



HAL
open science

The beam and detector for the NA48 neutral kaon CP violation experiment at CERN

V. Fanti, A. Lai, D. Marras, L. Musa, A. Nappi, R. Batley, A. Bevan, R.S. Dosanjh, R. Galik, T. Gershon, et al.

► **To cite this version:**

V. Fanti, A. Lai, D. Marras, L. Musa, A. Nappi, et al.. The beam and detector for the NA48 neutral kaon CP violation experiment at CERN. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 2007, 574, pp.433-471. 10.1016/j.nima.2007.01.178 . in2p3-00145549

HAL Id: in2p3-00145549

<https://hal.in2p3.fr/in2p3-00145549>

Submitted on 11 May 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Beam and Detector for the NA48 Neutral Kaon CP Violation Experiment at CERN

NA48 collaboration

V. Fanti, A. Lai, D. Marras, L. Musa ¹⁾, A. Nappi ²⁾

Dipartimento di Fisica dell'Università e Sezione INFN di Cagliari, I-09100 Cagliari, Italy.

R. Batley, A. Bevan ³⁾, R.S. Dosanjh, R. Galik, T. Gershon ⁴⁾, B. Hay, G.E. Kalmus,
S. Katvars, C. Lazzeroni, R. Moore ⁵⁾, D.J. Munday, M.D. Needham ⁶⁾, E. Olaiya ⁷⁾,
M.A. Parker, M. Patel ¹⁾, M. Slater, S. Takach, T.O. White, S.A. Wotton
Cavendish Laboratory, University of Cambridge, Cambridge, CB3 0HE, UK ⁸⁾

F. Bal, G. Barr ⁹⁾, G. Bocquet, J. Bremer, P. Brodier-Yourstone, P. Buchholz ¹⁰⁾, M.
Burns, A. Ceccucci, M. Clément, T. Cuhadar-Dönzelsmann ¹¹⁾, D. Cundy ¹²⁾, N. Doble
¹³⁾, V. Falaleev, F. Formenti, W. Funk, L. Gatignon, A. Gonidec, P. Grafström, B.
Hallgren, P. Kapusta, G. Kessler, W. Kubischta, W. Iwanski, A. Lacourt, G.
Laverriere, G. Linser, C. Ljuslin, A. Marchioro, M. Mast, J.P. Matheys, M. Morel, A.
Norton ¹⁴⁾, J.P. Orlic, B. Panzer-Steindel ¹⁵⁾, D. Schinzel ¹⁶⁾, W. Seidl, H. Taureg, J.-C.
Tarle, M. Velasco ¹⁷⁾, O. Vossnack, H. Wahl ¹⁴⁾, P. Wertelaers, J. Weterings
CERN, 1211 Genève 23, Switzerland

C. Cheshkov ¹⁾, A. Gaponenko ⁵⁾, E. Goudzovski, P. Hristov ¹⁾, A. Kalinin,
V. Kekelidze, Y. Kozhevnikov ¹⁾, D. Madigozhin, N. Molokanova, Yu. Potrebenikov,
A. Tkatchev, A. Zinchenko

Joint Institute for Nuclear Research, Dubna, Russian Federation

O. Boyle, I. Knowles, V. Martin, H. Parsons, K.J. Peach ¹⁸⁾, R. Sacco ¹⁹⁾, E. Veitch,
A. Walker

Department of Physics and Astronomy, University of Edinburgh, Edinburgh, EH9 3JZ, UK ⁸⁾

V. Carassiti, M. Contalbrigo, A. Cotta Ramusino, P. Dalpiaz, C. Damiani, J. Duclos,
P. Ferretti, P.L. Frabetti ²⁰⁾, A. Gianoli, M. Martini, F. Petrucci, M. Porcu, F. Rossi,
M. Savrie', M. Scarpa, C. Simani

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

A. Bizzeti ²¹⁾, M. Calvetti, G. Collazuol ¹³⁾, G. Graziani, E. Iacopini, M. Lenti,
F. Martelli ²²⁾, A. Michetti, G. Ruggiero, M. Veltri ²²⁾

Dipartimento di Fisica dell'Università e Sezione dell'INFN di Firenze, I-50125 Firenze, Italy

H.G. Becker, M. Behler, H. Blümer ²³⁾, D. Coward ²⁴⁾, C. Ebersberger, K. Eppard,
M. Eppard ¹⁾, H. Fox ²⁵⁾, K.H. Geib, A. Hirstius ¹⁵⁾, A. Kalter, K. Kleinknecht, U. Koch,
L. Köpke, P. Lopes da Silva, S. Luitz ²⁴⁾, P. Marouelli, L. Masetti, I. Melzer-Pellmann
²⁶⁾, U. Moosbrugger, C. Morales Morales, A. Peters ¹⁵⁾, B. Renk, J. Scheidt ²⁷⁾,
J. Schmidt, S.A. Schmidt, V. Schönharting, Y. Schué, J. Staeck, R. Wanke, R. Wilhelm,
A. Winhart, M. Wittgen ²⁴⁾, O. Zeitnitz

Institut für Physik, Universität Mainz, D-55099 Mainz, Germany ²⁸⁾

A. Dabrowski, T.M. Fonseca Martin ¹⁾

*Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208-3112,
USA*

J.C. Chollet, S. Crépe ²⁹⁾, C. De La Taille, L. Fayard, L. Iconomidou-Fayard,
G. Martin-Chassard, J. Ocariz ³⁰⁾, G. Unal, I. Wingerter-Seez

*Laboratoire de l'Accélérateur Lineaire, IN2P3-CNRS, Université de Paris-Sud, F-91898 Orsay,
France* ³¹⁾

G. Anzivino, F. Bordacchini, P. Cenci, E. Imbergamo, P. Lariccia, P. Lubrano, A.
Mestvirishvili ³²⁾, A. Papi, M. Pepe, M. Piccini, M. Punturo, C. Talamonti ³³⁾, F.
Tondini

Dipartimento di Fisica dell'Università e Sezione dell'INFN di Perugia, I-06100 Perugia, Italy

L. Bertanza, P. Calafiura ³⁴⁾, R. Carosi, R. Casali, C. Cerri, M. Cirilli ³⁵⁾, F. Costantini,
R. Fantechi, F. Fidecaro, L. Fiorini ³⁶⁾, S. Giudici, B. Gorini ¹⁾, F. Laico, G. Lamanna,
I. Mannelli, V. Marzulli, D. Passuello, G.M. Pierazzini, F. Raffaelli, M. Sozzi,
R. Tripicciono ¹⁴⁾

*Dipartimento di Fisica dell'Università, Scuola Normale Superiore e Sezione INFN di Pisa,
I-56127 Pisa, Italy*

S. Anvar, D. Béderède, F. Bugeon, J.B. Chèze, J. Cogan ³⁷⁾, M. De Beer, P. Debu,
D. Durand, S. Edard, J.L. Fallou, A. Formica, L. Gosset, R. Granier de Cassagnac,
J. Heitzmann, H. Le Provost, F. Louis, I. Mandzhavidze, E. Mazzucato, A. Migliori,
M. Mur, B. Peyaud, S. Schanne, O. Steinkamp, G. Tarte, R. Turlay, B. Vallage
DSM/DAPNIA - CEA Saclay, F-91191 Gif-sur-Yvette, France

M. Holder, I. Augustin, M. Bender, A. Maier ¹⁵⁾, I. Schwarz, M. Ziolkowski
Fachbereich Physik, Universität Siegen, D-57068 Siegen, Germany ³⁸⁾

R. Arcidiacono, P.L. Barberis, F. Benotto, F. Bertolino, C. Biino, O. Brunasso,
N. Cartiglia, M. Clemencic ¹⁾, D. Dattola, S. Goy Lopez, G. Govi ¹⁵⁾, R. Guida ³⁹⁾,
F. Marchetto, E. Menichetti, S. Palestini ¹⁾, N. Pastrone

*Dipartimento di Fisica Sperimentale dell'Università e Sezione dell'INFN di Torino, I-10125
Torino, Italy*

A. Chlopik, Z. Guzik, J. Nassalski, E. Rondio, M. Szleper ¹⁷⁾, W. Wislicki, S. Wronka ⁴⁰⁾
*A. Soltan Institute for Nuclear Studies, Laboratory for High Energy Physics, PL-00-681
Warsaw, Poland* ⁴¹⁾

H. Dibon, G. Fischer, M. Jeitler, M. Markytan, I. Mikulec, G. Neuhofer, M. Pernicka, A.
Taurok, L. Widhalm
*Österreichische Akademie der Wissenschaften, Institut für Hochenergiephysik, A-10560 Wien,
Austria* ⁴²⁾

submitted to Nuclear Instruments and Methods A

Abstract

The Beam and Detector, used for the NA48 experiment, devoted to the measurement of $Re(\epsilon'/\epsilon)$, and for the NA48/1 experiment on rare K_S and neutral hyperon decays, are described.

-
- 1) Present address: PH Division, CERN, 1211 Geneve 23, Switzerland.
 - 2) Dipartimento di Fisica dell'Università e Sezione dell'INFN di Perugia, I-06100 Perugia, Italy
 - 3) Present address: Oliver Lodge Laboratory, University of Liverpool, Liverpool, L69 7ZE, UK.
 - 4) Present address: High Energy Accelerator Research Organization (KEK), Tsukuba, Ibaraki, 305-0801, Japan.
 - 5) Present address: University of Alberta, Edmonton, Alberta T6G 2J1, Canada
 - 6) Present address: NIKHEF, PO Box 41882, 1009 DB Amsterdam, The Netherlands.
 - 7) Present address: CCLRC Rutherford Appleton Laboratory, Chilton, Didcot, Oxon, UK OX11 0QX
 - 8) Funded by U.K. Particle Physics and Astronomy Research Council
 - 9) Present address: Particle and Nuclear Physics, University of Oxford, Oxford, OX1 3RM, UK
 - 10) Present address: Fachbereich Physik, Universität Siegen, D-57068 Siegen, Germany
 - 11) Present address: University of British Columbia, Vancouver, BC Canada V6T 1Z4
 - 12) Present address: Dipartimento di Fisica Sperimentale dell'Università e Sezione dell'INFN di Torino, I-10125 Torino, Italy
 - 13) Present address: Dipartimento di Fisica dell' Università' di Pisa, I-56127 Pisa, Italy
 - 14) Present address: Dipartimento di Fisica dell'Università e Sezione dell'INFN di Ferrara, I-44100 Ferrara, Italy
 - 15) Present address: IT Division, CERN, 1211 Geneve 23, Switzerland
 - 16) Present address: ETH, 8093 Zurich, Switzerland
 - 17) Present address: Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208-3112, USA
 - 18) Present address: John Adams Institute for Accelerator Science, University of Oxford and Royal Holloway University of London
 - 19) Present address: Queen Mary, University of London, Mile End Road, London E14NS, UK
 - 20) Present address: Joint Institute for Nuclear Research, Dubna, Russian Federation
 - 21) Dipartimento di Fisica dell'Università di Modena e Reggio Emilia, I-41100, Modena, Italy
 - 22) Dipartimento di Fisica dell'Università di Urbino
 - 23) Present address: Universität Karlsruhe and Forschungszentrum Karlsruhe, D-76021 Karlsruhe, Germany
 - 24) Present address: SLAC, Stanford University, Menlo Park, CA 94025, USA.
 - 25) Present address: Physikalisches Institut, Universität Freiburg, D-79014 Freiburg, Germany
 - 26) Present address: DESY, D-22607 Hamburg, Germany
 - 27) Present address: Fachhochschule Hof, D-95028 Hof, Germany
 - 28) Funded by the German Federal Minister for Research and Technology (BMBF) under contract 7MZ18P(4)-TP2
 - 29) Present address: Laboratoire de Physique Subatomique et de Cosmologie, CNRS-IN2P3-Université J. Fourier, F-38026 Grenoble, France
 - 30) Present address: Laboratoire de Physique Nucleaire et de Hautes Energies, IN2P3-CNRS-Universités Paris VI et Paris VII, F-75252 Paris, France
 - 31) Funded by Institut National de Physique des Particules et de Physique Nucléaire (IN2P3), France
 - 32) Present address: High Energy Physics Department, University of Iowa, Iowa City, IA 52242, USA
 - 33) Present address: Dipartimento di Fisiopatologia Clinica, Università degli Studi di Firenze, I-50134, Firenze, Italy
 - 34) Present address: Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
 - 35) Present address: University of Michigan, Ann Arbor, MI 48019, USA
 - 36) Present address: Institut de Fisica d' Altes Energies, Facultat Ciències, Universitat Autònoma de Barcelona, 08193 Bellaterra, Spain

1 Introduction

The NA48 experiment [1] has been designed with the principal goal of detecting and quantifying the presence of direct CP violation in the neutral kaon decays into two pions, as expressed by the parameter $Re(\epsilon'/\epsilon)$, which is related to the double ratio of the four decay rates:

$$R = \frac{\Gamma(K_L \rightarrow \pi^0\pi^0)\Gamma(K_S \rightarrow \pi^+\pi^-)}{\Gamma(K_L \rightarrow \pi^+\pi^-)\Gamma(K_S \rightarrow \pi^0\pi^0)} \approx 1 - 6Re(\frac{\epsilon'}{\epsilon}) \quad (1)$$

In order to minimize systematic differences, the experiment should ideally be able to compare directly K_S and K_L decays of equal momenta occurring simultaneously with the same geometry relative to the detector elements and in the absence of significant background [2, 3, 4].

The relevant features of the experimental setup adopted are described in the following chapters referring to the different components. Here the basic elements and the resulting requirements for the experiment are outlined.

The basic elements were (Fig. 1):

- The use of a high-energy proton beam from the CERN SPS to produce a neutral beam from a target far upstream of the detectors, which was rich in K_L in the range of momenta $70 \text{ GeV}/c < P_K < 170 \text{ GeV}/c$. A lower intensity proton beam, selected by 'channelling' the protons not interacting in the K_L target in a bent silicon crystal, was transported to a second target, close to the beginning of the fiducial decay region. This second target was laterally displaced relative to the axis of the K_L beam and the K_S beam produced from it was aimed to overlap with the K_L at the longitudinal position of the main detectors. The K_S intensity was adjusted to produce a number of two-pion decays in the fiducial decay region somewhat larger than those from K_L . A time coincidence between an observed decay and a signal in a tagging counter, crossed by the low intensity proton beam on its way to the K_S target, was used to identify the decay as coming from a K_S rather than a K_L (Section 6). The absence of a coincidence with the signal from a plastic scintillator counter preceded by a photon converter, which was placed immediately after the end of the K_S beam collimator, sharply defined the beginning of the fiducial region. The neutral beams and their decay products were first transported inside a large vacuum tank. This was closed off by a thin *Kevlar* window to separate the decay region from the detectors. A pipe ran from this window through the centre of all following detectors to ensure that the beam be transported in vacuum (Section 2).
- A tank filled with helium at atmospheric pressure housing a spectrometer, composed of a large-aperture, vertical-field dipole magnet sandwiched between two pairs of wire chambers. These provide measurements of the angles of the two charged pion

³⁷⁾ Present address: Centre de Physique des Particules de Marseille, F-13288 Marseille Cedex 9, France

³⁸⁾ Funded by the German Federal Minister for Research and Technology (BMBF) under contract 056SI74

³⁹⁾ Dipartimento Interateneo di Fisica e Sezione INFN di Bari, I-70126 Bari, Italy

⁴⁰⁾ A. Soltan Institute for Nuclear Studies, Laboratory for High Energy Physics, PL-05-400 Swierk-Otwock, Poland

⁴¹⁾ Supported by the Committee for Scientific Research, grants 5P03B10120, 2P03B11719 and SPUB-M/CERN/P03/DZ210/2000

⁴²⁾ Funded by the Austrian Ministry of Education, Science and Culture under contract GZ 616.360/2-IV GZ 616.363/2-VIII and by the Fund for Promotion of Scientific Research in Austria (FWF) under contract P08929-PHY

The NA48 Detector




Figure 1: The NA48 detector

tracks before and after the magnetic deflection. The window between the vacuum decay tank and the spectrometer was made of *Kevlar* composite only 0.003 radiation lengths thick. To provide a safety enclosure, the helium tank was closed off by an aluminum window located just before the scintillation hodoscope (Section 3).

- A double walled, suitably segmented, plastic scintillator hodoscope for triggering on charged pion decays (Section 5.4).
- A calorimeter with high granularity and energy resolution, used to identify and measure the individual energies and transverse coordinates of multiple photons from $\pi^0 \rightarrow \gamma\gamma$ decays. The information concerning the gamma showers was needed to trigger on neutral pion decays and hence had to be available in real time (Section 4).
- A hadron calorimeter used at the trigger level to provide information for hadron identification (Section 5.1).
- A muon veto, with thick iron walls and scintillation counter hodoscopes, for muon identification (Section 5.7).
- Rings of anti-counters to help veto events with charged particles and/or photons

outside the aperture of the wire chambers and calorimeters (Section 5.3).

- Synchronous readout, digitization and pipelining of the information from all detectors to a central data-acquisition and analysis system. The event building and recording was carried out asynchronously (Section 7 and 8).

Specific requirements for each detector were identified at an early stage in the design of the experiment, often resulting in challenging specifications.

All the detectors need to stand high beam intensity, in order to allow the collection, in periods of the order of 100 days of running time, of more than one million $K_L \rightarrow \pi^0\pi^0$ decays, the channel having the lowest branching ratio of the four decay modes recorded. The time resolution between signals had to be a few nanoseconds for the trigger and 200-300 ps off-line to reject accidental background.

Weights were applied to the reconstructed K_L events, equal to the survival probability of K_S of the same momentum in the interval from the beginning of the decay region to the decay point. The weighting procedure, which provides an equalization of the distribution of K_S and K_L decays, and hence of the acceptance, allows the double ratio of the four rates to be evaluated simply in bins of kaon momentum, after integrating over longitudinal vertex position. This all-important gain in systematic uncertainty was obtained at a loss in statistical power of $\approx 35\%$.

MonteCarlo simulation showed that, to achieve a systematic accuracy of better than 10^{-4} in $Re(\epsilon'/\epsilon)$, the momentum scale for decays into pairs of charged and neutral pions had to be fixed to better than 10^{-3} and similarly for the length of the fiducial region. Exploiting the known values of the masses of the neutral kaon and of the charged and the neutral pion, this translated into the need to know the ratio of the transverse dimensions of the detectors relative to their longitudinal positions with an accuracy of about 10^{-4} . The choice was made to use a quasi-homogeneous liquified noble gas calorimeter with full containment (27 radiation lengths) of high energy electromagnetic showers, whereby the best affordable option was liquid krypton. To exploit fully the potential stability inherent in a vacuum-insulated cryogenically-controlled device, it was decided to immerse the high voltage coupling capacitors, the shaping amplifiers and calibration pulsers directly in the liquid krypton at the back of the calorimeter. A tower structure of electrodes was chosen to detect the initial current induced by the drift of the ionization, with cells of transverse dimensions of about 2×2 cm², covering 5.3 m² of the front face of the calorimeter. A solution for the structure had to be found to satisfy the stringent requirements of the experiment. In the end an array composed of thin Cu-Be ribbons kept under tension between precisely machined spacers and pointing to the centre of the fiducial region was designed and built with better than ± 0.05 mm local tolerance and ± 0.1 mm large scale tolerances.

