen to be similar, and the overlap among the sub-triggers is small in most cases.) In coordination with the trigger group (represented by J. Pretz [Pre04], in this case), the order of precedence for the physics triggers was defined as listed in Table 5.4. Events of non-physics triggers (all other triggers) are passed through Cinderella unfiltered.

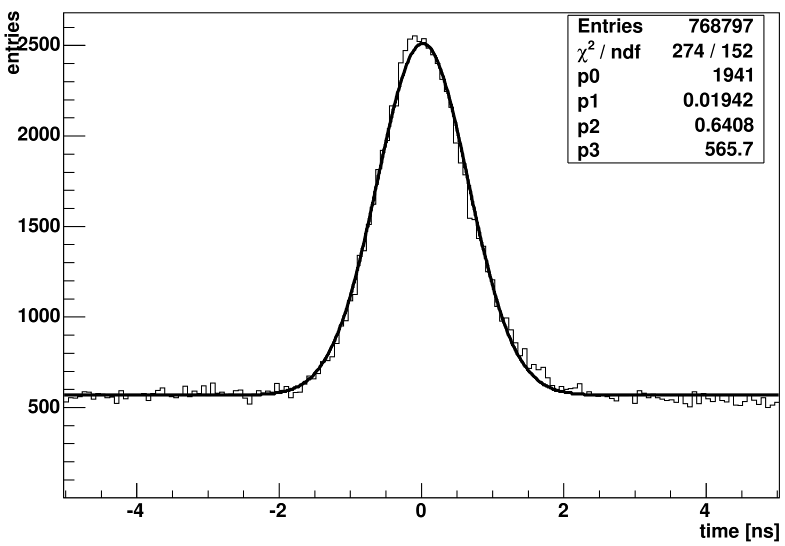

Figure 5.3: Fit performed by automatic calibration: Shown here is the timing of BM01P1 relative to Inner Trigger fitted with Gaussian distribution plus constant background. The uncertainty of time (combined for detector and trigger) is determined as 640 ps and an insignificant t_0 shift of 19 ps is observed.

5.4 Automatic Calibration

As explained in the previous sections, the positions and widths of the timing peaks of the detectors participate in the criterion of beam track reconstruct-ability and thus have a direct impact on the performance of the online filter. The offset of detector time against trigger time is determined by the individual detector experts for each individual channel and stored in the calibration database. However it cannot be ruled out, that the latency of a detector or a sub-trigger may drift (e.g. due to temperature changes) or leap (e.g. caused by exchange of a cable). In the case of movement of the timing peak outside of the acceptance of the online filter, the filter inefficiency would rise, degrading its performance.

To avoid this situation an automatic calibration procedure has been established for the detectors with TDC read-out (BMS and scintillating fibre). During filter operation, histograms of detector time relative to trigger time are generated for all combinations of planes and sub-triggers. At the end of every run, the histograms of the different Cinderella instances running on the event builders are transmitted to the *Calibrator*[9], where they are combined. Then a Gaussian distribution (plus constant background) is fitted to every histogram, an example being shown in Figure 5.3. The width and position of the timing peak determined in this way are stored as time uncertainty σ_t and as correction of t_0 with respect to the "official" calibration file. This correction has been termed *refinement*.

The implementation of an auto-calibration procedure for silicon detectors (APV read-out) still is pending. The safe and reliable implementation of the mechanism described

[9]During the 2004 data taking period, the Calibrator was running either on pccoeb01 or on pccoeb02.

above already was laborious and time-consuming. The calibration of silicon detectors (cf. section 3.5) involves fitting of two-dimensional histograms with complicated functions with many free parameters, which is very challenging and complex to implement, so that constraints in man-power forbade an engagement in that respect. Thus for the present, unfortunately the silicon detectors are not calibrated automatically, a fact that is accounted for by employing only comparatively loose cuts on silicon time information.

5.4.1 Results of Automatic Calibration

When interpreting the data obtained by the auto-calibration procedure, the main difficulty is the fact, that the measurement always consists of the sum of t_0 of detector plane *and* sub-trigger. Yet it would be desirable to be able to relate changes to either the detector or the trigger.

In an attempt to isolate the effects of the trigger system, in Figure 5.4 the t_0 refinements generated during the data taking of 2004 were plotted per sub-trigger but *averaged over all detector planes* (in black), compensating uncorrelated changes of the detectors. In the same figure, for comparison, the refinements of the individual planes are plotted in colours. While it is possible to think of systematic latency changes common to all observed detectors which would not be equalised by the averaging, the resulting plots are still useful.

When examining them, especially noteworthy is the drop in t_0 of $\sim$0.3 ns depth in the time between 7$^{\text{th}}$ to 20$^{\text{th}}$ of July 2004, corresponding to the time of reduced beam intensity. Also visible (mainly for IT, inclMT and CT) from 26$^{\text{th}}$ to 28$^{\text{th}}$ of June are the effects of an error in Cinderella (fixed on 28$^{\text{th}}$ June) which caused inaccuracies in the results of auto-calibration. Another interesting effect: the newly installed BM05 (pink) and BM06 (dark green) are retarded by nearly 1 ns in relation to the average, *but just for the calorimeter trigger*.

When plotting per plane t_0 refinements averaging over triggers, as in figures 5.5 and 5.6 for BMS and fibres, the t_0 constants look fairly stable. Catching the eye is the deviation of the Calorimetric trigger (dark green) with relation to the average.

For evaluation of the plots of time resolution it has to be considered—analogue to the t_0 data—that the obtained time resolution σ_t is that of sub-trigger and detector combined: $\sigma_t = \sqrt{\sigma_{trig}^2 + \sigma_{det}^2}$. With σ_{det} in the range of a few hundred picoseconds, the contribution of the detector may only be neglected for the Calorimeter Trigger, which with its broad timing peak may be called “slow” compared to the other triggers. The plots (figures 5.7 and 5.8) show reasonably stable behaviour for σ_t within limits of variation of about 10%.

Concerning the observed drifts and leaps of the t_0 constants of the triggers, most notably of the Calorimeter Trigger, the trigger group has provided no background information up to now.

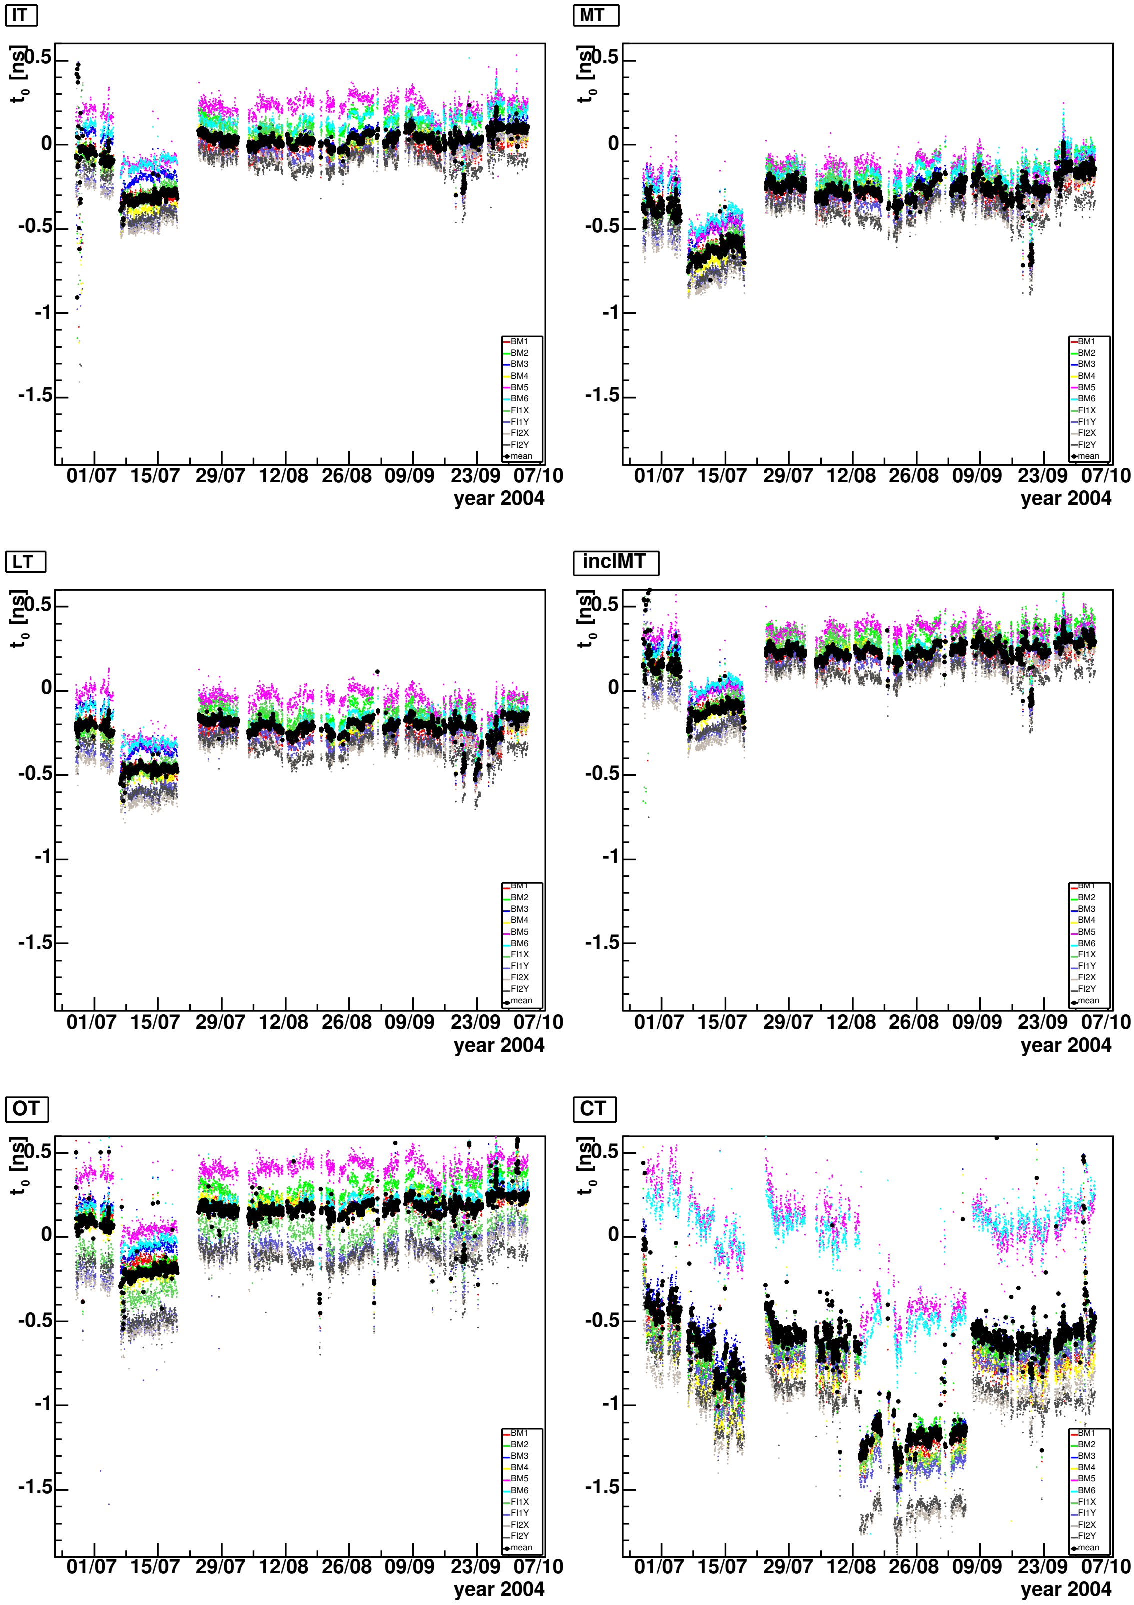

Figure 5.4: Refinements of t_0 calibration for every sub-trigger. The large black points show the average of all detector planes while the small coloured points display the t_0 shift with respect to the individual detector planes.

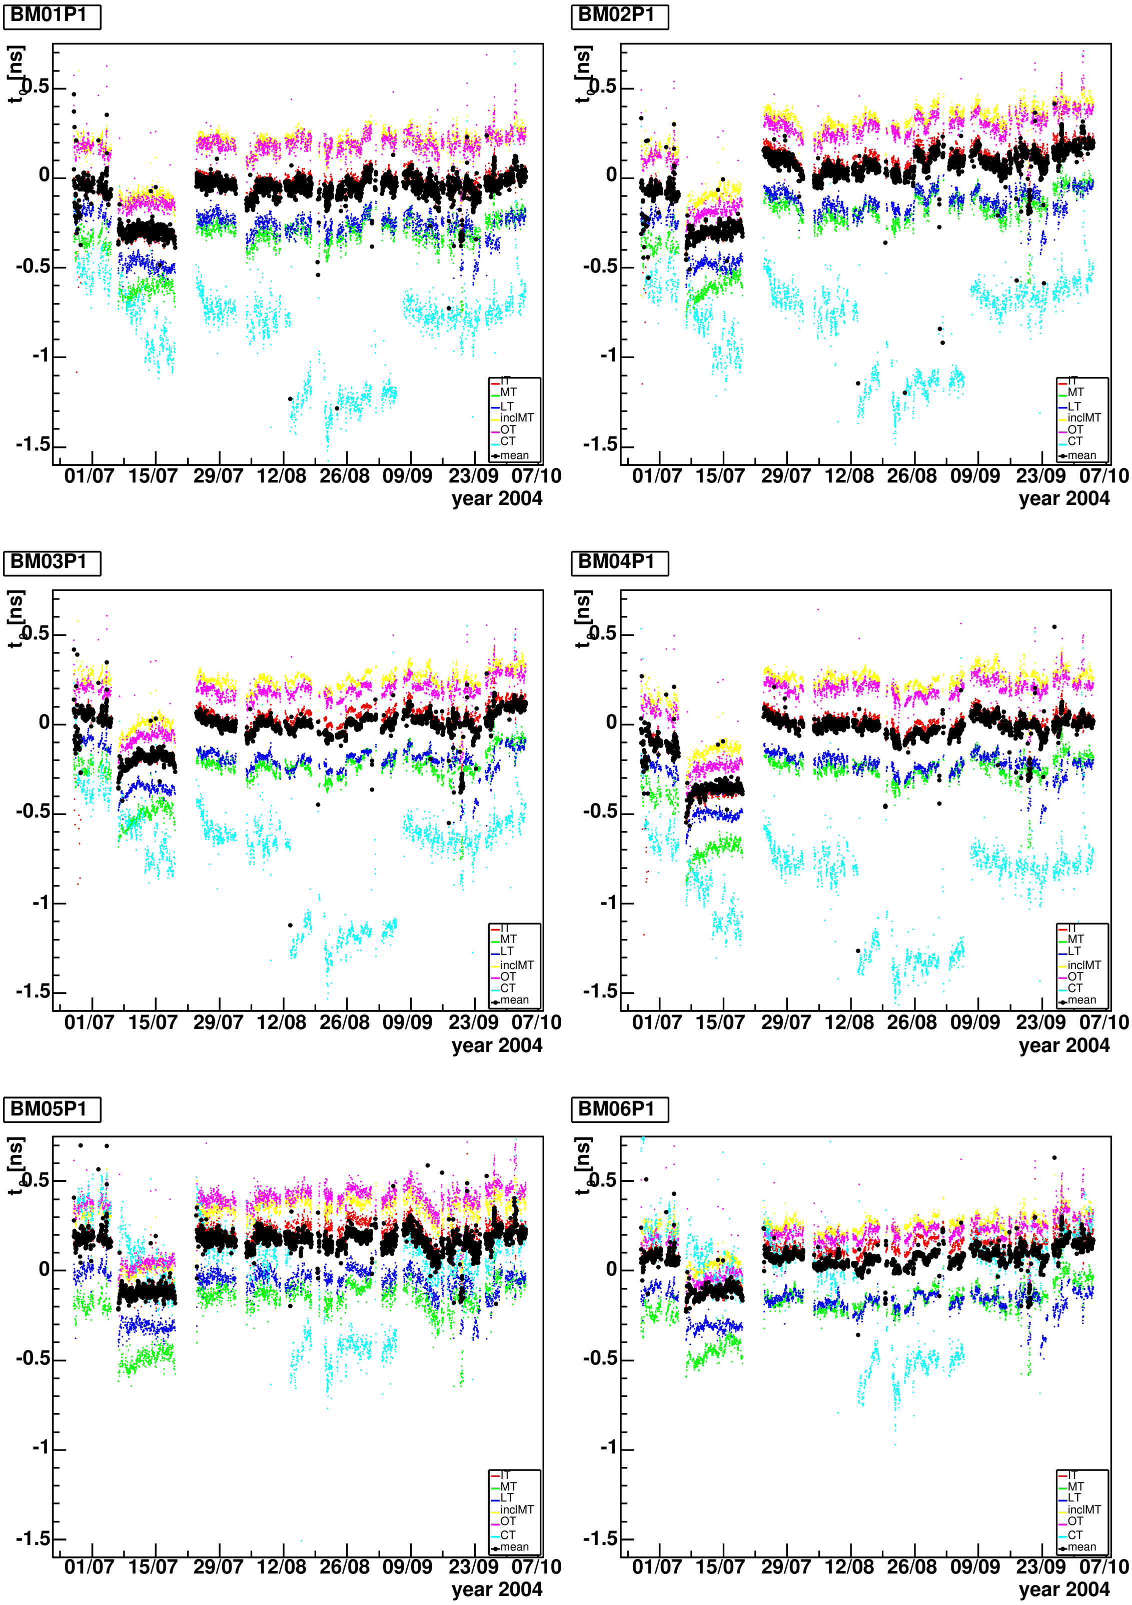

Figure 5.5: Refinements of t_0 calibration for every BMS detector plane. The large black points show the average of all sub-triggers while the small coloured points display the t_0 shift with respect to the individual sub-triggers.

