

3. The Soudan 2 Experiment

The Soudan 2 Experiment has been built by a Collaboration of physicists and engineers from Oxford University, the Rutherford Laboratory, the University of Minnesota, Argonne National Laboratory (Chicago) and Tufts University (Boston) – see Appendix A. The original aim of the experiment was to search for proton decay. The Collaboration also works on the composition of cosmic rays, active galactic nuclei searches, the construction of whole-sky maps of muon flux and the measurement of the atmospheric neutrino flavour ratio, the latter being the subject of this thesis.

The detector is situated in the disused iron mine at the Soudan Underground State Park, Minnesota, USA. The Soudan mine hosted the Soudan 1 experiment looking for proton decay in the 1980s and is a tourist attraction in the summer (averaging 30,000 visitors every year).

The experimental hall is located in a specially excavated cavern 714m of rock (2090 m.w.e.) underground in order to shield the detector from the products of cosmic ray interactions in the upper atmosphere. Only very high energy muons and the lightly interacting neutrinos reach the detector. The cavity was excavated in the period 1984-85 and installation of the detector elements started in 1986.


3.1 The Soudan 2 Detector

The Soudan 2 detector is a fine grain calorimeter optimised for interactions whose charged products have velocities below the Čerenkov threshold, and which are harder to separate from background in large water Čerenkov detectors. An example of such processes is the difficult super-symmetric proton decay modes, like $p \rightarrow K^+ \nu$ [34][36]. This design

provides high resolution, three-dimensional reconstruction of the event. In addition, the calorimetric nature of the detector provides dE/dx (ionisation) information used in the event analysis.

The Soudan 2 *main* calorimeter detector in its completeness consists of 224 identical *modules*, stacked in a rectangular array (of 8×2 high $\times 14$ modules) to form a continuous volume of 963 tons, measuring $8\text{m} \times 5.4\text{m}$ high $\times 16\text{m}$ (fig. 3.1). It is surrounded by a $\sim 2000\text{m}^2$ *active shield* of aluminium proportional tubes, often referred to as the *veto shield*, which is used in tagging charged particles entering the main detector. The veto shield solid angle coverage is almost 4π . The modular nature of the main detector allowed operation to start well before its completion date and has facilitated maintenance procedures. The first 64 modules of the detector began taking data in September 1988 and the detector has often been taken apart for refurbishment operations without major disruptions to its operation.

The main detector is a collection of 14 *walls* stacked one next to the other, each consisting of 8 modules across and 2 modules high, stacked with the drift tubes along the




• **Figure 3.1:** Layout of the Soudan 2 detector. The modular calorimeter is surrounded by veto shield counters on all sides – the floor panel is not drawn. A half-wall is highlighted. Adapted from [57].

short dimension; each wall is physically divided into two *half-walls* of 4 by 2 modules. Each half-wall is divided along the drift direction (z-axis) into two *physical looms*, which are read out by the wireplanes of the modules – see next section. The module position is numbered, starting at module 0 at the bottom south-east corner of the detector which is also the origin of the detector global co-ordinate system; the x-axis is defined from east to west, the y-axis upwards and the z-axis (the drift axis) from south to north.

Two events in the Soudan 2 detector are given as examples of the detector operation in §3.4. The reader may wish to read on the details of the detector function beforehand.

3.1.1 The Main Detector Module


The modular assembly of the Soudan 2 detector was a necessity because of the limited access to the bottom of the mine, the size of the elevator cage determining the module



• Figure 3.2: Details of module design. From [5].

dimensions: each module measures $1\text{m} \times 1.1\text{m} \times 2.7\text{m}$ and weighs 4.3 tons (fig. 3.2). The design, operation and performance of the modules are presented in great detail in [5] and [6]. Each module is made of alternate layers of corrugated steel sheets (measuring 1.6mm thick \times $1\text{m} \times 94.4\text{cm}$ and making up 85% of the detector's mass) and *bandolier*. The bandolier is a sandwich of two Mylar sheets containing drift tubes made of Hytrel, a resistive ($10^{12} \Omega\text{cm}$) plastic. The tubes are 1m long and have inner and outer diameters of 15mm and 16mm respectively. Bonded to the Mylar are twenty-one copper electrodes to provide a *linear* potential gradient from -9kV at the middle of the drift tube to 0kV at each tube mouth (180V/cm). As the steel sheets are stacked they form hexagonal "pockets" which house the drift tubes of the bandolier; 241 corrugated steel sheets allow for alternate layers of 31 and 32 tubes, the whole construction amounting to 7560 drift tubes per module. The steel-bandolier stack is compressed and the readout *anode wires* and *cathode boards* are installed at the front and back (in what will become the north and south faces) of the module before closing the structure in an airtight enclosure. The 63 anode wires run vertically down, passing 10 mm in front of the tube mouths and will be set at a potential of $\approx +2.2\text{kV}$, while, 5 mm behind them, the 239 horizontal cathode strips are mounted on boards. The anode wires and cathode strips define a *wireplane*.

