



The trigger for $K^0 \rightarrow \pi^0 \pi^0$ decays of the NA48 experiment at CERN

G. Barr^a, D. Cundy^a, F. Formenti^a, B. Gorini^a, B. Hallgren^a, W. Iwanski^{a,1},
P. Kapusta^{a,1}, G. Laverrière^a, M. Lenti^{a,2}, I. Mikulec^{a,3}, M. Velasco^{a,4},
O. Vossnack^a, H. Wahl^a, M. Ziolkowski^{a,1,5}, M. Porcu^b, F. Rossi^b, C. Avanzini^c,
P. Calafiura^{c,6}, M. Cirilli^{c,*},⁷, F. Costantini^c, F. Laico^c, G. Magazzù^c, F. Morsani^c,
G.M. Pierazzini^c, D. Rizzi^c, M. Sozzi^c, R. Tripiccione^c, H. Dibon^d, G. Fischer^d,
M. Jeitler^d, M. Markytan^d, G. Neuhofer^d, M. Pernicka^d, A. Taurok^d, L. Widhalm^d

^a CERN, EP Div. CH-1211 Geneva 23, Switzerland

^b Dipartimento di Fisica dell'Università e Sezione dell'INFN di Ferrara, I-44100 Ferrara, Italy

^c Dipartimento di Fisica, Scuola Normale Superiore e Sezione INFN di Pisa, I-56100 Pisa, Italy

^d Österreichische Akademie der Wissenschaften, Institut für Hochenergiephysik, A-1050 Wien, Austria

Received 9 August 2001; received in revised form 5 November 2001; accepted 19 November 2001

Abstract

The trigger used for the collection of the samples of $K^0 \rightarrow \pi^0 \pi^0$ decays in the NA48 experiment at CERN uses a novel pipeline design in order to satisfy the demanding specifications of a high rate kaon beam. The trigger algorithms, architecture and performance are described. © 2002 Elsevier Science B.V. All rights reserved.

Keywords: Trigger; Pipeline; High energy physics; NA48

*Corresponding author. CERN, EP Div., CH-1211, Geneva 23, Switzerland. Tel.: +41-22-767-1246.

E-mail address: manuela.cirilli@cern.ch (M. Cirilli).

¹ On leave from Institute of Nuclear Physics, Cracow, Poland.

² On leave from Sezione INFN di Firenze, I-50125 Firenze, Italy.

³ On leave from Österreichische Akademie der Wissenschaften, Institut für Hochenergiephysik, A-1050 Wien, Austria.

⁴ Present address: Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208, USA.

⁵ Present address: Fachbereich Physik, Universität Siegen, D-57068 Siegen, Germany.

⁶ Present address: E.O. Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.

⁷ Present address: Dipartimento di Fisica dell'Università di Roma "La Sapienza" e Sezione INFN di Roma, I-00185 Roma, Italy.

1. Introduction

The NA48 experiment at CERN is designed to study direct CP violation by evaluating the real part of the parameter ε'/ε . This is done by counting the number of K_L^0 and K_S^0 into $\pi^+ \pi^-$ and $\pi^0 \pi^0$ (where $\pi^0 \rightarrow \gamma \gamma$) decays to measure the double ratio

$$R = \frac{\Gamma(K_L^0 \rightarrow \pi^0 \pi^0) \Gamma(K_S^0 \rightarrow \pi^+ \pi^-)}{\Gamma(K_S^0 \rightarrow \pi^0 \pi^0) \Gamma(K_L^0 \rightarrow \pi^+ \pi^-)} \simeq 1 - \Re\left(\frac{\varepsilon'}{\varepsilon}\right). \quad (1)$$

The measurement is actually performed by triggering on the four decay modes and counting the ones which pass the containment and background rejection cuts which are made offline. Systematic

uncertainties due to changes in detector efficiencies are reduced by collecting the four decay channels simultaneously, while acceptance corrections are minimised by using almost collinear K_L^0 and K_S^0 beams and by weighting the K_L^0 events to make the decay length distribution the same as that of the K_S^0 events. The K_S^0 and K_L^0 decays are distinguished by tagging the protons which produce the K_S^0 . The limiting mode for the statistics is the decay $K_L^0 \rightarrow \pi^0 \pi^0$, which has a branching ratio of 9×10^{-4} . An overall sensitivity of 2.6×10^{-4} on $\Re(\epsilon'/\epsilon)$ has been reached in three years of data taking (1997–1999), with the high intensity K_L^0 beam delivering 500 k decays per second [1].

This paper is concerned with the sub-system known as “the neutral trigger”, i.e. the sub-system responsible for triggering on K_L^0 and K_S^0 decays into $\pi^0 \pi^0$ final states. The main background to these signal events comes from $K_L^0 \rightarrow 3\pi^0$ decays with overlapping or undetected photons. The $3\pi^0$ decays represent the dominant K_L^0 decay channel, with a $\sim 21\%$ branching ratio. In spite of this very unfavourable signal to noise ratio, the trigger design was mainly driven by the need of a remarkably high efficiency rather than an excellent purity. The reason for this choice lies in the basic principle of the double ratio technique: in order to perform a precision measurement of $\Re(\epsilon'/\epsilon)$ with this technique, the unavoidable systematic biases in the counting process must be symmetric between at least two of the four modes in Eq. (1). The requirements this places on the neutral trigger system are:

- a very high efficiency (above $\sim 99\%$), which can be evaluated to better than one part per thousand; this is required to avoid large correction to $\Re(\epsilon'/\epsilon)$ due to different kinds of inefficiencies for $K_L^0 \rightarrow \pi^0 \pi^0$ and $K_S^0 \rightarrow \pi^0 \pi^0$;
- a low dead time, because any dead time condition on a decay mode must be applied to all the others decay modes to maintain the principle of simultaneous collection;
- a well understood response to decays which occur close in time to other activity (accidentals);

- a good rejection of the high rate of uninteresting decays;
- a fast delivery of trigger decisions.

The neutral trigger is thus designed as a fully synchronous pipelined system. It continuously performs computations on the digitised and summed information from the 13 248 channels of the electromagnetic calorimeter. It reconstructs online the energy of the kaon (with a 3% resolution), the decay position along the beam (to 3 m over a distance of 100 m), the number of photons and their arrival times (to 3 ns). The proper decay time of the kaon is derived from the energy and decay position. The trigger implements these calculations every 25 ns in a dead-time-free pipeline, and performs a selection based on the reconstructed physical quantities. A weak transverse momentum cut is also applied. All these selection cuts are programmable, providing the possibility of setting trigger requirements for events other than the $K^0 \rightarrow \pi^0 \pi^0$ decays, like neutral kaon rare decays or $\eta \rightarrow 2\gamma$ decay. The neutral trigger also provides the total energy of the event by summing the energy from the hadron and the electromagnetic calorimeters. This is used to impose a total energy threshold on $K^0 \rightarrow \pi^+ \pi^-$ triggered events.

This paper is organised in the following manner. Section 2 gives a brief overview of the NA48 liquid krypton electromagnetic calorimeter (LKr). This section also includes a description of the so-called “neutral hodoscope” (NHOD), which is used as an independent trigger to measure the neutral trigger efficiency. Section 3 contains an overview of the neutral trigger chain, followed by a detailed description of each component of the computation pipeline. The implementation of the total energy trigger is described in Section 4. Section 5 describes the setting-up and operation of the trigger, while Section 6 reports performance measurements in the NA48 beams.

2. The neutral particle detectors

Kaons in the NA48 beams have an average energy of 110 GeV and decay in flight. The decay

products travel along the z direction for about 100 m before reaching the detectors. Photons from $\pi^0\pi^0$ decays are measured by the quasi-homogeneous LKr calorimeter: this is a liquid krypton ionisation chamber, with a 127 cm ($27 X_0$) deep active region. The Molière radius of the showers is 4.7 cm, which is somewhat larger than in sampling calorimeters. The read-out is arranged in a tower structure of copper ribbons [2] defining the electrostatic cells. The cells are not strictly parallel, but widen slightly by 10 mrad in order to point to the average position along the beam where the kaons in the accepted energy range decay. There are a total of 13 248 cells arranged to define a Cartesian grid of $2 \times 2 \text{ cm}^2$ cells on which photons impact (see Fig. 1). The x and y axis, respectively, define the horizontal and vertical directions in the plane orthogonal to the beam-line. The sensitive area of the calorimeter forms an octagon, and can be considered a 128×128 cell square grid with the corners missing.

The maximum drift distance within a cell is 1 cm. The electrostatic field within the cells is 3 kV/cm, and the corresponding electron drift velocity is 0.3 cm/ μs . The preamplifiers and switching circuits for the calibration system [3] are mounted directly on the ends of the cells inside the liquid krypton. The signals go through further amplification and shaping stages at the feedthroughs of the krypton cryostat and on the Calorimeter Pipeline Digitiser (CPD) readout modules (see Section 3.3).

The shaping circuits differentiate the signals so that only the fast initial rise of the current produces the pulse which is used in the neutral trigger and in the readout. This arrangement results in a fast detector (the signal pulse has a FWHM of ~ 70 ns) with good energy, position and time resolution. A 3% flat undershoot on the signal is present for the $\sim 3 \mu\text{s}$ following each pulse, during which electrons continue to drift across the electrode gap. At 25 GeV, the energy resolution is better than 1%, the spatial resolution is 1 mm and the time resolution on a single shower is 280 ps.