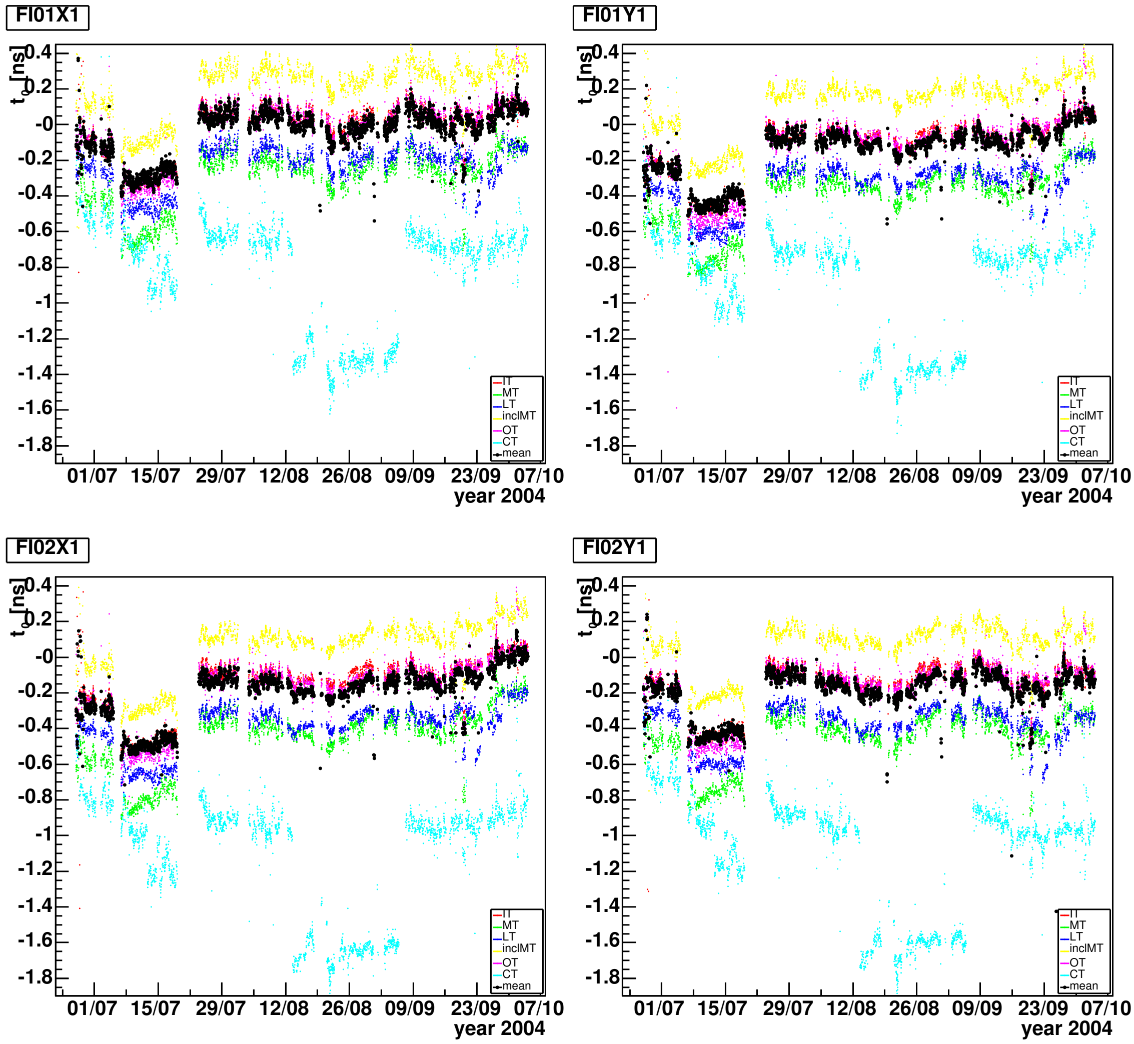

Figure 5.6: Refinements of t_0 calibration for every fibre detector plane. The large black points show the average of all sub-triggers while the small coloured points display the t_0 shift with respect to the individual sub-triggers.

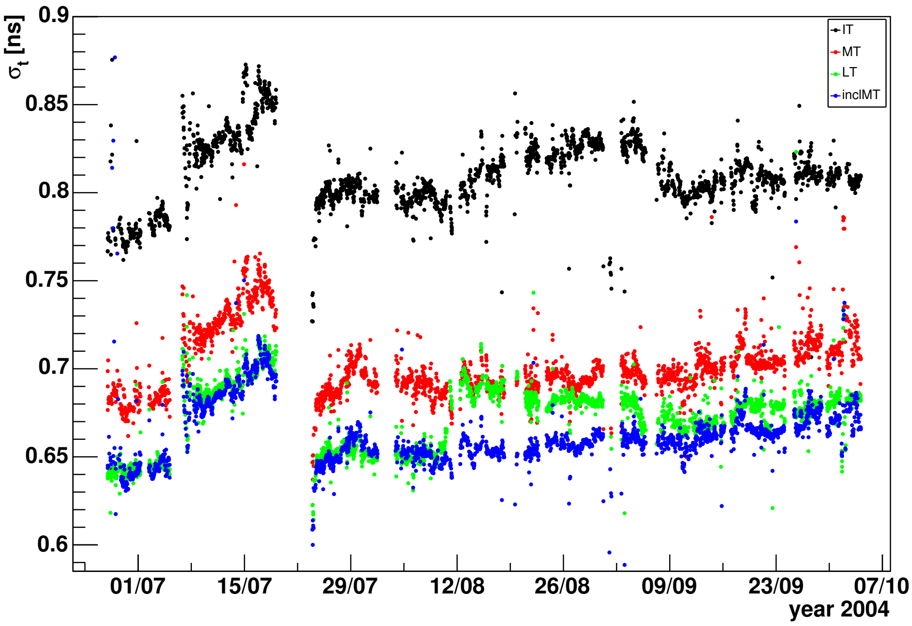

Figure 5.7: Widths of the timing peaks of different triggers, averaged over detector planes of BMS and scintillating fibre (FI01 and FI02), plotted for all of 2004 muon beam data taking.

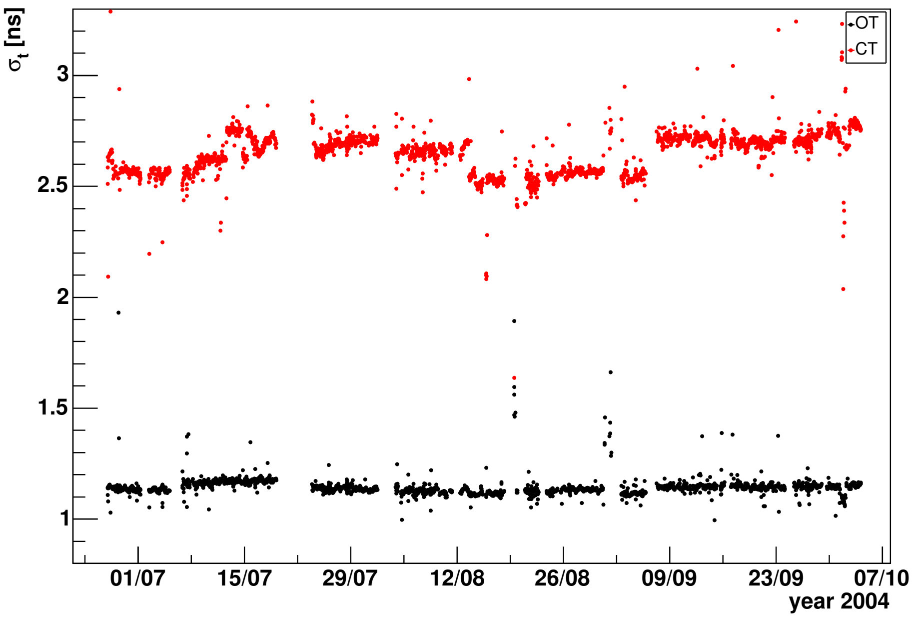

Figure 5.8: Widths of the timing peaks of different triggers, averaged over detector planes of BMS and scintillating fibre (FI01 and FI02), plotted for all of 2004 muon beam data taking.

5.4.2 Quality Criteria for Auto-Calibration Results

With a full set of refinements created for every run, the question arises, which of them should be used by Cinderella for the following runs. A selection criterion that guarantees that the refinement is *valid* as well as *recent* needs to be established. The demand is for Cinderella to act as intelligently as possible, refusing to apply bad refinements that result from detector failures or from the fitting algorithm[10] converging on a wrong peak. At the same time, Cinderella has to react immediately on legitimate changes to the refinements. For this reason a selection algorithm was chosen that operates mostly on *meta-data*.

Every fit that has converged without errors from the fitting package is subject to consistency checks. The position and height of the fitted peak are compared to the estimates of the peak-finding algorithm (which serves to provide initial values for the fitting algorithm) and the fit is flagged as bad if the deviation is larger than a certain threshold. A signal to noise ratio of less than 0.5 will flag the fit as bad, too. In the calibration database[11] are stored the position (t_0 refinement) and the width (time resolution σ_t) of the timing peak, the uncertainties reported by the fit algorithm for these quantities (in the following referred to as fit error or fit uncertainty), the reduced χ^2 (χ^2 divided by the number of degrees of freedom) returned by the fit algorithm as a measure for the quality of the fit, the number of entries of the histogram that was fitted (in the following referred to as statistics), and the flags of the sanity checks.

During the startup of Cinderella at the begin of every run, recent refinements are selected from the database according to a mechanism designed to avoid "bad" calibrations. For the two meta-parameters fit error and reduced χ^2, truncated mean and truncated standard deviation are determined for a sample that encompasses the last 24 hours, truncated by the highest and lowest 20% of each of the two meta-parameters. The obtained mean and deviation then is used to cut away from a sample comprised of the auto-calibration results of the last 6 hours all entries for which one of the meta-parameters is larger than its mean by more than two standard deviations. From the remaining refinements the *latest* is put to use. In the case that no refinement is left over after the cuts, Cinderella refuses to start and requests an *auto-calibration run*, during which only refinements are generated and no filtering is done.

The description above references the latest development status of Cinderella (as of version 2.0.10, installed 16$^{\text{th}}$ of September 2004) which includes more safeguards and logging than earlier versions. While the results of the next section indicate that the automatic calibration system worked reasonably well, it still is, to a certain extent, a work in progress. Detailed monitoring and analysis of the operation of the auto-calibration system had to be postponed until after the hadron run due to lack of manpower. Its results confirm the working of auto-calibration, yet some minor reconsideration is spurred, too:

Analysis has shown the relation between statistics and reduced χ^2 to be approximately linear (depicted in Figure 5.10). Paradoxically at first (naïvely red. χ^2 would be expected

[10] The Levenberg-Marquardt multidimensional derivative least-square fitting algorithm (see [Mor78] for details) is being used in the implementation of the GNU Scientific Library (GSL).

[11] table `tb_calib_refined` in database `runlb` on `pccoeb03`

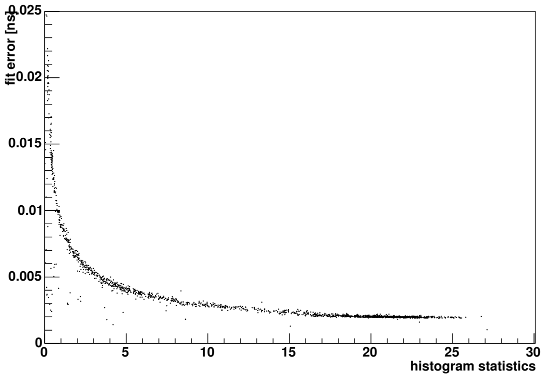

Figure 5.9: Uncertainty of the fit algorithm in its determination of the t_0 constant plotted against the number of entries to the histogram. As expected, the uncertainty of the determination of t_0 is proportional to the inverse of the square root of the statistics of the histogram.

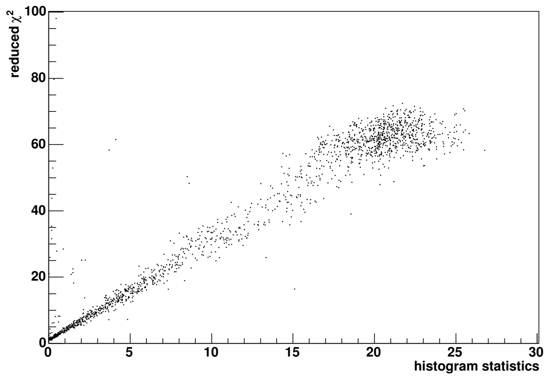

Figure 5.10: Reduced χ^2 of the fit displayed in relation to the number of entries to the histogram. The linear behaviour (opposite to intuition) is explained by small (but consistent) deviation from ideal Gaussian shape of the timing peak.

to *decrease* for higher statistics), this behaviour may be motivated by the consideration that the timing peaks do *not exactly* take the shape of an Gaussian distribution and thus with higher statistics the small but existing deviation from ideal shape get increasingly significant, expressed by rising red. χ^2. This realisation should initiate a change to the cut on reduced χ^2 described above for future versions of Cinderella.

The dependency of fit error on statistics however can be seen in Figure 5.9 as being proportional to the inverse of the square root of the number of entries to the histogram, very nicely in accordance to assumption.

5.5 Filter Performance

5.5.1 Availability

During the physics data taking with muon beam in 2004, spanning roughly the months of June, July, August and September, Cinderella was running *active* since end of June, when the first production version was deployed and configured. With respect to the whole period of muon data taking, Cinderella was used for filtering of 2105 out of 2602 recorded physics runs (81%), working on $11.0 \cdot 10^9$ Events out of $13.6 \cdot 10^9$ Events (81%).

Regarding only the time after the deployment of Cinderella (see Table 5.5), the overall availability was 96.5% of events. A portion of 1.5% of events were not filtered because Cinderella was auto-calibrating at that time, and the part of 1.9% of events was not processed with the online filter since Cinderella had been disabled by shift crew (mostly due to DAQ problems, of which some were, but most were not caused by Cinderella).

Runs	Runs (rel.)	Events	Events (rel.)	Filter Mode
2107	94.9%	$11.04 \cdot 10^9$	96.5%	filter-active
56	2.5%	$0.17 \cdot 10^9$	1.5%	filter-calibration
52	2.3%	$0.21 \cdot 10^9$	1.9%	mark-only
4	0.2%	$0.01 \cdot 10^9$	0.1%	pass-through
1	0.0%	$0.0003 \cdot 10^9$	0.0%	none
2220	100.0%	$11.44 \cdot 10^9$	100.0%	Total

Table 5.5: Survey of the number of runs taken and events processed in the different modes of the online filter during muon beam data taking of 2004 (taking into account only the period after deployment of Cinderella). The 2$^{\text{nd}}$ and 4$^{\text{th}}$ column contain the same information as the preceding columns in each case, yet calculated *in relation* to the entirety.

5.5.2 Fraction of Accepted Events

The first of the fundamental performance parameters of Cinderella is the fraction of accepted events. In Table 5.6 the averages of the fractions of accepted events and rejected/discarded events during the muon beam data taking of 2004 are summarised for the different triggers.[12] In average, Cinderella rejected 23.4% of all events, of which 0.8% were kept as part of the monitoring sample and 22.6% really were discarded. The rejection ratio varies widely for different triggers: 14–15% of events of the Middle Trigger (inclusive and exclusive) and the Calorimetric Trigger were rejected, the rejection ratio of the Inner trigger was 22% and Outer and Ladder Trigger even had average rejection ratios of 50–51%.

When interpreting these numbers, it is especially vexing that MT and inclMT show identical fractions of rejected events, differing by less than 0.03%. As the two variants of the Middle Trigger only vary with regard to inclusion of a calorimeter threshold, it would be expected that the higher impurity of the inclusive version also should lead to a higher fraction of rejected events in Cinderella.

As another interesting observation, the triggers LT and OT, which are utilising large hodoscopes comparatively far away from the beam, yield highest reduction rates in Cinderella. This may be explained by reasoning that they are especially susceptible to the effects of halo due to their size and position, and halo being the type of impurity that is reduced by Cinderella. While these arguments to a lesser extent also apply to Middle Trigger, this trigger is less prone to contamination by halo due to its two-dimensional target pointing.

Calculating with an average event size of 39 kByte, the utilisation of the online filter has saved 91 TByte of raw data. When the tape compression of $\sim$23% is taken into account, the amount of tape storage that has been saved may be quantified as 70 TByte. Finally

[12] This data, as well as that of the following plot, was generated from Cinderella's end-of-burst (EOB) statistics which were stored in the `beamdb` on `pccold05` and also serve as data source for the online display of acceptance. Due to small differences in methodology, the total number of events quoted here is slightly highly than that given in section 5.5.1, which was taken from `runlb` on `pccoeb03`.

Trigger	Accepted	Rejected	Discarded	Total events
ALL	76.579% ±0.0004	23.421%	22.640%	11,075,965,803
CT	85.177% ±0.0007	14.823%	14.329%	2,590,150,544
inclMT	86.160% ±0.0008	13.840%	13.379%	2,114,952,214
IT	77.778% ±0.0010	22.222%	21.481%	1,768,637,126
LT	48.609% ±0.0016	51.391%	49.678%	919,872,140
MT	86.133% ±0.0060	13.867%	13.405%	32,804,533
OT	50.467% ±0.0011	49.533%	47.882%	1,896,303,053

Table 5.6: Overall fractions of accepted and rejected events for physics runs recorded in *filter-active* mode in 2004. Due to continuing write-out of a pure monitoring sample of $1/30$ of *all* events, the ratio of *discarded* (deleted) events is slightly lower than that of *rejected* events (classified as "bad").

including the cost of the tape of 0.80 CHF per GByte [Mal04], Cinderella has saved the collaboration approx. 57,000 SFr (approx. 37,000 Euro as of January 2005) in expenses during the muon beam data taking of 2004.

Figure 5.11 shows the development of the ratio of accepted events over time. It can be observed that the reduction rates for the individual triggers are quite stable over the course of the three months of Cinderella run time, yet there are some notable structures: From 7th to 20th of July 2004 the SPS was running at half of its nominal intensity due to PS septum problems,[13] which manifests itself in larger error bars (best observable for Middle Trigger which has the lowest rate). On 22th of July cuts were tightened by 0.75 σ, resulting in a drop of ratio of accepted events of all triggers, made possible by the (thankworthy!) provision of more accurate calibration files for the silicon detectors by [Din04] and [Fri04a]. From 13th of August to 5th of September transversity data was being taken. As pointed out by [Pre04], the additional target dipole field somewhat diverts the beam, resulting in a changed trigger acceptance, which in turn leads to changed reduction rate.

