Technical Design Document

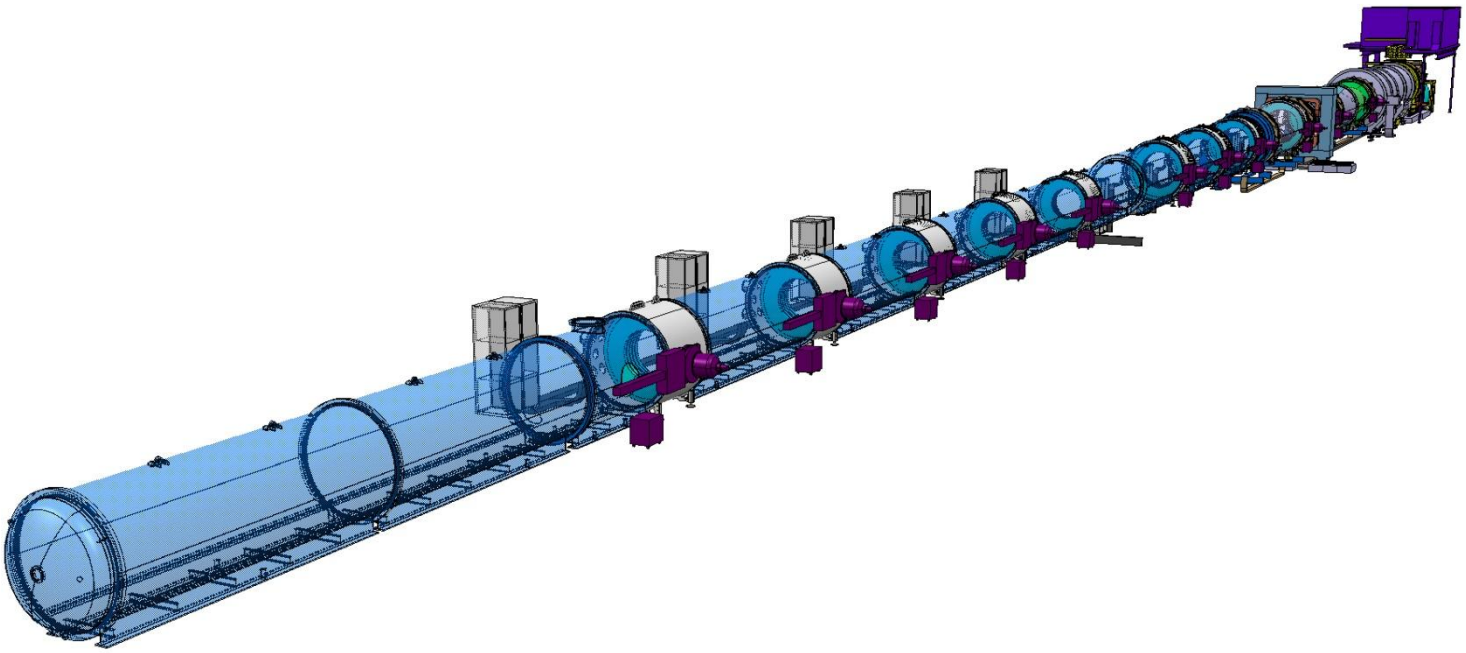

Document N°:	NA62-10-07
Last Modified:	December 2010
Prepared by:	NA62 Collaboration

Chapter Editors:

F. Ambrosino, A. Ceccucci, H. Danielsson, N. Doble, R. Fantechi, A. Kluge, C. Lazzeroni, M. Lenti, G. Ruggiero, M. Sozzi, P. Valente, R. Wanke.

Main Editor: F. Hahn

Acknowledgments:

NA62 collaboration wishes to thank all the technical staff involved in the design, testing and prototyping for their invaluable contributions.

We wish to express our thanks to the SPSC referees Lutz Feld and Christofas Touramanis, for their constructive criticism, advice and support.

Special thanks go to Andrea Fantechi for the artist view of the experiment on the cover page.

We acknowledge Valerio Bocci, head of electronics laboratory of INFN Rome (LABE).

The editors wish to thank warmly Veronique Wedlake for the great care and dedication she has shown in the editing and the formatting of the files.

And, special thanks to many colleagues for meticulous proof-reading.

NA62 Collaboration:

Belgium, Louvain-La-Neuve
Université Catholique de Louvain
Eduardo CORTINA GIL, Maria Elena MARTIN ALBARRAN, Georg Norbert NUESSLE, Sergey PODOLSKY

Bulgaria, Sofia
University of Sofia St. Kliment Ohridski
Faculty of Physics
Venelin KOZHUHAROV, Leandar LITOV

Germany, Mainz
Johannes-Gutenberg-Universitat Mainz
Volker BUESCHER, Manuel HITA HOCHGESAND, Konrad KLEINKNECHT, David LOMIDZE, Burkhard RENK, Rainer WANKE, Andreas WINHART

Italy, Ferrara
Università di Ferrara e INFN
Vittore CARASSITI, Stefano CHIOZZI, Angelo COTTA RAMUSINO, Pietro DALPIAZ, Federico EVANGELISTI, Alberto GIANOLI, Luca LANDI, Roberto MALAGUTI, Michele MELCHIORRI, Alan NORTON, Ferruccio PETRUCCI, Marco STATERA, Stefano SQUERZANTI, Heinrich WAHL

Italy, Florence
Universita e INFN
Andrea BIZZETI, Francesca BUCCI, Mario CALVETTI, Antonio CASSESE, Enrico CELEGHINI, Roberto CIARANFI, Enrico IACOPINI, Massimo LENTI, Michele VELTRI

Italy, Frascati
Istituto Nazionale di Fisica Nucleare (INFN)
Laboratori Nazionali di Frascati (LNF)
Antonella ANTONELLI, Emilio CAPITOLO, Cesidio CAPOCCIA ,Aldo CECCHETTI, Giovanni CORRADI, Bruno DULACH ,Rosario LENCI, Matthew MOULSON,Claudio PAGLIA, Mauro RAGGI, Massimo SANTONI,Tommaso SPADARO, Diego TAGNANI, Sauro VALERI

Italy, Naples
Universita e INFN
Fabio AMBROSINO, Alfonso BOIANO, Francesco CASSESE, Domenico DI FILIPPO, Paolo MASSAROTTI, Marco NAPOLITANO, Vito PALLADINO, Giuseppe PASSEGGIO, Lorenzo ROSCILLI, Giulio SARACINO, Giuseppe VITIELLO

Italy, Perugia
Sezione INFN di Perugia and Università degli Studi di Perugia
Giuseppina ANZIVINO, Patrizia CENCI, Evelina MARINOVA, Monica PEPE, Roberto PIANDANI, Mauro PICCINI

Italy, Pisa
Sezione di Pisa (INFN)
Bruno ANGELUCCI, Carlo AVANZINI, Luca BERRETTA, Sandro BIANUCCI, Claudio CERRI, Gianmaria COLLAZUOL, Flavio COSTANTINI, Luigi DI LELLA, Niels DOBLE, Riccardo FANTECHI, Stefano GALEOTTI, Stefano GALLORINI, Sergio GIUDICI, Gianluca LAMANNA[1], Guido MAGAZZU, Italo MANNELLI, Elena PEDRESCHI, Giuseppe PIERAZZINI, Fabrizio RAFFAELLI, Giuseppe RUGGIERO, Marco SOZZI, Franco SPINELLA, Stefano VENDITTI, Luciano ZACCARELLI

Italy, Rome
INFN Sezione di Roma Tor Vergata
Universita' degli Studi di Roma Tor Vergata
Vincenzo BONAIUTO, Adolfo FUCCI, Giovanni PAOLUZZI, Andrea SALAMON, Gaetano SALINA, Emanuele SANTOVETTI, Fausto SARGENI

Italy, Rome
Universita e INFN, Roma I
Sezione di Roma I (INFN)
Giacomo CHIODI, Giulio D'AGOSTINI, Emanuele LEONARDI, Riccardo LUNADEI, Marco SERRA, Paolo VALENTE,

Italy, Turin
Universita e INFN
Cristina BIINO, Brigitte BLOCH-DEVAUX, Giulio DELLACASA, Sara GARBOLINO, Pierre JARRON, Flavio MARCHETTO, Giovanni MAZZA, Ezio MENICHETTI, Angelo RIVETTI, Richard WHEADON

Mexico, San Luis Potosi
Universidad Autónoma de San Luís Potosi
Instituto de Fisica
Jurgen ENGELFRIED

Romania, Bucharest-Magurele
Horia Hulubei National Institute of Physics and Nuclear Engineering (IFIN-HH)
Cornelia COCA

Russia, Dubna
Joint Institute for Nuclear Research (JINR)
Vladimir ELSHA, Valeri FALALEEV, Levan GLONTI, Vladimir KEKELIDZE, Yuri KIRYUSHIN, Dmitry MADIGOZHIN, Milena MISHEVA, Natalia MOLOKANOVA, Serguei MOVCHAN, Irina POLENKEVICH, Yury POTREBENIKOV, Viacheslav SAMSONOV, Sergey SHKAROVSKIY, Ilya SLEPNEV, Andrey ZINCHENKO

[1] Present address: PH Department, CERN 1211 Geneva 23 CH

Russia, Moscow
Russian Academy of Sciences
Institute for Nuclear Research (INR)
Vladimir BOLOTOV, Evgeni V. GUSHIN, Aleksey KHUDYAKOV, Andrey MAKAROV

Russia, Protvino
Institute for High Energy Physics, Russian Federation State Research Centre (IHEP)
Gennady BRITVICH, Oleg CHIKILEV, Victor KURSHETSOV, Vladimir OBRAZTSOV, Iakov POPOV, Vitaly SEMENOV, Mikhail TIKHONOV, Oleg YUSHCHENKO

Slovakia, Bratislava
Comenius University
Tomas BLAZEK, Vladimir CERNY, Roman LIETAVA

Switzerland, Geneva
Conseil Européen pour la Recherche Nucléaire (CERN)
Gianluca AGLIERI-RINELLA, Spasimir BALEV, François BUTIN, Andrea CATINACCIO, Augusto CECCUCCI, Hans Olof DANIELSSON, Eric DAVID, Neil DIXON, Massimiliano FIORINI, Laurentius GATIGNON, Roberto GUIDA, Ferdinand HAHN, Stefan HAIDER , Mark HATCH, Olivier JAMET, Jan KAPLON, Alexander KLUGE, Philippe LENOIR, Peter LICHARD, Michel MOREL, Matthew NOY, Ans PARDONS, Francisco PEREZ-GOMEZ, Lukas PERKTOLD, Petra RIEDLER, Vladimir RYJOV, Antonino SERGI, Giovanna VANDONI, Piet WERTELAERS

United Kingdom, Birmingham
University of Birmingham
Bjorn Inge HALLGREN, Karl HARRISON, Marian KRIVDA, Evgueni GOUDZOVSKY, Cristina LAZZERONI, Karim MASSRI, David John MUNDAY, Angela ROMANO, Mark SLATER, Richard STALEY, Nigel WATSON

United Kingdom, Bristol
University of Bristol
H. H. Wills Physics Laboratory
Helen F. HEATH

United Kingdom, Glasgow
University of Glasgow
David BRITTON

United Kingdom, Liverpool
University of Liverpool
Oliver Lodge Laboratory
John DAINTON, John FRY

United States of America, Fairfax
George Mason University
Philip RUBIN

United States of America, Menlo Park
SLAC National Accelerator Laboratory
David COWARD

United States of America, Merced
University of California Merced
Roland WINSTON

Contents

1 Introduction

1.1 Physics Objectives and Detector Overview

The rare decays $K \to \pi\nu\bar{\nu}$ are excellent processes to study the physics of flavour because of their very clean nature. Thanks to the hard (quadratic) GIM mechanism, these decays are dominated by short-distance dynamics. Moreover, the short-distance amplitude is governed by just one single semileptonic operator whose hadronic matrix element is measured experimentally from the semileptonic kaon decays. This insures that the main theoretical uncertainties can be eliminated by data. The strong suppression to the Standard Model (SM) contributions and the remarkable theoretical precision of the SM rate makes these decays sensitive probes to possible new degrees of freedom. The SM prediction [1] [2] including the latest NNLO QCD corrections [3], electroweak corrections [4] and long distance [5] contributions, reads:

$$\mathrm{BR}\,(K^+ \longrightarrow \pi^+\nu\bar{\nu})\,(\mathrm{SM}) = (8.5 \pm 0.7)\cdot 10^{-11},$$

where the uncertainty is dominated by the current precision of the CKM mixing matrix parameters rather than by theoretical errors affecting the hadronic matrix element. Progress on the measurements of these CKM parameters is expected to further improve this SM prediction. Predictions of this branching ratio are available for various extensions of the SM including models with a 4th generation of quarks and leptons [6], Littlest Higgs [7], Randall and Sundrum mechanism [8] and supersymmetric flavour models [9].

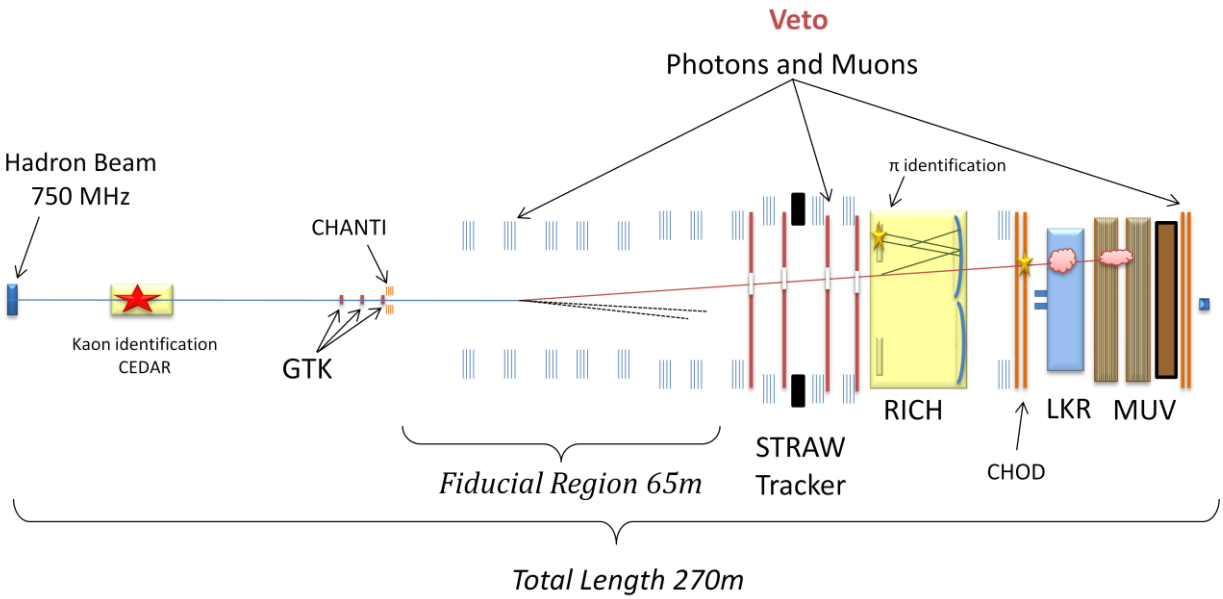

Figure 1 Schematic view of the NA62 experiment showing the main sub-detectors (not to scale).

On the experimental side, the most precise results have being obtained by the E787 and E949 experiments at BNL by studying stopped kaon decays [10]:

$$BR\ (K^+ \longrightarrow \pi^+ \nu\bar{\nu})\ (Ex) = \ (17.3\ ^{+11.5}_{-10.5}) \cdot 10^{-11}$$

The gap between the theoretical precision and the large experimental error motivates a strong experimental effort. Significant new constraints can be obtained with a measurement of the rate of this reaction at the level of 10% or better. This requirement sets the bar for the next generation of rare decay experiment.

NA62 aims to collect of the order of 100 $K^+ \to \pi^+ \nu\bar{\nu}$ events in about two years of data taking and to keep the total systematic uncertainty small. To this purpose, at least 10^{13} K^+ decays are required, assuming a 10% signal acceptance and a $K^+ \to \pi^+ \nu\bar{\nu}$ branching ratio of 10^{-10}. To keep the systematic uncertainty small requires a rejection factor for generic kaon decays of the order of 10^{12}, and the possibility to measure efficiencies and background suppression factors directly from data. Kaon intensity, signal acceptance and background suppression are, therefore, the driving criteria for the construction of the NA62 experiment [11], [12].

The CERN-SPS extraction line, already used for the NA48 experiments, can deliver the required intensity. A new beam line will provide a secondary charged hadron beam 50 times more intense than in the past, with only 30% more SPS protons on target. The NA62 experiment will be housed in the CERN North Area High Intensity Facility (NAHIF) where NA48 was located. Protons from the SPS at 400 GeV/c impinge on a beryllium target and produce a secondary charged beam. Considerations about signal acceptance drive the choice of a secondary positive beam of 75 GeV/c with 1% momentum bite and a divergence below 100 μrad, both in the x and y projections. Around 6% of secondary particles are K^+, the others π^+ and protons. The beam has a square shape of 60x27 mm^2 area and a total flux rate of about 750 MHz at the end of the beam line.

The use of a decay-in-flight technique to identify the K^+ decay products is the experimental principle of NA62 and it is well matched to the characteristics of the CERN SPS. The high momentum of the secondary beam improves background rejection and sets the longitudinal scale of the experiment. The main elements for the detection of the K^+ decay products are spread along a 170 m long region starting about ~100 m downstream of the beryllium target. Useful K^+ decays will be detected from a 65 m long fiducial region. The largest detectors have an approximately cylindrical shape and around the beam axis, whose diameter varies from 12 to 220 cm, in order to let the very intense flux of undecayed beam particles passing through. The overall rate integrated over these detectors is in the range of 10 MHz. A schematic layout of the experiment is shown in Figure 1. To achieve the required background suppression, different techniques have to be employed in combination. This allows one to measure the rejection factors from the data by inverting the cuts one at the time.

Since the neutrino-antineutrino pair is undetectable, the signature of the signal consists of a single π^+ track reconstructed downstream of the decay volume and matched to a K^+ track upstream. Timing, spatial and angular information are needed to match these two tracks. Generic K^+ decay modes can mimic a signal when just a single track is detected and the others particles escape detection. Beam particles interacting in the last station of the beam detector or in the residual gas can also contribute to backgrounds if they are accidentally matched to an incoming kaon track.

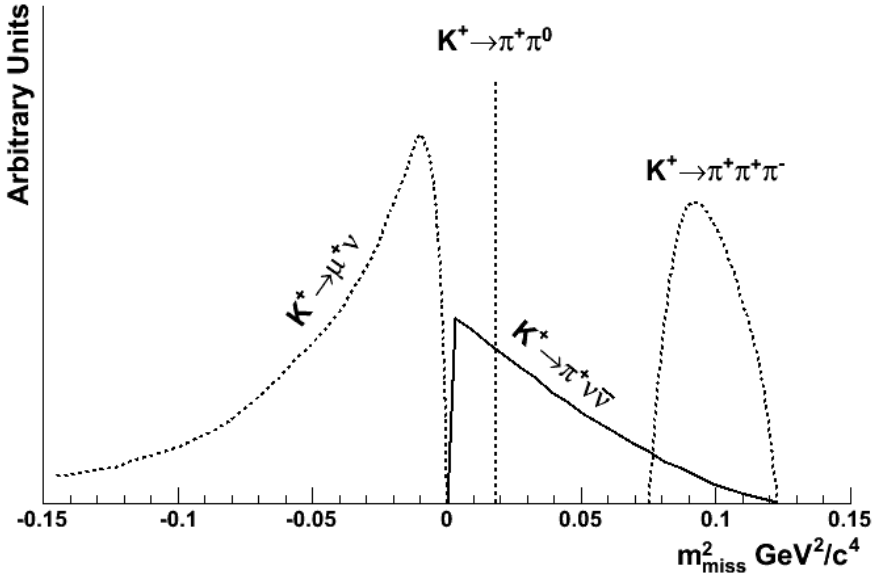

Figure 2 Shape of the $\boldsymbol{m^2_{miss} := (P_K - P_{Track})^2}$ *for signal (thick solid line) and background events under the hypothesis that the charged track is a pion. These background sources refer to decays which are kinematically constrained.*

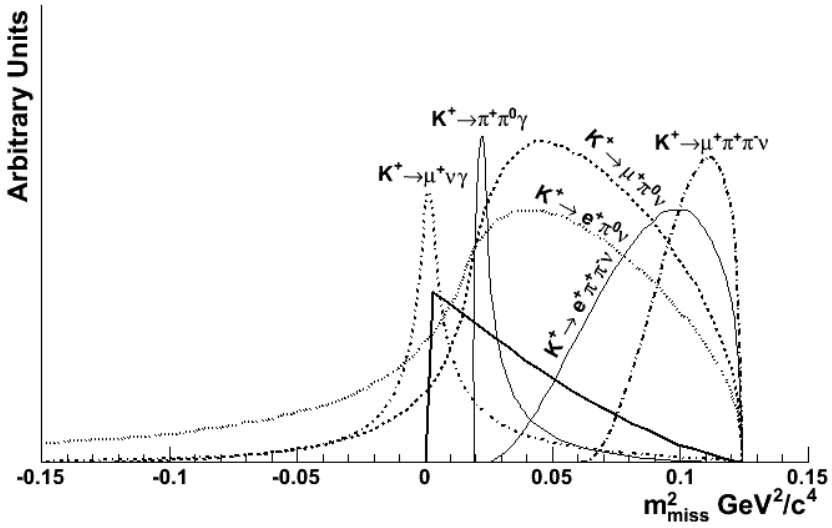

Figure 3 Shape of the $\boldsymbol{m^2_{miss} := (P_K - P_{Track})^2}$ *for signal (thick solid line) and background events under the hypothesis that the charged track is a pion. These background sources refer to decays which are kinematically not constrained.*

The experiment, therefore, needs tracking devices for both K^+ and π^+, and also calorimeters in order to veto photons, positrons and muons. In addition, particle identification systems to identify the incident kaons and to distinguish π^+ from μ^+ and e^+ must complement the tracking and veto detectors to reach the ultimate sensitivity and to guarantee redundancy. The guiding principles for the construction of the NA62 detectors are, therefore: accurate kinematic reconstruction, precise particle timing, efficiency of the vetoes and excellent particle identification.

1.2 Tracking System

The most discriminating variable to distinguish the $K^+ \to \pi^+ \nu\bar{\nu}$ signal from background is the squared missing mass: $m^2_{miss} \equiv (P_K - P_\pi)^2$. Here P_K denotes the 4-momentum of the parent particle assumed to be a kaon and P_π is the 4-momentum of the decay particle assumed to be a π^+. This variable rejects more than 90% of K^+ decays, as shown in Figure 2. There are two signal regions, one on each side of the $K^+ \to \pi^+\pi^0$ peak. One can define Region I (Region II), the signal region, in which m^2_{miss} is lower (larger) than $m^2_{\pi 0}$.

Backgrounds from kaon decays ($K^+ \to \pi^+\pi^0$, $K^+ \to \mu^+\nu$ and $K^+ \to \pi^+\pi^-\pi^0$, see Figure 2) can enter these signal regions because of kinematic resolution effects or because of the lack of a well defined kinematical constraint (e.g. kaon semi-leptonic decays as shown in Figure 3). The high branching ratios of decays like $K^+ \to \pi^+\pi^0$ and $K^+ \to \mu^+\nu$ allow the first source to compete with, or even be larger than, the second one. As a consequence the tracking requires low mass and high precision detectors. Simulations have shown that an overall material budget of around 2% radiation length, together with an intrinsic resolution on the track position of about 100 μm and a momentum resolution below 0.5%, correspond to a resolution on m^2_{miss} below 10^{-3} GeV^2/c^4. This is enough to suppress the Gaussian tails of the reconstructed two-body background from, while keeping the signal acceptance within the requirements

The tracking system consists of a kaon and a pion spectrometer. They are made of thin detectors orthogonal to the beam axis, with one or more magnets in between for momentum analysis. The different transverse size for the upstream and downstream detectors, together with the required performances, determines the technology choice. The kaon spectrometer, called Gigatracker (GTK), consists of three solid state micro-pixel stations matching the beam dimensions. These stations are placed in vacuum. The GTK is crossed by the full beam intensity (about 750 MHz) and must provide a time resolution of better than 200 ps in order to avoid the wrong association of a beam particle to the decay reconstructed in the downstream detectors. A wrong association might lead to the wrong calculation of the missing mass which in turn would decrease the kinematical discrimination against, most notably, the $K^+ \to \pi^+\pi^0$ decays. Four dipole magnets provide an achromatic (no net bending) spectrometer for particles of any momentum. The GTK provides a 0.2% RMS momentum resolution and an angular resolution of around 15 μrad on an event-by-event basis.

A pion spectrometer consisting of four straw tube chambers (STRAW) is operated directly in the vacuum tank. The first chamber is placed about 80 m downstream of the last station of the GTK. The choice to operate the STRAW chambers in vacuum minimises the multiple scattering by avoiding the use of a window to separate the decay from the tracking volume and the helium which would be otherwise required in between the chambers to operate them near atmospheric pressure. A careful technological choice for the straws can reduce the material budget down to 0.5% radiation length per chamber. The same magnet employed in the NA48 experiment will be used: it is located after the second STRAW chamber and provides a 270 MeV/c momentum kick in the horizontal (x-z) plane. This bending provides a sufficient pion momentum resolution. A region free of tubes, 12 cm wide, in the middle of the chambers forms a passage where the beam passes through undisturbed. The x-position of the beam hole in the chambers follows the 75 GeV/c positive beam path. This configuration minimizes the particle rate in the inner part of the chambers, due to the halo of the beam, and protects against backgrounds from K^+ decays, like $K^+ \to \pi^+\pi^+\pi^-$ and $K^+ \to \pi^+\pi^- l^+\nu_l$, in those cases when

only the positive pion is fully reconstructed. The position of the last chamber, 20 m downstream of the magnet, increases the lever arm to avoid those high momentum negative pions from decays like $K^+ \to \pi^+\pi^+\pi^-$ and $K^+ \to \pi^+\pi^- l^+\nu_l$ could escape detection. The spatial matching between the K^+ track, extrapolated into the decay region, and the pion track, extrapolated backward, largely protects against a miss-measurement of the pion direction due to a possible large angle scattering in the first chamber. Simulations of the whole tracking system, based on GEANT4, showed that non-Gaussian tails affecting the measurement of both the pion direction and momentum limit the ultimate rejection power of the m^2_{miss} requirement against $K^+ \to \pi^+\pi^0$ and $K^+ \to \mu^+\nu$ to about 2×10^4 and 2×10^5, respectively.

The tracking systems are also crucial against a more subtle, but equally important source of background. Beam hadronic interactions in the last station of the GTK material may cause the emission of a leading π^+ which enters the pion acceptance, while the associated multiplicity of soft particles can escape from the detector. A suitable cut on the reconstructed vertex position is crucial for the suppression of this accidental background. The limitation comes from a possible large scattering of the pion in the first plane of the straw spectrometer, which creates non-Gaussian tails and might shift the reconstructed vertex in the decay region.

1.3 Timing

The need to match the incoming K^+ and the π^+, and the high rate of particles (750 MHz) crossing the GTK detector set the timing performance of the experiment. Without precise timing, an accidental beam particle can be mistaken for the decaying kaon and be associated to the π^+. This association can spoil the kinematic resolution because the beam divergence is non-negligible. As a consequence, the kinematic rejection of the two body decays is weakened and more background enters the signal regions. With a 100 ps time resolution on the π^+ track and 150 ps time resolution for the beam tracks, the fraction of events with a wrong association can be kept to below 1% once appropriate spatial cuts (e.g. cuts based on the closest distance of approach between the pion and kaon track candidates) are applied. This is enough to suppress this source of background to the same level of the one due to non-Gaussian tails. Precise timing for the π^+ is provided by the RICH counter. As mentioned, the timing of K^+ requires a fast detector exposed to the full beam intensity (GTK). NA62 has performed [13] a full R&D study to demonstrate that excellent timing performance from micro-pixel silicon trackers can be achieved. It is also worth noting that possible inefficiencies of the Gigatracker complicate the picture of track miss-matching: because of the high rate environment, the maximum inefficiency per station allowed by the electronics and readout system must not exceed 2÷3%. Good timing performance are also required at the trigger level in order to veto the high flux of muon tracks originating from kaon and pion decays and to provide a fast Level 0 trigger.

1.4 Veto system

The kinematic rejection alone cannot provide the requested level of background suppression. An additional factor must come from vetoing photons and muons. The suppression of possible background from $K^+\to\pi^+\pi^0$ decays determines the design of the photon veto system. One requires an overall inefficiency of about 10^{-8} for the π^0 detection. It is crucial that the veto system covers an angle region from 0 to 50 mrad from the K^+ decay vertex with respect to the beam line. The NA62 design insists on the very high detection efficiency for high energy photons rather than attempting to capture photons of small energy going sideways. The reason for this approach is that in order to consistently detect photons at angles larger than 50 mrad one would need to install photon detectors along the

entire length of the vacuum tank without gaps, which would be a massive construction task. In addition, photons from kaon decays with angles larger than 50 mrad are characterised by low energies which would make their detection difficult and dependent on low energy thresholds.

The photon veto detectors are: a system of 12 Large Angle Veto (LAV) calorimeters covering an angle interval from 8.5 to 50 mrad, an electromagnetic calorimeter for the detection of photons between 1 and 8.5 mrad (LKR), and small angle calorimeters covering the region below 1 mrad (IRC and SAC). In the analysis, the π^+ momentum will be required to be less than 35 GeV/c. In this way, the momentum of the π^0 amounts to at least 40 GeV/c. Such a large energy deposit can hardly be missed in the calorimeters. The photon detection inefficiency has to be below 10^{-5} for photons in the 1-10 mrad region above 10 GeV, and less than 10^{-3} for photons above 1 GeV. These inefficiency requirements are a consequence of the angle-energy correlation of the photons from π^0 decays.

We reuse the liquid Krypton (LKR) electromagnetic calorimeter of NA48 as a veto for photons in the 1-10 mrad region. Measurements based on NA48 data have demonstrated the capability of the LKR to reach the required veto performance. These analyses allowed one to develop a method to extract the LKR efficiency directly from the data using fully reconstructed $K^+ \rightarrow \pi^+\pi^0$ decays with accuracy better than 10^{-5}, which is crucial for the control of systematics.

Since the LKR is equipped with a beam-pipe to accommodate the passage of the charged beam, an additional photon detector must cover the small angle region. For this reason, a rectangular calorimeter (SAC), built with “Shashlik” technology, is located at the end of the beam line just after a sweeping magnet that deflects the high intensity charged beam. The SAC inefficiency should be less than 10^{-5}, which is not a problem given the high energy of the photons in this region. An Intermediate Ring Calorimeter (IRC), located just in front of the LKR, complements the acceptance for photon detection in the interface region between the coverage of the SAC and the LKR.

The 12 LAV calorimeters are operated (with the exception of the last one) in the vacuum tank. The positions of the LAVs along the experiment are chosen to satisfy the required angular coverage. This configuration is the best compromise between acceptance and detector complexity, driven by the whole geometry of the experiment and the high energy of the beam. The LAV use the lead glass counters recovered from the barrel electromagnetic calorimeter of the OPAL experiment at LEP. They provide a level of inefficiency of about 10^{-4} for photon energies larger than 0.5 GeV, which is well suited for NA62.

Since the largest fraction of K^+ decays contains muons in the final state, a muon veto system is mandatory both on-line and off-line. Because of the high rates, a muon veto in the earliest stage of the trigger is crucial to reduce the data rate below 1 MHz, the maximum rate that can be handled by the readout system. Further muon suppression is required offline. The muon veto system (MUV) consists of two iron-scintillator hadronic calorimeter (MUV1 and MUV2). These detectors are located after the LKR. An iron wall separates this calorimeter from a plane of fast scintillators (MUV3) which provides both the trigger information and additional power in vetoing muons. Monte Carlo simulations predict an overall inefficiency for muon detection of about 10^{-5}, achievable by exploiting the electromagnetic and hadronic shower separation capability of the hadronic and LKR calorimeters together. Further muon rejection is provided by the Ring Imaging Cherenkov counter (RICH) described in the next section. Finally a veto for charged and neutral particles placed just downstream of the last GTK station, provides additional rejection of the accidental background coming from hadronic interactions of the

beam particles in the last GTK station, as previously discussed. This detector, called CHANTI, consists of scintillators assembled in a rectangular shape surrounding the beam.

1.5 Particle Identification

To provide an additional 10^2 suppression of backgrounds originating from the $K^+ \to \mu^+ \nu$ (γ) decays, a Ring Imaging Cherenkov (RICH) detector is used. The 17 m long RICH is filled with Ne at atmospheric pressure. It is placed after the last straw chamber and equipped with 2000 phototubes. Results based on a full length prototype [14] indicate that, for π^+/μ^+ separation, a miss-identification probability smaller than 1% can be achieved in the relevant momentum range. A 17 cm diameter beam tube passes through the whole RICH volume to avoid the interaction of the beam with the Ne gas. The RICH also distinguishes positrons from π^+, allowing the suppression of backgrounds with positrons in the final state, such as $K^+ \to \pi^0 e^+ \nu$, or channels with a π^0 Dalitz decay ($\pi^0 \to \gamma e^+ e^-$). The pion Cherenkov threshold of the Ne is around 13 GeV/c. Therefore, to insure full efficiency, a cut on the minimum π^+ momentum of about 15 GeV/c is required. The RICH has also remarkable timing properties: its time resolution, which is better than 100 ps, makes this detector ideal to measure the arrival time of the π^+ and to match it to the parent particle measured by the GTK. Thanks to the long lever arm between the STRAW chambers which follow the spectrometer magnet, the RICH is also able to provide full acceptance for high momentum tracks from kaon decays with more than one charged particle. Finally, the RICH can also provide a cross-check to the π^+ momentum measured by the straw spectrometer when the particle is assumed to be a pion. Although the RICH momentum resolution is a factor of four poorer than the one obtained from the magnetic spectrometer, this redundant measurement is useful to reject events in the tails in the m^2_{miss}.The material of the RICH may complicate the detection of photons because of conversions or photo nuclear interactions. To mitigate this inefficiency, a charged hodoscope (CHOD) placed downstream of the RICH and before the LKR is envisaged. Simulations showed that the RICH itself, the CHOD, the IRC and the last ring of the LAV together, keep this source of inefficiency at a negligible level. The CHOD, standalone or in combination with the RICH, also offers a convenient Level 0 trigger for charged particles.

The LKR is a powerful particle identifier, not only for muons, as stated before, but also for positrons and electrons, as demonstrated by the past experience of NA48 and the first phase of NA62. It is, therefore, a viable option to complement the RICH in the suppression of final states with $e^{\pm}$, provided that its excellent energy resolution is maintained.

In NA62 the positive identification of the K^+ is also important because about 93% of the beam particles are π^+ or protons. They can interact in the residual gas contained in the vacuum tank and produce a signature that can mimic the signal. These backgrounds cannot be rejected simply applying a cut on the reconstructed longitudinal vertex position because they can occur anywhere in the fiducial volume. The K^+ identification is achieved by a Cerenkov threshold detector, (CEDAR) operated to be blind to all particles but kaons of appropriate momentum. It is located upstream of the GTK. With a time resolution of about 100 ps, the CEDAR can flag the kaon time windows accurately and hence reject pion or proton beam gas interactions, which could be misinterpreted as signal events. The precise timing of the CEDAR relaxes the requirements for the residual gas pressure in the decay region to $\approx 10^{-6}$ mbar. The CEDAR will also improve the rejection of candidates originating from pion and proton interactions on the last GTK station and can also serve to further improve the time association between the K^+ and the π^+ tracks.

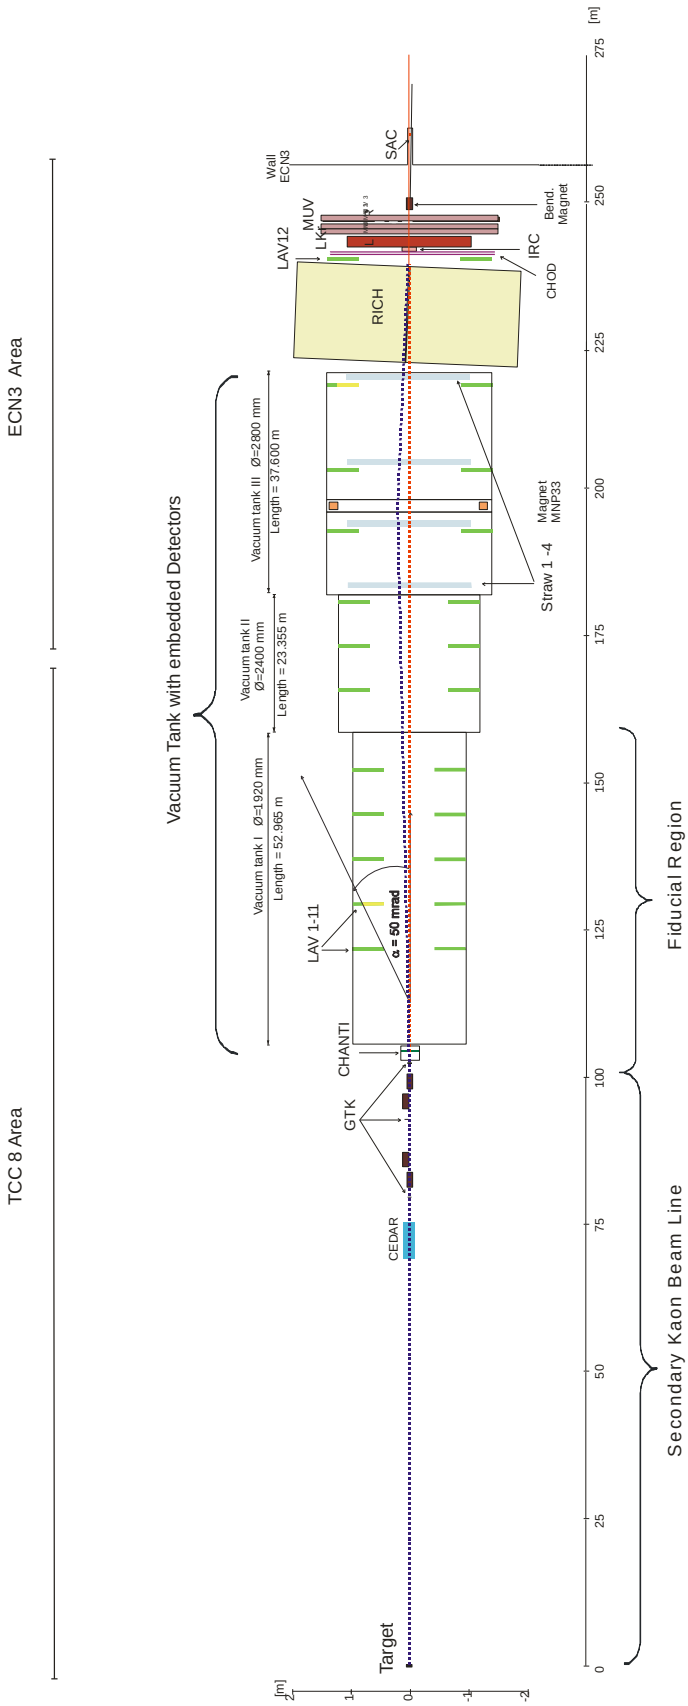

Figure 4 Longitudinal view of the Na62 Experimental setup (scaled model).

1.6 Summary of the proposed Detector Layout

For simplicity, a short summary of the NA62 layout is summarised here. The composition and layout of the NA62 beam and detectors are inspired by the experience gained with the previous Kaon decay experiment NA48 [15] performed at the CERN SPS. A secondary kaon beam line, ≈100m long, leads to a ≈65m long fiducial region, followed by a further length, over which the kaon decay products fly apart with respect to the beam line, so that they can be recorded in a series of detectors surrounding the beam.

The following detectors and features are incorporated in the NA62 layout (see Figure 1):

i) **The CEDAR** identifies the K^+ component in the beam with respect to the other beam particles by employing an upgraded differential Čerenkov (CEDAR) counter;

ii) The Gigatracker **(GTK)** is comprised of three Si micro-pixel stations measuring, time, direction and momentum of the beam particles before entering the decay region.

iii) **The STRAW** measures the coordinates and momentum of secondary charged particles originating from the decay region. To minimise multiple scattering the Chambers are built of ultra-light material and are installed inside the vacuum tank. The four Straw Chambers are intercepted in the middle by a large aperture dipole magnet (MNP33), providing a vertical B-field of 0.36T.

iv) The **RICH detector** is situated downstream of the last Straw chamber. It consists of a 17m long radiator filled with Neon Gas at 1 atm. allowing the separation of pions and muons between 15 and 35 GeV/c.

v) A system of Photon-Veto detectors provides hermetic coverage from zero out to large (~50mrad) angles from the decay region. This is assured by:
- the, existing, high-resolution Liquid Krypton electro-magnetic calorimeter (**LKR**),
- supplemented, at small and forward angles, by an Intermediate Ring (**IRC**) and Small-Angle (**SAC**) Calorimeters and,
- at large angles, by a series of 12 annular photon-veto (LAV) detectors **(LAV)**.

vi) The Muon-Veto Detectors (**MUV**) are composed of a two-part hadron calorimeter followed by additional iron and a transversally-segmented hodoscope. This system supplements and provides redundancy with respect to the RICH in the detection and rejection of muons.

vii) These detectors are complemented by 'guard-ring' counters (**CHANTI**) surrounding the last GTK station, and the charged-particle hodoscope (**CHOD**), covering the acceptance and located between the RICH and the LKR calorimeter.

viii) All these detector are operated and inter-connected with a high-performance **trigger** and **data-acquisition** (**TDAQ**) system.

1.7 Structure of this document

This Technical Design (TD) document is intended to be a concise but self-contained description of the NA62 experimental system. Further details can be found in many Technical Notes, referenced throughout.

A description of the high intensity K^+ beam and the beam defining detectors (CEDAR, GTK and CHANTI) is given in chapter 2. The downstream detectors (Photon Veto, STRAW Tracker, RICH, CHOD and MUV) follow in chapter 3. The Trigger and Data Acquisition System (TDAQ) is outlined in chapter 4.

2 Beam Line and Beam Defining Detectors

After the T10 target the experiment receives a high intensity hadron beam (750 MHz) with fixed momentum (75 GeV/c) and an average content of 6% kaons. Three sub-detectors are employed to measure the characteristics of the incoming decay particles:

- A threshold Cerenkov detector (CEDAR) is used to time stamp the K^+ particles in the beam.
- The Gigatracker (GTK) measures time, direction and momentum of the incoming beam particles. As it is placed in the beam it sees the full rate of incoming particles (750 MHz).
- The CHANTI is a charged particle veto detector –sitting directly behind the last the Gigatracker station- and detects unwanted interactions in the 3rd layer of the GTK.

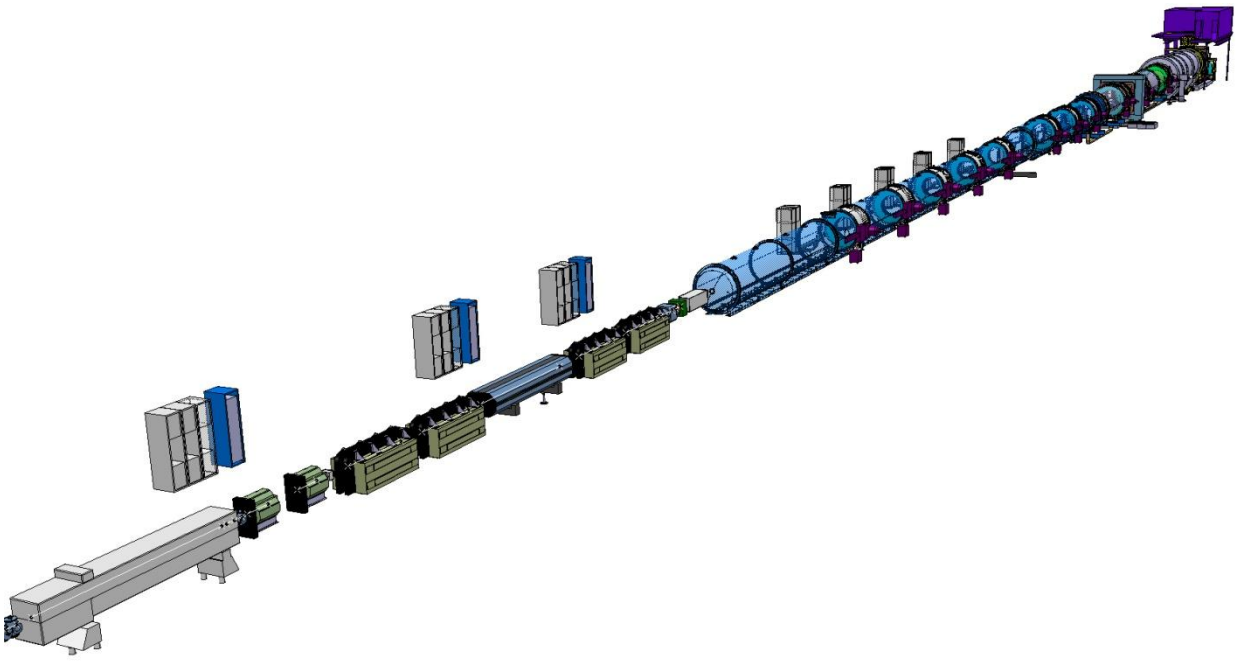

Figure 5 3D model of the experimental installation from the CEDAR.

2.1 The K^+ Beam Line

2.1.1 Rationale

2.1.1.1 Choice of Accelerator and Primary Proton Momentum

Distinct advantages can be gained from performing a charged kaon, ultra-rare decay experiment, $K^{+(-)} \rightarrow \pi^{+(-)} \nu \bar{\nu}$ at high energy and, in particular, at the 400 GeV/c SPS [2] :

From a simple empirical formula, which roughly fits the measured particle production data [16] , it can be deduced that, per primary proton of fixed momentum p_0, the maximum K^+ (K^-) production in a given momentum bite, $\Delta p/p$, and solid angle, $\Delta\Omega$, occurs at momentum: p_K ~0.35 p_0 (~0.23 p_0) and, at a constant fraction of the primary momentum p_K / p_0: K^+ (K^-) production increases as p_K^2 and therefore as p_0^2.

It follows that the number of K^+ (K^-) decays in a given fiducial length is maximum for p_K ~0.23 p_0 (~0.15 p_0) and, at constant p_K / p_0, the number of K^+ (K^-) decays in a given length increases as p_K and therefore as p_0.

Furthermore, the acceptance, resolution and rejection power of certain detectors, e.g. calorimeters and photon- and muon-veto counters, improve at higher energy.

The principal disadvantage of performing the experiment at high momentum lies in the difficulty of separating the kaons from other charged hadrons in the beam. Thus, in a beam employing radio-frequency cavities to provide transverse dispersion and subsequent cancellation of the deflection for π^+ and protons to be absorbed, the distance between cavities has to increase as p_K^2. However, the decay length, determining the survival of the wanted kaons increases only as p_K. [3]

However, even in a relatively short, 'unseparated' beam, it is possible to suppress the positron (electron) component, originating from the production and decay of π^0 in the target and subsequent conversion of the photons. The e^+ (e^-) component has been measured to be ~0.2 (~0.3) of the total particle flux for p_K = + (-)60 GeV/c and p_0 = 400 GeV/c [16] and varies roughly as $\sim(p_K/p_0)^{-2}$ relative to the π^+ (π^-) in the beam.

2.1.1.2 Choice of K^+

A choice of a positive rather than a negative kaon beam is motivated by the fact that, at a possible beam momentum of, for example, 75 GeV/c (Section 2.1.1.3), the ratio of production rates:

---

2 400 GeV/c is the highest proton momentum, at which the SPS can sustain a duty cycle of ≈0.3.

3 In the case of the 22 GeV/c K^+ beam proposed at FNAL [105]*, super-conducting r.f. cavities operating at a frequency of 4 GHz would be employed to suppress the π^+ and protons by at least an order of magnitude. The beam would have a total length of 216m, leading to a K^+ survival factor of only ≈0.27, compounded by losses inherent in the separation mechanism.*

K^+ / K^- per 400 GeV/c proton is = ~2.1 and the ratio: (K^+/π^+) / (K^-/π^-) = ~1.2, whilst the ratio: (K^+ / Total positively-charged particles) / (K^- / Total negatively-charged particles) = ~1.0.

2.1.1.3 Choice of Beam Momentum

The choice of p_K = +**75 GeV/c** as the central beam momentum is suggested with regard to the criteria listed in Table 1. The entries are shown for momenta of 60 and 120 GeV/c, at which complete measurements of particle production by 400 GeV/c protons on Be targets are available [16], and at 75 GeV/c, for which the relative particle composition was determined in 2007 and the absolute yield by interpolation of the data in [16] .

Moreover, the 'useful' range of momenta for a π^+ from K^+ decay to be distinguished from background μ^+ and hence to be accepted by the RICH detector is: ~15 < p_π < 35 GeV/c. In the decay of a K^+ at 75 GeV/c, particles, other than neutrinos, associated with an accepted π^+ must therefore carry an energy ≥ 40 GeV, sufficient to prevent their escape without detection.

Finally, 75 GeV/c proves to be near to the maximum momentum, for which the essential functions which are required of the beam (described in Sections 2.1.2.3 - 2.1.2.7) can be accomplished, with close packing of the elements, in the available length of ~102m from production target to the beginning of the decay fiducial region.

2.1.2 The High-Intensity K^+ Beam

2.1.2.1 Primary Proton Beam

The primary proton beam used for the production of the high-intensity K^+ beam has previously served to produce the K^0_L and the K^+ and K^- beams for experiments NA48 [15] and NA48/2 [17], respectively.

The 400 GeV/c, slow-extracted, proton beam from the SPS to the North Area is split into three branches to impinge on each of three targets (T2, T4, T6), to produce beams of secondary particles in an underground target cavern (TCC2). Target station T4 is the source of two secondary beams, H8 and H6. Magnets before and after the target allow the protons that have not interacted to be directed towards the front end of a third beam line, P42, for a wide range of choices of momenta in the two secondary beams[4]. The P42 beam is designed to transport the protons (through tunnels TT83, TDC8 and TT85) over a distance of ~823m to the entrance of the North Area High Intensity Facility (NAHIF) [18][5] . This comprises the enlarged underground target/beam tunnel, TCC8, and the

[4] *When required, different settings of the T4 target station magnets allow neutral particles to be directed towards the P42 beam line. Electrons (or positrons), from conversion of photons in a thin lead sheet, can then be selected with narrow momentum band (≈± 0.1 % $\Delta p_e/p_e$) around chosen central momenta (typically, 10 < pe < 100 GeV/c) and transported to the detectors via the P42 and K12 lines.*

[5] *An alternative path for the protons is available via target T6 (along a beam line, P62, which, after the first ≈130m, joins and thereafter follows the same path as P42), on condition that the front-end (pion decay section) of the M2 muon beam from target T6 is set to a momentum of -200 GeV/c for +400 GeV/c in P62.*

experiment cavern, ECN3. A drawing of the beam layout in the SPS North Area is available [19] and Figure 6 schematically shows the optics of the P42 beam.

Table 1 Criteria for the choice of the beam momentum.

Beam Characteristics		60 [GV/c]	**75 [GeV/c]**	120 [GeV/c]
Fluxes at production	p x 10^6	89	**171**	550
in 42 µster ΔΩ. % Δp/p acceptance	K^+ x 10^6	40	**53**	71
/ 10^{12} incident protons per s	π^+ x 10^6	353	**532**	825
	Total x 10^6	482	**756**	1446
Survival factor over 102m	K^+	0.797	**0.834**	0.893
	π^+	0.970	**0.976**	0.985
Fluxes at 102m from target	p x 10^6	89	**173**	550
in 42 µster ΔΩ. % Δp/p acceptance	K^+ x 10^6	32	**45**	63
/ 10^{12} incident protons per s	π^+ x 10^6	343	**525**	813
	Total x 10^6	464	**743**	1426
Decays in 60m fiducial length	K^+ x 10^6	3.9	**4.5**	4.1
/ 10^{12} incident protons per s [6]	π^+ x 10^6	6.1	**7.4**	7.2
K^+ decays / π^+ decays in 60m [7]		0.64	**0.61**	0.57
K^+ decays in 60m / Total hadron flux [8]	x 10^{-3}	8.4	**6.1**	2.9
$K^+ \to \pi^+ \nu \bar{\nu}$ Acceptance (Region I, no p_π cut)		0.08	**0.11**	0.11
Accepted $K^+ \to \pi^+ \nu \bar{\nu}$ / 10^{12} protons per s	x 10^6 x B.R.	0.31	**0.50**	0.45
Accepted $K^+ \to \pi^+ \nu \bar{\nu}$ / π^+ decays in 60m [7]	x B.R.	0.052	**0.067**	0.062
Accepted $K^+ \to \pi^+ \nu \bar{\nu}$ / Total hadron flux [8]	x 10^{-3} x B.R.	0.67	**0.67**	0.31

6 The K^+ decay rate relative to incident protons is related to the Signal with respect to the accompanying Background (notably neutrons and muons) resulting from primary protons interacting in the target and beam dump.