The NHOD consists of bundles of scintillating fibres placed vertically within the liquid krypton calorimeter at a depth of $\sim 9.5 X_0$. This is the depth where an average electromagnetic shower




Fig. 1. Front view schematic of the LKr calorimeter. Some individual cells are shown towards the right of the picture. Calorimeter signals from 2×8 single cells are added into supercells. Some x view supercells are shown along the left-hand side, while y view supercells are shown along the top. These subdivisions extend over the entire calorimeter. Supercell signals are summed up into 64 vertical and 64 horizontal projections (strips), which are shown along the bottom and right-hand side of the picture.

has its maximum development. The scintillators are read out with photomultiplier tubes. The NHOD is used to generate an independent and unbiased trigger on neutral decays; this trigger is generated by the coincidence of two hits in two half-planes of the hodoscope, either horizontally (left and right) or vertically (top and bottom), and runs at about 30 kHz. NHOD triggers are down-scaled by a factor 100 and collected in order to measure the $\pi^0\pi^0$ trigger efficiency.

3. The neutral trigger

3.1. Overview

The NA49 neutral trigger design has to satisfy the requirements already stated in Section 1, which are dictated by the experimental technique

adopted to measure ε'/ε . In addition, the trigger has to cope with the high rate of particles in the detector, which is of the order of 500 kHz. All these considerations led to the exclusion of the “classical” trigger design with several levels of filtering, realised with different technologies. The system is instead implemented as a single 128-steps synchronous pipeline running at the common experiment clock frequency of 40 MHz. This design extends the complexity of what is the classical first-level trigger, and removes completely the concept of second-level trigger. From the point of view of the trigger logic, the real challenge is the cluster-finding algorithm which performs non-trivial computations (linear interpolations) and has a cluster time resolution of 1/8 of the 25 ns sampling time-slice.

In the NA48 scheme, raw-data from the calorimeter are continuously flash-digitised and stored in a 200 μ s buffer, while the neutral trigger performs the complete event reconstruction for each 25 ns time interval irrespective of whether there was any physics-related activity in the detector. The various additions, conversions, multiplications etc., are performed by dedicated circuitry and the data are transformed into the trigger decision as they traverse the pipeline. The trigger decision is delivered synchronously with the detector signals, 25 ns \times 128 = 3.2 μ s after the first pipeline step. Although the circuitry is never idle, it spends a large fraction of its time computing the case when there was no detector activity. This approach has the advantages that there is no deadtime (the pipeline never has to wait for a preceding event), there are no limitations due to complex event topologies and the system is simple and stable to operate. The beams are debunched before being extracted from the SPS accelerator, so that the time distribution of kaon decays is roughly uniform within the 2.4 s machine spill. Whenever an event satisfies the neutral trigger cuts, the system identifies the time-slice at which the event occurred and generates a trigger relative to that time. The trigger decision is sent to the trigger supervisor system [4] which records the absolute event time.

The next section describes the algorithms used by the trigger to reconstruct the total kaon energy,

the decay vertex position, the centre of gravity and the number of peaks, and the way in which these computations were arranged to be suitable for a pipeline. This is followed by detailed descriptions of the implementation of each component of the pipeline.

3.2. Trigger algorithms

When a kaon disintegrates into n massless particles (photons) which impact on the calorimeter, the total energy E , centre of gravity C and vertex decay position D as measured by the LKr are given by

$$E = \sum_{p=1}^n E_p \quad (2)$$

$$C^2 = \frac{1}{E^2} \left(\sum_{p=1}^n E_p x_p \right)^2 + \frac{1}{E^2} \left(\sum_{p=1}^n E_p y_p \right)^2 \quad (3)$$

$$D^2 = \frac{1}{m_K^2} \sum_{p=1}^n \sum_{q \neq p}^n E_p E_q d_{pq}^2 \quad (4)$$

where E_p , x_p and y_p are the energies and impact positions of the photons in the calorimeter, m_K is the mass of the kaon, and $d_{pq}^2 = (x_p - x_q)^2 + (y_p - y_q)^2$ is the square of the distance between the impact positions of photons p and q . The centre of gravity represents the point where the kaon would have hit the front face of the calorimeter, if it had not decayed before. This virtual impact point is an estimate of the kaon’s transverse momentum.

The above expressions can be reformulated in terms of the total energy from the photons and of the impact positions on calorimeter strips oriented along the two axes x and y of the Cartesian grid formed by the read-out cells (see Fig. 1). Let $E_i^{(x)}$ and $E_i^{(y)}$ denote the energy on the i th strip in the x and y projections, respectively; E , C and D are then given by

$$E = \sum_i E_i^{(x)} = \sum_i E_i^{(y)} \quad (5)$$

$$C^2 = \frac{x_0^2}{E^2} \left(\sum_i E_i^{(x)} i \right)^2 + \frac{y_0^2}{E^2} \left(\sum_i E_i^{(y)} i \right)^2 \quad (6)$$

$$D^2 = \frac{1}{m_K^2} \left(E x_0^2 \sum_i E_i^{(x)} i^2 + E y_0^2 \sum_i E_i^{(y)} i^2 - E^2 C^2 \right) \quad (7)$$

where x_0 and y_0 denote the strip widths in the x and y projections. The importance of this rearrangement is that it allows these quantities to be reconstructed without a pattern recognition step to identify the photons. Indeed, the expressions are independent of the number of photons into which the kaon disintegrated. Eq. (7) gives an approximate expression of the vertex decay position, because the energy of the photon is spread over several adjacent strips instead of being concentrated at the photon impact position. The approximation is evident in the case of an event consisting of a single photon, for which D is zero in Eq. (4) whereas it is slightly positive in Eq. (7).

These quantities are then used to compute the proper decay time in units of the K_S^0 lifetime τ_S

$$\frac{\tau}{\tau_S} = \frac{m_K c^2 (z_0 - D)}{E c \tau_S} \quad (8)$$

where c is the speed of light and z_0 is the distance between the beginning of the decay fiducial volume (where $\tau \equiv 0$) and the front face of the calorimeter (where impact positions are defined).

Eqs. (5)–(7) summarise the actual way in which the trigger computation proceeds. The strips in each projection are built up from ‘supercells’ of 2×8 cells each (see Fig. 1). The summation of the energies in the 16 gain-equalised cells in each supercell is performed using analogue summing circuits; the supercell pulse is then digitised at the pipeline frequency of 40 MHz using 10-bit flash ADCs (2000 channels in total) and the output is filtered to exclude baseline and noise effects. The summation of the 16 supercells to build one strip is performed digitally. The digitisation, filtering and digital summations are performed in the Filter Module system (VFM, Section 3.4) while signal shaping, gain equalisation and analogue summation are done on the CPD modules which also house the calorimeter readout (Section 3.3). The x and y projections of the calorimeter are then used to compute the above quantities in the pipeline. The projection information is also used to perform

peak finding and to find out the arrival time of each peak. The construction of each sum in Eqs. (5)–(7) and the peak finding are performed in the PeakSum system (PSS, described in Section 3.6) and the final computations and trigger cuts are performed in the look-up table (LUT, Section 3.7) modules. A schematic of the complete trigger pipeline chain is shown in Fig. 2.

The trigger is also sensitive to rare neutral kaon decays which are composed entirely of photons, neutral pions and electrons, such as $K_L^0 \rightarrow e\bar{e}, \gamma\gamma, e\bar{e}\gamma, e\bar{e}\gamma\gamma, \pi^0 e\bar{e}$; since no minimum number of clusters is required, the trigger will detect all of them.

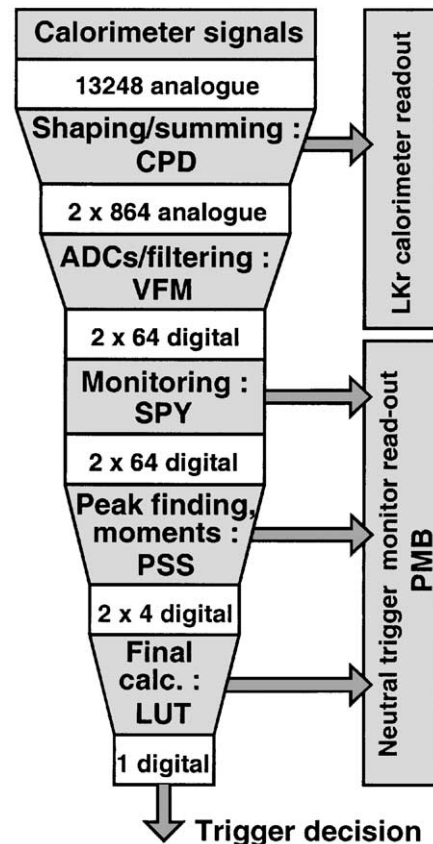


Fig. 2. Overview of the neutral trigger electronics chain. The processing occurs in a pipeline depicted from top to bottom in this diagram. The numbers give the number of signals (analogue) or number of pieces of digital information (typically 12–15 bits each) at each step.

3.3. The calorimeter pipeline digitiser

The CPD module is used to shape, digitise and store the calorimeter signals, and to read them out when an event is triggered [5]. The shaping produces a pulse from only the rapid initial rise of the current from the calorimeter signals. The module also performs the gain tuning and analogue summation to form the neutral trigger supercells out of groups of 16 individual cells. The shaping and trigger functions of the module are described in this section.

Each CPD module contains the circuitry to handle 64 calorimeter channels and is implemented as a three slot wide FASTBUS module. The 64 channels are arranged as an 8×8 square as shown on Fig. 1. Analogue sub-cards (CPDAS) contain the circuitry for two channels, mounted on both sides of each subcard. Fig. 3 shows the shaping and trigger parts of the CPDAS card. The main signal processing functions are implemented on a fully customized ASIC shaper called KRYPTON [6].

The analogue signals from the calorimeter line drivers are terminated with $100\ \Omega$ and differentiated with a 20 ns time constant. This differentiation reduces the dynamic range requirements of

the line receiver by a factor of two. The common mode rejection ratio is about 40 dB for frequencies up to 20 MHz. The differentiated signal is then processed by a Bessel filter implemented in the KRYPTON ASIC. For a 50 GeV energy deposit in a calorimeter cell, the input signal to the ASIC has an amplitude of $\sim 1\text{ V}$: this is the maximum signal that can be measured without saturation. The input signal has a typical rise time of 22 ns and a fall time of 3 μs . The Bessel filter has 9 poles, with a cutoff frequency of 10 MHz. The shape of the pulse after the Bessel filter is shown in Fig. 4: it is almost symmetrical, with a FWHM of 70 ns, a rise time of 43 ns and a fall time of 52 ns. The signal to noise ratio is 17 000–1 for the largest input signals. The output of the Bessel filter is common to the trigger and to the readout circuits. There is no significant difference between the digitised pulse shapes of the single cell calorimeter readout and of the neutral trigger analogue sums.