Deviating points scattered across the plot may be attributed to instable beam conditions. For example on 15th of July, beam line quadrupole QUAD33 tripped, massively increasing halo, leading to strongly reduced acceptance in Cinderella. A different problem in beam line magnet control did manifest itself in a similar way on 25th of September.

5.5.3 Inefficiency

The second of the fundamental parameters of the performance of Cinderella is the *inefficiency*, the fraction of all events that would have been be useful for physics analysis but are discarded by Cinderella.

[13]There was a leak of cooling water into vacuum. The spare septum was undergoing repairs (as it itself had failed at the beginning of the beam period of 2004), so that the only solution was to refrain from cooling temporarily, which in turn required reduction of intensity by reducing the number of injections into SPS. As a result, the intensity at COMPASS was reduced to approximately half of its nominal value while at the same time a periodic substructure (with period of ~25 us) was added to the intensity profile.

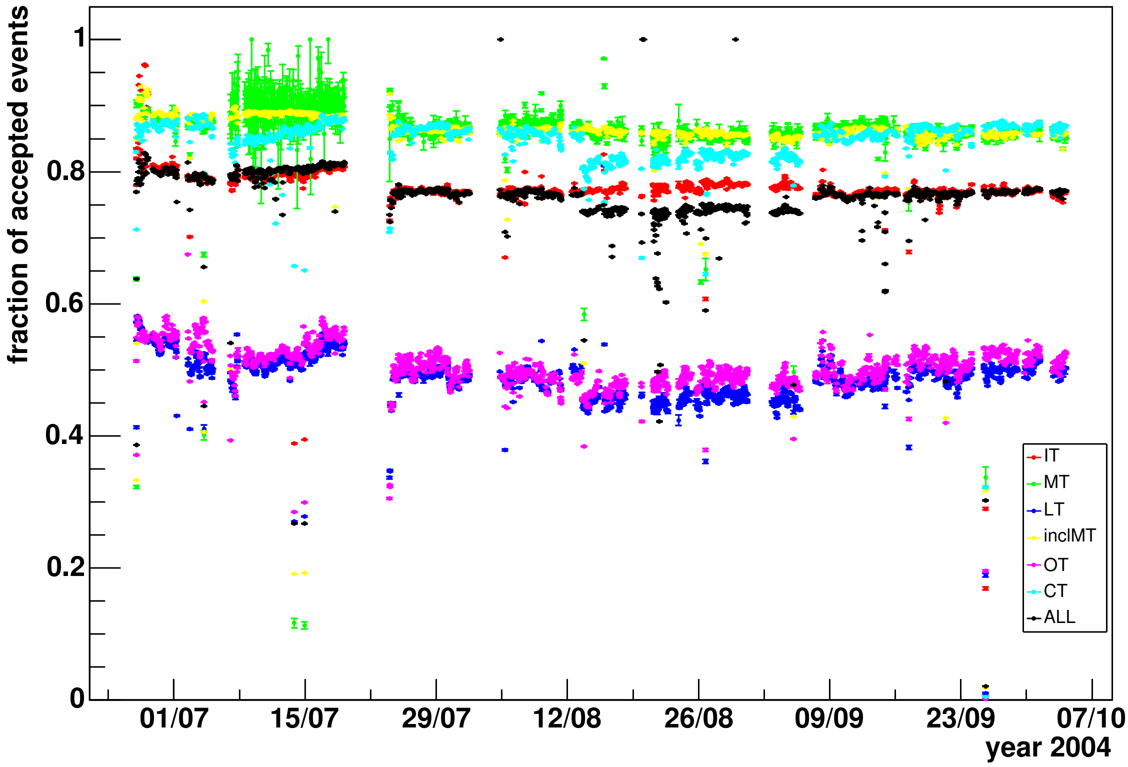

Figure 5.11: Ratio of accepted events during muon beam data taking of 2004, for physics runs with at least 225,000 events (corresponding to 5 average spills).

Currently CORAL includes all events in the mDSTs that have at least one reconstructed vertex. However for physics analysis a beam track is mandatory. Therefore the definition of an "event that may be useful" is one *having a beam track*.

In Table 5.7 the average inefficiencies of Cinderella are shown for all physics triggers and for different CORAL versions used for (pre-)production. The calculation of the inefficiencies has been done using the monitoring sample (1/30 of events). Since the results are very close to zero, it cannot be taken for granted that their (binomial) distribution is approximated well enough by the Gaussian distribution. Therefore the statistical errors are calculated asymmetrically from the binomial distribution, following [Sim01, Appendix A]. Systematic errors may stem from the fact that different sets of runs (only partly overlapping, if at all) were processed with the different CORAL versions.[14]

However runs processed with the same CORAL version share the same systematics so that the results within the same CORAL version may be compared without reserve. Concerning the latest (production) version of CORAL, Cinderella inefficiency is well below 0.4% for all triggers, nicely matching predictions. The individual triggers do not show large differences, the inefficiencies ranging from 0.21% (CT) to 0.31% (MT).

To allow meaningful comparison of different CORAL versions, Table 5.8 was compiled, using only data from runs that have been processed by more than one CORAL version. While the available statistics is limited and not all CORAL versions may be compared with each other, it still gives important insight: Cinderella shows significantly less inefficiency with respect to the later CORAL version "2004-11-17" (which is currently used for production) than relative to the "2004-06-09" version. (Changes between those two versions include improvements of parts dealing with the RICH and with tracking [Zha04].) This is especially interesting in consideration of the fact that the configuration of Cinderella was tuned using data processed with the "2004-06-09" version and thus naturally should perform best compared to it. This is a very reassuring result as it indicates that Cinderella was not designed and tuned to match CORAL and its very own characteristics but rather that the model developed for distinction between important and undesired events manages a sensible interpretation of the nature of the events.

[14]Only the latest CORAL version is used for production. Correspondingly, only small numbers of runs were processed with the earlier versions.

CORAL version	Trigger	Inefficiency	Monitored events
2004-06-09	ALL	0.330%±0.003	3,107,698
2004-06-09	CT	0.541%±0.009	723,766
2004-06-09	inclMT	0.248%±0.006	613,657
2004-06-09	IT	0.274%±0.005	1,164,524
2004-06-09	LT	$0.271\%^{+0.010}_{-0.009}$	294,609
2004-06-09	MT	$0.290\%^{+0.014}_{-0.013}$	161,120
2004-06-09	OT	$0.234\%^{+0.013}_{-0.012}$	150,022
2004-08-05	ALL	0.338%±0.002	8,091,463
2004-08-05	CT	0.235%±0.003	1,961,788
2004-08-05	inclMT	0.496%±0.006	1,189,864
2004-08-05	IT	0.272%±0.003	3,246,152
2004-08-05	LT	$0.469\%^{+0.008}_{-0.007}$	820,333
2004-08-05	MT	$0.575\%^{+0.012}_{-0.011}$	437,717
2004-08-05	OT	0.377%±0.009	435,609
2004-10-19-test	ALL	$0.399\%^{+0.016}_{-0.015}$	164,523
2004-10-19-test	CT	$0.519\%^{+0.037}_{-0.035}$	40,297
2004-10-19-test	inclMT	$0.413\%^{+0.040}_{-0.037}$	27,821
2004-10-19-test	IT	$0.268\%^{+0.022}_{-0.020}$	62,349
2004-10-19-test	LT	$0.469\%^{+0.057}_{-0.051}$	15,994
2004-10-19-test	MT	$0.607\%^{+0.086}_{-0.076}$	9,227
2004-10-19-test	OT	$0.396\%^{+0.073}_{-0.061}$	8,835
2004-11-17	ALL	0.259%±0.002	5,922,255
2004-11-17	CT	0.210%±0.004	1,251,978
2004-11-17	inclMT	0.274%±0.004	1,728,287
2004-11-17	IT	0.274%±0.004	1,921,728
2004-11-17	LT	0.261%±0.007	483,468
2004-11-17	MT	$0.307\%^{+0.011}_{-0.010}$	278,944
2004-11-17	OT	0.223%±0.009	257,850

Table 5.7: Cinderella inefficiency: The fraction of events that have a beam track in CORAL but would have been discarded by Cinderella. Since the different CORAL version have been used to process different runs, a systematic error of unknown magnitude (possibly large) exists *when comparing between different CORAL versions.* However comparison of different triggers *within the same CORAL version* does not suffer from any systematic error at all.

CORAL version	Trigger	Inefficiency	Monitored events
2004-08-05	ALL	$0.389\%^{+0.016}_{-0.015}$	162,630
2004-10-19-test	ALL	$0.399\%^{+0.016}_{-0.015}$	164,523
2004-08-05	CT	$0.513\%^{+0.037}_{-0.035}$	39,802
2004-10-19-test	CT	$0.519\%^{+0.037}_{-0.035}$	40,297
2004-08-05	inclMT	$0.417\%^{+0.041}_{-0.037}$	27,107
2004-10-19-test	inclMT	$0.413\%^{+0.040}_{-0.037}$	27,821
2004-08-05	IT	$0.260\%^{+0.021}_{-0.020}$	61,864
2004-10-19-test	IT	$0.268\%^{+0.022}_{-0.020}$	62,349
2004-08-05	LT	$0.465\%^{+0.057}_{-0.051}$	15,922
2004-10-19-test	LT	$0.469\%^{+0.057}_{-0.051}$	15,994
2004-08-05	MT	$0.596\%^{+0.086}_{-0.075}$	9,234
2004-10-19-test	MT	$0.607\%^{+0.086}_{-0.076}$	9,227
2004-08-05	OT	$0.299\%^{+0.065}_{-0.053}$	8,701
2004-10-19-test	OT	$0.396\%^{+0.073}_{-0.061}$	8,835
2004-06-09	ALL	$0.343\%^{+0.004}_{-0.003}$	2,783,648
2004-11-17	ALL	$0.290\%\pm0.003$	3,390,975
2004-06-09	CT	$0.553\%\pm0.009$	645,147
2004-11-17	CT	$0.241\%\pm0.006$	718,069
2004-06-09	inclMT	$0.260\%\pm0.007$	564,301
2004-11-17	inclMT	$0.296\%^{+0.006}_{-0.005}$	989,498
2004-06-09	IT	$0.291\%\pm0.005$	1,035,704
2004-11-17	IT	$0.313\%\pm0.005$	1,101,483
2004-06-09	LT	$0.285\%^{+0.011}_{-0.010}$	261,330
2004-11-17	LT	$0.292\%\pm0.010$	275,157
2004-06-09	MT	$0.298\%^{+0.015}_{-0.014}$	143,638
2004-11-17	MT	$0.345\%^{+0.015}_{-0.014}$	159,893
2004-06-09	OT	$0.243\%^{+0.014}_{-0.013}$	133,528
2004-11-17	OT	$0.249\%\pm0.013$	146,875

Table 5.8: Cinderella inefficiency: The fraction of events that have a beam track in CORAL but would have been discarded by Cinderella. The sections of the table (separated by double lines) are compiled from the same sets of runs, eliminating systematic error, so that values within a section may be compared with each other without reserve.

Chapter 6

Filtering of the Hadron Beam

For the hadron pilot run of 2004 including build-up and tests only six weeks were scheduled, leaving about two weeks for physics data taking. To make best use of the limited time, it was decided to record as much data as possible, up to full saturation of the 1GBit/s link from the experimental hall to the CDR. It was planned to employ the online filter only if the overall data rate would exceed the limit of the link, writing to tape as much data as possible.

Preparation of Cinderella for the hadron data taking was complicated by the fact that during build-up and also to some extent during operation the detector and trigger setup was very much a *work in progress*. Except for some very limited test runs in some of the previous years, there was no experience with the hadron beam at COMPASS. Useful previous data did not exist and until the very last moment trigger rates and (im)purities were subject to speculation and coarse estimates. Additionally, details of the spectrometer configuration were subject to discussion and change until well into hadron data taking.[1]

In the face of the entire experimental setup being a *rapidly moving target* before and during hadron data taking, the best approach for filtering of the hadron beam data was to prepare generic, flexible filter modules and then keep ready to join, configure and deploy them immediately upon demand. For some aspects of development it was necessary to simulate detector behaviour.

Yet in a way, the design decisions for Cinderella for the hadron programme were simplified by the lack of time for the build-up: They could be reached by exclusion. It was obvious that in the short time no new decoding module could be written, so that Cinderella would only be able to operate on the types of readout that already were implemented for the muon programme: scintillating fibre (TDC) and silicon (APV). Also it was self-evident that time would not suffice to implement any kind of geometric operations like tracking or vertexing as existing infrastructure did not include support for geometric properties at all.

[1]The topics subject to discussion and/or uncertainty did include trigger setup, number and mode of operation (cold/warm) of silicon detectors, multiplicity counter scintillator, (partial) removal of fibre detectors, and target thickness.

Position [m]	TBName	Type	Function
-5.35	SI01U1	silicon	beam definition
-5.35	SI01V1	silicon	beam definition
-5.34	SI01Y1	silicon	beam definition
-5.34	SI01X1	silicon	beam definition
-4.26	SI02U1	silicon	beam definition
-4.26	SI02V1	silicon	beam definition
-4.25	SI02Y1	silicon	beam definition
-4.25	SI02X1	silicon	beam definition
approx. -3.00		different materials	target
-2.01	SI03U1	silicon	vertex definition
-2.01	SI03V1	silicon	vertex definition
-2.00	SI03Y1	silicon	vertex definition
-2.00	SI03X1	silicon	vertex definition
-1.49	SI04U1	silicon	vertex definition
-1.49	SI04V1	silicon	vertex definition
-1.48	SI04Y1	silicon	vertex definition
-1.48	SI04X1	silicon	vertex definition
-1.02	SI05U1	silicon	vertex definition
-1.02	SI05V1	silicon	vertex definition
-1.01	SI05Y1	silicon	vertex definition
-1.01	SI05X1	silicon	vertex definition
32.50	HP01X1	hodoscope	Primakoff Hodoscope
33.43	EC02P1	calorimeter	ECAL2

Table 6.1: Selected elements of the spectrometer setup for hadron beam data taking as of 9th of November 2004.

Selected elements of the spectrometer setup as of 9th of November 2004 are listed in Table 6.1.[2]

6.1 Criterion: Track Multiplicity in Primakoff Hodoscope

One idea for filtering the hadron beam data emerged on very short notice. There had been ambitions to implement in trigger logic a cut on hit multiplicity in the Primakoff Hodoscope to reduce trigger rate for the Primakoff triggers. In principle it can be assumed that a Primakoff reaction has a signature of exactly *one* track (scattered beam hadron) in the Primakoff Hodoscope (PH). For Prim_2 trigger, which does not include PH in the trigger logic, there is also the possibility of zero hits to PH in cases of the scattered beam particle missing the hodoscope. Any track multiplicity in PH larger than *one*

[2] taken from `detectors.43035.hadron.dat`

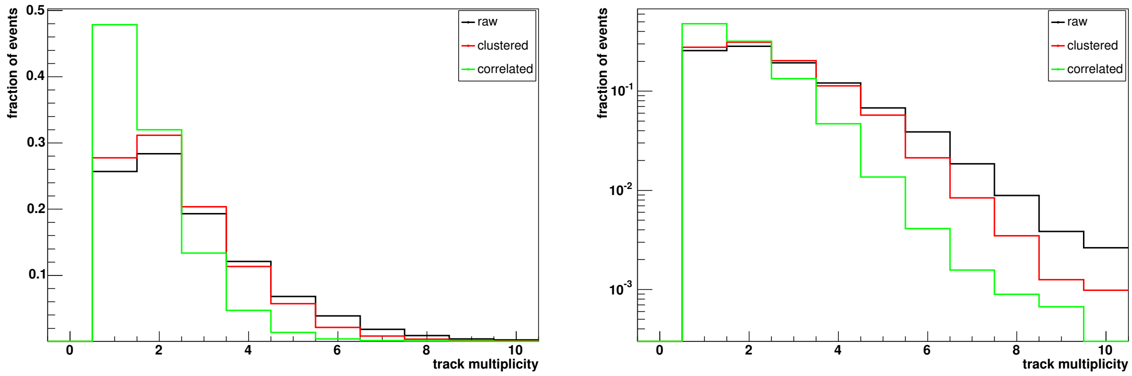

Figure 6.1: Multiplicity in Primakoff Hodoscope for Prim_1 trigger (linear resp. logarithmic scale)

is pointing towards the fact that there have been undesired interactions of beam particle and spectrometer structure, rendering the event useless for physics analysis.

After the aforementioned efforts of the trigger group have become generally known, it was immediately clear that Cinderella can duplicate the functionality of the hardware cut, yet with higher accuracy. Unlike hardware implementations, Cinderella is capable of clustering and time correlation of the hits to the individual hodoscope slabs. As adjacent slabs are mounted with a certain amount of overlap to improve efficiency, there is the possibility of one particle track passing through two slabs, which wrongly would be considered as two tracks by hardware logic. Thanks to clustering, the online filter computes the correct number of tracks in these cases. The second advantage of Cinderella is its flexible time cut, correlating the hits to the individual slabs using a narrow sliding window. This way, pile-up and noise may be rejected more efficiently, leading to less fake tracks compared to the fixed window approach of hardware logic.

Figures 6.1 and 6.2 show histograms of the track count computed by Cinderella with different processing logic in different colours for both Primakoff triggers. The black line depicts the raw, unprocessed hit counts to PH (which is similar to what would be evaluated as track count by an electronics-only trigger setup). The red line denotes the amount of hit clusters (tracks) using a hodoscope clustering algorithm specifically designed for this task. Finally the green line exhibits the number of simultaneous tracks, employing clustering *and* time correlation algorithms.

It is clearly visible that the number of tracks is reduced with every step of processing. As a direct consequence, a cut in track count realised in hardware would suffer from inefficiency to a much higher extent than the same cut implemented in Cinderella so that utilisation of the online filter would lead to *less deficit of useful events* compared to using a rendition of similar functionality in hardware.