During operation, the drift tubes are filled with 85% argon, 15% CO_2 gas mixture which is re-circulated through a purifier system. Traces of water are also fed into the system to preserve the Hytrel in good condition. A charged particle passing through a drift tube will deposit ionisation in the gas which, because of the high potential gradient, will drift to the mouth of the tube where an anode wire and cathode strip will read it out (fig. 3.3). The field from the anode wires ($\approx +2.2\text{kV}$) causes drift electrons to be amplified by an avalanche process. The positive ions left in the wake of the avalanche then drift towards the cathode pad; this induced current provides the signal in the anode wire and cathode strips. The drift speed in the tubes is $\approx 0.6\text{cm}/\mu\text{s}$ and ionisation takes a maximum of $\approx 83\mu\text{s}$ to reach the wireplane, which corresponds to the time needed to cover the 50cm from the middle of the



• **Figure 3.3:** Drift Tube cross-section. From [5].

tube to its mouth. In this fashion, the 2-dimensional position *and* the ionisation profile of the pulse are defined. The third (z) co-ordinate (along the drift axis) will be determined by measuring the pulse-time relative to the event-time. To determine the latter, we need the *event* T_0 , i.e. when the event happened relative to the trigger time. This information has to be extracted from the complete set of pulses in the event (see §3.3.3).

A charged particle crossing the wireplane will create large amounts of ionisation because of the lack of material and abundance of drift gas outside the tube mouths, producing a *blob*, a group of simultaneous, large pulses on neighbouring channels (both anodes and cathodes, depending on the direction of travel). Another feature of the design of the modules is *cathode cross-talk*, i.e. ghost cathode pulses that appear on either side of a real cathode pulse. This is due to a small fraction of the positive ion current drifting towards the cathode strips situated on either side of the tube mouth.

3.1.2 The Main Detector Electronics and the Trigger

The read-out of the detector is done at the level of *physical looms* where the anodes from the top module are bussed to the bottom module forming a 5m long wire and the cathodes are bussed across the 4 modules along the x -direction. A physical loom has $4 \times 64 = 256$ anode and $2 \times 128 = 256$ cathode read-out channels; each module however has 63 anode wires by

239 cathode strips and therefore anodes 64, 128, 192, 256 and cathodes 240-256, 496-512 carry no signal.

The signal from the anode wires and cathode strips are pre-amplified and then digitised by an array of 6-bit flash ADCs (Analogue to Digital Converters). To reduce the number of channels that need to be read out from a mostly quiescent detector (and at the same time save on resources) the electronics output has been multiplexed before digitisation by summing the signal from eight pre-amplifiers into one channel. The signal from the 56 physical looms is therefore compressed into 8 *logical looms* (numbered from 0 to 7), both in the anodes and the cathodes (the 8-fold multiplexing allows for a maximum of 64 physical looms but only 56 are present). The multiplexing scheme has been designed such that a half-tube is unambiguously located if a pulse on an anode channel is uniquely matched to a pulse on a cathode channel [18]. Moreover, care has been taken so that, for a typical neutrino interaction that is fully contained in the main detector, de-multiplexing produces a minimum of ghost images of reconstructed matched hits. After multiplexing, the pulses are sampled and digitised every 200 ns (the *detector time unit*, or dtu) and stored in a FIFO (First In First Out) register 1024 words deep. When the detector is triggered, data is read for a further 512 dtu and is transferred to the data-acquisition (DAQ) computer. Consequently, the detector information for each event is recorded for the *trigger window* of 1024 dtu (or 204.8 μ s), the trigger time being at 512 dtu by definition. Each hit in the event has a *hit time* associated with it, defined by this system of time registration.

The trigger of the Soudan 2 detector was designed and optimised for nucleon decay and neutrino induced events in the presence of noise (natural radioactivity from the rock surrounding the detector) and is fully described in [18]. The trigger accepts events satisfying a minimum number of groups of positive going pulses, or *edges*, in a given time window in multiplexed proximity. The completed Soudan 2 detector currently triggers at a rate of ≈ 0.5 Hz, approximately 2/3 of which come from energetic cosmic-ray muons that have travelled through the rock overburden above the detector. The remaining triggers are

due to naturally occurring radioactivity, noise or breakdown problems, DAQ malfunctions and other physics phenomena that we study, such as atmospheric neutrino interactions. Events are collected during detector *runs* of 1200-2000 events each; data taking is completed when a certain amount of data blocks³, has been written to disk.