2 The simultaneous K_L and K_S beams

2.1 Layout

A branch of the 450 GeV/c primary proton beam, slow-extracted from the CERN Super-Proton-Synchrotron (SPS) towards the North Area, was used to produce the neutral kaon beams. Due to the different mean decay-lengths of K_L and K_S , the beams were derived from protons striking two separate targets, respectively situated 126 m and 6 m upstream of the beginning of the common decay region. The longer length of the K_L beam was chosen to ensure that the K_S component in that beam had decayed to a proportion of $\sim 1 \cdot 10^{-6}$ at the highest momentum considered (170 GeV/c).

The layout of the simultaneous K_L and K_S beams is shown schematically, in vertical

Figure 2: Schematic layout of $K_L + K_S$ beams.

2.2 Primary proton transport to the K_L target

The 450 GeV/c primary proton beam, emerging from an upstream target station (T4), passed through a pair of dump/collimators to select the wanted flux. This beam was transported over a distance of 838 m and was finally focused and directed vertically downwards at an angle of 2.4 mr onto the K_L target. In particular, the optics incorporated an intermediate focus at a point, where a steering magnet (Trim 5) was placed, which allowed the angle of the beam onto the target to be adjusted horizontally (see Section 2.6). The choice of 2.4 mr, compared to zero production angle, reduced the total flux of neutrons in the experimental hall per useful K_L by a factor ~ 4 (to $\sim 10 : 1$), for only $\sim 25\%$ reduction in the useful K_L per proton.

2.3 The K_L target station

The detailed layout of the K_L target station is shown schematically, in vertical section, in Figure 3. The K_L target consisted of a 2 mm diameter, 400 mm long, rod of beryllium, made in four sections, each suspended between thin aluminium foils. It was precisely aligned along the K_L beam axis and was immediately followed by a water-cooled, copper collimator of 15 mm aperture, through which the neutral (K_L) and the remaining primary proton beams exited. A 10.8 T·m, vertically-deflecting, dipole magnet

¹⁾ Whereas for the runs from 1997 to 1999 the primary proton momentum was 450 GeV/c, it was changed to 400 GeV/c with a longer duty cycle of 4.8 s/16.8 s for the 2001 $Re(\epsilon'/\epsilon)$ run as well as for the high intensity K_S (HIKS) run in 2002. This entailed a factor of 0.81 in the yield of K_L and K_S per proton. The proton fluxes per pulse onto the targets was increased by a factor 1.5, resulting in an overall factor of ~ 0.63 in the instantaneous K_L and K_S rates per second of spill. This decrease in instantaneous rates served to demonstrate the absence of rate-dependencies in the detectors.

Table 1: Design Characteristics of the Simultaneous $K_L + K_S$ and High Intensity K_S Beams

Beam	$K_L + K_S$		HIKS
	K_L	K_S	
Primary protons per pulse on target	$1.5 \cdot 10^{12}$	$3 \cdot 10^7$	$5 \cdot 10^{10}$
SPS spill length(s)/Cycle time(s)	2.5/14.4		4.8/16.8
Proton momentum, p_0 (GeV/c)	450	450	400
Production angle of K^0 beam (mr)	2.4	4.2	4.2
Length of K^0 beam:			
Target to final collimator/AKS (m)	126.00	6.00/6.09	6.23
Target to front of e.m.calorimeter (m)	241.10	121.10	121.33
Angle of convergence to K_L (mr)	0.0	-0.6	-
Angular acceptance of beam (mr)	± 0.15	± 0.375	± 0.36
R.M.S. radius at e.m.calorimeter (mm)	~ 20	~ 30	~ 30
Useful momentum range, p_K (GeV/c)	70 - 170		
Fiducial length for decays, $\Delta z = 3.5\lambda_S$ (m)	~ 13	- ~ 32	-
K^0 flux per pulse at exit final coll.	$\sim 2 \cdot 10^7$	$\sim 2 \cdot 10^2$	$\sim 1.8 \cdot 10^5$ *
- decays between coll. and detector	$\sim 1.4 \cdot 10^6$	$\sim 2 \cdot 10^2$	$\sim 1.8 \cdot 10^5$
K^0 flux per pulse in useful p_K range	$6.4 \cdot 10^6$	$1.5 \cdot 10^2$	$1.4 \cdot 10^5$
- K^0 decays in p_K range and fiducial length			
- per pulse	$3.8 \cdot 10^4$	$1.5 \cdot 10^2$	$\sim 1.4 \cdot 10^5$
- per year	$1.1 \cdot 10^{10}$	$4.5 \cdot 10^7$	$3.5 \cdot 10^{10}$
-decays per pulse to $\pi^0\pi^0$ in p_K range and fiducial length	35	45	
Detector acceptance for $\pi^0\pi^0$ decays	~ 0.2		
Useful $K^0 \rightarrow \pi^0\pi^0$			
- detected per pulse	~ 7	~ 9	
- per year (100 days with overall eff. 0.5).	$2.1 \cdot 10^6$	$2.7 \cdot 10^6$	

* Taking into account a factor ~ 0.77 for absorption of K_S in the 24 mm long Pt photon absorber (see Section 2.9).

(B1 in Fig. 3) then served to sweep away charged particles, whereby the primary protons were deflected further downwards by 7.2 mr. At a distance of 10.95 m from the target, they impinged on a bent silicon crystal, centred 72 mm below the K_L beam axis. The crystal was designed to split off a wanted, small fraction of the protons by deflecting them through an upward angle of 9.6 mr, back to the horizontal. The majority of the protons and other particles, not channeled by the crystal, continued undeviated. With the exception of muons, they were absorbed in the water-cooled copper blocks of the following pair of beam-dump/collimators, which were fitted with tungsten-lined passages for each of the two wanted (K_L and proton to K_S) beams. The aperture for the K_L beam in the second beam-dump/collimator (8 mm diameter) limited the neutral beam to ~ 3 times the flux that was finally accepted.

2.4 K_L beam collimation

The neutral (K_L) beam was subsequently collimated in three successive stages by 'defining', 'cleaning' and 'final' collimators. The final collimator was located between




Figure 3: Layout of the K_L target station and of the proton beam trajectory via the bent crystal 'splitter' to the K_S - tagging station (vertical section). The dashed line shows the trajectory followed by the beam used to produce a High Intensity K_S beam from the K_S target (Section 2.9).

120.7 and 124.3 m from the target and had an aperture (varying from 54 to 57.5 mm in diameter), that fitted within the gap of the last sweeping magnet. It was preceded by the cleaning collimator, of such aperture as to prevent particles produced or scattered on the edges of the defining collimator from striking the final collimator. It was located ~ 20 m, upstream of the final collimator, to let K_S , which may be regenerated on its edges, decay away before reaching the fiducial volume. The condition that particles from the target should not strike the cleaning collimator directly, led to an optimum longitudinal position (~ 41 m from the target) and aperture (12.2 mm diameter) for the defining collimator that maximised the acceptance of the neutral beam.

The neutral beam was transported in vacuum over practically its full length. The only exception occurred at the location of the crystal and dump/collimators belonging to the target station, where the vacuum was interrupted by thin aluminium windows separated by a 4.6 m long passage through air (Fig. 3). The defining and cleaning collimators each consisted of a 1.2 m long steel block containing a cylindrical bore, fitted with a series of tungsten-alloy inserts of graduated aperture.

2.5 Muon background

A high flux of muons from the decay of charged pions (and kaons) produced in the K_L target and beam dump would naturally reach the detectors. The first magnet after the target (B1 in Fig. 3) deflected positive particles downwards and negative particles upwards. It was vitally important that this action was not cancelled by the following,

upward deflection needed to recuperate the protons. This deflection was provided by the highly-selective, bent crystal, in place of a magnet.

Downstream of the beam dump, use was made of the same-sign fields in the top yoke and in the aperture of successive vertical H- and C-shaped magnets belonging to the proton transport (B2 and B3 in Fig. 3), to increase the upward deflection of μ^- . The fields in their respective lower yokes were opposite and had little net effect on the μ^+ .

A 5 m long, toroidally-magnetised, iron 'scraper', surrounding the beam, was then designed to defocus any remaining μ^- , whilst focusing the μ^+ towards a series of three, 2 m long, H-shaped magnets of large cross-section. The space between their poles was filled with slabs of soft-iron, through which a 40 mm diameter hole had been bored to let the proton beam pass almost undeflected. The surrounding muons, however, encountered a mainly vertically-directed field integral of ~ 9 T·m and were hence deflected by an angle in the horizontal plane, which exceeded their r.m.s. multiple scattering in the iron by a factor ~ 10 and was sufficient to clear the downstream detectors.

2.6 The bent crystal

When charged particles enter a crystal within a small ('critical') angle to its planes, the collective Coulomb fields can cause them to be 'channeled'. If the crystal is mechanically bent, some of the particles still follow the planes and are thus deflected. The present application requires only a small fraction of the protons emerging from the K_L target to be deflected. The solution adopted has been described in detail [6] and is shown schematically in Fig. 4.

A mono-crystal of silicon, cut to dimensions of $[60 \cdot 18 \cdot 1.5]$ mm³, parallel to the (110) plane, was used. It was bent through an angle, $\theta_0 = 18.7$ mr over 56 mm of its length. The crystal holder was in turn mounted on a motorised goniometer, which allowed the crystal to be aligned on the incident beam with two transverse displacements and two angular rotations about axes perpendicular to the beam. When the crystal was rotated through an azimuthal angle, ϕ , by $\pm 28.7^\circ$ (about the vertical axis), the beam impinged on one side edge, traversed the crystal diagonally and exited from the opposite edge. The effective, fractional length of the bent crystal traversed, and hence the vertical angle θ of the beam deflection was then governed by adjusting the angle ϕ .

The experiment nominally required a fraction of $2 \cdot 10^{-5}$ of the protons incident on the K_L target to be deflected and transmitted to produce the K_S beam. The transmission could be regulated by increasing the width of the beam incident on the crystal. In fact, the focusing of the proton beam was arranged to give a large magnification in horizontal beam divergence at the K_L target. Another way to correct the transmission consisted of changing the horizontal angle of the incident proton beam about the centre of the K_L target, by means of an upstream steering magnet (Trim 5). This reduced the local intensity of protons that struck the crystal over the small width ($\sim \pm 0.2$ mm), within which the conditions for channeling were fulfilled.

Instantaneous rates related to K_S and K_L were recorded by the AKS counters located at the exit of the K_S beam collimator (Section 5.2 and Fig. 17 therein), a counter telescope (KSM) viewing the K_S target and by a counter array (BCTR) predominantly sensitive to the K_L beam at the monitor station at the downstream end of the beam line. The counting rates were displayed each pulse and were used when necessary to adjust the current in the upstream proton steering magnet (Trim 5), to render the K_S / K_L ratio nearly constant.

BENT CRYSTAL:

Si cut // to (110) planes
[60 mm x 18 mm x 1.5 mm]




Figure 4: Schematic arrangement of the bent crystal

2.7 Proton transport to the K_S target

The second of the beam dump/collimators downstream of the crystal (TAX 18) contained a set of tungsten-alloy inserts with graduated apertures of smallest diameter 2.4 mm, aligned on an axis 72 mm below and parallel to that of the K_L beam (Fig. 3). This passage was matched to the proton beam defined by the crystal. There followed the H- and C-shaped magnets of opposite polarity (B2 and B3), which successively deflected the transmitted protons back towards and then onto the K_L beam axis. In between these, the protons passed through the tagging station (see Section 6), which was installed in a space 36 mm below the K_L axis and was preceded by a protecting collimator, fitted with separate passages for the proton and K_L beams.

Once on the K_L beam axis, the protons were steered through the defining and cleaning collimators and were refocused by a series of quadrupole magnets to a point 109 m downstream of the crystal. Before reaching this focus, they were again bent upwards by additional dipole magnets to a point 72 mm above the K_L beam at an upward angle of 3.6 mr. Here they impinged on the K_S target, which had dimensions similar to those of the K_L target and was mounted in vacuum.

2.8 The K_S beam

The layout of the K_S beam is shown schematically in Figure 5. The vertical displacement of the target with respect to the K_L beam (72 mm) was chosen to ensure that the

Figure 5: Layout of the K_S target station and beam (vertical section).

The K_S target and the following collimator apertures were aligned along an axis pointing downwards at 0.6 mr, so as to intersect the (horizontal) K_L beam axis near the centre of the detectors, at a longitudinal distance from the K_S target of 120 m. This alignment was achieved by adjusting the transverse position of the target and steering the proton beam onto it, observing the centre of the neutral beam at the monitor station ~ 14 m downstream of the cross-over point. There the K_S axis should coincide horizontally with that of the K_L and be ~ 8.4 mm vertically below it. The resulting production angle of 4.2 mr for the K_S beam rendered the ratio of K_S to K_L decays in the fiducial region approximately equal at the two ends of the momentum range used.

The K_S target was followed by a 7.5 T·m, vertical sweeping magnet (B7), the gap of which was filled with tungsten-alloy inserts containing the passages for the two beams (Fig. 5). The passage for the K_S beam was shaped to absorb the remaining protons and charged secondary particles from the target, which were deviated by the field. The magnet was followed by a steel collimator block, which again contained precisely-bored passages for the two neutral beams, the one for the K_S being filled with inserts, graduated from a beam-defining aperture of 3.6 mm diameter at 4.8 m from the target to a final diameter of 6.0 mm at the exit, 6.0 m from the target. At this point, both K_L and K_S beams emerged into the common decay volume, pointing towards the detectors.

2.9 The High Intensity K_S beam

Subsequent to the NA48 measurement of $Re(\epsilon'/\epsilon)$, the beam line was adapted to provide a high-intensity K_S (HIKS) beam for a high-sensitivity investigation of K_S and neutral hyperon decays (experiment NA48/1) [7].

In the absence of interactions on the upstream (K_L) target and proton beam dump and without products from decays of the K_L beam passing through the detectors, it was

possible to increase the intensity of the K_S beam by a large factor ($\sim 10^3$) compared to that of the K_S component of the simultaneous beams.

The primary proton beam leading to the K_L target, at this time 400 GeV/c momentum, was attenuated and subsequently collimated to the required flux of $\sim 5 \cdot 10^{10}$ protons per SPS pulse in the upstream section of the beam transport line. The protons then travelled without further loss through the K_L target station with the target removed and by-passed the bent crystal, on a trajectory shown by the dashed line in Fig. 3. They were finally focused and directed onto the K_S production target at the desired angle with respect to the following K_S beam axis. Compared to the standard layout shown in Fig. 5, the following modifications were made to allow a high-intensity beam of K_S and neutral hyperons to be produced:

- The K_S target was displaced upstream by 0.23 m to make space for a photon absorber in the beam. This consisted of a choice of remotely-controlled and finely-adjustable platinum rods, 2.6 mm in diameter and, alternatively, 24 mm ($8.0 X_0$), 32 mm ($10.7 X_0$) or 40 mm ($13.4 X_0$) in length, which could be inserted into the HIKS beam line at a point 0.6 m after the target centre. Here, due to the chosen production angle of 4.2 mr, the primary proton beam had an upward displacement of 2.5 mm with respect to the K_S beam axis and hence missed the absorber.
- Below the HIKS beam line, the K_L beam passage through the final collimator traversing the sweeping magnet was blocked by a 1.5 m long cylindrical plug of bronze to absorb particles escaping in that direction.
- Due to the small diameter (6.0 mm at the exit) of the last part of the K_S collimator, K_S decays that were detected could not originate as far upstream in the collimator as was the case for K_L . Hence the length of the collimator following the magnet (over which it was required that the directions of charged decay products should not be disturbed by magnetic field) could be reduced to 0.6 m. A small, vertically-deflecting, sweeping magnet could therefore be added in place of the first part of the collimator block. The magnet was arranged so that the defining section of the K_S collimator was located inside its gap. Its field served to sweep away low-energy, charged particles (in particular e^+e^- pairs) created in this section.
- The AKS converter and counter system (see Section 5.2) was dismantled from the beam. It was replaced by a tube to allow the HIKS beam to travel in vacuum throughout its length from the target to the downstream end of the detectors.

The design characteristics of this High Intensity K_S beam are listed in the column headed HIKS of Table 1. A factor of $\sim 10^3$ could be gained in its intensity compared to that of the K_S component of the simultaneous $K_L + K_S$ beams, whilst maintaining comparable total counting rates in the principal detectors.

2.10 Vacuum system

Downstream of the last collimator (and the AKS counters, shown in Fig. 5), the vacuum chamber containing the two beams opened up into a cylindrical tank of total length 89 m and inner diameter 1.92 m increasing in two steps up to 2.40 m at the end of the decay region. The volume was evacuated to a residual pressure $< 10^{-4}$ mbar.

At its downstream end, the tank was closed off by a thin 'window' to separate the vacuum from the following volume, which was filled with helium gas, housing the drift chambers of the magnetic spectrometer. The window was made of a 3-layer composite of epoxy-impregnated *Kevlar* tissue, draped and cured in a mould of spherical shape with 1.3 m radius of curvature (Fig. 6). It was glued at its centre to an aluminium ring of

outer diameter 184 mm and at its circumference to a steel flange of inner diameter 2.30 m. Over the area between these two flanges, the mean thickness of the window had been measured to be 0.9 mm, corresponding to $\sim 3.1 \cdot 10^{-3}$ of a radiation length. The total leak rate through the window for helium gas was found to be < 1 mbar·l/s. At the end of the neutral beam line (254 m from the K_L target), the vacuum tube was closed off by a 0.1 mm thick *Mylar* window just in front of the monitor station.

2.11 Beams instrumentation and control

The primary proton beam line to the K_L target was equipped with secondary emission monitors (SEM) consisting of foils of aluminium housed in vacuum, on which the charge induced by the beam could be measured. The total flux of protons was measured on single foils (BSI) covering the whole beam. In other types of SEM the foil was split in half (BSP) or formed into a filament (BBS) that could be scanned across the beam, thereby allowing either the asymmetry in position or the profile of the beam to be measured upstream of the target. Downstream of the bent crystal, the flux of the transmitted proton beam was monitored on a pair of scintillation counters at the tagging station. A final beam monitor station was placed at the downstream end of the experiment just before the beam dump. In order along the beam, this station comprised a pair of filament scintillator (FISC) counters, a multi-wire-proportional chamber, a scintillation counter and a fibre beam monitor (BCTR)²⁾ which were used to adjust the K_L and K_S beams in turn and, with the simultaneous beams, to record essentially the K_L beam component.

During the data taking, two counter-based monitors were used for the beam intensity control: the K_S beam monitor (KSM) placed alongside the K_S target³⁾, and the K_L monitor (BCTR) just before the beam dump.