However at some point it was noticed that the Primakoff Hodoscope is mounted less than a metre before ECAL2, which may lead to fake tracks in HP from backscattering ("albedo") of ECAL2, constituting a source of inefficiency of unknown magnitude. A quick study based on physics data already taken was conducted by Alexey Guskov, how-

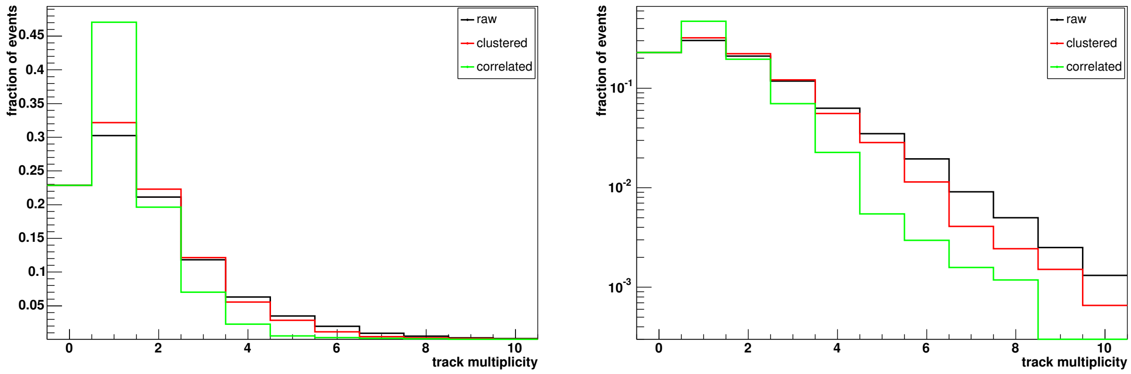

Figure 6.2: Multiplicity in Primakoff Hodoscope for Prim_2 trigger (linear resp. logarithmic scale)

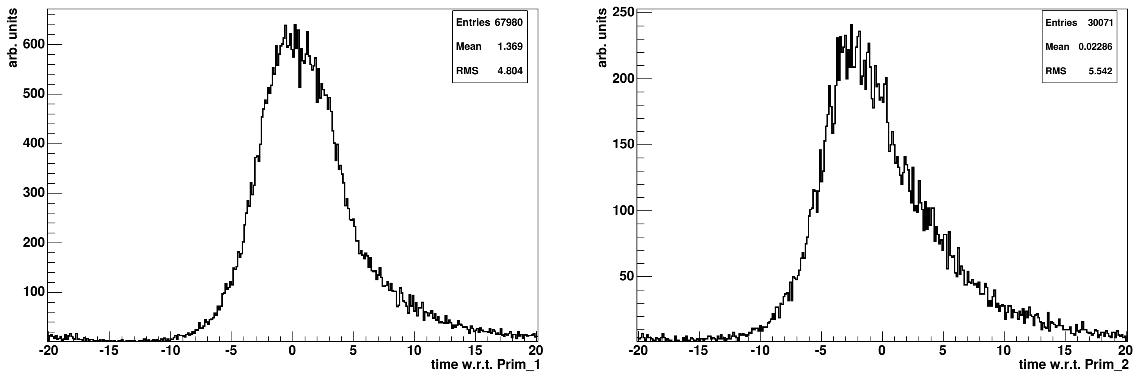

Figure 6.3: Timing of Primakoff Hodoscope with respect to Prim_1 and Prim_2 triggers

ever the size of the effect could not be determined sufficiently in the little time that was available.

The timing plots of the Primakoff Hodoscope (fig. 6.3) do not shed very much light on the issue, either. When hitting PH, the distance travelled by albedo particles is longer by at least $\sim$1.5 m compared to that of the scattered beam particle. Consequently their timing should be late by at least $\sim$5 ns. Especially for Prim_2 trigger the timing peak is skewed to the right which may indicate a significant share of albedo, but for an unambiguous determination of the contribution of backscattering the time resolution of the PH is not good enough.

In light of these facts, the conservative decision was taken and the cut in the Primakoff Hodoscope was abandoned. It was agreed that the risk of undesired effects was not worth the increase in statistics that would have been gained.

6.2 Criterion: Track Multiplicity in Silicon after Target

The other of the few open venues for a filtering algorithm was to determine track multiplicity from cluster counts in the silicon planes. As quantifying the number of tracks after the target may give evidence of the type of interaction, cutting schemes for all subtriggers can be envisioned.

For Primakoff triggers, the scattered beam hadron is the only particle taking part in the interaction which can be detected by the silicon detectors after the target. A reconstructed track multiplicity larger than *one* is a sure sign that an undesired interaction is being observed and that the event should be discarded.

For Diffractive Trigger the situation is a little bit more complicated. In diffractive interactions, short-lived hadrons are produced and the tracks of their daughter particles may be observed along with that of the scattered beam hadron. Charge conservation applies, so that the number of *visible* tracks leaving the target is odd.[3] Additionally, the track count is known to be larger than *one* as the produced hadrons decay by strong interaction already within the target or very close to it.

It is reasonable to assume that a substantial part of undesired “hard” hadronic interactions in the target do not trigger the Aperture Sandwich veto and thus represent a considerable share of the impurity of Diffractive Trigger. Large parts it may be rejected by the online filter using track counting, requiring more than two charged tracks after the target.

As the consensus among the collaboration was to refrain from using Cinderella on Primakoff data, the possibilities of utilising after-target track multiplicity to reduce trigger rate were not explored any further for the Primakoff triggers. Therefore the following sections concentrate on application of track counting to Diffractive Trigger.

6.3 Coincident Multiplicity Algorithm

Both criteria introduced in the last sections aim at deducing track multiplicities from hit multiplicities, albeit of different detectors. Thus, the basic systematic is shared. As addressed already partly in section 6.1, the main obstacles to this task are constituted on the one hand by cases in which one track causes multiple hits and cases where hits occur without (relevant) track, namely due to noise or off-time hits (pile-up), on the other hand. The former case is handled reasonably well by the detector-specific clustering algorithms for silicon and hodoscope. For the latter case, a mixed approach is appropriate. In a first step, noise and pile-up is suppressed by several quality cuts in silicon decoding. Then the coincident multiplicity algorithm serves to determine the track multiplicity in silicon and hodoscope after further reduction of noise and pile-up by its requirement of coincidence.

The coincident multiplicity algorithm draws heavily on the method for fast computation of coincidence established in section 5.2 for the conditional coincidence algorithm. Again

[3] except in extremely rare events containing doubly charged particles

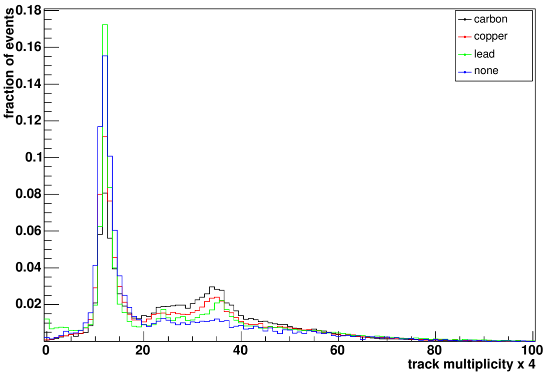

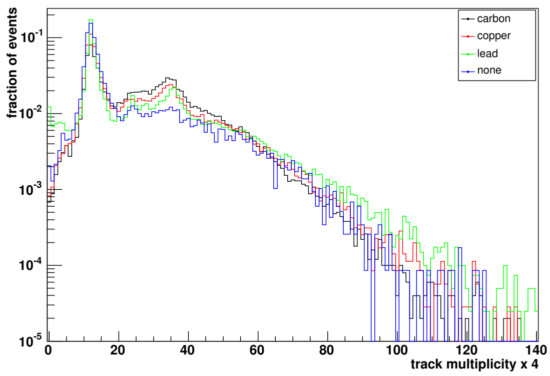

Figure 6.4: Total cluster count for all 12 silicon detector planes after target in Diff_1 events for different target materials (linear resp. logarithmic scale). This value *cautiously* may be interpreted as *track multiplicity* $\times$ *12*.

boundaries are calculated to both sides of the time points within the flexible distance of $c \cdot \sigma_t$. But when the sorted array of boundaries is traversed in chronological order, instead of testing a conditional expression, the sum of *coincident* clusters in all participating planes is recorded continuously. Now the track multiplicity is yielded as "average track count" by dividing the maximum of the cluster sum by the number of planes.

Ideally, all planes should report the same number of clusters exactly. However in reality, the number of clusters in the individual planes only roughly does correspond to the number of tracks through them. Still detector inefficiency or overzealous cuts may lead to less clusters than traversing tracks while noise spikes may bring about planes with more clusters than incident charged particles. Blindly averaging over all planes is not the approach most suited to that condition. For a plane with cluster count much different than that of the other planes it is quite certain that an anomaly has taken effect and that calculation is best served by completely excluding that plane from the average. This consideration spurs the introduction of the *truncated mean* as measure of track multiplicity.

The actual implementation avoids division and merely calculates the *truncated sum* of clusters, disregarding the n_l planes with lowest and the n_h planes with highest cluster count, where n_l and n_h are subject to configuration. The truncated sum obtained in this way finally is compared to a configured set of intervals and the event is discarded if the truncated sum falls within an undesired range. This rather general approach allows for selection of events with more (or less) than a certain number of tracks or even for cherry-picking of just some desired track multiplicities, like 3, 5, 7..., which is profitable for reducing Diffractive trigger.

For illustration the total cluster count for all 12 silicon planes is depicted in Figure 6.4 without truncation, using different colours for different target materials. (In a setting *with truncation* the corresponding data are exhibited in figures 6.5 and 6.6, for comparison.)

6.4 Geometric Acceptance as a Source of Inefficiency

The computational results of the coincident multiplicity module for Diffractive Trigger may be viewed in figures 6.5 and 6.6 in linear respective logarithmic scale. For these plots the configuration $n_l = n_h = 4$ was used,[4] so that, with a total of 12 silicon planes after the target, the sum of clusters of the 4 planes remaining after truncation is reported. This quantity therefore is to be interpreted as *track multiplicity* $\times$ *4*, as determined by online filter.

Looking first at Fig. 6.5 (linear scale), sharp peaks at multiples of four catch the eye. With up to $\sim$50 the ratio of peak height to background is very good for track multiplicity 1 and then decays for higher multiplicities to $\sim$3 (multiplicity 3) and $\sim$1 (multiplicity 5). However in Fig. 6.6 (logarithmic scale), peaks may be identified even up to track multiplicity 8, after which the signal finally is drowned by the background.

The causes of the background are manifold. Remains of noise/pile-up may skew peaks in direction of higher track multiplicity while the opposite, cuts taken too tightly, may skew peaks the other way. Examination of the aforementioned plots reveals both effects up to a certain extent, yet proves the latter more pronounced, especially in regions of higher track multiplicity.

Another type of background is due to an effect that has been neglected up to now. The angular acceptance of the silicon micro-strip detectors is smaller than that of the spectrometer: For tracks emanating from the target, station SI05 is limited to a half-angle acceptance of $\sim$13 mrad, while SI04 provides $\sim$17 mrad and SI03 sports $\sim$25 mrad of coverage in the vertical plane. In the horizontal plane the acceptances are larger by a factor of 7/5, which is explained by the rectangular shape of the silicon wafers (cf. section 3.3). As a consequence of this limitation, valid tracks may bypass some or all of the silicon detector planes and thus (partially) escape the notice of any conceivable track multiplicity algorithm based on the silicon detectors.

Yet, the situation is ameliorated by the fact that the largest share of tracks of desired interactions is emitted at very low angles. As by far the largest part of all tracks missing the silicon detectors belong to undesired reactions, the implication of the limited angular acceptance is that the track multiplicity may be reported wrongly for *some undesired* events but only for *very few desired* events. This means that the effect may introduce impurity to the online filter to some extent—which is tolerable—rather than causing a great deal of the much more dreaded inefficiency. In any case, tracks that exceed a certain angle hit the Aperture Sandwich veto and therefore prevent events containing them from being triggered.

The second "geometric" source of inefficiency stems from the fact that two tracks may be wrongly regarded as being one when they are too close. This problem occurs for almost any kind of detector, in the case of the COMPASS silicon detectors however it is somewhat worsened by the fact that Cinderella's silicon clustering has been optimised for one-track resolution.

[4] A series of tests with different n_l and n_h indicated this combination to yield the best separation of adjacent peaks.

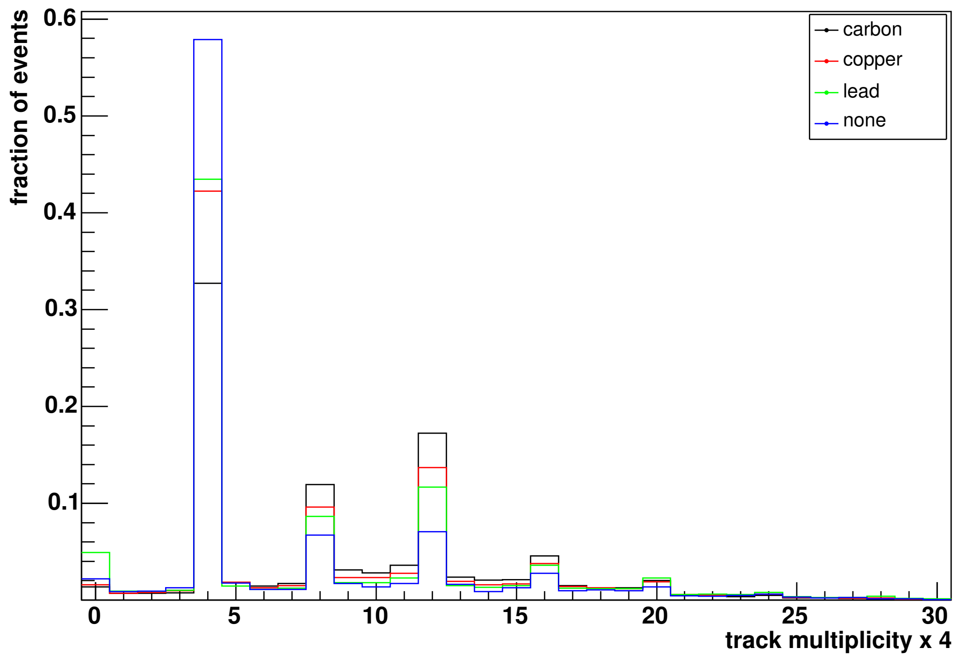

Figure 6.5: Eightfold truncated cluster count of after-target silicon planes in Diff_1 events for different target materials (linear scale). To be interpreted as *track multiplicity* × *4*.

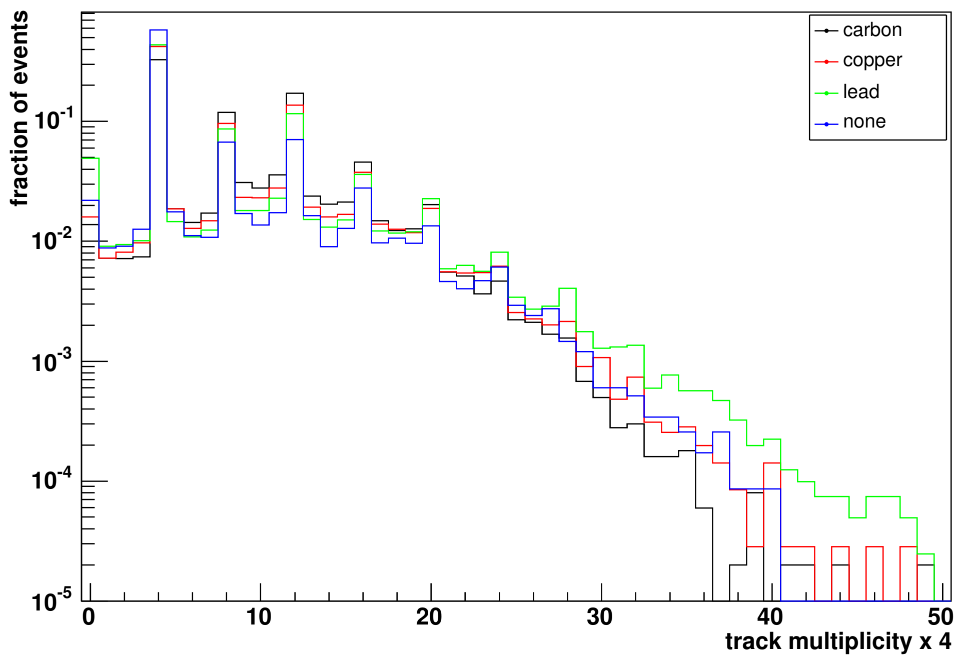

Figure 6.6: Eightfold truncated cluster count of after-target silicon planes in Diff_1 events for different target materials (logarithmic scale). To be interpreted as *track multiplicity* × *4*.

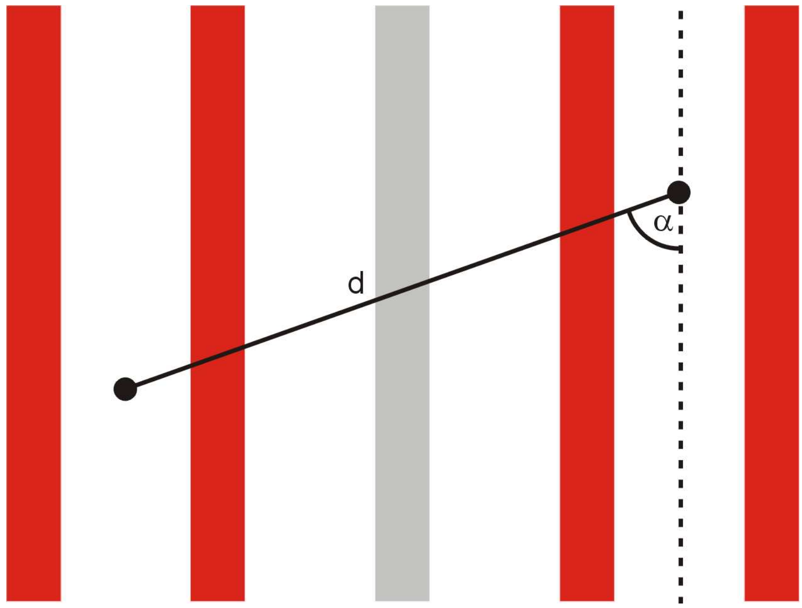

Figure 6.7: Illustration of estimation of average minimal distance between two tracks for the clustering algorithm to recognise them as separate. The black dots denote particle trajectories. The red bars represent activated strips of the silicon wafer whereas the grey bar marks the strip below threshold between them, which is necessary for separation.