7 The ratio of K^+ to π^+ decay rates is related to the Signal to Background of decay products in the detectors.

8 The K^+ decay rate relative to the total hadron beam rate is related to the Signal available for a given (limiting) particle flux through the beam-defining detectors used for particle identification (CEDAR) and tracking and momentum measurement (GigaTracKer).

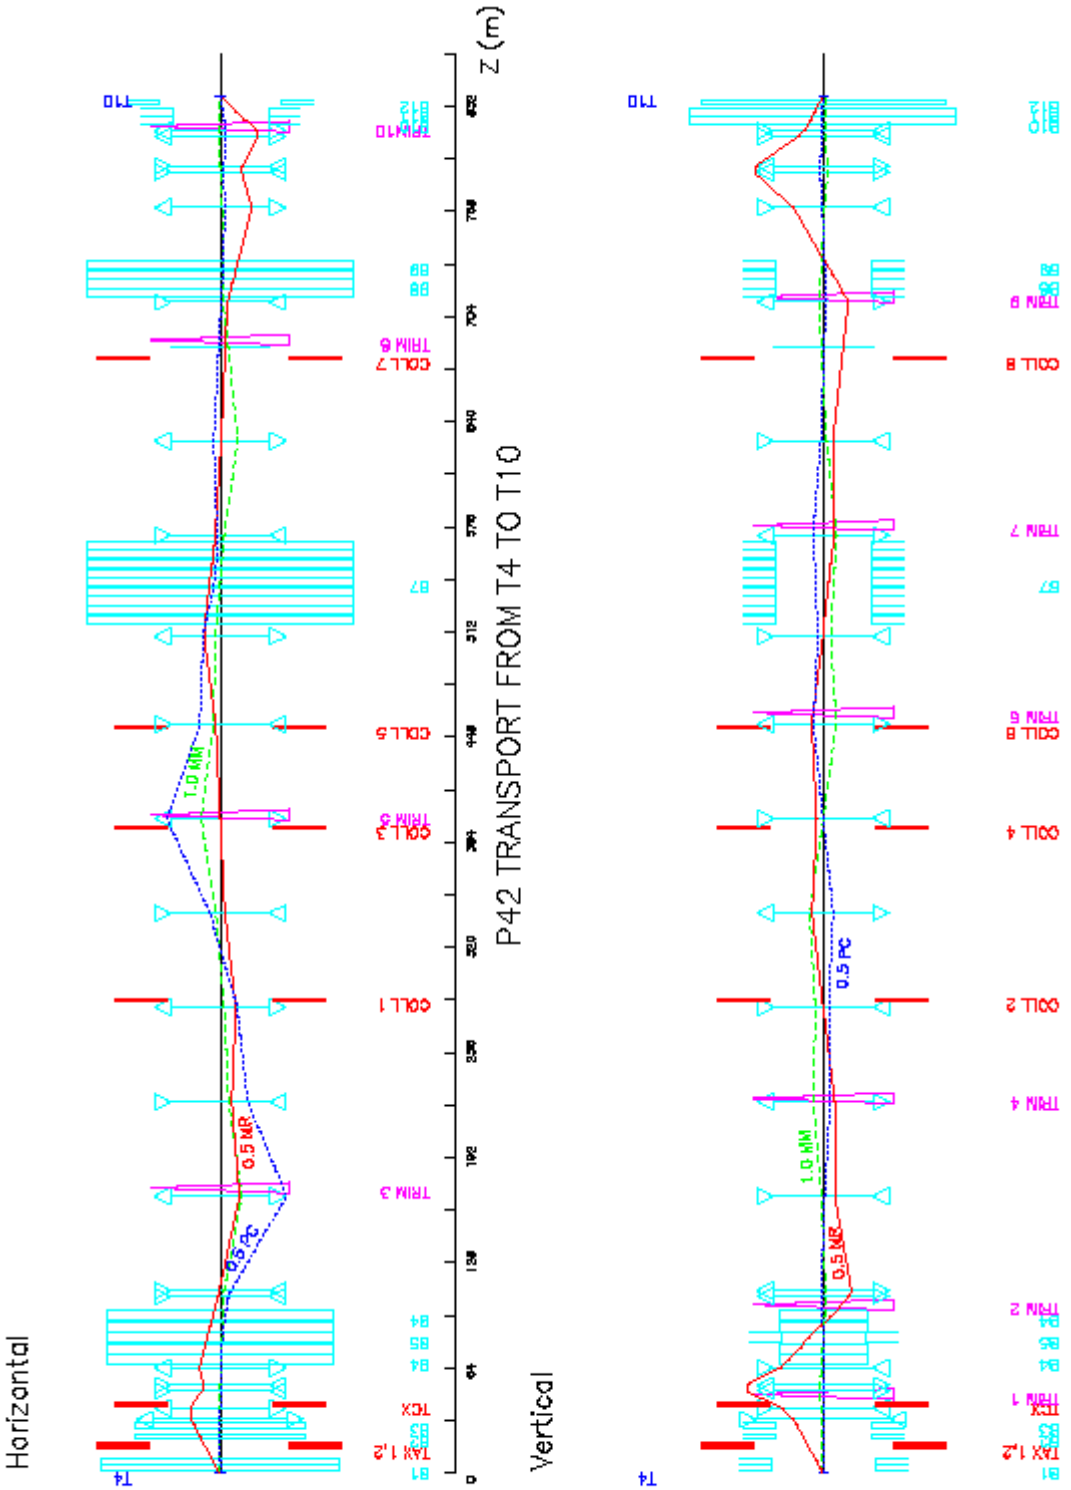

Figure 6 Schematic Optics of the primary proton beam [P42]

At the exit of target station T4, the transmitted proton beam passes through apertures in two vertically-motorized beam-dump/collimator modules, TAX 1 and TAX 2 of beam P42 (XTAX043018, XTAX043020), in which holes of different apertures define the angular acceptance of the beam and hence allow the flux of protons to be selected over a wide range. Moreover, access to the zones TCC8 and ECN3 is governed by closing TAX 1 and 2 in conjunction with switching off two magnets at the front of the beam. Table 2 lists the combinations of apertures available in TAX 1 and 2 for the proton beam and the corresponding transmission factors. The 'Range' refers to a possible restriction, which can be imposed on the movement of each TAX and hence on the apertures that may be selected, so as to guard against unwanted high proton fluxes being transmitted.

To prevent the primary beam from causing damage to components, a computer surveillance program (*P0-SURVEY*) allows the currents in the principal magnets of the P42 beam line to be monitored and, in case of error (or bad reading), automatically causes TAX 2 to close. This surveillance is usually extended to include critical magnets in the kaon beam line, K12, to prevent the secondary beam from passing into the detectors of the experiment.

Table 2 Apertures and Positions of TAX 1 and TAX 2 of Beam P42

Apertures and Vertical Positions of TAX 1			***TAX 2 of Beam P42***			
Diameter Ø (mm)	Position y (mm)	Range	Diameter Ø (mm)	Position y (mm)	Range	Transmission factor
14.0	-81	Medium	**12.0**	-21	Medium	1.0
14.0	-81	Medium	**10.0**	+19	Medium	~0.8
7.5	-21	Medium	10.0	+19	Medium	~0.5
7.5	-21	Medium	**4.0**	+59	**Small**	~0.12
7.5	-21	Medium	**2.0**	+99	**Small**	~0.03
--	+140	Small	--	+140	Small	*No Beam (Access)*

2.1.2.2 Secondary Beam Layout and Design

The secondary hadron beam, 'K12HIKA+', is designed to be derived from a high, but attainable, flux of 400 GeV/c protons (Section 2.1.2.10) in the underground North Area High Intensity Facility. The target/beam tunnel, TCC8, and the cavern, ECN3, where the detectors of experiment NA48 have been installed, have a combined length of 270m. It is planned to reuse the existing target station, T10, (located 15m from the beginning of TCC8), and to install the secondary beam along the existing (straight) K12 beam line, of length 102m to the exit of the final collimator, which marks the

beginning of the decay fiducial region and points to the NA48 detectors (notably the liquid krypton electro-magnetic calorimeter, LKR). A listing of the order and geometry of all elements along the new beam for experiment NA62 is contained in an updated *BEATCH* output file, accessible from [20] under the heading: K12HIKA+ and a drawing of the layout of the beam and detectors is available [21]. Figure 7 schematically shows the beam optics, calculated using the programme *TRANSPORT* [22], corresponding to the updated output file also accessible from [20].

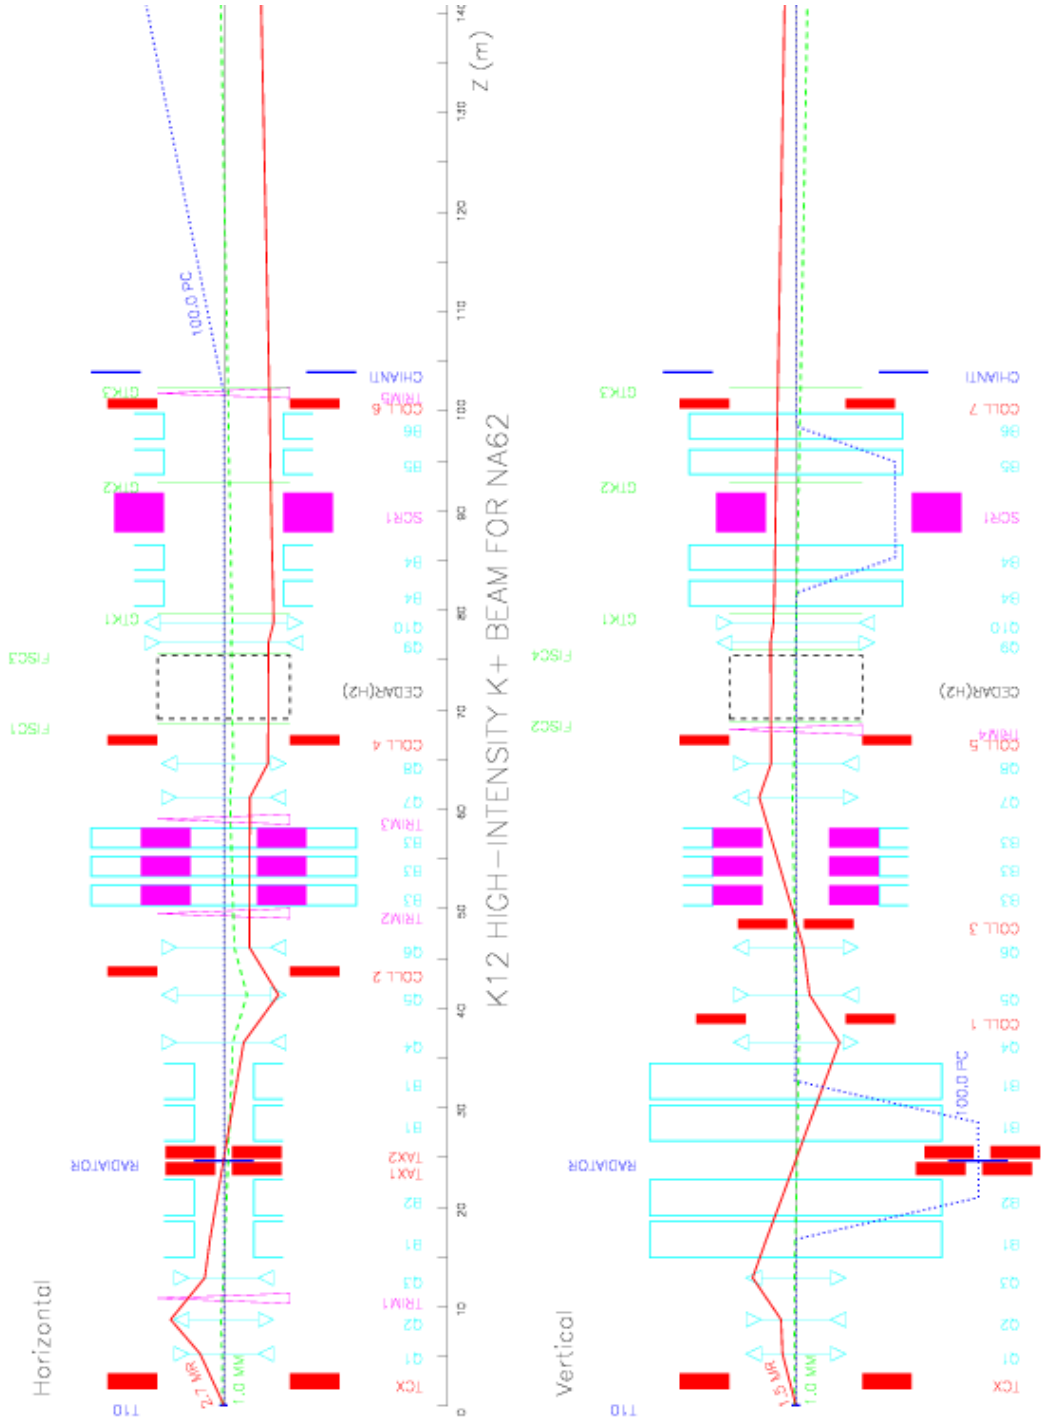

Figure 7 Schematic layout and optics of the high intensity K^+ Beam (K12HIKA+).

2.1.2.3 Targeting

The primary protons, transported via the P42 beam line, are focused and directed at zero angle [9] onto a 400mm long, 2mm diameter, beryllium target. This is suspended between thin aluminium foils and is cooled by forced convection of air in the T10 target station. The target is followed by a 950mm long, water-cooled, copper collimator, offering a choice of bores of different apertures. The largest, 15mm in diameter, is generally selected to transmit the wanted secondary particles as well as the remaining primary proton beam. Unwanted particles (notably $\pi^{\pm}$, $K^{\pm}$) issuing from the target at angles > ~6 mrad are thereby absorbed before they can decay to muons. Following the vacuum entrance window (0.1 mm thick aluminium), the opening angle for the transmitted beam is further reduced to ~4 mrad by a 1.6m long, collimator (TCX 101003) of 28mm diameter bore.

2.1.2.4 Acceptance, Momentum Selection and Positron Suppression

In contrast to the previous, simultaneous K^+ and K^- beams [17], the first active elements of the high-intensity beam are a triplet of radiation-hard, small-aperture, quadrupole magnets, QUAD 1, 2 and 3 (QNRB101005, QNRB101009, QNRB101013), which collect a large solid angle acceptance (±1.5 mrad horizontally x ±2.7 mrad vertically) at 75 GeV/c central momentum. The quadrupoles are fitted with cylindrical, stainless-steel vacuum tubes of 55mm inside diameter; those in QUAD 2 and QUAD 3 serve to define the acceptance angles in the horizontal and vertical planes, respectively. The drift-space between these two quarupoles serves to lodge a horizontal steering magnet, TRIM 1 (MDXH101011), used to centre the beam at the longitudinal position between two, vertically-motorized, beam-dump/collimator modules, TAX 1 and TAX 2 of beam K12 (XTAX101024, XTAX101026), where it is brought to a focus in both planes.

Front views of the two TAX modules, each 1.615m long, are shown in Figure 8. They are composed of copper-, followed by iron-blocks, mounted on water-cooled, copper tables, to enable them to absorb the remaining primary proton beam. TAX 1 + TAX 2 are located in a ~3.4m-long space between 0.2mm-thick aluminium vacuum exit and entrance windows, in the middle of a 'front-end achromat'. The achromat consists of four, vertically-deflecting, radiation-hard dipole magnets, BEND 1A, 2, -1B and 1C (MTRV101017, MTRV101021, MTRV101029, MTRV101033). The first and second magnets (with opposite polarity) give the 75 GeV/c beam a parallel downward displacement of 110mm, so as to pass through a set of graduated holes, of smallest diameter 10mm, contained in tungsten-alloy cylinders which are inserted in the blocks of TAX 1 and TAX 2. A list of these blocks and their inserts is shown in Table 3.

[9] *The choice of zero production angle provides the maximum yield (though not the highest fraction) of K^+ per incident proton. In this case, an angle around zero is practically imposed by the need to place strong acceptance quadrupoles of small aperture so close to the target, that the remaining primary proton beam must also pass through their apertures before being dumped.*

Table 3 List of blocks and cylindrical, W-alloy, collimator Inserts

TAX 1 (XTAX 101024)				
Block		Height from base (mm)	Inserts	
No	Mat		Outer Ø (mm)	Inner Ø (mm)
1	Cu	742	40	–
		522	40	15.0→13.5
		302	40	15.0→13.5
2	Cu	742	40	–
		522	40	13.2→12.2
		302	40	13.2→12.2
3	Fe	742	40	–
		522	40	12.0→11.0
		302	40	12.0→11.0
4	Fe	742	40	–
		522	40	10.8→**10.0**
		302	40	10.8→**10.0**

TAX 2 (XTAX 101026)				
Block		Height from base (mm)	Inserts	
No	Mat		Outer Ø (mm)	Inner Ø (mm)
1	Fe	742	60	**10.0**→10.8
		522	60	**10.0**→10.8
		302	60	–
2	Fe	742	60	11.0→12.0
		522	60	11.0→12.0
		302	60	–
3	Fe	742	60	12.2→13.2
		522	60	12.2→13.2
		302	60	–
4	Fe	742	60	13.5→15.0
		522	60	13.5→15.0
		302	60	–

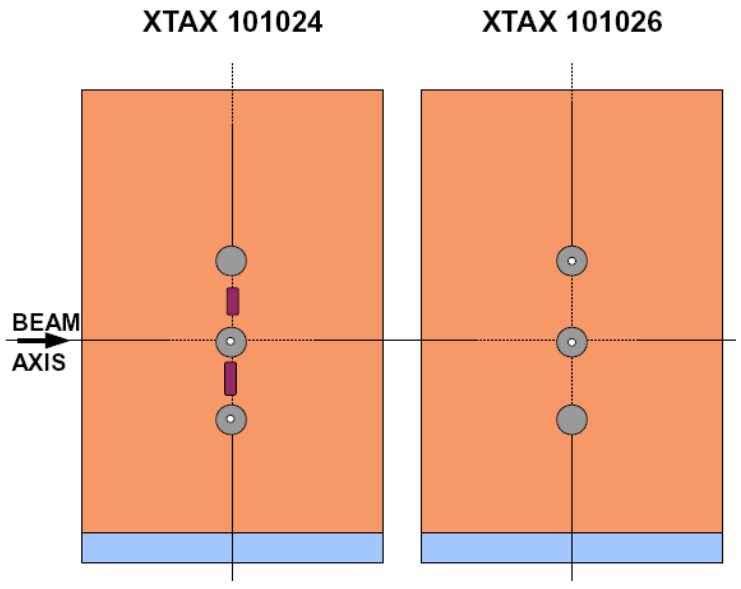

Figure 8 Front view of the beam-dump / collimator modules TAX1 and TAX2 of the K12 Beam.

The combination of vertical positions and permitted ranges of TAX 1 and TAX 2, used to select the 'K+ DOWN' path through the front-end achromat is indicated in Table 4 and is illustrated in Figure 9 (a), together with other combinations, which would allow for different choices of beam mode (discussed in Section 2.1.2.11)[10] . The chosen set of apertures redefines the beam at the focus in the horizontal plane and, by suitable off-set in the vertical plane (also indicated in Table 4), permits a narrow momentum band of ± 1.0 % Δp/p (r.m.s.) to be selected about the central momentum of +75 GeV/c. At this point between the two TAX modules, a motorized mechanism permits a 'radiator' (XCON101025), shown in Figure 10, made of tungsten plates of various thicknesses from 0 up to 5mm (~1.3 X_0), to be introduced into the beam. This is designed to cause e+ to lose sufficient energy by Bremsstrahlung to be rejected (with a survival factor down to ~10^{-3}) after the beam momentum has been redefined (see Table 5). Here the relatively large intrinsic divergences of the beam at the foci ensure that multiple-scattering in the radiator material does not lead to excessive loss of hadrons out of the beam emittance. The third and fourth magnets return the wanted particles onto the undeviated axis and thereby provide the second stage of momentum-definition.

Table 4: Vertical Positions of TAX 1 and TAX 2 for different Modes of Beam K12

TAX 1		***TAX 2***		
Position y (mm)	Range	Position y (mm)[11]	Range	K12 Beam Mode, *Path*
-110.0 -**107.5**	Large	+110.0 **+112.5**	**Small**	**K12HIKA+** : K^+ *DOWN* Δp/p (r.m.s.) = 1.0 %
+110.0	**Small**	-110.0	Large	K^- *UP*
+110.0	**Small**	+110.0	**Small**	Simultaneous K^+ *DOWN* + K^- *UP*
0.0	Medium	0.0	Medium	K12HIK0L : K^0_L *ON AXIS*
+140	**Small**	+140	**Small**	MUONS: *K-LESS* *UP+DOWN beams STOPPED*

[10] *In this case the TAX 'Ranges' are set to prevent both modules from being moved into a position near 0.0 , where the primary proton beam after the target might be transmitted along the subsequent beam line. (For the K0L on axis (Section 2.1.2.11, option 2)), this is prevented by imposing a minimum vertical angle of incidence of the upstream P42 beam – and hence a non-zero production angle for the K0L).*

[11] *Note for TAX 2, that the y-position is measured with the POSITIVE direction pointing DOWNWARDS !*

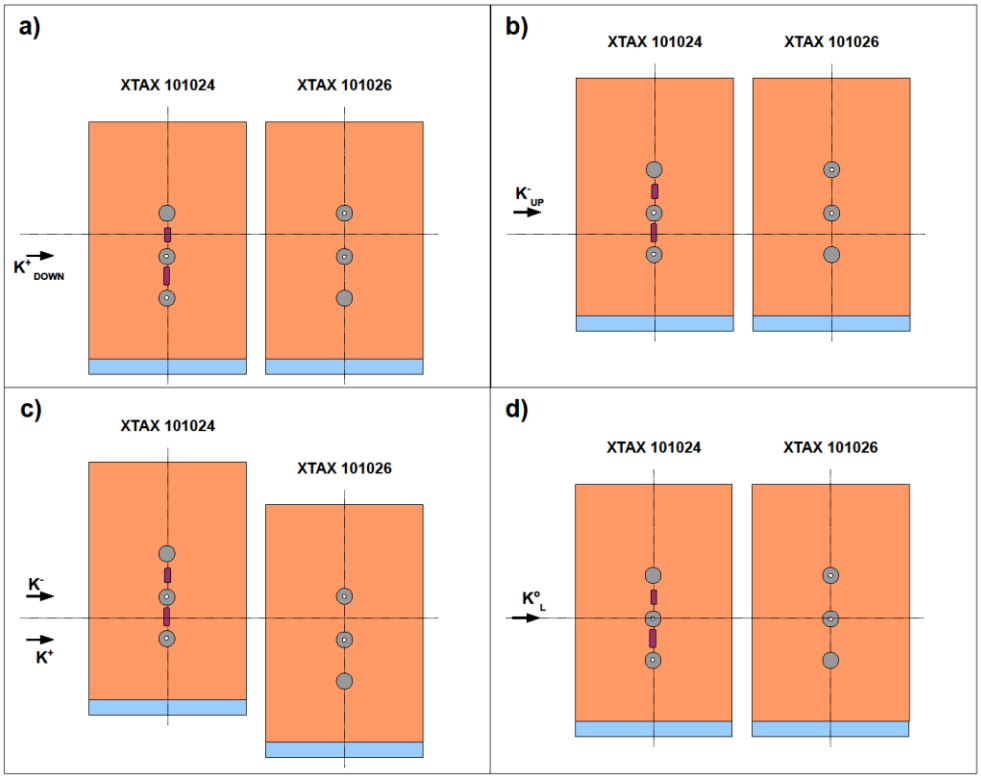

Figure 9 TAX 1 and 2 positions for different selections of K12 Beam modes and paths

(a) K12HIKA+ : K^+ *DOWN* (b) K^- *UP*

(c) Simultaneous K^+ *DOWN* + K^- *UP* (d) K12HIK0L : K^0_L *ON AXIS*

Table 5 Radiator positions and thickness with the corresponding positron (electron) survival factor.

Vertical Position: y (mm)		Thickness of W		$e^{\pm}$ Survival factor
For K^+, K^- beam(s):	For K^0_L beam:	(mm)	(X/X_0)	
±100	0	0	0	*1.0*
+60	-50	2.0	0.53	~0.06
+20	-90	3.0	0.80	~0.015
-20	+90	4.0	1.06	~0.004
-60	+50	5.0	1.33	~0.001

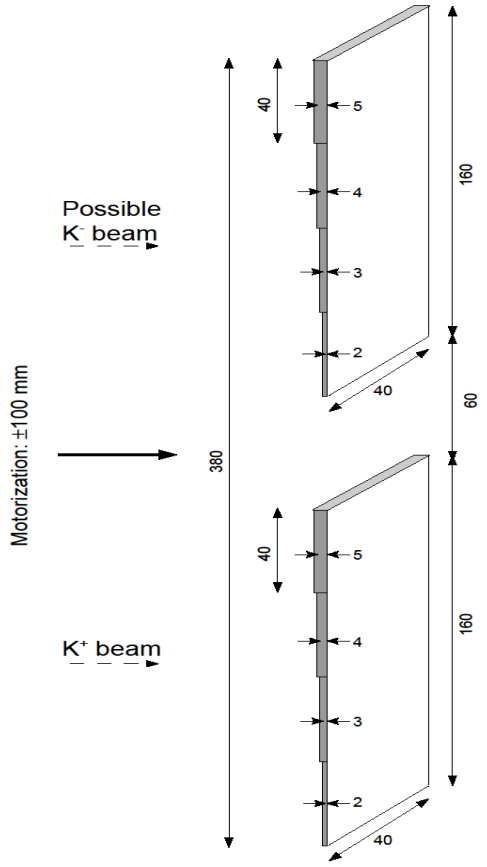

Figure 10 Schematic view of tungsten radiators that can be moved vertically into the beam to absorb positrons (electrons) in the K^+ (K^-) beam.

2.1.2.5 Beam Definition and Muon Sweeping

A triplet of quadrupoles, QUAD 4, 5 and 6 (QNL 101036, QNL 101041, QNL 101046), interspersed by adjustable, vertical and horizontal acceptance-redefining collimators, COLL 1 (XCSV101038) and COLL 2 (XCSH101043), serves to refocus the beam in the vertical plane and to make it parallel and of limited lateral extent in the horizontal plane. This allows it to pass through a 40mm diameter, almost field-free bore, drilled in iron slabs, which are inserted to fill the 200mm high gaps in three, 2m-long, dipole magnets, BEND 3A, 3B and 3C (MBPL101051, MBPL101054, MBPL101057). The vertical magnetic field in the iron surrounding the beam serves to sweep aside muons of both signs, whilst the deviation of the beam due to the small stray-field inside the bore can be cancelled by two steering dipoles, TRIM 2 (MDXH101049) before and TRIM 3 (MDXH101059) after the MBPL magnets. In the vertical plane, an adjustable collimator, COLL 3 (XCSV101048), with tapered jaws open to ~±1.2mm, shown in Figure 11, is used to redefine the beam at a second focus. At this point the e^+, that have been degraded in momentum, are displaced sufficiently to be absorbed.

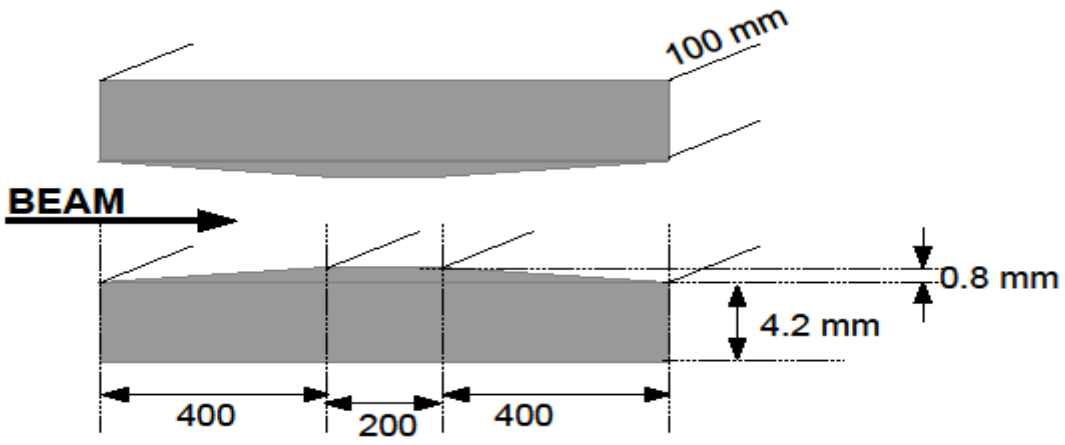

Figure 11 2-jaw collimator (XCSV101048) used to redefine the vertical beam spot

2.1.2.6 K^+ Tagging

A pair of quadrupoles, QUAD 7 and 8 (QNL 101061, QNL 101065), then renders the beam sufficiently large and parallel in both planes to match the requirements of an upgraded CEDAR differential Čerenkov counter (XCED101072) [23], which can be tuned to tag only the K^+ in the beam (Chapter 2.2). The CEDAR occupies a length of ~6.0m. To limit the material traversed by the beam to a total of <~7x $10^{-3}X_0$, hydrogen gas (at an absolute pressure of ~3.85 bar) is used to fill the CEDAR pressure-vessel. This is connected directly to the beam vacuum tube via flexible bellows at each end. The hydrogen gas volume is separated from the vacuum by thin entry and exit windows of aluminium-alloy, 75mm in diameter. To be sure that a rupture would not affect the detectors of the experiment, the downstream window is designed to be thicker (~0.2mm) than the upstream one (~0.1mm). The upstream window is in contact with the beam-line vacuum only, which is linked to an evacuated vessel of sufficient volume (~4m^3) to contain the hydrogen gas at a pressure below atmospheric.

The CEDAR is preceded by an adjustable, 4-jaw (horizontal and vertical) cleaning collimator, COLL 4 / 5 (XCHV101067), and a vertical steering magnet, TRIM 4 (MDXV101068). About 7m apart along the beam and straddling the CEDAR, are two pairs of filament scintillator counters, FISC 1 and 3 (XFFH101069, XFFH101075) and FISC 2 and 4 (XFFV101070, XFFV101076), which, when connected in coincidence respectively, permit the mean divergence of the beam to be measured and tuned to zero in each plane.

2.1.2.7 Beam Tracking and Momentum Measurement

A pair of weakly-focusing quadrupoles, QUAD 9 and 10 (QFS 101077, QFS 101079), matches the beam through the tracking and momentum-measurement stage, shown schematically in Figure 12 and limits the beam size through the apertures of the downstream detectors. The beam tracking system consists of three 'GigaTracKer' stations (Chapter 2.3), composed of Si-pixel arrays, each of active area: ±30mm (horizontally) x ±13.5mm (vertically), installed in the beam vacuum, where,

when needed, they can be isolated by separation valves (VXSV101080, VXSV101093 and VXSV Ø=320mm). The stations are arranged so that the space between GTK 1 and GTK 3 is occupied by a 'second achromat', composed of four, vertically-deflecting, C-shaped dipole magnets, BEND 4A, -4B, 5 and 6 (MCBV101082, MCBV101085, MCBV101095, MCBV101098). Whereas the yokes of the first three magnets are, as usual, below the beam (*Down*), that of the fourth is supported above the beam (*Up*). Use is thereby made of the return fields in the yokes of the third and fourth magnets to supplement the 'defocusing' action for μ^+, produced by a 5m-long, toroidally-magnetized iron collimator, SCRAPER 1 (XCMV101090), which surrounds the beam in the momentum-dispersed section between the second and third magnets. Here, its lower jaw, placed 20mm below the beam, can intercept and deflect away μ^+ of momenta < ~55 GeV/c, which are dispersed out of the 75 GeV/c parent π^+, K^+ beam. Moreover, the upper jaw blocks the undeviated and upper paths through the achromat, thereby ensuring that no punch-through, neutral (neutrons, K^0_L) or negatively-charged (π^-, K^-), hadrons can be transmitted to the end of the beam. GTK 2 is located in the same section, just after the magnetic collimator, where the +75 GeV/c beam has a parallel, downward displacement of 60mm and a dispersion of 0.6mm / % Δp/p that allows the track momentum to be measured with a resolution of ≈ 0.2 % (also indicated in Figure 12). GTK 3 is placed on the undeviated axis, after a final, adjustable, 4-jaw cleaning collimator, COLL 6 / 7 (XCHV101101), and is followed by a series of 'guard-ring' scintillation counters (CHANTI), which surround the beam in vacuum and are designed to veto interactions in the exit layer of GTK 3.

Between the cleaning-collimator and GTK 3, a horizontal steering magnet, TRIM 5 (MDXH101102), is used to deflect the beam away from the straight axis by an angle of +1.2 mrad (towards the Jura). This angle is adjusted, so that the subsequent -3.6mrad deflection, due to the single spectrometer magnet (MNP-33), directs the beam back through the central aperture of the LKR calorimeter, as shown schematically in Figure 13.[12]

[12] *The +1.2 mr horizontal deflection causes a small dispersion in angle as a function of momentum of the beam entering the fiducial region: Δx' (mr) = -0.012 Δp/p (%). For dp/p(r.m.s.) = 1.0 %, the contribution to the angular spread of the beam, Δx'(r.m.s.) = 0.012 mr, is small compared to its intrinsic divergence (*Table 6*). Using the measured momenta (*Figure 12*), the contribution to the error on the angle measurements of individual tracks, dx'(r.m.s.) = ≈0.003 mr, becomes negligible.*

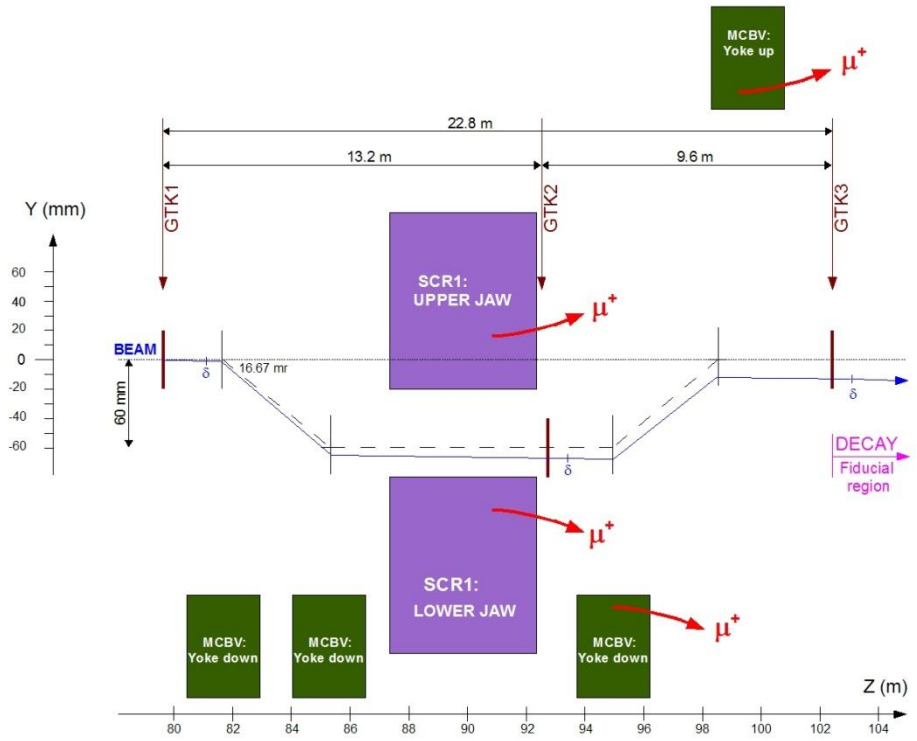

Figure 12 Schematic layout of the beam tracking and momentum-measurement stage (second achromat).

Spatial resolution of one GTK station, $\sigma(x, y) = 0.3 \text{ mm}/12^{1/2} \approx 0.087$ mm

Multiple scattering in one GTK station, $\delta(x', y') = 15 \text{ mrad}.(X/X_0 = 4.5\ 10^{-3})^{1/2} / p \text{ (GeV/c)} \approx 0.013$ mrad

Measurement Resolutions:

Angle : **Δx', Δy' (r.m.s.)** (mrad) = $[\ 2\ (\sigma/22.8)^2 + \delta^2 + (13.2\ \delta/22.8)^2\ ]^{1/2}$

Resolution: GTK1+3, Mult. scatt.: GTK3 GTK2

= **~0.016 mrad**

Momentum: $\frac{\Delta p}{p}$**(r.m.s.)** (%) = $[\ \sigma^2 + (9.6\ \sigma/22.8)^2 + (13.2\ \sigma/22.8)^2 + (9.6 \text{ x } 13.2\ \delta/22.8)^2\ ]^{1/2}$ x 100/60

Resolution: GTK2 GTK1 GTK3 , Mult. scatt.: GTK2

= ~0.22 %

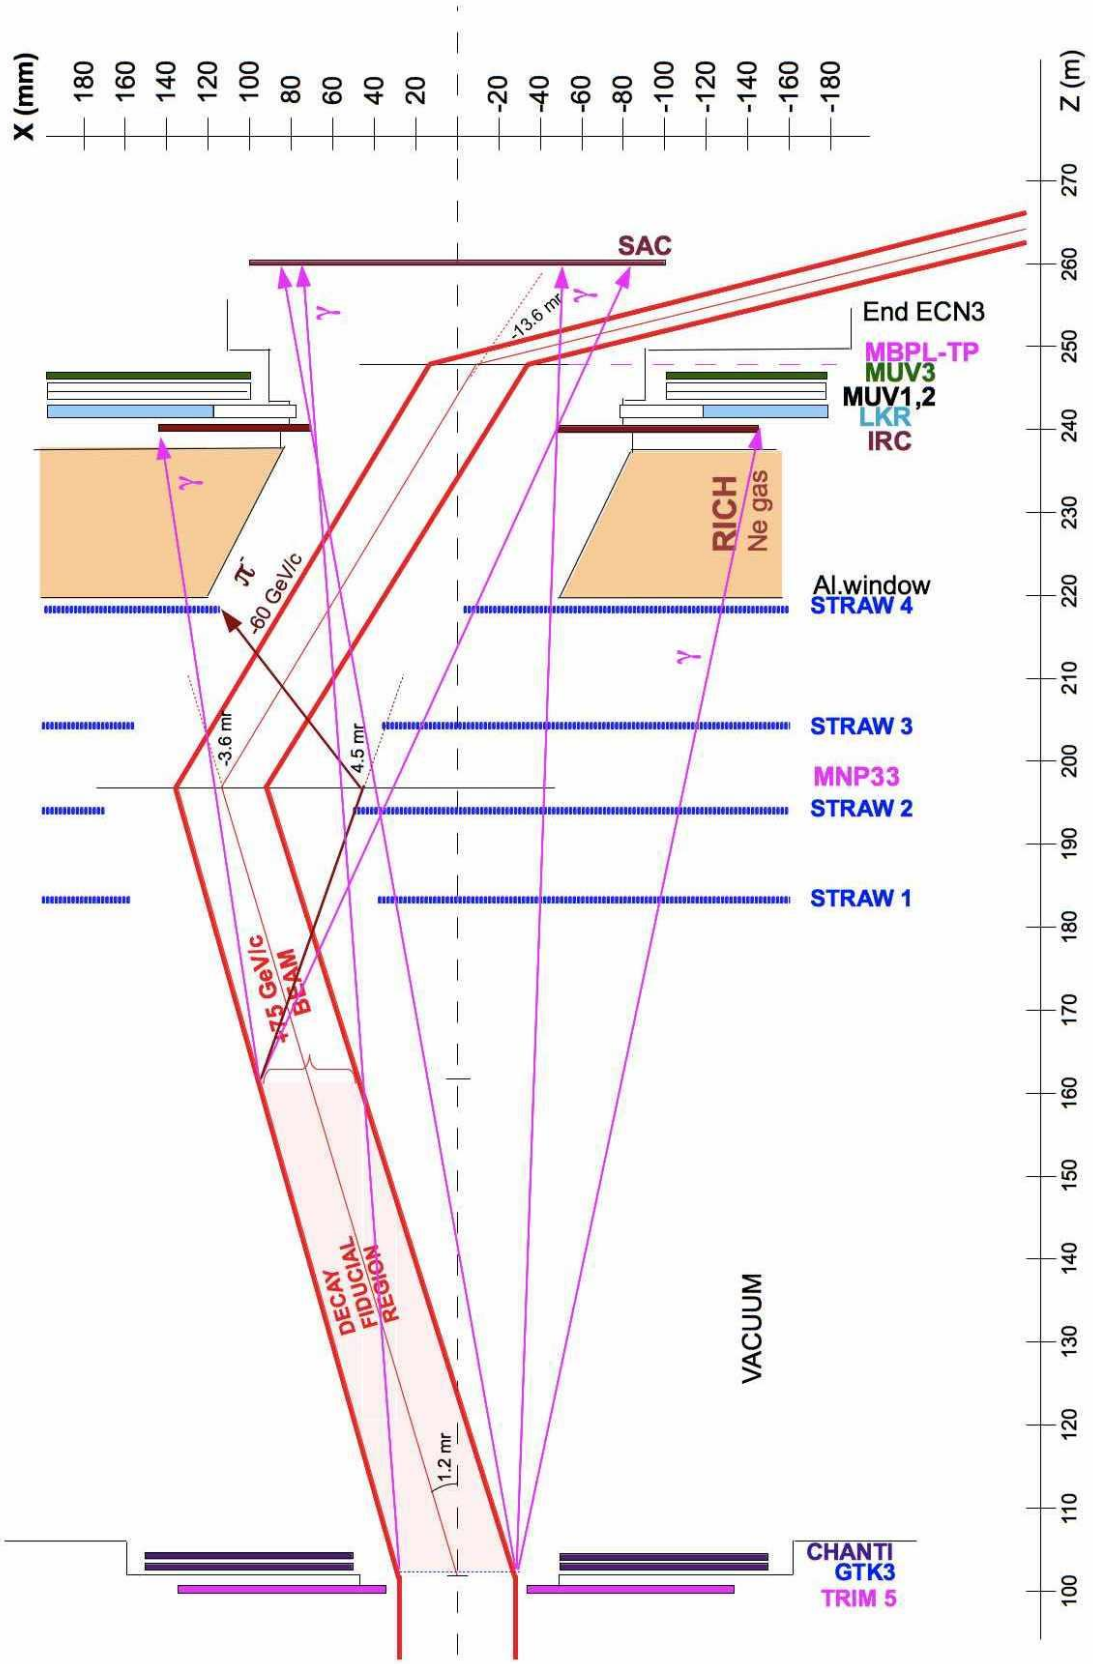

Figure 13 Schematic layout of the downstream part of the beam. Shown is the angular deviation of the charged kaon beam and the 2σ width of the beam profile.

2.1.2.8 Decay Region, Magnetic Spectrometer and Detectors

The decay fiducial region is contained in the first 60m of a large, ~117m long, evacuated tank (described in Chapter 3.1.1), which is closed off by a thin (~0.045 X_0) aluminium window, separating it from the neon gas of the following RICH counter. This window is off-set horizontally and its centre is traversed by a thin-walled aluminium beam tube (of inside diameter 155.8mm), which follows the trajectory of the beam, which is thus transported in vacuum through the downstream detectors (Figure 13).

The detector components are described in later Chapters. Of relevance for the beam is the single magnetic spectrometer, comprising four tracking chambers composed of straw-tubes (STRAW 1 - 4), covering the full acceptance outside a ~118mm-wide, empty passage around the beam path. These are interspaced by the existing, large-aperture dipole magnet (MNP33), which provides a horizontal p_T-kick of -270 MeV/c, thereby deflecting the 75 GeV/c beam by -3.6 mrad, so as to converge to the undeviated axis at a point just after the LKR calorimeter. As well as tracking candidate π^+ from K^+ decays, the layout of the spectrometer is designed to detect π^- (from K^+e4 decays) of momenta extending up to 60 GeV/c, deflected through angles > + 4.5 mrad by MNP33; (a limiting ray is shown in Figure 13).

Near the point of convergence of the beam with the axis following the LKR calorimeter, a pair of 'Big' filament scintillator counters, FISC 5 and 6 (XFBH101245, XFBV101246), allow the beam to be observed and steered correctly. The beam is then deflected to the side through a further angle of -13.6mrad by a 2m-long, 200mm-gap, tapered-pole magnet (MBPL-TP) to reach a point ~12m further downstream, where it clears a small–angle, photon-veto calorimeter (SAC). This is inserted by ~6m on rails into the beam vacuum tube, of 600mm inner diameter. This tube is itself installed in a larger pipe, which is buried in the ground beyond the end of the cavern, ECN3. The beam is finally absorbed in a beam dump composed of iron surrounded by concrete at the downstream end of this pipe, at a distance behind the detector, which allows space to shield against 'back-splash'.

As shown in Figure 13, the detection of photons from the fiducial region is rendered hermetic in the forward direction by an intermediate ring calorimeter (IRC) in front of the LKr calorimeter and a final small-angle calorimeter (SAC). At the IRC, the beam still has a displacement towards positive x with respect to the undeviated axis. The inner aperture of the IRC, chosen to be 120mm in diameter, can thus be centered on the beam with an off-set in x of +12mm. With appropriate alignment (at an angle of -2.0mrad) and a modest increase (to 170mm) in the inner diameter of the upstream end of the vacuum tube traversing the RICH, it thereby becomes possible to assure that photons pointing from the fiducial region towards the SAC will not encounter any material in their path. This feature is vital, since any e^+e^- resulting from conversion would escape detection in the SAC, due to the action of the preceding beam-deflecting magnet.

2.1.2.9 Beam Parameters and estimated Performance

The principal characteristics of the high-intensity K+ beam are listed in Table 6, where the parameters determining the flux yield are compared with those for the previous K12 simultaneous K^+ and K^- beam, designed for experiment NA48/2 [17]. The effective solid angle and momentum acceptance, as well as the beam sizes and divergences are calculated using the ray-tracing programme TURTLE [24], corresponding to the updated output file accessible from [20]. For example, the simulated momentum distribution and the spot sizes at the positions of the three beam tracking detectors (GTK1, 2 and 3) are shown in Figure 14 and Figure 15, respectively. The particle fluxes of the previous beam are taken from the measurements actually made at 60 GeV/c. The relative particle-composition of the 75 GeV/c beam is derived from a scan made with a CEDAR counter in a secondary beam (H2) at the SPS in 2007, whereas the total hadron flux is taken as the average of the values obtained by interpolation from 60 and from 120 GeV/c according to the empirical formula proposed in [16].

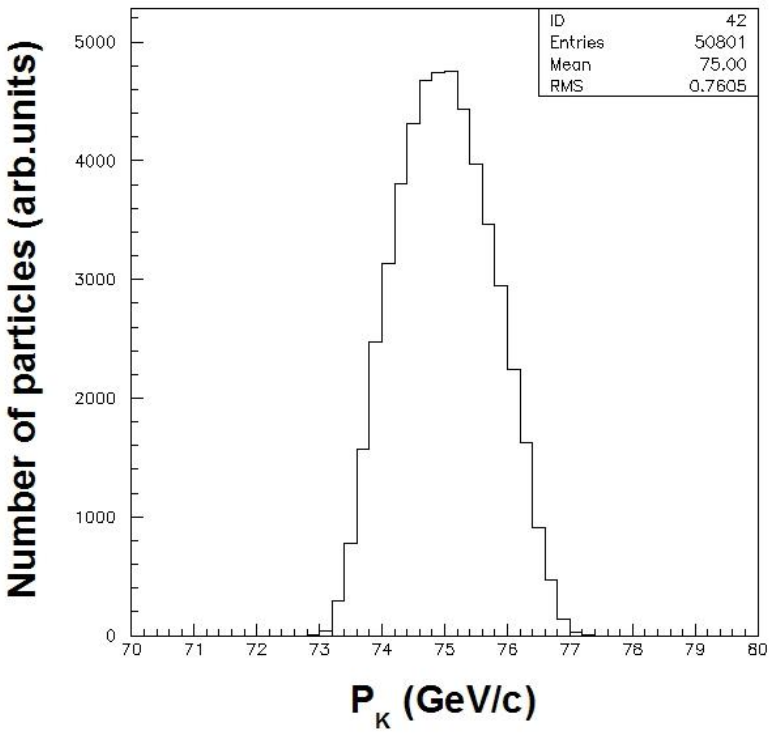

Figure 14 Beam momentum spectrum.

Table 6 High Intensity K^+ Beam for experiment NA62 with 1 X_0 W radiator to suppress e^+ in comparison with previous NA48/2 beam.

Beam Experiment	K12K+K- NA48/2	**K12HIKA+** NA62	Comparison FACTOR [13]
SPS Protons per s of spill length Instantaneous Proton Rate per effective s	~2 x 10^{11} 3.3 x 10^{11}	0.7 x 10^{12} **1.1 x 10^{12}** (equiv. 1.0 10^{12}) [14]	 3.0
SPS Duty Cycle (s / s) Effective Duty Cycle (s / s)	4.8/16.8 = 0.29 ~0.18	~0.3 ~0.2	 ~1.1
Beam Acceptance x_0', y_0' (mrad)	± 0.36, ± 0.36	± 2.7, ± 1.5	
Solid Angle (μsterad)	≈ 0.4	≈ 12.7	32
Mean K^+ Momentum $<p_K>$ (GeV/c)	60	**75**	K^+ 1.4 π^+ 1.5 Total Hadrons 1.6
Momentum Band: - Effective Δp/p (%) - r.m.s. Δp/p (%)	 ± 5 ≈ 3.7	 ± 1.65 **1.0**	 0.33
R.m.s. Divergence: x', y' (mrad) at CEDAR		0.07	
2 r.m.s. Beam Size (mm) Area at GTK 3 (mm^2) R.m.s. Divergence: x', y' (mrad)	r= ~15 ~700 ≈ 0.05, 0.05	x= ± 27.5, y= ± 11.4 ~980 0.09, 0.10	 ~1.4
Decay Fiducial Length: (m) Δz (τ_{K+}) Decay Fraction: ($1 - e^{-\Delta z}$)	50 0.111 0.105	**60** 0.107 0.101	 0.96
Inst. Beam Rate / s (MHz): p **K^+** π^+ e^+, μ^+ Total	2.9 1.0 11.1 ~3 , ~0.13 ~18	173 **45** 525 ~0.3 , ~6 **750**	60 45 47 ~0.1 , ~45 ~42
Fraction of hadrons in GTK 1 GTK 1 + 2 + 3 Mean Rate over sensitive area of KABES / GTK3 (kHz / mm^2) Maximum Rate per Si pixel (kHz / [0.3 x 0.3] mm^2)	 ~25	0.987 0.981 463 114	 ~19
Overall Efficiency x Running Time/year (days /yr.) Effective Spill Time/year (s /yr.)	0.5 x 120 = 60 9.3 x 10^5	0.6 x 100 = 60 1.0 x 10^6	 ~1.1
K^+ Decays per year in fiducial length	1.0 x 10^{11}	**4.5 10^{12}**	≈ 45

[13] *The comparisons refer to the positive beam only of the simultaneous $K^+ + K^-$ beams for experiment NA48/2.*

[14] *The equivalent instantaneous proton beam rate of (1.0 x 10^{12}) s^{-1} takes into account the 10% loss of secondary hadrons from the K12HIKA+ beam due to absorption and multiple scattering in the 1 X0 W radiator, introduced to suppress the e+ component. This equivalent proton rate is used for the comparison FACTOR.*

The muons accompanying a high-energy, high-intensity, secondary beam contribute a major part of the single-particle flux, to which the detectors outside the beam are exposed. The transport and decay to $\mu^{\pm}\nu$ of a wide spectrum of $\pi^{\pm}$ and $K^{\pm}$ originating in the target has been simulated using the programme *HALO* [25], corresponding to the four updated output files accessible from [20]. The programme tracks the parent particles and their decay muons inside the beam apertures and the 'halo' muons leaving the apertures through the vacuum tubes, magnet yokes and shielding surrounding the beam. Results of such calculations are given in Table 7 and Table 8. As an example, the distribution of halo muons traversing a plane of vertical straw tubes in STRAW chamber 4 is plotted in Figure 16.