Additional monitors (made of *Plexiglas* as the active medium and read by a 10-stage photomultiplier) were installed for the 2001 run. They were housed between the standard K_L monitor and the beam dump. The main advantage of the new monitors was the small time scale on which the instantaneous beam intensity could be studied. While for the traditional monitors the intensity was computed from the difference in counts between two consecutive triggers, leading to a time scale $\approx 170\mu\text{s}$, the additional monitors had a set of four integration times that were chosen to be 200 ns, 1 μs , 3 μs and 15 μs .

3 The magnetic spectrometer

The first major detector following the *Kevlar* window was the magnetic spectrometer, consisting of four multiwire Drift Chambers (DCH) and a dipole magnet.

3.1 Helium tank

The spectrometer was housed in a steel cylinder, about 23 m long and 2.8 m in internal diameter, filled with helium at atmospheric pressure and divided into five sections by the four DCHs. As the central section traversed the spectrometer magnet, this and the two adjacent sections were made of stainless steel. The downstream end cap was made of aluminum; the upstream section, closed-off by the *Kevlar* window, was of carbon steel. Along the axis of the tank, throughout its length, a thin-walled tube of 152 mm inside diameter allowed to transport the beam in vacuum (Fig. 1).⁴⁾ The tank volume was about 120 m³.

²⁾ Two orthogonal planes of scintillating fibers with a sensitive area of 18×18 cm²

³⁾ Four plastic scintillators parallel to the target. The K_S beam intensity was monitored counting particles emitted sideways from the K_S target.

For safety reasons due to the neighbouring large-volume vacuum system, the tank had been designed as a vacuum vessel. For this reason it was possible, in the year 2000, to run with no drift chambers and no *Kevlar* window, with the tank extending the vacuum volume. This run served to record only kaon decays with neutral particles in the final state.

The need to minimize multiple scattering drove the choice of gas and of the thin *Mylar* foils used for the DCH windows.

Contamination of the DCH gas mixture (Ar/ethane) by helium was prevented by covering both sides of each chamber with a second *Mylar* foil and filling the space between them with argon. The inner DCH volume was maintained at 1.0 mbar above atmospheric pressure by proper adjustment of the oil level in output bubblers. The argon between the foils was similarly maintained at an overpressure of 1.3 mbar, so as to set a pressure gradient moving outwards from the chamber volume.

The pressure regulation of the helium, which flowed in a closed circuit, was assured by an inflatable balloon located under the ceiling of the experimental hall, above the tank. It contained a maximum of 12 m³ of He, and was normally half-filled. A weight-counterweight system provided the required 0.3 mbar overpressure in the tank relative to the atmosphere and served as a 'lung' to absorb sudden barometric changes.

A further safety level was provided by two hydraulic guards, located in the experimental hall. They assured an air flow to the tank in case of underpressure and release of helium in case of overpressure. The underlying logic of these safety devices was that, even in the temporary absence of electric power, neither overpressure nor underpressure could reach values dangerous for the Drift Chambers.

Helium purity in excess of 95% was required in the tank during the experiment. This was obtained by filling the tank with pure helium at the beginning of data-taking and continuously purifying the gas during the run in a closed circuit. The purification process, performed by a dedicated device, was based on the smaller adsorption of helium at low temperature (196 °C, the boiling point of nitrogen) with respect to contaminant gases. The adsorptive elements consisted of zeolites.

During normal operation the gas contained in the tank represented a fraction of about $7.6 \cdot 10^{-3}$ of a radiation length.

Figure 6 shows a schematic view of the helium system, with the dimensions of each section of the tank.

3.2 Magnet

The iron yoke of the spectrometer magnet had an opening between its poles of 2.40 m vertically \times 3.20 m horizontally. The active fiducial aperture was determined by the helium tank which had an inner diameter, through the magnet, of 2.37 m. Two pairs of coils provided a magnetomotive force of $0.98 \cdot 10^6$ A-turns for a dissipated power of 3.1 MW[9, 10], giving a vertical field of 0.37 T in the center of the magnet.

The three components of the field (B_X, B_Y, B_Z) were measured using a set of Hall probes. Three maps were recorded for magnet currents corresponding to 100%, 87% and

⁴⁾ For the 1997 run this tube was made of stainless steel. Then, for the 1998-1999 runs, it was replaced by one made of a carbon-fibre composite, in order to diminish the generation of electromagnetic showers by decay products hitting the tube. At the end of 1999, for unknown reasons, the carbon-fibre tube imploded. This caused serious damage to the spectrometer drift chambers which had to be rebuilt. Following this event, the tube was replaced by one constructed of 1 mm thick aluminium, which was used in 2001 and for subsequent experiments.

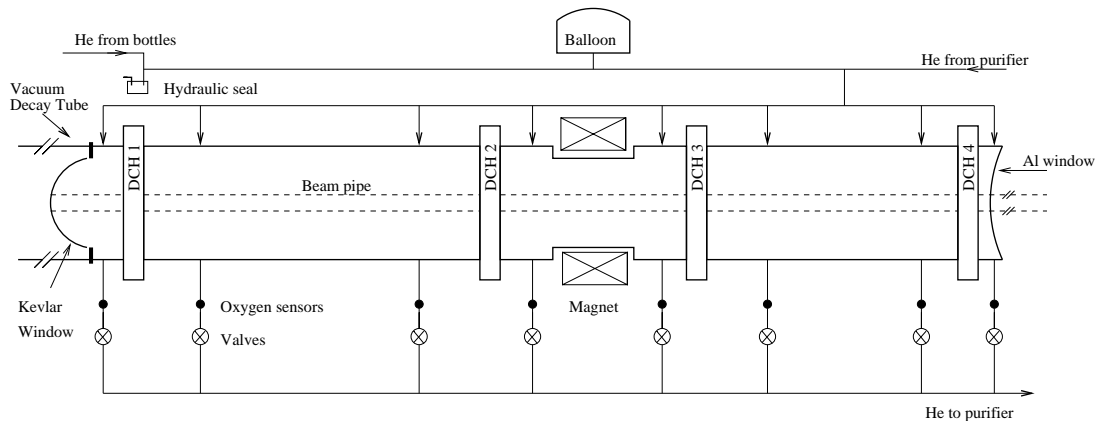


Figure 6: Schematic layout of the helium tank

42% of the nominal value. The precision on the relative field in a map was estimated to be 10^{-4} and limited to 0.3 G for low field values. An absolute calibration of the Hall probes was made using an NMR probe at the magnet centre. A three-dimensional calculation of the magnetic field was performed using a finite-element method to obtain solutions of the Poisson equation[11]. After applying a global correction of $6 \cdot 10^{-3}$, the computed and measured field values agreed within $3 \cdot 10^{-3}$ inside the fiducial volume. The field integral along the axis between the chambers DCH2 and DCH3 corresponded to a 257 MeV/c transverse momentum kick. The variation of the field integral over the aperture of the spectrometer did not differ more than $\pm 6\%$ from the value along the axis.

During data taking, the value of the magnet current and the responses of 2 Hall probes were recorded and the map used for on-line and off-line reconstruction was corrected for any variation of the field larger than $5 \cdot 10^{-4}$. The overall calibration of the spectrometer including the absolute value of the magnetic field and the chamber geometry was checked by the reconstruction of the kaon mass from the $K^0 \rightarrow \pi^+\pi^-$ decay.

3.3 Drift Chambers

The four drift chambers [12] used for the NA48 spectrometer had an octagonal shape with a transverse width of 2.9 m. Their fiducial area was about 4.5 m². In order to achieve high detection redundancy, each chamber contained 8 planes of grounded sense wires orientated in four different directions, orthogonal to the beam axis: 0° (X,X'), 90° (Y,Y'), -45° (U,U') and $+45^\circ$ (V,V'). Each so-defined view contained two staggered planes of wires to resolve left-right ambiguities. An additional feature of the chambers was a 160 mm diameter central hole for the beam pipe that required careful construction and positioning of an ensemble of rings on which the central wires of the different planes were soldered.

To minimize multiple scattering effects, the wire chambers were constructed with minimal amount of material along the beam direction. A schematic view of the geometry of one view is shown in Fig. 7. Sense wires had a diameter of $20 \mu\text{m}$ and were made of gold-plated tungsten. The electric field was created by applying a negative voltage on two planes of gold-plated Ti-Cu wires located on each side of the sense wire planes, at a distance of 3 mm. The field wires had a diameter of $120 \mu\text{m}$. The spacing between sense and field wires on a plane was 10 mm. Thin *Mylar* foils ($22 \mu\text{m}$) coated with graphite were used to shape the electric field in the drift cell and also to act as separating walls between




Figure 7: Drift cell geometry

two consecutive views.

All wires were stretched and soldered on thin (3.8 to 6.2 mm) *G10* frames using printed circuits glued on the frames. Wires were pre-tensioned on stiff steel frames before they were transferred to the final *G10* frames. The length of the wires varied between 1.28 m and 2.72 m. The nominal mechanical tension applied to the wires was 55 g for sense wires and 520 g for field wires. Such tensions ensured that field and sense wires hanging horizontally in the chamber had similar saggittas to within $30\ \mu\text{m}$. The precision obtained on the 1 cm gap between two wires was found to be better than $10\ \mu\text{m}$ (rms) and parallelism was achieved to better than $25\ \mu\text{m}/\text{m}$ (rms). The chamber frame support, on which the wire planes were mounted, consisted of two 12 cm thick octagonal frames made of stainless steel. Between these a total of 38 *G10* frames were stacked in a chamber. These included field and wire planes, *Mylar* foils and spacers. Precise positioning of the *G10* frames was ensured by 40 peripheral rods each of 32 mm diameter.

Two $50\ \mu\text{m}$ thick *Mylar* windows 2 cm apart in the beam direction were mounted on the entry and exit sides of the chamber to hold a 160 mm diameter aluminium ring used as a flange for the beam pipe. Inside the chamber, the *G10* rings which held the central wires and the separating mylar windows were tied together by four 5 mm diameter epoxy rods. The whole central structure was free to move inside the chamber since it was independent of the beam pipe. The gap between the internal rings and the beam pipe which traverses the chamber was about 2 mm. Given a total tension of about $3 \cdot 10^3\ \text{N}$ due to the central wires, the displacement of the internal structure was smaller than $100\ \mu\text{m}$.

The total number of wires per chamber was 6160 and the amount of material in one chamber, including the gas mixture, corresponded to about $4 \cdot 10^{-3}$ of a radiation length.

3.4 DCH readout

A readout time measurement accuracy of about 1 ns with no deadtime was achieved with specially designed amplifiers and TDC circuits and a readout architecture tailored to the low multiplicity events from neutral kaon decays [13]. The anode signals in the wire chambers were amplified, discriminated and transformed to logic pulses in electronic circuits located on the chambers. The output pulse was a logic 50 ns wide ECL pulse with

50 ns deadtime following the pulse, achieved by a delay element.

All anode wires in DCH1, DCH2 and DCH4 were instrumented with amplifier and TDC circuits while only the X and Y views of DCH3 were read out. The time of the wire hits was measured with respect to the experimental clock in specially developed TDC circuits. The 25 ns clock cycle was divided in 16 bins by sending the clock pulse through a 16 element delay chain; the position of the clock pulse in the chain at the moment of a hit was recorded. The delay chain was stabilized by a phase-locked loop; this ensures that the full delay equals the interval between clock pulses. Information on address and time for wires with a hit was stored in the output buffer of a TDC chip. An ASIC chip for 16 channels with a common output buffer for 128 hits containing words of 21 bits (17-bit time and 4-bit address) was manufactured in $1\ \mu\text{m}$ technology. The readout of the chip was initiated by a non-empty flag of the output buffer; it proceeds without interruption of the data input, as required for a dead-timeless system. The contents of the 16 TDC chips in one VME module, connected to the wires of one plane, were transferred to a circular buffer, where hits were stored according to their time in slots of 400 ns width.

In parallel to the data extraction from the TDC output buffers, a fast synchronous output (with 25 ns intervals) of the TDC chips, containing the instantaneous multiplicity, was used to calculate the plane multiplicity in the preceding 100 ns time interval. There was an option in the readout system to reset the TDC output buffers when this plane multiplicity exceeded a limit (normally set to 7). In order to facilitate the offline analysis, a special "overflow" word was added to the data stream whenever such a reset occurred.

The circular buffers to which data were transferred from the TDC chips were located in separate VME modules, one for each chamber plane. Data were transferred to these modules on a special high-speed (20 MHz) backplane. The circular buffers could be accessed by level 1 (LV1) or level 2 (LV2) triggers.

Information from the chambers was assembled by a DSP in the master crate. From there, formatted events were sent through an optical link to the central event builder of the experiment. The central intelligence of the readout system was provided by a single board computer⁵⁾ in the master crate; VME connections to the chamber crates were made by VIC bus.

The rate capability of the system was 1 MHz of 32 bit words. This was sufficient for charged K^0 decays which normally consisted of two tracks contained in an average of 120 words, including headers. A buffer for two events was inserted at the place where the data transfer speed was lowest, at the output of the circular buffer. A veto signal was given to the Trigger Supervisor as long as this two-event buffer was full, effectively delaying the readout of the next event. This condition however was almost never met.

3.5 Wire chamber operation and performance

The NA48 drift chambers were operated with a 50% Ar + 50% C₂H₆ gas mixture flowing through a bubbler filled with demineralized water. The water was maintained at a constant temperature of about 4°C in order to keep the fraction of water vapour in the argon-ethane mixture at a level of a few per mille. The use of a small concentration of water vapour as an additive in drift chamber gas mixtures has proved to be very efficient in improving the operational stability of the chambers by limiting the occurrence of Malter effects on the field wires. The gas mixture was also filtered with a molecular sieve to prevent deposits from oil contaminants inside the drift chambers. The gas continuously

⁵⁾ CES FIC 8234 single board computer

flowed through each chamber at a rate of about 75 liters per hour and was released into the atmosphere outside the experimental hall. The gas pressure inside the chamber was maintained at about 0.7 mbar above the pressure of helium in the tank surrounding the spectrometer. Its value was continuously monitored during data taking. In addition, two oil bubblers located at the entrance and exit gas outlets of each chamber were used as safety devices to prevent gas pressure exceeding safety limits. The gas flows of argon and ethane were controlled to a precision of about 1%, which guaranteed a stable and homogeneous gas mixture in the chambers.

The chambers were operated with a nominal high voltage value of -2300 V on the field wires and of -1440 V on the graphite planes. At the end of each SPS spill, the high voltage was ramped down by several hundred volts and then restored to its nominal value a few seconds before the beginning of the following spill. At the nominal beam intensity




Figure 8: DCH efficiency as a function of high voltage

of $1.5 \cdot 10^{12}$ ppp on the K_L target, the current drawn in the chambers varied between 20 and $40 \mu\text{A}$ per plane, depending on the location of the detector along the beam line. At the nominal high voltage values, the gas gain was measured to be about $8 \cdot 10^4$ and the average drift velocity of the electrons across the drift cell to be $45 \mu\text{m}/\text{ns}$.

Signals at the output of the preamplifiers were discriminated with a typical threshold value of 30 mV corresponding to a peak current of about $3 \mu\text{A}$ detected on the sense wires. With nominal high voltage and discrimination threshold settings, the efficiency of the sense wire planes was measured to be about 99.5% and found to be rather uniform all over the detection plane. Fig. 8 shows the dependance of the average chamber efficiency as a function of high voltage for two threshold values.

Due to the use of long wires in the chambers, the hit-time resolution depends rather strongly on the hit location along the wire. Fig. 9 shows a typical variation of the space




Figure 9: Space resolution in a view as a function of the hit distance from the preamplifier




Figure 10: Distribution of coordinate residuals in one projection

resolution provided by one view of a chamber as a function of the distance of the hit from the preamplifier card. Using the information from track hits in all eight sense wire planes of a chamber and correcting for the incidence angle of the tracks, the average space resolution in each coordinate (X or Y) was measured to be better than $100 \mu m$. A typical distribution of coordinate residuals obtained by comparing the position measured in one projection with the one from two other views in the same chamber is shown in Fig. 10.

The momentum resolution provided by the spectrometer was checked with electron beams of well defined momenta of 25 and 100 GeV/c and with the nominal magnetic field intensity. The resolution obtained was $\sigma_P/P = 0.48\% \oplus 0.009 P [GeV/c]\%$, where the first term was due to multiple scattering in the He tank and in the drift chambers, and the second term was due to the position resolution of the chambers.

The decay vertex position for two-track charged events was measured using only the two chambers located upstream of the magnet. The distribution of the closest distance of approach of the two tracks had an rms spread of about 7 mm and the resolution on the vertex position along the beam line was about 0.5 m at the AKS position. The invariant $\pi^+\pi^-$ mass resolution for reconstructed K^0 decays was better than $2.5 \text{ MeV}/c^2$ and the resolution of the vertex decay time, using only drift chamber hits, was about 700 ps.

Precise alignment of each plane of drift cells was obtained using muon beams or two-track events from K_L decays traversing the spectrometer with the magnet turned off. Precisions of the order of $10 \mu m$ and $10 \mu r$ have been obtained on the average transverse position and rotation of planes, respectively. Similarly, relative positions and orientations of wires inside the drift planes have been determined and taken into account in the off-line track reconstruction programme. Finally the position of each drift plane along the beam line was measured with respect to the positions of the final beam collimator to an accuracy of about 1 mm.

3.6 The DCH readout for the NA48/1 experiment

In 1999 tests were made to check the performance of the detector with high intensity beams, similar to those foreseen for the 2002 run (NA48-1: high intensity K_S beam)[7] and

2003 (NA48-2: simultaneous K^+/K^- high intensity beams)[8]. These tests highlighted the necessity of a new readout for the drift chambers to function with both the high event rate and the rate of high multiplicity events.

The requirements fulfilled by the new system were the following:

- continuous sensitivity, i.e. concurrent data recording and readout during bursts;
- capability to accept up to 64 simultaneous hits per plane;
- ability to sustain single wire rates up to 400 kHz;
- no dead time under the above conditions;
- buffer at least 204.6 μ s deep for level 1 and level 2 trigger decision;
- separate nondestructive data extraction from the above buffer for level 1 and level 2 time stamps.

Moreover, to facilitate the integration with the experiment trigger/data acquisition system, the new readout system was data-compatible with the old readout system and resembled the layout of the old one. This meant, for instance, measuring the drift time in bins of 1.56 ns (1/16 of 25 ns), and arranging the electronics in nested blocks depending on chamber number and plane number inside the chamber [14, 15].

Each chamber was read out by a system of eight TDC and Readout Buffer units (TDC&RB, one per DCH wire plane) and one Crate Service and Collect card (CSC), all housed in a VME-9U crate. To interface the read-out with the L1/L2 trigger system and the experiment timing signals to the four crates (one per chamber), a dedicated board (new Master Service Card, nMSC) was also made.

The nMSC was intended to receive the trigger system requests, the clock, and the experiment timing signals and to distribute them to the four drift chamber readout crates. The synchronous L1 triggers were queued by the nMSC and dispatched to the CSCs as the L1 drift chamber readout was ready. An XOFF signal blocking the L1 trigger system was generated internally if this queue was nearly full or if it was collected from the downstream system.