3.1.3 The Active *Veto* Shield


The main detector is surrounded by a 2040 m² active veto shield which tags charged particles. It is particularly useful in the study of events which have passed all containment cuts but have shield activity associated with them. These give us a handle in the study of the background contamination of the contained event sample from neutral particles entering the main detector.

The Tufts University High Energy Group has been in charge of the veto shield installation and running [69]. The majority of the shield construction had been completed by 1988. The shield is mounted flush on the Soudan 2 cavity walls, typically at a distance of 4 m from the top, east, west and south calorimeter faces, 1 m from the bottom and 12 m from the north face. The choice of the placement has been such as to detect cosmic ray muons (or the charged products of their interactions in the surrounding rock) travelling through the detector cavity, even if they do not cross the main detector. The area coverage of the device is currently above 99%. This figure used to be lower in the early days of the experiment. The major shield upgrade was the installation of the top of the north wall⁴ in Aug. 1991 and the completion of the bottom part in Dec. 1992. The small “holes” at the floor wall are due to steel beams supporting the main calorimeter modules. Finally, very small areas of the ceiling are still uncovered because they are not easily accessible.

The original shield elements consist of two layers of four hexagonal proportional counter cells each, filled with a 95% Ar 5% CO₂ gas mixture which is re-circulated and

³ A block is a unit of data storage, corresponding to 512 Mbytes.

⁴ The reader should note that the north wall is the one furthest away from the main detector.



• **Figure 3.4:** Element (Tufts module) of veto shield surrounding the Soudan 2 detector.

purified through a system similar to that used for the main detector (fig. 3.4). The maximum time for ionisation to drift from one of the corners of the tube to the central high-voltage anode cable is $0.8 \mu\text{s}$ [69]. For this reason, the time resolution of the shield is poorer than that of the main detector and is set at $1 \mu\text{s}$. These elements are custom built for the experiment by Tufts, come in various lengths and are installed with one (the *in* tube) four-cell tube facing the calorimeter and the other (the *out* tube) facing the cavity wall. Adjacent tubes of equal length interlock to form *panels*, the collection of which forms six *superpanels*, one for each face of the shield.

In order to further improve the shield efficiency additional layers of tubes from the defunct Harvard-Purdue-Wisconsin (HPW) experiment have been deployed on the ceiling superpanel, in a criss-cross fashion, such as to obtain a 3-dimensional point on the trajectory of each charged particle [37].

In addition, the muon tubes from the TASSO experiment have been assembled in a crossed fashion into panels which have been placed right on top of the main detector, covering the inter-modular gaps and making the identification of incoming charged particles an even easier job. This last upgrade of the shield comes from Oxford and was completed in 1996 [79].

Random Rates	Single Hits	747.0 kHz
	Adjacency Hits	10.0 kHz
	Overlap Hits	11.4 kHz
	TASSO Hits	0.5 kHz
Efficiency	Ceiling	99.0%
	Floor	93.6%
	East Wall	96.2%
	West Wall	92.7%
	North Wall	98.1%
	South Wall	97.6%
	TASSO array	82.8%

• **Table 3.1:** Veto shield random rates and efficiencies.

The smallest unit of activity in the shield, the *single* hit, occurs when only one layer of the tubes is hit. Because of natural radioactivity from the rock and concrete in the cavern, the singles random rate is high (Table 3.1). For this reason, two types of coincidences are used in the data analysis:

- An *adjacency*, defined by two adjacent in-out layers of Tufts tubes being hit at the same time, has a low random rate and is used in the analysis.
- Equally important, albeit having a slightly higher random rate, is the *overlap*, defined by the coincidence of singles or adjacencies of two crossed Tufts or HPW tubes, giving a 3D resolution of the order of 20 cm. The crossed TASSO arrays sitting right on top of the main detector also produce overlaps with spatial resolution of 8 cm. The efficiency of the TASSO panel is lower than that of the other panels but its random rate is low and the quality of the data is not compromised.

The shield random rates and efficiencies are summarised in Table 3.1. These numbers are for the latest period of data taking. They have been slowly improving over the years.

3.1.4 Detector Survey: The Soudan 2 Database

It is crucial for most aspects of the data analysis and for the simulation of the experiment to keep a record of various detector parameters as a function of time. This information is stored in the Soudan 2 database, the main copy of which sits at the Soudan mine computers while mirror copies at each collaborating institution are updated daily.

The modules of the main detector may vary in height (there is a spread of 5cm approx.) because of very slight variations in the thickness of the corrugated steel sheets or that of the bandolier. Before the anode wires and cathode boards are installed, the stack of steel and bandolier is surveyed every 20 planes using a measuring tape. The size, internal structure and mass of each module is recorded in the database.