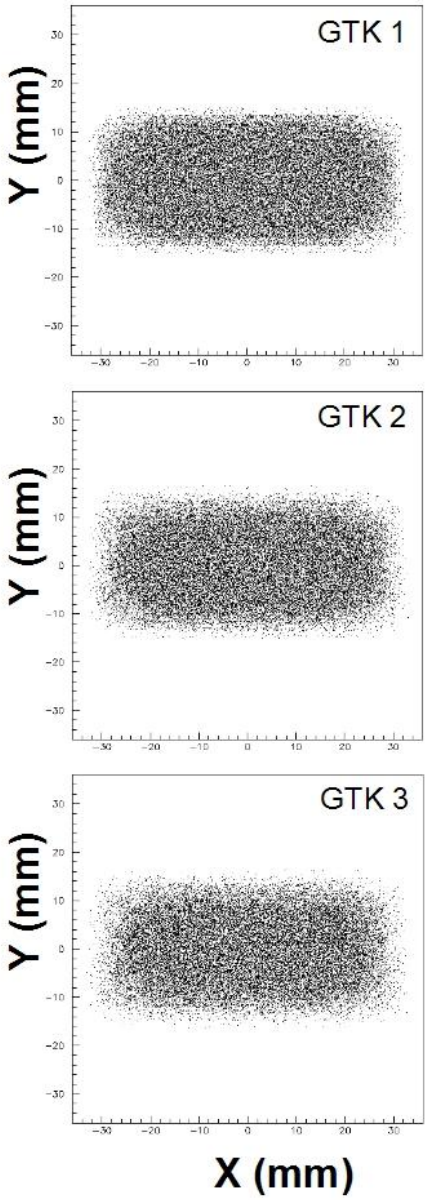

Figure 15 Spot-sizes of the high intensity K+ beam in the three Gigatracker stations.

Table 7 Estimated Muon Halo in the K12HIKA+ Beam

INSTANTANEOUS RATES x 10^6 from 1.1 10^{12} incident protons per eff. s [15]	**π^+** -> **μ^+** ν	**K^+**	**π^-** -> **μ^-** $\bar{\nu}$	**K^-**	**TOTAL**
π , K at GTK 3 (z = 102.4m)	525	45	0	0	570
μ in beam at GTK 3	4.38	0.46	0	0	4.84
μ 'HALO' at GTK 3 reaching LKR cal.	1.66	0.65	0.15	0.01	2.47
μ at CEDAR PMT's: 25<r<30cm (z=69.3m)	4.16	0.35	0.76	0.10	5.37
μ in GTK Guard Ring CHANTI: 4.5<|x|<15, 2.5<|y|<15cm ANTI-HALO counter : r> 18, |x|,|y|<60 cm (z=104.1m)	1.26 1.47	0.62 0.33	0.04 0.64	0.01 0.10	1.93 2.54
Large Angle Veto: LAV 1 (z=121.4m) LAV 1-12 - TOTAL LAV 1-12 - OR	1.13 5.88 1.93	0.25 3.52 1.52	0.28 1.26 0.46	0.05 0.25 0.12	1.71 10.91 4.03
μ per plane of STRAW 1 (z=183.5m) [210 x 210] cm^2, |x| > 5.9 cm [210 x 210] cm^2, |y| > 5.9 cm	 2.11 2.53	 4.00 4.13	 0.21 0.25	 0.02 0.02	 6.34 6.93
STRAW 2 (z=194.1m) [210 x 210] cm^2	2.22 2.58	4.40 4.59	0.21 0.22	0.01 0.01	6.84 7.40
STRAW 3 (z=204.5m) [210 x 210] cm^2	2.18 2.57	4.72 4.92	0.20 0.21	0 0.01	7.10 7.71
STRAW 4 (z=219.1m) [210 x 210] cm^2	2.89 2.77	4.82 5.08	0.15 0.16	0.01 0.01	7.87 8.02
RICH: 12 < r < 120 cm (z=237.35m)	2.27	5.67	0.15	0.01	8.10
IRC: 6 (x=+1.2cm) < r < 14.5 cm (z=239.7m)	6.62	1.02	0	0	7.64
CHOD/ LKR: 12<r< 120 cm (z=241.08m)	2.78	5.80	0.15	0.01	8.74
MUV: r >12,|x|,|y| <140 cm (z=246.5m)	4.39	6.30	0.18	0.02	10.89
SAC: [20 x 20] cm^2 (z=261.0m)	0.04	0.13	0	0	0.17

[15] *The statistics of the simulation are such that each ray tracked represents 10^4 particles.*

Table 8 Expected instantaneous muon rates in the detectors - for 1.1 x 10^{12} protons per effective second of spill

DETECTOR	Area (cm^2)	K12 HIKA+ BEAM TOTAL RATE (MHz)	MAXIMUM INTENSITY / effective s (kHz/cm^2)	Annual MUON dose rate (Gy per year) [16]
CEDAR PMT's: (25<r<30cm)	864	5.4	~10	~5
GTK 3 Guard Ring:				
- CHANTI: 4.5<|x|<15cm 2.5<|y|<15 cm	855	1.9	~4	~2
- ANTI-HALO counter: r>18, |x|,|y|<60cm	1.34 x 10^4	2.5	~0.2	~0.1
Large Angle Veto:				
- LAV 1	1.67 x 10^4	1.7	~0.1	~0.05
- LAV 1-12 TOTAL	26.0 x 10^4	10.9		
- LAV 1-12 OR	26.0 x 10^4	4.0		
STRAW Chamber 4:				
- per plane: 5.9 < |x|<105 cm 5.9 < |y|<105 cm	4.16 x 10^4 4.16 x 10^4	7.9 8.0		
- per Ø 1 cm STRAW: (5.9<|x|<6.9)x210cm (5.9<|y|<6.9)x210cm	210 210	0.5 0.4	~30	~15
RICH: (12 <r< 120cm)	4.48x10^4	8.1		
IRC: 6 (x=+1.2cm)<r< 14.5cm	547	7.6	~80	~40
CHOD / LKR cal.: - 12 <r< 120cm - 12 <r< 16cm	4.48x10^4 352	8.7 1.3	~4	~2
MUV: r>12, |x|,|y|<140cm	7.79 x 10^4	10.9		
SAC: (20x20) cm^2	400	0.17	~0.4	~0.2

[16] *1 Gy = 1 J/kg = 6.24 x 10^9 MeV/g = [6.24 x 10^9 / (dE/dx)] charged particles / cm^2.*
Taking dE/dx ~1.8 MeV/g cm^2 for minimum-ionizing charged particles traversing the detectors, 1 Gy/year = ~3.5 x 10^9 μ± / cm^2 / year.
This translates to a maximum intensity of ~2 x 10^3 μ± / cm^2 / effective second of spill integrated over a year of ~1.7 x 10^6 eff.s (100 days x 0.2 eff. duty cycle)

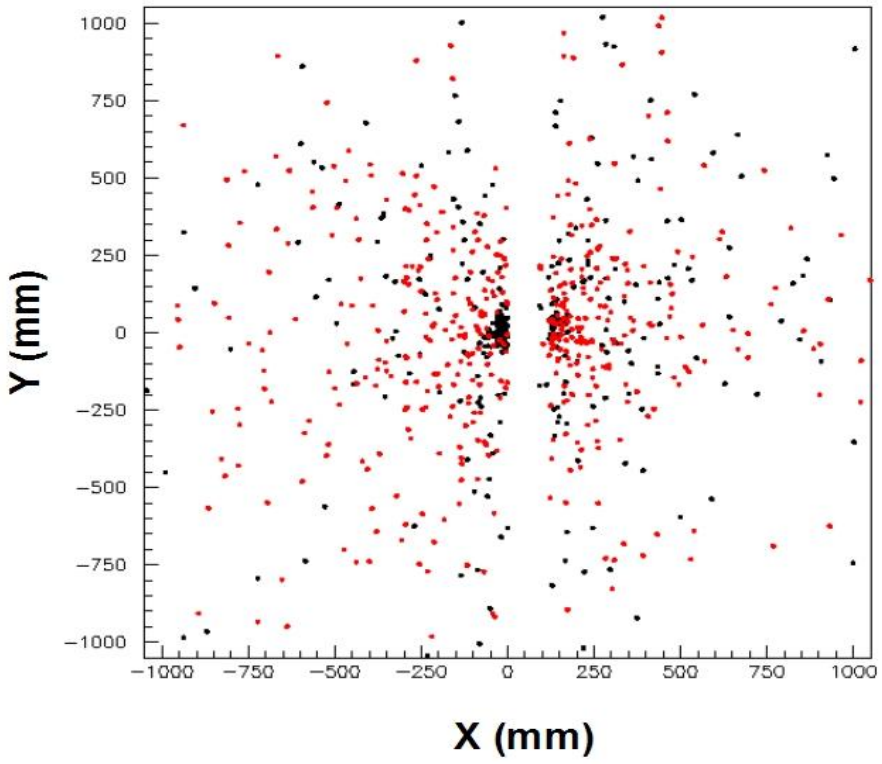

Figure 16 Simulated distribution of the muon halo - traversing a plane of vertical straw tubes (6 < |x| < 105cm) in STRAW chamber 4. One entry in the histogram represents 10^4 $\mu^{\pm}$ from π^+,(π^-) $\rightarrow \mu^{+(-)}\nu$ decays (black dots), K^+ and (K^-) $\rightarrow \mu^{+(-)}\nu$ decays (red dots) - per 1.1 10^{12} incident protons per effective second of spill.

-

2.1.2.10 SPS Scheduling and Proton Availability

From the parameters in Table 6 we conclude that, with effective overall data-taking efficiency (SPS and experiment) of 0.6, $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ branching ratio of 1 x 10^{-10} and detector acceptance of 0.1, **~45 events** could be accumulated in **~2.5 x 10^6 seconds** of 400 GeV/c slow-extracted proton beam on target (equivalent to ~1.7 x 10^6 effective seconds of spill). With a duty cycle of 0.3 (~0.2 effective), this could be provided in a year of ~100 days of scheduled fixed-target proton physics.

The instantaneous primary beam rate assumed is 1.1 x 10^{12} protons on target T10 per effective second of spill, corresponding to a mean rate of ~0.7 x 10^{12} per real second of flat-top for slow-extraction at 400 GeV/c momentum.

Thus, for example, with an SPS cycle giving 9.6 real seconds of flat-top per pulse, this would correspond to ~7 x 10^{12} protons per pulse (ppp) onto target T10. With a 300mm-long beryllium target generally used in the upstream target station T4, this would require ~1.8 x 10^{13} ppp to be directed into the branch from the 3-way split towards T4.

The proton flux needed at T4 might be reduced to ~1.0 x 10^{13} ppp, if the proton beam, instead of being focused onto the target in both planes, were to be rendered parallel in the vertical plane with an extent of ~10mm. Then the 2mm-high T4 target would intercept only ~20% of the protons, allowing the remaining ~80% to by-pass the target without absorption. A study has shown that the front-end optics of the P42 beam line to T10 could be adapted according to this scheme.

2.1.2.11 **Future Beam Options**

The layout and optics of the K12HIKA+ beam are optimised to provide the high-intensity, single-charge beam required by the $K^+ \to \pi^+ \nu \bar{\nu}$ experiment NA62.

However, in adopting a 'straight-line' layout from the existing target station T10 to the LKR calorimeter, two further options have been built in to the front-end of the beam. These have been included in the combinations of K12 beam dump/collimator (TAX 1 and 2) positions listed in Table 4 and are illustrated in Figure 9 (c) and (d), respectively :

1) Simultaneous, positively- (K^+) and negatively-charged (K^-) beams of hadrons with momenta up to ~± 75 GeV/c could be selected, following '*DOWN*' and '*UP*' paths through the front-end achromat (Figure 9 (c)). The first triplet of quadrupoles (with D-F-D polarities) [17] could be tuned to bring the two beams to foci, having similar, first-order, magnifications in both planes, at the point of momentum-selection in the middle of the front-end achromat. The second quadrupole triplet (with D-F-D polarities) could then bring the beams to a common, second focal point, where differences in the cumulative, second-order (notably 'chromatic') aberrations would almost cancel. A subsequent (new) sequence of four quadrupoles (with D-F-D-F polarities), placed upstream of the second achromat, could be arranged to give the two beams similar lateral sizes and (small) divergences entering the fiducial region leading to the detectors. Moreover, the $e^{\pm}$ radiator mechanism in the front-end achromat Figure 10 is designed to intercept both beams and to suppress the e^+ and e^- components by similar factors.

2) A high-intensity, neutral beam (K12HIK0L), with angular acceptance up to ~± 0.2 mrad in each plane, could be selected along the straight, undeviated axis (in Figure 9 (d)). The beam could be simply derived from the primary proton beam at zero production angle, or, if the first three quadrupoles and one steering dipole magnet are removed, it could be produced at a (vertical) angle of ~2.5 mrad, to yield a more optimum ratio of the K^0_L / neutron content.

[17] The focusing sequence of the quadrupoles refers to their action on the positive beam in the horizontal plane and is similarly valid for the negative beam in the vertical plane. Therefore, to obtain two beams of opposite charge having similar characteristics, the focusing action of the quadrupoles on each beam must be made similar in the horizontal and vertical planes. In addition, there is a weak effect due to the fringe fields at the ends of the rectangular poles of the vertically-deflecting dipole magnets in the front-end and in the second achromats. This effect results in a small extra focusing in the horizontal plane, which is similar for the two beams.

In both cases, however, the layout of straws, in each of the four STRAW chambers and the alignment of the RICH counter (Figure 13), would have to be modified to exploit such beams. Otherwise the beams would intercept these detectors.

2.2 CEDAR

2.2.1 CEDAR Detector Requirements

The disadvantage of high energy protons used by NA62 and, consequently, of a high energy secondary beam, is that the kaons cannot be efficiently separated from pions and protons at the beam level. The consequence is that the upstream detectors which measure the momentum and the direction of the kaons are exposed to a particle flux about 17 times larger than the useful (kaon) one. A critical aspect is therefore to positively identify the minority particles of interest, kaons, in a high rate environment before their decay. This will be achieved by placing in the incoming beam a differential Cerenkov counter, CEDAR, filled with hydrogen gas.

A FLUKA simulation was used to study the interactions of pions, kaons and protons with the residual gas in the vacuum decay tank and the probability that such an interaction can cause fake triggers was computed. The conclusion is that (in the absence of kaon tagging) the vacuum should be better that 6 10^{-8} mbar to keep the background to less than one fake event per year. This very challenging requirement can be relaxed by at least an order of magnitude by positively tagging the kaons by means of a CEDAR Cerenkov counter in the beam line, filled with hydrogen gas at an absolute pressure of about 3.6 bar. A necessary part of this kaon identification is the precise timing of the different components in order to guarantee a good rejection of the background due to the accidental overlap of events in the detector.

An upgraded form of the CEDAR built for the SPS secondary beams (CERN Report CERN-82-13) will be used, and will be insensitive to pions and protons with minimal accidental mis-tagging. The choice of the Hydrogen gas is dictated by the need to minimise material on the beam line, and hence reduce multiple Coulomb scattering – see Table 9. The foreseen thicknesses of the upstream and downstream aluminium vacuum windows are 150 µm and 200 µm, correspondingly.

Table 9: the thickness expressed in X_0 for different gases used in the CEDAR counter. The pressure values are chosen to have the same index of refraction.

	Helium	**Nitrogen**	**Hydrogen**
Pressure [bar]	10	1.7	3
Al window thickness [µm]	2x400	2x100	150+200
Al window thickness [10^{-3} X_0]	9.0	2.2	3.9
Gas thickness [10^{-3} X_0]	12	35	3.2
Total thickness [10^{-3} X_0]	21	37	6.6

2.2.2 Overview of the CEDAR operation

The CERN CEDAR counter (see Figure 17) has been designed to identify particles of a specific mass by

making the detector blind to the Cerenkov light produced by particles of different mass. For a given beam momentum, the Cerenkov angle of the light emitted by a particle traversing a gas of a given pressure is a unique function of the mass of the particle and the wavelength of the emitted light. The Cerenkov light emitted by particles of different mass is then not transported by the CEDAR optics through the diaphragm slit onto the light detectors but absorbed on its way, as explained in more detail below.

Table 10: Characteristics and resolutions of the CEDAR

Parameter	**Value**
Gas type:	H_2
n-1	~142 x 10^{-6}
Nominal pressure for kaons	3.85 bar
θK	30.9 mrad
Kaon rate	50 MHz
Time resolution	<100 ps
$\Delta\theta/\theta$	4.8 10^{-3}
$\Delta\beta/\beta$	5x10^{-6}
Number of photons per kaon	100
Rate per photomultiplier	3 MHz

The CEDAR consists of a pressure vessel filled with gas of a precisely controlled pressure. At the end of the vessel, the Cerenkov light is reflected by a spherical Mangin mirror onto a ring-shaped diaphragm of 100 mm radius with adjustable aperture width, located at the beginning of the vessel. The pressure of the gas can be adjusted so that only light of the wanted particle type will be transmitted through the diaphragm slit. Eight photomultipliers in its original version are placed behind the slit and the coincidence of at least 6 of them will indicate the passage of a particle of the wanted mass. A chromatic corrector lens, designed to match the dispersion curve of the gas and positioned in between the mirror and the diaphragm, ensures that the light of all wavelengths arrives at the same radius on the diaphragm plane. The advantage of this design is that light from unwanted particles hits the diaphragm plane at a different radius from that of wanted particles. It therefore does not pass through the aperture and does not contribute to the rate. Hence the device is well suited to tagging minority particles in a high-intensity beam. This only works if all rings are concentric, which requires all beam particles to be parallel to each other. This imposes certain constraints on the beam optics.

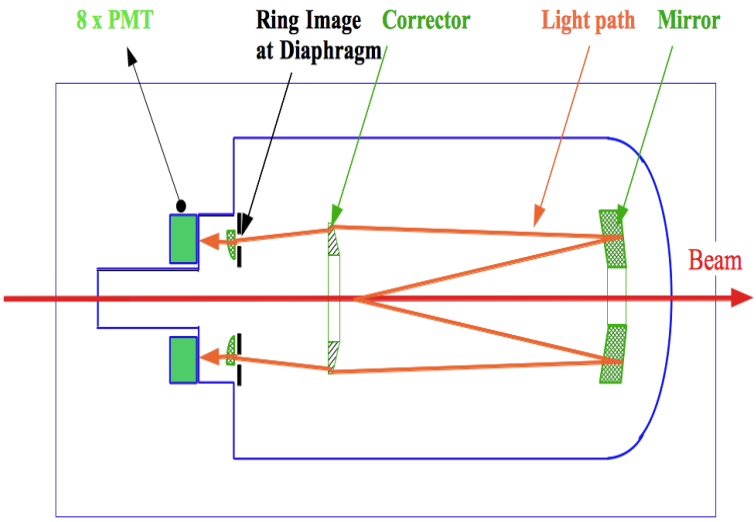

Figure 17 Schematic layout of the Standard West-Area CEDAR.

Two versions of the CEDAR counter have been built for use at the SPS (1). The North CEDAR, filled with Helium gas, is optimized for high energies and the West CEDAR, Nitrogen filled, for lower beam momenta. The difference is related to the Cerenkov angle, determined by the beam momentum and the refractive index of the gas, and the optical correction, which relates to the dispersion of the gas used. It has been verified by a ray tracing program that the West version of this instrument would function well for our application using Hydrogen at room temperature instead of Nitrogen, thus reducing significantly the scattering of the beam in the gas. The optical design minimises the dispersion of Cerenkov light and enables the aperture of the diaphragm to be reduced. Thus, photons produced by charged kaons pass through while light from pions and protons is blocked. During 2006, a test run was performed on one of the CEDAR-West Cerenkov counters (filled with N_2) and validated its ability to distinguish kaons from pions and protons in the NA62 experiment, as well as the light spot shape predicted by a simulation program. It was also found from the simulation that the upstream 1.2 metre section of beam pipe containing hydrogen contributes only marginally to the efficiency and can thus be replaced by an extension of the beam vacuum pipe. In addition to a small reduction in multiple Coulomb scattering, such a modification is helpful in the redesign of the optical system necessary to handle the increased photon flux. The main parameters of the proposed Hydrogen-filled CEDAR-W counter are listed in Table 10.

The main effects that broaden the light spot at the diaphragm are:

1. optical aberrations, limited to about 6 microns and therefore negligible;
2. chromatic dispersion, largely corrected for by the chromatic corrector;
3. multiple scattering of the beam during its traversal of the gas, minimised by the choice of Hydrogen gas;

4. inhomogeneity of the refractive index of the gas due to temperature gradients, minimised by thermal insulation and (in the case of NA62) by the location of the CEDAR in an underground area with inherently rather stable temperature;
5. beam divergence, which is one of the dominant factors, but limited to an RMS angle of about 70 μ rad in each plane according to full beam transport simulation;
6. noise and halo particles contributing to the scintillation light produced in the gas.

The rate from the kaon component in the high-intensity beam for NA62 is 50 MHz and the size of the light spot emerging from each quartz window is 30 x 8 mm (rectangular). The CEDAR detector is required to achieve a kaon tagging efficiency of at least 95%, with a time resolution of 100 ps. The present photomultipliers are too slow for this purpose, and an optical system with appropriate and multiple photodetectors will be developed to achieve a singles rate on each pixel not higher than 3 MHz mm^{-2}.

2.2.3 CEDAR West: Vessel and internal optics

The CEDAR counter is housed in a steel pressure vessel of 534 mm inner diameter and 558 mm outer diameter, welded on square flanges, to which a nose with eight quartz windows is bolted on the upstream end while a spherical head closes the downstream end. The overall length of the main vessel is 4500 mm. A 1243 mm long nose extends upstream of the quartz windows making the overall length of the counter 6200 mm. The thermal expansion is 0.05 mm/K. The beam windows at the entry and exit are 100 mm diameter, and made out of Aluminium alloy. The measurement of 6200 mm quoted above includes the thermal shielding and external envelope, while the value 5772 mm refers to the length of the gas volume including the full nose. The dimensions of the shortened counter are being decided, including shielding and infrastructure. There is also some length required for the diaphragm window. (See the overall drawing at url: http://na48.web.cern.ch/NA48/NA48-3/groups/beam/CEDAR/CedarDrawing.pdf).

Table 11: Relative positions of all optical elements with respect to the entrance window

Element	**Z (mm)**
Entrance window	0
Upstream end of quartz windows	1111
Condenser plane	1211
Diaphragm plane	1251
Chromatic corrector upstream end	2234
Mangin mirror exit	5772
Exit window	5772

The optical elements inside the counter comprise the mangin mirror, the chromatic corrector, the diaphragm and 8 sets of condenser lenses and quartz windows. A sketch of the overall layout and

the optical path is given in Figure 18 and the relative positions of all elements with respect to the entrance window are listed in Table 11.

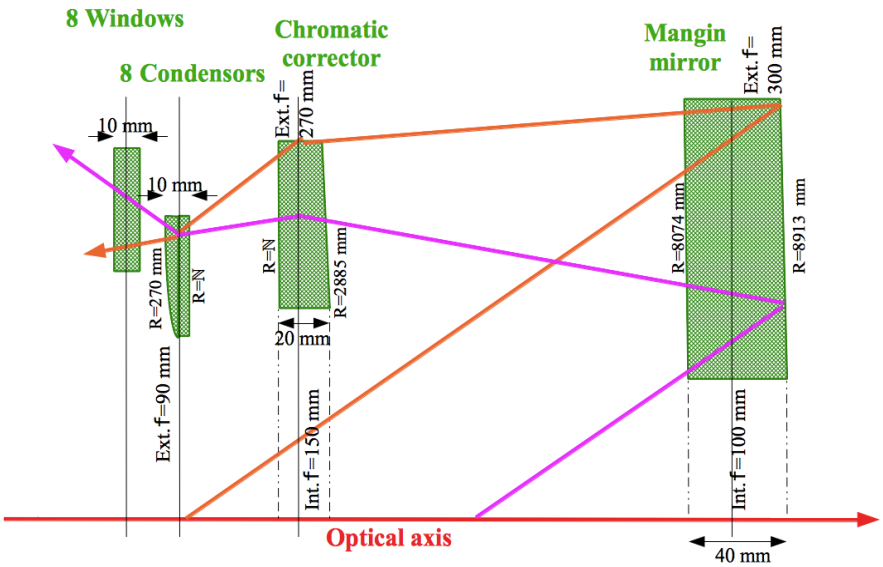

Figure 18: Schematic layout of the optical elements and the light path.

The whole counter is packed with 100 mm of polyurethane foam to provide thermal insulation and this ensures a temperature uniformity of ±0.1 degrees. The supporting feet, gas pipes and electric cables are designed to minimise thermal conduction to the outside world. Furthermore, thick Aluminium shells have been clamped to the outside of the big vessel and Copper shells are fastened on the nose. These modifications enable the impact of external temperature variations to be reduced by a factor of about 20. The vessel temperature is monitored by three platinum wire thermometers located at both ends and at the diaphragm plane.

The optical axis of the CEDAR counter must be precisely aligned with the beam axis. To that end the vessel is supported in three points. One point is fixed and located at the longitudinal position of the quartz windows. The two other points are located 4.347 m downstream and rest on a V-shaped surface under the mirror. The alignment system, made of cast Aluminium, has two chariots moved by right-left screw drive, opening the V-surface. By moving the chariots in opposite direction the vertical position can be adjusted, whereas another screw drive allows a motion in the same direction and thus a horizontal adjustment. The step resolution is 10 microns providing an angular positioning within ±2.3 μ rad. The X, Y range is from -4.3 to +4.3 mm.

The optical elements (mangin mirror, chromatic corrector, and diaphragm) are attached to an optical tower made as an isostatic, triangular assembly of tubes bolted onto spherical cups. This tower is straight to within ±0.2 mm over 4.5 metres and has a static sag of 0.1 mm. It is supported by a sphere under the diaphragm and is fixed without constraint to the pressure vessel so that pressure and temperature variations do not affect the optical axis. The ring diaphragm is fixed at the upstream end of the tower. The mirror is located at the downstream end and can be aligned in radial

and longitudinal directions. The chromatic corrector rests on the first intermediate ring of the tower and can be aligned in three directions.

The mangin mirror is made from high-quality Suprasil I quartz blanks. The outer diameter is 300 mm, the diameter of the inner hole is 100 mm. The thickness of the quartz mirror is 40 mm, with different radii of curvature for the reflecting surface, 8610 mm, and the refracting surface, 6615 mm. Lens and mirror surfaces were polished to give an optical path length homogeneous to the required $\lambda/8$. The figures quoted here are CEDAR-W, as used in the design. The dimensions quoted in Figure 18 are relative to the CEDAR-N, and should be taken only as an example.

The chromatic corrector is also made out of Suprasil I of 20 mm thickness. The outer diameter is 320 mm and the diameter of the hole 150 mm. The radius of curvature of the entrance surface is 1385 mm, while the exit surface is flat (R > 5000 km).

The 8 optical windows are made of quartz cylinders of 45 mm diameter and 10 mm thickness. Their ring frame is made of stainless steel and the binding is obtained by heating the rings at 150^{o}C. A layer of Silver was deposited on the circumference of the quartz disk before heating. After cool-down the silver layer is compressed and produces a perfect seal. In a destructive test, a pressure of 300 bar was reached.

All optical surfaces are coated with a quarter-wavelength thickness of MgF_2 for minimum reflection at 300 nm wavelength.

The diaphragm comprises a disk with 8 elongated apertures and 8 outer and inner segments moved by right-left screw drives. The segments are bolted onto high-precision guided chariots. The 8 screw drives are provided with gears and are turned simultaneously by a gear mounted on ball bearings in a V-groove on the periphery of the disc. The opening can be varied between 0.03 and 20 mm in steps of 0.01 mm with a motor located external to the vessel. The circularity of the aperture shows radial deviations below 0.02 mm. The azimuthal opening of the diaphragm is 8 times 42.6 degrees and covers 95% of the ring.

2.2.4 Adapting the CEDAR for NA62

NA62 will use the West CEDAR filled with hydrogen at up to 5 bar (note: up to 5 bar for pressure scans is needed to see full proton peak). The nose will be shorten from 1243 mm to about 600 mm. The design of the mechanics and cooling system must accommodate the different potential choices of the optical readout system, minimise the flux of muons and neutrons hitting the photo-detectors, and ensure high levels of optical efficiency and stability, as well as minimising temperature gradients in the vicinity of the CEDAR. New photodetectors and associate electronics are needed in order to operate at the required kaon rate. A detailed description of the necessary changes and adaptations is given in the following sections.

2.2.5 Light transport system

The optics of the CEDAR light collection must be re-designed to match the photodetectors necessary to handle the high beam flux in NA62. Reflecting the light through 90 degrees will also be necessary to allow the greater space required for photodetectors and preamplifiers and to locate them so as to

minimise damage from radiation. The final design of the light collection system is highly correlated with the technology choice for the photodetectors, with different rate, noise limitations and packing density. As a consequence the light may have to be focussed or de-focussed onto the photodetector planes using ellipsoidal mirrors as illustrated in Figure 19 and a lensing system to focus the light onto the active area of the photodetectors.

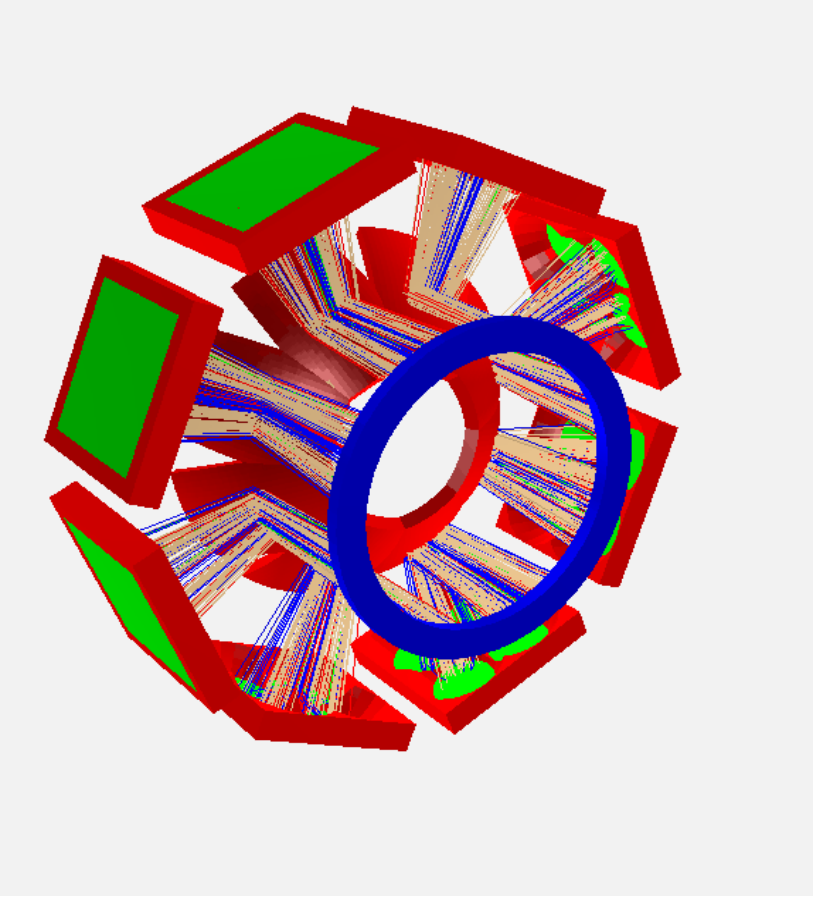

Figure 19: The Simulation of 100 events in the CEDAR nose cone. The ray tracing of the photons from (right to left) the quartz window (blue), the ellipsoidal mirrors (red), the light collection cones (red) and photodetector plane (green) is illustrated. The optical components are schematic and do not reflect a final geometrical arrangement. The photons are shown in fawn for those that are detected directly and blue if detected after reflection from the cone. Photons that are lost at reflection in the cone are shown in green whilst those photons lost by reflection at the PMT window are indicated in red.

One possible solution, adopted by the NA62 RICH detector, would be the use of light-collection cones. Again, the final design of such light collectors is dependent on the choice of photodetector technology. A proof of principle of the full optical system has to be developed for each of the photodetector options. Initial simulation studies have been undertaken to understand the design of the optical components. These preliminary studies indicate that the depth and design of the light-collection cones will be critical. A cone design that is too shallow will result in loss of photons at the phototube window due to Fresnel reflection and will suffer from light being reflected back out of the cone. With a cone that is too deep difficulties could arise in the manufacture process to ensure a

good reflective coating. The preliminary studies indicate that a photon collection efficiency of ~90% should be achievable. This design study assumed ellipsoidal collection cones, with a depth between 10-25mm; a conical design is currently under investigation. The simulation of 100 events overlaid is shown in Figure 19, the photons are generated via a full simulation of the CEDAR.

The mirrors will be fabricated from glass with radii of curvature of the ellipsoidal sections still to be optimised. The light collection system will be a bespoke solution with two approaches under consideration: machining the “cones” out of a solid block, or creating a mandrel and manufacturing the cones out of plastic, glass or carbon fibre. Further studies are in progress to ensure that the light losses in the cones due to reflections and non-normal incidence are kept to a minimum, acceptable level.

2.2.6 Mechanical Support, Cooling and Safety Considerations

The major mechanical considerations involved in modifying the Frontend of the CERN West CEDAR for use by NA62 are as follows:

- The photodetectors and optical components must be redesigned and located to maximise the capture of Cerenkov light, and mechanical rigidity must ensure optical stability.
- The preamplifiers and signal-shaping electronics must be located adjacent to the photodetectors with signal-processing electronics removed from hazardous radiation.
- The existing protective and thermally-insulated metal cover around the nose section of the CEDAR must be replaced. Cooling and insulation must be designed to minimise temperature gradients and thermal instability that might result in unacceptable local variations in the refractive index of the hydrogen gas.

The safety requirements resulting from the necessary mechanical modifications and the use of hydrogen gas are as follows:

- A nitrogen blanket around the optical-readout electronics and HV is required to eliminate any possibility of an explosion in the event of a hydrogen leak from the CEDAR. The mechanical design will incorporate sensors to monitor the flow of nitrogen and the temperature of the enclosure, as discussed below in the gas and safety sections.
- The CEDAR must be connected to the vacuum beam pipe at both ends in such a way that a hydrogen leak is not accompanied by any admixture of air in order to prevent any risk of explosion. An important secondary consideration is that damage to sensitive detectors must be minimised by mitigating the effects of hydrogen leaks and the shock wave resulting from the potential rupture of the window at either end of the CEDAR. The proposed way of doing this involves a single aluminium window between the high-pressure hydrogen and vacuum beampipe at each end of the CEDAR with a large vacuum volume to capture any escaping hydrogen. This task is being studied by the Beam group and is now completely separate from any modifications to the CEDAR nose. A preliminary schematic drawing of the pressure protection volume is shown in Figure 21.

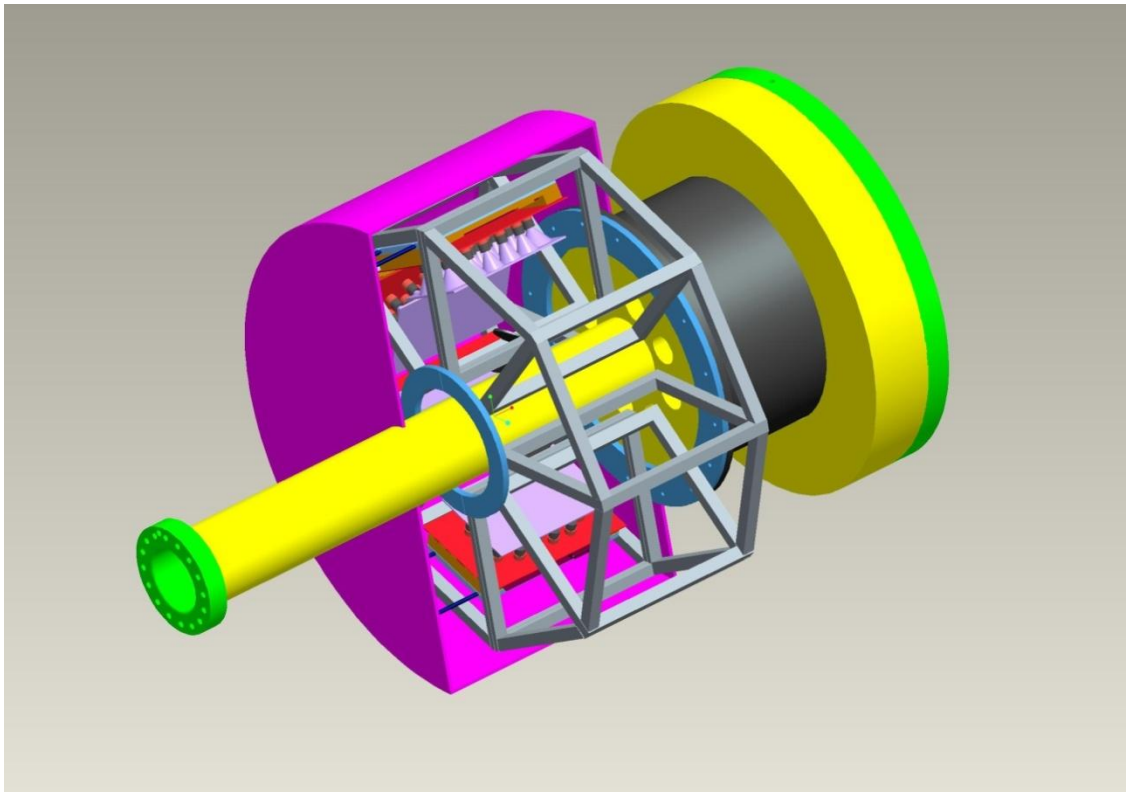

Figure 20: Conceptual Design of the Support Structure and Nitrogen Enclosure.

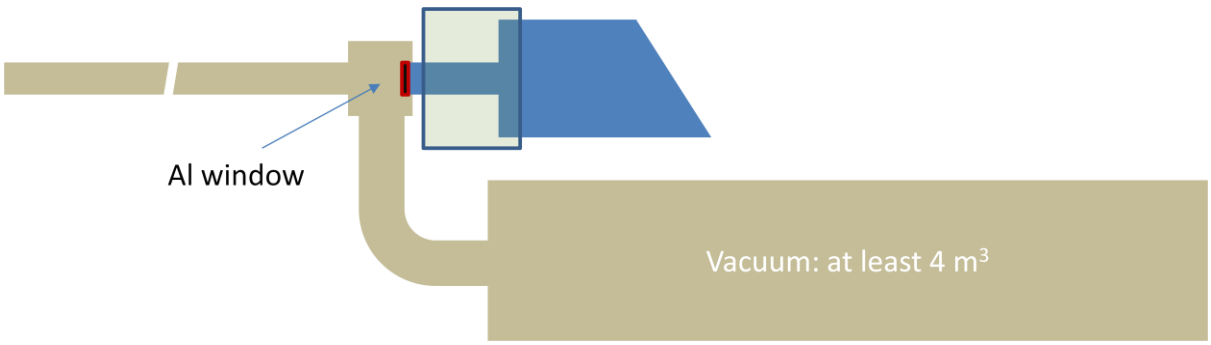

Figure 21: Conceptual drawing of the retention vacuum volume attached to the upstream end of the CEDAR(in blue). The aluminum window on the upstream side has thickness of 150 µm, the downstream window has a thickness of 200 µm.

2.2.6.1 Adapting the existing CEDAR nose

The proposed design involves replacing the current photo-detection system of 8 PMT's and the protective and thermally-insulated metal cover forming the front end of the CEDAR. All optical components internal to the CEDAR are already optimised for use with hydrogen, while the seals on the quartz windows are perfectly safe for use with hydrogen at 5 bar.

Because of their increased size and unavoidable dead space the 8 new sets of photo-detectors will need to be relocated away from the quartz windows in the nose, and light guides will be required to focus the Cerenkov light onto their active surfaces. The mechanical challenge is to preserve the current optical stability and ensure that all the Cerenkov light reaches the photo-detectors. This requires the precision mounting of the photo-detectors and optical components in a rigid, light-weight structure that can be precisely located relative to the quartz windows of the existing nose.

A disc-shaped lattice support structure housing the new photodetectors, optics, readout electronics and cooling will be cantilevered off a support cylinder bolted onto the CEDAR flange. This structure will support all cabling and cooling pipes and be enclosed in a gas-tight environmental chamber through which Nitrogen gas will be circulated. A separate protective metal cover [not shown] will enclose the environmental chamber and clip onto the metal cover enclosing the main body of the CEDAR.

In addition to removing the heat from the electronics it is important to maintain a constant temperature for all components of the CEDAR nose that are in thermal contact with the hydrogen gas in order to prevent local variations in density that will affect the refractive index. This will be done by a combination of heat removal, using chilled de-ionised water, and thermal insulation. Finally, it is important that the design ensures rapid access to, and replacement of, any PMT or electronic component that may fail during use.

2.2.6.2 The Photodetector Support Structure

Our conceptual design locates the 8 sets of photo-detectors and readout electronics into 8 pods within a lightweight, disc-shaped latticework support structure at a radius of about 30 cm from the beam. Studies indicate that the high-intensity neutron background radiation is rather uniform and that it will not be necessary to vary the radius or azimuthal angle of any of the 8 pods. A customised, ellipsoidal 45^{o} mirror to reflect the Cerenkov light from each of the 8 CEDAR quartz windows towards a set of photo-detectors will be located at the same radial position as the quartz window. A set of light guides, using either internal reflection from Plexiglas (male) cones or reflection from the surfaces of polished metal (female) cones, will channel the light onto the active surfaces of the photodetectors. The cones will be designed with variable light-collection areas to ensure approximate equalisation of the light intensity falling onto each PMT. Provision will be made for fine adjustments to the position and angles of the mirrors.

The photodetector support structure will be built in two halves that can be ‘clam-shelled’ around the beampipe and precisely located onto the support cylinder to facilitate installation. Figure 22 shows a cross-section of one half of the support structure with the envisaged layout of optical and electronic components and cooling plates. The cooling pipes and cabling [not shown] will be led around the outside of the structure and out of the environmental chamber to the floor, where appropriate patch panels are situated. Both the support cylinder and support structure will remain fixed in place after installation and the optics and electronics in each of the 8 segments will be mounted in drawers that slide forwards to enable access to the relevant components through the environmental chamber.

We envisage a separate, nitrogen-flushed, environmental chamber, made in two sections from lightweight carbon fibre, which surrounds the support structure and is supported off it. Simple non-

load-bearing seals around the beam pipe and support cylinder will ensure the necessary degree of gas-tightness. Ports will be incorporated to enable rapid nitrogen flushing to remove air from the system and to relieve excessive overpressure. The front face of the chamber will incorporate removable panels to enable rapid access to the electronics, cooling and optics within each sector. This chamber will also ensure that no extraneous light enters the optical system and will keep out any dust. The whole structure will be covered by a protective metal casing rigidly clipped to that of the main CEDAR. This casing will have front doors that can be removed or opened to enable unrestricted access to the front of the environmental chamber and hence to all components within the support structure.

The thermal environment must be controlled to ensure that the hydrogen gas in all parts of the CEDAR remains at a constant temperature and that any changes in temperature occur only very slowly and without local fluctuations. To this end the insulation within the protective cover surrounding the main body of the CEDAR will be upgraded with better performing, fire-retarding material. Thermal design of the new CEDAR Frontend will concentrate on removing the heat generated by the electronics and insulating those surfaces in thermal contact with hydrogen gas or with the main body of the CEDAR. The three regions that require particular attention are: i) the hydrogen-filled beampipe; ii) the curved surface and front face of the nose in which the quartz windows are situated; and iii) the flange onto which the support cylinder is bolted. Purpose-designed insulation will be incorporated into these different areas, as indicated in yellow in Figure 20. Details of the proposed temperature control and monitoring are given in section 2.2.6.3.

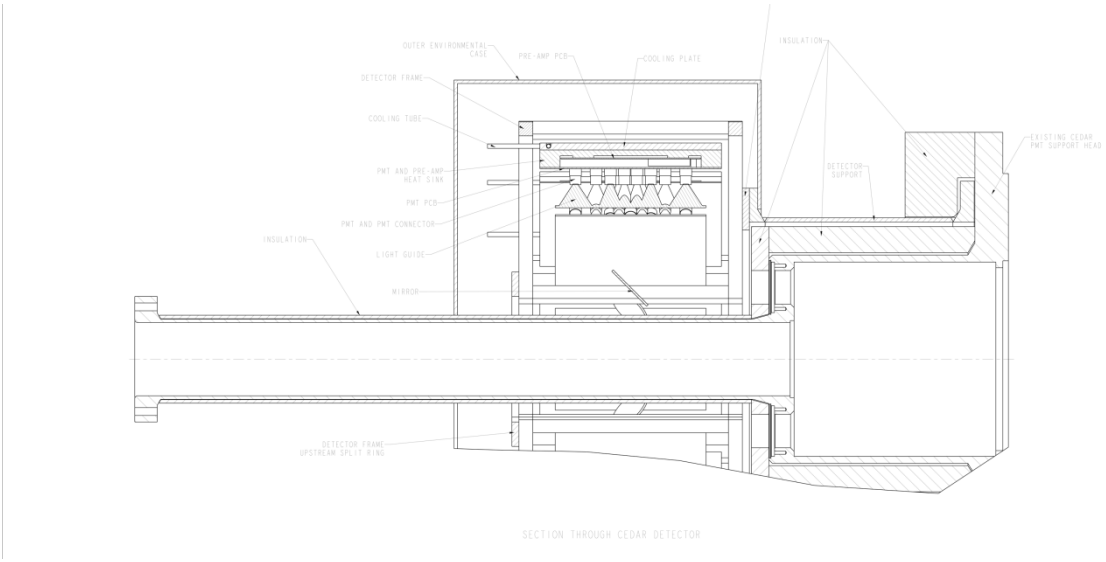

Figure 22 : Cross-sectional Layout of the Support Structure and Nitrogen Enclosure

An installation frame will be required to hold the two halves of the support structure as they are populated with optical, electronic and cooling components in the laboratory and to facilitate system tests prior to installation on beamline. It is likely that this frame will also be useful in supporting the two halves of the structure while they are being bolted to the support cylinder after being craned

onto the beamline and in facilitating the location of the environmental chamber. It is envisaged that the outer protective metal casing will be craned directly into place.

2.2.6.3 Safety Considerations, Cooling and Monitoring.

All custom made electrical equipment on the detector (photodetectors and electronics) will be enclosed within the environmental chamber that is filled with dry nitrogen gas to prevent the development of an explosive atmosphere in the event of hydrogen leaks. A small overpressure will ensure that a nitrogen atmosphere also surrounds the optical components in the support structure and prevents any possibility of moisture condensing from the atmosphere. If the nitrogen pressure drops, the power will be cut to the photodetectors and readout electronics to prevent the occurrence of any sparks within the enclosure.

The heat load from the photodetectors and front-end electronics is not expected to exceed 10 W per sector, or 100 W in total. This heat will be removed as near to source as possible by metal cooling plates, through which chilled de-ionised water flows, in thermal contact with heat sinks on the electronics PCBs. The temperature and flow of water through the cooling manifold will be controlled by electronics monitoring the temperature difference between the hydrogen beampipe close to the nose and that of the main body of the CEDAR and also the rate of change of temperature within the gas enclosure. The aim will be to maintain a stable temperature at the beampipe within $\pm 0.1^{o}$C.

It is important to monitor the long-term performance of the 45^{o} mirrors and Light Guides in case any of the surfaces should deteriorate. Early warning would enable replacement parts to be fabricated in good time. A system of 8 ultra-violet LEDs will be installed on the CEDAR nose symmetrically spaced between the 8 quartz windows. By flashing the LEDs and monitoring the output from the 8 groups of photo-detectors the overall response of the optical and electronic systems may be monitored. Factoring out the response of the electronics and photodetectors, obtained from separate monitoring, will provide a long-term record of the optical performance of all parts of the system to enable early warning of any deterioration of the optical components. Careful design is necessary to enable separation of the different contributions from the mirrors and light guides, requiring more than one signal from each group of PMT's.

All custom-designed electronics will require certification by CERN for flammable safety and all commercial sensors will need an ATEX certificate for zone 2 (non-sparking) equipment.

2.2.7 Gas system

Hydrogen gas has been identified as the most suitable radiator gas to operate the CEDAR detector. Physics performance arguments reveal two basic motivations for this choice:

- The use of Hydrogen limits the radiation length of the gas to 0.2 % of X_0 (instead of 3% for Nitrogen or 1% for Helium).
- The refractive index, its wavelength dependence and the optical properties of the gas must be suitable to provide 95% efficiency for the identification of 75 GeV/c momentum kaons.

This gas is very flammable and must not mix with Oxygen (or Air) in any place of the system. All zones exposed to the gas system will be declared as flammable gas areas (Zone 2) and all equipment used in these areas must follow the CERN flammable gas safety rules.

All exhaust gases lines must be connected to the gas extraction system in TCC8.

2.2.7.1 Outline of the Gas System Design

The Hydrogen gas supply is located in the gas surface building 920, all other gas system components are situated in the experimental area (TCC8) close to the CEDAR detector itself. The relevant detector and gas-system parameters are summarized in Table 12.

Table 12: Detector and gas system parameter

Radiator gas volume	1.1 m^3
Operating gas	100% H_2
Operating pressure	Between 2.5 and 3.5 bara[18]
Pressure scans	Vary pressure in small steps around peak
Smallest pressure step during pressure scan	5 mbar
Absolute pressure (density) accuracy	<< 0.1%
Detector leak rate	< 5.0 * 10^{-4} Std. cc/s

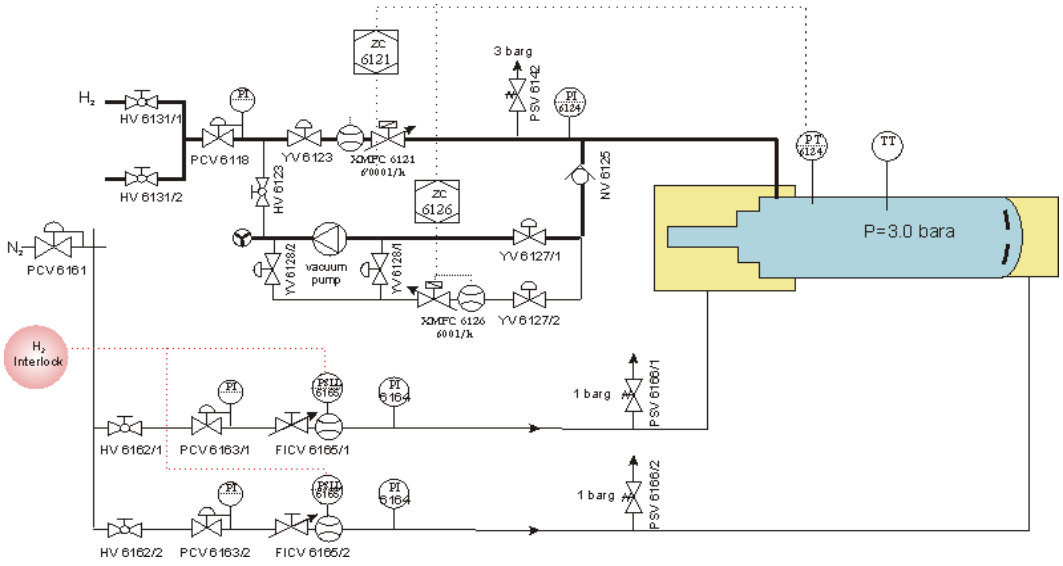

Figure 23: Layout of the CEDAR gas system (blue = CEDAR radiator volume, yellow = envelope volumes)

[18] The unit bara is used to denote that the scale refers to an absolute pressure referenced to vacuum.

The CEDAR radiator comprises a gas volume of about 1.1 m^3, which will be pressurized during operation to a constant pressure to detect kaons (≈ 3.6 bara). Unlike other gas systems, Cherenkov radiators do not need gas circulation; they can in fact operate with a fixed amount of gas without any gas renewal. This concept is advantageous because gas turbulence and temperature variations inside the detector will in this case not influence the gas density, providing temperature gradients are minimised.

The CEDAR gas system is shown in Figure 23. The input and output gas is metered using mass-flow-controllers, the absolute pressure inside the radiator is measured by means of a high accuracy transmitter, which has an absolute precision of 0.01% of full scale. The gas pressure and flows are monitored and controlled by a programmable control computer (PLC) allowing remote filling or emptying of the detector to any desired pressure. For safety reasons, some external parts of the detector will be enclosed by envelope volumes and be flushed with Nitrogen gas evacuating any Oxygen in the direct vicinity of the detector. In case of insufficient N_2 flow the Hydrogen fill flow will be interrupted.