The asynchronous L2 triggers were immediately dispatched without queuing them. An internal timeout checking the busy lines of the read-out generated an XOFF blocking signal to the Trigger Supervisor. The clock and synchronization signals were received from the experiment distribution system and were also sent to the CSCs as ECL differential signals.

The CSC performed as a supervisor for the TDC&RB modules in the crate. It received the experiment clock, synchronization signals and trigger requests and distributed them to the other boards in the crate using the P0 bus (timing) and private P3 bus (trigger). The board fed back to the nMSC two handshake signals (Busy and XOFF) for L1 and L2 triggers.

The CSC was also used to produce 8 VME programmable threshold voltages (12 bit resolution) for the drift chamber preamplifiers and to control a chain of 16 16-fold differential ECL pulsers used for testing purposes.

Moreover, it collected multiplicity and L2 data from the TDC&RBs. The multiplicity data, five bits of hit count information from each TDC&RB, were processed through a chain of VME programmable look up tables to produce a multiplicity trigger word which was sent to the experiment Trigger Supervisor. The L2 data instead were collected via a private 32 bit backplane bus operating at 10 MHz and forwarded to the experiment DAQ system via an external 32 bit bus used by the experiment DAQ[16].

The design of the new readout for the NA48 drift chambers was based on a TDC ASIC that was promptly available and field-tested at the time when the project started:

the TDC-F1 ASIC[17]. The chip had been especially designed for the COMPASS experiment at CERN, but had resources and flexibility that also made it fit the present requirements. It was decided to use the TDC Compact Mezzanine Card (CMC) daughter card specifically designed for the COMPASS experiment[18] and slightly modified for NA48 as the building block for the system.

A CMC holds 4 TDC chips for a total of 32 input channels, together with the logic needed to provide the reference clock, trigger signal and serial configuration bus to each chip. The chips were readout in daisy chain, token passing mode: data were stored in the CMC output FIFOs. These were two 18 bit parallel FIFO, 1024 words deep, for a 36 bit output data word, operating synchronously with the clock.

Each TDC&RB unit was made of two boards in a master/slave configuration, each board housing four CMCs so that one unit can cover the 256 wires of one chamber plane. Each CMC was interfaced with a “Front End and Buffer Block” (FEBB) composed of two Field Programmable Gate Arrays (“Input FPGA” and “Glue FPGA”) and two DPRAM memories, dedicated separately to L1 and L2 data. The Input FPGA was used to configure the TDC-F1 registers via a dedicated serial connection. It also provided input channel masking and a few counters for hit-counting purposes. The Glue FPGA was responsible for the formatting of the data into 11 bits to match the experiment resolution, and for processing independently the L1 and L2 requests, extracting from the DPRAMs the hit measurements within the trigger matching window. Moreover the Glue FPGA generated and distributed a 1.6 μ s periodic trigger to the CMCs, so that the TDCs were continuously sending out data.

L1 data were collected from the eight FEBBs of a master/slave pair by the MassBox Interface (MBI) FPGA and were sent via a 40 MHz, 16 bit, bus to an external associative memory device which was part of a processor farm (see Section 8) for tracking analysis and charged L2 trigger generation. L2 data were collected from the FEBBs by the DAQ Interface (DAQI) FPGA, which sends them to the CSC through a private 32 bit backplane bus operating at 10 MHz.

4 The Liquid Krypton Electromagnetic Calorimeter

One of the most important properties of a liquefied noble gas calorimeter used as an ionization chamber, is that, once a sufficient level of purity is maintained, it produces an intrinsically stable signal, well correlated to the energy of an incoming photon or electron. Such a signal can be amplified, shaped and recorded digitally with circuitry that can be calibrated accurately and continuously.

Liquid krypton was adopted, in a quasi homogeneous configuration, allowing for the full (> 99% containment) development of electromagnetic showers within a cylindrical volume with 5.3 m² cross-section and 127 cm depth. Thin Cu-Be ribbons were chosen to construct the electrodes in the form of longitudinal towers of about 2x2 cm² cross section. The active volume was divided into 13248 2x1x127 cm³ double ionization cells, with the central detector as anode. At the downstream end, the anodes (which together with the grounded side ribbons form a low impedance transmission line) were directly connected to preamplifiers via low inductance decoupling capacitors and, through high impedance resistors, to the high voltage (3000 V) which produced the drift field for the ionization. The mother boards housing the preamplifiers consisted of 18-layer printed circuits through which precisely calibrated amounts of charge could be injected into selected groups of cells. The preamplifiers were connected to external transceivers by individual 50 ohm coaxial cables via feedthroughs at room temperature [22].

The electrode structure, together with the attached preamplifiers, calibration system and cables, was housed in a vacuum-insulated cryostat consisting of a horizontal cylinder with a vertical cylindrical neck. The 260 cm diameter warm windows were formed from 4 mm thick aluminium sheets. The front cold window was double walled; its outer wall, which holds the pressure of the liquid, was a convex 2 mm thick stainless steel foil. The inner wall was flat and only 0.5 mm thick. A specially designed hexcel aluminium structure connected the two foils. The windows had a 20 cm diameter central aperture for the vacuum beam pipe. The total amount of material upstream of the sensitive volume corresponded to 0.58 radiation lengths.

Under operation the cold window is deformed by the LKr pressure, therefore causing the formation of a lens (with the sagitta of about 3 mm) of passive liquid upstream of the front plate of the electrode structure. In order to minimize this effect, compensating plates were manufactured according to the expected deformation from plates of the same composite material used for the electrode structure. Capacitive sensors applied on the front face of the compensating plate were used to monitor the residual gap under operation, which was determined to be in the range of 0.3 to 1.2 mm.

The active front surface extended from the beam pipe to an octagonal outer boundary with inscribed circle of 128 cm radius. The neck of the cryostat housed a double-walled, bell-shaped vessel which provided adequate expansion volume for the krypton. The cables for the preamplifiers, their power supply, temperature and pressure gauges passed through the space between the outside of the bell and the neck. The electrical feedthroughs were arranged on 16 vertical columns of four feedthrough plates each equipped with 12 vacuum-tight, glass-insulated, 50-pin *Cannon* connectors.

The total heat input to the cryostat was such as to produce 4 liters of gaseous krypton per second by evaporation. The evaporated krypton flowed from the top of the bell through a purifier and was subsequently recondensed and returned to the cryostat [19]. This arrangement allowed uniform and stable temperature conditions to be maintained over the whole liquid krypton volume, with a level of purity such that the mean life time for free electrons exceeded 300 μ s i.e. 100 times the drift time across a cell.

4.1 The cryogenic system

The cryogenic system consisted of the calorimeter cryostat, a storage dewar for the liquid krypton and two nitrogen dewars. The calorimeter was equipped with a cool down/warm up unit. Liquid transfer was generally achieved with a centrifugal pump. There were several purifiers for purification of the krypton either in the liquid or in the gas phase.

The vacuum insulated cryostat had a volume of about 9000 liters. The main static heat source (2 KW) was the solid conduction through the suspension of the inner vessel and through the signal and power cables. Additional 1.3 KW were produced by the preamplifiers in the liquid. Direct cooling of the krypton with nitrogen was possible in case of emergency, via heat exchangers in the gas space of the calorimeter. However, use of heat exchangers was not advisable if high stability was required, because of the large temperature difference between the two liquids, making the temperature control difficult. Instead, a krypton condenser using liquid argon as intermediate cooling was developed.

The krypton condenser was a separate unit outside the cryostat. The gas was brought in thermal contact with a heat exchanger cooled from a bath of saturated liquid argon at 10 bar and 117 °K, slightly above the triple point of krypton (115.95 °K). The argon was in turn cooled by liquid nitrogen the flow of which was controlled such as to

maintain the argon bath pressure constant.

The evaporated krypton gas (450 l/min NTP) was sent through a purifier before being reliquefied. The cold surface of the condenser (3 m² for a nominal cooling power of 4.5 KW) was made of 230 vertical pipes connected at the top with the argon reservoir. The krypton condensing on the outer surface of the pipes (0.7 l/min) drops down into a collector and flows back to the calorimeter via a vacuum-insulated line.

Commercial Oxisorb⁶⁾ purifiers were used to purify the krypton in both phases. However, purification in the liquid phase was avoided during data-taking periods to minimize the disturbance to the thermal conditions of the calorimeter. In the liquid phase, the liquid was taken from the bottom and passed through the purifier with a centrifugal pump, of which the cold part was hermetically sealed from the electric motor by means of a magnetic coupling.

Cooling down was achieved first by pumping the cryostat, then by filling it with krypton. The gas was subsequently circulated through an external purifier. In parallel, the calorimeter was cooled down gradually at a speed of 0.5 °K/hr by means of a flow of cold nitrogen gas through cooling channels inside the krypton vessel. Krypton was gradually added to maintain a small over-pressure. Warm up of the cryostat was achieved with the same system, which was equipped with electric heaters.

A 10000 liter storage dewar was available if it were needed to empty the cryostat. However it was seldom used, as the cryostat has been kept full for 8 years.

4.2 Signal Formation and Determination of Absolute Energy Scale

The total number of free electrons which result from the development of a shower, fully contained in the liquid krypton, was a direct measure of the energy E of the incoming photon or electron. In a multi-electrode chamber the ionization was spread over a number (in our case about 120) of cells. The initial current induced on the anode of a cell in which electrons with a total charge q (directly proportional to the amount of shower energy deposited in the cell) start to move with a drift velocity V_d is $I = \frac{q \cdot V_d}{W}$, where W is the width of the cell. For the 3 KV applied voltage the dependence of the drift velocity on the width of the cell can be neglected and the ratio between induced signal and ionization is inversely proportional to the width of the cell, with the same proportionality constant for all cells.

The lateral distribution of an e.m. shower was strongly peaked transversely to the direction of the incoming photon or electron and up to 40% of its total energy can be deposited in a single cell. Therefore, to obtain a ratio between induced initial current and shower energy deposited which was independent of the specific cells involved, it was imperative to maintain strict tolerances on the relative widths of the cells. It was possible to achieve a $\pm 45 \mu\text{m}$ accuracy in electrode spacing by using as electrodes 40 μm thick, 18 mm wide Cu-Be ribbons, to which a longitudinal tension of 2 N was applied and forcing them every 21 cm against spacers (5 mm thickness) precisely machined, constraining the transverse position of the ribbons to within ± 0.25 mm, from epoxy-fiberglass composite material⁷⁾ in a ± 5 mr zig-zag fashion (Fig. 11). This accuracy translated into a $\pm 0.5\%$ spread of the reconstructed energy of the shower relative to its real energy deposition in the calorimeter.

⁶⁾ Messer Griesheim GmbH

⁷⁾ Material 4411W, produced by Stesalit AG, CH4234 Zullwill. End-plates and spacer-plates were machined respectively by Cinel srl, I35100 Vigonza and Pantotech Italia srl, I24027 Nembro

The ribbon electrodes were constructed from Cu-Be reels manufactured according to tight specifications on flatness and straightness of edge cut⁸⁾. A dedicated device unrolled and cut the ribbon at the proper length, pressed corrugation in proximity of the end and provided proper termination for connection to end-plates and read out electronics. The corrugation near the ribbon ends (± 1 mm deep and 12 (25) mm long at the front (back) face) are required to provide the electrodes with the mechanical compliance (0.7 N/mm) necessary for proper tensioning, and to solve problems related to deformation of the end-plates (of 50 mm -front- and 64 mm -back- thicknesses) under the load of the ribbons and to thermal contraction [20].




Figure 11: LKr cell design.

It was important that the relative energy scale and the transverse dimensions of the fiducial region were established and maintained as accurately as possible; an error of only $2 \cdot 10^{-4}$ in the energy scale induced a systematic error in $Re(\epsilon'/\epsilon)$ of 10^{-4} , which was our design goal. To achieve this accuracy it was necessary, in addition to the uniformity of the cell width, to keep the transverse localization within a tolerance of 0.1 mm over 100 cm over the whole active volume. This was achieved using the spacer-plates already mentioned and mapping accurately the position of the slots. In addition, the numerically-controlled milling machines used to produce the spacer-plates rendered it possible to change the pitch of the cells at no additional cost; an increase of the pitch by 0.19% between successive spacers was chosen. In this way a projective tower structure was obtained diverging from a point 110 m in front of the calorimeter, close to the average longitudinal position of the kaon decay vertices. The result was that the angle between photon pairs originating at a known distance from the calorimeter (e.g. the sharply defined beginning of the K_S fiducial region) could be determined from the reconstructed transverse shower development, independently of longitudinal fluctuations. Since the invariant mass of the π^0 was known with high precision and $M_{\pi^0}^2 = E_1 * E_2 * \sin^2\theta_{12}$, the accurate measurement of the angle

⁸⁾ Manufactured by Matthey SA, CH 2520 La Neuveville

between the pair of photons directly translated into a determination of the energy scale for the calorimeter.

The uniformity and time stability, over a period of several years, of the ratio between the energy of incoming electrons or photons and the ionization produced by their showers, was found to be constant to better than 0.1%.

4.3 Cold electronics and calibration

Beside the fact that a cryogenic environment offers good thermal stability, the choice of cold preamplifiers was made for two other reasons:

- To minimize the charge transfer time to achieve speed and fast shaping of the pulses
- To exploit a property of silicon field-effect transistors, whereby the noise as a function of temperature has a minimum at the liquid krypton temperature

A cold charge-integrating preamplifier was used, based on a BNL design [21], with a restoring time of 150 nsec. The dispersion of the gains of the 13500 preamplifiers was 0.9%.

The use of DC levels and a digital control addressing system allows the distribution of accurate electronic calibration signals to any chosen group of 8 cells and hence the determination of the electronic gain of each channel. The determination of the pedestals and of their time stability was achieved by recording their values between and during beam bursts (with a frequency proportional to the instantaneous beam intensity). A global calibration constant was measured for each channel by injecting a known current into the input of the preamplifier and comparing the output signal with the one obtained pulsing the calibrator [22]. The ratio of induced current to the corresponding recorded signal was thus monitored and equalized to better than 0.8%.

4.4 Signal Processing

A system of transceivers installed on mother boards on the warm feedthroughs at the top of the calorimeter performed zero-pole cancellation, to eliminate the 150 ns restoring time constant of the preamplifiers, and drove the digitizing electronics via individually shielded 10 m long twisted pair cables [23]. The output signal had the shape of the induced current with a rise time of about 22 ns. In the subsequent 64 channel CPD (Calorimeter Pipeline Digitizer) modules the signals were integrated-differentiated to an approximatively Gaussian shape with 70 ns FWHM, a pulse height proportional to the initial current and a 3% negative undershoot, which lasts as long as the drift time of the electrons across the cell. In order to achieve the required dynamic range (from 3 MeV to 50 GeV) and to preserve the energy and time resolution of the detector, the calorimeter pulses were continuously digitized by 10 bit 40 MHz flash ADCs preceded by four dynamically switched attenuators, effectively resulting in a 14 bit dynamic range with 10 bit resolution [24, 25, 26, 27, 28].

4.5 Zero suppression of the data

The digitized information from all LKr cells was processed in a correlated way in order to perform the zero suppression: the data from cells not useful in reconstructing the cluster energy were removed from the data stream without compromising the high resolution. This function was handled by a specialized custom engine, the Data Concentrator (DC) system, which processed the data in parallel at high speed. It consisted of 32 boards, each containing 32 ASIC processing elements (mapping 16 calorimeter cells each); the data was received through eight optical links per board, each running at 20 MB/s.

The parallel processing elements worked in pipelined mode on the incoming data, performing pedestal subtraction and gain correction on each channel, looking for cluster seeds and marking neighbouring cells for readout in a multi-step process of "halo expansion". This process was expanded beyond chip and board boundaries with parallel intra-chip geographical connections and a custom bus. Pedestals, gain and seed threshold tables were loaded at initialization time and could be chosen on an event-by-event basis depending on the trigger type. This also applied to the halo algorithm used, which ultimately determined the number of cells being transferred to the output stage and could be selected among a choice of eight programmed into the controller memory. The processing of the data occurred in parallel in all elements, paced by an external micro-controller board with a constant 10 kHz pipeline clock. Several events were simultaneously present within the DC, at different stages of processing, and all the samples selected for readout, belonging to cells at a variable radius around seeds, were stored in dynamic buffers on each processing board, obtained from the 320kB on-board RAM. Data output of assembled event fragments was sent through a single optical link per board, similar to the four input ones.

The control of the ASIC processors was performed by a commercial on-board micro-processor (Motorola MC68EC040) with a minimal custom operating system running cross-compiled code, through a client-server mailbox communication protocol implemented in a dual-ported RAM. All the boards were in 9U VME format and housed in two crates, one of which also hosted the controller board; optical inputs were located on the front of the boards and outputs on the back, where the inter-board bus was also implemented on a custom backplane.

In the year 2002, the LKR read-out underwent a software upgrade to almost double, the read-out speed. The key point of the new read-out mode, called SuperEvent read-out, was the collection of two consecutive triggers to build a SuperEvent at the level of the CPD linear buffer, before sending the data to the Data Concentrator for zero suppression. The SuperEvent was then treated as a single event by the rest of the read-out chain. When reaching the PC farm, at the Event Builder level, the SuperEvent was split into the original events to allow the complete event building.

A SuperEvent was transferred and treated by the DC as a single one. The zero suppression algorithm was applied on the basis of the OR-ing of the trigger words of the original events in the pair. Because of this, the DC suppressed only the cells which were empty in both events, those being the result of the spatial superimposition of the two events. As a consequence, in every event, groups of cells were read-out even though they were empty, just because they were the image of clusters existing in the partner event, the so called "Ghost Clusters".

After splitting and decoding the LKR data in the Event Builder PCs, the cells were passed through the GhostCluster Suppression Algorithm and were afterwards written according to a new data format. This suppression algorithm looped over all cells read-out, creating a hit-map for every cell containing the difference between the maximum and minimum pulse heights in the 8 timeslots. A seed cell was defined when this difference was above a certain (programmable) threshold. An outward spiral expansion with a predefined box-radius was performed around every seed cell, in order to flag the accepted cells. A GhostCluster resulted in a group of cells (typically a $\sim 10 \times 10$ box) which contained only pedestal samples.

For efficiency studies, a downscaled number of events with flagged but not suppressed cells were passed through the system.

4.6 Shower Reconstruction

The pulses were reconstructed using the digital filter method [29]. The average accuracy with which a calibration pulse was reconstructed at high energy was $\approx 10^{-3}$ in energy and ≈ 130 ps in time. Usually, three samples (centered around the maximum) were used to measure the energy and the time. In case one sample was not correctly measured, only two samples were used [30].