The trigger part of the ASIC contains a programmable attenuator and two amplifiers with high-impedance current outputs (transconductance amplifiers) used for the trigger summing. The programmable attenuator is used to equalise the response of the calorimeter cells before summing. It consists of a 6 bit modified $R/2R$ ladder network with $R = 7\text{ k}\Omega$ and switches made with CMOS transistors. The attenuation is adjustable in the range 0.48–0.82, corresponding to a $\pm 26\%$ gain variation (0.8% per LSB); this allows




Fig. 3. Block diagram of the trigger part of the CPD analogue subcards.




Fig. 4. Signal after shaping by the Bessel filter.

the channel-to-channel gain differences to be adjusted before building the analogue sum. A separate control bit allows to disable the trigger channel, in order to exclude noisy calorimeter channels. Each transconductance amplifier at the output of the KRYPTON ASIC consists of a differential amplifier using two operational amplifiers to improve the linearity, as shown in Fig. 3. The gain is set with a $525\ \Omega$ resistor: the matching between the two amplifiers is typically $\sim 0.2\%$, even though the tolerance of this resistor is 5% .

The summing of signals for the 2×8 supercells is done in two steps. First there is a partial current summing of the two channels on each CPDAS card as shown in Fig. 3. The transistor used for the common base has a typical β of 290, which is constant within 5% for the actual operating current. The second step of the summing is performed by AC coupling the eight CPDAS cards to a differential summing bus. The bus has a distributed resistance of $70\ \Omega$ in parallel with the impedance of the bus. The voltage on each summing bus is sensed and amplified by a differential transconductance amplifier (MAX435 [7]) which has a high common mode rejection. There are two MAX435s in parallel in order to supply $\pm 20\ \text{mA}$ signals to the VFM. Small printed circuit boards mounted on the auxiliary backplane of the CPD FASTBUS crate are used to distribute the analogue signals to the VFMs.

3.4. Filter Module system

The Filter Module system processes the analogue 2×8 supercell signals from the CPD modules and produces digitised, filtered and summed outputs. Each output corresponds to the calorimeter energy deposited in one horizontal or vertical strip. This is done in three stages:

- Digitisation of the 2×8 supercell signals.
- Digital filtering to reduce baseline and noise effects.
- Summation of 16 supercell energies belonging to one strip.

There are 128 strips overall, 64 in each projection. The system comprises 64 VFMs, each containing all the circuitry to handle two adjacent strips, and

is housed in four 9U VME crates. Clock, control and test cycle signals are distributed from a filter control module (FCM) in each crate.

3.4.1. The Filter Module

A functional view of the VFM is shown in Fig. 5. The filter operates as a pipeline running at $40\ \text{MHz}$, as does all of the following circuitry in the neutral trigger system. Inputs are band-pass filtered before entering the $40\ \text{MHz}$ 10 bit flash ADCs (Philips TDA8760). A signal of $1.1\ \text{mV}$ at the input of the FADC corresponds to one LSB. The input sensitivity of each channel is adjustable within a range of $\pm 6.8\%$ using a DAC. This is used in conjunction with the gain adjustment on the CPD, primarily to correct the gain variations

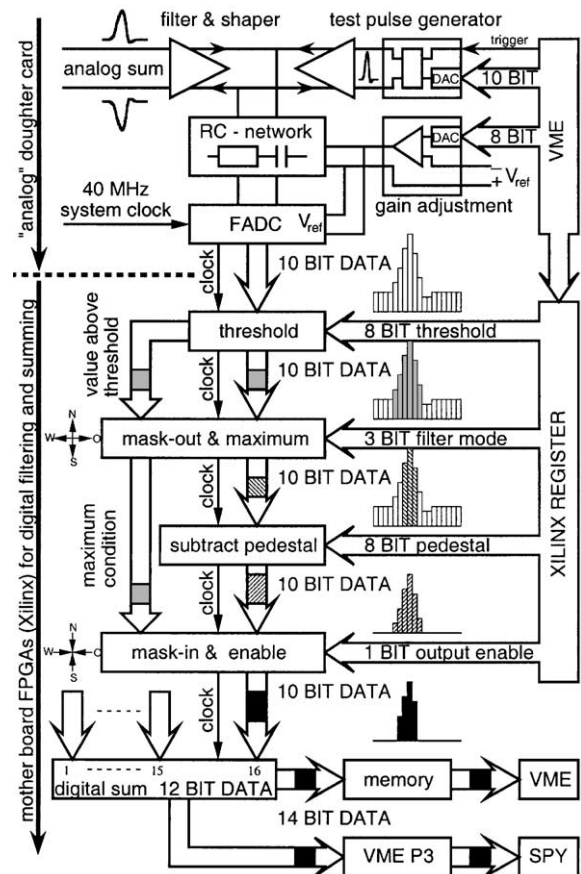


Fig. 5. Functional diagram of the Filter Module.

between the x and y projection summing circuitry on the CPD itself. All of the above circuitry is housed on individual daughter cards for each channel.

Filtering and summation are performed in field programmable gate arrays (FPGA) manufactured by Xilinx Corporation. The main 9U VME motherboard houses eight XC4010D filter FPGAs, with four channels each, and two XC4004A adder FPGAs. Each filter channel has two adjustable values (a threshold and a pedestal) and various mode selection bits, which are programmed when the trigger is software initialised. The normal operating mode is shown in Fig. 5. The filter sets to zero all ADC values except those associated with a calorimetric energy pulse. The filtering algorithms work as follows: The digitised signal must be higher than the preset threshold value in two consecutive time-slices of 25 ns. Requiring two time-slices reduces the sensitivity to electronic noise.

- Four time-slices are passed through the filter: the maximum, the two slices before it and the one after. Optionally, other combinations of samples around the maximum may be selected. For calibration and debugging purposes, the filter can also be programmed to pass all the time-slices (transparent mode) or nothing at all.
- The preset pedestal value is subtracted from each ADC value, setting to zero all negative results.
- An additional “mask logic” can be activated to recover the energy information from supercells belonging to the halo of the shower. These supercells usually do not fulfil the filter conditions above. When the mask logic is enabled, the supercells adjacent to one supercell above threshold pass automatically through the filter. This “shower leakage” correction has not proven very advantageous in the NA48 conditions, so the mask logic has been disabled since the 1998 data taking (see Section 6).

The signals from the 4 channels of each filter FPGA are then summed and passed as a 12-bit word to the adder Xilinx for the strip; a single adder FPGA handles the data of four filter FPGAs. Each of the two outputs of a VFM

represents the sum of 16 input channels in a 12-bit form, which is continuously updated at 40 MHz. An additional control bit (AR bit) allows a “bad channel” flag to be set in order to exclude some supercells from the peak-finding algorithm. This feature turned out to be particularly useful during the 1997 run, when a whole column of the calorimeter was without high voltage. The parity of the 12-bit data-word is calculated and sent along with the flag bit and the result of the summation to the subsequent stage of the trigger chain (the SPY system); data are transmitted via ECL twisted pair cables.

The level of noise and the pedestal stability are two critical aspects of the system, since they determine how low the filter thresholds can be set. The electronic noise is shown in Fig. 6, where the RMS of the pedestal for each VFM channel is histogrammed. The threshold value, selected for each channel individually, is five times the RMS value plus the mean pedestal value. Most channels have the threshold set to 3 ADC counts, and the noisiest ones to 5 ADC counts. Pedestal stability is shown in Fig. 7, showing the distribution of differences between pedestals as measured on two

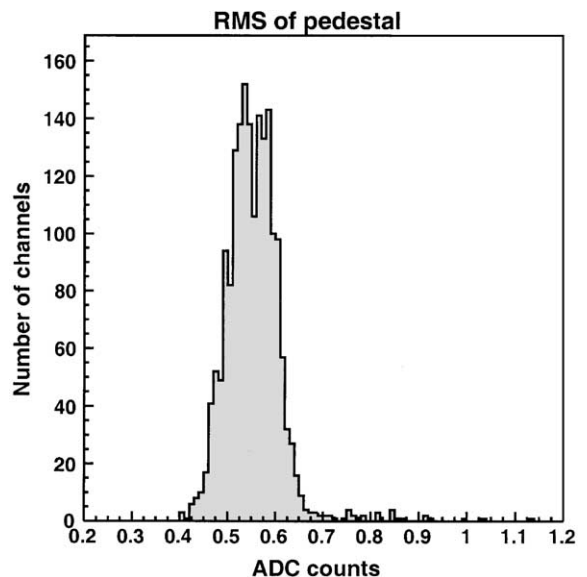


Fig. 6. Noise: distribution of the RMS of the pedestal for each channel.

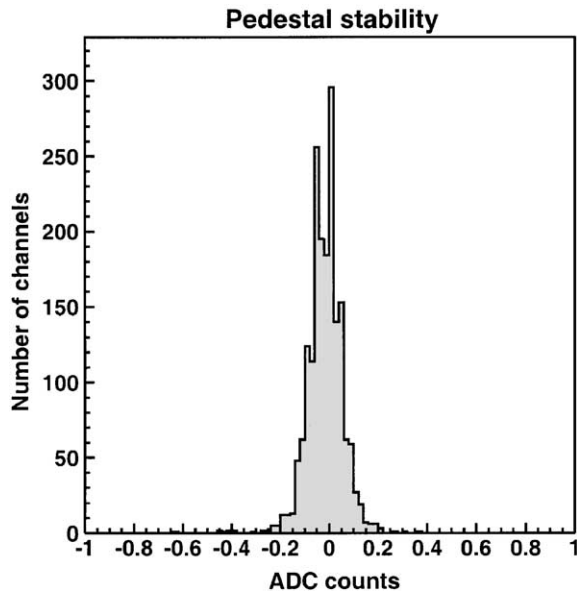


Fig. 7. Pedestal stability: distribution of the change in pedestal over a period of 14 weeks.

occasions 14 weeks apart. Pedestal shifts are much less than one ADC count.