2.2.7.2 Filling and Emptying of the Detector

Before introducing any H_2 gas, the entire system (including the detector) is filled with clean and dry N_2 gas to evacuate air from the system. The detector is then pumped to vacuum using a dry membrane vacuum pump. If the detector has reached a vacuum pressure below 5 mbar the Hydrogen gas can be admitted to the detector and be filled to the desired pressure. At the end of a physics run the Hydrogen is pumped out and the detector is refilled with nitrogen.

2.2.7.3 Running mode and Pressure Scan

During physics operation the radiator pressure will be set accurately to the desired operation pressure, e.g. 3.73 bara. The gas density is then left constant by closing all valves to the radiator volume. To guaranty sufficient stability for long physics runs the leak rate of the detector must be smaller than 5.0 * 10^{-4} Std. cc/s. If this leak tightness cannot be reached the system can be upgraded with a density feed-back loop, so that the gas system compensates the gas losses automatically.

It is foreseen to perform pressure scans at regular intervals to calibrate the detector response around the kaon peak. To do this, the chamber pressure is varied in small steps ($\Delta p \geq 5$ mbar) between two limits to determine precisely the pressure with the highest counting rate for kaons. To maintain the best possible temperature stability, the pressure scan starts at the highest pressure and releases successively small amounts of gas from the detector.

2.2.8 Photodetectors

The CEDAR detector is required to achieve a kaon tagging efficiency of at least 95%, with a time resolution of 100 ps/kaon. The intrinsic rate due to the kaon component of the beam is approximately 50 MHz, and each kaon produces an average of 100 photons at the exit of the quartz window (i.e. losses due to internal mirrors and lenses are already included). The light spots produced by the CEDAR optics are a set of eight 30 x 8 mm^2 rectangular areas at the position of the quartz exit

windows, as illustrated in Figure 24. These photons will be guided onto the active surfaces of the photodetectors, as described earlier.

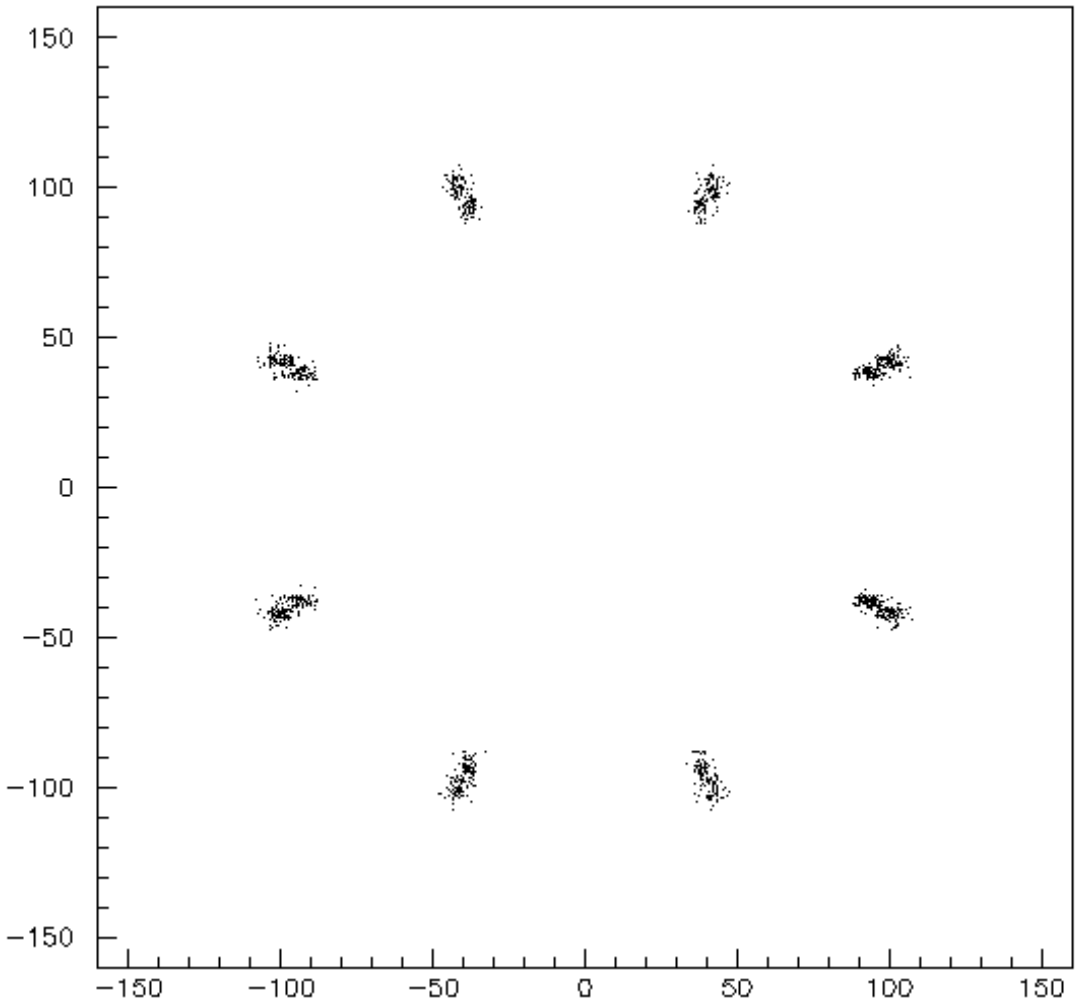

Figure 24: Distribution of photons on the 8 quartz windows at the exit of the CEDAR

At present the baseline choice of photodetector to be used for the CEDAR readout is conventional photomultipliers, particularly of the type used for the RICH detector: Hamamatsu R7400U-03. The number of photomultipliers will be chosen to maximise the kaon efficiency, while minimising the dead time and accidental noise.

Current technology limits the mean rate of single photons that can be detected efficiently by a detecting surface to be of the order of 4 MHz mm^{-2}. Assuming a photodetector with a gain of 10^6 and a quantum efficiency of 25% at the peak for ~420 nm (~20% average), this would correspond to an equivalent anode current of 0.1 $\mu A\ mm^{-2}$.

The light impinging on a spot is not uniformly distributed but is concentrated in the centre of the spot, as shown in Figure 8. However, if 30 PMTs are used to cover the spot, the hottest region will experience a rate of 4-5 MHz per PMT, which is within the PMT tolerance.

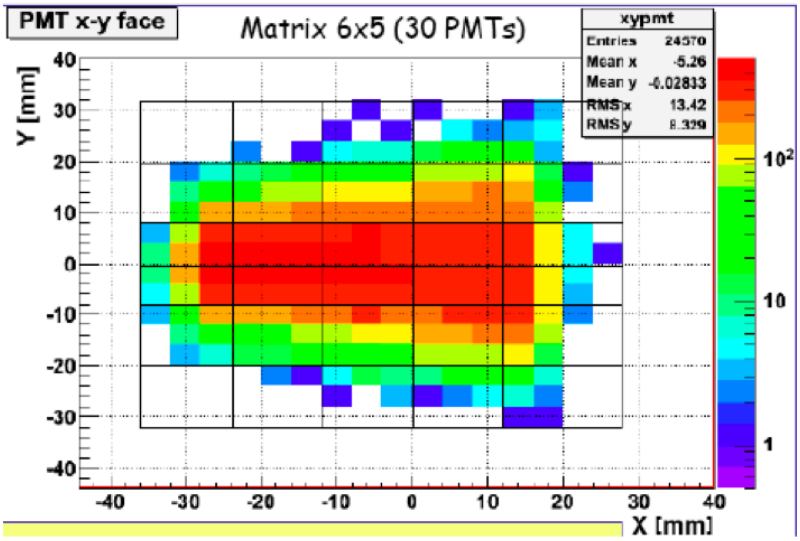

Figure 25: Distribution of photons on 1 window at the exit of the CEDAR

With 100 photons/kaon distributed among 8 spots, there are an average of 24 photons when convoluted with the quantum efficiency as a function of the wave length. Therefore an average of 3 photo-electrons per spot would be detected for a typical detection efficiency of 20%. For a typical photo-detector time resolution of 300ps the single kaon crossing time would therefore be measured with a resolution of 300ps/√24 =60ps. When taking into account ~80% efficiency for additional lenses, mirrors and cones, the kaon detection efficiency with 3 photo-electrons/spot and a coincidence of various spots should easily exceed the required 95%. The coincident detection of photons will also address concerns about the effect of PMT thermal noise.

Vacuum photomultiplier tubes have a high resistance to damage by radiation and a relatively low dark-count rate. In addition PMT's with quartz windows are sensitive to short wavelengths. The large sensitive area simplifies the optics, since the sensitive region of each PMT is large enough that light collection cones can be used to channel the light, thereby reducing the effect of dead areas and making optimal use of the quantum efficiency. However the required combination of ~100 pico-second timing, high rate and long service life push the requirements to the limit of what is currently possible. Careful evaluation of other vacuum PMT options will be necessary, as well as an optimization procedure to establish the minimum number of PMT's necessary to achieve the desired performance.

Other options that could be considered are a single micro-channel-plate (MCP) photomultiplier and traditional multi anode technology. Table 13 lists some of the characteristics of the different PMT options. All devices have quartz, or UV-glass, faceplates with a short wavelength cut-off of approximately 160nm, or 185nm, respectively, and a quantum efficiency peaking at about 420nm.

The R5900 device is also available with an ``Ultra Bialkali'' photocathode with a peak quantum efficiency of 43% at 330nm.

Table 13 Comparison of alternative phototube options with respect to the Hamamatsu R7400P (first row).

Design Option	***Tube***	***PMT per port***	***TTS (FWHM)***	***Peak QE***	***Peak λ***	***Max anode current per device***	***Nominal Gain***	***Anode per device***
Array of single-anode PMT's	Hamamatsu R7400P	30	280ps	18%	420nm	100μA	10^6	1
Multianode PMT	Hamamatsu R5900U-00-L16	8	180ps	21%	420nm	100μA	$4x10^6$	16
Multianode MCP-PMT	Photek PMT240-96	1	100ps	15%	420nm	1.0μA	$2.5x10^5$	96

Using a multi-anode device would reduce the count rate on each channel and avoid issues with inefficiency due to pile-up. With 8 multi-anode PMT's, a photo-cathode quantum efficiency of ~25% and a gain of ~10^6 the total anode current will be approximately ~3 μA, well within the 100 μA limit. *To verify the real gain of this device, more studies would be advisable.*

Microchannel plate PMT's offer high timing accuracy and fast rise times. Transit time spreads of ~30 ps have been reported for single anode devices, although the response time for multi-anode device is ~100 ps; the mulit-anode device is also current limited. The response of MCP-PMT's is not linear, however, and there are local dead-spots after a pulse as the MCP recharges. Careful investigation of gain with pulse rate would therefore be needed. A MCP lifetime can be expressed in terms of C/cm^2 and is typically about ~1 C/cm^2. However, a tube from Photonis has been shown to give a significant reduction in gain after only 0.4C/ cm^2. The annual requirements from NA62 are ~3 C/cm^2, well in excess of the expected limitation. It is known that this limit can be increased significantly using thin film fabrication techniques. The lifetime of any proposed MCP-PMT would have to be evaluated, especially since the PDE (typically 15%) is rather low and any reduction due to aging would be problematic.

Considering the various characterists outlined below, the R7400 is the default solution. However, ideally all PMT options should be evaluated and optimised against:

Test Regime. The photodetector efficiency (PDE) should be evaluated in the laboratory, specifically for light with a Cherenkov emission spectrum. Characterisation of the chosen devices should be measured and the dark count rate verified for the widest range of working conditions anticipated at NA62. The ability to distinguish two kaon signals close in time should also be carefully addressed,

using a pulsed laser, since this will have an impact on the choice of the front-end readout electronics for the photo-detectors. Additional tests will be needed to measure gain stability with respect to count rate, and accelerated lifetime testing will also be performed. Timing resolution measurements will also be carried out.

Anode Current. Although the passage of a kaon only produces an average of about 3 photons at each optical port, the high kaon rate (50 MHz) results in an illumination of 600 million photons per second. Photomultipliers are rated at a maximum anode current of about 100 μA and a gain in the region of 10^6. Coping with such high illumination implies tiling several photomultipliers together to reduce the photon flux at each device.

Count rate per channel. Assuming a double pulse resolution of <10 ns, a count rate per channel of less than 10 MHz is needed in order to keep the inefficiency due to pile-up below 5%. Hence, it will be necessary to divide the light between several PMT's per port. In addition, multi-anode readout will tested as a way to reduce the count rate on each channel to a lower level.

2.2.8.1 Front-End electronics

The typical signal output (produced by one photon) from the baseline choice of PMT's to be used for the CEDAR is about 100 μA peak, with a rise time of about 0.6 ns and a FWHM of about 2 ns. In the Front End Electronics, this signal will be amplified about 50 times, and then discriminated and time-stretched by a NINO ASIC to produce a LVDS signal. The NINO was designed for analogue signals from devices faster than PMT's, and requires a differential input. A suitable preamplifier/shaper board has been designed to adapt the PMT signals to the NINO input for the RICH detector. It will be placed on a PCB mounted on the PMT's themselves, and must be radiation hard[19], while the rest of the electronics will be placed far enough from the beam not to require radiation hardness.

The measurement of the Time Over Threshold, implemented using the NINO chip and the time measurement of both pulse leading and trailing times in the common TDC/TEL62 readout system, will be used to extract information on the pulse amplitude, both to correct the time-slewing induced by amplitude fluctuations, and to discriminate against pile-up. Considering the PMT analogueue signal to be of Gaussian shape, it can be shown that edge time and time over threshold are linearly correlated. It is possible in principle also to predict the correlation between time over threshold and amplitude.

Using 240 PMT's and a quantum efficiency of 25%, the expected kaon and photon fluxes lead to an average photo-electron rate of 3.8 MHz on a single PMT channel. The average number of detected photons per kaons is 18, while the probability of having more than one photon on a single PMT is of the order of 8% in one event and 9% in two consecutive events spaced by less than 20 ns.

[19] The expected radiation levels near the PMT's are approximately 4 Gy/year.

Considering that the PMT intrinsic time resolution is 300 ps, and that a time resolution of about 100 ps on the single kaon is needed in order to suppress the accidental background, a minimum of 9 photons need to be detected for each kaon. The dead time losses are of the order of 20 ns (11ns from NINO and 5 ns from HPTDC, and allowing for ~4ns duration of the signal). The double pulse resolution of the whole system is dominated by the electronics and is of the order of 10 ns.

2.2.8.2 Readout Electronics

The readout for the CEDAR will be based on the common TDC/TEL62 system. The digital signals will be sent from the front-end electronics to the readout electronics with a few metres of high-quality twisted pair cable. The times of both the leading and trailing edges of the pulse will be measured.

Concerning dead time losses, the NINO chip introduces a stretching time of 11ns, while the HPTDC has a deadtime which can be set to be 5ns but it does not contribute to the total deadtime due to the NINO mode of operation. The major limitation arises from the finite size of the HPTDC hit channel buffers , which can sustain a maximum of 40 MHz over a set of 8 channels, or a maximum of 10 MHz for a single channel, whichever is lower, on the assumption of 1 word written for each signal (or 5 MHz for 2 words).

The 16 HPTDCs in a single TDC/TEL62 boards are arranged in sets of 8 channels which share some group channel buffers. By using only 2 channels for each set (1/4 of the total number available) the rate will be reduced to a sustainable level to guarantee a small probability of hit loss, at the price of using only 32 channels per TDC board (8 channel per TDC; 128 channels per TEL62) . Hence the readout system will consist of 16 TDC boards and 2 TEL62 boards. The hit loss probability estimated in this configuration is 1%, which corresponds to a detection inefficiency per kaon of a few % for a minimum of 10 photons distributed in at least 6 spots. However, further studies are in progress to ensure that the light distribution on the spot is compatible with a rate per channel not exceeding 5 MHz.

The contribution of the readout system to the time resolution is estimated to be of the order of ~70 ps, based on the performances of a prototype system used in the RICH tests), which gives a negligible contribution when combined with the intrinsic PMT resolution of 300-400 ps.

The readout time window will be 75 ns, and an average of 4 kaons are expected to appear in such time window (with a tail up to 12 kaons). With an average of 18 hits per kaon, and 32 bit per photon, this corresponds to an average of ~288 bytes in the readout window (and a maximum of 864 bytes). Assuming a trigger rate of 1MHz, this corresponds to an average of ~0.3 GB/s of readout rate (and a maximum of 0.9 GB/s). This will be distributed on 2 TEL62.

Data arriving in an interval of 1 ms will be saved in the TEL62 buffer at any time. Given an average of 9 hits per kaon per TEL62, with 32 bit per photon and a kaon rate of 50 MHz, a maximum input of ~40 bytes is expected per kaon per TEL62, which corresponds to ~200 KB in the buffer of 1ms - well below the maximum buffer of TEL62.

The processing done in the TEL62 will include counting of the number of PMT's fired per spot and the number of fired spots, to enable the development of algorithms that use the multiplicity and the

geometry to suppress the background. It is also foreseen to record the photon times in 1ns slots inside the 75 ns readout window to allow for a precise time coincidence with the trigger signal.

2.3 Gigatracker (GTK)

2.3.1 Detector Requirements and Performance

The GigaTracKer (GTK) is a core part of the NA62 experiment. It is a spectrometer that provides precise measurements of momentum, time and angle of the incoming 75 GeV/c kaon beam. Beam measurement is essential for the selection of the $K^+ \to \pi^+ \nu\bar{\nu}$ decay and the background reduction will be achieved by the kinematical constraint based on the missing mass variable:

$$m^2_{miss} \approx m^2_K \left(1 - \frac{|P_\pi|}{|P_K|}\right) + m^2_\pi \left(1 - \frac{|P_K|}{|P_\pi|}\right) - |P_K||P_\pi|\theta^2_{\pi K}$$

In the equation above $|P_K|$ is entirely derived from GTK while the angle, $\theta_{\pi K}$, is derived from the measured directions of the decaying kaon given by the GTK and of the pion from the decay, which is measured with the downstream spectrometer, the STRAW detector, which also provides the measurement of $|P_\pi|$. This is symbolically indicated in the decay sketch in Figure 1 on page 11.

The Gigatracker is composed of three stations (GTK1, GTK2 and GTK3) mounted around four achromat magnets as shown in Figure 26.

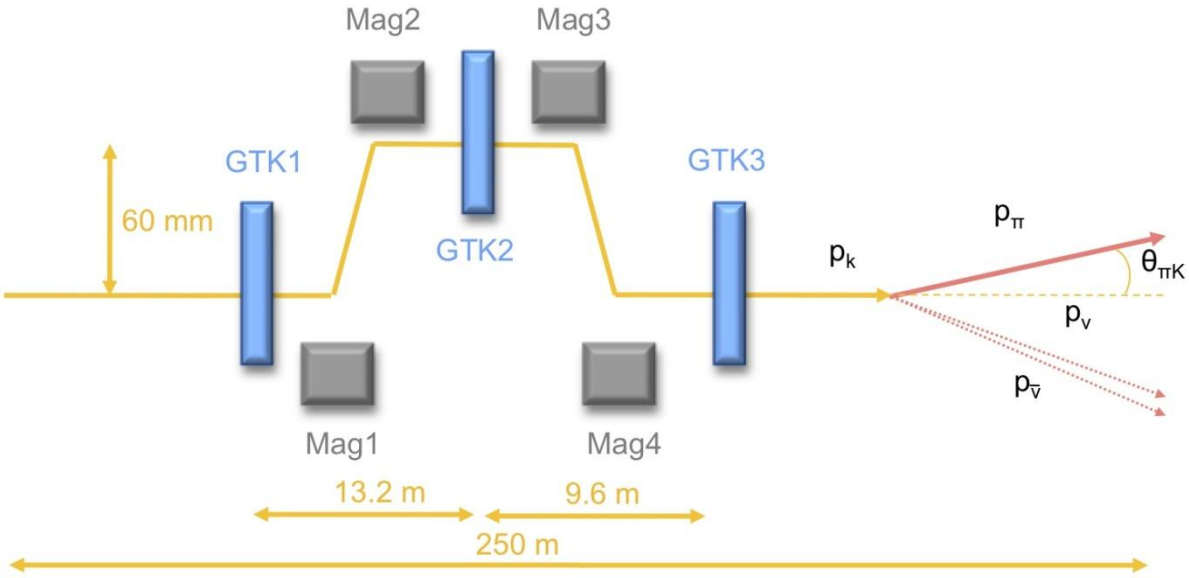

Figure 26 Layout of the Gigatracker stations and sketch of the decay.

This setup is placed along the beam line inside the vacuum tank, just before the fiducial region in the decay vacuum pipe. It has to sustain a high and non-uniform beam rate and has to survive in high radiation environment.

The overall simulation of the kaon decays in the detector (see section 1.1) has determined the beam track momentum and direction resolution requirements. From this study and taking into account the expected STRAW resolution, it has been derived that the Gigatracker has to measure the momentum with a relative resolution of $\sigma(p_K)/p_K \sim 0.2$ % and the direction with a resolution of the order of 16 μrad. A pixel size of 300 μm x 300 μm is sufficient to achieve the required resolution.

Finally the beam spectrometer has to sustain a high and non-uniform beam rate of 0.75 GHz in total, hence the name Gigatracker, with a peak of 1.3 MHz/mm^2 around the centre.

To meet these resolution requirements the beam spectrometer is installed in vacuum with a minimal amount of material crossed by the beam to preserve the beam divergence and to limit beam hadronic interactions. The spectrometer is composed of three hybrid silicon pixel stations. Each station is made of one hybrid silicon pixel detector with a total size of 63.1 mm x 29.3 mm containing 18000 300 μm x 300 μm pixels arranged in a matrix of 90 x 200 elements. With this configuration the detector matches the expected beam dimensions of ~60 mm × 27 mm. The pixel dimensions and the distances between stations are adapted to deliver the required momentum and direction resolution.

The amount of material crossed by the beam at each station influences the angle measurement. The chosen sensor thickness of 200 μm corresponds to 0.22 % of a radiation length, X_0. The design efforts take into account to minimize the material as much as possible as the physics performance strongly depends on a low material budget. Giving an equivalent budget to the read-out and to the cooling the total amount of material per station has been required not to exceed 0.5% X_0.

Furthermore, due to the high intensity, the required time resolution on every single track using all three stations is 150 ps (rms).

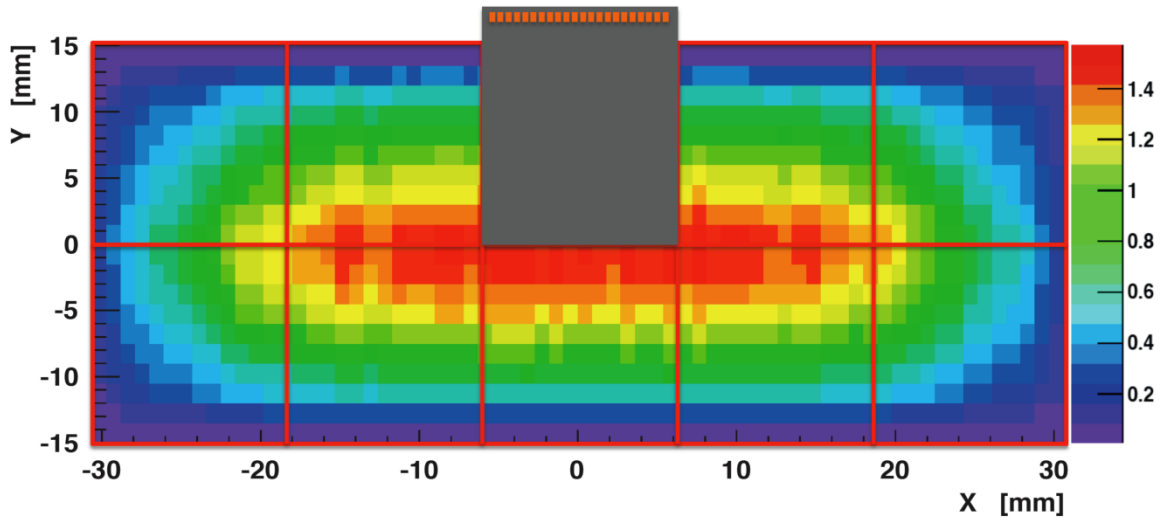

Figure 27 Beam intensity distribution over GTK station 3 (units are MHz/mm^2). One of the 2 x 5 read-out chip is drawn.

This is a novel requirement for a silicon pixel detector and none of the existing systems has such a capability.

The beam region is completely covered by one single, fully active silicon sensor bump-bonded to 2 x 5 read-out chips (ROCs), which each deliver a time resolution of better than 200 ps for each pixel hit, allowing to achieve a track resolution better than 150 ps. A description of the system is reported in an internal note [26]. Figure 27 shows the intensity distribution in station 3.

The expected fluence for a typical 100 days run time year is 2 × 10^{14} 1 MeV n equivalent per cm^2 in the central region of the sensor. This value is comparable to those expected in the inner layers of the LHC trackers over ten years of operation and requires the sensors and the application specific

integrated circuits (ASIC) to cope with the high radiation environment. Section 2.3.3 describes the measurements performed on sensors. In order to reduce the radiation induced leakage current on the relatively large sensor the operating temperature is set to 5^{o}C or below. The very low mass of the detector, the operation in vacuum and the need of limiting the leakage current increase due to radiation damage demands a very efficient and reliable cooling system.

Table 14 shows the principle detector specifications. These parameters are to be seen as acceptable values. The research efforts described in the following sections aim for improvements with respect to these values especially for the time resolution, the material budget and the efficiency and will result in a largely improved physics performance. In the following sections specifications derived from these parameters described.

Table 14 Detector performance requirements.

Material budget per station	≤ 0.5 % X_0
Overall efficiency per station	> 99 %
Dynamic range	0.6 -10 fC; 5'000 – 60'000 e^-
Pixel size	300 µm x 300 µm
Time resolution	< 200 ps rms

.

2.3.2 Detector Elements and Architecture

Each component of the hybrid silicon pixel detector is optimized in terms of material budget and performance. Hybrid pixel detector modules establish electrical connections between the silicon sensor and the read-out chips using Sn-Pb solder bumps. For the GTK, the sensor thickness is 200 µm and the thickness of the read-out chips is 100 µm. The stations are located in vacuum. For the on-detector system integration the following parameters are optimized.

- Material budget in active beam area.
- Read-out chip size and active matrix size.
- Accessibility to power and I/O connections on the chip.
- Particle rate per pixel.
- Heat dissipation and cooling.

In order to minimize material and maximize geometric efficiency in the active beam area any connections to the read-out chip are outside the beam area. The beam profile has been adapted in such a manner that two adjacent rows of read-out chips cover the beam area. Figure 27 illustrates the principle. The ASICs are accessed only from the top or the bottom of the assembly. Thus wire bond connections to the read-out chip are placed only on one side of the ASIC. The active pixel matrix extends to the other 3 sides of the chip, allowing a maximum of efficiency in the beam centre with the highest particle flux.

Two options for the heat transfer are under investigation. One is based on a silicon substrate with micro cooling channels. The other is based on placing the module in a vessel with gaseous cooled nitrogen flowing. The module is placed in vacuum and the heat dissipated by the 10 read-out chips is estimated to be less than 2 W/cm^2 for a total of 32 W. The size of the active pixel matrix on each chip is 13.5 mm x 12 mm, corresponding to 45 rows x 40 columns of pixels. On the chip and outside the active pixel matrix the end-of-column circuitry performs data processing, control, configuration and off-chip data transfer.

An important design aspect of large pixel matrices, where power supply is provided from one side only, is the power drop along a column. The power supply metal plane layout has to be carefully considered. For chips in 0.13 µm CMOS technology and a column length of 13.5 mm with a power consumption of 2 W/cm^2 simulation studies indicate that powering from one side can be done.

In the configuration presented the material budget is uniform over the entire active area. Table 15 summarizes the material budget of the active components in the beam. The sensor and system parameters are discussed in sections 2.3.3 and 2.3.4.

Table 15 Material budget of sensor, read-out chip and bump bonds.

Component	**Material**	**Thickness [µm]**	**X_0 [%]**
Sensor	Si	200	0.22
Bump Bonds	Pb-Sn	≈25	0.001
Readout Chip	Si	100	0.11

Figure 28 shows the system block diagram. The GTK assembly (sensor and 10 pixel chips) is electrically connected via wire bonds to the GTK assembly carrier printed circuit board (PCB), which connects the ASICs to high/low voltage power, and to the serial control and to the status and data interfaces of each chip using dedicated signal traces. Special care must be taken with the layout as the power supplies (LV, HV up to 600 V), the common station ground plane, on-chip thermal sensor connection and up to 40 Gbit-serial links are in close proximity. The assembly carrier also acts as a vacuum feed-through. Inside the vacuum, only filtering capacitors, a temperature and humidity sensor but no active components are placed on the assembly carrier. In order to keep the stations electrically isolated from the control room and increase noise immunity, the data transmission to the off-detector read-out electronics is over optical fiber. Thus optical/electrical converters for the serial links to the control room are placed outside the vacuum. Local LV regulators are being considered. However, this will limit or complicate the remote control capabilities of the LV supply of the read-out ASICs. Depending on the radiation effects on the ASICs, compensation on the power supply voltage might be appropriate.

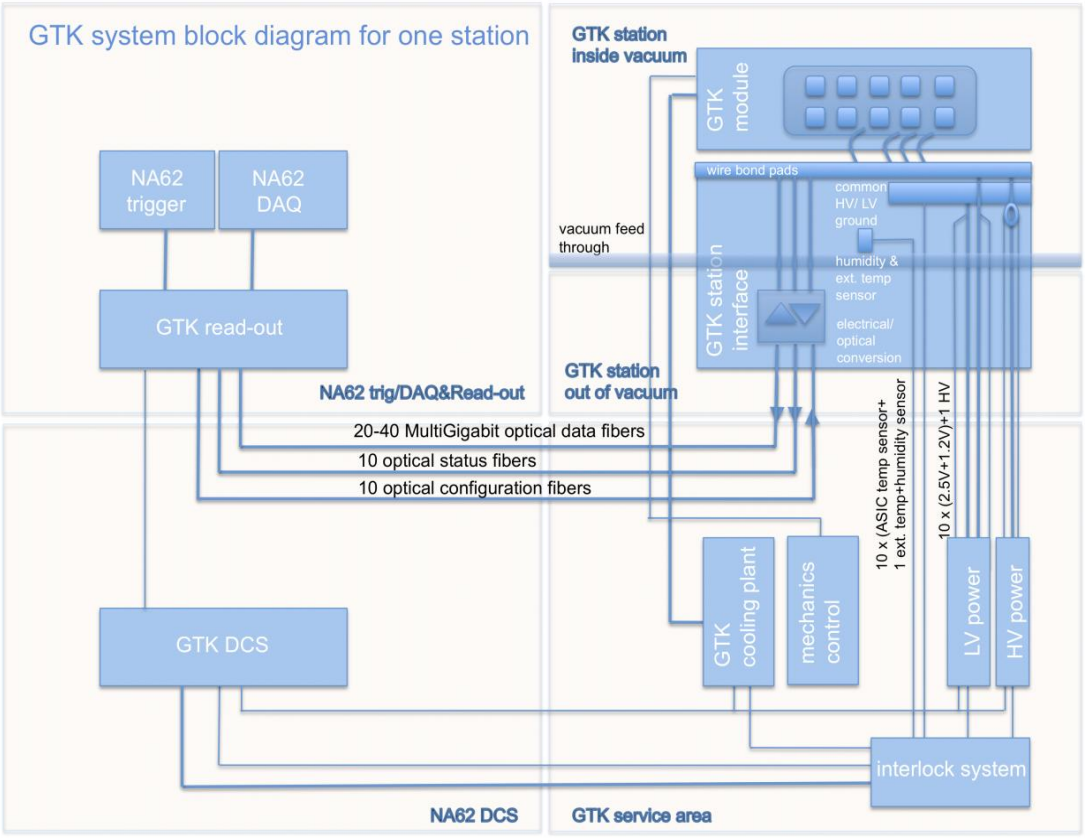

Figure 28 GTK system block diagram.

2.3.3 Silicon Sensor and Bump-Bonding

2.3.3.1 Silicon Sensor

The silicon sensor thickness has been studied in simulation, see section 2.3.4, and an overall thickness of 200 µm was found to be compatible with the signal requirements for precise timing information together with the material budget constraints.

A detailed study of the radiation damage effects was carried out using test-diodes produced by the manufacturer that will also deliver the final GTK sensor wafers (FBK-Trento, Italy). Irradiations were carried out up to a fluence of 2 x 10^{14} 1 MeV n_{eq} cm^{-2} and detailed annealing studies investigated the effects on leakage current and full depletion voltage. The results have been summarized in an internal NA62 note [27]. The maximum acceptable operating temperature for the GTK stations has been set to 5°C which allows an operation time of approximately 50 days using standard p-in-n sensors as measurements on irradiated diodes indicate that the leakage current of a full size sensor, at an operating temperature of 5°C and 50 days of operation, will reach a maximum acceptable value of approximately 270 µA. However, efforts are being made to design the cooling system to operate at -20°C, which allows an operation of 100 days. Continuous monitoring of the sensor leakage current will allow the radiation induced leakage current increase to be followed and also provide an estimate of the overall fluence. A precision of 100 nA or better seems to be adequate.

The configuration of the experiment will permit the exchange of the Gigatracker stations at the end of their operation period. The baseline choice for the sensor material is currently p-in-n due to

simpler processing procedures, availability and experimental arrangement. However, the front-end design will allow the study of the connection of other sensor types, such as n-in-p or similar materials.

A mask design layout has been defined and a batch of p-in-n sensors has been delivered as 4" wafers. The layout of the wafer is shown in Figure 29, indicating the large silicon sensor for the full-size stations in the centre, surrounded by test sensors foreseen to be used with the prototype electronics and general test structures. A detailed list of the individual sensor parameters is given in Table 16.

The wafer substrate is n-type (phosphorous doped) with a resistivity of about 4-8 kΩcm. The wafer thickness is (200±10) μm and the bow (after processing) is less than 30 μm to comply with the flip chip bonding requirements. The wafer back side is metalized to allow wire bonding contacts for applying the reverse bias to the sensor.

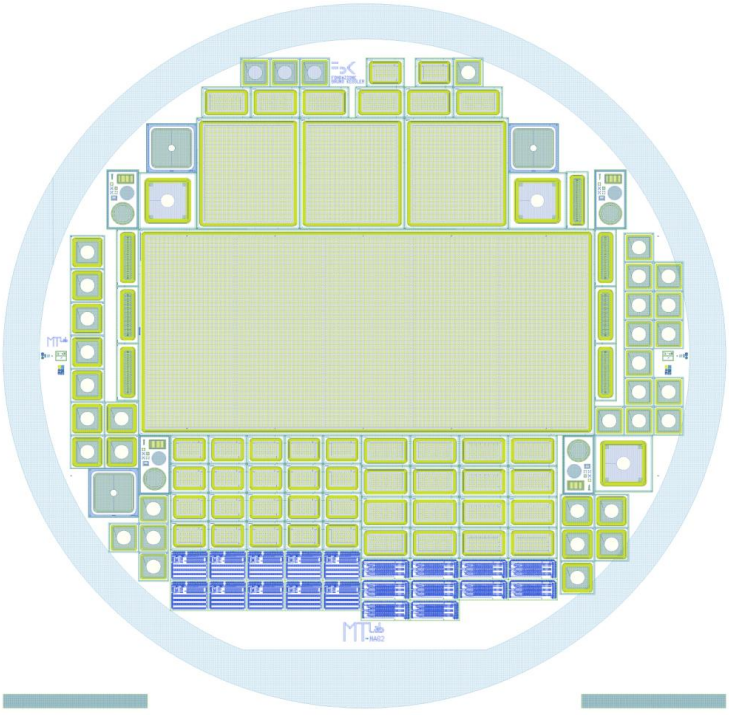

Figure 29 CAD drawing of sensor wafer.

Table 16 Main sensor elements on wafer

Number	Structure
1	**Full-size sensor:** • 18000 pixels (90 × 200 matrix) • 27.0 × 60.8 mm^2 active area • 29.3 × 63.1 mm^2 effective size on wafer

3	**Single-chip size sensor:**
	• 1800 pixels (45 × 40 matrix)
	• 3 types (see Figure 29):
	  • A: 300 µm × 300 µm pixels everywhere
	  • B and C: enlarged pixels in lateral columns (300 µm × 400 µm)
	• 13.5 × 12.0 (12.2 for B and C) mm^2 active area
	• 15.8 × 14.3 (14.5 for B and C) mm^2 effective size on wafer

22	**EOC prototype sensor:** • 60 pixels (6 × 10 matrix) • 1.8 × 3.0 mm^2 active area • 4.1 × 5.3 mm^2 effective size on wafer
22	**P-TDC prototype sensor:** • 105 pixels (7 × 15 2.1 × 4.5 mm^2 active area • 4.4 × 6.8 mm^2 effective size on wafer

In order to achieve the required timing precision it is necessary to operate the GTK sensors at high over-depletion. A target operating voltage of approximately 500 V has been defined and a multi-guard ring structure is implemented in every sensor in order to ensure stable operation at relatively high voltage.

Charge collection efficiency, speed and sensor capacitance influence the overall timing performance of the GTK station. Ideally the sensor characteristics should not degrade timing. However, under irradiation of the sensor the collected charge degrades. This degradation is non-uniform from chip to chip and pixel to pixel and may introduce non-uniform pixel time response.

2.3.3.2 Bump Bonding

Each pixel cell contains one bump pad to connect it with the corresponding cell in the front-end chip. In addition the guard ring of each sensor needs to be connected to the corresponding pads on the front-end chip using bump bonds. The bump-bonding pads have octagonal shape (26 µm with 20 µm opening in the passivation) and the bump pad centre is located at about 50 µm from the pixel edge. The bump-bonding pads are arranged in a mirrored scheme on neighboring columns (Figure 30) with enlarged pixel cells (300 µm x 400 µm) in the interchip region. This allows a fully sensitive area over the full sensor size to be maintained. A schematic drawing of the implemented geometry is shown in Figure 30 and Figure 31.

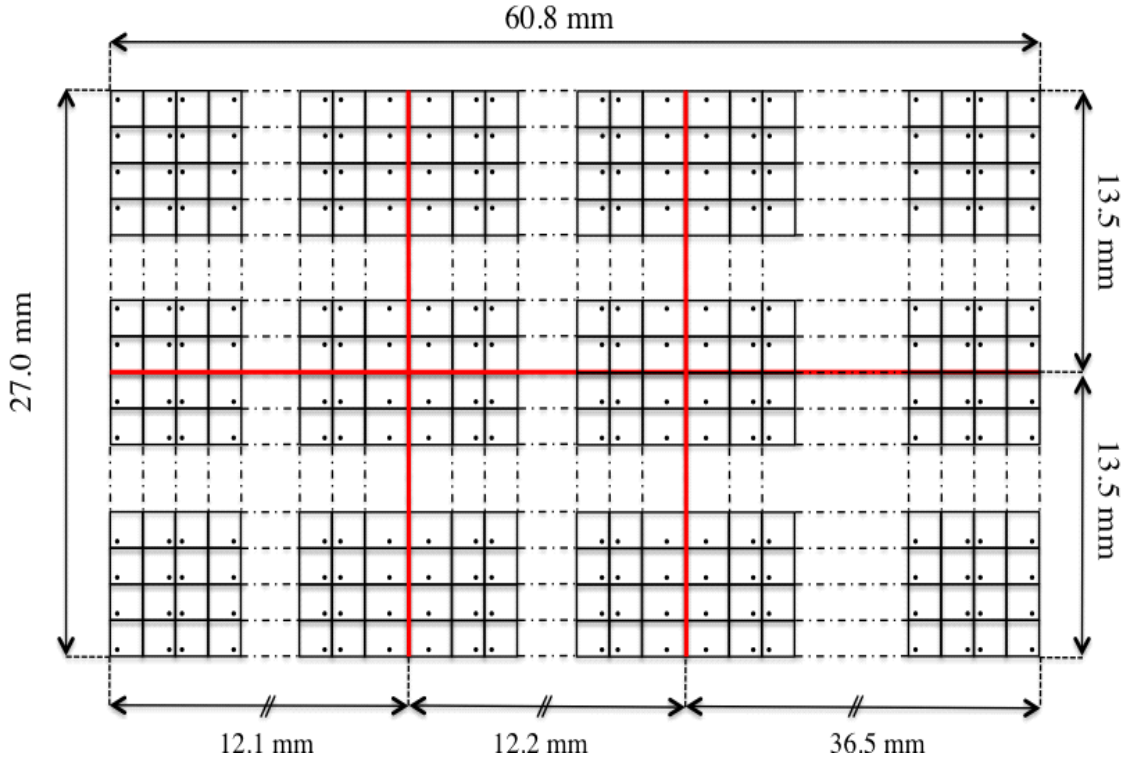

Figure 30 Bump pads arrangement scheme.

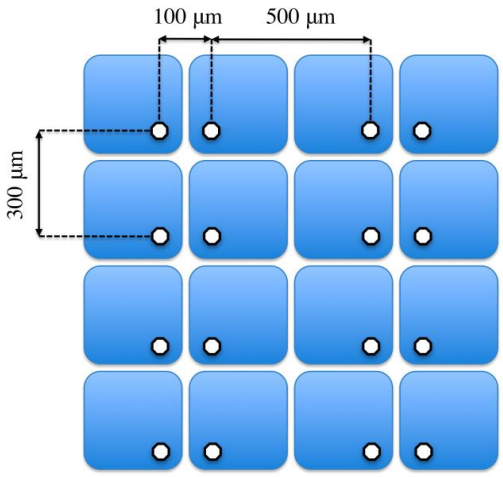

Figure 31 Schematic layout of the full-size sensor. Enlarged pixel cells (300 µm x 400 µm) are put in the border region between neighboring read-out chips.

Prototype assemblies with demonstrator ASICs containing only 1-2 columns and final GTK assemblies with ASICs containing the full matrix will differ in the way that the prototype chips will not be available on full size wafers. Generally the sensors will be made available on 4" wafers of 200 µm thickness. The prototype front-end chips will be available as single dies from multi project wafer runs whilst the final GTK front-end chips will be available on 200 mm diameter wafers. This will require two different technical solutions to build flip chip bonded assemblies for the prototype and for the final GTK planes.

A dedicated market survey was launched in July 2009 to identify companies capable of flip chip bonding the prototype assemblies and the final GTK assemblies. The two companies that fulfilled the

requirements defined the survey carried out at CERN (IZM Berlin and VTT Helsinki) have been invited to reply to a price enquiry for the prototype production. Following the price enquiry the contract for the production of prototype bump-bonded assemblies was awarded to IZM Berlin.

Specific emphasis is laid on the thinning of the final GTK readout chips in order to meet the material budget requirement. A final thickness of 100 µm or less will comply with the requirements, but will also demand the development of specific process steps during production.

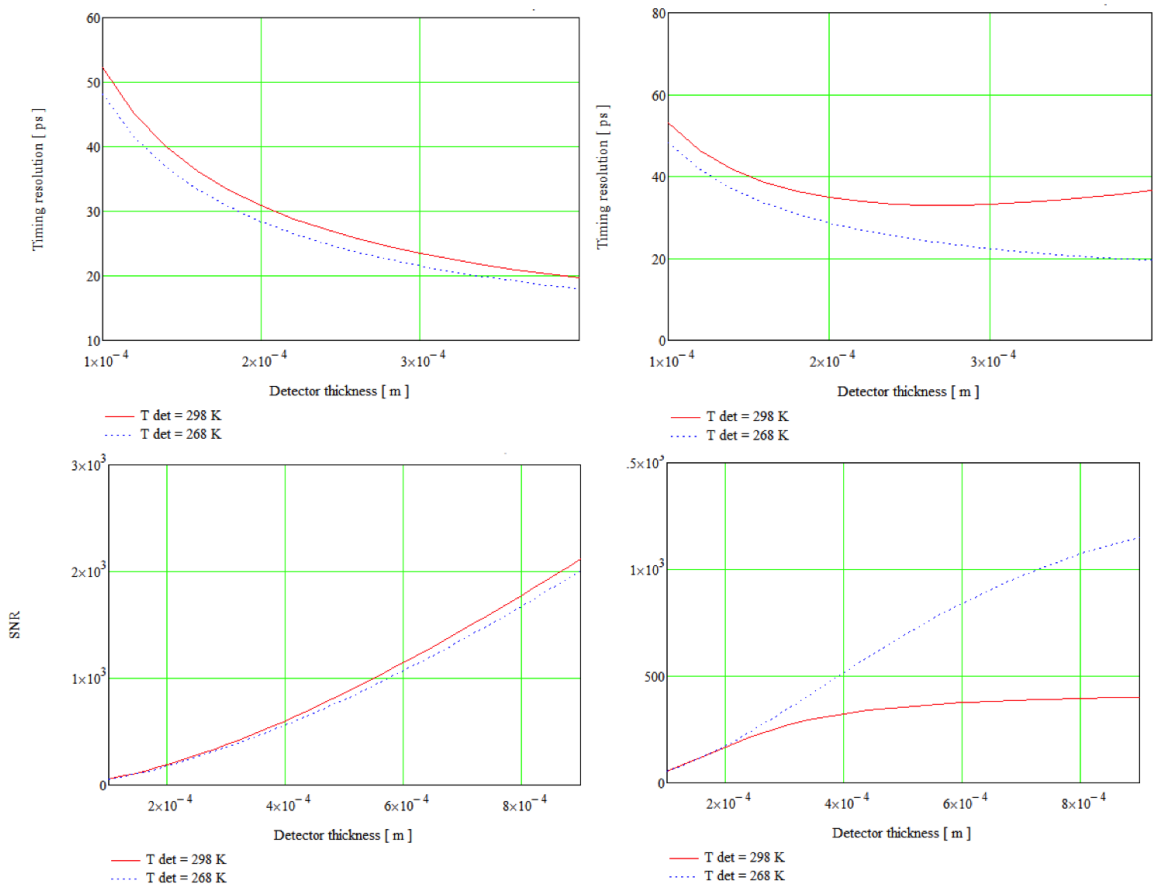

Figure 32 Simulation of time resolution and signal to noise ratio over sensor thickness.

2.3.4 Readout ASICs

The design of the pixel read-out chip poses unprecedented challenges in the combination of the required pixel cell size of 300 x 300 µm^2, the particle arrival time measurement resolution of better than 200 ps and the material budget of less than 0.5 % X_0 of radiation length per station. The low material budget translates to a constraint on the power consumption of less than 2 W/cm^2 as the cooling system needs to be kept to a minimum. A trade-off between signal amplitude, signal speed and acceptable material in the active area defines the thickness of the sensor to 200 µm.

Figure 32 gives the theoretically achievable time resolution and the signal to noise ratio over the sensor of the expected detector signal and the front-end electronics thickness for the most probable input signal of 2.4 fC for two leakage current values representing the sensor state before and after

irradiation. The study [28] takes measured noise values from a 0.25 µm ASIC technology into account, but similar results were obtained with the selected 0.13 µm ASIC technology. The shaping time is assumed to be identical to the charge collection time of the respective detector thickness. The power consumption in the first input transistor is assumed to be 325 µW.

The time resolution can be expressed as

$$\sigma_t \propto \frac{\mathrm{ENC} \cdot \tau_p}{\mathrm{Q_{in}}}$$

where ENC is the Electronic Noise Charge, τ_p represents the speed of the preamplifier and Q_{in} is the input signal charge.

For low leakage currents the white noise dominates, which decreases when the shaping time is increased with thicker sensors (Figure top left). As the signal and speed change approximately in the same way, time resolution is determined mainly by the noise. Thus the time resolution decreases with sensor thickness. For high leakage currents after irradiation and at room temperature increasing the shaping time does not decrease the noise anymore sufficiently and the parallel noise, which increases with shaping time becomes dominant (Figure top right). But cooling the detector reduces the parallel noise and brings us back to a situation similar to the pre-irradiation one. The two plots in the bottom of the figure show the signal to noise ratio (ratio of input signal divided by the rms noise values) before irradiation (left) and after (right), when increasing detector thickness the input signal is increased and thus the signal to noise ratio. When the parallel noise due to the increase of the leakage current after irradiation becomes dominant (Figure 32 bottom right plot), the signal to noise ratio still increases for thicker detectors, but at a slower pace. In conclusion a 200 µm thick detector is a good compromise between achievable time resolution and low material budget. The study also shows that cooling irradiated detector improves the time resolution.

Figure 33 shows the expected charge distribution (Landau) with the most probable value around 2.4 fC. The available signal charge (5000 to 60000 electrons or 0.8 to 10 fC) and the requirements on time resolution and jitter need analogue front-end time-walk compensation.

The integration of the analogue front-end together with time-to-digital converters and the high speed on chip data acquisition and read-out circuits need to be done with the utmost care in order not to disturb the analogue performance. Two different approaches have been considered to tackle this problem. One is to integrate the time-to-digital converter (TDC) together with the analogue front-end directly in the pixel cell. The TDC will be active only during a hit in that particular cell. A reference clock needs to be distributed over the entire pixel matrix posing a possible noise source to the analogue front-end electronics. The other approach is to keep the pixel cell as simple as possible, containing only the analogue front-end, and send the discriminator signal to the end-of-column area of the pixel matrix inside the ASIC where the TDC evaluates the arrival time. In this configuration the pixel cell stays simple, no clock needs to be distributed, but the discriminator signal needs to traverse the pixel matrix to be evaluated in the end-of-column.

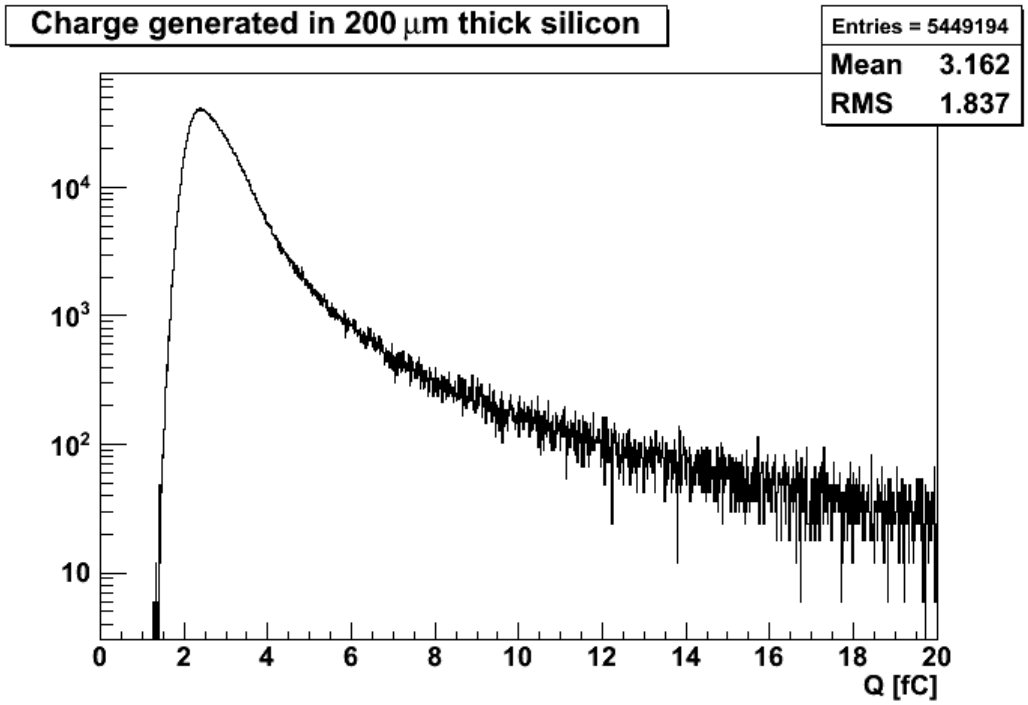

Figure 33 Landau distribution of input charge with most probable value around 2.4 fC.

The specified time resolution requires efficient time walk compensation. Two time walk compensation architectures have been investigated. a) Constant Fraction Discriminator (CFD), which directly gives a compensated discriminator signal, but uses sensitive analogue circuitry and b) Time-over-Threshold (ToT) discriminator, which delivers two trigger events to be time tagged but uses simpler and less sensitive circuitry. Compensation by sampling the signal shape with a Flash ADC has been excluded for size and power consumption reasons.

As simulations cannot show the performance and drawbacks of the architecture and time-walk compensation schemes in real life, it was decided to produce test structures for both architectures and compensation schemes.