For detector hits of two proximal tracks to be taken as two separate clusters it is necessary that there is at least one strip with signal amplitude lower than the level of zero suppression in between the strips being hit by the two particles. The average minimal distance between two tracks for them to be recognised as separate tracks by the clustering algorithm is estimated coarsely in the following discussion. An illustration is given in Figure 6.7.

With readout pitch of $\sim$53 μm and average cluster size below 2, this condition in average is satisfied for track distance $d > 159\,\mu$m, assuming d perpendicular to the readout strips. As the readout strips on both sides of the wafer are orthogonal, the connecting line d between the two points of intersection of the particle tracks and the silicon plane always includes an angle larger or equal to 45° with the direction of the strips of one side of the wafer. Thus for two tracks to be separated at least in one of the two projections of the wafer, the average angle α may be taken as 67.5° and the average minimal track distance d needs to be enlarged by a factor of $1/\sin(67.5°)$ compared to the condition of perpendicularity. Now for two tracks to be treated as separate in at least every other plane, $d > 172\,\mu$m is required. Thus the minimal included angles for two tracks to be treated as separate in at least every other plane may be calculated as 0.17 mrad, 0.11 mrad and 0.09 mrad for stations SI03, SI04 and SI05 using the detector positions listed in Table 6.1.

6.5 Configuration

After the discussion of the limitations of the determination of track multiplicity from silicon cluster multiplicities, figures 6.5 and 6.6 are revisited to constitute the ranges for in- and exclusion of events. Before the online filter was used, the limitation of transfer rate required the down-scaling of Diff_1 trigger by factor 2. Therefore the natural approach was to aim at reducing Diff_1 trigger by the same factor in Cinderella, while retaining as much of useful events as possible, so that hardware prescaling could be avoided.

For Diff_1 trigger, after-target track multiplicities of 3, 5, 7... define useful events due to physics considerations (cf. section 6.2). Examination of the aforementioned plots however shows that the goal of data reduction by $\sim$50% already may be reached by rejecting events with zero or one tracks. In the name of minimising inefficiency such a setting of course is much preferred to a tighter cut. As there is a "distance" of two tracks towards valuable track multiplicities this setting exhibits great robustness towards all sources of inefficiency.

In terms of Cinderella configuration, the cut for events with Diff_1 trigger finally was fixed to *track multiplicity* $\times$ *4* larger than 5, effectively cutting for "more than 1.25 tracks" after the target.

As a full calculation of inefficiency was impossible due to lack of time,[5] the only way to gain an impression of it was to consider incidents that would lead to rejection of a legitimate event:

For an event with 3 tracks after the target to be rejected, any of the following conditions need to be satisfied in full (or a combination of these needs to be satisfied partly):

- all three tracks share an included angle of less than $\sim$0.11 mrad
- two out of three tracks exceed a vertical angle of $\sim$17 mrad or a horizontal angle of $\sim$24 mrad
- more than 6 out of 12 planes of stations SI03–SI05 fail completely

For events with 5 or more tracks after the target, the requirement for rejection are still more drastic, correspondingly.

From all this it may be concluded that the inefficiency should be very small. These considerations also did convince the COMPASS collaboration so that Cinderella was switched to active with the discussed setting on 9$^{\text{th}}$ of November 2004.

[5] It would have required CORAL production of hadron data *plus* elaborate analysis of the results.

Runs	Runs (rel.)	Events	Events (rel.)	Filter Mode
103	73.0%	$0.95 \cdot 10^9$	74.3%	filter-active
38	27.0%	$0.33 \cdot 10^9$	25.7%	mark-only
141	100.0%	$1.27 \cdot 10^9$	100.0%	Total

Table 6.2: Survey of the number of runs taken and events processed in the different modes of the online filter during hadron beam data taking of 2004 (taking into account only the period after deployment of Cinderella). The 2^{nd} and 4^{th} column contain the same information as the preceding columns in each case, yet calculated *in relation* to the entirety.

Trigger	Accepted	Rejected	Discarded	Total events
ALL	89.538% ± 0.0010	10.462%	10.113%	949,491,141
Diff_1	53.732% ± 0.0039	46.268%	44.725%	165,597,373

Table 6.3: Ratios of accepted and rejected events for recorded physics runs with *filter-active*. Due to continuing write-out of a pure monitoring sample of 1/30 of *all* events, the ratio of *discarded* (deleted) events is slightly lower than that of *rejected* events (classified as "bad").

6.6 Filter Performance

6.6.1 Availabiliy

In October of 2004, Cinderella was ready for deployment at the beginning of hadron data taking, in fact it was with Cinderella that the very first "physics" plots (track count after target) of hadron beam data have been produced, just hours after the first run has been taken. For historical reasons these plots are reproduced in Figure 6.9.

On 9^{th} of November 2004, the Diff_1 trigger had been setup and the use of the online filter had been agreed upon by the collaboration so that Cinderella could be switched on. Of the remaining five days of beam time, $1\frac{1}{2}$ days were used by MD[6] and muon beam (for systematic studies). Unfortunately the shift crew forgot to enable the online filter after return to hadron beam on 12^{th} of November and it was not until the next day that the mishap was noticed and corrected by Stefano Panebianco [RLB, comment #12036]. So altogether only about 100 runs with about 1 billion events were taken with Cinderella in *filter-active* mode (see Table 6.2).

6.6.2 Fraction of Accepted Events

The fraction of accepted events is summarised in Table 6.3. Although the settings for filtering of the Diff_1 trigger have been chosen very conservatively, approx. 46% of the events could be rejected, allowing to dispense of pre-scaling, which had been necessary

[6] **M**achine **D**evelopment

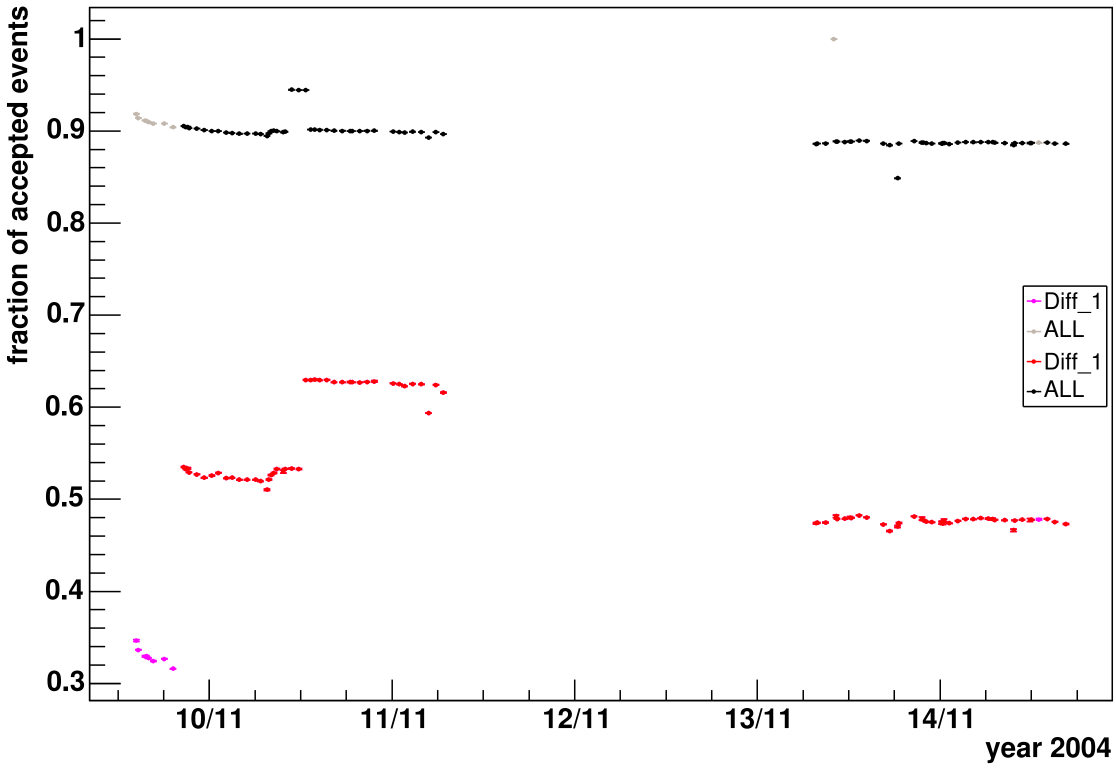

Figure 6.8: Fraction of accepted events during 2004 hadron data taking for runs with at least 400,000 events (corresponding to 5 average spills). Physics runs are plotted in red/black while non-physics runs are depicted in violet/grey.

before with a factor of 2. Thus utilisation of Cinderella effectively doubled the available statistics for Diff_1 trigger, while it was running. In the approx. $2\frac{1}{2}$ days in which the online filter has been enabled, it has reduced the volume of recorded data by approx. 2.4 TByte.

Figure 6.8 shows the development of the fraction of accepted events over time. It is plotted for physics runs as well as for non-physics runs to elaborate the behaviour of the online filter for different targets. In stable conditions, the fraction of accepted events is observed to be stable, too. Outstanding details are well explained by changes to setup or operation of the experiment.

In the beginning on afternoon of 9 November 2004, no target is present so that the fraction of accepted events is low, naturally. At 21:28, 9 November a copper target is inserted, leading to an increase in accepted events. At 9:52, 10 November, a prescaling factor of two was introduced for the Diff_1 trigger which at 14:18 was removed again. This action lead to a prominent bump in the fraction of accepted events for all triggers while that of Diff_1 is not influenced. On the same day at 14:16 the target was switched to carbon which manifests itself in a higher fraction of accepted events. The runs after 17:09, 12 November were taken with a lead target resulting in yet another change.

6.6.3 Inefficiency

Since for hadron beam data Cinderella employs properties of the *physics process* for its decision taking (instead of properties of the beam particle as for muon beam data), the analysis of its inefficiency takes a new quality. A general inefficiency may not be quoted anymore as it now is dependant on the physics channel that is being analysed.

To determine the inefficiency of Cinderella for a given physics channel, the distributions of examined physical properties may be compared for pure "accepted" samples and pure "rejected" samples. When fitted for signal s and background, the inefficiency I of Cinderella for the regarded physical property may be calculated as

$$I = \frac{s_{rej}}{s_{acc} + s_{rej}}, \tag{6.1}$$

where s_{acc} and s_{rej} denote the signal heights determined for pure accepted and pure rejected samples. Attention has to be paid to the fact that both samples need to stem from the same basic population, which is achieved most easily by either using a *mark-only* runs and/or by utilising the pure sample of *filter-active* runs. (Please refer to section 7.5 for instructions how to obtain these samples.) For quantities of interest to physics analysis, the inefficiency is expected to be very low for reasons discussed in section 6.5.

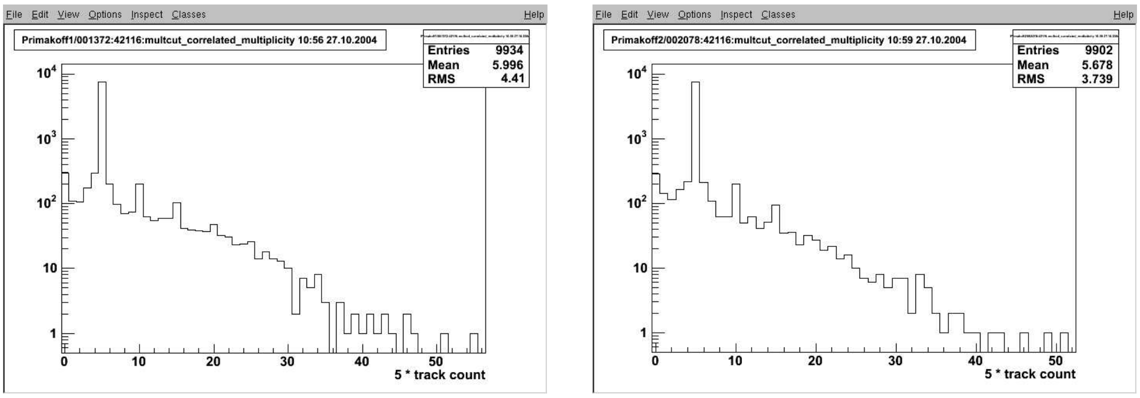

Figure 6.9: First "physics" plots of hadron beam data of 2004, generated from run #42324, the first run with hadron beam and non-empty target (3 mm lead), reproduced as posted in [RLB, comment #11582]. The plots are to be interpreted as *track multiplicity* $\times$ *5* for Prim_1 and Prim_2 triggers, respectively.

Chapter 7

Integration in the COMPASS Experiment

After the previous chapters have given in-depth information on the *functioning* of Cinderella, the chapter at hand takes a practical approach, focussing on *interfaces* rather than internals. Here the information is concentrated that is of vital interest for all members of the COMPASS collaboration, even if they only take casual interest in the online filter.

7.1 Integration in the DAQ

The machine code of Cinderella is executed on the event builder computers `pccoeb12` to `pccoeb21` where it shares CPU, memory and I/O resources with the residing `eventbuilder` processes of DATE. Receiving raw events from the `eventbuilder` processes directly after assembly, the filtered data by Cinderella are written to the local RAID array where they are buffered until collected and transferred to CASTOR by CDR.

The mode of operation of Cinderella is selected by the shift crew in the Run Control software[1]. Errors or warnings encountered by Cinderella are forwarded to the log window of the Run Control (as depicted in Figure 7.1) where they are available to the shift crew for diagnosis. Fatal errors pop up in an extra window that is impossible to miss (see Figure 7.2). Additionally, detailed logging is available on the individual event builder computers in `/tmp/cinderella/logs/`.

For the current run, status information of Cinderella are displayed in the "GDC" section of the Run Control, one column per event builder. A typical screen-shot is reproduced in Figure 7.3 (redundant event builder columns having been omitted). Of the information listed there, the last six lines refer to the online filter and are described in order:

[1] User interface for controlling COMPASS DAQ, originating from ALICE DATE (cf. [ALI99]) but heavily modified for use at COMPASS by Lars Schmitt [Sch04b].

Figure 7.1: Cinderella diagnostic messages in Run Control

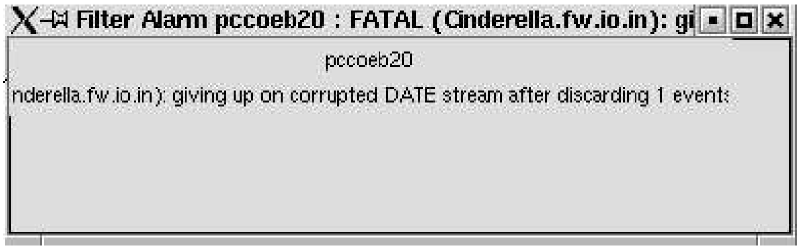

Figure 7.2: Run Control pop-up window for fatal Cinderella errors

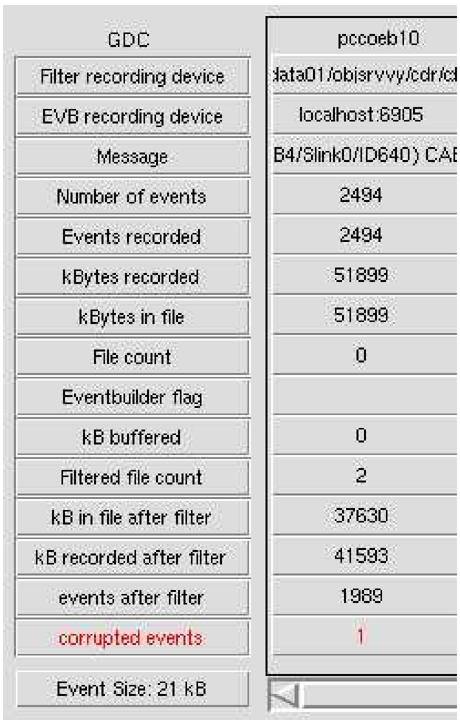

Figure 7.3: Cinderella status information reported in Run Control

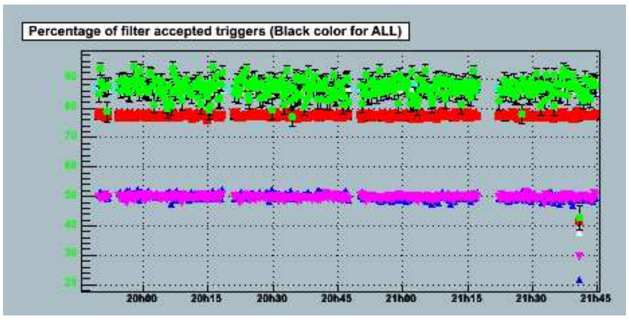

Figure 7.4: Ratio of accepted events in Jiawei's online monitor

kB buffered: shows the amount of derandomisation buffer of Cinderella that is occupied by events waiting in queue to be filtered.[2] During normal operation the buffer level is "breathing" in sync with the SPS duty cycle: In the 4.8 s of spill, the level is rising sharply as data are received faster than they are filtered. During the off-spill time of 12 s duration, the buffer level is receding back to zero with moderate speed as buffered data is processed continuously without new data being received.

Filtered file count: gives the number of output files that the filtered data have been split to (avoiding single files larger than 1 GByte).[3]

kB in file after filter: lists the amount of data that have been filtered and written to disk *for the current output file*.

kB recorded after filter: specifies the *overall* amount of data that have been filtered and stored on disk.

events after filter: quotes the number of events that have been accepted by Cinderella. (This is especially interesting in comparison to the "Number of events" displayed 10 lines higher, which shows the number of events as recorded by DATE *before* they are processed by Cinderella.)

corrupted events: exhibits the number of events that have been rejected due to fundamental errors (eg. faulty sub-event structure).

The acceptance ratios of Cinderella for all triggers are included in Jiawei's online monitor, being displayed in real time during data taking. (Figure 7.4 gives a screen shot.) However it has to be noted that the ratios exhibited there only reflect the *decisions* taken by Cinderella. Whether or not the decisions *are applied* and events actually are discarded depends on the *mode of operation* selected by the shift crew (see next section).

The mode of operation of Cinderella is stored in the `runlb` database along with other meta-information like run number, date and time, and beam line configuration. It may be gathered from the COMPASS Run Logbook [RLB] for later reference and analysis.