The position of the calorimeter modules is also stored in the database and is surveyed in two ways: (i) by a measuring tape where the frame of reference is an array of wires stretched across the detector cavern along all three dimensions and (ii) by following the path of energetic cosmic ray muons travelling through the main detector. The second approach is also used to calibrate the drift speed of each module as well as the half-module parameter; the latter is a time pedestal that compensates for the slightly different drifting behaviour of each module-half. This effect is due to small variations in the position of the readout anode wires and cathode boards relative to the tube-mouths.

The efficiency for each of the two module wireplanes is determined by processing the pulses from energetic cosmic ray muons travelling through the detector and is recorded as a function of time in the *half-tube database*. This survey has proved very useful in order to pinpoint inefficient modules and have them rectified during detector upgrades.

Finally, the survey of the active veto shield components is made using similar techniques to the above; the results are also stored in the database.

3.2 The Analysis Chain

The analysis chain is divided into two sections:

- The software event reconstruction (SOAP) leading the first stage of event selection performed by processor FILTER – see next section.
- The second stage of event selection and the Physics analysis of the data. Two approaches have been adopted here: (i) the scanned analysis and (ii) the automated analysis, which is the subject of the present thesis.

The software event reconstruction as well as the Physics analyses that follow will be highlighted below.

3.2.1 SOAP

The Soudan 2 collaboration has developed SOAP (Soudan 2 Offline Analysis Programs), a set of computer processors which assist the physicist with the data reconstruction and analysis. The programming language at Soudan 2 is FORTRAN 77 and the BOS memory manager is used throughout. The algorithms deal with the veto shield and the main detector information that is collected in each trigger.

The raw data from the veto shield is stored in the so-called SHLD banks. The processor SPLASH produces the SHLA and SHLO banks for adjacency and overlap shield hits respectively.

The raw data from the Soudan 2 main calorimeter detector (as well as from the Soudan 2 Monte Carlo simulation package, see chapter 4) are stored in the so-called 4BCB banks, containing the compressed raw (FIFO) readout for each detector crate of electronics. The first processor to run on the data is CREPLR; its function is to unpack the compressed data, thus producing the *raw* pulses. Processor SOFTPL follows, producing the *soft* multiplexed anode and cathode pulses, where detector noise (due to electronics problems or breakdowns) has been mostly removed and cathode cross-talk pulses have had their size

drastically reduced. The PMT (Pulse MaTching) processor [63] de-multiplexes the soft pulses by means of their timing and their pulse-height profiles, producing reconstructed 3D hits with definite x and y co-ordinates and a z co-ordinate (defined along the drift axis) dependant upon the event T_0 , i.e. when the event happened relative to the trigger time (see §3.3.3). The event T_0 is estimated by processor TODET in collaboration with processor SEARCH [64], which fits tracks in multiplexed anode-time and cathode-time projections and then attempts to combine them to form tracks across loom boundaries.

Finally, the processor FILTER makes a rough selection of contained events, rejecting those failing the fiducial volume cut, using matched SEARCH tracks and 3D hits from PMT. FILTER reduces the size of the event sample by a factor of $\sim 10^3$, leaving some 10^4 events to be analysed every kTon-year.

3.2.2 The Scan Analysis

From there on two completely orthogonal approaches have been adopted within the Soudan 2 collaboration. Historically, the events were scanned by hand, using STING, a custom built graphics interface to SOAP [7][47]. This procedure was widely accepted in the early days of the experiment when the physicist would also evaluate the detector performance and operation. Scanning involves manual selection of the events (rejecting noise events and cosmic ray muons which fooled FILTER), followed by manual reconstruction (interaction vertex estimation, pulse-matching and track/shower fitting) and flavour classification into tracks (quasi-elastic ν_μ interactions), showers (quasi-elastic ν_e interactions) or multi-prongs (events with associated hadronic production). The data results are compared to the Standard model expectation by means of a Monte Carlo (MC)-simulation. In order to minimise systematic errors, the MC data has to be scanned exactly the same way as the real data,

making this method very labour intensive. The exposure of the MC sample analysed is only 5.9 times that of the real data, making this method prone to statistical errors.⁵

The scanning method has been the subject of long discussions and internal checks within the Soudan 2 collaboration, the understanding of systematic effects due to the human factor being our main concern. The greatest part of the data has been scanned with the MC data mixed in; the scanner is not aware if the event is real or MC-simulated. At one point in time two scan committees were analysing the same data sets and arrived to consistent results. In Soudan 2 we rest assured that scanning is a viable and reliable method.