2.3.4.1 ASIC Specification

The GTK pixel electronics is very demanding and requires performance close to the limits of the 130 nm CMOS technology used. The circuit is optimized in view of trade-offs between acceptable power dissipation of 2 W/cm^2, timing precision of less than 200 ps, centre pixel particle rate of 124 kHz/pixel and an output data rate per readout chip of up to 6 Gb/s including data transmission overhead. Table 17 presents a summary of the specification.

The beam particle rate is 750 MHz. It corresponds to the following rates for the centre chip. The hit rate in the centre pixel is 114 kHz. The hit rate of the centre column is 2.7 MHz or 0.66 MHz/mm^2. The average hit rate per pixel on the centre chip is 58 kHz with a chip hit rate of 105 MHz. Table 17 summarizes the rates.

Assuming a 32 bit word per hit, and considering the average chip hit rate, a data rate of more than 4 Gb/s needs to processed and sent off each read-out chip. On-chip triggering is ruled out as the earliest trigger signal arrives 1 ms or later after the event, accumulating 4 Mbit of data during the trigger latency. The area of silicon needed to store the data and process the trigger on chip would be too large. Consequently the data are shipped off the chip in a triggerless, data driven architecture. In order to reduce digital noise and connectivity a low number (2-4) of Gbit serial links are used for each chip.

Table 17 GTK ASIC specification

General	Max. material budget per station	0.5 % X_0
	Overall efficiency per station	> 99 %
	Pixel size	300 µm x 300 µm
Architecture/geometry	Chip active area	13.5 mm x 12 mm
	Chip size	19.5 mm x 12 mm
	Number of pixels per chip	40 columns x 45 rows = 1800
	Size of pixels	300 µm x 300 µm
	Thickness of read-out chip	100 µm
Pixel analogue front end	Input dynamic range	0.6 -10 fC; 5000 – 60000 e^-
	Electronic noise in pixel input with sensor	200 e- rms
	Time resolution of the discriminator	100 ps rms
	Discriminator time walk after correction	80 ps rms
	Overload detection threshold	10 fC; 60000 electrons
	Maximum dark current per pixel tolerance	1 µA
	Power dissipation per pixel	1 mW max
	TDC time binning	100 ps
Data rate	Beam rate	750 MHz
	Average centre chip hit rate	105 MHz
	Average particle rate per pixel in station	42 kHz
	Average particle rate per pixel in centre chip	58 kHz
	Peak rate per pixel /beam centre	114 kHz
	Average centre column hit rate	2.7 MHz or 66 MHz/cm^2
	Centre chip data rate (40 bit per hit word, no contingency, no encoding)	4.3 Gb/s
	Center chip data rate (incl. 8b10b encoding, 30% contingency)	6.4 Gb/s
	Serializer data rate (4 used per chip)	1.6 Gb/s
	Read-out efficiency	99% (98% in beam centre)
	Operating temperature in vacuum	-20 ^oC to +5 ^oC
Radiation	Centre pixel total dose in 1 year	2.10^{14} particles/cm^2, $6*10^4$ Gy
Trigger	Number of levels	1 (L0)
	Latency	~ 1 ms initially, future upgrade several ms.
Off chip trigger	Trigger window	75 ns

In order to identify high charge deposit in the sensor and the ASIC provides information if a charge of higher than 10 fC has been deposited.

2.3.4.2 ASIC Architecture

The active pixel matrix of 13.5 mm x 12 mm in size arranged in 45 rows x 40 columns of pixels with dimensions of 300 µm x 300 µm, results in 1800 pixels per chip. The chip area outside the active beam area used for data processing and control is estimated to be in the order of 4.5 to 6 mm deep resulting in a total chip size of 18 to 19.5 x 12 mm, see Figure 27. The communication to and from the pixel read-out chip is done via high speed differential signals. This reduces the number of connections and thus increases module reliability. The number of I/O pads is limited, as the connections to the module are done on one side of the chip only, see Figure 27. A high number of single ended data lines would increase the digital system noise. The clock signal and the serial control signals are connected on two high speed serial input lines. Configuration data and settings are transmitted on one slower serial input port and one parallel input port. The pixel hit data output ports consist of two to four multi gigabit serial ports for the data and one slower serial port for status information. For test purposes a number of single ended outputs are foreseen. Figure 34 illustrates the I/Os of the chip.

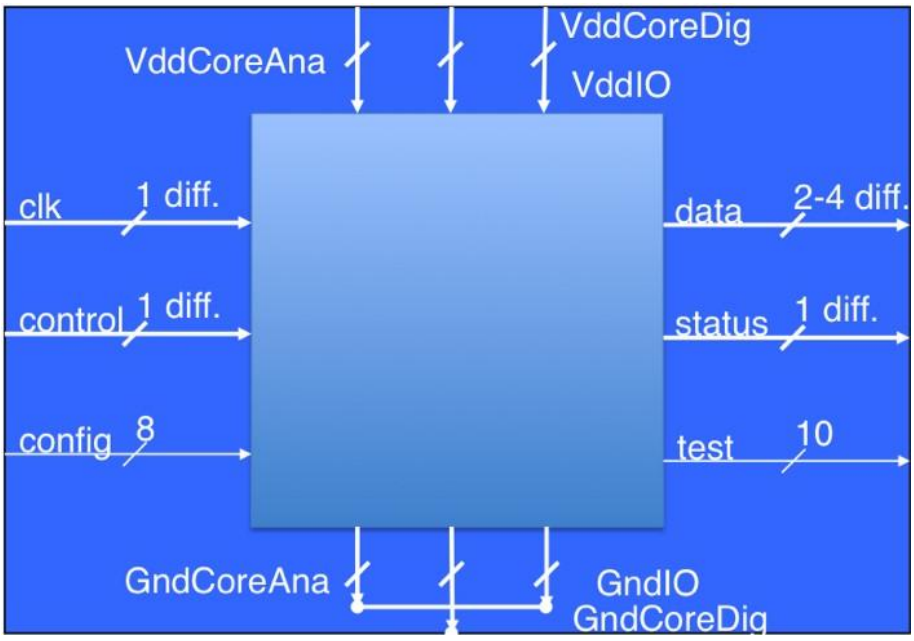

Figure 34 I/O connections of the GTK pixel read-out chip.

The design and layout of the digital processing units must be done so that the digital noise has little effect on the analogue front end. The implementation of the pixel read-out chip GTK is done in 0.13 µm CMOS technology. This technology is fully available and design tools have been provided. First studies indicate that the technology might be sufficiently radiation hard by itself for the expected average total dose of 10^5 Gy/year. However, a large variation of dose is expected between chips of a GTK station and between chip edge and centre. This dose gradient might introduce variation of

signal response between pixels. To circumvent this effect, the trimming in each pixel used for process mismatch can also partially compensate for the dose gradient.

Critical registers and nodes are protected against single event upset effects. The 130 nm has a relatively high substrate resistivity when compared to 250 nm technology. This makes it easier to isolate digital and analogue blocks by means of guard-rings. However, the high sensitivity and large bandwidth of the analogue front-end require the use of architectures to reduce the digital noise such as differential analogue and digital circuits, which tend to reduce both propagation of ground and substrate noise and improve immunity of the analogue front-end.

2.3.4.3 Design Options

Both architectures described above and time-walk compensation techniques have been developed and integrated into two demonstrator ASICs. In one the CFD and the TDC are implemented inside the pixel cell – TDC per pixel (P-TDC). In the other chip each pixel contains a time-over-threshold discriminator and a transmission line driver to send the discriminator signal to the end-of-column (EOC). Arrival time of leading and trailing edge of the hit signal are time tagged in delay locked loop (DLL) based TDCs in the EOC.

Both demonstrator chips contain at least one full, folded column containing the final analogue front-end, the full TDC functionality but reduced read-out capability compared to the final pixel chip design. The results of the qualifications are described in sections 2.3.4.5.4 and 0. The characterization of the demonstrator front-end chips in view of building the final full pixel matrix and selecting the final building blocks and architecture will address the following questions in general:

a) time resolution; b) charge collection efficiency (sensor/ASIC assembly); c) read-out efficiency.

The following paragraphs outline the measurements in detail. If applicable all tests are done with and without sensor. For the silicon assembly (sensor/ASIC) the parameters also are measured with respect to varying high voltage and position of charge injection on the pixel. Furthermore the absolute charge collection gain is determined.

- Arrival time resolution
  (for 2.5 ns charge injection duration and read-out clock activated)
  - Full chain time walk or compensation error over input charge and folded over Landau charge distribution.
  - Full chain jitter over input charge and folded over Landau charge distribution.
  - Influence of read-out clock on full chain resolution/jitter - digital clock/noise separation
  - Behaviour of long transmission lines (clock or signal)
  - Influence of signal charge injection duration on full chain resolution/jitter
- Analogue frontend behavior
  (for 2.5 ns charge injection duration and read-out clock activated)
  - analogue fronted time walk or compensation error over input charge and folded over Landau charge distribution
  - analogue frontend jitter over input charge and folded over Landau charge distribution
  - Front-end noise
  - Gain distribution

  - Offset distribution before trimming
- TDC behavior
  - TDC resolution
  - TDC INL/DNL
  - TDC jitter
- Information about deposited charge in the sensor
- System robustness
- Simulated efficiency for full matrix chip
- Simulated power of individual blocks and measured power as far as possible
- Size of individual blocks
- Temperature stability
- Estimated and possibly measured SEU rate of individual blocks

The P-TDC option uses a constant fraction discriminator in the pixel already compensating for the time walk. In the EOC column option a time-over-threshold discriminator is used which provides signal amplitude information by measuring the pulse width allowing the compensation of the time walk.

In the P-TDC option, a dual slope Wilkinson TDC in each pixel converts the discriminator edge into a time stamp synchronized with a master clock distributed to the entire pixel array. In the pixel a buffer and derandomizer is used to increase the read-out efficiency before the data are sent to the column bus and subsequently off the chip. In the EOC pixel circuit only a driver follows the discriminator stage to send the hit signal onto the column bus. There the hits are time tagged using a delay locked loop based TDC before being sent off the chip.

The features of the two options are summarized below. In the P-TDC, noise might be injected in the analogue circuits in the pixel array due to the digital activity in the pixel cell. The clock needs to be distributed over the entire matrix with a minimum of jitter to avoid non recoverable errors on the time tagging. In the EOC option each individual discriminated pixel signal needs to be sent to the EOC, however, the transmission lines are only active once a hit has occurred. The ToT compensation scheme imposes the measurement of two time tags increasing the output data volume and rate. The ToT discriminator is better adapted to measure the signal amplitude to detect high energy transfer in the sensor. The CFD discriminator would need to be modified to include large signal detection.

Table 18 summarizes the 2 architecture blocks. Figure 35 and Figure 36 show the basic architectures of the P-TDC and the EOC respectively.

Table 18 Summary of the design option for the p-TDC and EOC architectures.s

Chip Architecture	Time walk compensation	Pixel Circuit	Column Bus	EOC	Backend
P- TDC	CDF discriminator	TDC, Buffer, derandomizer	Synchronous bus	Buffer Derandomizer, Coarse counter drivers, read-out control	Serialiser FIFO logic
EOC	ToT discriminator	Hit driver	Asynchronous bus	EOC TDC, buffer derandomiser	Serialiser FIFO logic

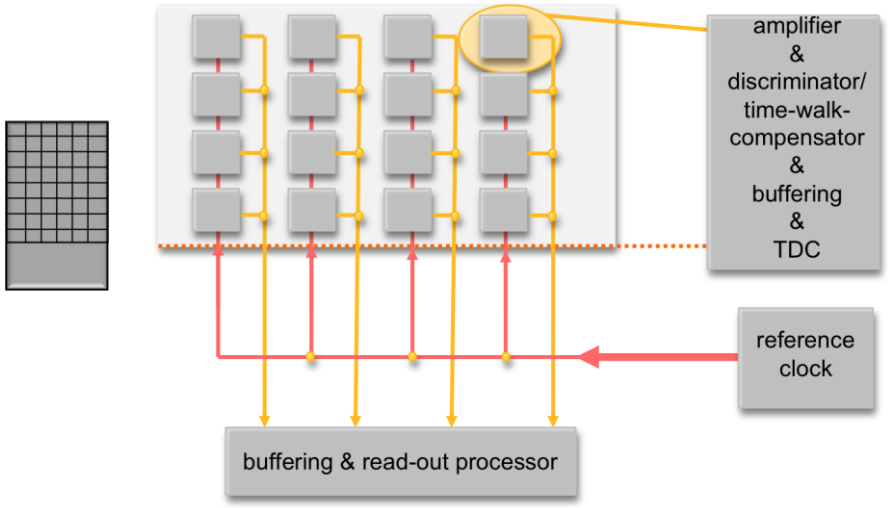

Figure 35 Block diagram Pixel in TDC option (pTDC).

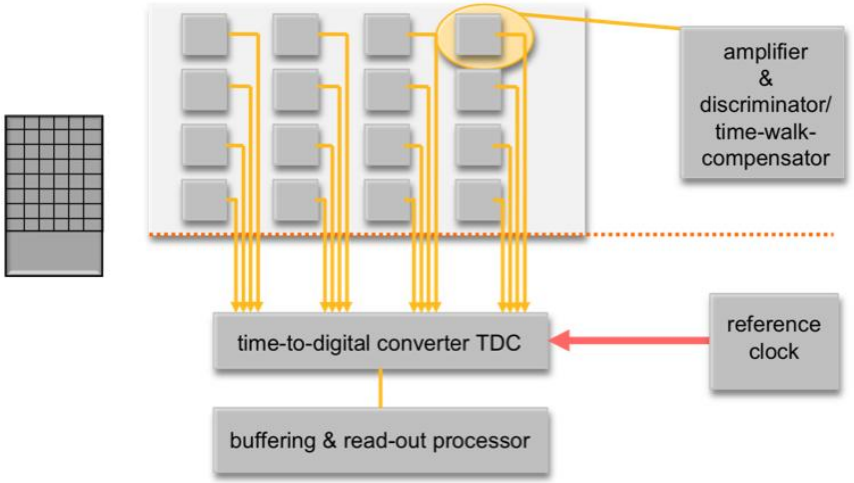

Figure 36 Block diagram End-of-column option.

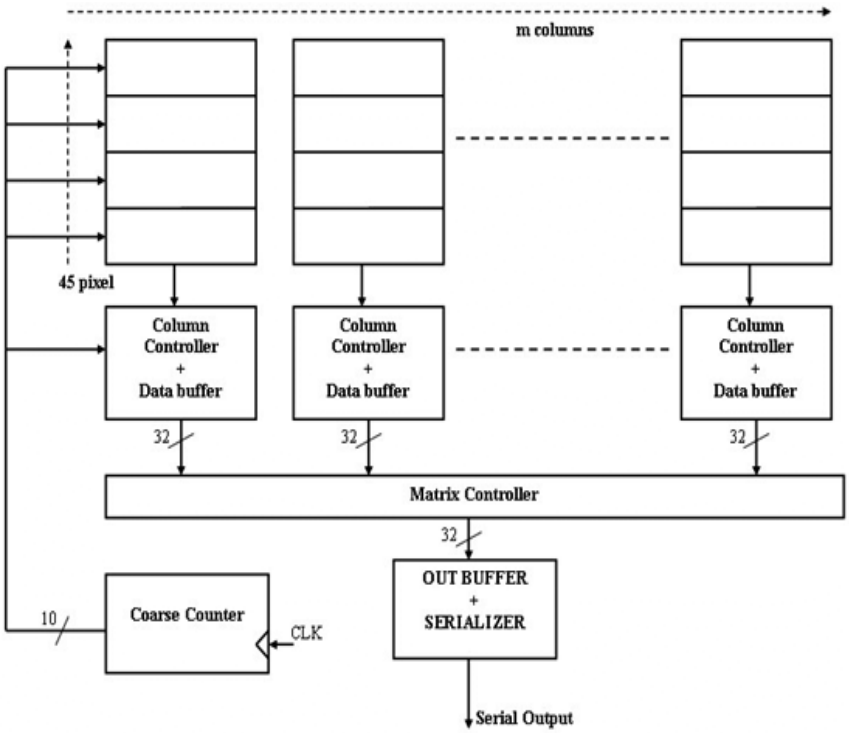

Figure 37 Block diagram of the P-TDC option.

2.3.4.4 TDC per Pixel (P-TDC) Architecture Overview

In the P-TDC most of the signal processing, including time to digital conversion and multi-event buffering, is performed inside each pixel. Figure 37 shows a simplified block diagram of the P-TDC architecture, while a scheme of the pixel cell is shown in Figure 38. Each column, 45 pixels, is read-out by a dedicated controller. The data are then merged and prepared for the output serializer. In this approach only digital signals are exchanged between the pixels and the periphery. The most critical signal is the clock, which must be distributed with high precision to the whole matrix. A dedicated transmission line for each pair of column is used for this purpose.

The time walk error is corrected at the analogue level with a Constant Fraction Discriminator (CFD). The master clock of the chip has a design frequency of 160 MHz. The clock pulses are counted and the resulting word is Gray-encoded and sent to all pixels. When a hit is flagged by the CFD, the value present on time-stamp bus is latched into local registers, providing a first coarse estimation of the event time. In each pixel, a TDC based on a Time to Amplitude Converter (TAC) generates a voltage proportional to the interval between the hit and the next clock edge by charging a capacitor with a constant current source. The resulting voltage is then digitized, providing the required fine time measurement. Figure 38 shows a block diagram of the P-TDC pixel cell.

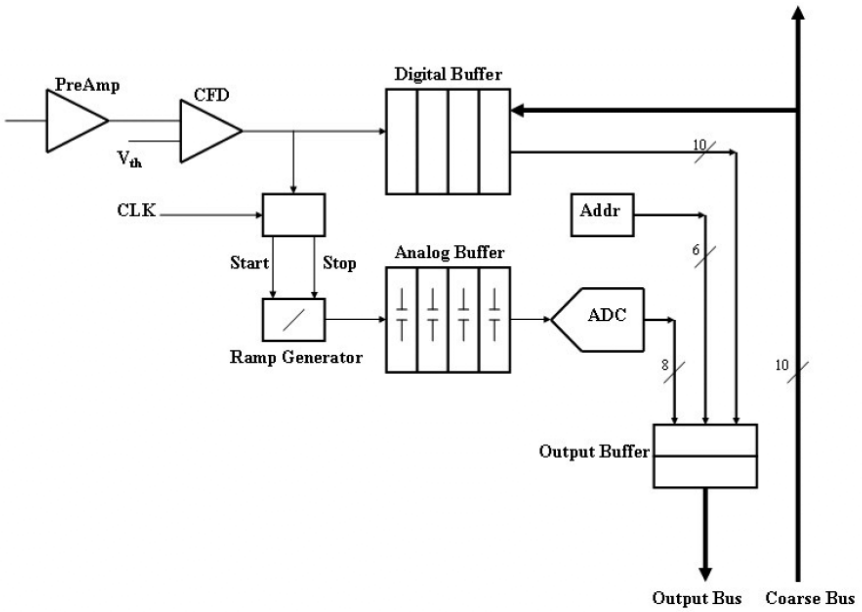

Figure 38 Block diagram of the P-TDC pixel cell.

2.3.4.4.1 P-TDC Preamplifier and Constant Fraction Discriminator

The analogue section can be divided into three main blocks: the front-end amplifier, the CFD and the TDC. The front-end has been designed to be compatible with sensors of either polarity. In order to minimize the risk of digital noise pick-up a fully differential topology is used for both the front-end amplifier and the CFD. The preamplifier employs a telescopic cascode with a class AB output stage. The latter is needed to drive the CFD with good power efficiency. A low-frequency feed-back, implemented with an RC low-pass filter followed by a slow op-amp, allows the compensation of the sensor leakage current. Dark currents up to +/- 200 nA can be accommodated while keeping the resulting output offset below +/- 10 mV. The input stage of the Constant Fraction Discriminator is formed by a passive RC filter, in which both the delayed copy of the signal and the signal fraction are generated. The filter is followed by a high gain differential amplifier, made by cascading two differential cells with resistive load. This amplifier drives a third stage, composed by two differential pairs sharing the same load. One pair is driven by the signal while the other one is used to apply the threshold. A hysteresis is introduced at this level to avoid excessive noise triggering. Finally, a differential to single ended converter generates a logic pulse of 1.2 Volt, suitable to drive the digital part.

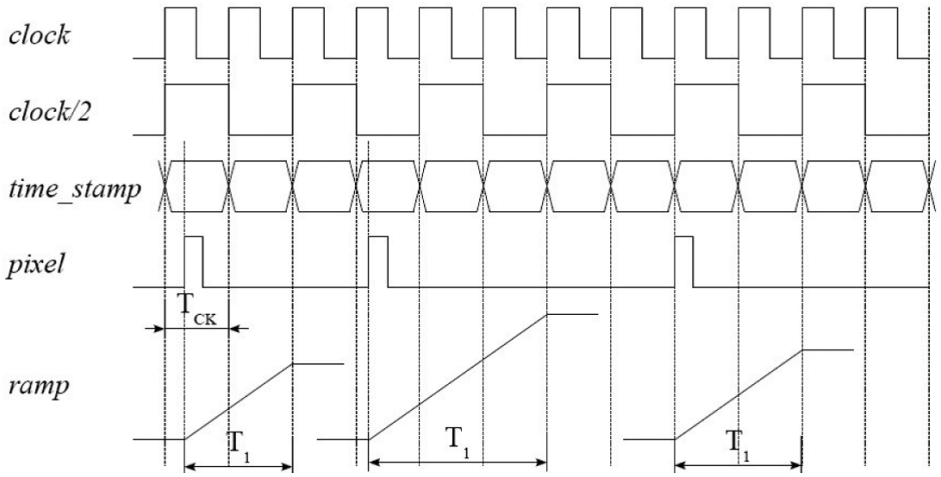

Figure 39 Generation of the analogue ramp for the TDC.

2.3.4.4.2 P-TDC Time-to-Digital-Converter

In each pixel, a TDC based on a Time to Amplitude Converter (TAC) generates a voltage proportional to the interval between the hit and one of the clock edges by charging a capacitor with a constant current source. The resulting voltage is then digitized, providing the required fine time measurement. The relevant waveforms are sketched in Figure 39. One must take into account potential misalignments between the coarse and the fine time logic. For this reason, the latter works at half the clock rate. With such an arrangement, the LSB of the coarse counter overlaps the MSB of the TDC, thus making digital error correction possible. Furthermore, the analogue ramp is stopped on the first rising edge following a falling edge. This introduces a time offset, which allows a proper settling of the circuits generating the ramp. The ADC employs a Wilkinson topology: the capacitor of the TAC is discharged by a small current and the clock pulses necessary to restore the baseline are counted. The time needed for the fine time measurement may reach up to 1 μs in the worst case. Since the maximum expected rate per pixel is 140 kHz, a multi-buffering scheme is mandatory and it has been implemented with two FIFOs, one analogue and one digital. The digital FIFO allows the storage of up to four values of the coarse counter, while the analogue one, formed by four capacitors, accommodates the corresponding fine time measurements.

The P-TDC approach allows an optimal use of the relatively large pixel area and a simplified design of the End of Column logic (EoC). This can be sized according to the average event rate since data are already derandomized in the pixel cell.

Figure 40 shows a simplified schematic of the TDC. It is formed by two PMOS current sources, four TAC cells (which implement the analogue FIFO), a capacitor and a synchronous comparator.

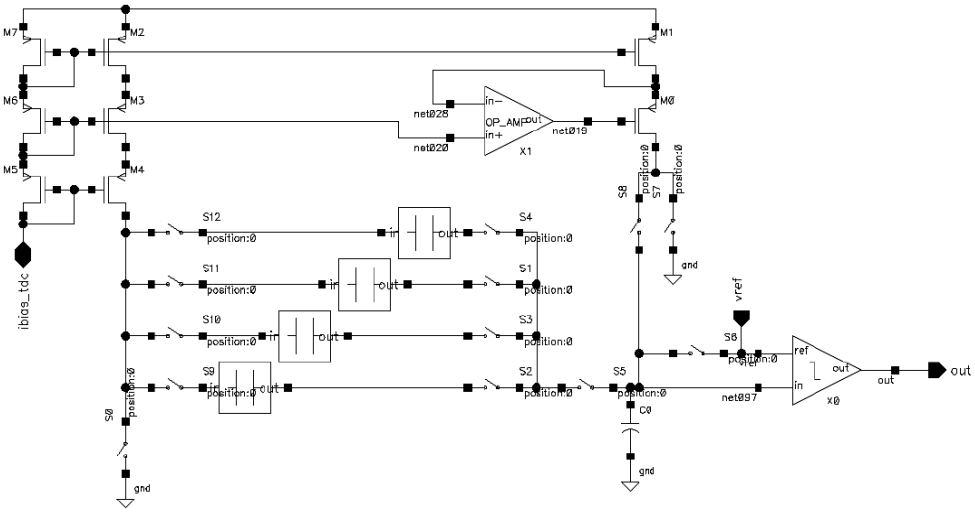

Figure 40 Simplified schematic of the TDC.

The TDC works as follows. When the CFD fires, the current source formed by M2-M4 is connected to one of the TAC cells, as shown in Figure 40. The capacitor is put in the feed-back path of a single-ended cascode amplifier. The purpose of this circuit is two-fold. First, it keeps the drain of M4 at virtual ground, thus avoiding a modulation of the current through the finite output conductance of the transistors that would introduce nonlinearity. Second, it provides an inversion of the signal, which is stored on capacitor C9. In other words, while the TAC capacitor is *charged*, C9 is *discharged*. Therefore, the baseline value at the input of the comparator can be restored with a current having the same polarity as the one delivered by M2-M4. Both current sources are hence implemented with the same type of transistors, which is very important because the matching between them is critical. Additionally, the value of C0 is four times the one of the TAC capacitors. In this way, to have a 7 bit resolution, the current provided by M1-M0 only needs to be 32 times the one given by M2-M4, resulting in a much more compact circuit Figure 41 shows the schematic of the TAC.

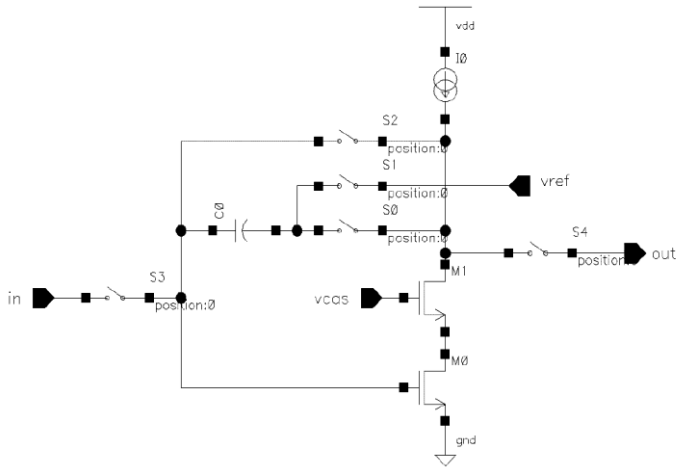

Figure 41 Schematic of the Time to Amplitude Converter. Four of these circuits form the analogue FIFO.

2.3.4.4.3 P-TDC Digital Section

The digital section of the pixel is composed by the digital FIFO that stores the coarse counter values, the logic controlling the TAC, the TDC and the interface to the End of Column readout. One of the possible drawbacks of the P-TDC option is that the many digital gates in the pixel are directly on the beam trajectory. The protection against Single Event Upset is therefore a primary concern and is addressed through the systematic use of Hamming encoding, as detailed in one of the next sections.

2.3.4.4.4 P-TDC - End of Column Logic (EoC)

The schematic of the column controller is shown in Figure 42. A state machine sends a read enable signal to the pixels, using a token ring mechanism. The read enable is then passed from one pixel to the next. When the read enable is high each pixel writes out the data content and its empty FIFO flag. A busy signal alerts the End of Column FSM if there are data available in at least one of the 45 pixels. In order to avoid ambiguities in the time of reconstruction the events belonging to the same cycle of the time stamp counter value are grouped together. Adding an extra bit to the time stamp does this rearrangement. The additional bit is used to tag events of type 0 and events of type 1 in two different FIFOs. A merger circuit puts the data frame in the right order, adding a header and a trailer. The header contains the information of the frame counter value. To separate the frames, the circuit uses the information provided by the pixel logic. Each pixel sets an “old-data flag” if the frame changes when it still contains data.

This flag is propagated through the whole column with a fast or. The depth of the FIFO has been chosen to be 32. VHDL simulations have shown that at the maximum nominal rate of 3.3 MHz per column, the maximum FIFO occupancy is 20. The layout area of the EoC is 1500 μm x 300 μm.

2.3.4.4.5 Single Event Upset Protection

Given the high particle flux, protection against Single Event Upset (SEU) is a serious issue for the GTK front-end electronics. In the P-TDC option the following solutions have been adopted:

- Pixel cell: Hamming encoded state machines and registers with auto-correction (single error correction and double error detection).
- End of Column: Hamming encoded state machines with auto-correction (single error correction and double error detection).

Concerning the EoC FIFOs the following considerations have been made. The particle flux in the EoC area is 1.3 10^6 /cm^2 per second assuming that in the EoC area the flux is 1 % of the peak flux. The SEU cross section in silicon is 10^{-8} cm^2/bit [29] and the collision probability is: $P(x) = e^{-x/L} = 3.3\ 10^{-6}$, where x is the thickness of the active silicon containing transistors (1 μm) and L is the mean free path in silicon (30 cm). The flux of particles giving SEU is therefore 1.3 10^6 x 3.3 10^{-6} = 4.3 particles per square centimetre per second. One FIFO of 32 words with 32 bits per word has ~1200 flip-flops including control logic and hamming code protection. The number of SEU per second is then 1200 x 10^{-8} x 4.4=5.14 10^{-5}. The data stay in the FIFOs for at most 12 ms, so the upset probability is 6 10^{-10} for single error and 10^{-19} for double errors. For these reasons the FIFOs have a Hamming encoder at the input and a Hamming decoder with single error correction and double errors detection at the output registers only and are protected by triple mode redundancy (TMR).

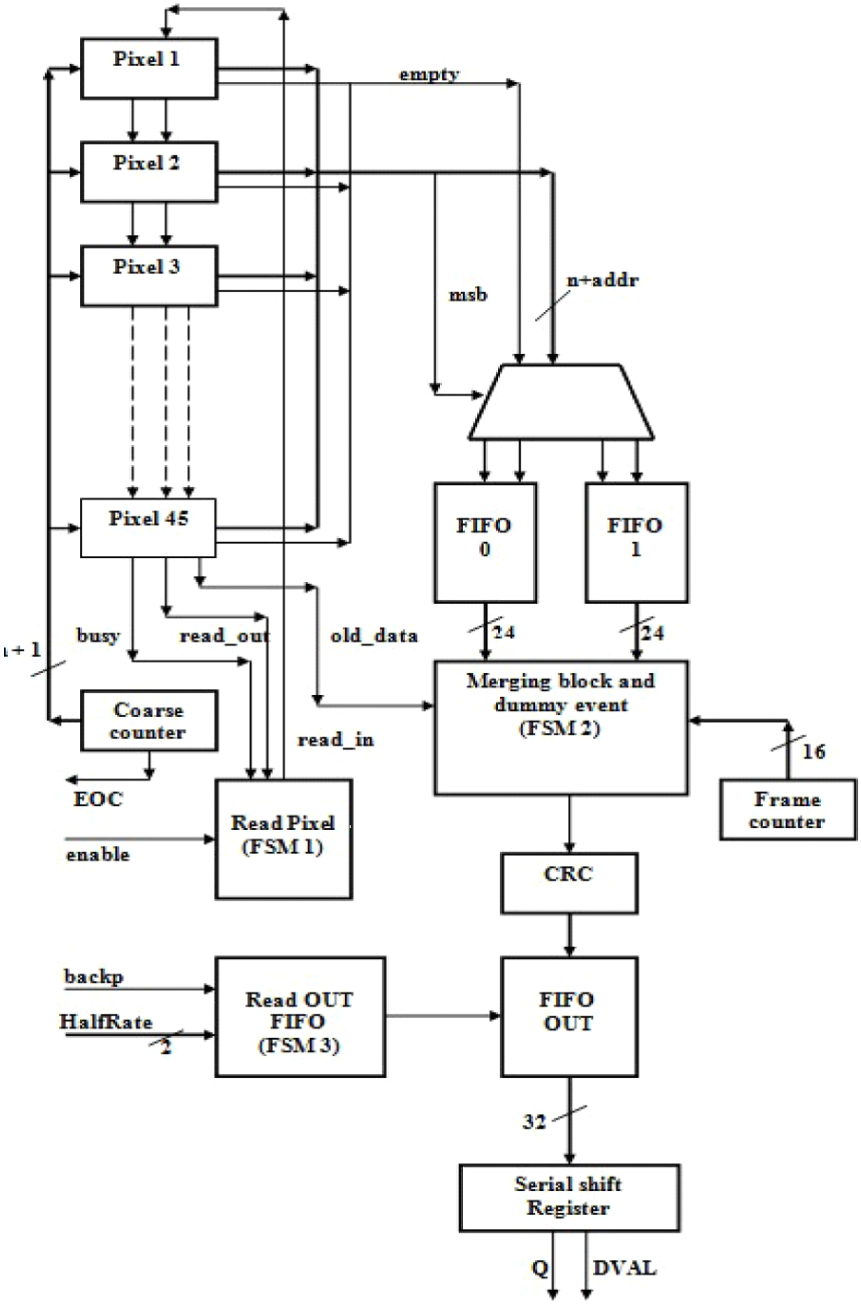

Figure 42 Block scheme of the End of Column logic.

2.3.4.4.6 Demonstrator for the P-TDC Option

Prototype implementation.

The key goal of the design was to have a chip as complete as possible, so that all the critical issues associated to this option could be evaluated and a future extension to a larger, final ASIC could be done with a minimum of risk. Therefore, all the blocks contained in the pixel cell (front-end

amplifier, constant fraction discriminator, DAC for threshold setting and full in-pixel logic) are implemented in their final form. The only circuit, which is missing in the pixel cell is a “watch-dog” to flag hits that deliver a very high amount of charge in the sensor. However, due to the moderate resolution required (2-3 bits) the implementation of such a circuit does not entail critical issues. The design of the EoC logic is final and includes the SEU protection strategies discussed in the previous section.

The digital circuitry that merges the data coming from different columns and generates the serial output stream has not been included in this version. The prototype ASIC contains two long columns with the final number of pixels (45) and one short column with 15 pixels. Two spare pixels are provided for debugging purposes. In order to have an acceptable form factor, the long columns have been folded in three segments. One of the critical aspects of the P-TDC option is the quality of the clock distributed to the pixels. To assess this point the clock and the digital signals are propagated to the long columns with folded lines, which have a length slightly higher than the one expected for the final layout (14 mm). Figure 43 shows a picture of the prototype ASIC.

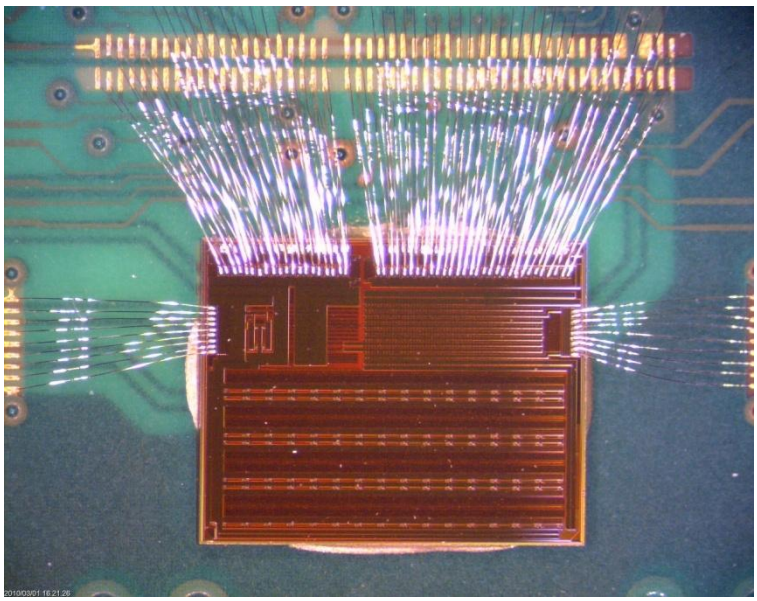

Figure 43 Picture of the p-TDC ASIC prototype wire bonded to the test printed circuit board

Test set-up

The bare die is wire bonded to a custom designed PCB. A combined data-pattern generator/logic state analyzer unit is used as the main testing instrument to provide the test vectors and read-back the chip response. A computer running routines in lab-view controls the test equipment and is employed for data analysis.

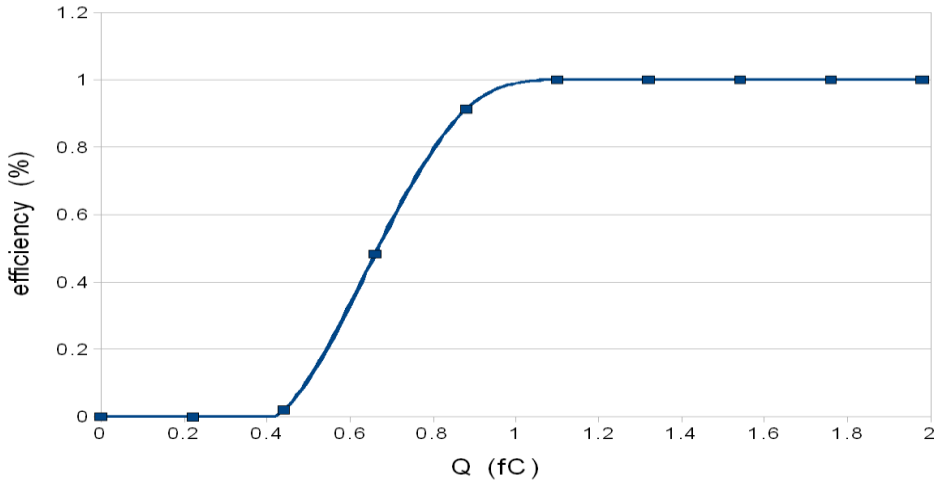

Figure 44 Efficiency curve for one pixel.

Test results

For testing purposes, a calibration signal can be injected through a 22 fF nominal capacitor. One leg of the capacitor is connected to the preamplifier input. An external trigger signal switches the other leg between the preamplifier virtual ground and a voltage level controlled by an external DAC. The difference between these voltages multiplied by the capacitor value gives the amount of injected charge. The switching leg of the capacitor is also connected to a low-pass filter. The purpose of this circuit is to smooth the input signal, providing a shape as similar as possible to the one that will be given by a real detector. Figure 44shows a typical S-curve obtained for one pixel.

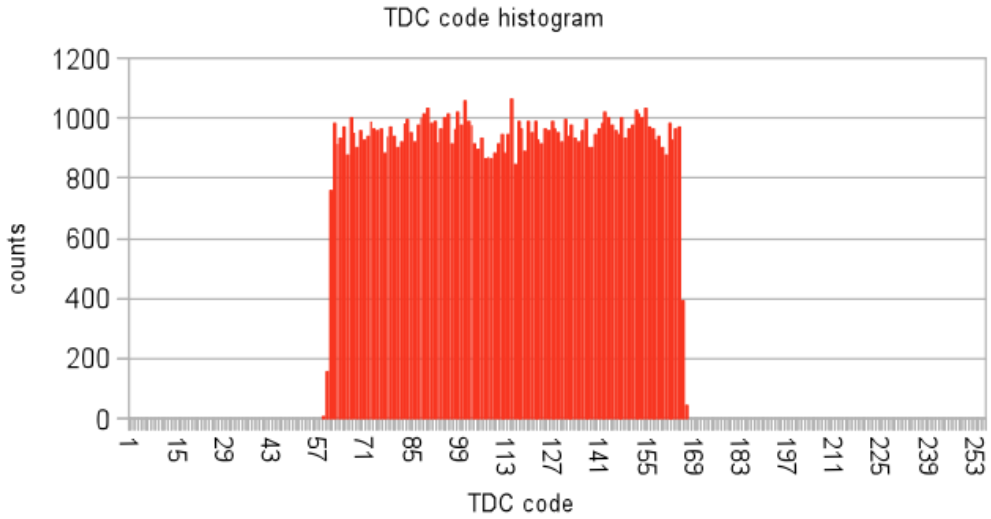

Figure 45 Example of raw code distribution used in the calculation of TDC DNL and INL from a data sample of 100.000 events.

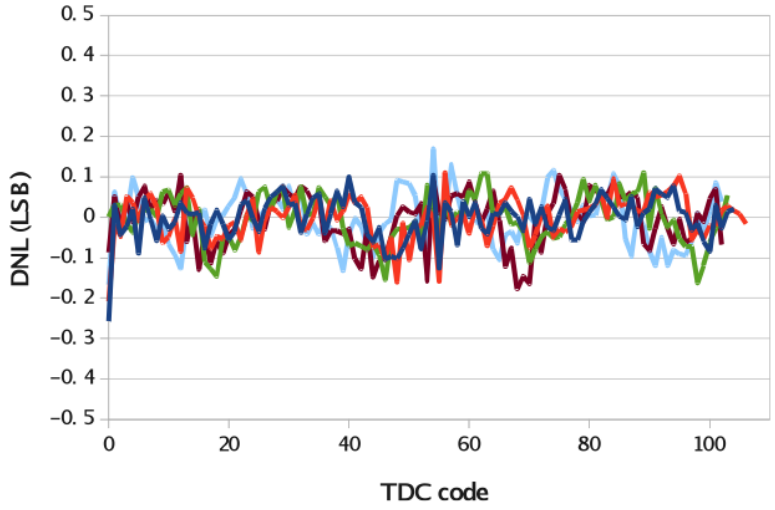

Figure 46 Differential non-linearity of five different TDCs from the matrix.events.

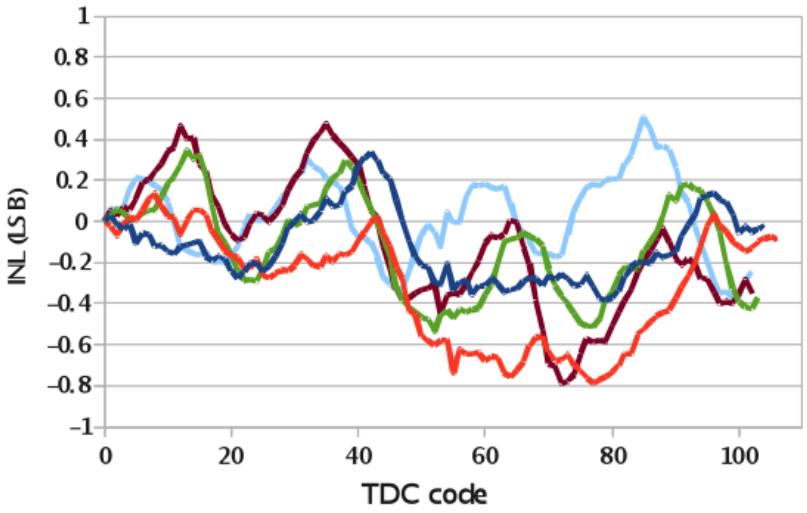

Figure 47 Integral non-linearity of five different TDCs from the matrix.

The TDC can be tested independently from the front-end part. In fact, it is possible to trigger the TDC either by an external digital pulse or by an internally generated one. The latter can have a length either of one or three clock cycles. In this way it is possible to perform code density tests (with the external pulse) and to calibrate the full-scale range of the TDC (with the internal one). Due to an undersized buffer in the EoC logic the maximum working frequency is at present 128 MHz. Therefore, the expected time bin is 122 ps instead of the design value of 97 ps. The experimental average measured value is 140 ps, with the extra 18 ps coming from a slightly higher value of the discharged current due to a systematic mismatch in the current source. While both effects that lead to the increased bin sized are well understood and can be easily fixed in a final version, a bin value of 140 ps gives a quantization noise of 40 ps rms, so its impact on the overall time resolution would be still acceptable. Figure 50 shows a raw data plot used to calculate the DNL and the INL of the TDC. This plot is obtained by feeding to the TDC an external test pulse with a frequency uncorrelated with the one of the clock. Figure 51 and Figure 52 report the DNL and INL profile, respectively. For clarity, the results of five different TDCs randomly selected from the three different columns have been superimposed in these plots.

Figure 53 shows a measurement of the residual time walk of the CFD taken at different clock frequencies (blue squares=no clock, yellow triangles=160 MHz, red circles=320 MHz). Figure 54 shows the jitter measured in the range 1.5-10 fC. The measurement was done at the clock frequency of 128 MHz.

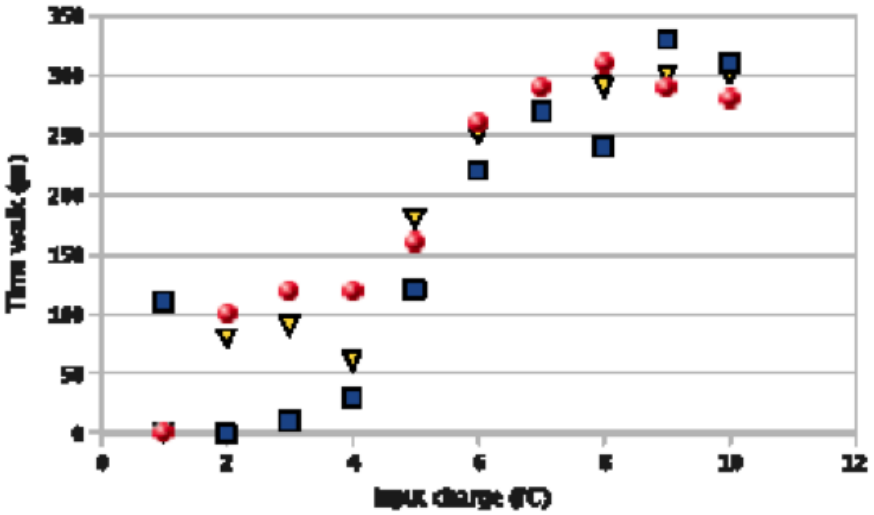

Figure 48 Residual CFD time walk at different clock frequencies.

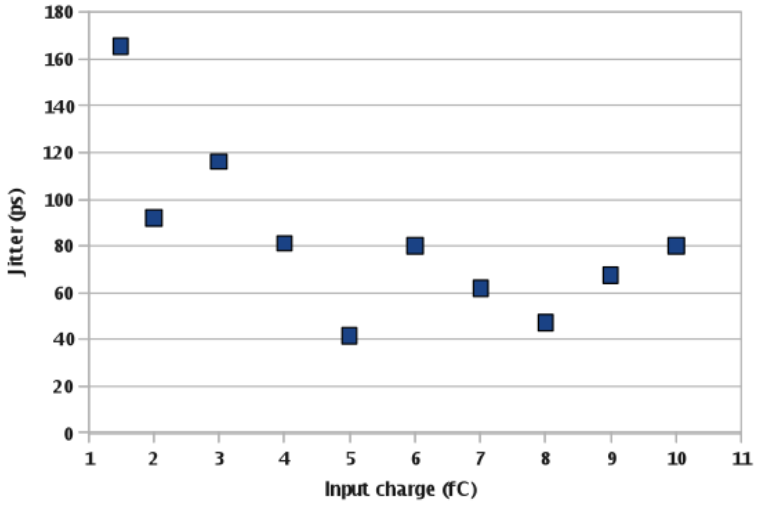

Figure 49 System jitter at 128 MHz clock.

The accuracy of the present test results is limited by the statistic, which can be achieved in a reasonable measuring time with the set-up. At the time of writing a new system based on a state-of-the art FPGA in under commissioning.

The most critical point of the P-TDC option is the risk of coupling digital noise to the amplifier inputs, since a lot of logic is resident in the pixel. This risk may be enhanced by the presence of the detector capacitance, so a full and realistic assessment of the performance of this option can be done only when the prototypes bump bonded to a sensor will be available.

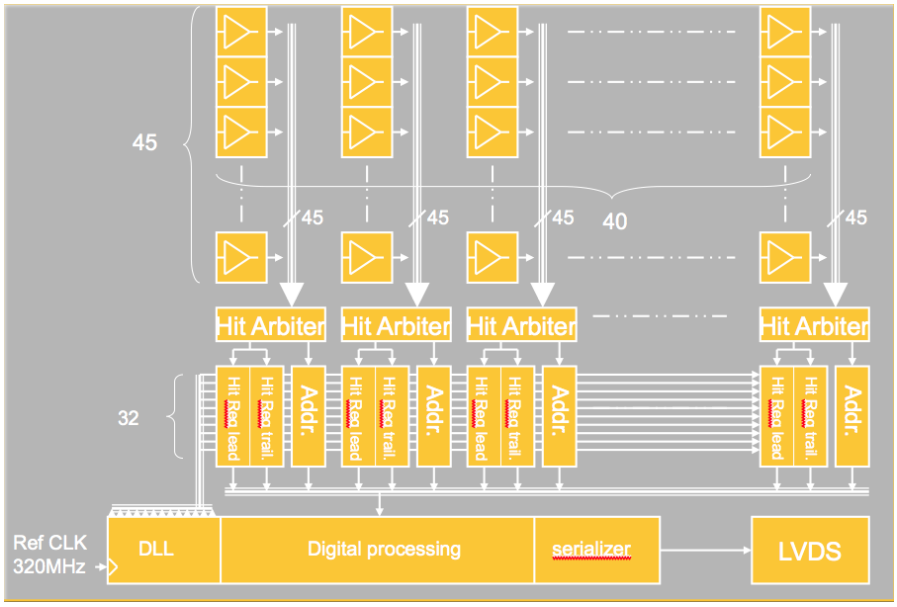

Figure 50 EOC architecture block diagram.

2.3.4.5 End-of-Column (EOC) Architecture Overview

The time walk compensation is based on time over threshold correction (ToT) measurements. Thus the rising and the falling edges of the discriminator signal are time tagged. The leading edge time stamp provides information of the hit arrival time and the trailing edge provides input charge amplitude information used to correct time walk. At the time of the leading edge, a synchronous coarse clock counter value is latched to increase the dynamic range of the TDC to 12.8 µs. The double time stamp with coarse counter and address information are stored in a buffer, before they are serialized and sent off the chip.

Figure 50 shows a block diagram of the final EOC chip architecture [30] [31]. The end-of-column (EOC) option employs simple pixel cells containing the amplifier, time-over-threshold (ToT) discriminator and a transmission line driver per pixel cell, which sends the discriminated signal to the end-of-column region (EOC) of the chip. There a delay locked loop (DLL) based TDC [32] [33] time tags the rising and falling edge of the discriminator signal with a time bin of 100 ps. In the chip, 5 not-neighbouring pixels in a column are multiplexed together and connected to one TDC using a combinatorial hitArbiter circuit. Simulation shows that the hit efficiency for the worst-case centre column stays above 99.5 %. Pixel hit addresses and rising and falling time tag words are sent to a buffer FIFO before the data are serialized and sent out of the chip. The time walk compensation using the two time tags is performed off-line allowing the evaluation of the input signal charge for each hit. In this architecture no clock signals are distributed over the clock matrix, reducing possible interference with the analogue electronics and data are only transmitted to the EOC circuits once a hit has been registered. The digital processing and transmission units are geographically separated from the analogue front-end electronics.