[2]For most of 2004 the buffer size was configured to be 131072 kByte (128 MByte).
[3]The first file only contains configuration and status information and thus is much smaller.

mode	input checks	auto-calibration	filtering	cutting
filter-active	yes	yes	yes	yes
mark-only	yes	yes	yes	no
filter-calibration	yes	yes	no	no
pass-through	yes	no	no	no
none	no	no	no	no

Table 7.1: Modes of Operation

To log incidents pertaining to Cinderella the domain "online filter" has been added to the "Comments" section of the Run Logbook.

Finally, a manual for the operation of Cinderella has been added to the documentation folder (and stuck to the wall as well) of control room HNB 411 for reference by the shift crew. It is reproduced in appendix C for completeness.

7.2 Modes of Operation

In the 2004 beam time there have been five different modes of operation of Cinderella. An overview is given in Table 7.1 and a detailed description follows:

filter-active: Filtering (decision taking) and cutting (enforcing decisions) is enabled and auto-calibration is conducted after every run. This is the default setting for physics data taking.

mark-only: All data are kept as decisions are taken and recorded but not enforced. Auto-calibration is executed after every run, just as in **filter-active** mode. This setting is intended for testing of Cinderella or as a safe fall-back if there is any doubt whether Cinderella is working correctly.

filter-calibration: No filtering or cutting are being done, but auto-calibration is active. In principle, the functionality of **mark-only** is duplicated, but the making and recording of hypothetical filter decisions is omitted.

pass-through: Only input checks and buffering are being used. All other functionality is disabled. This mode is useful for detector and trigger tests.[4] It is preferred with regards to **none** since input checking, buffering and blocked writing of Cinderella improve the general stability of data taking and its logging may help to pinpoint detector or DAQ problems.

none: Cinderella is not executed, the data write-out is handled by the DATE `event-builder` process.

[4]However if all ROBs are disabled except for only *one*, this setting cannot be used and **none** should be employed instead.

7.3 Event Modification by Cinderella

Apart from the obvious fact that Cinderella is discarding a certain fraction of events, it should be noted that there is a restricted set of modifications being done by Cinderella to the original "raw" DATE events. Great care was taken not to interfere with existing applications.

Formerly unused, the two highest bits of the DATE trigger mask[5] now are employed by Cinderella to convey the filtering decision and the monitoring flag (cf. section 4.3) for that event. Unfortunately it could not be avoided to break softwares that rely on the implicit assumption that the 20 highest bits of the trigger mask remain unused. Tests of the trigger mask now are required to ignore the two highest bits.[6]

Into the DATE event structure a new equipment with Source ID 1 is subjoined, communicating more detailed filter status information. For every event, the elapsed time during execution and the decision that has been taken by every filter module is logged. Additionally, spill-wise summaries of the same information are stored in end-of-burst events[7]. This mechanism enables fast online monitoring of rejection ratios (only one event per spill and event builder needs to be evaluated) while retaining *full statistics*.

7.4 Detector-Related Aspects

When comparing detector plots of filtered and unfiltered data, it is necessary to bear in mind that the online filter serves to reduce pile-up and noise. This effect should be noticeable in all detectors, yet supposedly it is most outstanding for detectors that actually took part in the filtering algorithm. When comparing two samples with identical number of events, the *filtered* shows better signal-to-noise ratio than the *pristine* sample. This effect is demonstrated exemplarily in the timing plot of scintillating fibre tracking detector FI01X1 in Figure 7.5.

For the detectors that it decodes, Cinderella is reading calibration information from the official calibration database. Care should be taken for the data contained therein to be of high quality. Changes to calibration files during data taking are likely to influence the filtering properties of Cinderella, although this effect is diminished by the automatic calibration of the online filter. Thus updates of the calibration files of detectors that are used by Cinderella should be discussed with the appropriate coordinators beforehand and scheduled for period boundaries, and completely avoided during transversity data taking.

[5] the ninth field of the DATE event header structure as of [Fis03]

[6] To test for events with exclusive Ladder Trigger exemplarily, the formerly possible `(mask == 4)` now yields wrong results. The more thoughtful variant `(mask & 0xfff == 4)` is as correct now as it has ever been.

[7] the last event of a spill

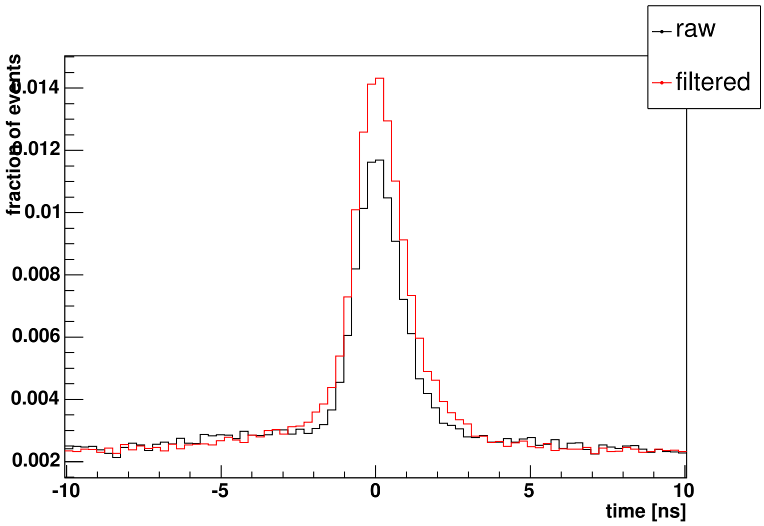

Figure 7.5: Timing plot of FI01X1, run #37928, Inner Trigger (exclusive): Comparison between pristine and filtered samples normalised to event count.

7.5 Physics Analysis

The COMPASS analysis package, Sergei Gerassimov's PHAST, includes two functions to query Cinderella's categorisation of an event. It is important to note that *the output of these functions only may be interpreted in the context of the filtering mode* (cf. section 7.2) of the run!

The function `PaEvent::OnlFltAccepted()` may be used to check the filter decision for a particular event. It returns `true` if the online filter has classified the event as "good" and `false` if Cinderella has rated it as "bad". (It also returns `true`, if the online filter refrained from filtering because the trigger of the event was configured not to be filtered, like for Primakoff events during the hadron run.) Generally speaking, the output of `OnlFltAccepted()` only is useful, when event categorisation has been done by Cinderella, i.e. in runs taken in **filter-active** and **mark-only** modes in events whose trigger was configured to be filtered.

Most events of an mDST file taken in one of the two categorising modes of Cinderella naturally return `true`. (Cinderella and CORAL mostly agree on what they think is worthy to keep.) However for both modes of operation there exist mechanisms by which events considered "bad" by Cinderella may slip into an mDST. Two requirements have to be satisfied for this to happen: Obviously, the event needs to be considered useful by CORAL, but "bad" by Cinderella. This can either happen because CORAL in some aspects has a broader criterion for inclusion than Cinderella (One reconstructed vertex anywhere in the spectrometer is sufficient for CORAL, while the requirements of Cinderella (cf. sections 5.1 and 6.2) are somewhat stronger.), or because Cinderella took a wrong decision, the sum of which constitutes its inefficiency (cf. sections 5.5.3 and 6.6.3). The second requirement is for such an event to get written to tape so that it is processed by CORAL. In **mark-only** mode, certainly this is granted for every event. In **filter-active** mode however, events categorised as "bad" by Cinderella usually are discarded already prior to tape writing. Yet the existence of the monitoring sample ensures that also in this mode some "bad" events are included in the mDST.

The function `PaEvent::OnlFltMonitor()` returns `true` if the event is part of the monitoring sample. It has to be noted, that the result is useful only for runs that were taken in **filter-active** mode. For runs of all other filtering modes, the concept of a "monitoring sample" does not exist.[8]

From the previous explanation it may be inferred that for a run in **filter-active** mode, `OnlFltAccepted()` may only be `false` if `OnlFltMonitor()` is `true`.[9] For runs with **mark-only** setting, all four possible combinations may occur, however the output of `OnlFlt-Monitor()` has no widely useful meaning there.

[8]Though every n^{th} event still is flagged "monitor" in most other modes.

[9]Yet under some *very obscure* circumstances it may happen that a **filter-active** run returns `false` for both functions. This is either when an event is larger than 4 MByte (!) or when a DATE start-of-run event or a TCS first-in-run event contains physics data. Up to now, this condition has been observed only in 23 events of the slot 2 production of run 37420.

Chapter 8

The Cinderella Software Project

The process of preparing programs
for a digital computer is especially attractive,
not only because it can be economically and scientifically rewarding,
but also because it can be an aesthetic experience
much like composing poetry or music.
— Donald E. Knuth, in [Knu67]

At the beginning of the development of Cinderella it was already palpable that the project would need to meet high demands in correctness of physics, stability, and speed, which consequently have become its main goals. Further objectives include reproduce-ability and portability. To achieve these aims, the whole process of software development has been geared towards them, a sustained effort being spent on *design* and *structure*.

This chapter serves as an introduction in the structure and philosophy of the source code of Cinderella.

8.1 Software Engineering

The demands being made on Cinderella first of all are reflected in the selection of programming language. Naturally, the choice was restricted among high-level languages within widespread application. Compiled languages were preferred over interpreted languages as it was felt that their huge advantages in speed of execution—fast by default and including the possibility of optimisation ranging from variation of compilers and their options to employing processor's SIMD[1] capabilities, should the need arise—could not be recouped by the favourable properties of interpreted languages that include rapid

[1] **S**ingle **I**nstruction, **M**ultiple **D**ata, describing a concept of parallel execution of identical commands, is a method to accelerate homogeneous calculations that is present in most modern x86 CPUs (but not restricted to this architecture), best known in the brand names MMX (Intel), 3DNow! (AMD), SSE (Intel), and AltiVec (Motorola).

development and deployment. (In respect of the given argumentation, although being some kind of a hybrid concerning the differentiation, Java needs to be associated with the interpreted languages.)

The preceding argumentation in favour of popular, compiled languages narrows the choice to a ruling between C[2] and C++[3]. The benefit of C surely is its maturity and its simplicity of expression—constituting perfection of the sort that leaves nothing to take away rather than nothing to add—which implies an ever so slight advantage in speed of execution. Yet its heavy reliance on pointers makes C prone to related errors that sometimes are hard to find. In stark contrast C++ sports rich syntax[4], encouraging the creation of well-structured code as especially is desirable for achieving stability in a complex project. Its power of expression however constitutes a two-edged sword as the high inherent complexity of the language also favours more complex errors that are harder to find.

In the end, it came down to a judgement call and it has been the sum of a few small advantages has tilted the scale towards C. Its higher maturity, less complexity and somewhat higher speed seemed more beneficial than the object orientated approach of C++, following the opinion that the structural soundness of a software project much more than on the programming language depends on the determination and discipline of the developers. Further on the existence of the Linux kernel (written entirely in C) proves C to be well suited for development of large software projects showing landmark stability and speed.

As primary compiler the GCC[5] (supported versions range from v.2.95 of 1999 to v.3.3.5 of 2004) was chosen, as it constitutes the de facto standard, possessing a healthy development community and support of a huge number of hardware platforms. A portable coding style was adopted that facilitates the use of different compilers and the targeting of platforms other than x86[6], which is currently used for the event builder computers. This capability is vital to allow later extension of the current event builder cluster to a "filter farm", which may be comprised of different hardware altogether. The prosperity of this effort has been demonstrated by successful tests of Cinderella on Mac OS X running on a PowerPC[7] processor by Roland Kuhn [Ku04].

Cinderella is licensed to the public by its authors under the GNU General Public License (GPL), as it is best suited to meet scientific customs of publication and global collabora-

[2] C was developed as a successor to B by Brian Kernighan and Dennis Ritchie in the early 1970s. An improved version was standardised as ANSI C in 1989 and updated in 1999, termed C99.

[3] C++ was conceived by Bjarne Stroustrup in the 1980s as an improvement upon C, adding features like object orientation, overloading, templates, and exceptions. It was standardised by a joint ANSI-ISO committee in 1998, revised in 2003. While C++ compilers have matured considerably during the last decade, there still is no compiler to fully support the standard.

[4] The difference in complexity is illustrated very well by the fact that the C reference [KR88] is dwarfed by its C++ counterpart[Str04] by approx. a factor of 2 in page size and a factor of 3 in page count.

[5] **G**NU **C**ompiler **C**ollection

[6] The processor architecture most commonly used in desktop PCs, based on the instruction set of the Intel 80386 CPU (first shipped in 1986).

[7] A processor architecture established by IBM and Motorola in the early 1990s that is mainly used in Apple PCs and embedded environments.

tion. The GPL grants permission to anybody to use, modify and redistribute the software on the condition that distributed modifications are licensed with the GPL, too. The full text of the GPL is enclosed with the Cinderella sources and it may be obtained from [GPL].

For the sake of performance and stability, it was decided to refrain from linking to the complex libraries of DAQDataDecoding, CORAL and ROOT. As a direct consequence, Cinderella needed to be written from scratch entirely, duplicating parts of the functionality of the aforementioned libraries.

The sources of Cinderella are available from the subversion[8] source repository[9] or may be browsed at the web site of Cinderella [CIN]. At the time of writing, they comprise ~22,000 lines of code, split about equally among the framework, the filter modules, and infrastructure and tools—the three major sections the online filter is organised in.[10] An overview of the different versions of Cinderella that have been used for physics data taking in 2004 is given in Appendix A.

The main tasks of the framework are input and output of event data, status and error logging, reading of configuration, and management of the filtering process. The filter modules serve as building blocks—providing capabilities as time reconstruction, conditional coincidence evaluation, and coincident multiplicity calculation—that are assembled and executed according to configuration by the framework. The category infrastructure and tools encompasses important auxiliary programs as the Calibrator (managing automatic calibration), Cinderella's own fast histogramming package, a visualisation program, and assorted tools for testing and debugging.

8.2 The Framework

Constituting the central part of Cinderella, serving as a pivotal structure and overseeing administrative tasks, the framework is described by its components in this section:

Event Buffer: The core of the framework of Cinderella is formed by the event buffer[11], also called derandomisation buffer, serving a multitude of purposes. It is accessed concurrently by the input, output, and filter threads, avoiding in-memory copying of data to the highest possible extent. Managed in circular fashion (FIFO), incoming data are inserted at the next free position, events are being processed by filter threads at the next "unfiltered" position and events are being moved to the output buffer[12] from the next "filtered" position, all at the same time.

[8] a novel source code management system, successor to the popular CVS, cf. [BCSP04, SVN]

[9] currently located at `svn+ssh://hamlet.e18.physik.tu-muenchen.de/opt/svn/cinderella`

[10] This statistics was generated using David A. Wheeler's 'SLOCCount'.

[11] Sized 128 MByte during most of 2004, but may be enlarged for the future.

[12] This is the only time that event data is being copied inside of Cinderella. It has been shown that disk write performance is increased distinctly by consolidating the output data into a continuous buffer and handing them over to the operating system in a single instance, compared to issuing many `write()` system calls on segmented data that would be necessary to avoid in-memory copying. The whole issue, having been

I/O System: The most elegant way to avoid losing time waiting for input or output is to profit from the scheduling capacity of the operating system and to establish separate threads dedicated to these tasks. In Cinderella, apart from writing to the event buffer, the input thread has the task of performing a series of checks to the consistency of the (sub)event structure[13]. Likewise, the output thread, additionally to writing the accepted events to disk, handles appendation of status information to the processed events.

Filtering System: The central part of the framework is represented by the filtering system. At startup of Cinderella, the filter modules are initialised according to configuration and an ordered list of filter modules is prepared—dubbed the *filter chain*—following dependency information derived from the configurations of the individual filter modules. During normal operation, the events are passed through the filter chain one after another, the modules being applied in order. As soon as a module of the chain is capable of taking a decision on a particular event (be it positive or negative), processing of the event is ceased, the decision recorded and the next event inserted at the top of the chain. At the end of the run, the filter modules are destroyed in reverse order of their initialisation, dumping histograms of variables of interest.

The filter system is designed to allow parallel execution of multiple filter threads to be able to take full advantage of all CPUs of the host system. This capability was confirmed in mid-September 2004 after the deployment of version 2.0.10, when the online filter was configured to running two parallel filter threads for the rest of the beam time, matching the number of processors of the event builder machines. Yet two is not the upper limit to the number of filter threads and the performance of Cinderella is expected to scale well with the number of processors, at least true for the range of computer systems that may be considered economically sensible to acquire for a possible filter farm.

Watchdog: The duty of the watchdog thread is to monitor the filter threads, and in the unlikely event of their failure, to cancel and re-spawn them. A failure is assumed, when a filter thread takes more than a certain interval (usually 1000 ms) to filter a single event. It has to be stressed that this facility is merely an *additional* measure to improve stability. It is by no means *necessary* for stable operation, proven by the fact that it was only introduced with version 2.0.10 in mid-September 2004 and the stability of Cinderella having been generally acknowledged long before.

Configuration System: Very early during the development of Cinderella it became apparent that a major challenge would be the management of complexity. This perception lead to the conception of a hierarchically structured configuration system, based on XML[14] for configuration file syntax, taking advantage of readily available parsing libraries and a hierarchic structure perfectly matching the topology of

raised by the provision for inclusion of Cinderella status information with the events which enlarges them slightly, however in the framework lays the foundation for data reduction by *event post-processing*.

[13] as defined in [Fis03]

[14] e**X**tensible **M**arkup **L**anguage, a W3C standardised, portable markup language possessing good readability for humans as well as machines

the configuration system. Features include representation of integers, floats and strings, as scalars or in arrays.

Messaging System: For purposes of diagnosis and logging, a likewise hierarchically structured messaging system was devised. Messages are assigned priorities ranging from `DEBUG3` (lowest) to `FATAL` (highest). Configuration of the desired level of logging may happen at compile-time or at run-time, the former bearing the advantage of totally eliminating the logging code for messages below threshold, allowing for detailed debugging in the tightest of loops without harming the performance after recompiling with higher threshold. Possible targets for the logging stream are files (including `stdout`) and the messages window of the COMPASS Run Control.

DATE Interface: For communication with DATE[15], its shared memory interface is used to transfer meta-information like the status of Cinderella to the Run Control or to request the end of run in case of fatal errors.