3.2.3 The Automated Analysis

The aim of this thesis is to describe the alternative method, developed in Oxford over the last three years, and based on the principle of minimal human intervention in the analysis of the data. The selection of the events is accomplished through an almost automatic process called RINSE (Rapid Isolation of Neutrino Soudan 2 Events), described in chapter 5. The flavour of the events is then extracted statistically and not on an event-by-event basis. The flavour estimators and the methodology of the Physics analysis are described in chapter 6. Finally, and most interestingly, the results from the analysis are presented in chapter 7.

⁵ The MC-sample used in the scanned analysis is significantly smaller than the one analysed in the present thesis, although the Physics is the same in both – more on the MC in following chapters.

3.3 Elements of Software Event Reconstruction

This section introduces elements of the event reconstruction that will be used in the length of this thesis.

3.3.1 The Event Box

The *event box* determines both which reconstructed 3D hits comprise the event and what is the location of the event in the detector. It is the rectangular parallelepiped, aligned to the detector axes, that contains the *largest* cluster of 3D hits, the size of the cluster being determined by the number of hits. The (mutually exclusive) boxes are bounded on all six faces by *either* the detector edge *or* a 50 cm apron which is empty of 3D hits. The prime factor that determined the choice of this width of the apron is the existence of inter-modular cracks: events crossing module boundaries will typically re-appear further than 20 cm away. The distance of 50 cm has been chosen such as to optimise the inclusion of hits belonging to the event relative to noise hits and roughly corresponds to three radiation lengths in the Soudan 2 detector. Under certain conditions, the largest event box will be joined with the second largest box – in particular when these are found on either side of the same inter-modular crack.

The 3D hits in the event box are classified as physical or unphysical, according to their registered time and T_0 (see below). 3D hits in the box are also marked as *isolated* if their nearest neighbour is more than 30 cm away. Finally, 3D hits are also marked as *small* if their integrated pulse-height is less than 25 ADC units – the average integrated pulse-height of a hit in an event is around 200 ADC units. The pulse-height of all 3D hits in the present analysis has been corrected for drift attenuation by using the information in the half-tube database.

The collection of *physical, non-isolated and non-small 3D hits* in the event box defines the event in the detector. In the analysis described in this thesis all reference to “3D hits” will neglect all other 3D hits.

3.3.2 Clustered 2D Pulses

The Clustered 2D Pulses algorithm attempts to define the set of multiplexed soft 2D pulses that are part of the event. The adopted procedure only considers the channel number and time of each 2D pulse and is independent of pulse-matching performed by PMT.

For each of the anode and cathode multiplexed looms, pulses are clustered in channel and time space. In the anode domain, clusters are defined by pulses not more than 13 channels (20 cm) apart. Similarly, in the cathode domain the minimum separation between clusters is 20 channels (20 cm). In the time domain, pulses of different clusters are separated by at least 50 dtu ($10 \mu s$) in a head-to-tail fashion, i.e. one is interested in the time difference between the end of the first pulse and the start of the second. 2D clusters must have at least 5 pulses. All pulses of integrated pulse-height below 25 ADC units are rejected.

A small fraction of events is characterised by breakdowns that are usually restricted to a single module, and rarely light up a whole loom, and whose pulses occur simultaneously. This *snake*-pattern is easily recognisable, even by software,⁶ and these pulses are ignored in the event reconstruction, i.e. they are not pulse-patched or used in track fitting. A 2D cluster will not be considered any further if at least 60% of its 2D pulses belong to a snake.

The clusters that have been selected in the procedure highlighted above hold the set of *clustered 2D pulses*. These will be used in the selection of events, which is presented in chapter 5. One will be interested in how many of the clustered 2D pulses are matched by PMT. Large discrepancies would hint at reconstruction failures.

3.3.3 T_0 and the T_0 Window

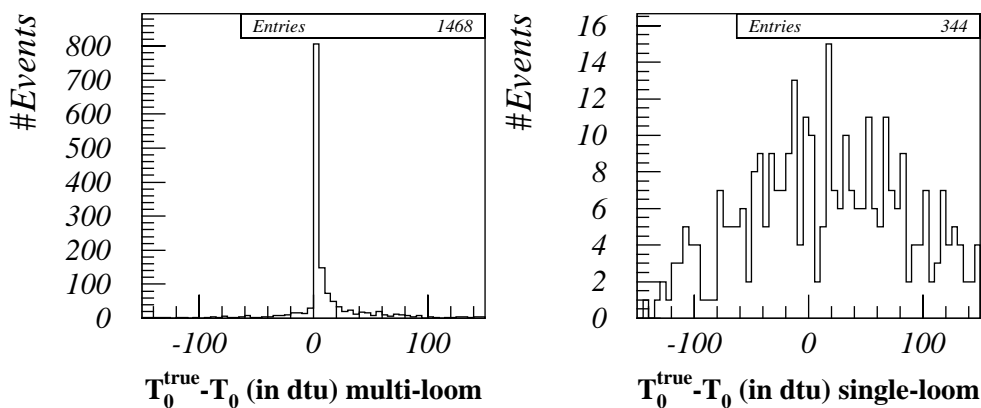
The main calorimeter is a drift chamber and the event will in general have occurred at a variable time with respect to the trigger: an event has first to drift to the wireplane for the