A pulse generator which supplies test pulses to the daughter cards is included on the main motherboard. The nominal clock frequency of the system is 40 MHz, however the clock is distributed to the components on the card in a way that allows the modules to work properly up to 62 MHz; above this frequency, the components start to develop excessive heat.

3.4.2. Implementation

The analogue and digital power distribution on the PCB are completely separate: the separation line runs through the FADC, and they are connected only on the crate power supply rail.

The mechanical layout of the system was designed to simplify the layout of the cabling from the CPDs and to minimise cost and space requirements. The modules are housed in a 9U VME crate. Due to the non-standard mounting of components on both sides of the board, it was possible to house 32 channels on a single module. As a consequence, the front panel is shifted to the left with respect to the VME standard. To avoid the loss of one slot in the crate the single-sided

FCM was designed thin enough to accommodate the extra thickness of the adjacent VFM. The double sided boards also create difficulties in meeting the cooling requirements, because the distance between two daughter boards of neighbouring modules is only about 2 mm. Using a fan unit with an airflow of $\approx 1500 \text{ m}^3 \text{ h}^{-1}$ and static pressure equivalent to a 15 mm water column, the difference between inlet and outlet air temperature is around 20°C.

3.5. SPY and PMB modules

The section of the pipeline following the ADC modules provides a complete self-testing and online monitoring facility for the downstream components of the chain. This is done using SPY and Pipeline Memory Board (PMB) modules.

SPY modules have been developed into a commercial module (VN552, [8]) by the C.A.E.N. firm in collaboration with NA48. Each SPY has four digital 14-bit inputs, and handles the data from four VFMs. This implies the use of 32 SPY modules, 16 for each of the two projections, housed in two 9U VME crates. The modules can be configured via the VME bus to work in acquisition or in pattern generation mode.

In acquisition mode, the data from each input are fanned out to 4 main outputs (sent to the subsequent stage of the trigger chain) and to 4 monitoring outputs (sent to PMB modules). Two of the four inputs are also fanned out to auxiliary outputs. This allows the correct handling of the borders between cards in the cluster finding algorithm implemented in the PS system. The relative time skew of all the SPY module output bits is guaranteed to be smaller than 2 ns.

Each SPY module is equipped with an acquisition memory that allows the complete input data for 131072 clock cycles to be recorded and subsequently read via VME. This acquisition facility may be either run continuously or triggered by an external signal. The modules can be programmed to perform an online zero-suppression (i.e. record only data above the VFM system threshold), or to select data to be recorded according to a consistency check on the input patterns. The content of a clock counter is stored

in parallel to the data to allow the recovery of the information collected with one of the filtered modes.

In pattern generation mode, the SPY memory locations can be filled via the VME bus with predetermined patterns. The modules can be instructed to cycle these patterns on the outputs at 40 MHz, in order to test the downstream parts of the neutral trigger pipeline.

The PMB modules are used to read out many NA48 subdetectors (they are described in Ref. [9]). The modules are in 9U VME format. Each PMB has 8 digital inputs of 24 bits. The input data are temporarily stored in local buffers at 40 MHz. The buffers can store up to 204 μ s of data while trigger decisions and derandomisations are performed. Whenever a trigger arrives, the information in a time window around the trigger time-slice is read out and fed into the NA48 data acquisition system in order to be recorded with the event. The PMB modules used for the neutral trigger receive (a) all data from the SPY modules, (b) copies of the outputs of the Final Recombinator (FR, Section 3.6) and (c) copies of the output of some of the LUT modules (Section 3.7). This way, how the data coming from the VFM system are processed by the trigger can be checked and the reasons for trigger inefficiencies can be discovered. The PMB modules have programmable offsets which allow compensation for the difference in arrival time of the pertinent data from different parts of the pipeline for a given trigger.

3.6. The PeakSum system

The PeakSum (PS) system is the stage of the neutral trigger pipeline chain where the total energy and the weighted energy sums (moments) are computed from the individual strip energies, and where clusters are detected. The algorithm for cluster (peak) finding is more delicate and complicated than the one for moments calculation. A cluster is identified when a peak is simultaneously detected in space and time in a given projection. Spatial peaks are detected as peaks in energy, i.e. they correspond to strips where the deposited energy is greater than in neighbouring strips. A temporal peak is identified when the energy in a

strip at a given time-slice is higher than in the previous and following time samples. The time of each cluster is reconstructed with a granularity of 3.125 ns.

These calculations are performed independently and in parallel by two identical systems, one for the horizontal and one for the vertical projection. The major part of the computation is implemented in ASIC chips called NTP chips.

3.6.1. Moments algorithms

For each projection, the system computes the following sums, which are components of the expressions (5), (6) and (7):

$$ES = \sum_{i=-32}^{31} E_i \quad (9)$$

$$W = \sum_{i=-32}^{31} E_i i \quad (10)$$

$$QS = \sum_{i=-32}^{31} E_i i^2 \quad (11)$$

where E_i represents the energy on strip i of a projection.

The slight shift in the origin for the moments calculation is due to a construction limitation: it was not possible to position the centre of a strip at the exact centre of the calorimeter. This shift has no effect on the vertex reconstruction, since Eq. (7) is independent of the choice of origin; this is true as long as the origin is the same for both WS and QS. Concerning the centre of gravity computation, the asymmetry shifts the origin of the centre of gravity by 2 cm. The trigger cut in this quantity is anyway rather loose (15 cm), to allow for the difference in divergence of the K_S^0 and K_L^0 beams. Thus the shift can be safely neglected.

3.6.2. Peak-finding algorithm

The peak-finding algorithm operates on the strip energies to identify calorimeter clusters. A cluster is defined by the simultaneous detection of a *spatial* and a *temporal* peak. Let, $E_i(t_j)$ represent the energy on strip i from time-slice j . A spatial peak is detected when all of the following

conditions are met:

$$E_{i-1}(t_j) < E_i(t_j) \geq E_{i+1}(t_j). \quad (12)$$

The conditions for detecting a time peak are as follows:

$$E_i(t_{j+1}) < E_i(t_j) \geq E_i(t_{j-1}) \geq E_i(t_{j-2}) \quad (13)$$

and in addition,

$$E_i(t_j) > T \quad (14)$$

where T is a programmable threshold which can be set to a value higher than that of the thresholds in the VFM filters. A threshold in the peak-finding algorithm can avoid counting noise fluctuations as calorimeter clusters, while keeping the lowest possible threshold in the VFM system; this is desirable since the VFM filter thresholds affect the moments calculation. Since the 1998 data taking, a threshold of 16 ADC counts (2 GeV) has indeed been applied in the PSS to reduce the contamination from clusters generated by high-energy muons from $K_{\mu 3}$ decays.

3.6.3. Fine-time reconstruction

The clusters identified as above are then classified according to their arrival times in eight bins within each 25 ns time bin, giving a time reconstruction granularity of 3.125 ns. This information is used to allow a tight coincidence on the peak counting cut to minimise the loss of good events due to overlapping accidental particles. Care has been taken in the design of the calorimeter, in the cabling and in the electronic design and layout to ensure that the dispersion of the signal delays from one channel to the next is kept to a minimum, since there are no ways to adjust the time of the individual channels.

The reconstruction proceeds by analysing the shape of the pulse on the strip which passed the peak detection criteria defined in Eq. (13) at the time j (strip i in Eq. (12)). The algorithm looks at data from two earlier ($E_i(t_{j-1})$ and $E_i(t_{j-2})$) and one later ($E_i(t_{j+1})$) time bins. A reference value

$$E_{\text{ref}} = \alpha \cdot E_i(t_j) \quad (15)$$

is computed, corresponding to a fraction α of the maximum pulse height value associated with the

peak. This fraction is given by

$$\alpha = \frac{1}{2} + b_8 \cdot \frac{1}{8} + b_{16} \cdot \frac{1}{16}, \quad (16)$$

where b_8 and b_{16} are 1-bit programmable registers which allow the value of α to be changed slightly. The system then estimates the time t_{ref} at which the pulse height is equal to E_{ref} on its rising edge. This is done by performing a linear interpolation between $E_i(t_{j-2})$ and $E_i(t_{j-1})$ (see Fig. 8). α is chosen to be close to $\frac{1}{2}$ so that the time of the pulses is reconstructed from a region where the pulse derivative is a maximum, in order to get the best time resolution. Simulation studies showed that this linear interpolation gives a better intrinsic time resolution with respect to a technique which performs a parabolic interpolation over the 3 time samples around the pulse maximum; this is particularly true for small pulses, where the FADC quantisation error is important.

A limitation of the algorithm is that the interpolation is not extended to time-slices earlier than t_{j-2} . Whenever $E_{\text{ref}} < E_i(t_{j-2})$, then t_{ref} is assigned to be t_{j-2} without further calculations. This becomes relevant if the pulses are much longer than the sampling period. During the trigger design phase, it was still unclear whether the pulses would have been sufficiently short to avoid this problem: the b_8 and b_{16} registers were thus included in the logic to tune this part of the interpolation procedure.