8.3 The Filter Modules

The general design concept of filter modules foresees data input, data output, and decision taking for all modules. The iteration of the framework over the filter chain may be envisioned as pipelining input and output of successive modules, handing over event data from one to another, expanding the comprehension of the event with every step. The decision of every module is expressed in ternary logic. Either the event is accepted *or* rejected and in either case processing ended, *or* the module concludes that it cannot take a decision and defers the responsibility to the next module which continues evaluation, aided by the output data of the previous module. (When the last module of the filter chain returns an undecided vote, the event is accepted by the framework, faithful to the device *in dubio pro reo*.)

This scheme is sufficiently general to allow for decoding modules (that only process data, always returning "undecided") and decision taking modules (that produce no output and always return positive or negative decision). Yet the universality of the design pays off, as a module designed for decoding may reject single events due to fatal shortcomings, or a module designed for decision taking may defer the decision to the next decision module chained behind it.

The now following descriptions of the individual modules (roughly in order of execution) are kept rather brief since most of the filter modules already have been elucidated in previous chapters.

Gen_TOC: This module always is set to the first position of the filter chain as all other modules depend on it, directly or indirectly. The **Gen_TOC** module is parsing the hierarchical event structure, beginning with DATE event and sub-event structure,

[15]**D**ata **A**cquisition and **T**est **E**nvironment, cf. [ALI99]

continuing through equipment blocks and S-Link multiplexing down to the individual CATCH blocks. The Source IDs of and the pointers to the CATCH blocks then are stored together, creating a *table of contents* (TOC) permitting rapid, random access to the individual detector's data by the subsequently executed decoding modules. (The assignment of Source IDs to detectors by the decoding modules is read from the mapping files.)

Trigger: The **Trigger** module serves to decode the time of the trigger, both in high and low resolution (multiples of 64 ps and 128 ps, respectively). Additionally, the trigger mask is parsed and the TCS[16] phase[17] is determined. The information gathered is stored for access by subsequent filter modules.

Fibre: The task of the **Fibre** module is the decoding of time information of hodoscopes and scintillating fibres (based on TDC[18] read-out). For use with the hadron pilot run, clustering was implemented specially for hodoscopes. The output of this module is a list of time and error pairs, the error being constant, either computed from the width of the timing peak during auto-calibration or read from configuration file. The time is quoted relative to trigger time.

Silicon: This module has the role of decoding the time information of the silicon microstrip detectors (based on APV[19] read-out) and includes a custom-tailored clustering algorithm. Similarly to the **Fibre** module, the output is returned as list of time and error pairs, time values specified relative to trigger time. However in this module, the error in time is computed individually for every cluster.

Anticipating its future utilisation in CORAL and COOOL, the module is written in a fashion that allows easy disjointing from Cinderella and re-using in a different environment. The underlying algorithm for time reconstruction is detailed in section 3.5.

Additionally, this module contains provisions for *simulating* silicon digits, which during preparation for the hadron beam served as substitute for the at that time not yet existing hadron data.

TriggerCut: Here a first cut of the time and error pairs, gained from the previously described decoding modules, is undertaken. Time and error pairs outside a fixed window around trigger time may be weeded out. More selective however, time and error pairs may be discarded that fail to approach the trigger time by less than a (configurable) multiple of their own errors. The **TriggerCut** module takes time and error pairs for input as well as for output, reducing their numbers to a varying degree, depending on configuration.

TimeCut: This module implements decision taking based on the conditional coincidence algorithm, evaluating the coincident completion of a configurable logical condition

[16]**T**rigger **C**ontrol **S**ystem, cf. section 2.2.5

[17]offset of the trigger time relative to the last tick of the TCS clock, necessary for time reconstruction of silicon and GEM detectors

[18]**T**ime to **D**igital **C**onverter, cf. [Bra99, Fis01]

[19]**A**nalogue **P**ipeline, **V**oltage-type, cf. [Gru01]

of multiple detector planes, as described in section 5.2. Its input are time and error pairs, usually somewhat diminished after traversing the **TriggerCut** module. The sole output is the decision.

MultCut: The decision taking grounding on determination of coincident multiplicity, as described in section 6.3, is implemented in this module. Similar to the previous module, its input are time and error pairs, the output being the decision.

Prescaler: This module represents the precise duplication of a hardware prescaler, accepting every n^{th} event and rejecting the rest. It has been conceived exclusively for debugging purposes, yet as an example of a *minimal* filter module it may serve as a starting point for colleagues without previous experience with Cinderella who would like to implement a filter module.

8.4 Infrastructure and Tools

The category of *infrastructure and tools* contains all parts of Cinderella that neither are filter modules nor vital to the concept of the framework. The *infrastructure* intentionally was not subjoined with the *framework* to emphasise the effort of building a modular structure, different functionalities being segregated to different units, yielding distinct advantages in maintainability. It is by no means implied that elements of this section may be considered dispensable.

Calibrator: Very much like the framework of Cinderella, the **Calibrator** represents the framework for automatic calibration. Its task is listening on network sockets[20], accepting data from the Cinderella instances running on the event builder computers, combining it and passing it to the auto-calibration modules that are constituted by dedicated parts of the corresponding decoding modules.[21] The results of automatic calibration, referred to as *refinements*, are stored in the `runlb` database in the table `tb_calib_refined`, from where they may be read during the next initialisation of Cinderella.

Histogramming System: This subsystem consists of a light-weight implementation of one- and two-dimensional histograms with integer bins one, two or four bytes wide. It includes compressed I/O to file and database and provides conversion functions to `TH1` and `TH2` objects of ROOT which are used for visualisation.

cat_date: The Swiss army knife for operations on DATE streams, **cat_date** contains provisions for their concatenation, truncation, selection (by trigger mask or burst number) and reduction (stripping undesired Source IDs). A valuable tool for testing of Cinderella, its versatility and especially its reduction capability render it profitable for other applications, eg. in detector analysis.

[20] During the 2004 data taking period, the **Calibrator** was executed on either `pccoeb01` or `pccoeb02`.
[21] At the time of writing, only the **Fibre** module contains provisions for automatic calibration.

hist_draw: This multi-purpose visualisation tool may be used for display of histograms generated by Cinderella (reading them from file or database) as well as graphs and histograms created on the fly from database queries. A plethora of command-line switches allows tiled and superimposed arrangement of plots and serves to control many other aspects of layout and presentation. Most figures in this thesis were generated using **hist_draw**.

Chapter 9

Conclusion and Outlook

In 2004 the online filter Cinderella for the first time has been utilised for regular data taking of the COMPASS experiment. Two entirely different filtering criteria have been applied to muon beam data and hadron beam data, respectively. To ensure safe operation of the online filter, a procedure of automatic calibration of scintillating detectors has been adopted. For purposes of analysis and verification, a pure sample of 3% of all events has been kept regardless of the filter decision. The stability of the Cinderella software proved to be remarkable.

During the muon run, the ability to reconstruct track and momentum of a beam particle was asserted by Cinderella for events to be kept. This was accomplished by evaluation of coincident hit multiplicities in 22 detector planes of scintillators and silicons. With inefficiency well below 0.4% for all sub-triggers, the main goal of *conservation of meaningful data* to avoid distortion of the delicate asymmetry measurements, has been reached without cutting back. At the same time, the overall data rate could be reduced by 23%, saving 90 TByte of raw data, corresponding to approx. 57,000 SFr in expenses for tape media.

In the hadron pilot run, the track multiplicity after the target was used as filtering criterion for the Diffractive Trigger. Time information of hits to 12 planes of silicon micro-strip detectors positioned after the target was analysed and the track multiplicity extracted as sophisticatedly truncated mean of plane hit multiplicity. In that way, the rate of the Diffractive Trigger could be reduced by 45%, allowing for removal of the previously existing pre-scaling factor of 2, during operation of Cinderella effectively doubling the statistics taken with that sub-trigger. An algorithm for reduction of the Primakoff sub-triggers had been ready for deployment but was not utilised because the intended trigger rates could just about be handled without.

As a prerequisite, a method for precise reconstruction of silicon time was implemented, exploiting the known shape of the signal together with clustering to achieve an accuracy of time measurement well below the sampling interval of 26 ns. In normal operation of Cinderella the time resolution was approx. 4 ns (n-side) and approx. 5 ns (p-side). (With specific cuts however, much lower values have been observed.)

With the current state of Cinderella, a sound foundation has been laid for online filtering in the COMPASS experiment that may be built upon in the future. The inclusion of *tracking* in the region upstream of the first spectrometer magnet is envisioned to allow halo suppression and improved after-target track counting. Event compression by (partial) in-filter decoding or noise suppression may be imagined for calorimeters or the RICH.

Its structured and flexible design have prepared Cinderella to run on a ROB, on a large multi-processor machine or in the environment of a filter farm. It has the potential very well to be a key player in the effort of increasing trigger rate after 2005, provided that its development is continued at high pace.

Appendix A

Cinderella Versions

The following table lists all versions of Cinderella that have been used for physics data taking in 2004. The first column gives date of installation, the second lists the first run taken with the new version. The version number is specified in the third column while the fourth column quotes the subversion[1] revision which serves—very much as a CVS tag would do—to uniquely identify the version in the source repository. In the last column a short description of changes introduced with the new version is contained.

Date	Run	Version	Revision	Comment
2004-06-26	36744	2.0.2	1241	first version suitable for production use
2004-06-28	36837	2.0.3	1244	maintenance release, bugfix for auto-calibration
2004-07-08	37315	2.0.4	1253	maintenance release, improved startup speed
2004-07-09	37357	2.0.5	1255	maintenance release, improved startup speed
2004-07-12	37466	2.0.6	1256	maintenance release, improved stability for malformed event data
2004-07-14	37544	2.0.7	1260	maintenance release, improved error reporting
2004-07-22	37799	2.0.8	1269	maintenance release, improved input checks and logging
2004-09-16	40617	2.0.10	1377	major maintenance release, enabled parallel filtering in two threads, added watchdog thread, improved refinement selection (auto-calibration)
2004-10-28	42475	2.1.1	1429	some new implementations for hadron beam data

[1] a novel source code management system, successor to the popular CVS, cf. [BCSP04, SVN]

Appendix B

Time-Line

This table lists selected dates that have been mentioned in this text or have been significant for Cinderella in some other way. (It is not intended as a chronicle of the COMPASS experiment.)

Date	Run	Comment
2004-05-14	34558	First 2004 “physics” run
2004-06-27	36744	First run with *filter-active*
2004-07-07		Running at $\sim 1/2$ SPS intensity due to septum problem
2004-07-15 01:00		trip of QUAD33
2004-07-20		End of septum problem
2004-07-22	37544	Cuts tightened by 0.75 sigma
2004-08-13 16:53	38991	First “Transversity” Run
2004-09-05	39988	Last “Transversity” Run
2004-09-25		beam line magnet control problem
2004-10-04	41397	Last muon run before change to hadron
2004-10-27	42321	First “Primakoff” Run
2004-10-27 03:28	42324	First Hadron Run with target
2004-11-09 14:00	43036	Target empty
2004-11-09 15:24	43037	First Cinderella Hadron Run
2004-11-09 21:38	43052	Target Copper 3.55mm
2004-11-10 11:47	43084	Diff_1: increased pre-scale factor from 1 to 2
2004-11-10 14:16	43086	Target Carbon 23.5mm
2004-11-10 14:18	43087	Diff_1: reduced pre-scale factor back from 2 to 1
2004-11-11 11:27		Target Lead 2mm + 1mm
2004-11-11 12:12	43132	Running Muons in Hadron setup
2004-11-12 14:45		Target removed
2004-11-12 16:09	43188	Last Muon run in Hadron setup
2004-11-12 17:09	43189	Target Lead 2mm + 1mm
2004-11-12 17:09	43189	First run after switch back to hadron
2004-11-13 15:21	43247	Electron Converter in
2004-11-13 17:33	43251	Electron Converter out
2004-11-14 22:10	43323	Last 2004 “physics” run

Appendix C

Online Filter Shift Crew Instructions v2.2

(For the sake of completeness the manual for the COMPASS shift crew regarding operation of the online filter is reproduced here in its latest revision.)

C.1 Filter Configurations

filter-active: During normal physics data taking, use this setting. Filtering is active and Cinderella is auto-calibrating continuously.

filter-calibration: If Cinderella refuses to start with message "A FILTER CALIBRATION RUN IS NEEDED!"[1] take two runs using this setting and then switch back to *filter-active*. During calibration, filtering is disabled and Cinderella solely runs for the purpose of calibrating itself.

mark-only: If there is a problem running the filter in *filter-active* mode, switch back to *mark-only*. This is a safe default, in which all filtering steps are done, but instead of discarding bad events they are tagged as bad and written to disk anyway.

pass-through: This is the right setting for running of detector or trigger tests[2]. Only some input checks (see below) are executed, no decoding and no beam track filtering takes place. Running with *pass-through* is better than using *none* because Cinderella's disk write buffering and input checking is improving the stability of the DAQ.

none: Only in the case that Cinderella doesn't work in pass-through, this setting should be used. Then Cinderella is switched off completely.

[1] You might need to scroll down in the runControl messages window to find that notice.

[2] If only one ROB is used, the setting *none* must be used.

C.2 Filter Monitoring

C.2.1 Jiawei's Online Monitor

After every spill, the filter accept ratio for every trigger is displayed in the online monitor. (The black points show total accept ratio and not that of Inner Veto.) When is spectrometer is in stable operation, the accept ratios should also be fairly stable and not vary by more than a few percent points. (Except Middle Trigger, where statistical fluctuations of 10% can be observed.)

Any spikes in trigger rate will cause corresponding spikes in rejection rate. This is perfectly alright and an indication for the fact that the filter is working as intended. But if there is a jump by more than 5% in one of the accept rates and it lasts longer than a few spills, most probably it is a problem in one of the detectors in beam telescope (SI01-03, FI01-02) or beam momentum station (BM01-06) or a trigger problem. Naturally Cinderella is quite sensitive for problems in these areas.

If the problem is not resulting from detector or trigger, Cinderella should be switched to *mark-only* mode and an DAQ-on-call expert be contacted.

C.2.2 runControl EVB window

To protect itself and also the software which will read the raw data, Cinderella has an input filter which immediately discards corrupted events. Event integrity is checked down to the SLink level, which means that event size mismatches (which may happen when a CATCH goes south) will lead to the rejection of the event. The count of events discarded is displayed as Error Count in the runControl Status Display and the type of corruption is printed in an error message in the runControl log. As such errors tend to persist until the end of spill, the Error Count in the runControl Status Display usually increases by a few thousand. This is, however, no reason to stop the run, unless the problem does not go away at the next start of spill.

C.2.3 runControl Messages window

The most interesting place to look for messages is the log window in the runControl. Only error messages are printed there, so normally you should not see anything from Cinderella. However, if anything goes wrong, check that log for errors first (this might involve some scrolling). If some equipment had errors, it will be named explicitly in the message. Other sources for Cinderella logs are (in increasing order of verbosity) `/pro/site/main-2004/logFiles/filter`, which includes log messages from all Cinderella instances, and finally the local log on each EVB, which is in `/tmp/cinderella/logs`[3].

[3] The link `last` points to the last opened log. To view the log of run XXXXX, use "`less *:XXXXX:*`".

Even if there are no log messages in runControl, always have an eye on the online monitoring screen!

C.3 What if the run does not start?

If the start of run is aborted because of the Cinderella, there will be an error message indicating the reason in the log window of the runControl. This will happen when Cinderella needs calibration. Common causes for such a lack of filter-calibration are no beam for more than 6 hours, a detector calibration file in the COMPASS calibration database having been updated or a detector calibration file missing from COMPASS calibration database. In the first two cases a *filter-calibration* run must be taken (see below). In the last case the responsible detector group has to provide a calibration, and until then Cinderella cannot run.

C.4 How to do a filter auto-calibration run

Switch the "Filter Config" to *filter-calibration* and take two runs with normal physics conditions. These don't have to be 200 spills long, 50 spills each should be enough. Then try to switch back to *filter-active*. If it still complains about missing calibration refinements, call an expert[4].

C.5 Cinderella background information

For those who don't know yet: Our online filter (aka 2$^{\text{nd}}$ level trigger) is named after a person from a fairy tale—Cinderella—who has among other tasks to sort the good peas from the spoilt ones.

C.5.1 Input filter

Another feature of the input filter is the rejection of events which contain only data from one ROB as this signals an eventbuilding failure during normal operation. This failure usually is due to missing data from one or more crashed/hanging ROBs, so the run should already have been stopped automatically in this case. If not, you can try to stop it, but be patient, the eventbuilder will then have to process a lot of (corrupted) data until all is flushed from the functioning ROBs.

It is considered an error if an event only has information from one ROB, so if you really want to take data with only ***one ROB****, be sure to switch the Filter Config to* none*!*

[4]The calibrator is running on `pccoeb02` and has a log similar to Cinderella in `/tmp/calibrator/logs/last`. You can try to check which detector gives trouble and investigate further by looking at the COOOL timing histograms for BMS, FI01/02 and Silicon

C.5.2 Beam track filtering in a Nutshell

At the moment, Cinderella does no tracking to find a beam track. All it does is correlate hit times from the BMS and the beam telescope (FI01/02 and SI01/02/05) to find "time clusters". There are minimum requirements on such a time cluster, which ensure that enough BMS measurements (two in total, at least one of them before BEND6) are available for momentum determination and enough position information is measured for beam tracking. For the latter, different settings are used for different triggers, with different time windows and minimum hit counts, but the principle is that typically three out of the four projections ($0°$, $5°$, $90°$, $95°$) have enough hits to support a space track.

Because of this algorithm timing is extremely important for a stable operation of the filter.

C.5.3 Filter auto-calibration

To allow for calibration shifts (day/night, etc.) and to ensure that the timing calibrations are always up-to-date, Cinderella has an auto-calibration procedure. After each run the accumulated raw timing histograms are sent from each EVB to the calibrator process, which merges them and derives a new calibration. For scintillating fibre type detectors, this can be easily done via a fit to the T_0-peak, because the background is very flat. For the silicon detectors, this procedure is not yet ready, so the calibrator only looks for the maximum in the histogram and checks that it is consistent with the old calibration. You can look at these histograms with `/date/cinderella/bin/hist_draw`[5], they are stored at `pccoeb02:/tmp/calibrator/histograms`.