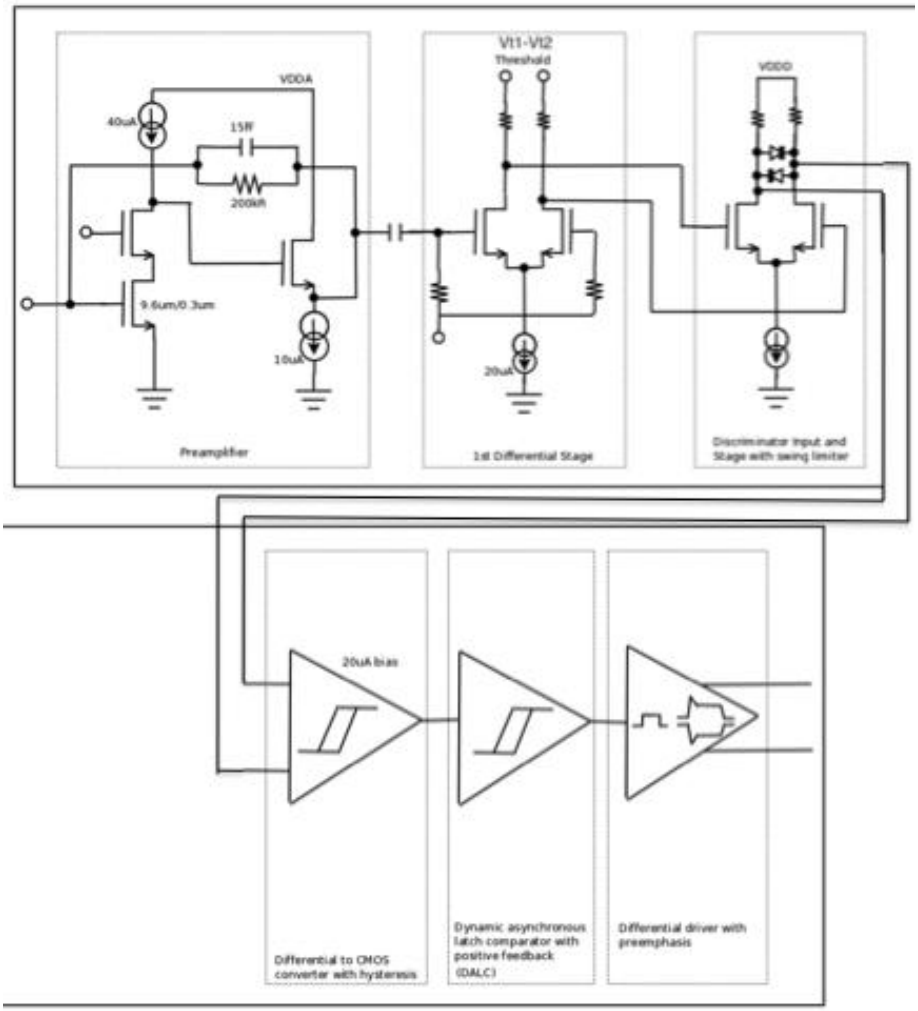

Figure 51 Block diagram of EOC pixel cell.

2.3.4.5.1 EOC Analogue Pixel Frontend

The pixel cell contains 7 functional blocks:

- a transimpedance preamplifier;
- a differential post amplifier;
- a first stage of the discriminator;
- a second stage of the discriminator with hysteresis;
- a dynamic asynchronous latch comparator with transitional positive feedback;
- a differential transmission line current driver with pre-emphasis; and
- a coplanar transmission line.

These functional blocks are shown in Figure 51 with the exception of the transmission lines. The preamplifier uses a cascode amplifier with an NMOS input transistor biased at 40 µA of dimensions 9.6 µm width and 300 nm length with a simple resistive (200 kΩ) and capacitive feedback defining the pulse gain at a level of 30 mV/fC. The dimensions of the input transistor were optimized for the detector capacitance of 250 fF. The low value of the input capacitance allows the use of a simple cascode stage offering very high bandwidth (1 GHz gain bandwidth product) at very affordable power consumption (60 µW). The open loop gain is in the range of 45 dB, the input impedance of the preamplifier stays in the range of 1 to 2 kΩ for 1 GHz bandwidth, which is sufficiently low to provide efficient charge collection from the detector. Crosstalk between neighbouring pixels is expected to be less than 4 %.

The first differential pair applies the external differential threshold voltage $V_{T1} - V_{T2}$ to the input signal. The fully differential structure of the comparator provides very good rejection of common mode noise from the digital power supply and good threshold uniformity. The entire preamplifier-shaper circuit has a gain of 70 mV/fC with CR-RC3 characterics and a peaking time of around 5.5 ns.

The following three stages of the comparator provide high sensitivity together with very high speed at a power of 50 µW. The simulated time jitter (transient noise simulation) of the full front end stimulated with 3 fC signal is in the order of 30 ps RMS. The predicted ENC values are shown in Figure 52. These simulation results show that for a detector capacitance of 250 fF the predicted ENC is below 200 e^- .

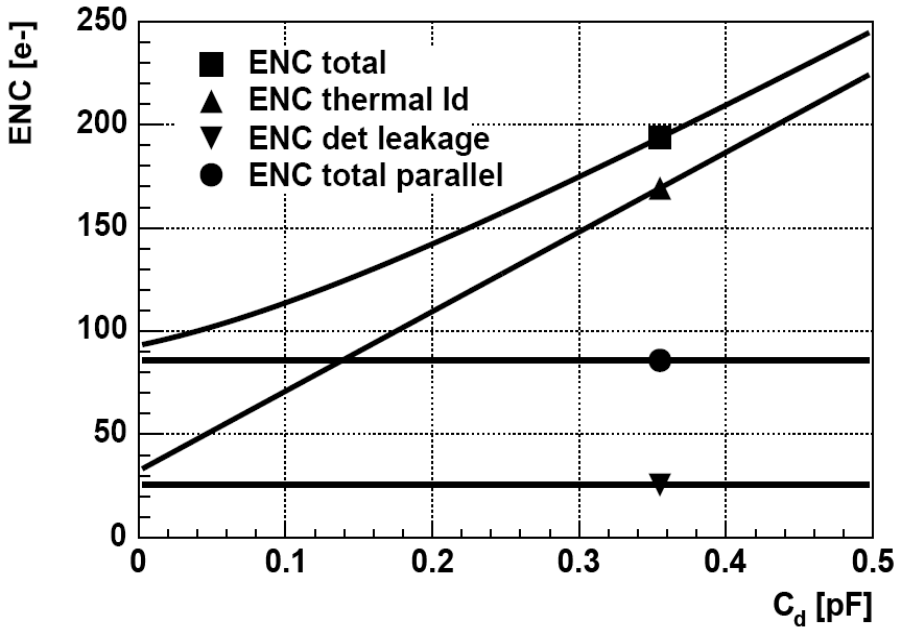

Figure 52 Noise analysis for 0.13 µm CMOS technology.

Figure 53 shows simulations of the time walk and Figure 54 of the time over threshold pulse width with a detector capacitance value of ~ 250 fF, giving a time walk of ~2.1 ns for an injected charge of 5.25 fC. The ToT discriminator pulse width simulation gives values from ~8.45 ns for an input injected charge of 1 fC to ~14.4 ns for an injected charge of 5 fC.

The input of each pre-amplifier is connected to the sensor bump bond pad and a calibration capacitance of 20 fF. The other terminal of the capacitance is connected to a calibration line accessible outside the chip. A voltage step on that line injects charge into the preamplifier. The rise time of this voltage step controls the lengths of the charge injection duration.

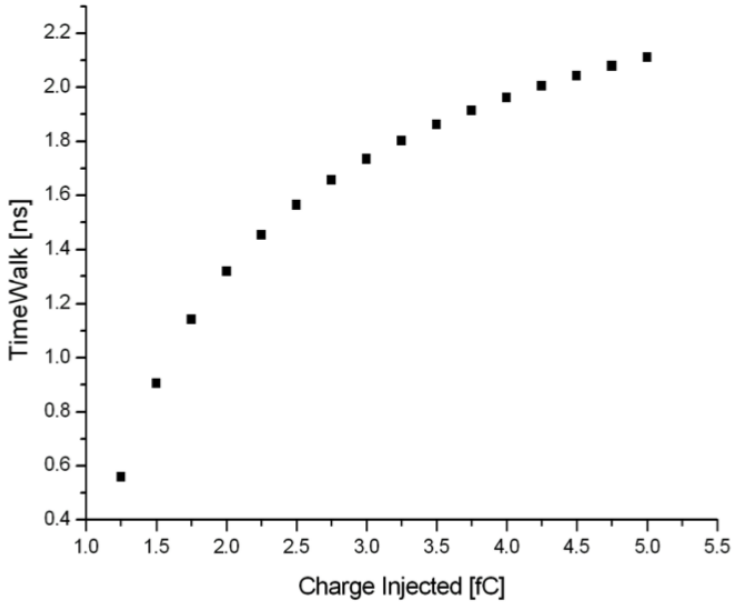

Figure 53 Simulated time walk for a threshold of 0.7 fC, reference is 1 fC. Reference where time walk is 0 if for 1 fC injected charge.

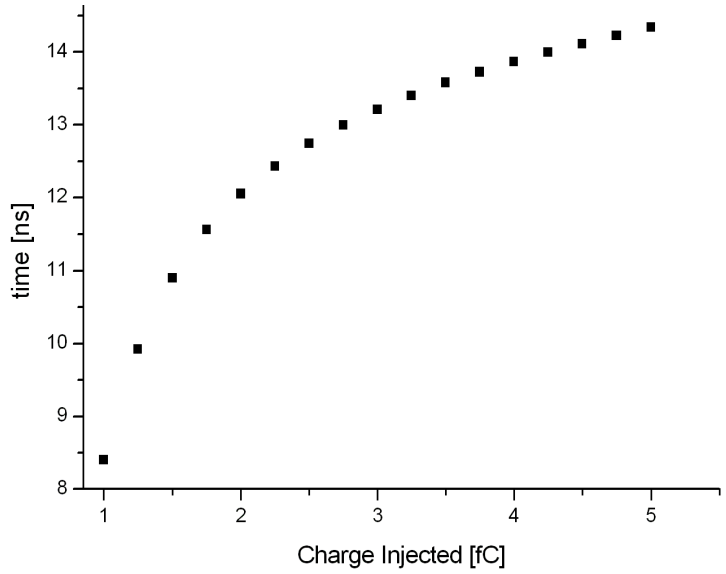

Figure 54 Simulated time over threshold pulse width of the ToT discriminator for a threshold of 0.7 fC.

2.3.4.5.2 EOC Bus System

During the transmission of the discriminator signal to the EOC region the steep rising and falling edges must be preserved to keep jitter values low. A current driver in the pixel cell is connected to differential transmission lines. The traditional use of CMOS inverters as repeater stages is avoided because they introduce delays that scale linearly with the wire capacitance and hence the wire length. Also the big CMOS swing limits the transmission speed. The use of a transmission line limits the resistive energy losses, where the high frequency components of the signal travel more quickly than the low frequency [34] [35].

Each pixel discriminator is connected to a differential driver, which switches a current source with a value of 100 µA when a hit occurs. A second DC current source equalizes the DC current during steady state time. The current signal swing defining the logic level in coplanar waveguides has been specified to ±100 µA, with a differential swing of 240 mV. The series resistance of the line varies for pixels close to the EOC and pixels at the other side of the matrix (distance=13.5 mm). However, this effect does not disturb signal discrimination because the far end amplitude does not change with the pixel to receiver distance thanks to the current drive operation. The high series resistance of lossy transmission lines in the CMOS process has an impact to the signal integrity of the bus. It degrades signal edge in slowing down rise time signal after 13.5 mm signal propagation. To circumvent this effect a current pre-emphasis is implemented in the line driver. Injecting a fast current pulse on the edges of the hit pulses does the signal pre-emphasis, achieving a rise time of 1 ns for the farthest pixels.

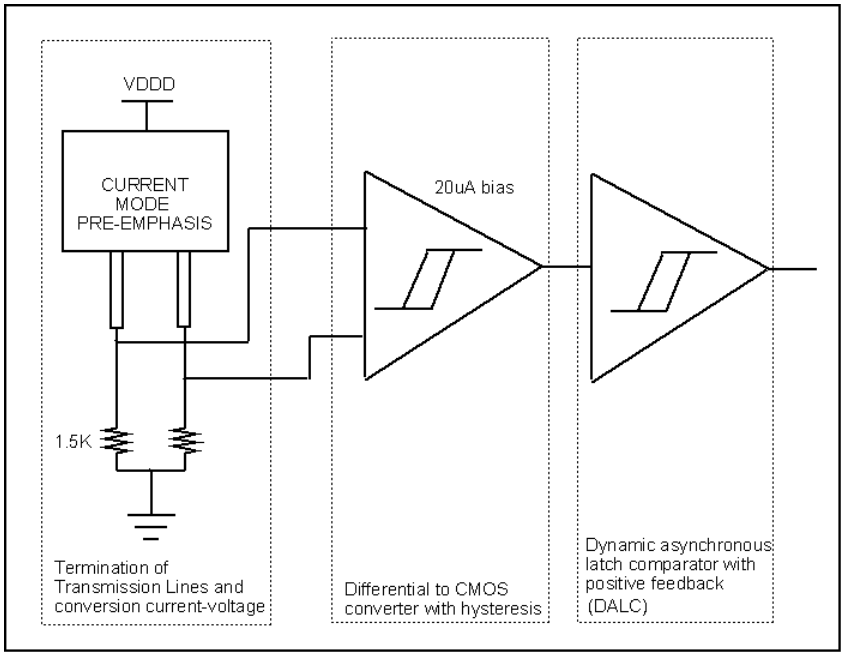

Figure 55 Schematic of the transmission line receiver.

The transmission line receiver consists of 2 different parts. The first one is a differential to CMOS converter with hysteresis and a 20 µA bias current. The second part is a dynamic asynchronous latch comparator with positive feedback (DALC). The DALC generates the fast transition signals without consuming static power by using positive feedback during transition to increase transition speed. A circuit that senses the differential current with a differential common gate cascode stage has been

applied. The block diagram of the receiver is presented in Figure 55. The receiver input stage biasing is provided directly by the static current of the transmission line of about 150 µA. The output of the receiver input stage is sensed with a broadband differential to single ended amplifier stage whose output is converted into a fast digital CMOS level signal by a DALC stage which generates pulses edges of 50 ps to drive the TDC inputs.

2.3.4.5.3 End of Column Logic

In the EOC architecture the arrival time is measured using delay locked loop (DLL) based TDCs, where buffers connected in series are used as base elements. The TDC circuits used in the end of column logic are based on previous developments done in 0.25 µm CMOS technology [32] and in 0.13 µm CMOS technology [33]. The propagation speed of these buffers is controlled by a voltage. The base frequency is sent to the first buffer and the phase relationship between the input of the first buffer and the output of the last buffer is monitored. If the phase difference is positive (the output is too slow) the control circuit of the DLL will increase the delay cell control voltage and thus the speed of the buffers and inversely decrease it in case of a negative phase difference. The phase detector continues to measure the phase of the input and output clock signals throughout the operation of the system continuously correcting the speed of the buffers. Then the individual buffers divide the base clock period by the number of buffers in the chain. Upon arrival of the hit the hit registers capture the state of these buffers and the state of a synchronous base frequency clock counter. Figure 56 shows the basic principle. For the EOC architecture a base frequency of 320 MHz is applied and 32 buffers divide the period of 3.125 ns in bins of 98 ps.

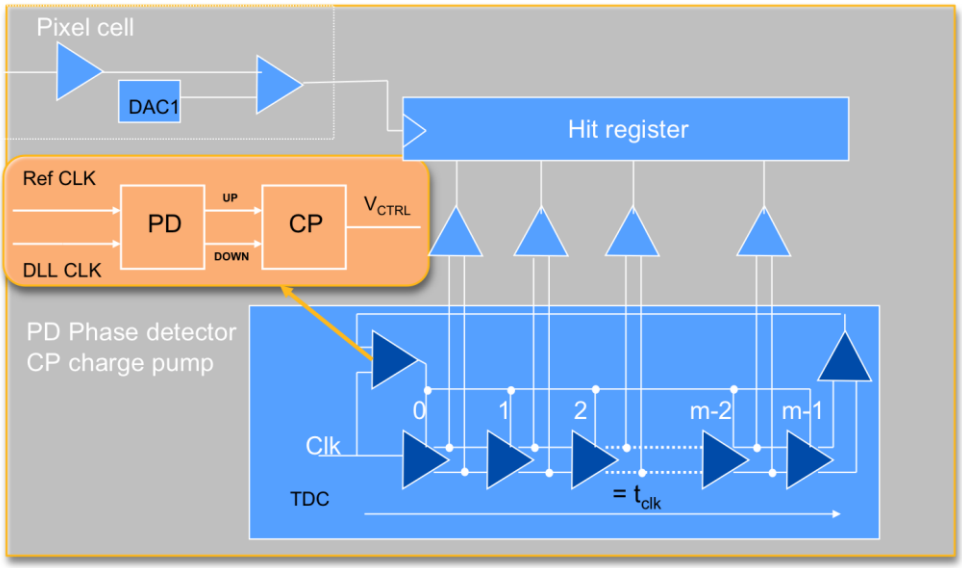

Figure 56 DLL based TDC principle.

It is not required that each pixel cell has its own hit register. An efficiency for the highest rate centre column of 99.5 % can be achieved by multiplexing 5 pixels to one hit register. A digital arbiter circuitry has been designed which sends hit pulses to the TDC banks and masks the rare simultaneous signals from the same group of 5 pixels. These overlaps are indicated with the address of the masked pixel in the output data stream. In order to process clusters and pixel charge sharing, the pixel group feeding each hit arbiter is formed from pixels that are not adjacent, see Figure 57.

The 45 pixels of each column are grouped in 9 sets of 5. Correspondingly 9 rising edge hit registers and 9 falling edge hit registers are used for each column.

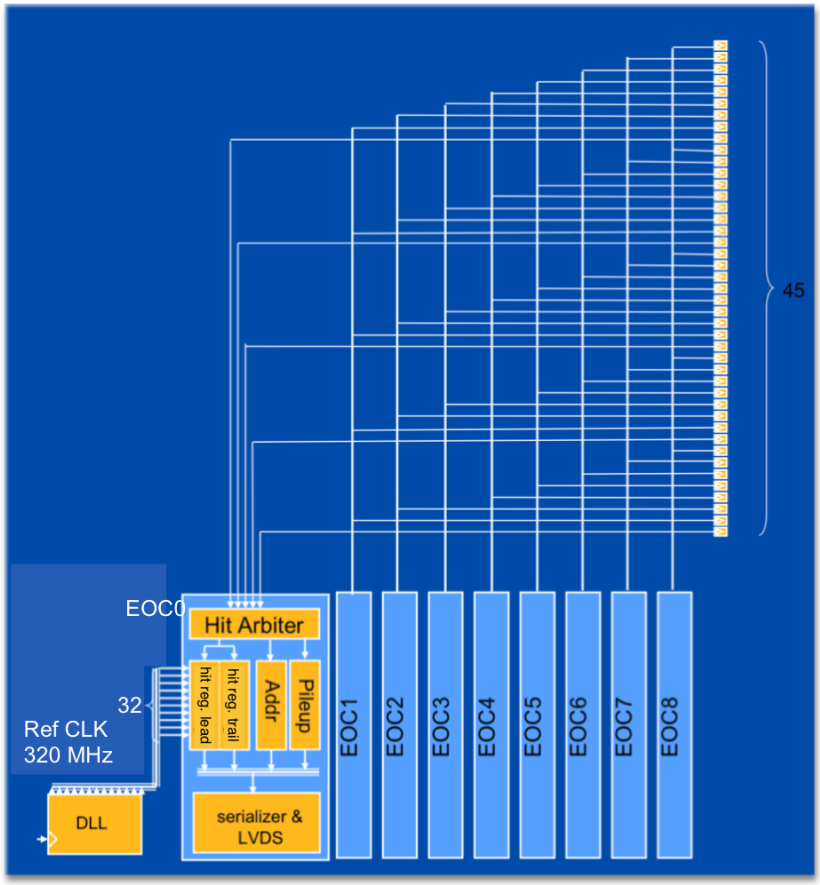

Figure 57 Block diagram of one column with transmission line scheme.

The GTK read-out chip will contain a number of DLLs and the buffered DLL signals are distributed to the hit registers. The output of the DLL buffer is differential and is converted to single-ended CMOS at the input buffer of each TDC bank. The hit registers in the TDC are built using 32 D-type flip-flops with rising edge trigger. The output of the receiver cell provides a rising edge trigger for both leading and trailing edge hit registers.

The sensitive analogue circuits are in the pixels and the digital power supply domain is geographically well separated from the analogue supply, reducing noise influence on the pixel cells. EOC circuits have been designed so that the receiver bank and the hit register bank for a full column and the digital column read-out circuits fit inside the 300 µm wide envelope dictated by the pixel size.

Figure 58 shows the block diagram for the full chip with the EOC architecture.

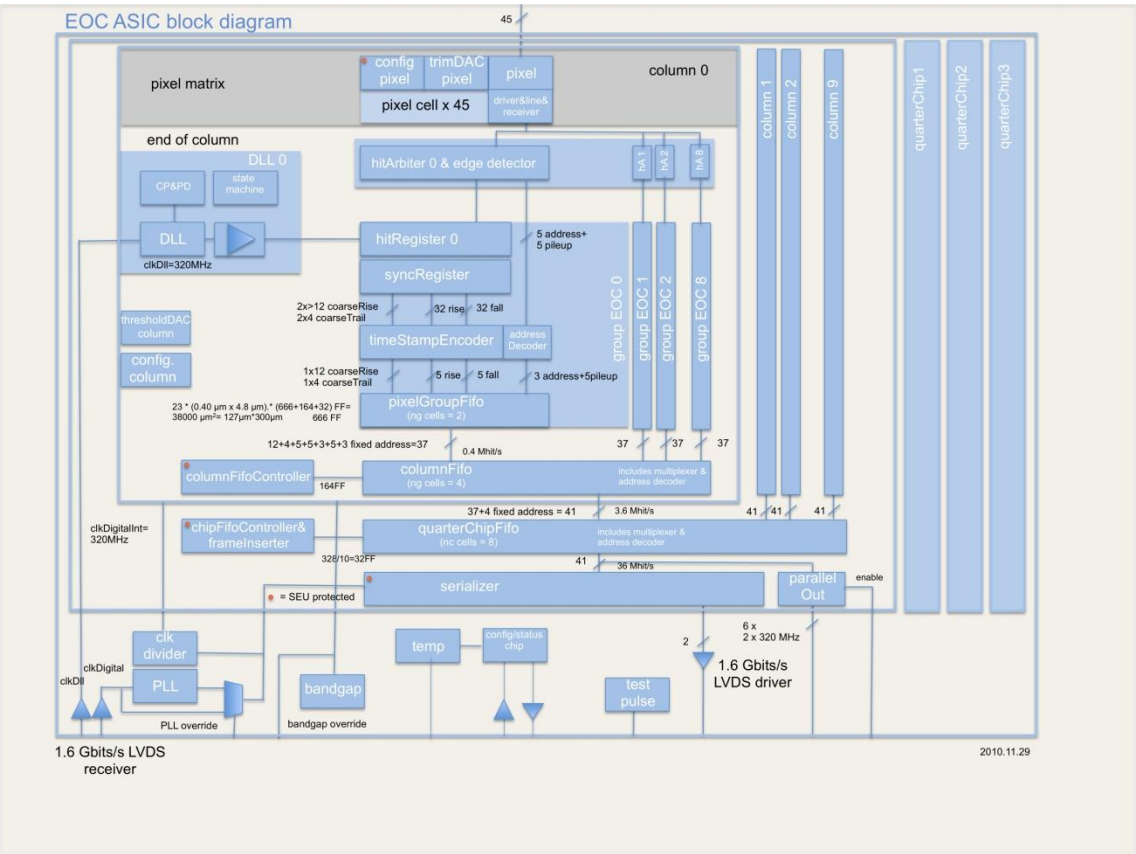

Figure 58 Full chip EOC block diagram.

2.3.4.5.4 EOC Demonstrator

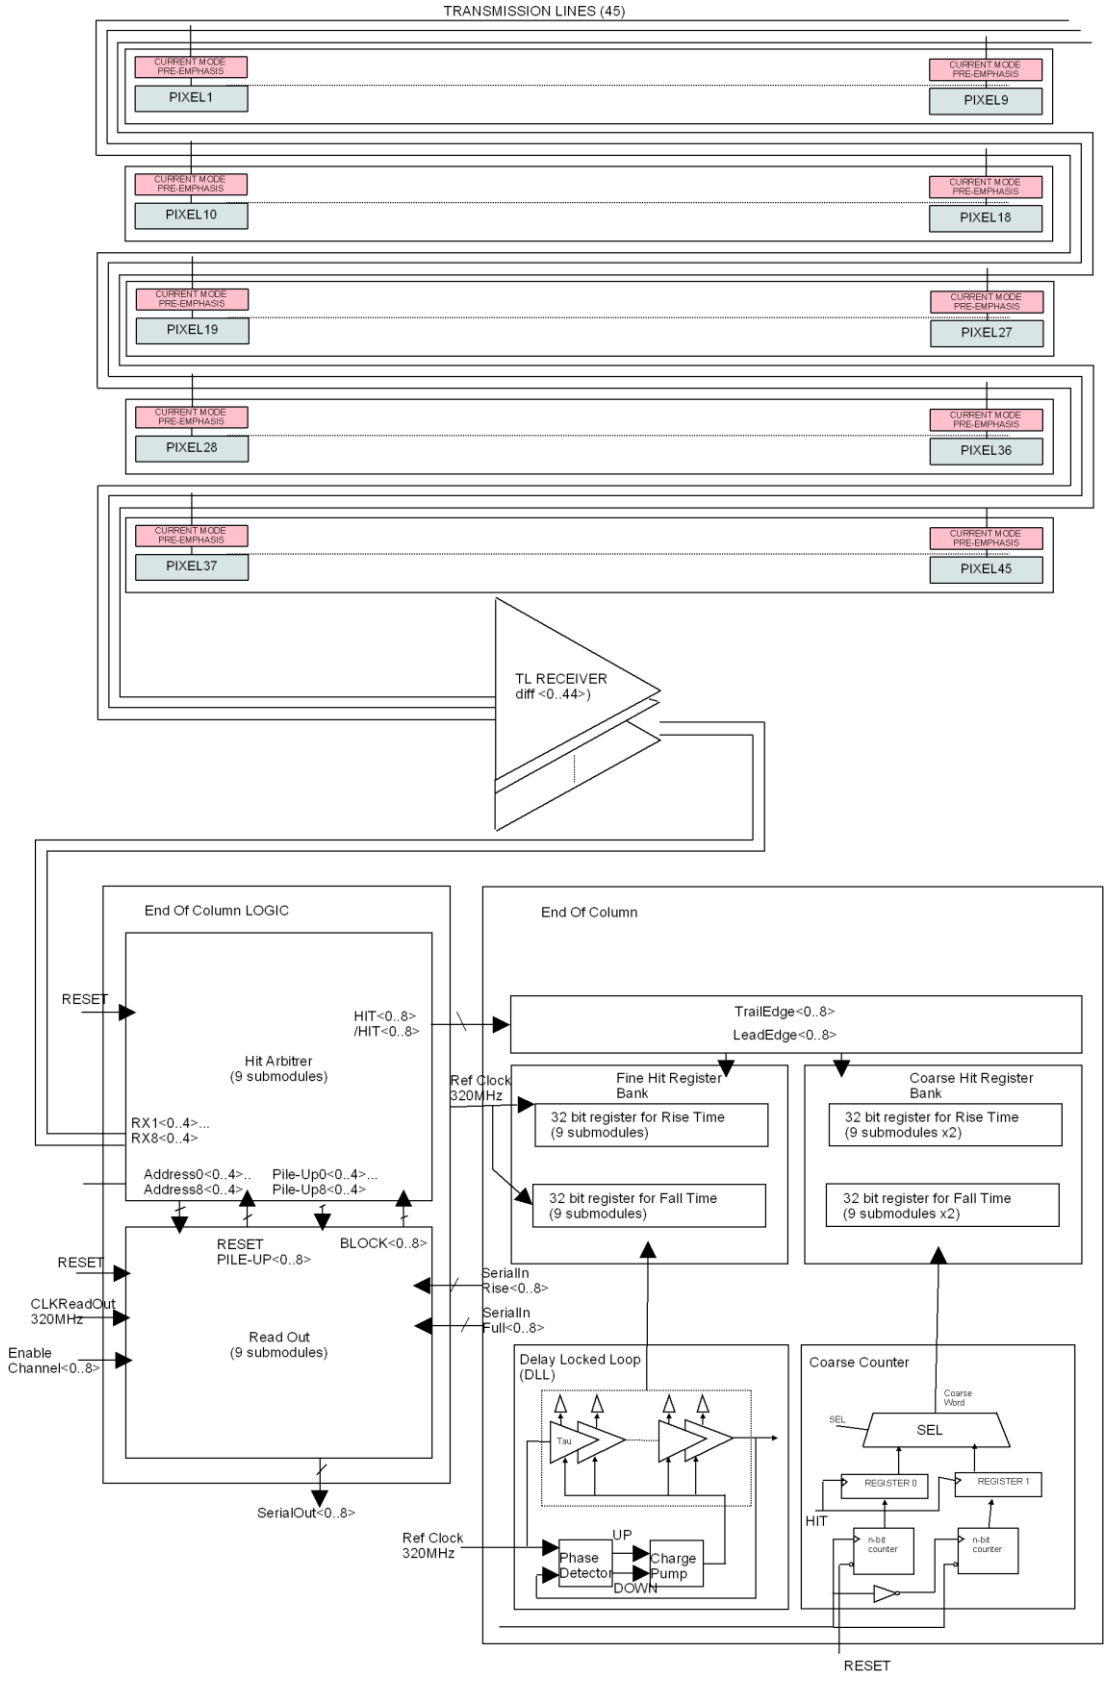

Figure 59 Block diagram of the 45 pixel demonstrator.

A demonstrator ASIC for the EOC column architecture with the ToT time walk compensation technique has been developed. Figure 59 shows a block diagram of the demonstrator. The demonstrator contains one full 45 pixel column folded back on itself to save silicon area. In addition to this column, test structures with 15 pixels have been included for the characterization of the analogue pixel cell. This demonstrator serves experimentally as proof of principle for the analogue front-end, the ToT time walk compensation technique, the low-jitter transmission of the signal along the column and the time tagging of the rising and falling edge maintaining a time walk compensated arrival time measurement resolution of less than 200 ps. The homogeneity of the pixel matrix will be evaluated. The entire chain from analogue pixel input, via the amplification stages and discriminator to the transmission line amplifier, transmission line and receiver, the hit arbiter logic and the TDC has been fully realized as for the final chip. The layout of the EOC circuits largely follows the 300 μm pitch defined by the width of the column. The data encoding and read-out from the TDC has been simplified as the implementation is straight forward.

The demonstrator is organized in such a way that the 45 pixel column feeds a set of 9 hit arbiters and 18 hit registers connected to one DLL. Rising and falling edge hit registers of the 9 pixel groups are read out using 9 x 320 MHz serializers independently. In order to asses the influence of digital noise on the DLL and analogue parts, the read-out clock can be switched on and off for each of the 9 streams independently. A set of 9 test pixels is connected to a second instance of the EOC column logic and operates independently. Test outputs along the processing chain allow qualification of individual blocks. Figure 60 shows the layout of the prototype.

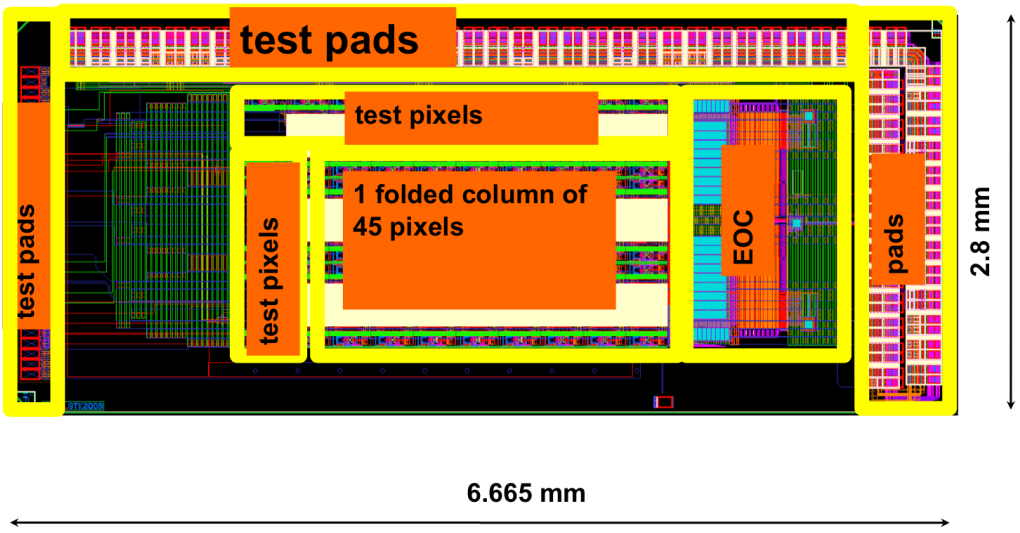

Figure 60 EOC column demonstrator layout.

The qualification of the demonstrator includes the testing of each individual block, the full processing chain, the full 45 pixel column and its uniformity using programmable test pulses on each individual pixel.

After the bonding of the demonstrator chip to the sensor, laser and beam tests were made to qualify the assembly of sensor and read-out chip (see section 2.3.5 on page 119).

Test setup

A test system was built up around the EOC GTK demonstrator ASIC. Figure 61 shows a schematic of the test bench containing the most important components. This system is based on an Altera Stratix III FPGA development card. A daughter card was designed to host the EOC ASIC and interface to the development card. This services card provides power, current and voltage references, analogue and digital signal buffering and SMA/SMC connectors to bring in and out test signals. Control of references is done through a single I2C bus upon which there are 8 voltage output DACs, each one providing a signal between 0 V and 1.25 V to 12 bit precision. Where necessary, current source and sink conversion is done using transistors. An EEPROM provides non-volatile digitally accessible identification of the card. A temperature sensor gives a local and potentially remote measurement of the temperature to about 1^{o}C precision. The test setup includes a Tektronix AFG3252 pulse generator and a LeCroy WaveRunner or WavePro oscilloscope. The test system uses a private gigabit Ethernet network to connect the various components together. A C++ software framework provides control and configuration of all the components of the setup, allowing automated sweeps to be performed in a detailed manner over all pertinent parameters. The collection and handling of the data output from the setup is also dealt with in the framework using a run oriented approach, in which a unique run number is generated for each instance of a measurement. This permits a given measurement to be repeated as many times as desired without overwriting any pre-existing data. The high precision DLL clock is provided by the clock generator (SRS CG635), whereas the readout clock comes directly from the FPGA

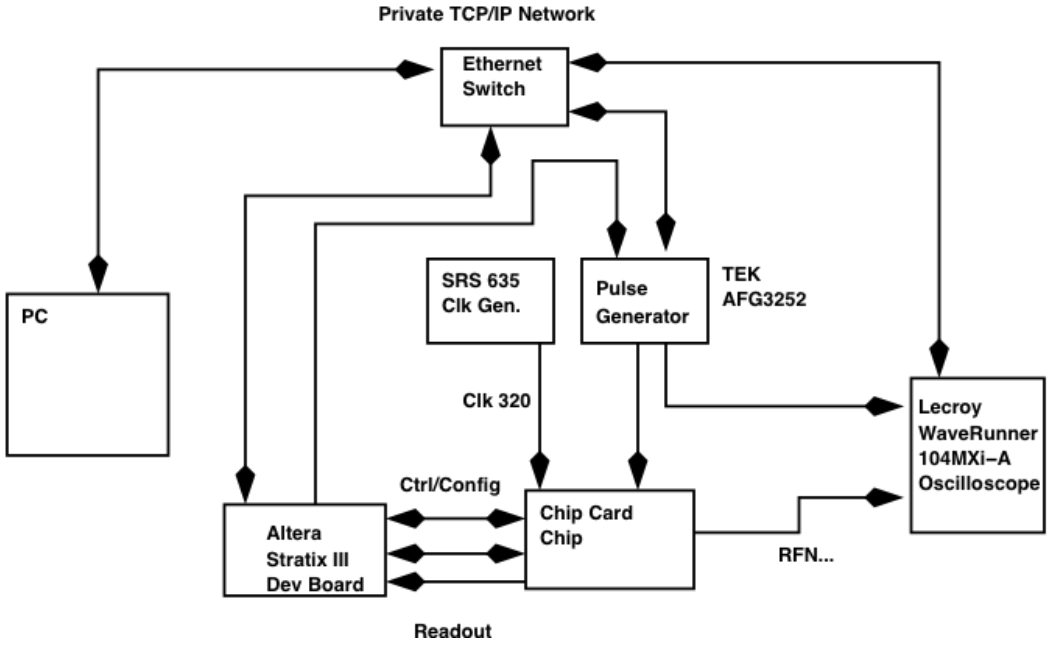

Figure 61 A Schematic view of the automatic measurement test setup.

Test results

Test principle:

All tests have been done with the nominal 320 MHz read-out and TDC reference clock frequency, corresponding to a TDC time bin of 97.7 ps.

Figure 62 shows the definition of particle arrival time t_0, the leading edge discriminator output time t_1, the trailing edge discriminator output time t_2 and the time-over-threshold pulse width t_{tot}.

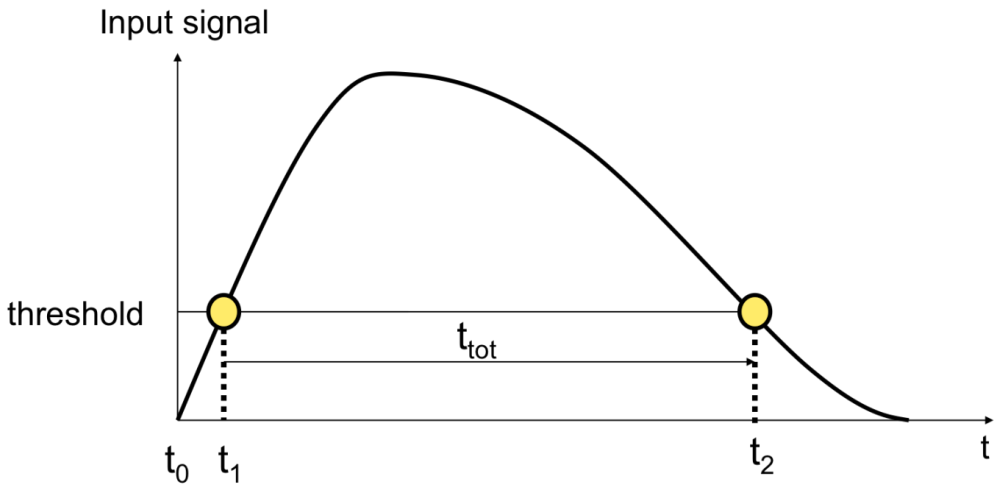

Figure 62 Definition of signal times: t_0 particle arrival time, t_1 leading edge discriminator output time, t_2 trailing edge discriminator output time and t_{tot} time-over-threshold pulse width.

The principle behind the testing is to present charge in varying amounts at the input of the front end amplifier. In the absence of a detector this is achieved by the inclusion of a capacitor in the pixel, one terminal of which is connected to the front end amplifier, the other being connected via an isolation circuit to a global, low resistivity, calibration net. The capacitance of this capacitor is expected to be approximately 20 fF, with an uncertainty of about 10 % on a given chip and a channel to channel spread of about 1-2 %. When a given pixel is selected, a voltage step present on the global calibration net causes a charge to be injected at the front end of that pixel. The orientation of the voltage step decides the polarity of the charge, and the rise time determines the duration of the charge pulse; the current being the time derivative of the voltage. Since the expected detector technology is p-in-n, the results presented are acquired with a positive going voltage step. The duration of the test pulse (rise time of injected charge) is varied from 2.5 ns to 6.5 ns in order to emulate the expected elongation of the current pulse from the detector during its lifetime resulting from the accumulated radiation damage. The total amount of charge injected is varied over the expected range of 1-10 fC. This requires a voltage step on the calibration line of between 50 mV and 500 mV. Once the charge injection pulse is complete, the signal level is returned to the starting value over a period of time of greater than 1 microsecond. This is long compared to the front end amplifier peaking time, and as such, avoids the opposite polarity pulse occurring at the output of this amplifier. This behaviour is what would be expected with the detector. The repetition rate of this pulse can be controlled via the test setup software. It may be run continuously, triggered by the chip controller card, or by software. Additional synchronisation with the DLL clock is possible to the level of around 15 ps (RMS).

In the following when referring to full-chain measurements an electrical test pulse has been injected into the pre-amplifier and the digital time stamp of the TDC bank is evaluated, including all components along the chain – pre-amplifier, ToT discriminator, transmission line driver, transmission line, transmission line receiver, TDC and hit registers. Analogue front-end measurements evaluated the pre-amplifier and ToT discriminator, whereas TDC measurements are done by injecting digital

input pulses directly to the TDC bank and reading the digital time stamp. Figure 63 shows graphically in the block of the EOC processing chain and the scope of the measurements performed.

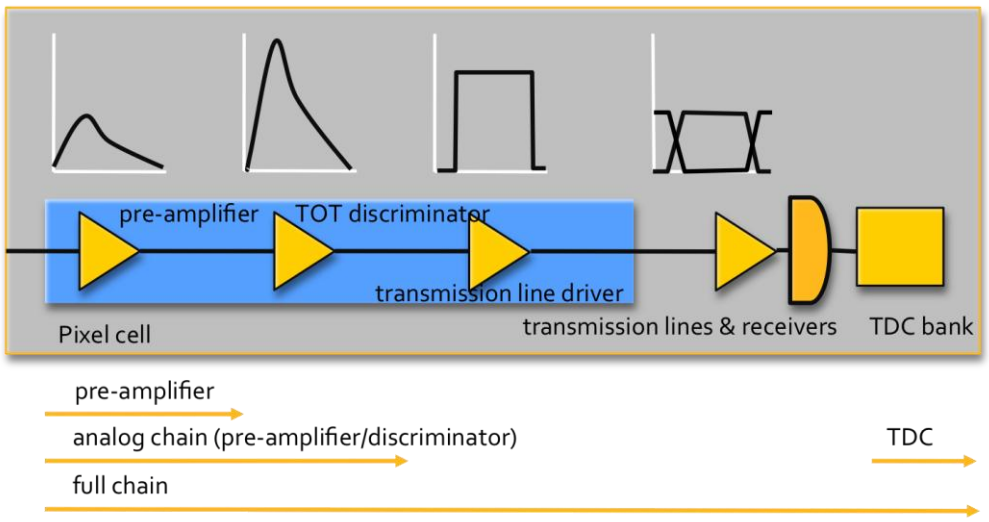

Figure 63 Block diagram of EOC processing chain.

Analogue frontend behavior

Unless stated otherwise, the charge injection duration is 2.5 ns and the read-out clock is activated. All measurements concerning the analogue frontend behaviour in this section are taken from the discriminated test pixel output.

- Analogue frontend time walk and jitter over input charge

Figure 64 shows the analogue front-end t_1 and t_2 responses over input charge and several discriminator thresholds. For a threshold of 0.7 fC the full input charge range time walk of t_1 is 2.2 ns and the ToT-pulse width is 7 ns. From this plot we also can see that the relationship between t_1 and time over threshold (t_2-t_1) is monotonic which required for a time walk correction.

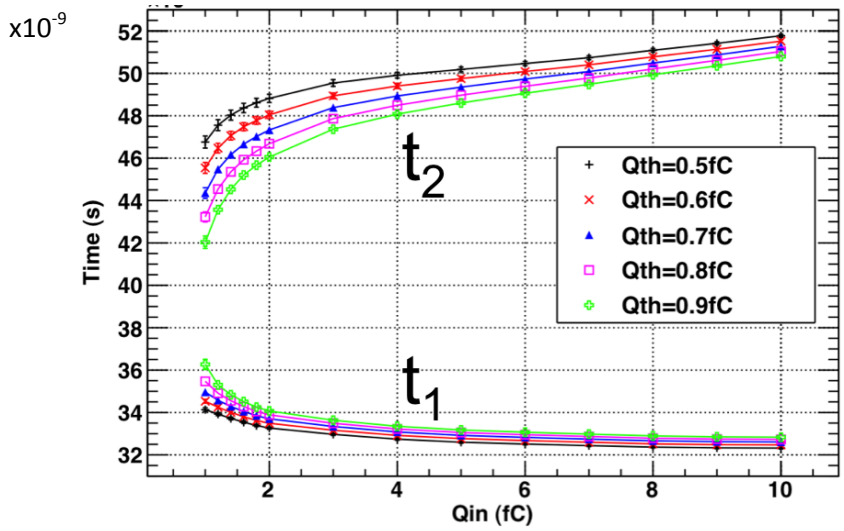

Figure 64 Analogue front end response for t_1 and t_2 over input charge for 2.5 ns input charge charge injection duration.

Figure 65 and Figure 66 show the analogue front-end t_1 and t_2 rms jitter respectively over input charge for several threshold values and input signal charge injection duration of 2.5 ns. These data are taken with 10000 traces per point, reducing the statistical error to a negligible level on the scale of interest. The behaviour shows that as the charge increases, the jitter appears to approach an asymptote. As the charge nears the threshold level, the jitter grows sharply. This can be explained as the noise on the signal being translated into the time domain by the finite gradient of the waveform as it crosses the threshold. A smaller gradient leads to a longer lever arm for this action.

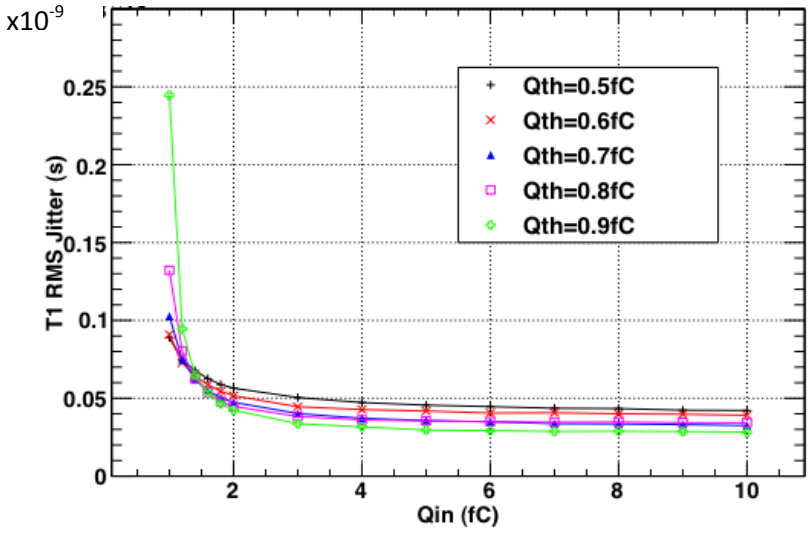

Figure 65 The t_1 rms jitter[ns] versus input charge for several thresholds and input signal charge injection duration of 2.5 ns.

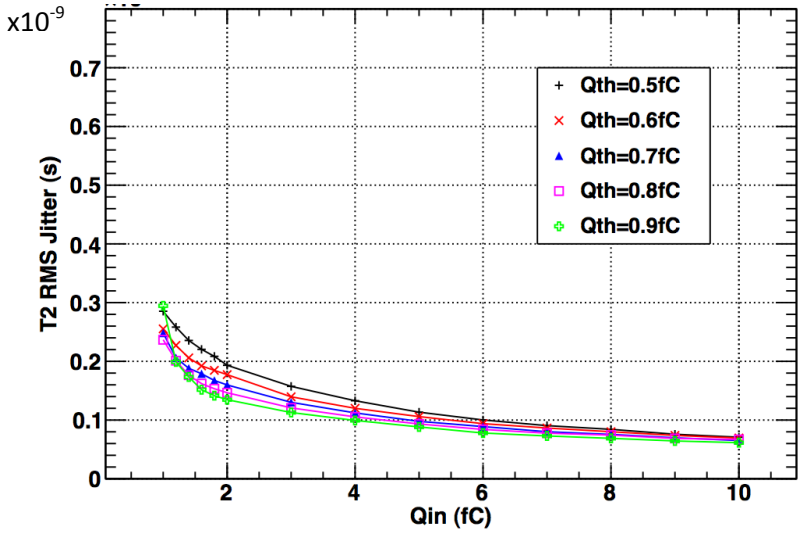

Figure 66 Analogue front-end t_2 rms jitter [ns] over input charge for several threshold values and input signal charge injection duration of 2.5 ns.

The t_2 jitter exhibits the same features, although the values are approximately a factor of 3 higher. This is consistent with the ratio of the gradients of the leading and trailing edges of the signal. The jitter of t_1 around the most probable charge value of 2.4 fC is around 55 ps. The jitter of t_2 for that charge is 150 ps. These jitter values do not contribute equally to the time walk compensated jitter,

but the t_2 jitter contribution needs to be weighted by the gradient of the time over threshold t_2-t_1 and t_1, which is around 1/3.3. The jitter depends on the level of the threshold. The preamplifier output is steepest in the mid section. Thus for high thresholds and charges higher than 2 fC the performance is consistently better. At lower charges, the higher threshold settings are closer to the pulse peak where the lower slew rate causes a corresponding degradation in performance. Thus the curves cross over each other, which can be seen at approximately 1.3 fC. This indicates that the threshold can be optimised for the charge spectrum issued from the detector.

- Behaviour of pixel at design rate:

Figure 67 shows how the t_1 RMS might be expected to vary with the frequency at which the pixel is firing for the nominal charge threshold of 0.7 fC and nominal charge injection duration of 2.5 ns. During this measurement, fixed frequencies between 10 kHz and 150 kHz, in steps of 10 kHz, were used over the standard 1-10 fC input charge range. The behaviour is largely flat with some small systematic variations occurring at high charges. These are almost certainly due to signal components with longer time constants appearing as the input charge becomes large, causing baseline fluctuations that contribute to the noise on the subsequent pulse. These results indicate that the pixel will indeed operate in a well behaved manner up to the maximum required by the experiment.

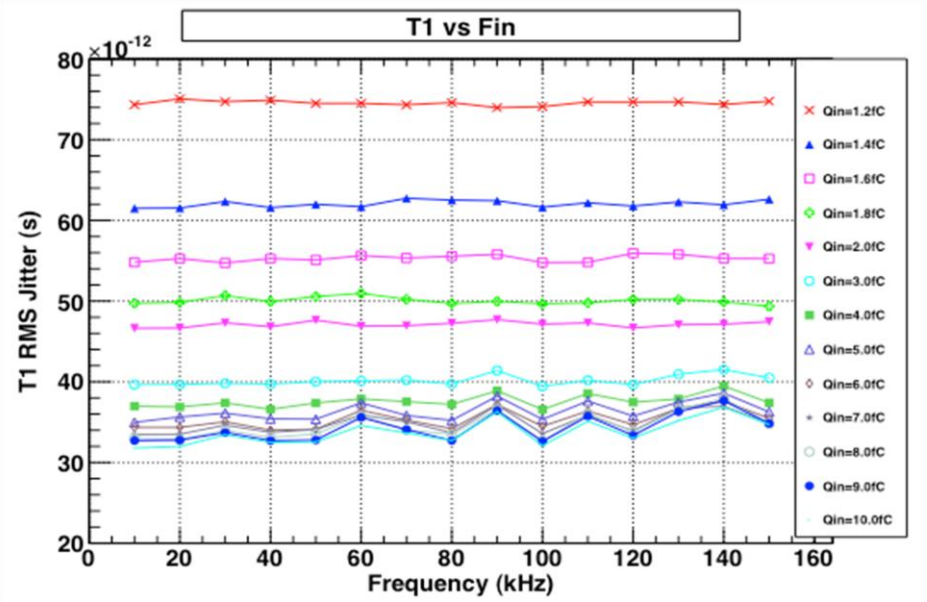

Figure 67 The t_1 rms jitter over pixel hit frequency for a threshold of 0.7 fC and various values on input charge.

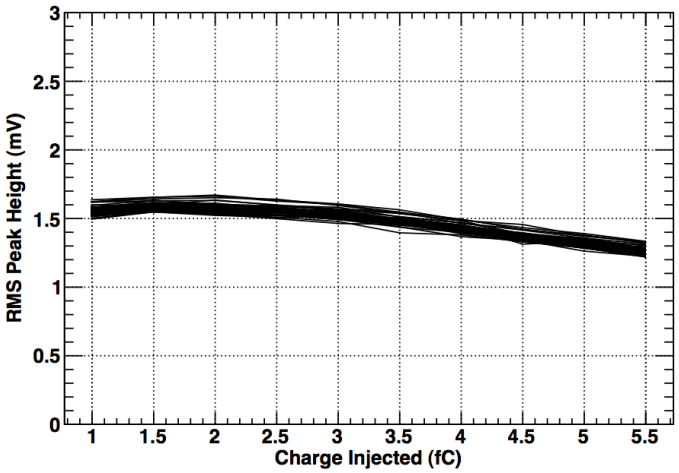

Figure 68 Rms peak height over injected charge.

Noise front-end

Figure 68 shows the rms of the peak height as a function of charge. With a gain of 70 mV/fC this corresponds to approximately 130 electrons. The measurement has been performed by varying the threshold for a given charge injection and evaluating the discriminator response (measuring the s-curve). The same noise value can be obtained by evaluation of the preamplifier output pulse without input signal. Figure 69 shows the rms of the base line to be 576 µV. This value corresponds with a to approximately 132 electrons (the gain of 70 mV/fC on the ASIC together with the external instrumentation amplifier yields a gain of 27 mv/fC).

Figure 69 Noise level on preamplifier output without input signal.