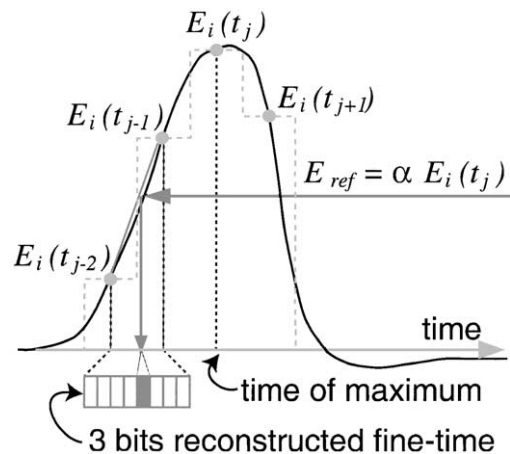


Fig. 8. The peak time reconstruction algorithm.

Additional features of the peak algorithm are:

- A masking procedure to veto the possibility of detecting peaks in the two neighbouring strips of a previously detected peak with the reference time reconstructed in the same or the following time-slice. This means that if a peak is detected in strip i with time in time-slice j , the algorithm disables the possibility of detecting a peak in strip $i+1$ or $i-1$, with a reference time reconstructed in time-slice j or $j+1$. This procedure is provided to avoid double counting of peaks due to small time skews between the signals on neighbouring strips.
- A way of excluding from the peak-finding algorithm the regions of the calorimeter with known problems (i.e. dead cells). This is done by flagging the supercells with problems in the VFM system which sets the AR control bit for each strip if any flagged supercell has been used in compiling the strip energy. Dangerous situations occur when strip i has its AR bit set and strips $i-1$ and $i+1$ both have detected peaks. This could either be two real peaks, or a single peak in which the energy has been split by the dead cell(s). When this dangerous situation is detected, an “anomalous peak” flag is sent on to the LUTs.

3.6.4. Implementation

The PS system is implemented as two identical independent chains for the x and y projection. The system for each projection consists of eight “PeakSum” boards (PSB) and an FR board, as shown in Fig. 9, and is housed in a 9U VME crate. Each PSB is responsible for reconstructing energy sums and peaks occurring on the strips connected to its eight main inputs. The eight boards are daisy chained in two groups of four to sum up the energy sums and peak information from the two halves of a projection. The FR module performs the final summation step from the two groups and sends the results (ES, WS, QS and number of peaks in the sub-time-slices) to the LUTs.

Each PSB has eight main 14-bit inputs (12 signal bits, a channel flag and a parity bit) and two auxiliary ones. The input signals are supplied from the SPY modules. The main inputs correspond to

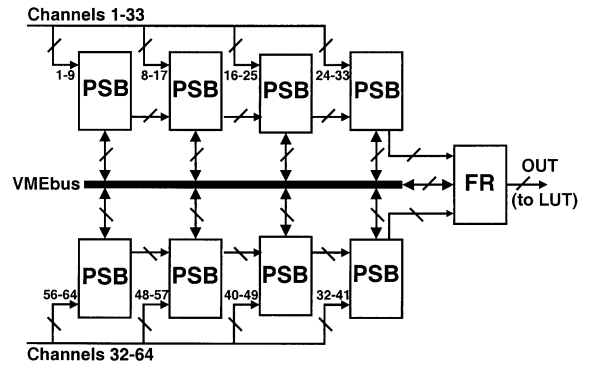


Fig. 9. PSS scheme for a single projection. Module to module connections are made customizing the J2 and J3 VME busses.

the information from 8 adjacent calorimeter strips. The auxiliary inputs contain duplicate information of the neighbouring cards, so that the module actually has information from ten adjacent strips. The data from the two additional strips is used in the peak-finding algorithm to find clusters centred at the edges of the modules (see later).

As with the rest of the neutral trigger, the PS operations are fully pipelined so that for every clock cycle the system accepts time-slice t_j information for the 64 strips, and produces results for time-slice t_{j-n} where n represents the overall latency of the system.

The single PS module itself is organised as two parallel daisy-chains (see Fig. 10), one devoted to peak-finding algorithms and the other to moments calculations. The base element of these local daisy-chains is the Neutral Trigger Processor (NTP) chip. The NTP chip is a VLSI semi-custom CMOS standard cell device specifically developed for NA48. The schematic of this chip is shown in Fig. 11. To avoid multiple development costs, all functions are implemented on a single NTP chip and the appropriate functional blocks are activated depending on the location of each individual chip on the boards. There are six different configurations: local peak, local EWQ, interboard peak, interboard EWQ, FR peak and FR EWQ. Peak and EWQ, designate, respectively, the mode of operation for peak-finding (and fine-time reconstruction) and energy momenta calculation.

Local EWQ configured chips handle four strip channels each; they are configured differently

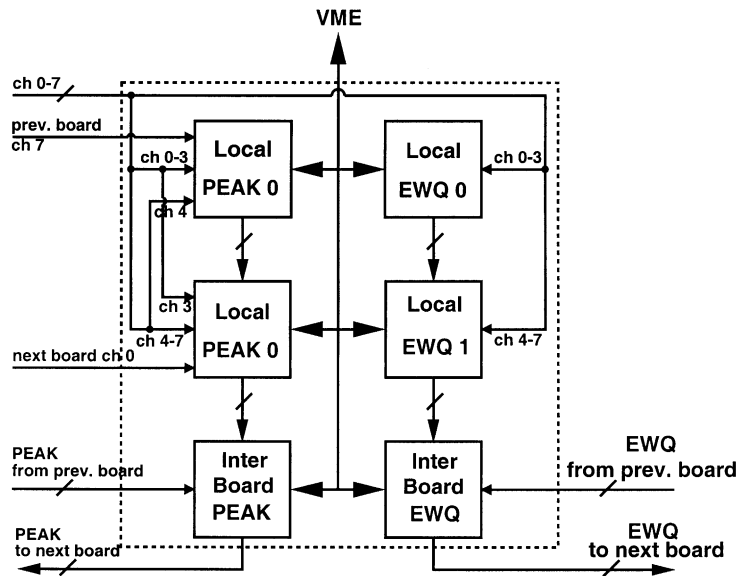


Fig. 10. Scheme of a single PS board.

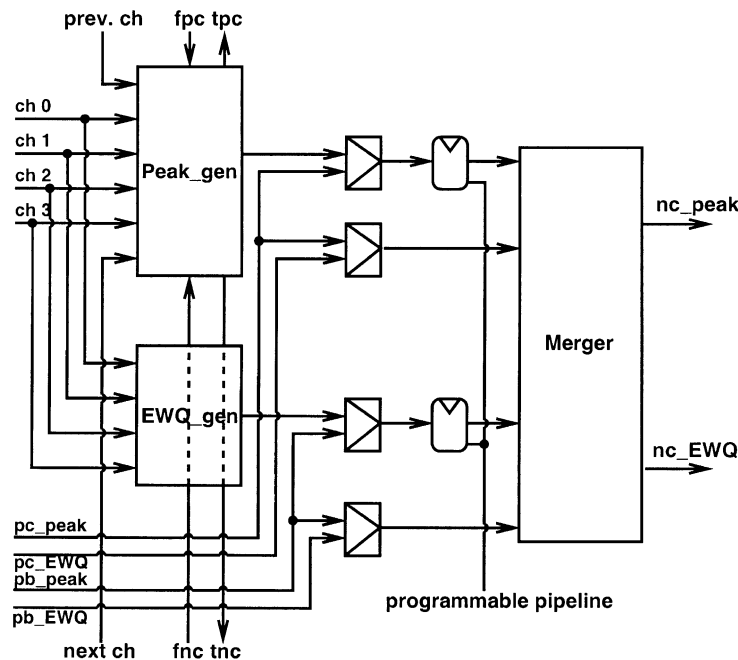


Fig. 11. Scheme of the NTP chip internal structure. EWQ-gen and Peak-gen blocks perform peak detection and momenta calculation. The merger block combines results from different chips or modules. Programmable pipelines allow temporal alignment before merging.

according to the weight needed in the moments calculation. The chip configuration is set in the 4-bit register *W* and in the two bits *CROSS* and *PN*. The *W* register has values from 0 to 15 and defines the

position of the four strip channels in the calorimeter. The *PN* bit takes into account the previously mentioned slight shift in moments calculation. The *CROSS* configuration bit inverts the succession of

input channels, and is to be set together with the PN bit for channels on the calorimeter side with negative weights assigned. The equations defining the digital computations in the chips are

$$ES = E_0 + E_1 + E_2 + E_3 \quad (17)$$

$$WS = f_0 E_0 + f_1 E_1 + f_2 E_2 + f_3 E_3 \quad (18)$$

$$QS = f_0^2 E_0 + f_1^2 E_1 + f_2^2 E_2 + f_3^2 E_3 \quad (19)$$

where $f_i = 4W + i + \text{PN}$ when $\text{CROSS} = 0$. When $\text{CROSS} = 1$, f_i is modified to be $4W + (3 - i) + \text{PN}$. Note that although Eq. (6) uses both positive and negative weights, the computation is performed locally with only positive weights and the FR performs the last step of combination by subtracting the contributions from the two sides of the calorimeter.

The intermediate sum results are transported between chips and along the daisy chain in a floating point format with leading-bit suppression, using a 9-bit mantissa m (10-bit effective) and 3, 4 and 5-bit exponent e for ES, WS and QS, respectively. The value is given by $(m + 2^9) \cdot 2^{(e-1)}$ unless the exponent is zero, in which case the value is given by m . This format is also used to transmit the final result to the lookup table modules. The results from the two local EWQ chips on each board and the daisy chained results from the previous board are combined using an NTP chip configured in interboard EWQ mode. This is done in a small on-board daisy chain as shown in Fig. 10. The use of this form of daisy chaining in the system minimises the pin count on the NTP chips.

The peak-finding algorithm is executed by an NTP chip in local peak mode for each group of four channels. Each chip receives copies of the information on the two strips adjacent to the group. The output is coded into eight 3-bit numbers representing the number of peaks in each of the 3.125 ns time intervals within the 25 ns clock cycle. A ninth 3-bit number counts the number of peaks associated with an AR bit flagged strip within the entire 25 ns window. A parity bit over the entire 27-bit data-word is generated. As for the EWQ data, the information from the two local peak chips and the daisy chained results from the

previous board are combined as shown in Fig. 10, using an NTP chip configured in interboard peak mode.