A cluster was defined if one cell was a local maximum (i.e. had more energy than any of the eight surrounding cells), if it had an energy larger than 0.25 GeV and if its energy E (in GeV) satisfied the relation:

$$E > (0.18 + 1.8E_{ave})$$

where E_{ave} was the average energy (GeV) of the eight surrounding cells. Once a cluster was found, its position was estimated using the centre of gravity (COG) of the energy deposition in 3×3 cells and a first estimate of its energy was obtained from the central cell. If only one cluster contributed to a cell and if the cell was within the cluster radius (11 cm) of the COG, its energy was directly added to the first estimate of the cluster energy. If the cell received energy from two (or more) clusters at the same time, the energy of this cell was shared among the clusters according to the formula $E_i = E_{cell} \frac{W_i}{W_j}$, where W_j was the expected energy in the cell from the cluster j , computed using the first estimate of the cluster energy, the distance of the cluster from the cell and the expected energy profile (taken from GEANT Monte-Carlo). If the cell received contributions from clusters at different time, a time sharing of the pulse was made. By means of this sharing procedure, new estimates of the cluster energy were obtained. The position was also recomputed from the COG of the energy distribution in 3×3 cells. Corrections to the position and the energy were applied and the sharing procedure was iterated if two clusters were separated by 11 cm or less or if there was a very low energy cluster close to a higher energy cluster. In this case, the low energy cluster was assumed to come from a splitting of one single photon and it was dropped from the cluster list. Ad hoc corrections were applied if the shower landed in the proximity of a dead cell. A fiducial cut of 2 cm around the centre of a dead cell was later applied for the physics analysis.

The position was estimated from the COG of the energy deposition in 3×3 cells around the seed. This COG was not an unbiased estimator of the true position because of the finite cell size, therefore a correction was implemented to relate this COG (which was in cell units) to the cluster position (in cm).

The following corrections applied to the energy were independent of the energy, so they do not intrinsically change the linearity: correction for the energy variation along y (due to electric field variations between electrodes), correction for the energy variation along x (due to the reduced charge collection efficiency near the anode). In 1997 the calorimeter was operated at the reduced bias high voltage of 1.5 KV/cm. This required a small space-charge correction to be applied to the data [31].

4.7 Calorimeter performance

The electronic noise was a significant source of degradation of the intrinsic energy resolution of the calorimeter. In the design of the electronics, a trade-off between fast read-out (needed to cope with the high rate of incoming particles) and small equivalent noise energy was chosen [24]. The total noise per shower was evaluated from the average of

the first two FADC samples which precede the signal summed over the ~ 100 cells which belong to a cluster. The R.M.S. (σ) value amounted to $\sim 121(116)$ MeV per shower. The coherent noise per channel amounted typically to ≈ 0.12 LSB (1 LSB = 4.2 MeV) and the incoherent noise per channel was roughly 2.0 LSB. For a shower box of ~ 100 channels the incoherent noise dominated, indicating the high level of noise immunity achieved in constructing and operating the calorimeter electronics.

Assuming that the drift-gap spread was quite small by construction and that the temperature distribution inside the liquid krypton active volume was uniform, a source of uncertainty on the cell-to-cell response of the calorimeter was due to the precision with which the parameters of the calibration system were measured. The charge injection calibration reduced the gain spread of the electronics from $\sim 3\%$ to $\sim 0.8\%$. Since up to 40% of the energy was deposited in one tower, a 1% smearing of the calibration constants should contribute less than 0.5% to the cell-to-cell response. We used the E/p ratio between the energy as measured by the calorimeter and the momentum as measured by the magnetic spectrometer, calculated for electrons from K_{e3} decays, to cross-check the cell-to-cell response of the calorimeter. The E/p ratio distribution R.M.S. (σ) of the distribution was 0.58% (0.44%). The spread was consistent with the precision ($\approx 0.6\%$) of the measurement of the calibration system parameters. The cell-to-cell spread contributed to the constant term of the energy resolution. To reduce this error we applied a correction factor to each cell.

The K_{e3} electrons were used to measure the linearity of the calorimeter. The $\approx 0.8X_0$ of passive material in front of the active krypton created a nonlinearity that could be described by a 50 MeV/E term in the E/p distribution. The detailed GEANT simulation showed that 50 MeV was the effective difference in energy lost for electrons with respect to photons. This difference was due to different dE/dx and radiation losses between photon and electron showers in the dead material in front of the active krypton. The E/p ratio for electrons deviated from 1 by 0.1% at most (Fig. 12).

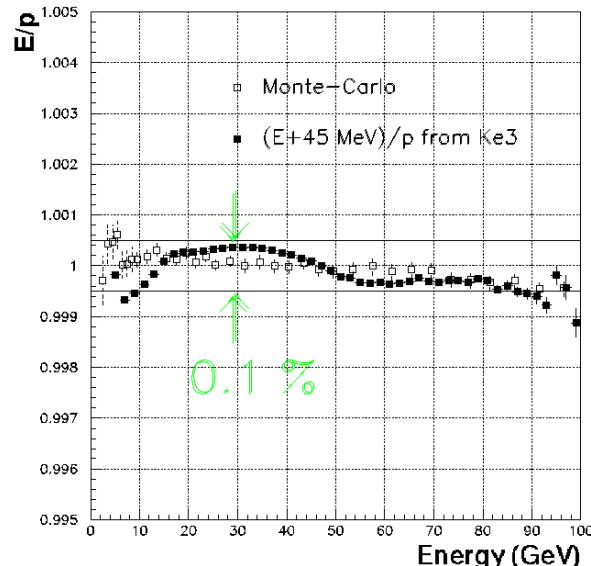


Figure 12: Linearity of the energy response

The energy resolution of the calorimeter was measured using an electron beam of different momenta (15, 25, 50, 100 GeV/c) and 0.1% momentum bite, which could be

deflected to cover the full height over a central band ± 15 cm in width (Fig. 13).




Figure 13: Energy resolution of the calorimeter

The σ parameters of Gaussian fits to the $\frac{\delta E}{E}$ distributions were between 1.3% (for electron energy 15 GeV) and 0.6% (100 GeV).

The time reconstructed from the pulses was corrected for differences in time offset among different cells. This offset was also taken into account when the time of the central cell was used as a first estimate of the phase to reconstruct the other cells of the cluster. The times of the central cell and of the most energetic lateral cell in x (which were stored in the output bank) were corrected for the y dependence inside the cell due to cross-talk effects. An upper limit for the time resolution was obtained studying the difference between the time calculated from the hit cell and the time from the side cell with the highest energy deposit. The σ parameters of such distributions were measured to be between 540 ps (for electron energy 15 GeV) and 270 ps (100 GeV) (Fig. 14)[32].

The granularity of the calorimeter allowed a good position resolution to be obtained. The position resolution could be measured for electrons by comparing the centre of gravity (COG) of the shower over a 3×3 grid, corrected for its non-linearity, with the position reconstructed from the two drift chambers (DCH) placed upstream of the calorimeter (Fig. 15).

In summary the performance of the liquid krypton calorimeter was (in units of GeV, cm, ns)[33, 34, 35]:

$$\frac{\sigma_E}{E} = \frac{0.032}{\sqrt{E}} \oplus \frac{0.09}{E} \oplus 0.0042 \quad (2)$$

$$\sigma_{X,Y} = \frac{0.42}{\sqrt{E}} \oplus 0.06 \quad (3)$$

$$\sigma_t = \frac{2.5}{\sqrt{E}} \quad (4)$$




Figure 14: Time resolution of the calorimeter




Figure 15: Position resolution of the calorimeter

5 Trigger counters

5.1 The Hadron Calorimeter (HAC)

The Hadron Calorimeter (HAC) was used to measure the energies and positions of particles showering in the material of this detector. The HAC was mainly used as a trigger device. It gave energy sums for the four quadrants and a total energy sum, which were used in the fast pre-trigger and as a veto in the trigger for neutral particles.

The HAC consisted of an iron-scintillator sandwich of 1.2 m total iron thickness (6.7 nuclear interaction lengths), with alternating layers of iron and plastic scintillator. It was divided longitudinally into two separate modules (front and back), each consisting of 24 iron plates, 25 mm thick, of dimensions $2.7 \times 2.7 \text{ m}^2$ (Fig. 16).

Each scintillator plane, inserted between the iron plates, consisted of 44 strips. Each strip spanned only half the calorimeter so that each plane was comprised of two half planes. The two central strips of each half-plane were shaped on one end to wrap around the central hole for the beam pipe so they had a width of 108 mm with a length on the short side of 1194 mm and 1243 mm on the long side. In the back module only, the two




Figure 16: Structure of the hadron calorimeter

outer strips in each half-plane were 1150 mm long and 108 mm wide. All other strips were 1300 mm long and 119 mm wide. The thickness of each scintillator was 4.5 mm. There were 24 planes in the front module and 25 in the back. In consecutive planes the strips were alternately aligned in the horizontal and vertical directions. In each module the strips with an identical alignment were coupled to the same photo-multiplier using a *Plexiglas* light-guide. In total there were 176 channels, 88 for each module.

The 176 channels were split into four groups of 44. Each group was connected to one CPD. The CPDs were developed especially for the readout of the liquid krypton calorimeter. To be able to use the same unmodified CPDs, the photomultiplier signals had to be processed to obtain a signal shape which was comparable to those of the LKr calorimeter. Therefore each PM pulse with a rise-time of ~ 10 ns and a fall-time of ~ 50 ns was modified to obtain a signal with a risetime of ~ 50 ns and a falltime of ~ 2.5 μ s. These signals were shaped again and digitized in the CPDs.

In 1995 the energy resolution was determined to be $\frac{0.69}{\sqrt{E}}$ for hadronic showers.

5.2 The AKS counter

The measurement of $Re(\epsilon'/\epsilon)$ required defining with high precision the fiducial length where K_S and K_L decays were accepted. This was accomplished, as far as the beginning of the fiducial region for K_S was concerned, by the AKS, a dedicated detector placed on the K_S beam line, able to detect π^\pm and γ from π^0 originating from upstream kaon decays. The AKS vetoes both charged and neutral K_S decays which happen upstream of its position.

The AKS detector was a system of four scintillator counters and a photon converter (Fig. 17) located in a short interruption of the vacuum on the K_S beam at the exit of the K_S collimator. The first (upstream) scintillator was a 30×30 mm² square plate, 4.5 mm thick, with a 3 mm radius hole cut in the centre and two photomultipliers (AKS1, AKS2) attached via light guides. Following the photon converter, there were three scintillator disks of radius 8 mm each, 4.5 mm thick, each attached via a light guide to a PM (AKS3, AKS4, AKS5).




Figure 17: Structure of the AKS counter.

The hole and the three counters were aligned among themselves on an optical table, and relative to the K_S collimator axis. The counter (AKS4) used to define the beginning of the decay region was the central one of the three circular counters; all the other counters were used to monitor its efficiency, which was measured to be 0.9993 ± 0.0003 .

In between the square scintillator and the three scintillator disks was placed the photon converter that could be moved remotely into/out of the K_S beam. The use of a crystal converter was particularly suited to this application because, for photons with incidence angles near crystalline axes, the pair production cross section is enhanced due to effects of coherence [36]. Different crystals have been considered for the AKS converter and tested in a tagged photon set-up at the CERN SPS [37].

The converter used in the AKS was an iridium crystal, 2.96 mm thick, corresponding to $0.9786X_0$ of amorphous iridium; if aligned properly, it corresponded to $1.79X_0$ for an average photon from a K_S decaying upstream of the AKS. The alignment of the crystal was found by maximizing the conversion of photons in the K_S beam and was routinely monitored during data-taking using the photons produced by π^0 s present at high rate in the K_S beam. The advantage of the enhanced pair production cross section is that the thickness seen by the K_S beam is minimized, therefore reducing the undesired scattering of the beam, while achieving an efficient photon conversion. The effective position of the AKS for neutral events was fixed within the converter at the depth corresponding to 50% photon conversion probability, while for the charged events it was located in the scintillator used as a veto (AKS4). Hence there was a difference in Z of 21 mm between these two positions and a correction was applied to account for kaon decays occurring in this space.

The end of the fiducial region for both K_S and K_L , as well as its beginning for K_L , were defined by software cuts.

5.3 AKL

The design of the large-angle Anticounters (AKL) were optimized to detect photons outside the LKr calorimeter acceptance. The AKL consisted of scintillation counters arranged in seven pockets of steel placed along the decay tank. Each pocket (Fig. 18) was longitudinally segmented into two similar layers of counters with a 35-mm thick steel plate in front, which acted as photon converter. The resulting total photon conversion efficiency was $\approx 95.5\%$.




Figure 18: Structure of an AKL pocket.

The counters, with trapezoidal or rectangular shape, were made of 10-mm thick NE110 scintillator plate. They were arranged to cover an annular shape around the calorimeter acceptance. The polar angle covered by the detector was from 10 to 50 mr for decays occurring along the z-axis in the fiducial region.

The number of counters per layer was twelve for the first four pockets and eight for the remaining ones.

The typical dimension of each counter was approximately $2000 \times 250 \text{ mm}^2$. Each counter was read out at the two ends. At each end the signal was discriminated with a Constant Fraction Discriminator (CFD) to compensate for the jitter due to the amplitude range. The variation in position of a hit along a 2-m long counter led to several nanoseconds of difference between the signals at the two ends of a counter. To correct for this effect a mean timer (MT) was implemented at the hardware stage to average the time between the two discriminated signals. Given the small ratio between width and counter length, the average time was almost independent of the hit position across the counter. The signals from the MTs of a given pocket layer were aligned in time and split into two formed signals to generate the inputs for a fast trigger and to the individual readout. To minimize the random veto probability due to the high rate expected in the counters, the signal length was constrained to less than 10 ns. For the fast trigger chain the counters of a pocket layer were OR-ed together to provide a signal associated with a hit in any of the counters in a layer. These ORs were finally aligned in time and summed up to constitute the overall AKL signal to be used in veto. The second set of MT signals and the signal amplitudes

from each phototube were digitized through the PMB system (see Section 5.6). At the offline stage it was then possible to calibrate the time alignment and furthermore to apply trigger rejection based on any selected counter hit configuration.

The AKL time resolution was found to be ≈ 550 ps (Fig. 19).




Figure 19: AKL time resolution.

5.4 The Charged Hodoscope

The NA48 trigger hodoscope was a system of scintillation counters which provided, with high accuracy, the t_0 signal for K_S and K_L decaying into charged pions. The detector (Fig. 20) consisted of a set of scintillators arranged into a plane of 64 vertical counters followed by a plane of 64 horizontal counters. The two planes were separated by 74 cm in order to distinguish prompt signals from calorimeter backscattering. Each plane was divided into four quadrants of 16 counters, designed to contain a circle of 121 cm radius. The thickness of a single counter was 2 cm, and its width was 6.5 cm close to the beam pipe and 9.9 cm further out. The length varied from 121 cm to 60 cm. The radius of the central hole was 12.8 cm.

The counters were made of NE110 plastic scintillator, later replaced by BC408 with similar fast light output and attenuation properties.

The scintillation light from each counter was collected from the downstream side via a *Plexiglas* light guide, with a fishtail shape, and recorded by a XP2262B photomultiplier. An absorbing black coat was painted on the opposite side to avoid reflected light hitting the PM.

All the counter surfaces were polished, except for the one on the side opposite to the PM, which was roughed before being coated with the black paint. Both the counter and the light guides were wrapped with an aluminized *Mylar* sheet of 25 μm thickness. Finally, each plane of the detector was wrapped with a black cover.

A passive box split the 128 signals coming from the individual PMs. The typical PM signal from a hodoscope counter was 300 mV high and 30 ns long. One third of this signal




Figure 20: Charged hodoscope.

was sent to the PMB readout system (see Section 5.6), the rest to updating discriminators with no dead time.

A time resolution better than 200 ps per counter was measured from the data. Fig. 21 shows the distribution of the time difference between the hodoscope and the tagger (see Section 6); Fig. 22 shows the distribution of the time difference between two vertical slabs of the hodoscope.

5.5 The Neutral Hodoscope

The experiment has demanding time resolution requirements in order to assign the charged and neutral pion decays to the correct K_S or K_L parent particle. The time-of-flight was measured between the tagger (see Section 6) and the main detector: the charged hodoscope, for charged events, or the Liquid Krypton calorimeter, for neutral events. Errors in time measurement, especially non-Gaussian tails in the time-of-flight distributions, could affect the precision of the experiment. In order to improve the time resolution for neutral events, a completely independent time measurement was performed using the neutral hodoscope.

The neutral hodoscope was a plane of scintillating fibres vertically installed in the Liquid Krypton Calorimeter at a depth of about $9.5X_0$, where the electromagnetic shower has its maximum. Scintillating fibres (*Bicron*, 1 mm in diameter) were preferred over massive scintillator rods because, for the given geometry, they offered better timing properties and more flexible light transport to the photomultipliers. The hodoscope was positioned at the second spacer plate of the calorimeter. The bundles (5.5 mm in diameter, ~ 20 fibres per bundle) were inserted into epoxy-fibreglass tubes, with inner and outer diameter of 7 and 8 mm, respectively, that were fixed to the spacer plate between planes of ribbons (Fig. 23). Fibre bundles were then grouped together, polished at one end and sent to photomultipliers located on the calorimeter front plate, inside the Liquid Krypton,




Figure 21: Tagger-hodoscope time difference




Figure 22: Time difference between two vertical slabs

but outside the active volume of the calorimeter. The optical contact between fibres and photomultiplier was made by mechanical pressure, produced by stainless-steel springs.

The neutral hodoscope was divided into 4 quadrants, each quadrant having 8 read-out channels for a total of 32 channels. The four channels closest to the beam pipe of each quadrant covered a vertical region 10 cm wide, the remaining four were 22 cm wide. Each channel collected signals from bundles of scintillating fibres.




Figure 23: Neutral hodoscope.

The read-out was based on the same PMB system (see Section 5.6), as used for the charged hodoscope, which records time and amplitude of the signals. The signals were also sent to the pretrigger system that provided a time measurement for neutral events, independent of that obtained from the calorimeter. This signal, called T0N, was a loose coincidence between opposite halves of the detector.

Event time resolution was of the order of 260 psec (Fig. 24) and single counter time resolution reached 200 psec. 95% efficiency was reached for photon energies above 25 GeV, 99% above 35 GeV.




Figure 24: Neutral hodoscope event time resolution.

5.6 Hodoscope Readout - The Pipeline Memory Board (PMB)

The deadtimeless readout of the charged particle hodoscope and other scintillation counters of the experiment was performed using a general purpose module, the "pipeline memory board" (PMB), developed for the experiment. The basic module provided the following functions for 8 channels of 24 bits:

- continuous storage of the 40 MHz input in an 8K circular buffer, the "pipeline memory", allowing trigger latencies in excess of 200 μ s. This function was performed by a specifically designed ASIC (PMChip [38]);
- transfer of trigger associated data from the pipeline memory to a multi-event asynchronous output buffer to give negligible dead times for data throughput of the order of 10 MB/s;
- readout of the output buffer with optional suppression of empty channels and sparsification.

The input to each channel was provided by daughter boards which plugged into appropriate connectors of a mother board. Different applications were allowed for by the use of three different daughter boards:

- the "ADC/TDC", that provided sub-nanosecond timing and pulse height information for scintillator signals;
- the "pattern unit", to store digital data produced by external, trigger related, sources;
- the "pipeline scaler" that allowed scaler data to be associated with an event, for measurements of instantaneous rates or the delay of the trigger with respect to a reference pulse.