TDC behavior

TDC INL/DNL

The TDC has a digital trigger input to permit the evaluation of its performance in a manner decoupled from that of the front end and the transmission lines. The linearities shown in Figure 70 and Figure 71 were measured by applying a large number of triggers (28 million) from a flat distribution on the time scale of a clock cycle (3.125 ns). An ideal TDC would exhibit an equal probability for a given trigger to fall in any of the bins, thus by estimating the probability each bin has of receiving a trigger from the ratio of the bin content to the total trigger count, an estimator for the bin width can be obtained. This estimated bin width is then used to calculate the differential and integral non-linearities with the values of 0.17 and 0.27 of an LSB of 97 ps. The non linearity will corrected for by means of a look up table off line.

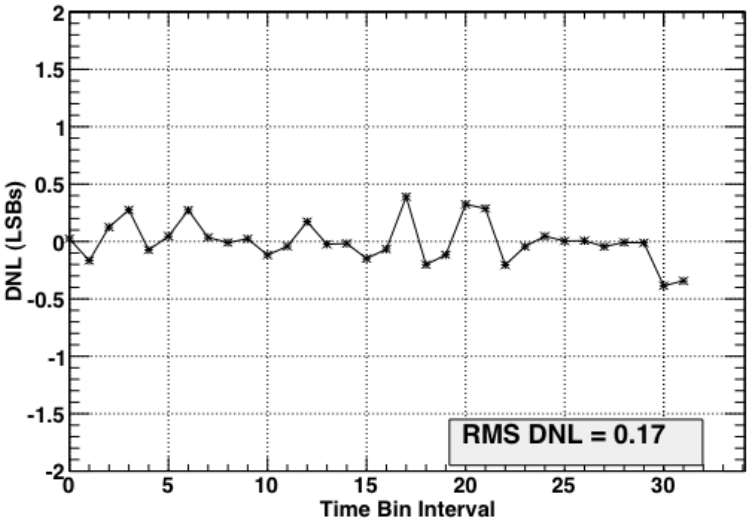

Figure 70 TDC differential non-linearity.

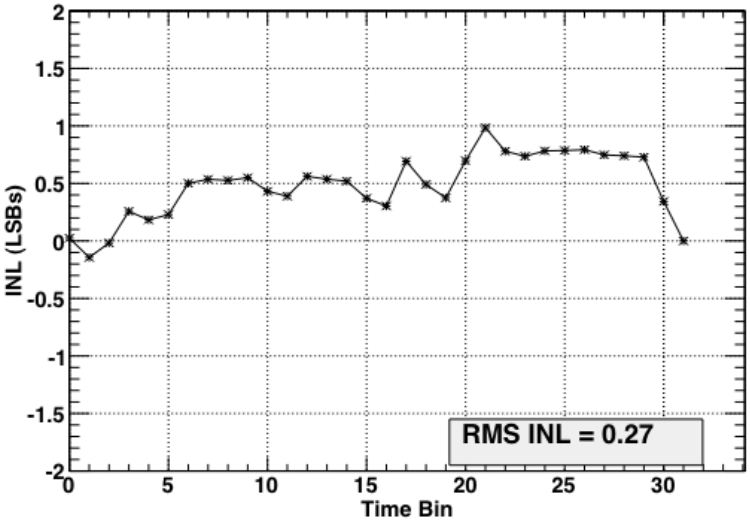

Figure 71 TDC integral non-linearity.

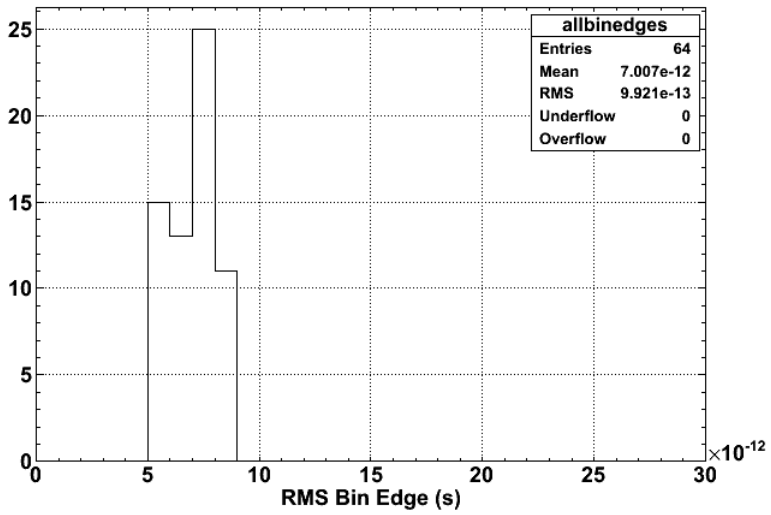

Figure 72 TDC jitter measurement.

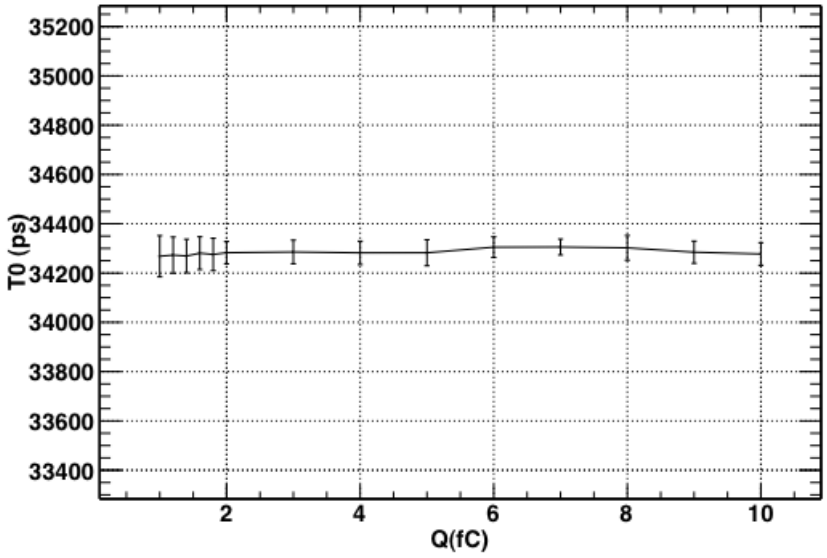

Figure 73 The t_0 - full chain time walk after compensation.

TDC jitter

By measuring the phase of the TDC digital trigger input to the DLL clock at the input to the chip on a case by case basis, two measurements of the same event are recorded: one from the oscilloscope and the other from the TDC. Taking the time given by measurements on the oscilloscope data as being the "true" time of the event, and histogramming these as a function of the TDC fine time bin, the tails on the distributions can be used to give an estimate of the jitter introduced by the TDC: i.e. when the TDC reported a given time code but should have reported the preceding (or subsequent) one. The jitter of the TDC has been measured to be 7 ps rms on average (see Figure 72).

The low jitter value shows the appropriate design and stability of the DLL control loop and the delay buffers. The performance of the TDC is dominated by the quantisation error.

Full chain arrival time resolution

Unless stated otherwise, the charge injection duration is 2.5 ns and the read-out clock is activated.

- Full chain time walk or compensation error and jitter over input charge.

The resolution of the time walk compensated time measurement is comprised of the systematic error given by the time walk compensation and the statistical error given by the jitter of the electronic components. The time-over-threshold (ToT) method gives in addition to the signal arrival time (t_1) the amplitude of the input signal by evaluation of the ToT-pulse width (t_2-t_1).

The systematic error of the time walk compensation can be eliminated. As the TDC transfer curve is known, the non-linearities can be compensated for each time code even improving the time measurement. Figure 73 shows the time walk after compensation, t_0, for a selected time offset of 1000 ps after the rising edge of the TDC clock. The error bars on the corrected time of the input charge correspond to the weighted sum of the full chain jitter of t_1 and ToT pulse width. Figure 74 gives a histogram of the t_0 residuals averaged over all phases and charges. The RMS width of the distribution is 10 ps corresponding to the uncertainty added by the time walk compensation. Figure

75 shows the jitter on t_0 from actual measurements with electrical test pulses. The mean error on the t_0 measurements for the most probable charge of 2.4 fC is around 60 ps.

Figure 76 and Figure 77 shows the rms t_1 and ToT jitter for all 45 pixels overlaid over input charge. The t_1 jitter for the most probable input charge of 2.4 fC is 55 ps. This shows the uniformity across all 45 pixels.

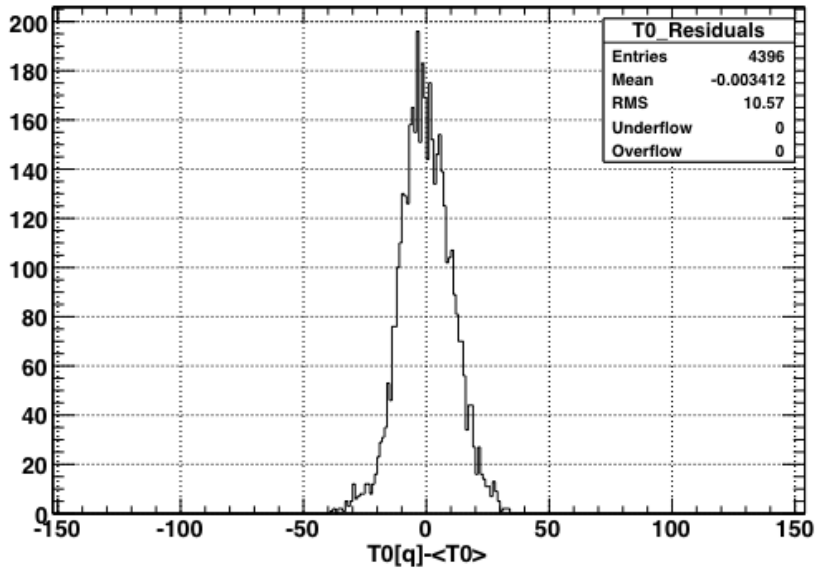

Figure 74 The t_0 residuals for all phase and all charge measurement points.

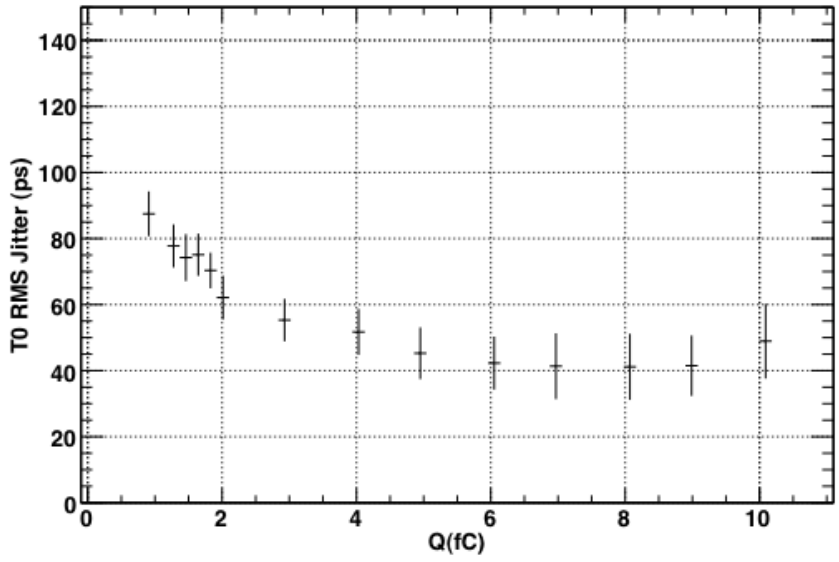

Figure 75 Jitter t_0 or time walk correction error measured with electrical test pulses.

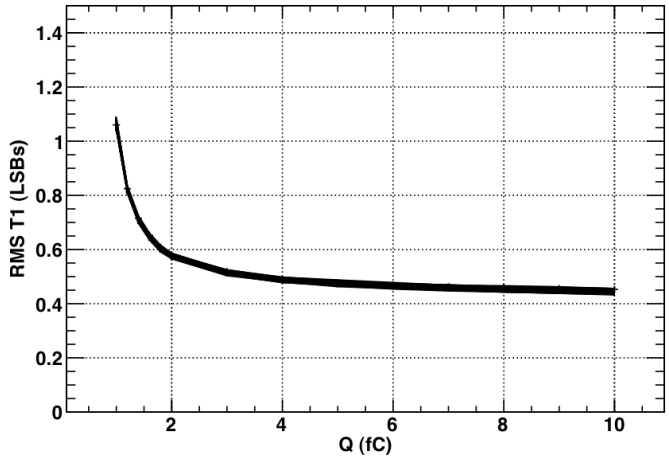

Figure 76 Full chain t_1 jitter over input charge for 2.5 ns input charge charge injection duration.

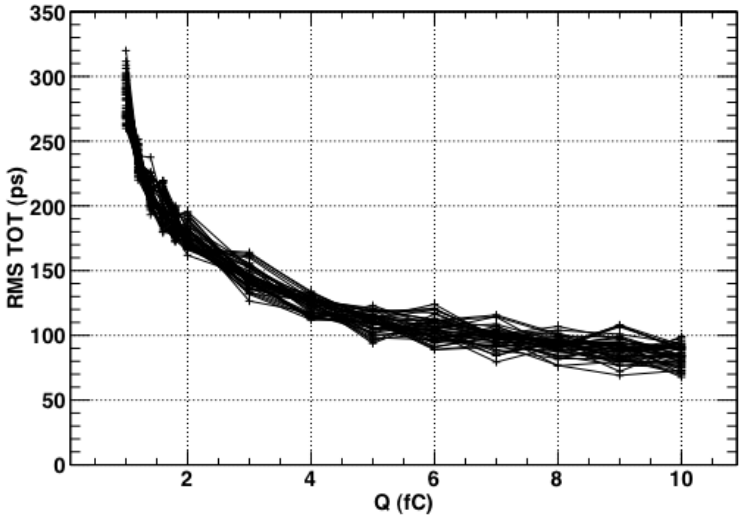

Figure 77 ToT pulse width rms jitter over input charge for 0.7 fC threshold.

Figure 78 Pre-amplifier signal output trace.

Influence of read-out clock on full chain resolution/jitter - digital clock/noise separation

During the measurements presented above the read-out clock was switched on. In order to get a measure of the influence of the digital activity on the analogueue front-end and to show a possible coupling of the DLL and the 320 MHz EOC read-out clock to the front-end a charge of 2.4 fC was injected to the front-end of the test pixel, once with the clock on and once off. It should be noted that the test pixel is at the far end of the EOC area. Figure 78 shows the signal trace after the pre-amplifier. No clock influence can be seen.

S-curves - gain and offset distribution

Control over the global thresholds is given by a 12 bit DAC producing an output from 0-1.25 V, corresponding to a step size of 305 µV. With the design gain of the front end amplifier being 70 mV/fC, this represents a step of approximately 0.00436 fC or 27 electrons. Stepping the threshold across a (large) fixed number of pulses generated from a given input charge and measuring the turn-on curve permits the front end amplifier offset, gain and noise to be determined, though a contribution to the noise value from the pulse generator is included. The integrated noise profile of the pulse can be accurately fitted using a complementary error function or cumulative frequency curve. The mean and standard deviation of the underlying Gaussian give the point where 50 % of the injected charges evoke a response from the discriminator and an estimate of the front end amplifier noise respectively. Repeating this procedure for charges in the range from 1.0 fC to 5.5 fC in steps of 0.5 fC results in the family of curves shown in Figure 79. From these measurements, the gains of all the pixels in the main array have been estimated using a straight line fit to the region from 1.0 fC to 3.0 fC. Figure 80 shows the distribution of the pixel gains for one full column with a mean of 72 mV/fC – the simulation value is 70 mV/fC.

Projecting the gain curves back to the axis permits the offset for each pixel to be estimated. The distribution of pixel offsets for the main pixel array is shown in Figure 80. This gives crucial input into the design of the threshold trim (see section Trim DAC requirement on page 118).

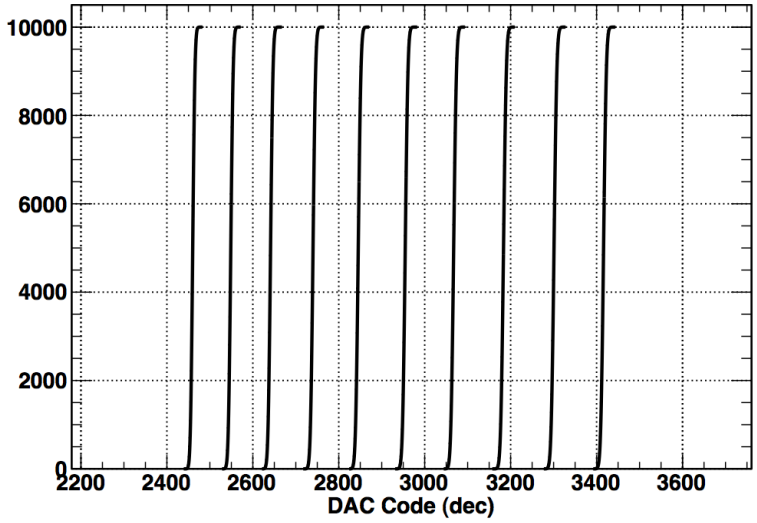

Figure 79 S-curves for several thresholds and 2.5 ns input signal charge injection duration.

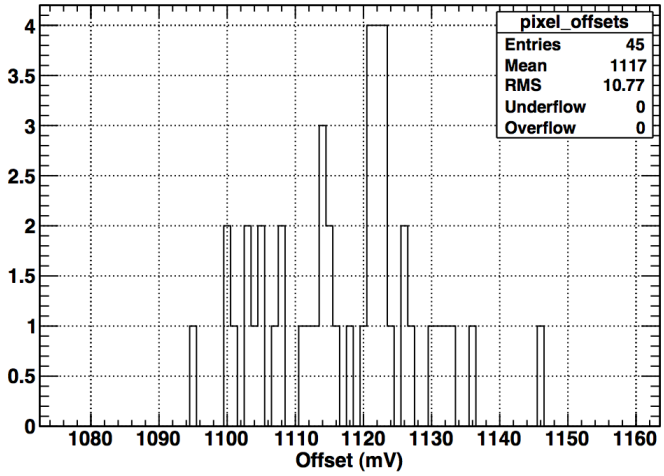

Figure 80 Distribution of pixel offsets for a full column.

Figure 80 shows the peak to peak distribution of the pixel offset for a full column before trimming of ~60 mV.

Behaviour of long transmission lines

In the EOC architecture the signals from the ToT discriminator are sent from the pixel cells over differential transmission lines to the end-of-column area. Depending on the position of the pixel this transmission length ranges from 0 to 14 mm. The maximum lengths exceed the value of 13.5 mm (45 * 0.3 mm) for the final columns as the column in the demonstrator has been folded and the curves of the folded structure add 0.5 mm. Figure 81 shows full chain rms jitter t_1 pulse width over the distance from the end-of-column region. They show that there is no dependency of the position of the pixel with respect to clock and it can be concluded that the transmission does not add to the measurement error.

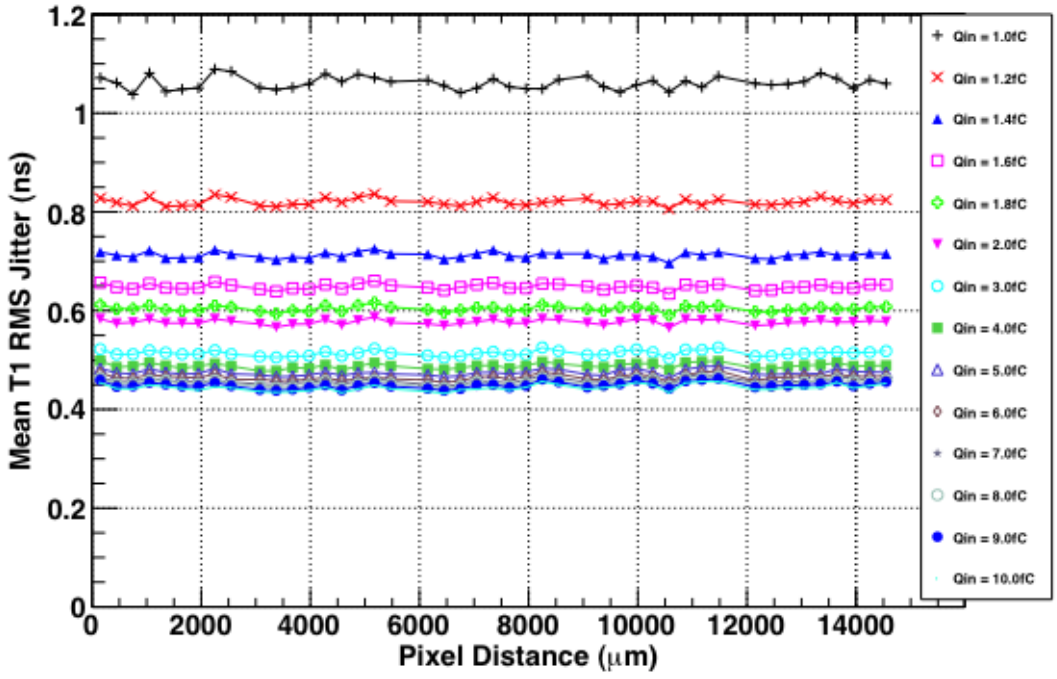

Figure 81 Full chain t_1 jitter over distance from end-of-column

Influence full chain resolution/jitter on signal charge injection duration

As the sensor signal shape will change with radiation, a study on the influence has been done. The nominal test pulse charge injection duration of 2.5 ns was varied up to 6.5 ns.

A summary of the gain as a function of the charge injection duration can be seen in Figure 82; as the distance between the curves for different charge injections do not change the plot also shows that the offset is preserved. The lower effective gain results in a lower signal gradient at the output of the front end amplifier, which in turn causes a degradation in the jitter performance. This can be seen in Figure 83, which shows the t_1 jitter behaviour as a function of injected charge for the same range of charge injection durations. Note that the threshold here is set to 0.7 fC, including an adjustment for the new effective gain.

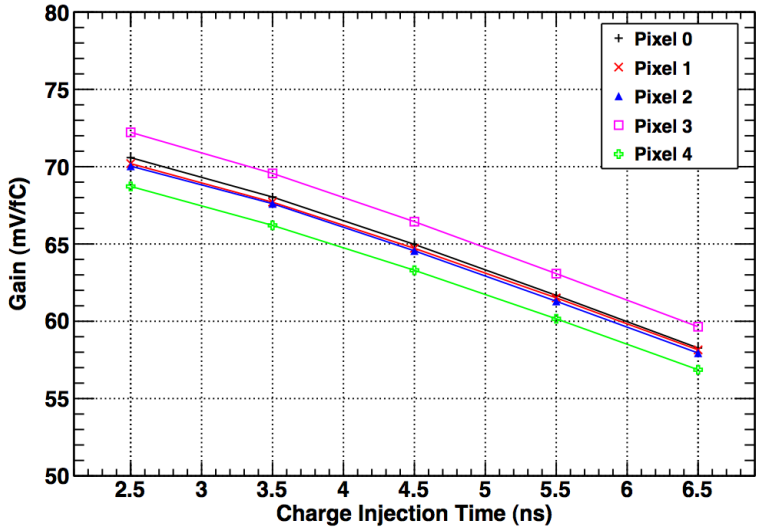

Figure 82 Pixel gain versus charge injection time.

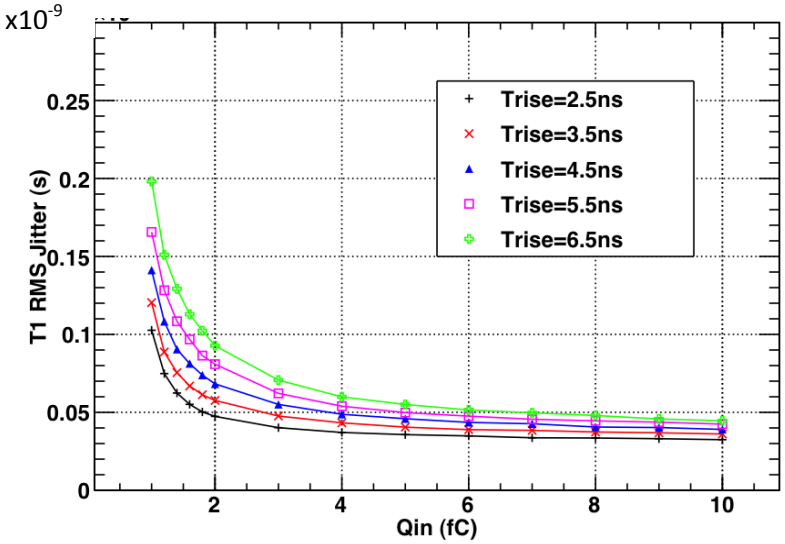

Figure 83 shows t1 jitter over signal charge injection duration.

Figure 83 and Figure 84 show the t_1 jitter and ToT pulse width of the discriminator over input signal charge injection duration. It is assumed that the time walk compensation is calibrated for each of the signal charge injection duration values. As the silicon sensor accumulates total dose over weeks of

operation, the signal shape and duration will change. The ToT scheme allows recalibration to changing sensor signal shapes.

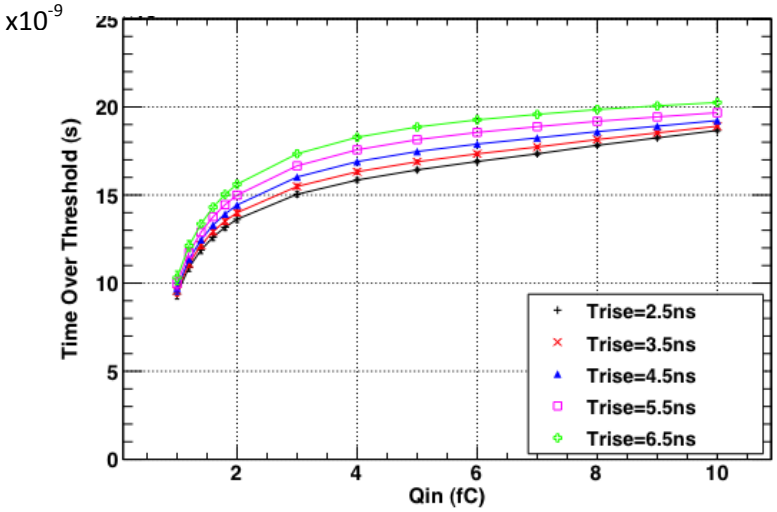

Figure 84 ToT pulse width versus input charge for different charge injection durations.

First measurement of the silicon sensor bonded to the read-out ASIC (silicon assembly)

The silicon assemblies where the sensor is bump bonded to the read-out chip have been produced and delivered in July 2010. Test with radioactive source indicate that the bump bonding process and the charge collection is effective.

Noise analysis with the silicon assembly has been conducted. Figure 85 shows the equivalent noise charge ENC over high voltage for different operation modes for the test pixel with analogue output. It shows that the noise for a high voltage value of 50 V is ~ 180 e-. The values have been obtained by evaluation of the pre-amplifier output signal. Analysis of the s-curve, see Figure 86 and Figure 87, yields a value of ~ 180 e- as well.

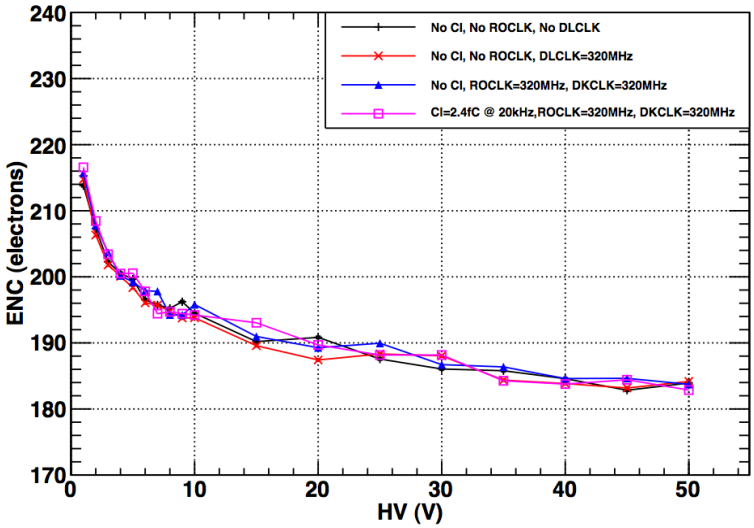

Figure 85 Equivalent noise charge versus high voltage (CI=charge input, ROCLK=read-out clock, DLCLK=digital clock).

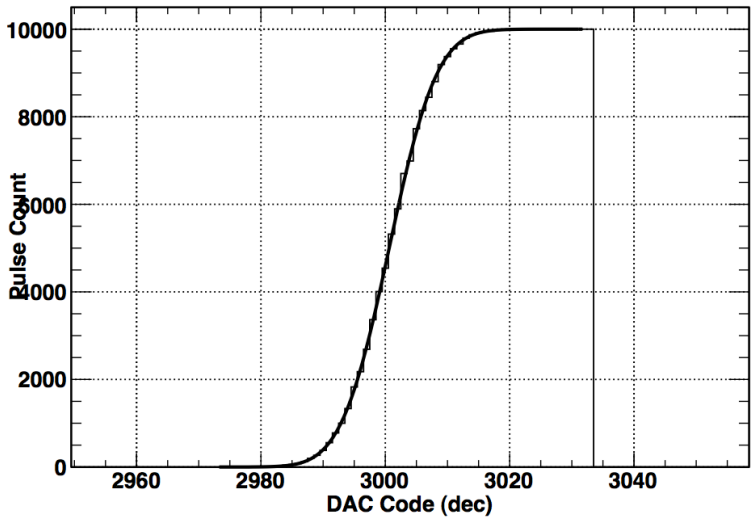

Figure 86 S-curve for the silicon assembly (DAC point = 0.305 mV).

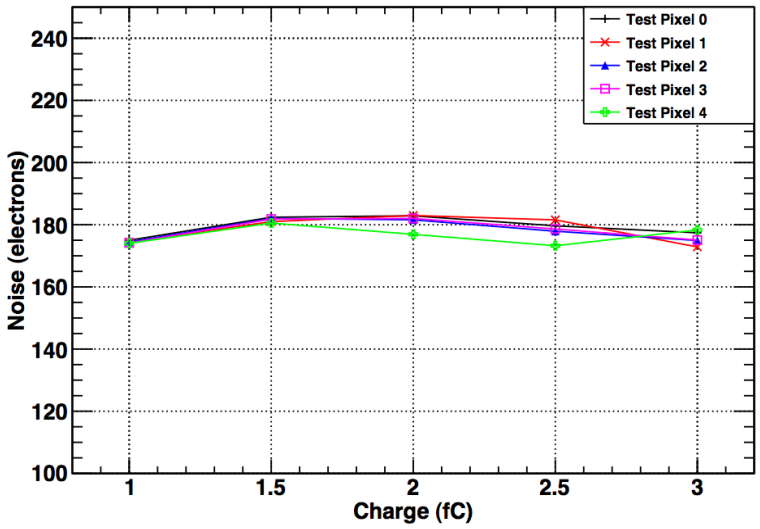

Figure 87 Noise versus electrical charge. HV=50V.

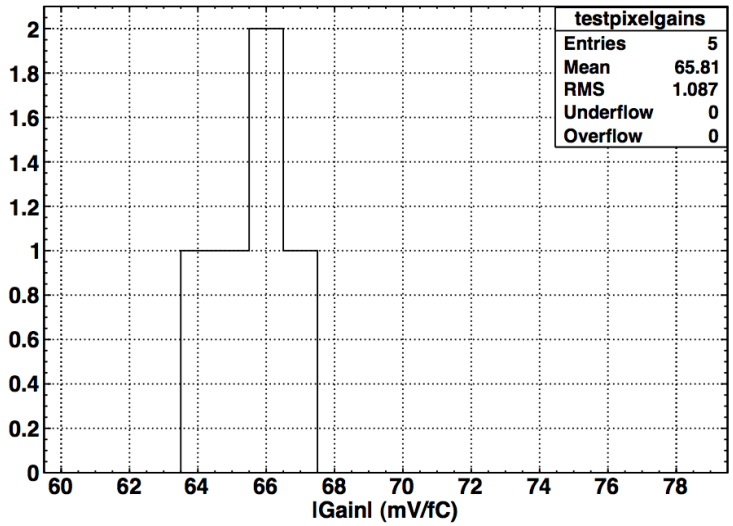

Figure 88 Analogue test pixel gain distribution. HV=50V.

As expected and in coherence with simulation the gain is 66 mV/fC. The reduction of the gain compared to the gain without sensor (70 mV/fC) is due to the additional capacitive load on the input. Figure 88 shows the gain distribution for the test pixels with analogue output.

Trim DAC requirement

In the final implementation the threshold of each pixel cell will be set by on-chip DACs. In-pixel DACs located in each pixel cell are used to compensate the offset in each pixel to equalize all pixels. A global DAC in each column will set the operating threshold for the offset compensated full column. The DACs are also used to scan the threshold during a s-curve scan and measure the turn-on behaviour.

The noise of the analogue front-end has been measured to be 130 electrons. Assuming a worst-case scenario the noise with the sensor is estimated to be below 300 electrons, corresponding to 0.05 fC. Providing a 5 σ separation from that noise estimate the minimum threshold is set to 0.25 fC. The maximum threshold is set to 4 fC in order to permit accurate determination of offset and gain by threshold scanning across known charge values (s-curve). The offset distribution earlier in this document in Figure 80 shows a peak-peak distribution of 50 mV for a 45 pixel column. This and an appropriate contingency is the range, which the in-pixel DAC needs to cover. The bin size of the in-pixel DAC is defined by the need to scan the s-curve over the turn-on region. The width of the turn-on region corresponds to the front-end noise. It is planned to apply a 5 bit in pixel DAC with programmable LSB 1.5 to 4 mV with 8 discrete values. With a minimum binning of 1.5 mV and a conversion factor of 70 mV/fC, the corresponding LSB equivalent value is 0.021 fC.

In order to scan the turn-on region of the s-curve with at least 6 points the LSB of the global 8 bit DAC has to be less than 1 sigma of the noise. For this reason a value of 1.2 mV was chosen and the data supporting this can be seen in Figure 86. The corresponding full range is 300 mV or 4.4 fC.

Table 19 shows the specifications for the DAC in the final chip. The minimum values in the paragraphs above have been adapted to accommodate tolerances.

Table 19 TRIM DAC specifications

	Global DAC	**In-pixel DAC**
# of bits	8	5
LSB	1.2 mV / 0.017 fC	1.5 - 4 mV / 0.021 - 0.057 fC
Min. value	0	0
Max. value	307 mV / 4.3 fC	45 - 128 mV / 0.064 – 0.183 fC

2.3.5 Gigatracker Test Results

2.3.6 Laser test

The EOC sensor/ASIC assembly was tested in a dedicated laser setup. The laser setup consists of a 1064 nm laser head mounted on an automated xy-stage with a focusing point of 7 µm (Gaussian sigma). The electronic test setup was identical to the setup used for the electrical charge injection tests described earlier with the exception that the test pulse trigger signal was sent to the laser, allowing full timing qualification. The laser energy injected into the sensor was measured by means of an optical splitter.

Using ^{109}Cd and ^{241}Am sources the laser setup was calibrated and a conversion table between laser energy and injected charge was built up. Figure 89 shows the energy pulse-height spectrum obtained with the Americium source. The output of the analog test pixel was histogrammed.

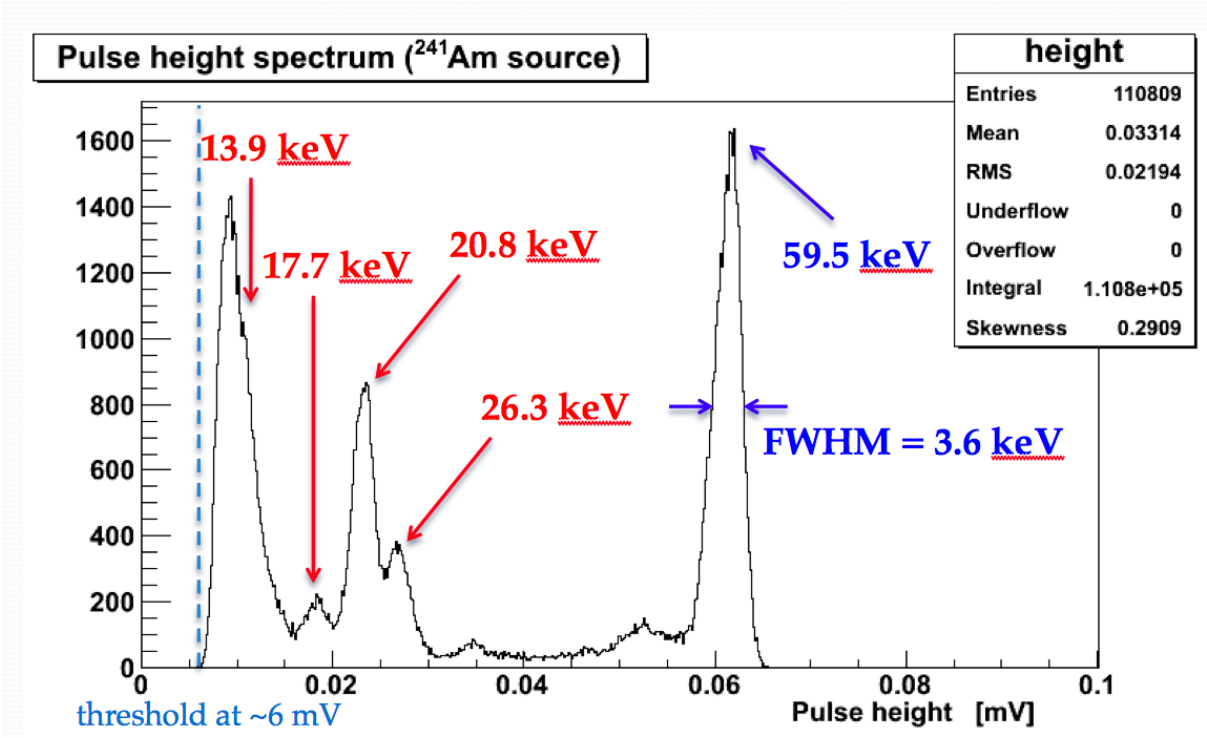

Figure 89 Pulse height spectrum for a ^{241}Americium source.

The laser setup was used to determine the timing behaviour of the full processing chain. Figure 90 shows the t_0 jitter for a pixel at the far end and close end of the column. The time walk compensated jitter t_0 for the most probable charge input of 2.4 fC is around 70 ps. For this measurement the laser light was centred in the pixel cell. Further measurements will be conducted to evaluate possible degradation when the charge is deposited off centre. The difference of the jitter between far and close pixel is very small. This again is an indication that neither the signal transmission of some 14 mm to the EOC logic nor the close proximity of the digital logic in the EOC contributes significantly to the jitter.

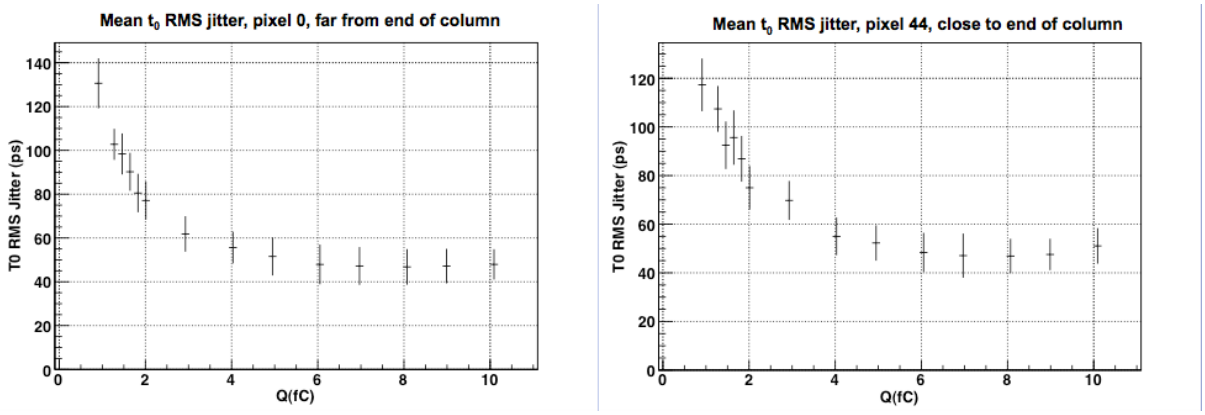

Figure 90 t_0 rms jitter for pixel 0 (far from EOC, left) and pixel 44 (close to EOC, right).

2.3.7 Gigatracker test-beam

The demonstrator bump-bonded assemblies were characterized in a dedicated test-beam. The setup of the experiment run in T9 (PS East Hall) is shown in Figure 91: a low rate 10 GeV/c hadron beam (mainly π+ and p) traversed consecutive GTK assemblies (four EOC and four p-TDC) installed on precisely aligned mechanical supports and scintillators with ~40 ps time resolution as time reference.

The assemblies were continuously operated at 300 V bias voltage for three weeks and both the assemblies and the DAQ showed very stable performance.

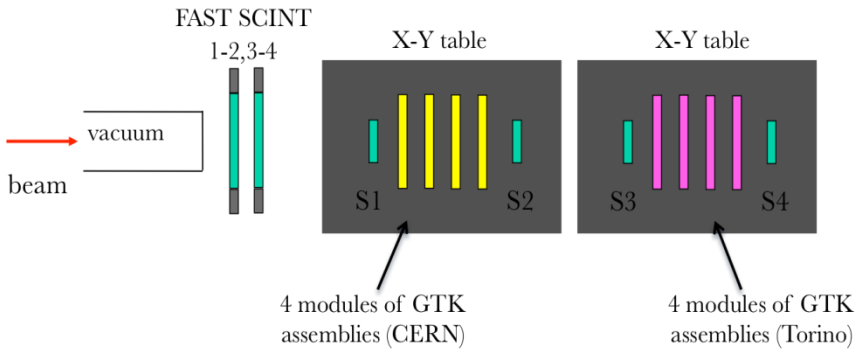

Figure 91 Schematic layout of the test beam setup in the T9 beam line.

The test beam data analysis was at the writing of this document ongoing. A first crude analysis with several data cuts applied is described below.

The time-walk correction is applied using the fast scintillators data, which provide the reference time. The initial simple analysis corrects the time walk but cuts the tails of the Time over Threshold distribution and rejects consecutive hits from pixels in a given 5-pixel group, as the absolute delay of these hits is different. Due to these cuts more than 20 % of the data was rejected. The cuts were applied to simplify and speed up the analysis; the final analysis will take all hits into account and applies a correction for the different delays.

After these cuts and the time-walk correction, the distribution of time differences between hits in the same pixel of two consecutive GTK stations is shown in Figure 92. The width of the distribution is ~250 ps rms, which corresponds to ~175 ps for a single GTK station. The analysed data on the EOC architecture indicate also a very low level of noise, as no pixel hits were recorded out-of-spill period.

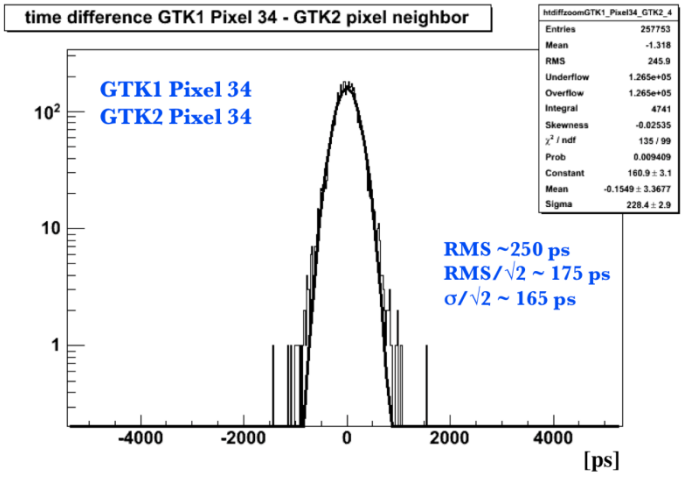

Figure 92 Plot of time correlation between hits in two consecutive GTK stations (EOC architecture).

2.3.8 Cooling

The cooling design is driven by the following GTK requirements:

- Material budget:
  - radiation length in active beam area ≤ 0.5% of X_0,
  - minimum amount of material in the sensor area (60x27mm) plus a 10 mm safety zone,
  - outside that zone no material budget constraints.
- Detector operating temperature 5°C or lower;
- Best possible uniformity of the temperature across the sensor area;
- Power by one GTK station 32 W;
- System operating in vacuum.

Due to the low mass of the detector module thermal run-away and destruction needs to be avoided. Furthermore the cooling system for the Gigatracker reduces the radiation damage of the sensor and consequently increases the module life time. Although the upper limit of the operation temperature has been set to 5°C, a lower operation temperature is highly desirable, as the module life time would increase from 50 days to 100 days when cooled at -20 °C. The material budget and the operation in vacuum are main integration challenges of the project. Mechanics and cooling integration can profit from the fact that no material budget restrictions exist for the area 10 mm outside the beam. From a cooling performance point of view the anticipated 2 W/cm^2 power dissipation by the active electronics are not difficult to meet. However, the cooling and mechanics support architecture must introduce only a minimum of material into the beam area and at the same time assure mechanical stability of the module and allow access to high speed electrical connections. At this time of the project the working group is pursuing two independent options: 1) a micro-channel structure; 2) a gas cooling vessel with mylar walls. **Error! Reference source not found.**

2.3.8.1 Micro-Channel Cooling

2.3.8.1.1 Motivation and Goals

One viable option to obtain a local cooling system satisfying the stringent requirements mentioned is the development of a thin silicon micro-channel cooling plate circulating liquid C_6F_{14}. The development aims for an operating temperature of -20°C to -30°C. Figure 93 shows a schematic of the hydraulic layout of the micro channels: the coolant will enter and exit the straight channels via manifolds positioned on top and bottom. The channels, distribution manifold and openings for the inlet and outlet connectors are etched into a silicon wafer, which is then coupled to a second wafer closing the hydraulic circuit. The final goal is to have both wafers in silicon bonded together by fusion bonding to produce a monolithic cooling element [36].

Recent results obtained in two different fields of development provide guidelines to solve the issues in engineering this cooling device, but dedicated R&D is nevertheless unavoidable for the specific application under study. On the one hand, micro-channel cooling devices have started to be actively studied for future applications for high power computing chips or 3D architectures [37], [38]. For these applications, where the power densities are extreme, the mass of the device (hence its material budget) is an irrelevant parameter. On the other hand, thin and light micro-fluidic devices in silicon are in development for bio-chemical applications [39], but the typical values of the flow rate

and pressure are much lower. Furthermore the presence of a low temperature fluid and of a high radiation level [40] is unique to the application in a HEP detector.

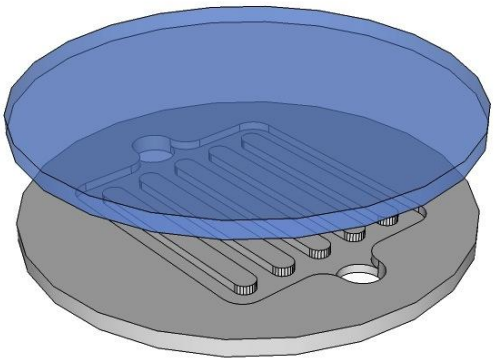

Figure 93 - Schematic of the micro-channel cooling device.

The procedure followed to tackle the different challenges and to converge in a limited time to a single device satisfying all the requirements is to move in parallel along different lines of R&D in a "matrix" approach, where the intermediate results of one line are used to steer the parallel developments. In this way, the study of the optimal fabrication technique provides information on which geometrical layouts are reliably attainable and which process is best suited for a particular layout. Analytical calculations and numerical CFD simulations are performed in order to explore a number of possible layouts and define optimal targets from the point of view of the thermo-hydraulic performance. Numerical structural simulations set the pressure limits to be considered for each channel dimension and wafer thickness. Experimental tests are conducted to validate the numerical model and to allow for a safe extrapolation of the numerical forecast to new configurations.

The maximum temperature difference of the coolant between the inlet and outlet has been set to $\Delta T=5$ K to guarantee a relatively uniform temperature distribution and minimize all thermal stress and related deformations. Also, to limit the mechanical stress in the micro channel cooling plate, the maximum acceptable pressure drop inside the cooling plate has been set to $\Delta p=5$ bar. Based on the above criteria the micro-channel dimensions could be as follows:

- 90 μm deep channels etched into a silicon wafer of 120 μm thickness
- Covered by a fusion-bonded silicon wafer with 30 μm thickness over the active area.
- Outside this area, the thickness of both silicon wafers is 520 μm, allowing for 400 μm deep outlet and inlet manifolds as well as good structural stabilization of the whole device

This design leads to an additional material budget 0.16% of X_0.

An alternative design, in case of technical difficulties with the fusion bonding process, relies on a flat Pyrex cover 50 μm thick anodic-bonded to the silicon wafer carrying the hydraulic circuit [41]. On top of this flat plate, an additional silicon frame (surrounding the beam area) will again be anodic-bonded. In this way the global structure of the cooling wafer will be symmetric, the effects of differences in the coefficient of thermal expansion (CTE) between silicon and Pyrex will be minimized and the same resistance to pressure and manipulation as in the baseline case will be attained.

For technological reasons and for optical access to the channels in the test and development phase, the first prototypes consist at the moment of a silicon wafer housing the channels and a closing Pyrex wafer joined to it by anodic bonding. In the following sections, details about the status of each R&D line will be given and the results leading to the proposed design will be discussed.

2.3.8.1.2 Device Fabrication

The fabrication of the cooling devices is performed at EPFL CMI[20] (Center of MicroNanoTechnology) in collaboration with the Microsystems Laboratory[21]. The micro-fabrication process is shown in Figure 94. It starts with a Czochralski silicon wafer polished on both sides (4'' diameter, 380 µm thick, 0.1-0.5 ohm-cm p-type). A layer of 1 µm of oxide (SiO_2) is grown on both sides of the wafer (Figure 94 (a)). Clariant AZ-1512HS photoresist is spin coated on one side of the wafer at 2000 rpm and lithography is performed to obtain an image of the channels in the photoresist (Figure 94 (b)). The chrome masks for the lithography used in this process are fabricated at CMI using the Heidelberg DWL200 LASER lithography system. The patterned photoresist is baked at 115°C for 50 s on a hotplate. Dry etching of the top layer oxide is used to transfer the micro-channels pattern (Figure 94 (c)). This structured oxide layer will subsequently be used as a mask to transfer the pattern of the channels in the Si wafer by deep reactive ion etching (DRIE). A second lithography is performed with frontside alignment to image two fluid transfer holes, 1.4 mm diameter, for fluid injection and collection from the two manifolds (Figure 94 (d)). DRIE is used to etch the access holes partially down to 280 µm (Figure 94 (e)). The photoresist is stripped in Microposit Remover 1165 at 70°C (Figure 94 (f)) and DRIE is used to etch anisotropical 100 µm deep channels separated by 25 µm wide structures in silicon (Figure 94 (g)). Subsequently the oxide layers are removed by wet etching in BHF 7:1 for 20 min at 20°C (Figure 94 (h)).

At present, the processed Si wafer and an unprocessed Pyrex wafer (4" diameter and 525 µm thick) are then cleaned in a Piranha bath (H_2SO_4 + H_2O_2) at 100°C and anodic bonding is performed to close the channels with the Pyrex wafer (Figure 94 (i)). The bonding is performed at ambient pressure and the temperature is raised to 350°C before being lowered to 320°C. At this stage a constant voltage of 800 V is applied between the Si and Pyrex wafer (Figure 95). The current is monitored and the bonding process is stopped when it reaches 10% of the initial value, typically in the order of 1 mA.

In the final production both the processed and the unprocessed wafers will be of 525 µm thick silicon. At this stage of the fabrication the two wafers will be dispatched to a specialized fusion-bonding laboratory after the execution of the specific surface preparation required. The bonded wafer will then be received back at CMI for the final processing. The related preliminary discussions with two specialized laboratories at present being pursued.

The basic fabrication process-flow described above is used to produce prototype devices to validate CFD thermo-fluid dynamics simulations and ANSYS structural calculations. The results provided with such validated numerical simulations allow predicting the behaviour of thinner devices with

[20] http://cmi.epfl.ch/

[21] http://lmis4.epfl.ch/

improved design to be predicted, thereby reducing the number of tests and the types of samples to be produced. A more complex micro-fabrication process is at present in preparation. This includes an additional step of photolithography to allow for different heights of manifolds and a final local etching to be applied to obtain a thinner region in the beam acceptance area.

Following any additional final processing (e.g. local thinning over the minimum material area), the bonded wafer is then diced according to alignment marks previously etched in Si to obtain a cooling plate with precise external references for integration into the electromechanical assembly.