⁶ Processor REJSNK: a snake is required to have at least 40 anode or 40 cathode 2D pulses within 2 dtu.

electronics to register and the detector to trigger. T_0 is defined as the time the event happened in the trigger window.

By choosing a wrong T_0 value, it is possible to move the z co-ordinate (defined along the drift axis) of hits that belong to the event out of the physical drift region. Conversely, random hits that do not belong to the event may appear as unphysical even when T_0 is correct. Processor T_0DET attempts to determine T_0 by minimising the dispersion of the event along the drift axis while maximising the number of 3D hits in the event box in the physical drift region. Moreover, the adjustment of T_0 can bring parts of events that cross physical looms into alignment. This is the principle of $SEART_0$, a component of processor $SEARCH$, which collaborates with T_0DET in order to obtain the final answer. The event T_0 can then be determined with good accuracy for events that cross looms, as illustrated by a MC study (fig. 3.5). One can only place upper and lower limits on T_0 for events that do not cross loom boundaries (single-loom events). They are placed in the middle of the loom *unless* a blob is detected by processor $BLOBER$ and the event is moved on the wireplane (fig. 3.5).

We often need to know the acceptable range of T_0 values for the event, referred to



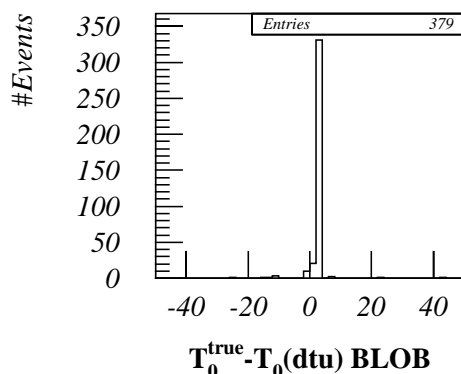
• **Figure 3.5:** Monte Carlo study of the performance of the T_0 algorithm. Distribution of difference of estimated value from true value, where the latter is obtained from the MC truth, for multi-loom and single-loom events.

as the T_0 -window, ranging from T_0^{\min} to T_0^{\max} . This is particularly true in the separation of contained events (neutrino candidates) from cosmic ray induced events (background candidates) using the information provided by the veto shield. In general, a shield hit occurring within the T_0 -window will flag the event as background-induced (see §5.3.7).

T_0^{\min} is defined by how much the event can be allowed to slide towards the module centre (i.e. the start of the loom) without making any hit in the box unphysical over that boundary. Similarly, T_0^{\max} is defined by how much the event can be allowed to slide towards the wireplane (i.e. the end of the loom) without making any hit in the box unphysical over the wireplane boundary. In addition, an allowance of 30 dtu is made on either end of the T_0 window.


The above algorithm is not applied when T_0 has been defined by a blob by both TODET and SEARCH. In this case, the T_0 window is defined by $T_0 \pm 50$ dtu, because T_0 is known with good accuracy, thanks to the wireplane loom crossing. This procedure is justified by the MC analysis (fig. 3.6).

According to the above algorithm, an event with T_0 defined by a loom crossing at the module centre (where a narrow T_0 window would be correct) will have T_0^{\max} defined by how far the earliest hit is from the wireplane, completely disregarding the T_0 information.



• **Figure 3.6:** MC study of the difference between the true and estimated T_0 for events where a blob on the wireplane was used.

Similarly, for an event with a wireplane loom crossing (and where the blob is not obvious) will have T_0^{\min} defined by the distance of the latest hit to the module centre. This conservative approach has been adopted because there is a 10% chance for accepting an *erroneous* loom crossing due to random hits in the detector, which would result to an erroneously narrow T_0 -window. In this fashion, the



• **Figure 3.7:** Monte Carlo study of the performance of the T_0 -window software algorithm. The distributions of the difference of true T_0 , obtained from MC truth, to the calculated values of T_0^{\min} and T_0^{\max} are shown. Less than 1% of the data fall outside the T_0 -window.


error involved in the separation of signal to background is minimised. The true (MC) T_0 is in the reconstructed T_0 -window 99% of the time (fig. 3.7).