The front end modules (PMB's) were housed in large size (VME 9U 400mm deep) crates, controlled by a "data routing module" (DRM), which broadcasted the trigger information to the individual PMB modules and steered the readout procedure. The first function was implemented via a custom bus connected to the undefined pins of the VME J2 connector. The readout of data the VME bus was used, but with a nonstandard

protocol, in which the transfer of data words was synchronous and modules containing valid unread data were enabled to assert data on the bus via the VME-bus "BG0" daisy chain rather than by individual addressing. The DRM modules also provided buffering of data and word counts.

The interfacing of the DRM modules to the trigger supervisor and to the readout PCs followed the standard data acquisition architecture of the experiment. Several DRM's were connected in daisy chain, via two independent ECL DTbuses, to a "TIC" processor (from which they received the trigger information) and to a "ROC" processor (to which they sent the data). Both processors were implemented on commercial "RIO" processors equipped with custom interface boards to the DTbuses and to the "TAXI" serial link with the trigger supervisor. Data were finally moved via VME bus from the "ROC", acting as VME master, to the "VFIFO" driver of the optical link to the data merger.

The readout of an event was performed in two separate phases, which were consecutive for a given event, but could otherwise be active concurrently. In a first phase the trigger information was received by the TIC and broadcast to all modules in the system and in each channel caused the transfer of a programmable number of words (≤ 16) from the pipeline memories to the 256 word output buffer. This phase has strict time limits and was performed promptly. In the second phase, information was moved from the output buffers to the ROC with optional suppression of empty channels. The time needed to complete the second phase depended on the size of the event, but other triggers could be processed during the readout of previous events, as long as the output buffers in the PMB's and the DRM buffer were not nearly full. Dead time could be generated (by sending an "XOFF" signal to the trigger supervisor) only when this condition occurred.

5.7 Muon Veto Counters

The primary function of the muon veto counters in the experiment was to reduce the Level 1 trigger rate by vetoing $K_{\mu 3}$ decays (27% of K_L decays). In addition, the counters provided identification of single and multiple muons in rare decay investigations, calibration and detector studies.

The complete muon veto counter array was composed of slightly overlapping long flat strips of NE110 plastic scintillator arranged in three planes, each plane being preceded by a 0.8 m thick iron wall to provide absorption of hadrons. Planes 1, 2 and 3, together with their respective absorbers were placed successively along the beam axis proceeding downstream from the HAC. The array occupied a longitudinal space of 4.3 m including an additional 0.4 m wall behind plane 3 to absorb backward-going products of interactions at the rear of the beamline.

Planes 1 and 2 were used in the veto and trigger logic and were each about 2.7 m square. The constituent strips for these planes were oriented horizontally and vertically respectively (Fig. 25). They were each 2.7 m long, 0.25 m wide and 1 cm thick, connected via a semiadiabatic light guide to EMI 9813KB photomultiplier tubes at each end, except for the central strip which was divided to allow the beam pipe to pass through the plane via a hole approximately 22 cm square. Plane 3 was used purely for monitoring and operational studies. It covered $2.7 \times 2.65 \text{ m}^2$ and was composed of six strips, each 2.7 m long, 0.45 m wide and 0.6 cm thick, with a semi-adiabatic light guide and Philips 56AVP photomultiplier tube at each end. The central strips were notched to allow for the beam pipe.

Signals from the PMT's were split into two branches. One branch passed to the veto/trigger logic after being further split to provide inputs to the monitor. Trigger signals

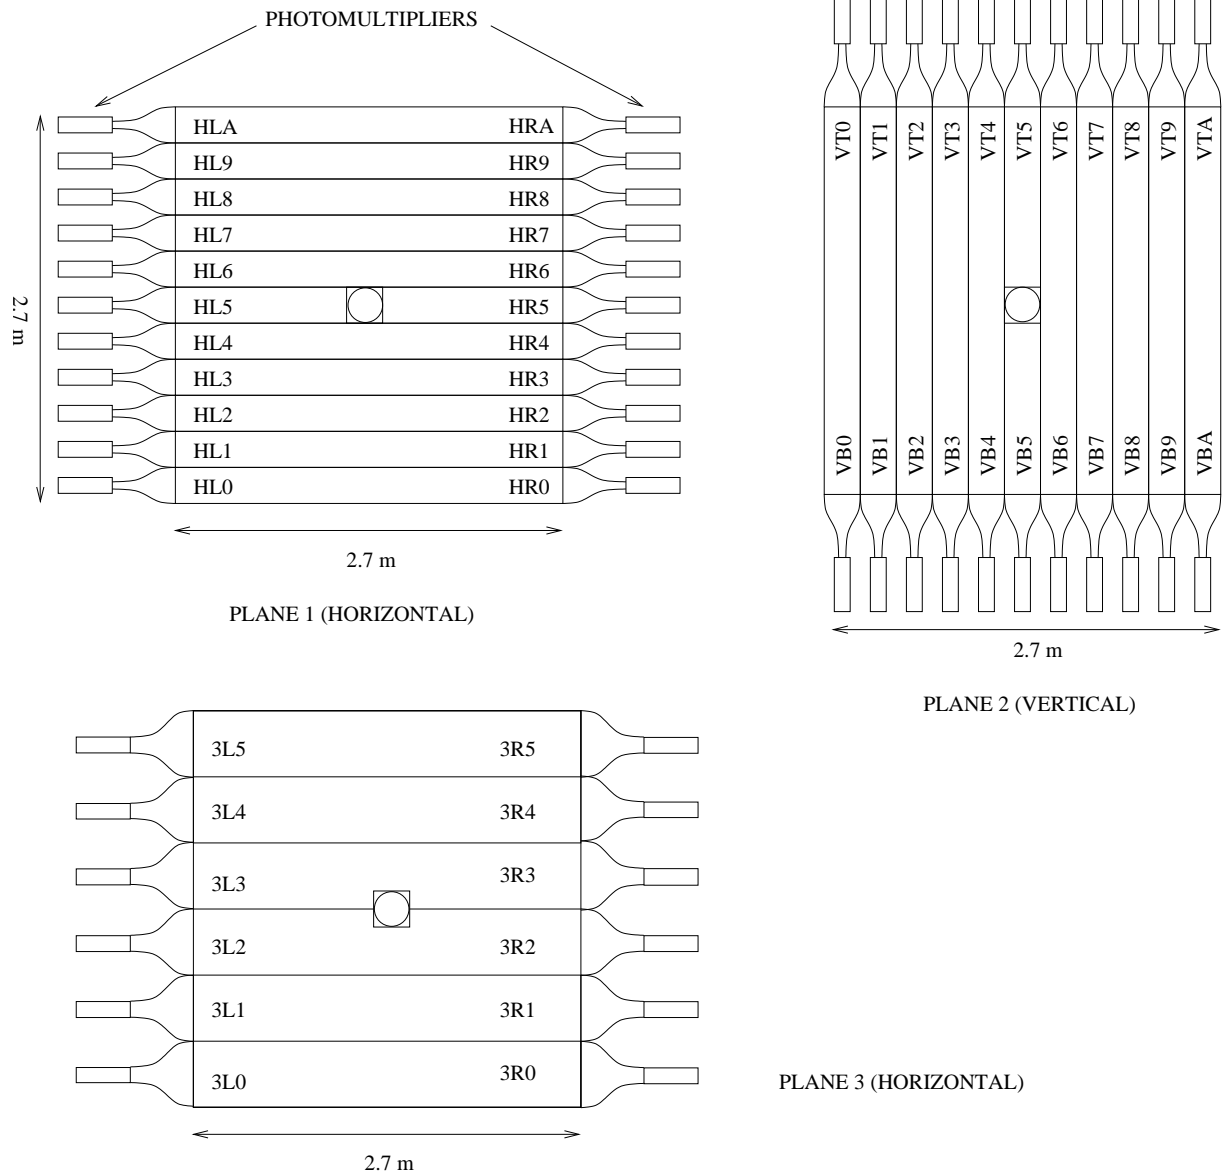


Figure 25: Structure of the muon counters.

were discriminated at the level of -40 mV and proceeded through standard CAMAC programmable delays and logic units to form coincidences defining muon hits in the $120, 25 \text{ cm} \times 25 \text{ cm}$ boxes formed by overlapping horizontal and vertical strips in planes 1 and 2, respectively. A single, purpose-built CAMAC muon trigger module [39], employing programmable delays and 22V10 Logic PALs, produced single and multiple muon triggers which were fed to the Level 1 trigger. This module could be configured and tested remotely with computer-generated signals, and its stability monitored.

The second branch from each of the MUV PMTs passed to the MUV readout electronics so that the PMT signals were recorded for each event along with the data from the other detectors. Precise timing information was later extracted from these data to provide offline reconstruction of muon hits. The readout electronics comprised 7, 8-channel, programmable threshold, 1 GHz sampling, pipelined TDC modules [39] plus readout controller and trigger controller. On receiving a trigger, the TDC modules sampled each input channel simultaneously at a rate of 1 GHz and provided a record of the time

above threshold for a period, typically 75 ns long, around the nominal interaction time. After zero-suppression in hardware, the data fragments were assembled and formatted before being transmitted to the event-building PCs. The data from each channel provided precise information about the time of each muon hit relative to the trigger time and the time above threshold, being correlated to the signal amplitude from the PMTs, was used in the timing calibration.

For single-muon events, the reconstruction efficiency was found to be $>99\%$, with time resolution in the TDC of 780 ps (Fig. 26, left) dominated by photoelectron statistics [40]. The large photo-electron statistics allowed a slewing correction to be performed, reducing the time jitter from ~ 700 ps to ~ 350 ps (Fig. 26, right). A trigger time resolution of 1.2 ns was achieved, limited primarily by residual channel-to-channel timing differences in the trigger logic.




Figure 26: Time resolution of the muon counters.

6 K_S tagging and the experiment clock

The K_S and K_L beams were separated vertically by 68.4 mm at the exit of the last collimator. This was sufficient to identify the origin of kaon decays into 2 charged particles

based on the reconstructed vertical vertex position. For decays into neutral particles this was not possible. The K_S or K_L assignment was therefore made by measuring the time difference between the passage of a proton in the tagging counters [41] upstream of the K_S target and the event time in the detector. The latter was taken from the trigger hodoscope in case of charged particles and from the LKr calorimeter in case of neutral mode decays.

6.1 The tagging detector

At the position of the tagging counter the projected proton beam density could be approximately described by Gaussian distributions. The total widths were ~ 1 mm for the horizontal and ~ 6 mm for the vertical direction. The beam was sampled by two sets of staggered scintillators arranged in two ladders which intersected in space both vertically and horizontally (Fig. 27). The orientations alternated along the beam direction. In each set there were 12 scintillators, which had a geometrical overlap of nominally $50 \mu\text{m}$ between successive detectors to avoid gaps in case of small misalignments. Fast photomultipliers and pipelined FADC readout ensured good time resolution and minimal dead time at particle rates exceeding 20 MHz.




Figure 27: The tagger.

Both sets of scintillators were mounted on a single carbon-fibre composite material structure. The support structure was machined to $\pm 5 \mu\text{m}$ accuracy and held the scintillator edges by means of spring-loaded pins and quartz washers. This ensured an area of $8 \times 8 \text{ mm}^2$ for the beam to pass through scintillator only. All scintillators were 15 mm long and 4 mm thick along the direction of the beam (1% of a radiation length). The scintillator width perpendicular to the beam varied from a minimum of 0.2 mm in the centre to 3.0 mm at the outside to keep the proton rate approximately equal in all channels.

The scintillators were made of BC-418⁹⁾, selected for its suitable decay time constant and radiation hardness. The central plastic pieces were exposed to roughly 10^{13} protons during 100 days of running. Radiation damage manifested itself in a slow light yield decrease of up to 50% over this period of time. It was balanced by an appropriate change of high voltage. All tagging scintillators were replaced by new ones after one year of operation.

Each scintillator was viewed by a Hamamatsu R-2076 photomultiplier located opposite the support structure. The rear faces of the plastic were coated with white optical paint to double the effective light output. The light was transported through Al-coated *Mylar* foil light guides into a short rod of UV transparent *Lucite* onto the photocathodes, which were only a few centimetres away from the scintillators. The photoelectron yield was about 300 p.e./proton. The phototube signals were transmitted to preamplifiers and fed into the custom-designed FADCs described below [42].

A nitrogen laser was used to establish and monitor the timing between the tagger and the other detectors. The UV light pulses (0.5 ns duration at 337 nm wavelength) were transported via quartz fibres to the single detectors, where they were locally distributed into short secondary fibres to illuminate individual photomultipliers. Furthermore the laser system served as an online monitoring for the tagger, controlling the functionality of each of the counters during data taking.

The scintillators of the tagger must be aligned precisely in the horizontal and vertical direction to ensure complete coverage of the proton beam. A remote-control motor system provided two rotational and two translational degrees of freedom with an accuracy of 0.01 degrees and 0.4 μm , respectively. The correct alignment was achieved in an iterative procedure: the beam was scanned in the vertical and horizontal directions in steps of 0.1 mm while the counting rates in the scintillators were recorded as a function of position. The resulting beam profiles in each scintillator thus provided angular and positional information simultaneously, which was used to correct the tagger position after each scan.

6.2 A 1-GHz Flash-ADC module for the tagging system

An 8-bit 1-GHz Flash-ADC module [42] was developed for the tagging system to digitise photomultiplier pulses and to recognise superimposed double pulses of protons crossing the tagging detector. The module contained two independent channels with 500 MHz sampling rate, which could be interleaved to obtain one channel with an effective sampling rate of 1 GHz. Both FADC channels ran continuously and data samples were stored for over 500 μs in a ringbuffer memory until they were overwritten by new data. During their live time data could be extracted from the ringbuffer without disturbing the write cycles and thus without creating any deadtime. A non-standard VME readout combined with the suppression of empty tagging channels accelerated the data collection.

The analogue input signal was sent through a low-pass filter, inverted, amplified, and transferred to the two FADC channels. (Alternatively, when running in 500-MHz mode, each channel received a separate input signal.) Each of the two channels used a commercial 8-bit, 500-MHz FADC chip of type TKAD10C, which consisted in turn of two interleaved 8-bit/250 MHz monolithic FADCs. The external 1-GHz clock (ECL) was divided by two and then used as the sampling frequency for both FADCs. The external reset signal was synchronised to the clock to give a precise starting point for the timing

⁹⁾ Bicron Corp., USA

information. An arbitrary number of FADC modules could be synchronised to each other by using one source for clock and reset.

The digitized data were written continuously into a ring buffer memory using interleaved write and read cycles. Each data sample entered the next free location of the memory, and when the memory was full the next point was stored at location one again and overwrote the old content. To reduce the frequency of the write cycles, each of the two continuous 250-MHz data streams of an FADC chip was split into four output streams of 62.5-MHz ECL signals, each of which was in turn distributed into four memory chips.

When a trigger arrived, an external VME master loaded the desired time interval and control bits for the data extraction (such as the zero-suppression threshold) and then sent the start address into the ringbuffer address counter. A read-out controller retrieved the data of interest from the ringbuffer, suppressed data points below an adjustable threshold (zero suppression) and wrote the data together with format words and timing information into a 16-bit FIFO memory. The VME slave interface was optimised to read data from the FIFO memories with a peak rate of up to 32 MB/s depending on the VME master. To accelerate the data readout a special subaddress allowed all cards inside one crate to be accessed (“broadcast write”) within 2 VME write cycles, and the same subaddress was used to read the FIFO memories from all cards inside a crate by one block transfer. Modules without relevant data could be excluded automatically from the readout procedure. An empty signal produced by the last module was sent to the interrupt input of the VME master and stopped the block transfer.

6.3 Reconstruction and performance

For each event trigger the corresponding data were extracted from the ring buffer using a 100 ns long read-out window. This was the starting point for the subsequent off-line reconstruction of pulses and, finally, the $K_S - K_L$ assignment for any given event. The reconstruction tried to find interesting regions where digitized pulse heights exceeded a certain threshold over the baseline. A simple peak-finding algorithm looked for pulses and prepared starting values for a fit of the photomultiplier pulses.

The signal $I(t)$ was described by a convolution of an exponential light decay and a Gaussian resolution function. Up to three overlapping pulses could be treated in a combined fit, allowing two close pulses to be resolved down to 4-5 ns. The reconstructed time per counter had a resolution of ~ 140 ps [3].

The Gaussian resolutions relevant in the tagging method were better than 200 ps for $K_0 \rightarrow \pi^+\pi^-$ (Fig. 28) and about 250 ps for $K_0 \rightarrow 2\pi^0$ decays. The efficiency was monitored continuously using the vertex position in charged-mode kaon decays, cross checked with Dalitz decays in the neutral mode. Depending on different proton rates in different years of data taking, the inefficiency of the tagging system was found to range between $1.1 \cdot 10^{-4}$ and $1.6 \cdot 10^{-4}$ [2, 3, 4], dominated by the inefficiency of the tagging detector itself.

6.4 The NA48 Clock

The NA48 clock [43] supplied a unique timing signal to the different subdetector systems. The tagger FADCs ran at a frequency of 960 MHz while all other subdetectors used 40 and/or 80 MHz. Therefore the central clock system, which produced both the original 960 MHz and a derived 80 MHz signal, was installed close to the tagger.

Between the tagger and the hodoscope, which were 220 m apart, a time resolution of 300 ps had to be achieved and its continuous monitoring to be ensured. To do this,




Figure 28: Time resolution of the tagger.

a special high-quality cable transmitted an 80 MHz pulse train from the central clock system near the tagger to the other subdetectors. This pulse train also transported the information of the start and the end of an accelerator burst by switching to 40 MHz for three or one 40 MHz periods, respectively.

The signal propagation speed was monitored by sending the clock signal back from one of the other subdetectors to the central clock system and by measuring the phase shift between the original and the return signal (accuracy 60 ps). For returning the signal, the same transmission line and the same type of electronics was used as for sending the original signal. This avoided ambiguities as to which line was suffering from variations in propagation speed.

The original 960-MHz signal was produced by a commercial pulse generator¹⁰⁾ while all other tasks (fan-out of 960 MHz to all tagging FADCs, derivation of the 80 and 40 MHz frequencies, integration of start/end of burst information into the 80-MHz pulse train, long-distance transmission, fan-out of 80 MHz to other subdetectors and comparison of outgoing and incoming clock phase) were carried out by custom-built NIM modules.

7 The Trigger System

7.1 The NA48 Pretrigger System

The NA48 pretrigger (Level 0) was based on logic signals coming from the charged hodoscope.

The signals from the hodoscope PMT were discriminated and individually timed within 1 ns. In this way, the time resolution was only limited by the physical dimensions of the counters and was about 5 ns.