Please note that the fitting procedure takes quite some time, so in order not to lengthen the pause between two runs, the auto-calibration is done asynchronously. This means that the new calibration values will only be available for the run after the next run. Currently the fitting is started when the calibration data from the next run arrive (i. e. when the next run ends) or after 15 minutes, whatever comes first.

[5]This probably does not work from the `onl` account, which is severely broken. Try with your own account and set ROOTSYS to /afs/cern.ch/sw/root/v3.10.02/rh73_gcc296/root

Bibliography

[Ada98] G. S. Adams et al. Observation of a new J(PC) = 1-+ exotic state in the reaction $\pi^- p \to \pi^+\pi^-\pi^- p$ at 18 GeV/c. *Phys. Rev. Lett.* **81**, 5760–5763 (1998).

[Ale95] Yu. Alexandrow et al. CHEOPS: CHarm Experiment with Omnipurpose Setup. *CERN SPSLC 95-22, SPSLC I202* (1995).

[Ale04] Vadim Alexakhin. Longitudinal polarization of Λ and $\bar{\Lambda}$ hyperons in deep-inelastic scattering at COMPASS. *Conference Contribution to SPIN 2004, 16th International Spin Symposium, Trieste* (2004).

[ALI99] CERN ALICE DAQ group. ALICE DATE User's Guide, DATE 3.5. *Alice Internal Note 1999-46* (1999).

[Alt04] C. Alt et al. Evidence for an Exotic S = -2, Q = -2 Baryon Resonance in Proton-Proton Collisions at the CERN SPS. *Phys. Rev. Lett.* **92 4**, 042003 (2004).

[Ant85] Y. M. Antipov et al. Experimental estimation of the sum of pion electrical and magnetic polarizabilities. *Z. Phys.* **C 26**, 495–497 (1985).

[Bar98] P. Barberis et al. Test of front-end chip for COMPASS MWPC. *COMPASS Note 1998-9* (1998).

[Bau96] G. Baum et al. COMPASS: A proposal for a COmmon Muon and Proton Apparatus for Structure and Spectroscopy. *CERN SPSLC 96-14* (1996).

[BCSP04] Brian W. Fitzpatrick Ben Collins-Sussman and C. Michael Pilato. *Version control with subversion*. O'Reilly, Inc., 2004.

[Ber04] Colin Bernet. The gluon polarization $\Delta G/G$ at COMPASS. (2004).

[Boh92] Markus Bohn. *Analyse Kernphysikalischer Experimente auf Parallelrechnern (Transputern)*. Diploma thesis, Johannes Gutenberg-Universität Mainz, 1992.

[Bov82] C. Bovet et al. The Cedar counters for particle identification in the SPS secondary beams: A description and an operation manual. *SPS-Note-82-40* (1982).

[Boy97] Owen Boyle et al. The S-LINK Interface Specification, .

[Bra88] Gilles Brassard & Paul Bratley. *Algorithmics – Theory and Practice*. Prentice-Hall Inc., 1988.

[Bra99] G. Braun et al. F1: An eight channel time-to-digital converter chip for high rate experiments. (1999).

[Bro04] Grzegorz Brona. *Search for pentaquark states in the COMPASS experiment*. Diploma thesis, Uniwersytet Warszawski, 2004.

[CAS] CASTOR Web Site, `http://castor.web.cern.ch/castor/`.

[CDR] CDR Web Site, `http://cdr.web.cern.ch/cdr/welcome.html`.

[CER] CERN Web Site, `http://www.cern.ch/`.

[Cho03] S.-K. Choi. Observation of a Narrow Charmoniumlike State in Exclusive $B^{\pm} \rightarrow K^{\pm}\pi^{+}\pi^{-}J/\psi$ Decays. *Phys. Rev. Lett.* **91**, 262001 (2003).

[CIN] Cinderella Online Filter Web Site, `http://www.e18.physik.tu-muenchen.de/research/compass/cinderella/`.

[Col05] Marialaura Colantoni. Measurement of electric and magnetic polarizabilities with Primakoff reaction at COMPASS. *Czech. J. Phys.* **55** (2005).

[COM04] The COMPASS Collaboration. COMPASS Status Report 2004. *CERN SPSC 2004-011, SPSC-M-714* (2004).

[COO] COOOL Web Site, `http://cbernet.home.cern.ch/cbernet/Coool/`.

[COR] CORAL Web Site, `http://coral.cern.ch/`.

[DeM04] Rita De Masi. *Development of a cryogenic silicon detector system and study of strange particle production in deep inelastic scattering*. Ph. D. thesis, Technische Universität München, 2004.

[dHo02] Nicole d'Hose et al. Possible measurements of GPDs at COMPASS. *Workshop Contribution to Future Physics @ COMPASS, CERN, 2002* (2002).

[Din04] Anna-Maria Dinkelbach, `Anna-Maria.Dinkelbach@e18.physik.tu-muenchen.de`. Private communication, 2004.

[Dor02] Valery Dorofeev. Simulation of the eta-pi Diffractive Production and Detection at COMPASS. *Workshop Contribution to Future Physics @ COMPASS, CERN, 2002* (2002). `http://wwwcompass.cern.ch/compass/publications/2004_yellow/`.

[Fau04] Peter Fauland. *The COMPASS Experiment and the RICH-1 Detector*. Ph. D. thesis, Universität Bielefeld, 2004.

[Fer05] Andrea Ferrero. Soft hadronic interactions in the COMPASS experiment. *Czech. J. Phys.* **55** (2005).

[FII] Fraunhofer Institut Integrierte Schaltungen. MPEG Audio Layer-3, `http://www.iis.fraunhofer.de/amm/techinf/layer3/`.

[Fis01] Horst Fischer et al. Implementation of the dead-time free F1 TDC in the COMPASS detector readout. *Nucl. Instrum. Meth.* **A461**, 507–510 (2001).

[Fis02] Horst Fischer et al. CATCH Users Manual (Draft). *COMPASS Note 2002-7* (2002).

[Fis03] Horst Fischer et al. The COMPASS Online Data Format Version 4. *COMPASS Note 2002-8* (2003).

[Fri04a] Jan Friedrich, `Jan.Friedrich@ph.tum.de`. Private communication, 2004.

[Fri04b] Jan Friedrich. Measurement of transverse Λ polarisation at COMPASS. *Conference Contribution to SPIN 2004, 16th International Spin Symposium, Trieste* (2004).

[Fuc03] Anna-Maria Fuchs (now married Dinkelbach). *Setup of a low Temperature Silicon Detector for the COMPASS Experiment*. Diploma thesis, Technische Universität München, 2003.

[GPL] GNU General Public License, `http://www.gnu.org/copyleft/gpl.html`.

[Gri15] Jakob & Wilhelm Grimm. *Kinder- und Hausmärchen*. Gutenberg Projekt, 2004, first edition compiled 1812-1815, `http://projekt.gutenberg.de/`.

[Gri84] Jakob & Wilhelm Grimm. *Household Tales*. Project Gutenberg, 2004, EText #5314, first German edition compiled 1812-1815, translated to English by Margaret Hunt in 1884, `http://www.gutenberg.org/`.

[Gru01] Boris Grube. *The Trigger Control System and the Common GEM and Silicon Readout for the COMPASS Experiment*. Diploma thesis, Technische Universität München, 2001.

[Huf52] David A. Huffman. A Method for the Construction of Minimum-Redundancy Codes. *Proceedings of the I.R.E.* , 1098–1102 (1952).

[Ilg03] Christoph Ilgner. *Fertigung und Inbetriebnahme einer Strohdriftkammerstation für das COMPASS-Experiment*. Ph. D. thesis, Ludwig-Maximilians-Universität München, 2003.

[Jac99] John David Jackson. *Classical Electrodynamics*. John Wiley & Sons, Inc., 1999.

[Ket03] Bernhard Ketzer et al. A fast tracker for COMPASS based on the GEM. *Nucl. Phys. Proc. Suppl.* **125**, 368–373 (2003).

[Knu67] Donald E. Knuth. *The art of computer programming*, volume 1. Addison-Wesley, Inc., 1967.

[Kon01] Igor Konorov et al. The Trigger Control System for the COMPASS experiment. *Nuclear Science Symposium Conference Record, 2001 IEEE* **1**, 98–99 (2001).

[Kon05] Igor Konorov, `Igor.Konorov@cern.ch`. Private communication, 2005.

[KR88] Brian W. Kernighan and Dennis M. Ritchie. *The C programming language*. Prentice-Hall, Inc., 2nd edition, 1988.

[Ku04] Roland Kuhn, `Roland.Kuhn@cern.ch`. Private communication, 2004.

[Ku05] Roland Kuhn. *Thesis In Preparation*. Ph. D. thesis, Technische Universität München, 2005.

[Lam00] Bodo Lampe and Ewald Reya. Spin physics and polarized structure functions. *Phys. Rept.* **332**, 1–163 (2000).

[Leb02] Mario Leberig. *Das COMPASS-Triggersystem zur Messung des Gluonbeitrags ΔG zum Protonenspin*. Ph. D. thesis, Johannes Gutenberg-Universität Mainz, 2002.

[Leo94] W. R. Leo. *Techniques for Nuclear and Particle Physics Experiments*. Springer, 1994.

[M2] M2 Beam Line, `http://sl.web.cern.ch/SL/eagroup/beams.html#m2`.

[Mag02] Alain Magnon et al. Tracking with $40 \times 40\,\mathrm{cm}^2$ MICROMEGAS detectors in the high energy, high luminosity COMPASS experiment. *Nucl. Instrum. Meth.* **A478**, 210–214 (2002).

[Mal04] Gerhard Mallot, `Gerhard.Mallot@cern.ch`. Private communication, 2004.

[Moi03] M. A. Moinester et al. First Observation of Doubly Charmed Baryons. *Czech. J. PHys.* **53**, B201–B213 (2003).

[Mor78] J. J. More. The Levenberg-Marquardt Algorithm: Implementation and Theory. *Lecture Notes in Mathematics* **630** (1978).

[Nak03] T. Nakano et al. Evidence for a Narrow S = +1 Baryon Resonance in Photoproduction from the Neutron. *Phys. Rev. Lett.* **91**, 012002 (2003).

[Nap95] E. Nappi et al. Semi-inclusive Muon Scattering from a Polarised Target. *CERN SPSLC 95-27, SPSLC I204* (1995).

[Ney04] Damien Neyret. Results on exclusive ρ_0 production from COMPASS. *Conference Contribution to SPIN 2004, 16th International Spin Symposium, Trieste* (2004).

[Pag04] Paolo Pagano. Measurements of Collins and Sivers asymmetries at COMPASS. *Conference Contribution to SPIN 2004, 16th International Spin Symposium, Trieste* (2004).

[Pei92] A. Peisert. *Instrumentation on High Energy Physics*, chapter Silicon Microstrip Detectors. World Scientific Publishing Co., Singapore, 1992.

[Pes04] D. Peshekhonov. Inclusive spin-dependent asymmetry A_1^D. *Conference Contribution to SPIN 2004, 16th International Spin Symposium, Trieste* (2004).

[PHA] PHAST Web Site, `http://ges.home.cern.ch/ges/phast/`.

[Pre04] Jörg Pretz, `pretz@mail.cern.ch`. Private communication, 2004.

[Pri51] Henry Primakoff. Photo-Production of Neutral Mesons in Nuclear Electric Fields and the Mean Life of the Neutral Meson. *Phys. Rev.* **81**, 899 (1951).

[RLB] COMPASS Run Logbook, `http://wwwcompass2.cern.ch/runLogbook/dirphp/`.

[Sch04a] Christian Schill. Measurement of the gluon polarisation $\Delta G/G$ at COMPASS. *Conference Contribution to SPIN 2004, 16th International Spin Symposium, Trieste* (2004).

[Sch04b] Lars Schmitt, `Lars.Schmitt@ph.tum.de`. Private communication, 2004.

[Sch04c] Lars Schmitt et al. The DAQ of the COMPASS Experiment. *IEEE Transactions on Nuclear Science* **51 3**, 439–444 (2004).

[Sim01] Frank Simon. *Commissioning of the GEM Detectors in the COMPASS Experiment*. Diploma thesis, Technische Universität München, 2001.

[SINa] SINTEF Web Site, `http://www.sintef.no/`.

[SINb] Norway SINTEF, Oslo. Technical Specifications HERA-B Double Sided.

[Str04] Bjarne Stroustrup. *The C++ programming language*. Addison-Wesley, Inc., 2004.

[SVN] Subversion Web Site, `http://subversion.tigris.org/`.

[Tak02] Naoki Takabayashi. *Polarized target for the measurement of the gluon contribution to the nucleon spin in the COMPASS experiment*. Ph. D. thesis, Nagoya University, 2002.

[vH02] Martin Frhr. v. Hodenberg. *A First Reconstruction of COMPASS Data*. Diploma thesis, Albert-Ludwigs-Universität Freiburg, 2002.

[vH04] Martin Frhr. v. Hodenberg, `mvhodenb@axfr01.physik.uni-freiburg.de`. Private communication, 2004.

[Wag01] Robert Wagner. *Commissioning of Silicon Detectors for the COMPASS Experiment at CERN*. Diploma thesis, Technische Universität München, 2001.

[Web04] Richard Webb. *First Measurements of Transverse Spin Asymmetries through Single Pion Production at the COMPASS experiment*. Ph. D. thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg, 2004.

[Wei03] Quirin Weitzel. *Triple GEM Detectors in COMPASS*. Diploma thesis, Technische Universität München, 2003.

[Wie04] Michael Wiesmann. *A Silicon Microstrip Detector for COMPASS and A First Measurement of the Transverse Polarization of Λ^0-Hyperons from Quasi-Real Photo-Production*. Ph. D. thesis, Technische Universität München, 2004.

[Zha04] Jiawei Zhao, `zhaojw@mail.cern.ch`. Private communication, 2004.

Own Contributions

The beginnings of the online filter date back to the COMPASS proposal of 1996, yet practical work on it did not begin until November of 2002, on when Roland Kuhn and I—in the position of a working student at that time—chalked up the first structure and wrote the first lines of code. In 2003 we continued to share the workload, each of us making a part-time effort so that in January of 2004, when I continued the work in the form of a diploma thesis, Cinderella was functional as a whole and had been tested in *mark-only* mode during parts of the 2003 data taking. At that point Roland ceased immediate work on Cinderella to be able to address the topic of his Ph. D. thesis, so that in 2004 the burden of work rested on my shoulders almost entirely. Yet I still could rely on his supervision during that time, in addition to that of Dr. Lars Schmitt and Prof. Stephan Paul.

As the COMPASS experiment was evolving (introducing additional planes to BMS and beam telescope) in 2004, Cinderella needed to be adapted to this changes. At the same time, the online filter lacked sophistication in many areas, a condition which did not destine it for production use, yet.

To allow Cinderella to be adapted flexibly to the changing demands of the experiment, I conducted a major overhaul of the topology of the online filter, adding layers of abstraction and increasing generality. In a next step, I created the module for time reconstruction of the silicon detectors. Further on, after many fruitful discussions with Roland, I designed and implemented the Conditional Coincidence Algorithm and the Coincident Multiplicity Algorithm, which constitute the back-bone of decision taking for data of muon and hadron beam, respectively.

I presented the online filter in plenary talks at COMPASS collaboration meetings in Paris and at CERN. I took three weeks of duty as expert-on-call for DAQ at the experiment at CERN (available day and night), most of which I spent on matters only peripherally related to Cinderella. My work on reduction of run startup/shutdown time (increasing beam utilisation by $\sim$2%) was highly appreciated by the colleagues. For setup of Cinderella I spent another week at CERN after which the online filter was enabled for the remaining three quarters of muon beam time. I was recognised as primary contact person for all issues concerning Cinderella among the COMPASS collaboration.

For optimised configuration of Cinderella I have devised a scheme of automatic variation and evaluation of combinations of configuration options. As no previous data had been available in the run-up to the hadron beam, I implemented a detailed simulation of the silicon detectors to assist the design and development of filtering algorithms for hadron beam data.

All analysis presented herein has been conducted by myself, with the exception of two plots in the "Silicon" chapter which have been labelled as such.

The underlying ideas to the filtering *criteria* (but not to their implementation) have been spurred by Prof. Stephan Paul and Dr. Lars Schmitt who possess a much deeper understanding of the experiment and its goals than I am capable of.

The framework of Cinderella in 2003 was implemented by Roland Kuhn alone, though the concepts have been devised jointly. Likewise the development of the filter modules in 2003 was primary my task. However this distinction never was enforced strictly and Roland has provided some enhancements to "my" modules (e.g. for decoding of trigger time) as well as I have improved upon "his" framework (e.g. during restructuring early in 2004 and by adding the watchdog thread in summer 2004).

For silicon time decoding, the invertible parametrisation of time as a function of digit ratio was provided by Dr. Jan Friedrich along with many helpful insights in the functioning of this type of detector. The calibration of the silicon detectors with regard to the aforementioned parametrisation was performed by Jan and Anna-Maria Dinkelbach.

The tool `cat_date` was developed by Roland and most of the "Shift Crew Instructions" reprinted in the appendix also have been written by him.

The interfaces to the experiment mostly have been contributed by specialists of the affected areas: The DATE interface as well as integration into the Run Control was provided by Lars; general aspects of the integration of Cinderella into the DAQ were mostly handled by Lars and Roland. The online monitoring of Cinderella was developed by Roland and Jiawei Zhao. The integration into the Run Logbook has been handled by Damien Neyret.

Acknowledgements

Foremost, for her ongoing love and support, I'd like to express my deepest thanks to Cecily. She made all my efforts worthwhile.

Greatest thanks go to Roland for the very fruitful collaboration on Cinderella. I did tremendously benefit from his knowledge. And it was fun to have him around, too.

I'd like to thank Lars for his friendly supervision.

Prof. Stephan Paul I'd like to thank for his benevolent supervision and for bringing me in touch with the interesting COMPASS experiment.

I'd like to thank Jan for many and fruitful discussions and for just being himself.

Further thanks go to Quirin for philosophy and physical exercises, Sebastian for tea and company, Annemarie for calibration, Boris for trouble-shooting and Michael for chocolate and bagpipe.