3.3.4 The Event Axis

The *event axis* is defined as the axis about which the moment of inertia of the event is *minimum* and is obtained by diagonalising the Moment of Inertia (MoI) tensor of the event,

$$\text{MoI Tensor} = \left(\sum_{r=1}^N \mathbf{r}_i \cdot \mathbf{r}_i \right) \times \mathbf{I}_3 - \begin{pmatrix} \sum_{i=1}^N x_i x_i & \sum_{i=1}^N y_i x_i & \sum_{i=1}^N z_i x_i \\ \sum_{i=1}^N x_i y_i & \sum_{i=1}^N y_i y_i & \sum_{i=1}^N z_i y_i \\ \sum_{i=1}^N x_i z_i & \sum_{i=1}^N y_i z_i & \sum_{i=1}^N z_i z_i \end{pmatrix} \quad (3.1)$$

where \mathbf{I}_3 is the 3×3 unity matrix, $\mathbf{r}_i = x_i \mathbf{i} + y_i \mathbf{j} + z_i \mathbf{k}$ is the position vector of the i^{th} hit about the Centre of Gravity (CoG) of the event and the event box contains N 3D hits. The MoI eigenvectors define a rotated orthogonal co-ordinate system, $x'-y'-z'$, where the origin coincides with the centre of gravity of the event and z' is defined along the *event axis* (fig. 3.8). The minimum moment of inertia eigenvalue of the event about this axis is given by



• **Figure 3.8:** Schematic of the orthogonal rotated co-ordinate system defined by the MoI tensor eigenvectors. The moment of inertia of the event about the z' axis (the event axis) is minimum.

$$\text{MoI}_A = \sum_{i=1}^N (x_i'^2 + y_i'^2), \quad (3.2)$$

while the other two larger moment of inertia eigenvalues are given by

$$\text{MoI}_B = \sum_{i=1}^N (x_i'^2 + z_i'^2), \quad (3.3)$$

$$\text{MoI}_C = \sum_{i=1}^N (y_i'^2 + z_i'^2). \quad (3.4)$$

The event axis will be used as an estimator for the direction of travel of the event in the detector, both in the selection of the events and the flavour analysis of the data, whereas the moments of inertia eigenvalues will be used in constructing a flavour estimator, Γ , in §6.1.2.


3.4 Soudan 2 Events

Two Soudan 2 events are presented in this section in order to give the reader a feeling of the detector operation. The pictures have been extracted from STING, the graphical interphase to SOAP, where one views the three detector projections along the x - z (view from the top), y - z (view from the east side) and y - x (view from the north side) planes. The colour coding of