The FR completes the daisy chain using FR, EWQ and FR Peak configured chips to combine the data from the two halves of the calorimeter.

“Local” chips have an internal computation pipeline latency of 11 clock cycles and a merging latency of 3 clock cycles. “Interboard” merging operations take 3 clock cycles. The total latency within a board is thus 17 clock cycles. The latency within the FR is 5 clock cycles, leading to a latency for the entire system of 31 clock cycles (775 ns).

3.7. Look-up tables

The LUTs implement the final part of the computation shown in Eqs. (5)–(8), perform event selection and define the trigger time. This is done by combining the outputs of the peak sum systems from the x and y projections. The system is implemented in a tree of identical LUT modules, as shown in Fig. 12, housed in a single FASTBUS crate. These modules can be programmed to perform diverse computations by the appropriate configuration of FPGAs and LUT memories.

3.7.1. Algorithms

The following describes the functionality of each LUT module in the tree shown in Fig. 12.

A0M The A0M module computes the average of the total energies from the two projections: $m_0 = (ES_x + ES_y)/2$. The incoming data are encoded as floating point numbers as described in Section 3.6.4, and the addition is done by adjusting the normalisation of the smaller number, adding the mantissæ directly and then encoding the result into 9 bits to be sent to the ECT module.

A1M The A1M module combines the first moment information of the two projections for the centre of gravity calculation: $m_1 = \sqrt{WS_x^2 + WS_y^2}$. The technique is similar to the one used in A0M.

A2M The A2M module combines the second moment information of the two projections for the D calculation: $m_2 = QS_x + QS_y$. The same technique as A0M is used.

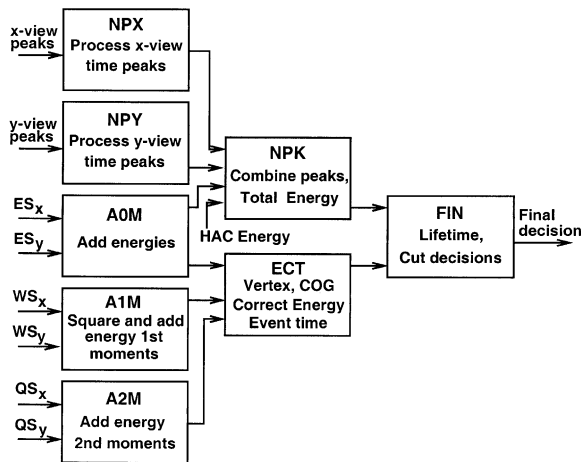


Fig. 12. Schematic of the network of eight 40MHz LUT modules.

NPX This module processes the peak data from the x projection, which consists of the number of peaks seen in each 3.125 ns sub-time bins. The module inspects the data and determines the maximum number of peaks seen within a programmable time window (3–5 sub-time bins). This time window is allowed to begin at any of the eight possible positions within one 25 ns pipeline step and uses some of the data from the following pipeline step where necessary.

NPY A copy of NPX for the y -projection peaks.

NPK This module computes the final peak-trigger decision based on the outputs of NPX and NPY. It also combines the data from the hadron and electromagnetic calorimeters to obtain the total energy for the event used in the $K^0 \rightarrow \pi^+\pi^-$ trigger (see Section 4). Several decision bits are used for each of the above cuts in order to defer the final decision to later stages of the pipeline.

ECT This module performs three functions: (1) it performs a parabolic fit between the m_0 values of three consecutive time-slices to determine the time at which any energy was deposited; (2) it computes the centre of gravity by dividing m_1 and m_0 , and (3) it performs the majority of the D computation as described in Eq. (7).

FIN This module completes the D calculation, computes τ/τ_S according to Eq. (8) and performs the actual decision on whether a trigger should be

generated. It also contains logic to avoid errors in matching up the different trigger conditions for events which occur on the boundaries between two time-slices.

3.7.2. Implementation

Each module consists of a series of FPGAs and $256k \times 12$ bit memories as shown in Fig. 13. Data are clocked through this circuitry at 40 MHz. Software has been written to generate the FPGAs routing definition files and the LUT programming functions from a configuration file. This configuration file is also used to automatically generate a Monte-Carlo simulation of the LUT tree, which is used to verify the programming. A basic statement in the configuration file defines a group of signals (a bus) and specifies where these signals are routed. Additional tokens can be added to the statement to perform combinatorial logic functions. The AND, the OR or the XOR (parity) of a group of bits may be computed and bits may be inverted. Pipeline step delays may be specified on statements in the configuration file. For convenience (in visualizing the algorithms), a ‘delay’ of -1 step (i.e. an advance) is also permitted (the software handles this by automatically delaying all the other signals by one step).

Each module contains a recorder unit with similar functionality to those on the SPY modules

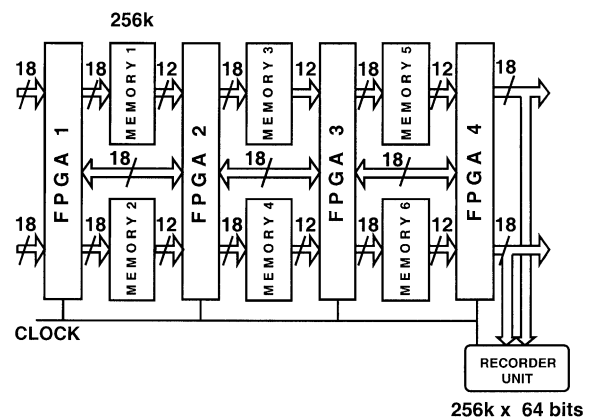


Fig. 13. Block diagram of one of the LUT modules. The six memories and the four FPGA devices which are used to route the signals between LUTs in a programmable way are shown. A recorder unit is also included on the board and is used for monitoring purposes.

(Section 3.5). The modules can operate up to 43 MHz.

4. The total energy trigger

A condition based on the total energy of the event is used in NA48 as a component of the trigger for charged decays. The total energy is computed summing the energy from the electromagnetic calorimeter (as given by the AOM LUT) with the information from the iron-scintillator sampling hadron calorimeter (HAC) installed downstream of the LKr calorimeter. The energy from $\pi^+\pi^-$ events is collected in both the electromagnetic and hadron calorimeters depending on the way in which each individual hadronic shower develops. The computation of the total energy proceeds as follows:

- The scintillator strips in the HAC are read by an array of 172 photomultiplier tubes. The electronic signals from the phototubes are shaped like the ones from the electromagnetic calorimeter (Fig. 4).
- The signals are fed into four standard CPD modules (44 photomultiplier tubes for each CPD, the remaining 20 CPD channels not being used). The CPDs are used in a very similar fashion to the LKr calorimeter CPDs to (a) produce analogue sums (12 overall) of the hadron energies and (b) read out the HAC for offline studies.
- A single extra VFM card is used to accept the analogue sums from the HAC and digitise and filter them in the same way as for the electromagnetic calorimeter supercells. The two digital outputs from the VFM correspond to the total energy in the front part and the back part of the HAC.
- The two digital outputs are fed directly to LUT FAC and HAC modules, which apply an appropriate intercalibration constant and add the two components of the HAC energy together. They then delay the data by 31 clock cycles to keep it synchronised with the electromagnetic calorimeter data which has passed through the PSS. This total HAC energy is then summed with the

electromagnetic energy (output of the AOM module) in the NPK module.

In the first year of data taking, the total energy signal supplied to the $\pi^+\pi^-$ trigger was simply a bit that indicated in which 25 ns time-slices the total energy was above threshold. The system was upgraded for subsequent years by reprogramming the LUTs and FPGAs to (a) detect the time-slice containing the peak of the total energy (above a given threshold) and (b) supplying a bit which is valid for a fixed number of time-slices before and after each peak.

5. Set-up, calibration and operation

The complete chain of the neutral trigger electronics is installed in racks in an electronics control room situated immediately adjacent to the liquid krypton calorimeter detector. The CPDs occupy 28 FASTBUS crates, the VFMs four VME crates, the SPY and PS systems two VME crates each, the LUTs one FASTBUS crate and the PMB system one VME crate. The installation of the system was completed in spring 1997.

Due to the synchronous nature of the system, the phase of the 40 MHz clock distribution to the crates is tuned so that the data of each input module are clocked in at a part of the clock phase when the data are stable. Digital data connections from one system to the next are all cut to the same fixed length. The transmission quality is then systematically studied as a function of the clock cable length over the entire 40 MHz clock cycle, and a suitable clock cable length is selected for all the five separate interconnections within the neutral trigger. Three further adjustments are needed for the HAC VFM to LUT connection and for two of the connections from the pipeline to the PMBs. Typically, the unusable part of the 40 MHz cycle is <25% of the full cycle, so this set-up procedure is not difficult and, once done, transmission remains stable. The system continues to operate reliably even if the transmission delay in the data cable is more than one complete clock cycle long. An example of this is shown in Fig. 14, where the fraction of successful transmissions is

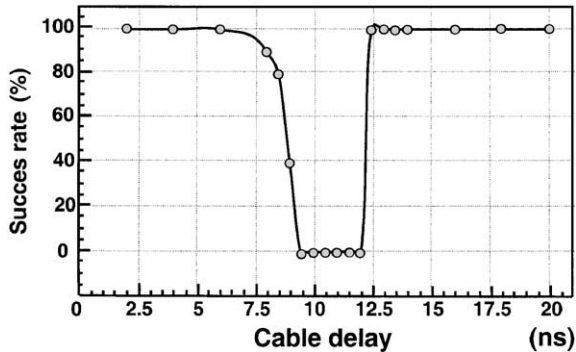


Fig. 14. Plot showing the success rate of a LUT board test as a function of the interconnecting data cable length. There is an unsafe region between 7 and 12.5 ns where the set-up and hold times of the receiver latch are not respected. For cable delays below 7 ns, data corresponding to the same pipeline stage are latched. For cable delays above 12.5 ns, the data which are latched correspond to the previous pipeline step, i.e. an extra pipeline stage is included in the transmission along the cable. Both modes of operation (short cables or long cables) are reliable.

shown as a function of the data cable length for a loopback test on a LUT module. The x and y outputs of the CPDs have propagation delays which differ by 5 ns relative to each other, due to differing topologies needed in the summing busses of the two views on the CPD mothercard. A correction for this is applied by adding an extra 5 ns to the clock cable of the two VFM y projection crates. Since a sufficiently large part of the clock phase is good for transmission to the SPYs, the x and y projection SPY systems can be run synchronously.