The coincidence of counters in opposite quadrants relative to the beam pipe was used as trigger signal when a kaon decayed into a charged pions pair. The rate of this signal, called QX , was about 400 kHz at nominal beam intensity. To avoid pretrigger inefficiency when the plane of the decay pion pair was close to either the vertical or horizontal direction, the definition of the quadrants was extended to include the first counter of the neighbouring quadrant. The inefficiency for the detection of track pairs having the geometrical configuration expected for $K^0 \rightarrow \pi^+\pi^-$, measured at the level of the pretrigger system (hence including detector and fast logic electronics), was at the per mille level.

¹⁰⁾ Hewlett-Packard Model 8656B

The pretrigger system provided the timing for the subsequent stages of the trigger. Its logic signals were aligned in time and were available after about 200 ns. The total rate of these signals was less than 1 MHz.

7.2 Trigger for charged particles

The purpose of the Level 2 Charged Trigger System (L2C) was to reduce the rate of candidate $K^0 \rightarrow \pi^+\pi^-$ events to a few kHz. The principle of the L2C was to perform a fast online event reconstruction using the hits detected in chambers 1, 2 and 4 of the spectrometer [44, 45]. Based on this reconstruction, it flagged all events, the hits topology of which was compatible with a $K^0 \rightarrow \pi^+\pi^-$ decay in the fiducial region. The L2C was made of 5 subsystems: the Coordinate Builders (CB), the Event Builder and Dispatcher (EBD), the Event Workers (EW), the Monitoring and Interface Service Cards (MISC) and the Control Workstation (see Fig. 29).

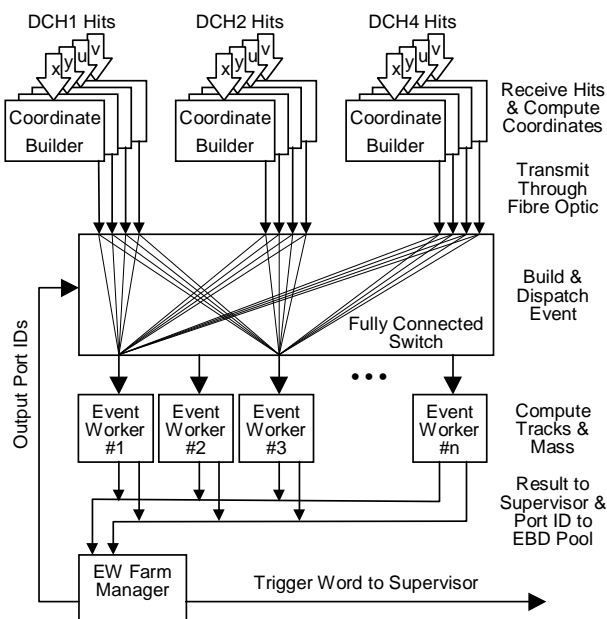


Figure 29: Block diagram of the Level 2 Charged Trigger system

There was one CB per chamber view. The function of each CB was to associate hits in the two staggered planes of one chamber view to build a list of hit coordinates. It received a block of data from the Storage Ring Cards of the corresponding two staggered planes each time a request from the Level 1 Trigger Supervisor was sent to the DCH readout system. For each wire hit, the wire number and the hit time were also sent when this hit time was in a window around the event time (typically -50 ns to +150 ns).

The data were processed through a 40 MHz pipelined algorithm implemented on firmware (Xilinx). For every hit, the drift time was computed, taking into account a possible offset for each group of wires attached to a single preamplifier card. Neighbouring hits of the staggered planes were matched to form coordinates. The drift distance was extracted from a Look Up Table taking as input the 2 drift times and the coordinate was obtained using the wire numbers. The list of reconstructed coordinates was sent to the Event Builder and Dispatcher through an optical link. The typical processing latency for one event was 1.2 μ s.

The Event Builder and Dispatcher (EBD) packed together the twelve lists of coordinates of a given event and sent them to an available Event Worker (see below). Like the Coordinate Builders, the Event Builder and Dispatcher (EBD) was a homemade set of 4 VME modules based on firmware. The role of the EBD was to pack together the coordinate lists coming from the Coordinate Builders and to send them to an available Event Worker. This was a programmable fully-connected switch with 16 input ports, 16 output ports and 16 internal data channels.

For each event, the EBD received the token of the next available Event Worker. Then data packets from the active input ports were read, serialised, and sent to the specified output port. Up to 16 transfers could be active in the switch: as soon as the packet of the first input port was read, the next packet for the following event could be read while other ports were read successively - and so forth for each port.

An Event Worker received the twelve sets of coordinates from the EBD. It associated coordinates in the different views of one chamber to produce a list of space points for each chamber. All combinations of space points in chambers 1 and 2 were tried to pairs of tracks originating from the decay region. For each such pair, space points in chamber 4 were used to look for a topology leading to an invariant mass close to the kaon mass when the tracks crossing the spectrometer were assumed to be due to a $K^0 \rightarrow \pi^+\pi^-$ decay.

The z position at which the tracks were closest was determined. If the closest distance of the tracks was less than 5 cm and if the z position was within the fiducial region (between -1 m and +48 m from the K_S target), this pair of tracks was used for the calculation of the $\pi^+\pi^-$ invariant mass and of the proper time of flight from the beginning of the decay region.

To do this both tracks were extrapolated in the y projection to the position of chamber 4; if a space point was found with a y coordinate close (within 5 mm) to this extrapolation, the x deviation of the track was computed and the momentum of the particle was obtained, assuming a uniform vertical magnetic field; if both tracks could be extrapolated to a space point in chamber 4, the invariant mass was obtained with a simple formula using the small angle approximation. If this mass was within ± 25 MeV/ c^2 of the K^0 mass, the proper time of flight between the beginning of the decay region and the decay vertex was computed. An event was flagged as a $K^0 \rightarrow \pi^+\pi^-$ candidate if this proper time was less than 4.5 times the K_S lifetime. The mass resolution for reconstructed K_L and K_S is shown in Fig. 30

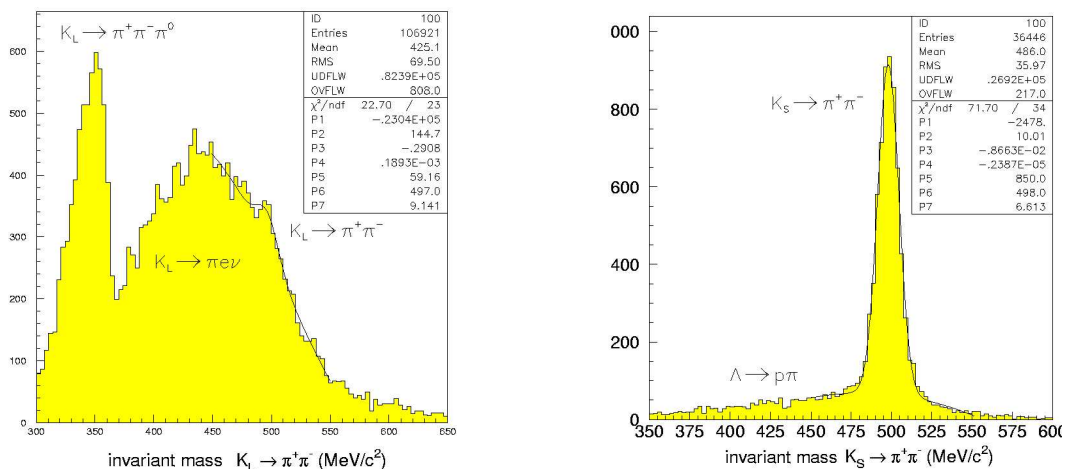


Figure 30: K_L and K_S mass resolutions from the Charged Trigger System

Once the algorithm was finished, the EW issued a trigger word which contained the time stamp at which the L1 request was issued and summarized basic properties of the event. This word was sent to the Monitoring and Interface Service Cards. Finally, the EW wrote summary files for a subset of events for monitoring and debugging purposes. These files were taken and read by the Control Workstation and/or appended to the burst summary recorded on tape for offline analysis.

The Monitoring and Interface Service Cards (MISC) received the trigger words from all the EWs and transmitted them to the general Trigger Supervisor of the experiment. In addition, it limited the input rate from the Level 1 Trigger Supervisor by using signals indicating the various queue levels in the Drift Chamber Readout and L2C chains. The MISC also received the clock signals of the experiment and allowed debugging and emulation of the L2C.

The mean efficiency of the L2C was found to vary with time, going from about 87 % at the beginning of the run to about 94 % at the end, after proper tuning of the relative time offsets of the different trigger signals and the Look Up Tables in the Coordinate Builders and after increasing the time allocated to issue the decision. The inefficiency was higher when one of the tracks was close to the beam pipe. This was due to higher accidental activity in the chambers around their central hole.

7.3 Trigger for neutral particles




Figure 31: Block diagram of the Neutral Trigger System

The Neutral Trigger [46, 47] performed computations based on the distribution of energy deposition in the liquid krypton calorimeter. The design was optimized to trigger efficiently on $2\pi^0$ candidates in the fiducial volume having total energy above 70 GeV. The majority of the background came from $3\pi^0$ events, partially contained events, low energy events and electron tracks from K_{e3} decays. Dead time was kept to an extremely low level by employing a pipelined design which performed the complete Neutral Trigger calculation every 25 ns. Fig. 31 shows the Neutral Trigger chain flow.

The Neutral Trigger reconstructed the following quantities to be used for trigger decisions: the energy of the kaon (to 3% of its value), the centre of gravity radius (corre-

sponding to transverse momentum), the number of photons and their arrival times (to 3 ns) and finally the estimated longitudinal decay position (to 3 m over a distance of 100 m). The decay position was obtained, as explained in detail below, from the first and second moments of the energy deposition by the photons, assuming the kaon mass value for their invariant mass. The proper decay time τ was also derived from the kaon energy and decay position.

The energy deposited in the 13248 cells of the calorimeter was summed up in two distinct steps to form 64 strips in each of the x and y projection of the calorimeter. The first part of the summation was performed on analogue signals from 16 single calorimeter cells to form the total energy of a supercell of 2×8 cells. This was done independently for the two projections. The calorimeter cells were individually equipped with programmable attenuators to allow the gains to be equalized. These signals were then digitized using 40 MHz 10 bit flash ADCs. From this point on, the signals entered the pipeline and were treated digitally.

The supercell information had the pedestal subtracted and was filtered before entering the second part of the summation process. The filter operated independently on each supercell: it consisted of a gate which passed through signals in the proximity of energy depositions (photon candidates) in the calorimeter. This excluded regions where there was no activity which would otherwise substantially have increased the noise level. Signals were also passed through in the time slices immediately surrounding the maximum of a pulse. The supercells which were selected in this way were passed to the second part of the summation, which digitally produced the twelve bit sums $\sum E_i^x$ and $\sum E_i^y$ for the 64 strips in each projection, using an engine based on a specific ASIC [48]. The following sections of the pipelined hardware reconstructed the number of photons, their arrival times, energies, first and second moments of the transverse energy distributions. From this information and imposing the kaon mass value to the overall invariant mass of the photons, the distance of the decay from the calorimeter was obtained and from it the proper lifetime of the K^0 after entering the fiducial region as well as its transverse momentum relative to the beam axis.

While the experiment was collecting data, the trigger was operated to require a minimum energy E of 50 GeV (the offline minimum energy cut was 70 GeV) and a maximum proper time of 4.5 K_S lifetimes (the offline maximum was 3.5). Events were rejected if the number of peaks seen in either view was six or greater. To avoid loss of events due to accidental activity, this cut was only made if all reconstructed peak times lay within 9.375 ns and there were no other peaks within approximately 150 ns of the event time. A weak centre of gravity cut requiring a radius < 15 cm was also applied.

The performance of the Neutral Trigger was monitored by reading out the calorimeter strip sums in a PMB readout (see Section 5.6). The intermediate reconstruction computations were also read out for monitoring. The efficiency of the trigger was monitored using downscaled neutral hodoscope triggers, which were generated independently of the Neutral Trigger hardware. The $2\pi^0$ trigger efficiency was measured using a control sample of events triggered by the Neutral Hodoscope. The efficiency was found to be 99.9 %, quite stable with time and with negligible difference between K_S and K_L events.

7.4 Trigger supervisor

The tasks of collecting, formatting, time-stamping and delivering trigger requests to individual Read Out Controllers were performed by the Level 1 and Level 2 Trigger Supervisors.

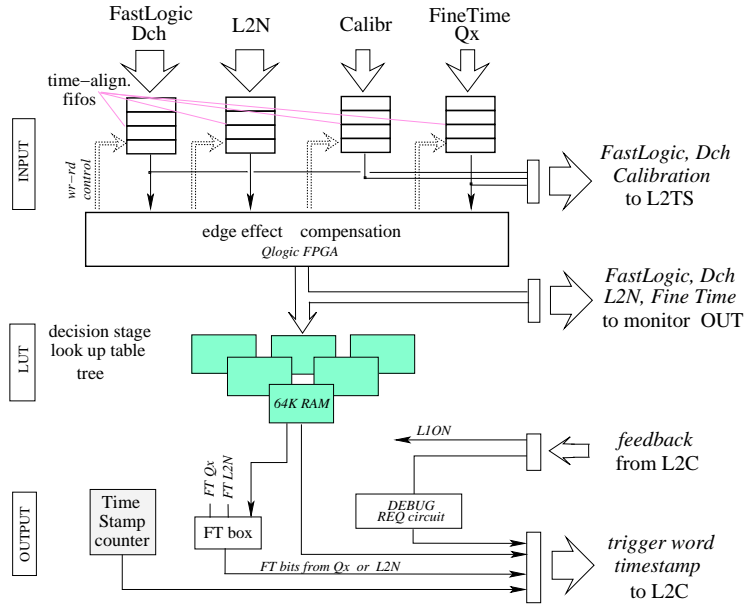


Figure 32: The L1 Trigger Supervisor architecture

The purpose of the Level 1 Trigger Supervisor (L1TS) was to provide suitably synchronised and preselected trigger requests to the L2C (Fig. 32). The signal sources were the charged hodoscope, the AKL anticounters, the hadron calorimeter and the muon veto. Signals coming from these sources, as well as several calibration pulsers, were first clock-synchronised and then fed to the Level 1 input, together with several lines coming from the Neutral Trigger and from the Drift Chamber Multiplicity logic. After being time-aligned and widened in order to overcome coincidence inefficiency at the pulse edges, all the signals were passed through a look-up table tree, the output of which was a three bit trigger code. This was sent to the drift chamber readout controller together with a 13-bit time stamp, in order to enable event data to be retrieved from front end buffers of the drift chambers. Data were then transferred to the L2C input queue. Two more time stamp bits (the Fine Time), obtained by a special high speed logic from the hodoscope signals, were also provided in order to have a better time measurement and to reduce the L2C computation time. Dead time was easily recorded and measured with a simple XON/XOFF protocol regulating the flow of requests according to the status of the L2C input queue.

The Level 2 Trigger Supervisor [49] collected clock-synchronised trigger signals from the L1TS, the L2C, the Neutral Trigger and other sources, computed time stamps and trigger codes, and finally distributed trigger requests to all the Readout Controllers. Its architecture allowed for both synchronous and asynchronous input trigger sources. Fig. 33 shows a schematic view of the Trigger Supervisor system.

For each of the four input sections, data signals (up to 24 per section) were written to port 1 of a 8K deep dual port fast RAM, of which the 13-bit address (Low Time Stamp, or LTS) was obtained either from an internal 40 MHz, 30-bit counter (Internal Time Stamp Counter, or ITSC), or from the data source itself. The remaining 17-bit High Time Stamp (HTS) was stored in parallel in a second memory at the same LTS address. Data and HTS were synchronously extracted from port 2, addressed by another 40 MHz, 30-bit counter (Master Time Stamp Counter, or MTSC), which, by appropriate presetting, ran late by some large amount (like 4000 clock periods) with respect to ITSC.




Figure 33: Block diagram of the Trigger Supervisor

This delay was loaded after each SPS burst, and was identical for all the four subdetector boards. The difference between ITSC and MTSC presets approximately represented the delay between the time the trigger data had arrived and the time the trigger decision would be taken.

The HTS extracted from the second memory was compared to the HTS from the MTSC. Whenever a match was found, the extracted data and HTS, as well as the MTSC LTS, were allowed to pass, since this meant that the data had not previously been extracted and therefore came from a new trigger.

After extraction, a special circuit individually corrected each data set for edge effects. Finally, in the time interval between two bursts, the memory was cleared. Since the MTSC presetting was the same for all the subdetector boards, by proper individual ITSC presetting, all synchronous subdetector boards could be aligned among themselves, as was required by the need to make coincidences, vetoes etc., as well as to obtain trigger-independent offset settings in all subdetector readout systems. Up to 96 input data bits from the 4 subdetector boards could be combined into a final 16-bit trigger word. Trigger word formation was obtained with extended flexibility by use of a large Look Up Table (LUT) tree, made by a combination of a Field Programmable Gate Array and a RAM. Each decoded bit of the trigger word could be software enabled and downscaled by a factor 1 to 64K, and was counted as well. The LUT logic also received up to 12 XOFF maskable, internally OR-ed signals from as many subdetectors, the purpose of which was to stop L2TS sending triggers when readout problems were found by at least one subdetector. Trigger words and time stamps generated by the LUT tree, after gating by XOFF, were written into the derandomiser stage of the chain. This behaved like a FIFO with programmable depth, with the purpose of buffering fast trigger sequences while guaranteeing a programmable Minimum Time Interval between triggers issued to the experiment. The FIFO, behaving like a fall-through with negligible delay (25 ns) between Write and Read, could be programmed to any depth between 1 and 16.

Trigger word and time stamp were finally transmitted to the trigger distribution section of the L2TS. There, a sequential event number was generated and added to the trigger word and time stamp. The resulting trigger packet, 64 bit long, was sent by the transmitter as a high speed serial stream to ten different Read Out Controllers, either over coaxial cable or optical fibre. The transmission time overhead was about 900 ns.

8 Data flow

8.1 Data Acquisition System

The NA48 data acquisition (DAQ) system was organized as a series of data-push links connected to a PC farm which merged event fragments from the various subdetectors into complete events using a switched Fast Ethernet network [50, 51, 52]. In the farm, data were then stored on local disks, from where they were transferred to the CERN computer centre. There third-level filtering, tape archival and online reconstruction were performed. The general architecture of the system is shown in Figure 34.