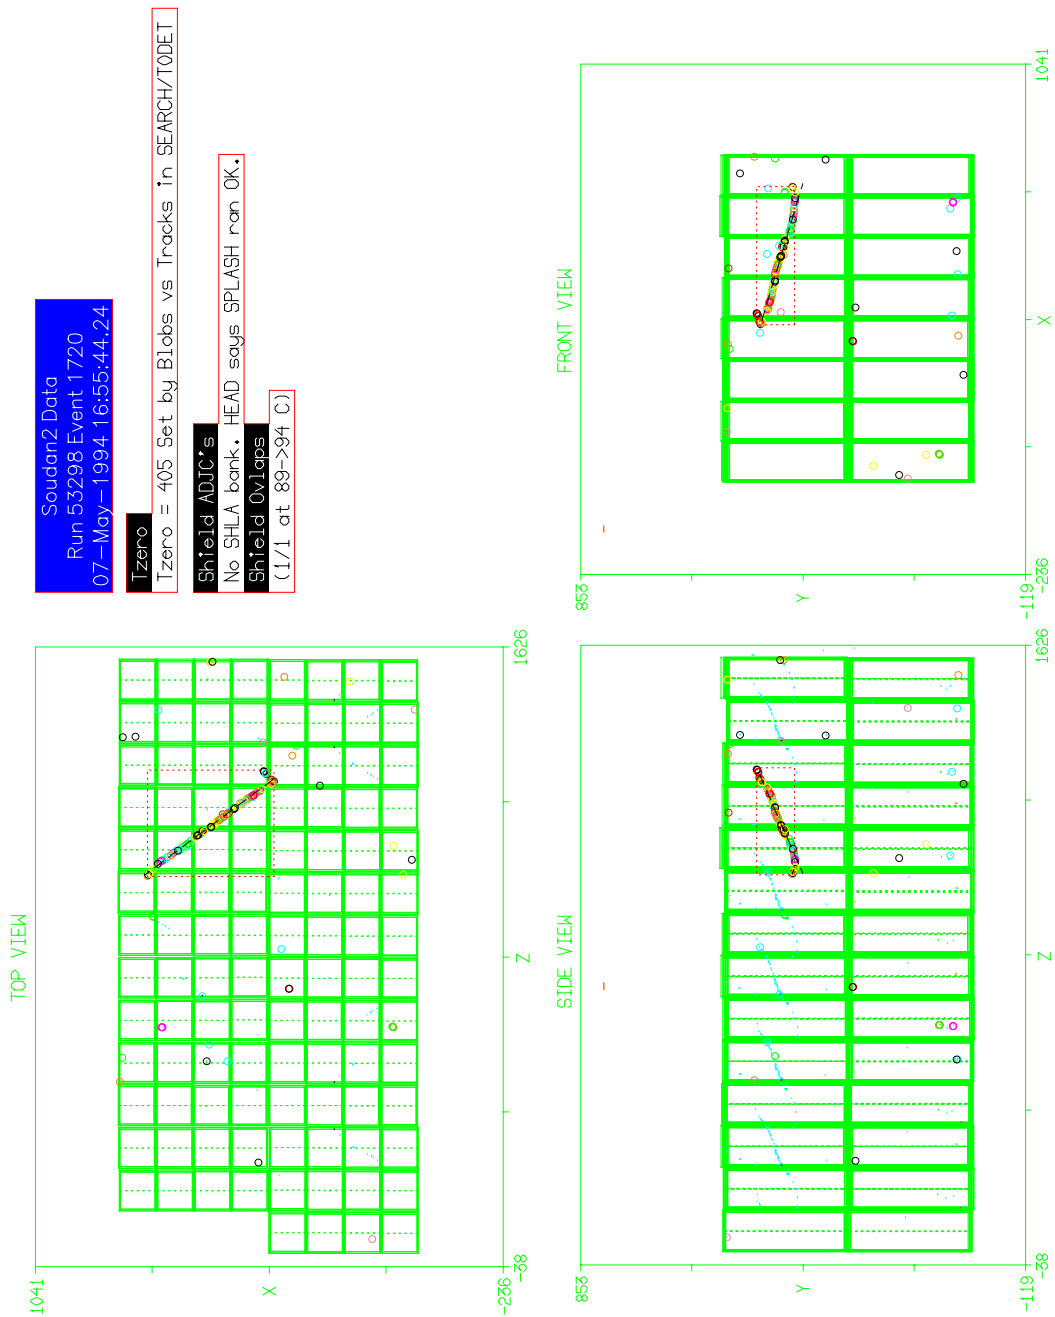
the pulses is related to their integrated pulse-height. The T_0 and the profile of shield elements that have been hit are shown.

The first event is a through-going cosmic ray muon (fig. 3.9). The events dates from March 1996 when the TASSO shield upgrade on top of the main detector was in operation. The muon nicely triggers all the shield layers it crosses: its passage produces 2 adjacency hits, one in the ceiling and the other in the floor, and 11 overlap hits. All the shield hits are coincident, at a time of 474-479 dtu (the shield time resolution is $1\mu s$, or 5 dtu). A random shield overlap at a time of 254 is also present in the floor. The presence of three TASSO overlap hits on top of the detector is noticeable; one also observes ionisation deposited beside the muon's path at the top of the detector. This extra activity can be attributed to either (i) the very high energy cosmic ray (CR) muon emitting bremsstrahlung radiation as it reaches the detector or (ii) to the products of the CR-muon scattering in the rock right above the detector. Inside the main detector, the ionisation deposited by the muon has been successfully reconstructed by SOAP. The event T_0 has been correctly defined at 480 dtu by the presence of blobs: the event enters the detector through the wireplane of a module. In this particular instance the T_0 algorithm has confirmed its agreement with the veto shield timing. This is a typical CR-muon event. Some 27 million of them have been analysed in a hotspot sky survey [49].


The second event is a high energy atmospheric muon neutrino interaction in the main detector (figs 3.10 and 3.11). There is no associated shield activity (the overlap hit at the ceiling is widely out of time) and the event is nicely contained within the fiducial volume. The SOAP reconstruction was successful and wireplane crossings were used in the determination of T_0 . The dashed box corresponds to the outline of the event box in each view. The event consists of a long muon track and a proton recoil. The latter can be identified by the energetic pulses and the short range and shows that the vertex is at the upper end of the event. The muon track multiple-scatters at the end of its path.



• **Figure 3.9:** STING view of a through-going cosmic ray muon.



• **Figure 3.10:** STING picture of an atmospheric muon neutrino interaction.



• **Figure 3.11:** STING zoomed-in picture of the muon neutrino interaction of fig. 3.10.