The clock delay on the HAC VFM is adjusted so that the peak of signals from the HAC occur in the same part of the clock phase as the peak signals from the electromagnetic calorimeter.

The neutral trigger system is controlled by a VME based microprocessor which is connected to the VFM, SPY and PSS VME crates via a VIC [10] bus, and to the LUT FASTBUS crate via FVSBI [11]. The configuration constants are downloaded at the beginning of each run from the microprocessor. The LUT contents are computed each time the modules are downloaded. Cut values are read by the LUT content computation

routines, so they can be altered without requiring a recompilation.

The neutral trigger is calibrated in three subsequent steps. The first step consists in setting the initial values for all the gain adjustment parameters in the CPDs and VFMs; this is done using the pulse-height slopes measured during the tests of each CPD and VFM module. When the system is fully operational, a channel-to-channel intercalibration is performed to compensate for gain differences in the various stages of the electronics chain of the trigger, and to identify dead channels. The intercalibration procedure uses the LKr pulser to collect data with a series of special patterns. The patterns are selected using different masks on the CPDs, pulser and VFMs, and are chosen to allow the independent measurement of the slope for every channel. Finally, physics data ($K_L^0 \rightarrow \pi^\pm e^m \nu$ and $K_S^0 \rightarrow 2\pi^0$) are used to check the overall energy scale. The average energy calibration is measured with the calibration pulser is 134 MeV/ADC count. This value changes to 140 MeV/ADC count when measured with physics events, because of the lateral energy loss from the shower due to supercells below the filter threshold of the VFM system.

Fig. 15 shows the difference in the calibration constants before and after the 8-weeks running period in 1997. The stability is better than 0.5%. The tails in the distribution are due to electronics channels which have been substituted and recalibrated.

6. Performance

The neutral trigger system has successfully operated since the first NA48 physics run in 1997. The online trigger cuts are summarised in Table 1, along with the corresponding selections applied in the offline analysis. The peak cut is implemented by summing up the number of peaks within a 9.375 ns wide time window (see NPX description in Section 3.7). The cut requiring <6 peaks is not applied when there are peaks surrounding this time window, since this is evidence that the event was accompanied by other (accidental) clusters not associated with the

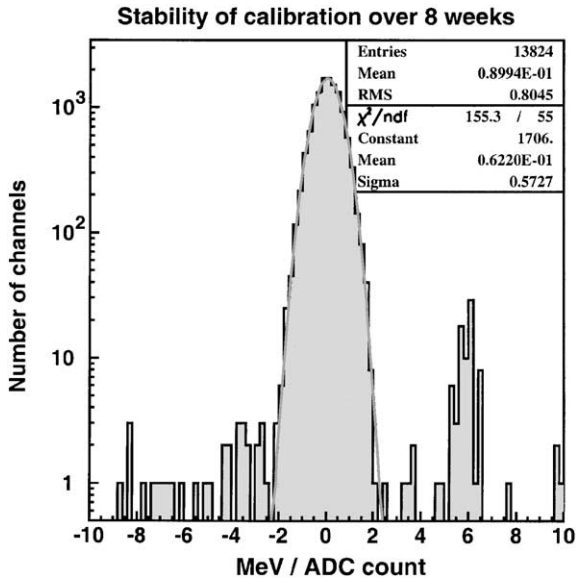


Fig. 15. Change in calibration constants over an 8 week data taking period.

Table 1

Summary of the variables computed online by the $\pi^0\pi^0$ trigger, along with their cut values and the corresponding cuts applied in the offline analysis

Cut	Trigger	Offline
Energy	> 50 GeV	> 70 GeV
Centre of gravity	< 15 cm	< 10 cm
Lifetime	< 5 τ_S	< 3.5 τ_S
x projection peaks	< 6 peaks within 3 \times 3.125 ns	4 photons ^a
y projection peaks	< 6 peaks within 3 \times 3.125 ns	4 photons ^a
Total energy (for $\pi^+\pi^-$)	> 35 GeV	> 70 GeV

^aNo fifth photon above 1.5 GeV within ± 3 ns of the event.

triggered event. The rate of particles in the detector is ~ 500 kHz, and the output rate of the neutral trigger is ~ 2 kHz. The number of $\pi^0\pi^0$ triggers collected in each burst is of the order of 5000; about 30 of these events satisfy the full offline $\pi^0\pi^0$ selection criteria (detailed information on the neutral event selection is given in Ref. [1]).

Trigger performance studies are done by analysing the data received by the PMB monitoring

system. Information collected by the PMBs, or “trigger variables”, are compared with the “offline variables”, reconstructed using the data collected by the main read-out system. As mentioned in Section 2, downscaled events triggered by the NHOD are collected during ϵ'/ϵ data taking. These NHOD triggers provide a sample of unbiased events for neutral trigger efficiency studies.

The neutral trigger energy resolution for single photon detection is checked performing the ratio between offline and trigger cluster energies in bins of 5 GeV. The mean and the sigma of the distribution in each bin are then obtained from a Gaussian fit and are plotted in Figs. 16 and 17, respectively. Events where multiple photons overlap in one projection are not included in the plots to avoid cases where a low energy shower is hidden by the tail of a very energetic shower.

It is clear from Fig. 16 that the neutral trigger underestimates the energy in the case of low energy photons, while high energy clusters are overestimated. This non-linearity is mainly caused by the energy dependence of the lateral shower leakage into supercells which are below threshold in the VFMs. The same plot also shows that linearity can be improved by enabling the mask logic in the VFM system, as explained in Section 3.4; the mask logic allows the supercells adjacent to a supercell above threshold to pass through the filter, thus correcting for the shower leakage. The neutral trigger has an energy resolution (Fig. 17) of about 3% in the high energy range. The resolution degrades when the mask logic is enabled because the additional supercells increase the noise, particularly for photon energies around 20 GeV. Overall, given the level of noise in the system and the energy distribution of the photons, it is somewhat more advantageous to run without the shower leakage correction: the mask logic has thus been disabled since the 1998 data taking period.

Fig. 18 illustrates the trigger performance with respect to the reconstruction of the kaon proper decay time τ . The difference between the trigger and offline proper decay time is plotted, in units of the nominal mean lifetime τ_S of the K_S^0 . The distribution is obtained using NHOD triggered




Fig. 16. Mean value of the ratio between offline reconstructed energy and neutral trigger energy, as obtained from a Gaussian fit. The mean value is plotted as a function of the cluster energy, in the case of single photons. Open and closed circles represent, respectively, data with and without the shower leakage correction enabled in the VFMs.




Fig. 17. Sigma of the ratio between offline reconstructed energy and neutral trigger energy, as obtained from a Gaussian fit. The sigma is plotted as a function of the cluster energy, for single photons. Open and closed circles represent, respectively, data with and without the shower leakage correction enabled in the VFMs.

events. The shaded area contains the $\pi^0\pi^0$ candidates with four reconstructed clusters, whose distribution has an RMS of $0.46 \tau_S$, with negligible tails. The neutral trigger underestimates τ by about $0.5 \tau_S$ on average: this is mainly due to a small left-over offset in the energy scale calibration, which ensures high efficiency on all events. The vertex position needed to compute the proper decay time is reconstructed with an accuracy of around 3 m over a distance of 100 m. The unshaded area in Fig. 18 represents the same distribution of differences between the trigger and offline values of τ with the inclusion of $\pi^0\pi^0$ candidates with extra (accidental) clusters. The distribution has a tail only on the side where the trigger lifetime is too low, meaning that it is shifted on the safe side with respect to the trigger cut value. This implies that accidentals do not cause inefficiencies while triggering on $\pi^0\pi^0$, which was a design goal of the neutral trigger system.

A crucial feature of the neutral trigger system is the online evaluation of clusters' times performed

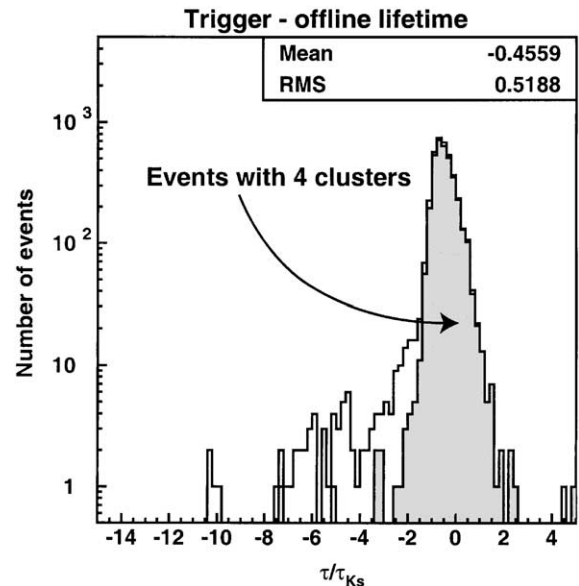


Fig. 18. Difference between lifetimes reconstructed by the trigger and offline. The shaded area is the distribution of the $\pi^0\pi^0$ candidates with four reconstructed clusters.

by the PSS. This is essential to reduce the amount of accepted $K_L^0 \rightarrow 3\pi^0$ decays while avoiding large inefficiencies due to accidental activity in the detector. Clusters are considered as coincident when they fall within a ~ 10 ns time window (see Table 1). The width of the time window is chosen on the basis of the time resolution that the system can achieve: the better the PSS is able to reconstruct the cluster times, the smaller the coincidence window can be defined. A small time window reduces the losses due to accidental decays superimposed to the 4 photons of a $K^0 \rightarrow 2\pi^0$ event.