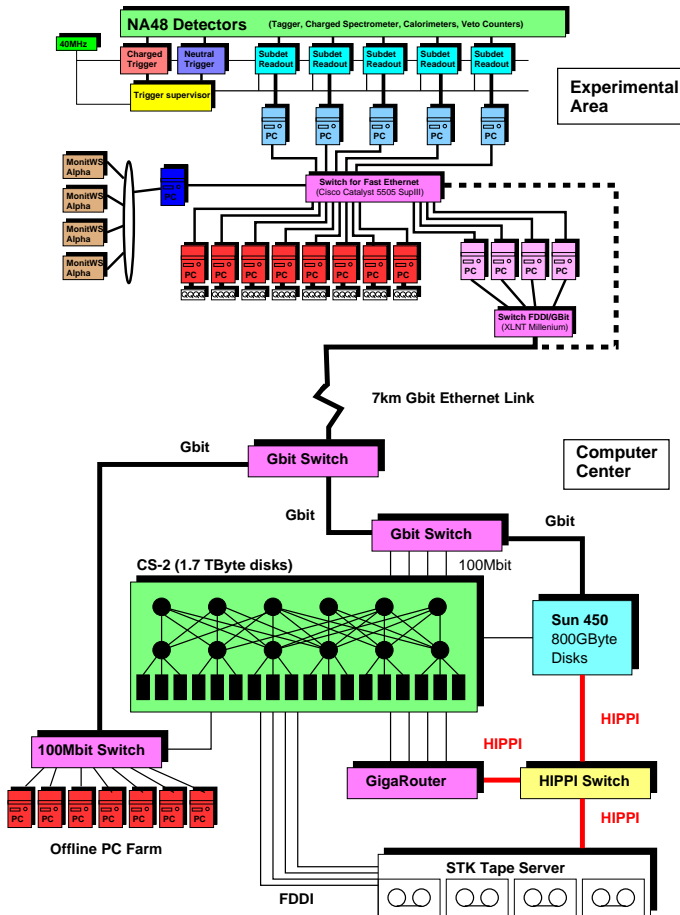


Figure 34: NA48 data flow

An important feature of the experiment was that all channels were read out continuously, synchronized by the global clock. The trigger circuitry detected desired events and issued a trigger, which resulted in a particular time window being selected. The read-out controller (ROC) of each subdetector selected data pertaining to that time window and passed it in form of an event fragment to its corresponding subdetector PC. Data rate control was achieved by an XOFF signal which was propagated back to the Level-2 Trigger supervisor. Each subdetector readout system was controlled by a subdetector computer which configured and monitored the ROCs but did not directly participate in data flow. Since core data flow only handled fast data transport, traditional tasks like detector control, configuration management and error reporting were handled by additional

subsystems which were loosely coupled.

During a SPS spill all subdetectors sent their event fragment streams into their subdetector PCs (SDPC) via DT16 links¹¹⁾, with a peak rate of up to 150 MByte/s through a DT16/PCI interface to the SDPC RAM [53]. In the interburst gap the SDPCs distributed their event fragments through a switched 100 MBit/s ethernet fabric to a group of event building PCs (EBPC)¹²⁾, which assembled complete events and wrote data out to local disks, randomizing the data flow into a continuous data stream. From the disks, data were transferred to the computer centre through four additional PCs acting as routers to distribute the traffic over four FDDI links, which were then switched into a single 7 km long Gigabit ethernet fiber link connecting the experimental area to the CERN computer centre.

As soon as a burst fragment file had been successfully transferred, it was deleted from the EBPCs disk buffer. In the computer centre the burst fragment files were combined into burst files, processed by a software filter and fed to the online reconstruction programme. Eventually, raw or filtered burst files and the results of the reconstruction were written to tape using the tape robot in the computer centre.

8.2 The Run Control and Online Data Base

The independent Run Control system was responsible for controlling the subdetectors in a coherent way and providing the necessary information to group data into periods of data taking under constant physical conditions ("Runs"). The online data base managed the sub-detector and trigger configurations and kept track of which configuration was used for which time periods. A simple data collection system gathered out-of-band data from the beam instrumentation, summary scalers and equipment that did not participate in data flow and sent them along with the burst data. Fast online monitoring allowed subsamples of full events extracted from the PC farm to be analysed and gave fast feedback on data quality and detector performance. Error messages from all subsystems were routed, displayed and archived by an error logger based on EMU [54].

Any system in the NA48 readout was required to have a control computer¹³⁾ connected to the network of the experiment. Its tasks were to initialize, configure, start and monitor the readout system of the subdetector; normally it did not take part in the data taking itself, which was handled by hardware and autonomous processors.

The Run Control system viewed the experiment as a set of finite state machines which were controlled by the run control programme (RCP). The participating computers communicated via the ControlHost message-passing package [55]. For each state transition the RCP sent the respective command to all subscribed computers. On these computers a standardized "harness" programme interpreted the commands and executed user tasks that corresponded to the required transitions. The user tasks reported success or failure to the harness via an exit code, which was then forwarded to the RCP. The RCP considered a state transition as globally successful when all subscribed systems had reported successful completion.

8.3 Central Data Recording

The online workstation farm was directly connected through fibers with a computer system in the computer centre. During the run period the data immediately flowed from

¹¹⁾ DT16 was a 16 bit wide parallel ECL bus with a maximum transfer rate of 30 MB/s.

¹²⁾ All the PCs were Intel-based industry-standard computers running the Linux operating system

¹³⁾ typical VME single board computer running OS-9

the online system to the parallel file system at the receiving end, a MEIKO CS-2 parallel computer. The transfer and the file system were optimized to improve the performance. The average transfer speed during the run in 1997 was about 10 MB/s per stream. The maximum aggregate measured via one FDDI link was between 8 and 10 MB/s.

It should be noted that, throughout the lifetime of the experiment, the equipment used both in the experimental area and in the CERN Computer Centre evolved following the available technology, but the architecture of the NA48 Dataflow remained the same.

After the successful arrival of the data on the CS-2, autonomous daemons started the transfer of the data onto magnetic tape (50 GB STK tapes in 1997). In addition, reconstruction and filter programs were started.

The disk buffers available in 1997 were about 24 GByte for the online system and about 800 Gbytes on the CS-2. Files were only deleted inside the diskbuffers if the space became low and the corresponding files had been saved onto tape. Before the data were stored, they were classified and filtered using the information from the reconstruction of charged and neutral particles in a software "level 3" filter and real-time reconstruction, running in a distributed way on 60-80 nodes of the CS2 parallel processor computer.

The action of the programme consisted of reading a set of raw data files, one for each burst, calling a set of decoding and reconstruction routines, performing a series of cuts and splitting or copying events to different output streams. The parameters of the routines, cuts, and output streams were dynamically chosen based on the properties of the event, e.g. the trigger information. They could be modified as the event was being processed. In this way, only the routines needed to classify the event were called. The information on the desired handling of the events as well as some general parameters to control the programme was stored in the online database using an ASCII configuration file.

The level 3 trigger acted as a filter using coarse and fast selection algorithms to reject quickly some classes of unwanted events. The purpose was to decrease the amount of data that was stored and that had to be handled in future reconstruction passes. It also loosens the time restrictions on the full real-time reconstruction. For each cut, events were appropriately down-scaled in order to measure offline the efficiency of the cut.

The real-time reconstruction used the standard off-line routines with calibration constants set to the best-known values at that time. Events that fulfilled tight criteria for the $\pi\pi$ final state and rare kaon decays were selected and written to disk in a separate stream for easier handling. In order to optimize the speed of the programme, the selection routines were generally based on the information of the trigger, the drift chambers and the liquid krypton calorimeter only. These selected events, as well as a fraction of all events passing the level-3 filter, were then subject to the full reconstruction of all detector parts. The reconstructed information was stored in a compact data structure (COMPACT) and written to disk in yet another stream.

The data from this stream were processed for the purposes of on-line monitoring of the data taking (hit maps for detector efficiency, stability of time offsets, etc.) and of the physical parameters (kaon mass, energy, decay vertex distributions, as well as the stability of the average values of many other parameters). More than 200 histograms for each run were available for inspection both in the control room and outside. About 30 of them were selected for permanent display in the control room as indicators of successful current data taking.

9 Acknowledgements

Our special thanks are due to Erik Uggerhoj and S.P. Moller of Aarhus University, Denmark. They provided the silicon and iridium crystals and, together with colleagues of the Aarhus-CERN-Strasbourg 'Bent Crystal Collaboration', fostered our understanding of particle channeling in crystals, which found a unique application in the simultaneous $K_L + K_S$ beams.

We are indebted to many technical staff in CERN, who contributed to the construction, installation and control of the beam line and its components. Patrice Pierre was foremost among these and we wish to dedicate the simultaneous $K_L + K_S$ beams to his memory. Further we would especially like to thank G. Dubail, C. Ferrari, B. Tomat, V. De Jesus, R. Allegrini, B. Chauchaix, A. Bonifas (planning, installation, vacuum and control), K.D. Lohmann, C. Beugnet, G.P. Ferri (beam instrumentation), S. Peraire (targets), G. Arduini (SPS operation), M. Chambardon (Survey), M. Mathieu, R. Rey (the *Kevalar* window) and G. Stevenson (radio-protection). G.P. Grillet was responsible for the adaption of the spectrometer magnet to the experiment. For the design and the construction of the LKr cryostat thanks should be given to J. Armand, G. Fersurella, A. Goncalves Martins de Oliveira, G. Juban, P. Lecossec, Y. Lesenechal, R. Truhan, D. Voillat, P. Wicht and W. Wilkens. B. Bel and Y. Camp supported the collaboration in organizational and computing matters.

We would like to acknowledge the contributions to the construction of the LKr calorimeter of M. Boldini, G. Molinari, R. Pilastrini, L. Maselli, S. Bigoni, M. Melchiorri, L. Landi, A. Magnani, L. Milano, F. Evangelisti and G. Laurenti from INFN Ferrara and Bologna, C. Avanzini, A. Basti, A. Bechini, L. Berretta, A. Del Colletto, S. Galeotti, F. Gherarducci, M. Favati, M. Iacoponi, V. Merlin, F. Profeti, D. Rizzi, P. Salutini, A. Tazzioli, S. Tolaini and L. Zaccarelli from INFN Pisa, D. Aisa, E. Babucci, C. Campeggi, L. Farnesini and A. Piluso from INFN Perugia and M. Tuveri from INFN Cagliari.

A special expression of thanks should be given to the memory of Gabriele Bonora from INFN Ferrara, who passed away in 1999, of Gianfranco Pagani from INFN Pisa and of Giuliano Abbrugiati of INFN Torino, who left us in 2003.

Thanks should be given to V. Radeka for the help in the design of the LKr cold preamplifier and to S. Rescia for the support in the implementation of the preamplifier circuits; we acknowledge the support of F. Bourgeois and P. Ponting in the design and construction phase of the CPD system.

We would like to thank M. Bizzarri and A. Alaimo from INFN Perugia for their respective development of the PMB system and the setting up of the pretrigger system.

INFN Ferrara acknowledges and thanks the work of S. Chiozzi and R. Malaguti for their determinant contribution in all the phases of the construction of the new Drift Chamber readout for NA48/1.

INFN Cagliari would like to acknowledge the contributions of M. Arba and M. Tuveri to the charged hodoscope mechanics.

We thank A. Fletcher and K. Baker from Cambridge University for their work on muon veto mechanics and electronics respectively.

References

- [1] G.D. Barr et al., CERN, Edinburgh, Mainz, Perugia, Pisa, Saclay, Siegen, Torino, Vienna Collaboration, "Proposal for a Precision Measurement of $Re(\epsilon'/\epsilon)$ in CP-violating $K_0 \rightarrow 2\pi$ Decays", CERN/SPSC/90-22, SPSC/P253, 20 July 1990.
- [2] V. Fanti et al., NA48 Collaboration, Phys. Lett. B **465**, (1999) 335-348.

- [3] A. Lai et al., Eur. Phys. J. C **22**, 231 (2001)
- [4] J.R. Batley et al., Phys. Lett. B **544**, 97 (2002)
- [5] C. Biino, N. Doble, L. Gatignon, P. Grafström and H. Wahl [NA48 Collaboration], CERN-SL-98-033-EA
- [6] N. Doble, L. Gatignon, P. Grafström, Nucl. Instr. and Meth. B **119**, (1996) 181.
- [7] R. Batley et al., Cambridge, CERN, Dubna, Edinburgh, Ferrara, Firenze, Mainz, Perugia, Pisa, Saclay, Siegen, Torino, Warsaw, Wien Collaboration. Addendum2 to P253. A high sensitivity investigation of K_S and neutral hyperon decays using a modified K_S beam. CERN/SPSC 2000-002, SPSC/P253 ADD.2, December 10, 1999.
- [8] R. Batley et al., Cambridge, CERN, Dubna, Edinburgh, Ferrara, Firenze, Mainz, Perugia, Pisa, Saclay, Siegen, Torino, Warsaw, Wien Collaboration. Addendum3 to P253. A precision measurement of charged kaon decay parameters with an extended NA48 setup. CERN/SPSC 2000-003, SPSC/P253 ADD.3, January 16, 2000.
- [9] E. Griesmyer et al., Nucl. Instr. and Meth. A **361**, (1995) 466
- [10] G. Arduini, CERN-PPE-92-025
- [11] Vector Fields Ltd, The OPERA 3D Reference Manual, Oxford OX5 1JE, England (1992).
- [12] D. Bederede *et al.*, Nucl. Instrum. Meth. A **367**, 88 (1995).
- [13] I. Augustin *et al.*, Nucl. Instrum. Meth. A **403**, 472 (1998).
- [14] R. Arcidiacono *et al.*, IEEE Trans. Nucl. Sci. **51**, 1470 (2004).
- [15] R. Arcidiacono *et al.*, Nucl. Instrum. Meth. A **518**, 493 (2004).
- [16] J. Alexander, "Eurogam Project: DT32-Bus Specification", 15-10-1998, Nuclear Physics Support Group, CLRC, Daresbury Laboratories, Daresbury, Warrington, Cheshire, UK WA4 4AD, United Kingdom.
<http://npg.dl.ac.uk/Eurogam/documents/edoc075/edoc075.ps>
- [17] TDC F1 Functional Description, Scientific Version, 14-12-2001, Acam-Messelectronic GMBH.
http://www.acam.de/Documents/English/DB.F1sc_e.pdf
- [18] G. Braun, H. Fischer, J. Franz, A. Grünemaier, F. H. Heinsius, K. Königsmann et al., "Draft Standard for CATCH Mezzanine Cards: CATCH/OPEN".
<http://hpfr02.physik.uni-freiburg.de/projects/compass/electronics/notes/2001-7.ps.gz>
- [19] J. Bremer *et al.*, CERN-LHC-98-7-ECR *Given at 17th International Conference on Cryogenic Engineering (ICEC 17), Bornemouth, England, 14-17 Jul 1998*
- [20] G.D. Barr et al., NIM A **370** (1996) 413.
- [21] B. Aubert et al., CERN/DRDC/90-31 DRDC/P5.
- [22] C. Cerri, *Prepared for 6th International Conference on Calorimetry in High-energy Physics (ICCHEP 96), Rome, Italy, 8-14 Jun 1996*
- [23] C. de la Taille, G. Martin-Chassard and N. Seguin-Moreau, LAL-RT-97-05 *Prepared for 6th International Conference on Calorimetry in High-energy Physics (ICCHEP 96), Rome, Italy, 8-14 Jun 1996*
- [24] Hallgren B. et al., The NA48 LKr Calorimeter Digitizer Electronics Chain, presented at the Wirechamber Conference 1998, Vienna, Austria, February 23-27, 1998 / CERNEP/98-48.
- [25] O. Vossnack, *Given at 7th International Conference on Calorimetry in High-Energy Physics (ICCHEP 97), Tucson, AZ, 9-14 Nov 1997*
- [26] R. Fantechi [NA48 Collaboration], *Given at 2nd Workshop on Electronics for LHC Experiments, Balatonfured, Hungary, 23-27 Sep 1996*

- [27] Z. Guzik, A. Chlopik, F. Formenti, A. Gianoli and O. Vossnack, Nucl. Instrum. Meth. A **427**, 574 (1999).
- [28] A. Gianoli *et al.* [NA48 Collaboration], IEEE Trans. Nucl. Sci. **47**, 136 (2000).
- [29] W.E. Cleland, E.G Stern, Nucl. Instrum. Meth. A **338**, 467 (1994).
- [30] G. Unal, NA48 Note 98/01
- [31] S. Palestini *et al.*, Nucl. Instrum. Meth. A **421**, 75 (1999).
- [32] S. Crépe, Nucl. Instrum. Meth. A **409**, 575 (1998).
- [33] G. D. Barr *et al.* [NA48 Collaboration], Nucl. Instrum. Meth. A **370**, 413 (1996).
- [34] G. Unal [NA48 Collaboration], arXiv:hep-ex/0012011, Proceedings of the IX International Conference on Calorimeter in High Energy Physics, Annecy, October 9-14, 2000.
- [35] M. Jeitler [NA48 Collaboration], Nucl. Instrum. Meth. A **494**, 373 (2002).
- [36] G. Diambri Palazzi, Rev. Mod. Phys. **40**, (1968) 611.
- [37] R. Moore *et al.*, NIM B **119**, (1996) 149.
- [38] A. Lai, L. Musa, IEEE Trans. Nucl. Sci. **42**, (812) 1995.
- [39] M.D. Needham, RAL Thesis, RAL-TH-1999-009 (1999).
- [40] T. Gershon, PhD Thesis, RAL-TH-2000-006 (2000).
- [41] T. Beier *et al.*, Nucl. Instr. and Meth., A **344**, (1994) 149-155; P. Grafström *et al.*, Nucl. Instr. and Meth., A **344**, (1994) 487-491; T. Beier *et al.*, Nucl. Instr. and Meth., A **360**, (1995) 390-394.
- [42] H. Bergauer *et al.*, Nucl. Instr. and Meth., A **373**, (1996) 213-222; H. Bergauer *et al.*, Nucl. Instr. and Meth., A **419**, (1998) 623-631.
- [43] M. Jeitler *et al.*, NA48 - NOTE 95-14. M. Jeitler *et al.*, NIM A **400**, (1997) 101.
- [44] S. Anvar *et al.*, Nucl. Instrum. Meth. A **419**, (1998) 686.
- [45] S. Anvar *et al.*, Proceedings of the IEEE NSS/MIC 98, Toronto (Canada), November 11-14, 1998.
- [46] G. Barr *et al.*, Nucl. Instrum. Meth. A **485**, (2002) 676.
- [47] G. Fischer *et al.*, Nucl. Instrum. Meth. A **419**, (1998) 695.
- [48] C. Avanzini *et al.*, IEEE Trans. Nucl. Sci. **43**, (1996) 1789.
- [49] R. Arcidiacono, P. L. Barberis, F. Benotto, F. Bertolino, G. Govi and E. Menichetti, Nucl. Instrum. Meth. A **443**, (2000) 20.
- [50] M. Wittgen, Nucl. Instrum. Meth. A **461**, (2001) 478.
- [51] A. Peters [NA48 Collaboration], The NA48 online and offline PC farms, presented at International Conference on Computing in High-Energy Physics and Nuclear Physics (CHEP 2000), Padova, Italy, 7-11 Feb 2000
- [52] M. Wittgen *et al.*, The NA48 event building PC farm, IEEE Trans. Nucl. Sci. **47**, (2000) 348.
- [53] Z. Guzik, A. Chlopik, F. Bal, F. Formenti and A. Lacourt, RISC mezzanines for controlling data acquisition in the NA48 experiment at CERN, Nucl. Instrum. Meth. A **452**, (2000) 289.
- [54] <http://emu.web.cern.ch/emu>
- [55] ControlHost - A Distributed Data Handling Package. User guide and reference. Copyright 1993-1995 CASPUR Consortium, Rome, Italy.