Fig. 19 shows the time resolution of the PSS. The difference between the individual time of each trigger cluster and the offline global event time is plotted for all the K^0 decays satisfying the full offline $\pi^0\pi^0$ selection criteria. The $\pi^0\pi^0$ sample used to produce Fig. 19 is selected with the additional requirement that no accidental activity is present over many clock cycles around each event. This ensures that each peak found by the trigger corresponds to one of the 4 photon clusters of the $\pi^0\pi^0$ event. In fact, it is not possible to

match the trigger peaks with the calorimeter cluster on the basis of their geometrical position, since the PSS does not provide such information.

The PS resolution is better than 2.5 ns in both views (Fig. 19 shows just one projection). This is an excellent performance, since this resolution is obtained with no geometrical selection criteria on the photon impact positions and thus is affected by the effects of possible skews between different calorimeter channels.

The neutral trigger efficiency is evaluated using the control sample of events triggered by the NHOD which survive the full offline $\pi^0\pi^0$ selection. Events are defined as inefficient if the neutral trigger bit in the trigger supervisor pattern unit is not set in a window of ± 2 time-slices around the triggered event.

Fig. 20 shows how NHOD events are distributed as a function of the proper decay time in units of τ_S ; inefficient events are also shown on the same plot. The events have been identified as K_S^0 or K_L^0 decays using the tagging detector, which tags the proton hitting the K_S^0 target. The K_S^0 distribution contains a contamination of $\sim 10.6\%$ of the true K_L^0 events (mistagging). The cut on τ/τ_S applied by the neutral trigger is set to provide high efficiency for events with $\tau/\tau_S < 3.5$ which are accepted in the ϵ'/ϵ offline analysis. In addition, K_L^0 events are weighted according to the K_S^0 lifetime distribution, in order to have the same decay spectra for both beams (see [1] for details on the tagging procedure and the offline analysis).

During the 1997 run, the LKr calorimeter operated at half the nominal 3 kV drift voltage, and a 4 cm column in the x projection was without high voltage (due to a failure in the HV connection). The neutral trigger demonstrated its flexibility in this difficult experimental situation: the AR control bit in the VFMs was set for this whole column and the number-of-peaks cut in the x view was loosened to ≤ 6 peaks. The trigger efficiency for $\pi^0\pi^0$ decays during the 1997 run was $(99.88 \pm 0.04)\%$, with no differences between K_S^0 and K_L^0 .

Since 1998, the high voltage of the calorimeter has been set at 3 kV and beam intensities have been increased. The overall neutral trigger efficiency during the 1998 and 1999 data taking was

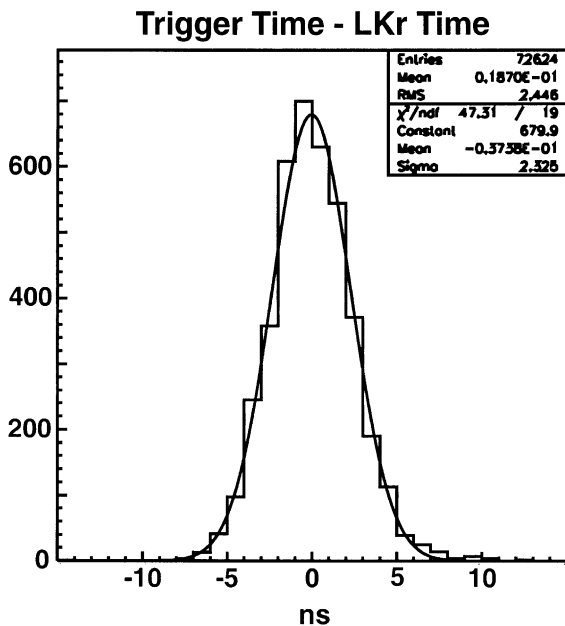


Fig. 19. Time resolution in the PS on the x -view. The figure shows the difference between the times of the trigger peaks and the global event time as obtained from the LKr calorimeter.

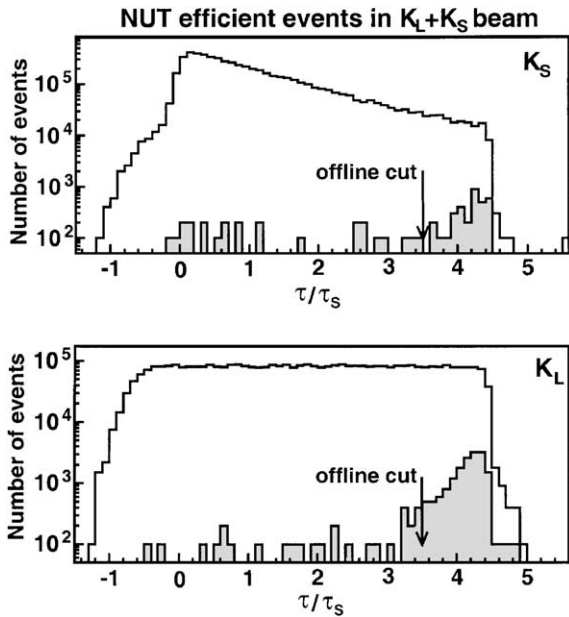


Fig. 20. Distribution of efficient and inefficient events in the neutral trigger monitoring sample, as a function of τ/τ_S . The top plot is for K_S^0 decays and the bottom one for K_L^0 decays. The blank histograms give the distribution of τ/τ_S for all events in the monitoring sample and the shaded histograms indicate the inefficient events.

$(99.920 \pm 0.009)\%$, and was the same for K_S^0 and K_L^0 events: this fact is of utmost importance for the $\Re(\epsilon'/\epsilon)$ measurement, since it implies that no correction has to be applied to the double ratio 1 to correct for biases in the neutral trigger behaviour. The trigger efficiency for these two run periods is plotted in Fig. 21 as a function of the kaon energy, for K_S^0 and weighted K_L^0 decays.

Event losses have been studied by looking at the trigger and offline information recorded along with each inefficient event. The main source of inefficiencies is the occurrence of two distinct classes of accidental decays, namely:

- (1) High energy Λ -baryon decaying far upstream in the K_L^0 beam-line within ~ 40 ns of the good $\pi^0\pi^0$ event. The energy deposited by the isolated protons from these decays shifts the centre of gravity reconstructed by the trigger towards higher values, so that the event fails the 15 cm trigger cut on this quantity.




Fig. 21. Trigger efficiency as a function of the kaon energy. Open circles represent K_S^0 events, while full square boxes represent K_L^0 weighted events.

- (2) Accidental $K_L^0 \rightarrow 3\pi^0$ decay occurring within 10 ns (the width of the peak finding time window) of the good $\pi^0\pi^0$ event. This causes a higher number of reconstructed peaks in the trigger with respect to the cut value.

Another cause of inefficiency is an occasional energy mismeasurement due to the bad response or mis-calibration of a cell.

7. Conclusions

The neutral trigger for the NA48 experiment has successfully operated over several years. The novelty of this system is largely its architecture, where the entire trigger computation is performed in a single continuously running pipeline. This approach has proven to be very successful: trigger operation is convenient and stable, with no unforeseen difficulties during set-up since a sufficient amount of hardware diagnostic circuitry had been implemented. The feature of having no deadtime is most important in the ϵ'/ϵ analysis. The system was also sufficiently flexible to allow rare decay triggers to be added. The trigger

efficiency was carefully studied and is well understood. The overall measured efficiency for $K^0 \rightarrow \pi^0 \pi^0$ events is $(99.920 \pm 0.009)\%$; it is the same for K_L^0 and K_S^0 events, thus bearing no correction on the double ratio measurement.

Acknowledgements

The neutral trigger is only one of a large number of subsystems which make up the NA48 experiment. It has been a great pleasure to work with our colleagues on the NA48 experiment, and we are most grateful for their help and advice. We would particularly like to thank the following CERN summer students who worked on the neutral trigger project: J. Anderson, K. Harder, G. Quirico. We are indebted to the Austrian Federal Ministry for Science and Transport for the financial support of this project (GZ 616.360/2-IV/6/95).

References

- [1] A. Lai, et al., A precise measurement of the direct CP violation parameter $\Re(\epsilon'/\epsilon)$, CERN-EP/2001-067.
- [2] G.D. Barr, et al., Nucl. Instr. and Meth. A 370 (1996) 413.
- [3] C. Cerri, The NA48 liquid krypton calorimeter: electrode structure, front-end electronics and calibration, VIth International Conference on Calorimetry, Frascati, Italy, Frascati Physics series, Vol. VI, 1996, p. 841.
- [4] R. Arcidiacono, et al., Nucl. Instr. and Meth. 443 (2000) 26.
- [5] B. Hallgren, et al., Nucl. Instr. and Meth. A 419 (1998) 680.
- [6] B. Hallgren, et al., IEEE Trans. Nucl. Sci. 43 (1996) 1605.
- [7] MAX435, 250 MHz wideband transconductance amplifier with differential output, Maxim Corp, document 94NR 8-5.
- [8] VN552 User's manual, EPSI monitor board, C.A.E.N., 1995.
- [9] A. Lai, L. Musa, IEEE Trans. Nucl. Sci. 42 (1995) 812.
- [10] VIC 8251, VIC to VME Interface with mirrored memory, Creative Electronic Systems SA.
- [11] FVSBI 9210, FASTBUS to VSB Interface, Creative Electronic Systems SA.