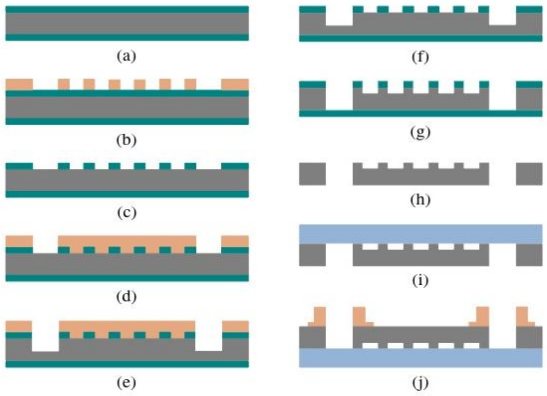

Figure 94 - Fabrication process-flow.

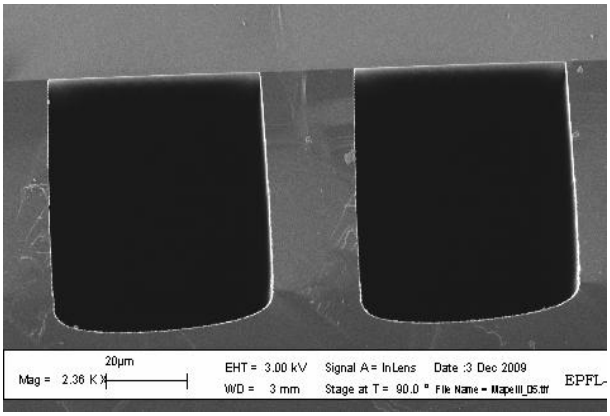

Figure 95 - Scanning Electron Microscope image of the cross-section of micro-channels etched in silicon bonded to a Pyrex wafer.

Finally, PEEK connectors (NanoPort assemblies from Upchurch Scientific) are aligned, with a gasket and with a preformed adhesive ring to the inlet and outlet on the silicon before clamping. They undergo a thermal treatment at 180°C for 2 hours to develop a complete bond between the NanoPort and the silicon substrate. The supplier guarantees correctly attached NanoPort connectors to withstand pressures of 34.5 bars. However we tested them repeatedly up to pressures of 70 bars before failure.

The choice of the cooling fluid circulating in the micro-channels has naturally been oriented towards perfluorocarbon fluids (C_nF_{2n+2}), which are widely used as coolant medium in LHC detectors. In particular C_3F_8 and C_6F_{14} are used in the inner tracking detectors of ATLAS and CMS. They exhibit interesting properties for cooling applications in high radiation environment such as thermal and chemical stability, non-flammability and good dielectric behaviour. Radiation resistance studies have been performed [42] by irradiating samples of C_6F_{14} with gammas from a ^{60}Co source up to an accumulated dose of about 5.6×10^4 Gy. At the location of the GTK stations, accumulated doses in the order of 6×10^4 Gy are expected over one year of operation. At such doses the sample C_6F_{14} “Flutec PP1” (iso-perfluorohexane from F2 CHEMICALS) shows the best radiation hardness and chemical stability. It has thus been selected as the cooling fluid for the micro-channel devices.

2.3.8.1.3 Thermo-Fluid Dynamics Design, Simulation and Validation of Micro-Channels

The relevant properties of C6F14 for thermo-fluid dynamics calculations are summarized in Table 20.

Table 20 C_6F_{14} properties.

Properties	**C_6F_{14} @ -25°C**
Density r [kg/m^3]	1805
Viscosity n [10^{-7} m^2/s]	8.2
Heat capacity c_p [J/(kg K)]	975
Thermal conductivity l [10^{-2} W/(m K)]	6.275

The required C_6F_{14} mass flow to extract the 32 W dissipated by the readout chips with the assumed temperature difference of 5 K is 7.325 x 10^{-3} kg/s. As a consequence, the Hagen-Poiseuille-law applies for the analytical pressure drop calculations, as the flow inside the channels is laminar in all configurations considered (Reynolds-Number between 100 and 300).

In all calculations the heat source is considered uniformly distributed over a surface of 60 mm x 40 mm, corresponding to the area of the readout chips. The hydrodynamic and thermodynamic length of the channels are considered coincident, L = l = 40 mm. The wall thickness between the channels has been set to half the channel width, to withstand mechanical stress due to the pressure in the channels and to provide enough bonding surface.

Analytical design

Analytical calculations were performed to obtain an initial estimate of the feasible values for the micro channel dimensions. For a given pressure drop, Figure 96 plots on the vertical axis the temperature difference between inlet and outlet as a function of the channel geometry (height and width). The continuous surface represents the ΔT in the coolant and the mesh surface the ΔT in the channel wall. It appears necessary to adjust both channel dimensions in order to keep a temperature difference of 5K. The reason is that in channels with a rectangular cross section, the smaller of the two dimensions dominates the pressure drop equation with a cubic power dependence [43]. For a fixed pressure drop, the ratio between height and width should typically be kept between 0.5 and 2. It is also visible from the plot that the temperature difference between the fluid and the wall, represented by the vertical separation of the two surfaces, rises with an increase of the channel size due to a higher line load per channel and a decrease in heat transfer surface. This has to be taken into account when selecting the channel cross section, as it will have an impact on the gap between the operating temperature of the coolant and the minimal temperature attained on the surface of the sensor.

It has to be considered that the channel height translates into a requirement for the thickness of the silicon wafer, where the channel will be etched. Therefore, channel height values larger than 120 μm cannot be considered for the present application. Figure 97 relates the reachable mass flow to the channel width, for a fixed pressure difference of 2 bar and a fixed channel height of 90 μm. One can see that the required nominal mass flow is reached for a channel width of approximately 120 μm. It

is also visible that the gain in mass flow is favourable up to an increase in width to around 150 µm. For wider channels the width-to-height ratio becomes too large and the channel height starts to be dominant for the mass flow.

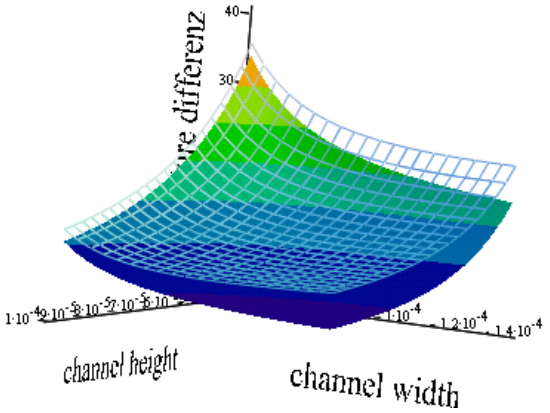

Figure 96 - Temperature differences between inlet and outlet in the coolant and in the channel wall.

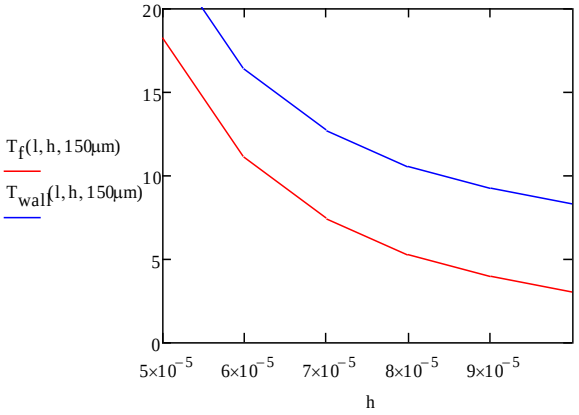

Figure 97 - Reachable mass flow related to the channel width for a fixed channel height.

The results from the analytical calculations performed indicate that the suitable ranges of the micro channel dimensions are the following:

- Width: between 100 µm and 150 µm
- Height: between 80 µm and 120 µm
- Fin width: between 25 µm and 75 µm
- Between 300 and 500 channels to cover the area

CFD simulations

Two manifolds will feed and drain the micro channels. The manifold design must aim at a small total pressure drop and a uniform distribution of the coolant over the channels. In particular, the total pressure drop gives the necessary pressure at the inlet of the cooling plate and is therefore an important factor for the mechanical stability of the design. A good hydraulic design of the manifolds requires the use of full 3D Computational Fluid Dynamic (CFD) simulations.

Different configuration and sizes of manifolds have been simulated in connection with channel geometries in the range identified above. The calculated total pressure drops span from 15 to 2 bar, mainly depending on the cross sectional area of the feeding manifold. The details of the best performing solution provided by the simulations are shown in the box of Figure 98: it is a wedge-shaped manifold, providing the advantage of a more uniform flow distribution inside the channels with respect to a simple straight geometry [44].

Figure 98 shows the pressure map for the following specifications and boundary conditions:

- Channel cross section 100 µm x 100 µm
- 2 NanoPort connectors placed on opposite sides of the inlet and outlet manifolds

- Wedged manifold with maximum width 1.6 mm on the inlet and 2.1 mm on the outlet
- Manifold depth 400 µm
- Inlet: mass flow 7.325 x 10-3 kg/s (ΔT 5 K)
- Outlet: pressure 1 bar

The pressure drop distribution across the channels is very uniform, thus indicating a uniform flow distribution of coolant across the device surface. The total pressure drop is 3.1 bar.

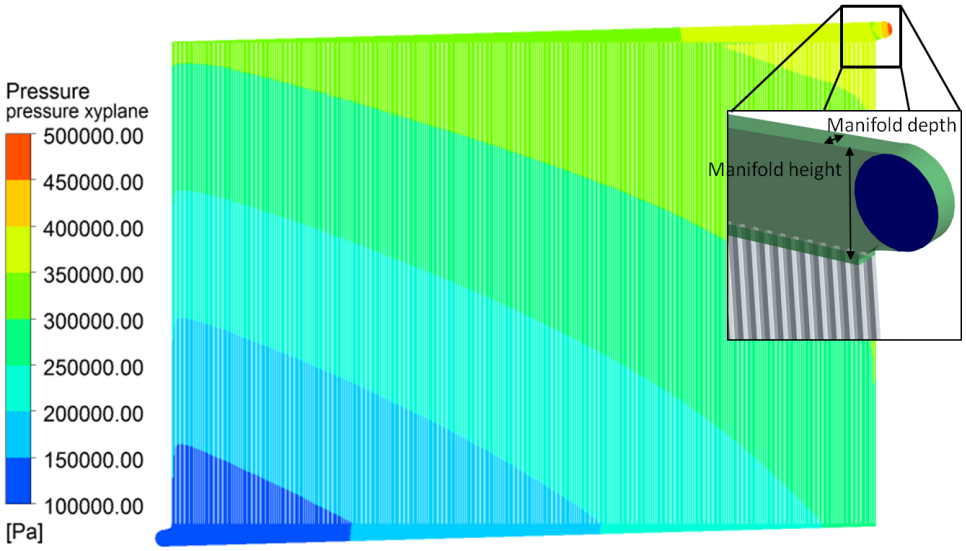

Figure 98 - Pressure [Pa] in micro channels and manifold with inlet and outlet connector on opposite sides.

Validation on full-scale prototypes

The analytical and computational simulations have been validated on full scale prototypes. A slightly simplified design has been used for these tests. It consists of an array of 480 channels, uniformly distributed over a 60 mm wide area. The channels have a cross section of 100 µm × 100 µm, are 48 mm long and are separated by 25 µm wide silicon walls. Two 60 × 1 mm rectangular distribution manifolds, 100 µm deep, connect all the micro-channels together at both ends. Injection and collection of fluids is performed through PEEK capillaries screwed into a NanoPort connector placed at the centre of each manifold. This geometry is not optimized for the overall pressure drop, but allows for a faster and less expensive production.

Figure 99 compares the pressure drop between inlet and outlet of the CFD simulation and the wafer in the test stand. One can see that the prediction of the simulation is very close to the experimental data. This validates the numerical model developed for the CFD simulations and lends reliability to the results obtained with this model for the detailed design of the micro channels and manifolds, thus largely reducing the number of samples to be produced for testing purposes.

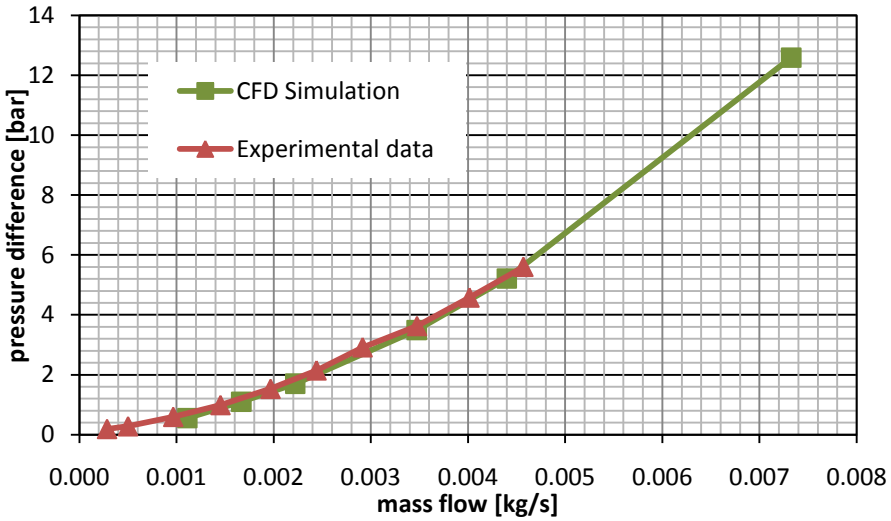

Figure 99 - Pressure drop comparison between simulation and tests.

Figure 100 shows a sample of full scale prototype wafer with a heating foil simulating the distributed thermal load glued on one side (left) and two NanoPort connectors bonded on the opposite side (middle). A typical mounting on the test stand for room temperature tests is also presented (right).

Figure 101 shows a cool down test of the wafer with the simplified geometry. The pictures are taken with a thermal camera. The light yellow rectangle on the left picture is the powered dummy load with approximately the same surface area as the read-out chips. The picture in the middle shows the moment when the cooling liquid enters the manifold and channels. The picture on the right shows the rapid temperature drop on the surface of the wafer and the dummy heater few seconds after the coolant circulation is established. A preference of the cooling liquid for the middle channels is clearly visible. The improved wedge-shaped manifold design will correct this effect.

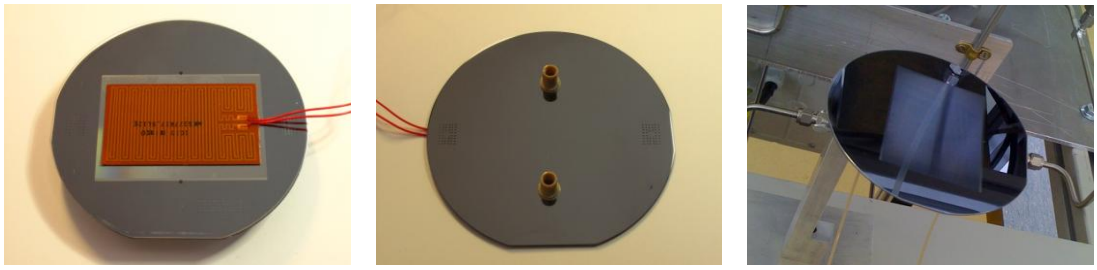

Figure 100 - Cool down test.

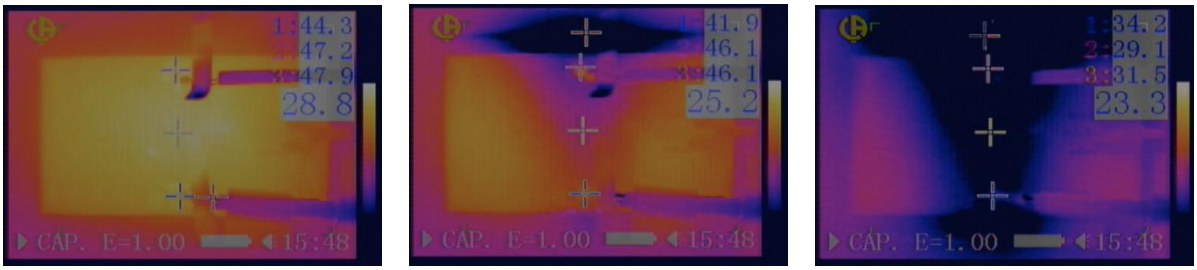

Figure 101 - Cool down test.

2.3.8.1.4 Pressure Tests and Structural Simulation

The maximum pressure safely attainable inside the cooling device without risk of collapse is a crucial parameter, as this determines the minimum amount of material needed for a given flow rate, to attain a given temperature uniformity on the detector surface. Accordingly, an experimental determination of this structural limit would require the production of an extremely large number of samples with different wall thickness. The problem has been approached in a similar way as for the analysis of the thermo-fluid dynamics properties: a numerical model has been developed and one set of samples has been produced for testing purposes. The comparison between the results provided by the simulation and the test performed allows the numerical model to be refined and validated; this can then be used to extrapolate the results obtained to different configurations.

Simplified experimental samples characterized by a 60 mm long rectangular manifold of different width w = 0.2, 0.25, 0.5, 1 and 2 mm have been produced by anodic bonding of a 380 μm thick silicon wafer and a 520 μm thick Pyrex wafer. The samples can be connected to a controlled pressure device through a NanoPort connector similar to those used for the complete cooling device. The pressure in the sample is raised until failure occurs and the limiting pressure is then determined from the recorded data while images from a high-speed camera help understanding the failure mechanism.

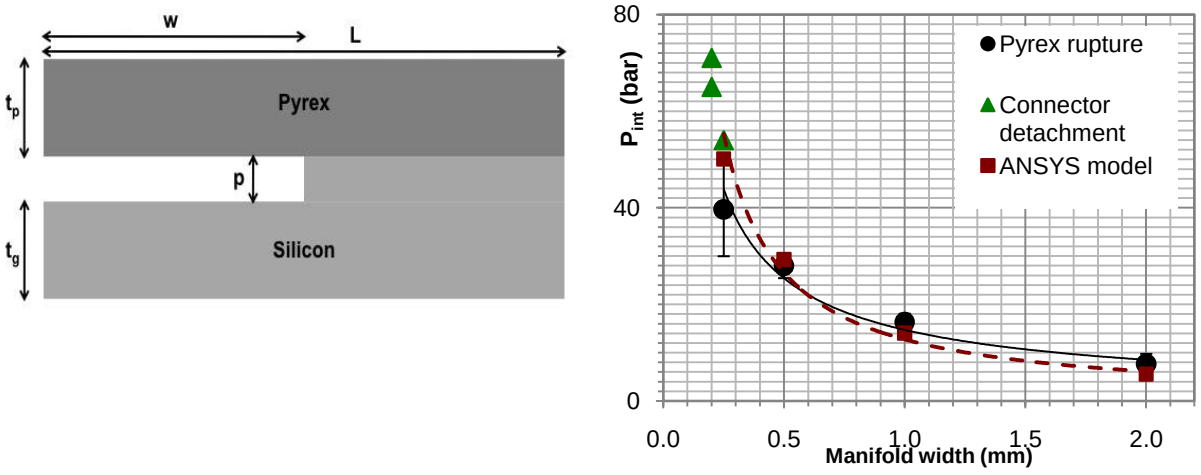

Figure 102 Left side: Geometry of the sample used in the simulations. Right side: Rupture pressure [bar] versus manifold thickness. Comparison between experimental results and simulations.

A 2D ANSYS parametric model simulating the geometry under test has been developed as shown in Figure 102: one layer of Pyrex (thickness T_p) bonded onto one layer of silicon (thickness T_g) in which a channel is engraved (width W and depth p). The total width of the sample is L. A symmetry boundary condition has been applied on the left edge and the right and bottom edges are fixed. A constant pressure P_{int} is applied in the channel.

For a defined Pyrex thickness t_p, the tensile strength was calculated in the Pyrex layer for different channel width W as a function of the applied internal pressure P_{int}. Delaminating of the bonded

interface was not taken into account in this study: indeed, as found in the literature, the Pyrex layer is identified as the weakest point in our sample. The maximum yield strength of Pyrex is 25 MPa[22] (165 MPa for Silicon). This value is then considered as the maximum stress admissible in the Pyrex layer before breakage.

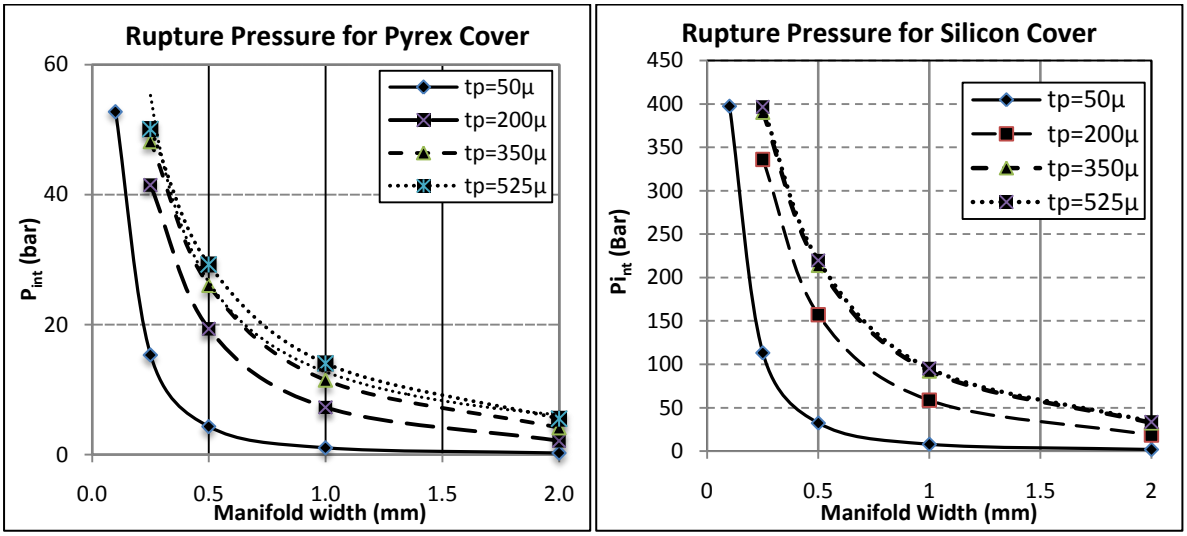

Figure 103 Simulation of the maximum sustainable pressure as functions of manifold width and thickness (left: for Pyrex; right for Silicon).

Figure 103 shows the pressure test results together with the calculated pressure for which the global stress (von Mises formulation) in the Pyrex layer is higher than 25 MPa for a Pyrex thickness of 525 µm and different values of W. The apparent good correlation between tests and simulation validates the ANSYS model and allows is to be used to forecast the behaviour of a Pyrex layer of different thickness. The slight divergence between calculated and measured data indicates a slight over-estimation for smaller values of W and under-estimation for larger values of W.

The simulated behaviour of various Pyrex thicknesses allows the maximum attainable pressure to be determined depending on the channel (or manifold) width and the Pyrex thickness desired.

Figure 102shows for the 100 µm wide channels in the beam area of the device, that even if one has a Pyrex thickness of 50 µm, the P_{max} allowed gives a large safety factor (more than 10) compared to the 3 to 4 bar pressure expected in this region. For the larger dimensions of the order of 1000 µm on the manifolds, the simulations indicate that if we want to keep a good safety factor and be able to withstand an absolute pressure of 5 bar to 7 bar, one should go to a thicker Pyrex layer. A 525 µm Pyrex thickness could hold a least 14 bar, thus providing a safety margin of 2 to 3 on the maximum allowed pressure. The manifolds must therefore be placed outside of the beam area, where a increase of the wafer thickness is allowed.

[22] As found in CES material database, www.grantadesign.com/products/ces.

Similar calculations are shown on the right hand plot of Figure 103 for a full silicon solution, for which the limiting value of the global stress considered is 165 MPa. The extremely high pressure values attained suggest that local de-bonding phenomena should not be entirely neglected in the calculations for this case. However, the higher mechanical properties of silicon clearly allow for a further reduction of the wall thickness to the desired 30 μm level. Should manipulation or production issues prevent this further reduction in wall thickness, the full silicon solution would safely allow for higher pressure levels corresponding to reduced channel dimensions. Thus preserving the target material budget might be preserved.

2.3.8.2 Gas Cooling

2.3.8.2.1 Configuration

The other cooling option is the construction of a vessel housing the GTK module and cooling it via a flow of cold gaseous nitrogen. The nitrogen will enter the vessel at a temperature of 100 K. The cold flow will keep the operation temperature of the module less than 5°C. In this design the GTK is installed between two cylindrical Kapton walls, 40 m thick, supported by an aluminum frame, see Figure 104. A second pair of thinner (10 μm) Kapton walls is inside the vessel surrounding the detector and chip assembly, providing a smoother flow over the detector surface. The total thickness crossed by the beam is (2 x 40) + (2 x 10) = 100 μm Kapton. The radiation length for this polymide film $(C_{22}\ H_{10}\ N_2\ O_5)_m$ is 28.6 cm[23] so the total material budget of the vessel is 0.035% X_0.

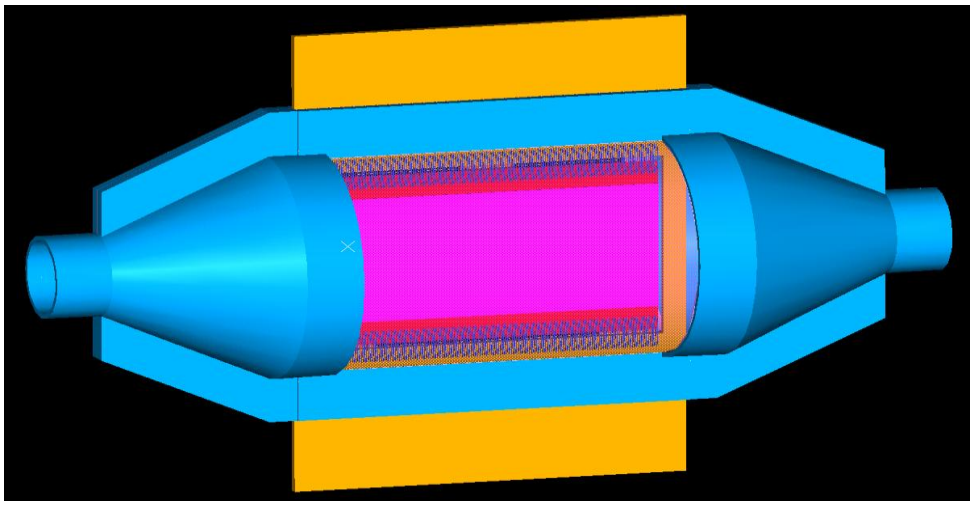

Figure 104 Cooling vessel for the GTK module.

The detector assembly is electrically connected to assembly carrier printed circuit board, which provides power and signal connections. However, the assembly is thermally decoupled from the assembly carrier board; the sliding support compensates for the different thermal expansion coefficients of the materials.

[23] Review of Particle Physics, http://pdg.lbl.gov/2010/AtomicNuclearProperties

The thermal analysis indicates a good uniformity of the temperature distribution across the detector area with a max value of 4 °C corresponding to a convection coefficient of 44 W/m^{2}K, see Figure 105.

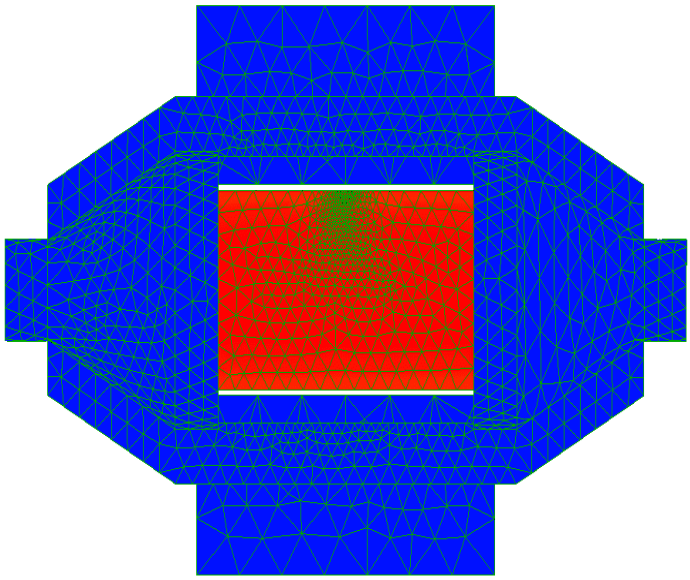

Figure 105 Temperature distribution as from simulation.

A structural analysis has been performed and takes the inner pressure due to the vacuum and the pumping system, the different elongation of the materials and the fixed constraints into account. It gives low mechanical stress values and displacements both for the vessel and the detector (< 2 MPa), see Figure 106.

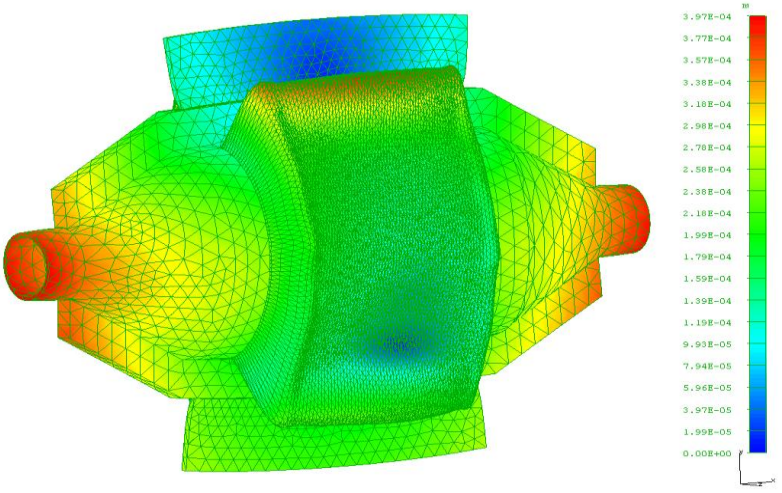

Figure 106 Map of thermal stress.

2.3.8.2.2 Design and Construction of a Full Scale Prototype

Following the results coming from the simulations, a full scale prototype has been designed and built in October 2009. Figure 107 shows the assembled prototype.

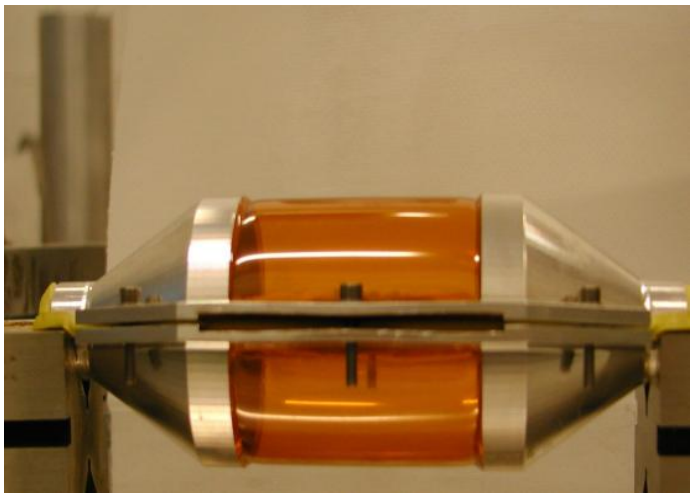

Figure 107 Full scale prototype.

The materials chosen for the construction are listed in Table 21.

Table 21 Material properties.

PART	MATERIAL	DENSITY (Kg/m^3)	THERMAL COND. (W/mK)	CTE (K^{-1})	ELASTIC MOD. (Pa)	POISSON'S RATIO
PCB	GLASS FIBER	2076	0,83	6,7E-6	45E+9	0,28
VESSEL WALL	KAPTON TYPE H	1420	0,16	45E-6	2,5E+9	0,34
VESSEL FRAME	ALUMINUM	2700	237	23,5E-6	70,6E+9	0,345
RESIN	EPOXY (ARALDITE 2012)	1180	0,22	30E-6	2,5E+9	0,34

The prototype was used to compare the results coming from the simulation to a full scale model working at the nominal operation parameters.

The following tests were done:

- structural test
- cooling test.

2.3.8.2.3 The Structural Test

The structural test concerned the pressure inside the vessel and the aluminum – kapton joint .

Kapton failure pressure test

The simulation and the standard calculations give a breaking pressure of 0,49 MPa (4,9 bar) for the 50 µm thickness kapton wall. A test, checking the failure pressure of the vessel, was set up, see Figure 108. The results of the test gave a breaking load of 0,52 MPa (5,2 bar) in a good agreement with the calculation. Considering the pressure inside the vessel of 2 bar given by the vacuum and the pumping system, the safety factor is 2.5.

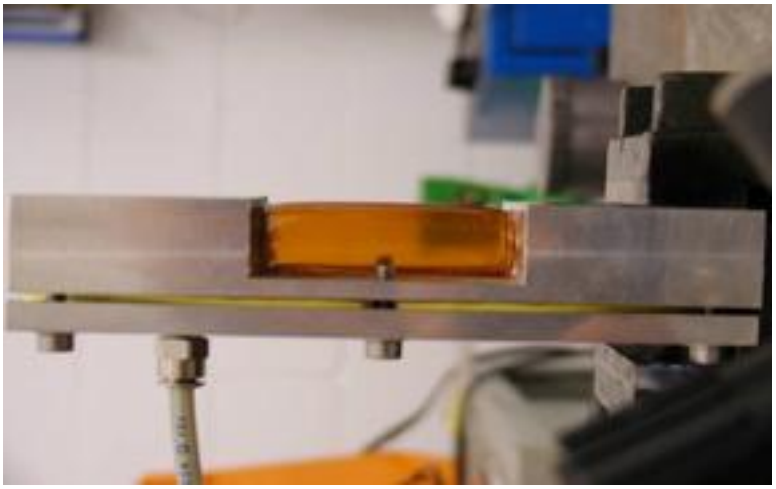

Figure 108 The Kapton failure pressure setup.

Quality of the joint Kapton – resin – aluminum

The overlap of the joint Kapton – aluminum has been designed to withstand a stress 2 times the breaking load of the kapton; the safety factor for the resin is 5. Fig. 6 shows the good behaviour of the joint after the breaking test.

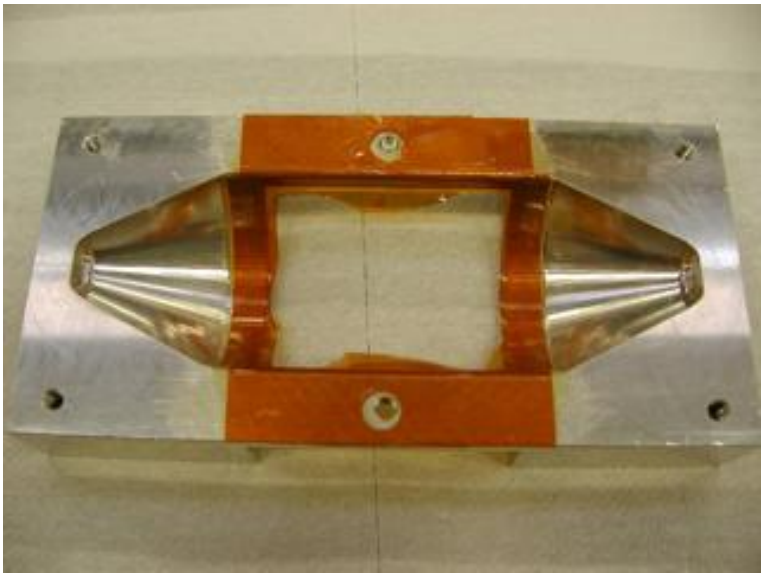

Figure 109 The Kapton - resin - aluminum joint.

2.3.8.2.4 The Cooling Test

Following the good results coming from the structural tests, a cooling system (see Figure 110, Figure 111 and Figure 112) was setup in February 2010.

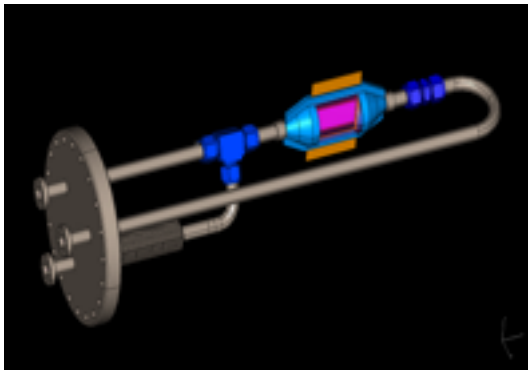

Figure 110 Cooling circuit design.

Figure 111 Cooling circuit implementation.

The system parts are listed below:

- vacuum chamber housing the prototype; pressure inside is 10^{-6} mbar;
- vacuum pumping system : scroll = 5 m^3/h and turbo 70 l/s;
- power and thermocouples feed-through;
- flowmeter;
- nitrogen in/out circuit

The test was limited by the maximum nitrogen flow rate of 7 m^3/h allowed by the pumping system.

In the test the detector was simulated, see Figure 112, by 16 resistors 222 Ω - 2 W with a total power of 32 W. The temperature was measured by 3 K-type thermo couplers mounted on the detector. A data monitor recorded the temperatures and the flow rate during the test.

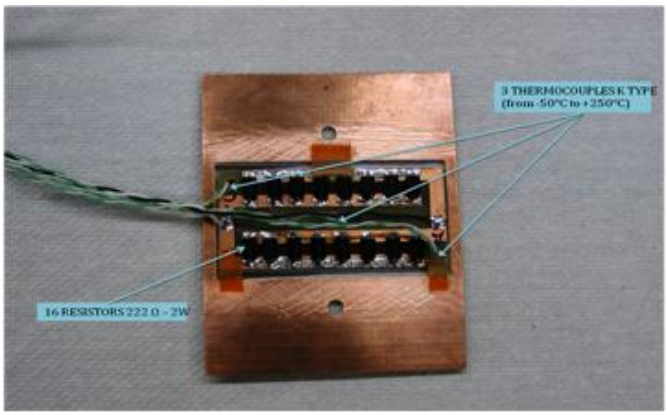

Figure 112 Detector mock up.

2.3.8.2.5 Cooling Results

The cooling results of the test are shown in Table 22, where

- P (W) is the power set on the detector
- Fr (m^3/h) is the nitrogen flow rate
- Tin (°C) is the measured temperature on the detector from the nitrogen inlet
- Tmid (°C) is the measured temperature on the middle of the detector
- Tout (°C) is the measured temperature on the detector from the nitrogen outlet
- Taverage (°C) is the average of the measured temperatures

Table 22 Measured temperature versus power & flow rate.

P (W)	Fr (m^3/h)	Tin (°C)	Tmid (°C)	Tout (°C)	Taverage (°C)
0,14	2,2	13,8	14,2	14,1	14,03
0,56	2,2	19,3	20,3	19,7	19,77
1,25	2,2	24,6	26,4	25,2	25,40
1,25	3,7	16,8	18,8	17,7	17,77
1,25	7,2	-17,2	-16,1	-16,6	-16,6
3,48	3,7	27,4	32,4	29,6	29,80
3,48	7,2	-13,8	-10,0	-11,90	-11,9
13,9	3,7	74,3	90,6	79,2	81,40
13,9	7,2	16,3	19,9	17,7	18

The results show a quite good uniformity of the temperature, with the maximum in the middle zone. The 7 m^3/h flow rate limited the power dissipation at 14 W. After the first test all the parts were checked and there was no evidence of damages due to the temperature excursion (from -17 °C to 80 °C).

In a second test in March 2010 the cross section of the cylindrical vessel was reduced of a factor 3 by adding two inner parallel flat walls at a distance of 12 mm. The amount of material was the same: flat walls thickness 10 µm and round walls thickness 40 µm, with a breaking safety factor 2.

The aim was to increase the thermal exchange at a reduced flow rate, improving the laminar flow and to compare these results with those of the first test.

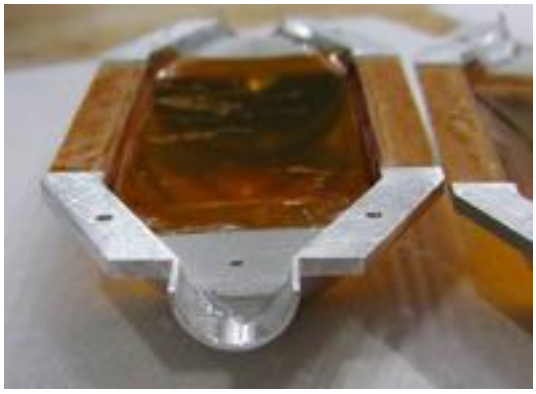

Figure 113 Flat Kapton wall on half of the vessel.

The cooling results of the test are shown in Table 23, where the symbols are the same as described for Table 22.

Table 23 Measured temperature versus power & flow rate

P (W)	Fr (m^3/h)	Tin (°C)	Tmid (°C)	Tout (°C)	Taverage (°C)
3,5	3,3	-	19,2	20,1	19,7
8,9	5,1	-	13,2	14	13,6
20	5,3	-	53,2	55,3	54,2
31,3	7,2	-	22,5	23,5	23,0

Figure 113 shows one flat wall glued on half of the vessel . The pressure between the flat walls and the round walls was balanced by two holes connecting the volumes.

The Tin is not reported because the first thermocoupler was not thermally connected to the detector.

The temperature shows a quite good uniformity and the 7m3/h flow rate allows the operation of the detector at nominal power.

Comparing the results of the two tests the reduced cross section given by the flat walls improve the cooling, increases the flow speed and reduces the needed flow rate.

2.3.9 Electro-Mechanical Integration

Each of the three GTK stations will be installed inside a dedicated vacuum vessel at three different positions along the NA62 experiment beam line. The vessels for GTK1 and GTK2 are identical and are 292 mm x 239 mm x 190 mm in size, whereas GTK3 is 550 mm x 560 mm x 2000 mm as it contains the CHANTI detector planes[24] as well, see Figure 114 shows the vessel design for GTK1 and GTK2.

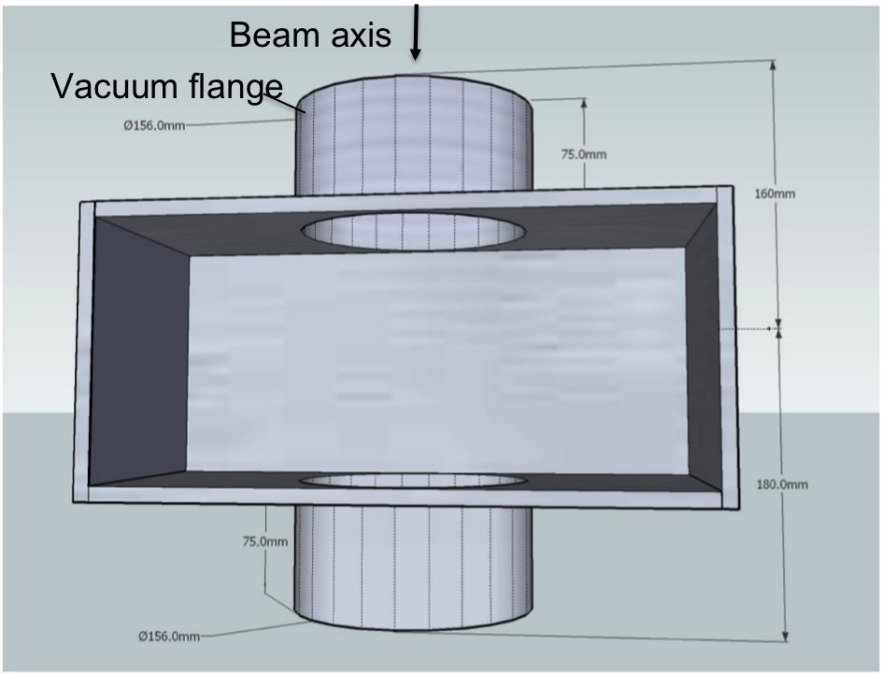

Figure 114 GTK 1 and 2 vessel.

The vessels will be fixed and precisely aligned on the beam axis. The vessel for GTK3 and the CHANTI will be fixed in place whereas GTK1 and 2 vessels are mounted on an xy-table for automated alignment to GTK3 and the beam axis. Connections with the beam pipe are made using bellows allowing the vessels to move.

Due to the radiation damage in the sensor, the modules need to be replaced on a regular basis, see chapter 2.3.3.1. The electro-mechanical integration is based on the principle that one GTK module – the assembly and its services - is an integral and compact component, which can be inserted and removed with a minimum intervention in the vessel. The Gigatracker assembly carrier is a printed circuit board which supports mechanically the assembly and provides all ASIC and sensor power supplies, control and read-out signal connections on differential signals and carries optical

[24] The CHANTI detector layout is described in section 2.4.2.

components outside the vacuum vessel and acts as vacuum feed-through for the electrical lines and the cooling connections. Figure 115 shows a drawing of the GTK assembly carrier for the micro-channel cooling option. On the left hand side the PCB has an opening for the sensor assembly and the micro-channel cooling plate. In the middle a flange with an O-ring is tightly glued around the PCB acting as vacuum feed-through. On the right hand side optical components, already outside the vacuum, are placed.

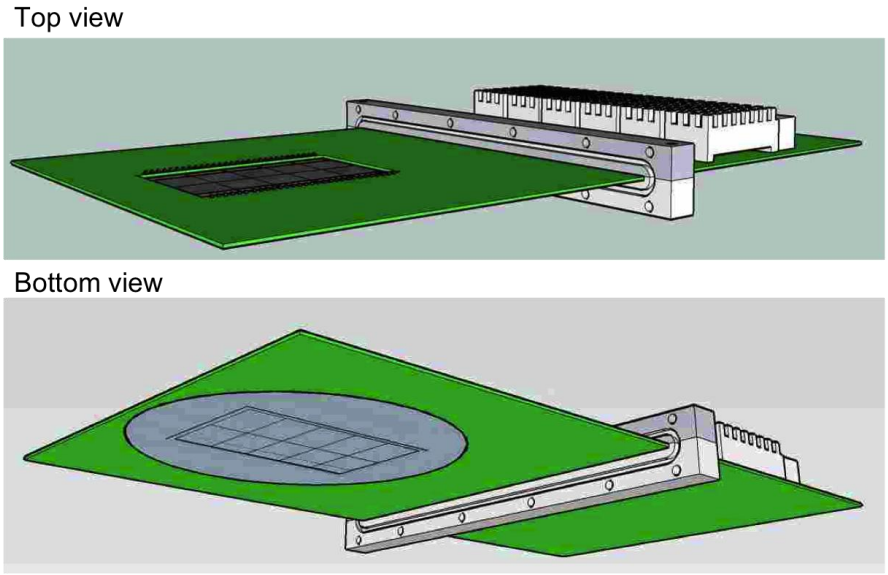

Figure 115 GTK assembly carrier.

Mechanical guides assure precise alignment of the assembly carrier and thus the silicon assembly with the vessel. Figure 116 shows the vacuum feed-though and the positioning pins. Figure 117 shows the module inserted in the vessel.

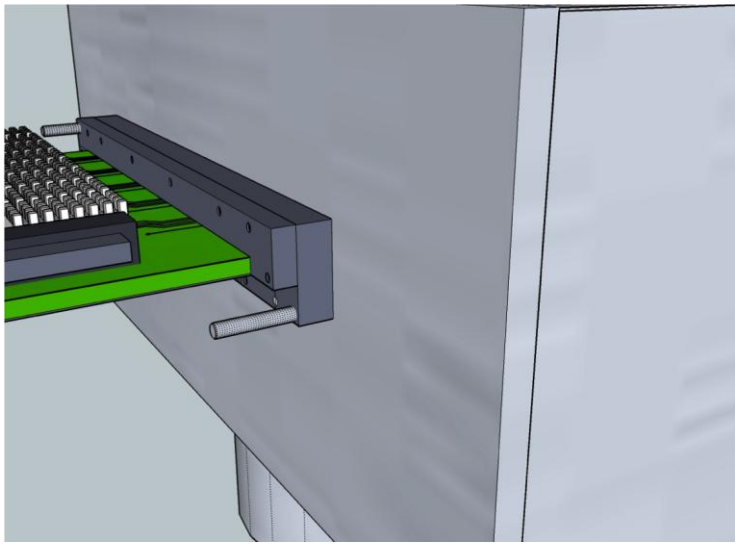

Figure 116 Vaccum feed-through and positioning pins.

Wire bonds provide the electrical connections between the ASIC wire bond pads, the sensor bias connections and the GTK assembly carrier. Each ASIC requires three different power supplies 1.2 V digital, 1.2 V analogue and 2.5 V I/O. In order to minimize the electrical interferences, each ASIC will be powered individually with separate power supplies. Measurements will show whether the analogue and digital supply of an individual chip can be combined on the assembly carrier. For the optical components and clock distribution buffers the card is connected to 3.3 V. This means a total of up to 31 power supply connections for each GTK interface board are required. In order to avoid the need of a precise pitch adapter for the wire bond connection to the printed circuit board, a staggered wire bonding scheme is proposed where the pitch in the PCB does not need to match the ASIC pitch of 73 µm. Figure 118 illustrates the principle where the power connections and the signal connections go to different rows on the PCB.

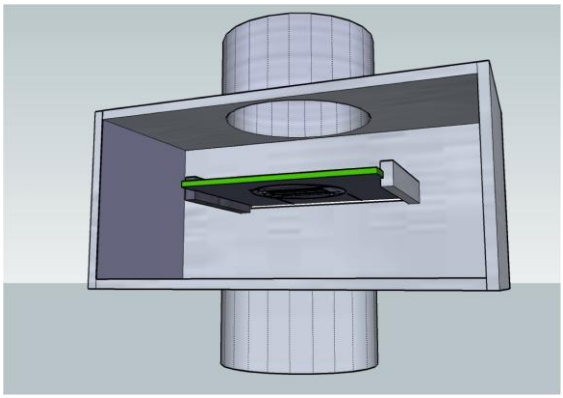

Figure 117 GTK module inserted in the vessel.

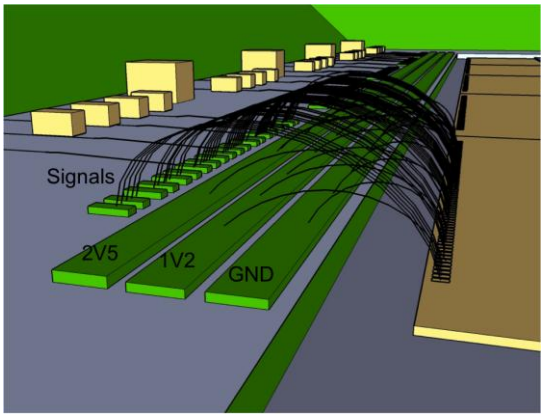

Figure 118 Staggered wire bonding scheme avoids pitch adapter.

Figure 119 shows a cross section of the module. The FR4 acts as rigidifier as well as carrier for the printed circuit. The figure also shows the micro-cooling connection (Nano Port). This design connects the electrical connection on the top side, whereas the cooling is connected on the bottom side avoiding mechanical interference.

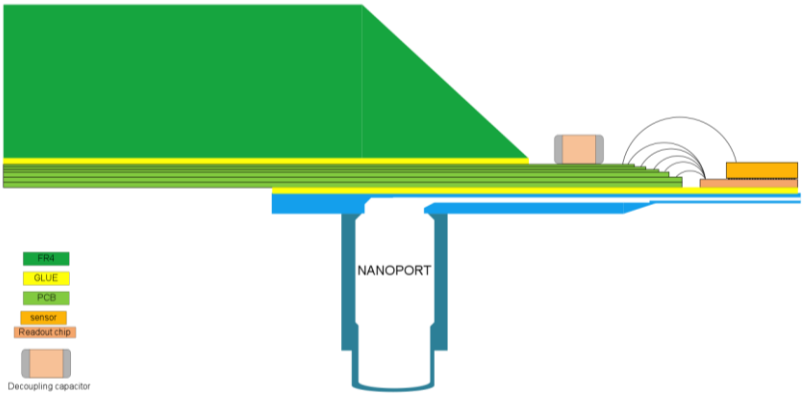

Figure 119 Cross section of GTK module.

Each of the 10 ASICs sends 6Gb/s data on differential lines to the optical components on the assembly carrier. The TDCs inside the ASICs require a low clock jitter signal. The electrical layout of the PCB is driven by these two constraints and requires a study and simulation of signal integrity and

special attention in the routing. Prototype GTK assembly carriers are going to be produced in order to optimize this layout.

Due to high variation of temperature in the cooled operation mode and in the un-cooled stand-by mode, the silicon components are not glued to the GTK assembly carrier, but are placed in a precisely machined opening on the GTK assembly carrier. Carbon fiber clips keep the silicon component in position (see Figure 120) and avoid mechanical stress due to different thermal expansion coefficients of silicon and FR4.

The design of the assembly carrier is carried out so that it is compatible with both cooling options. An integration design has been conducted for the gas cooling option. Figure 121 shows the design.

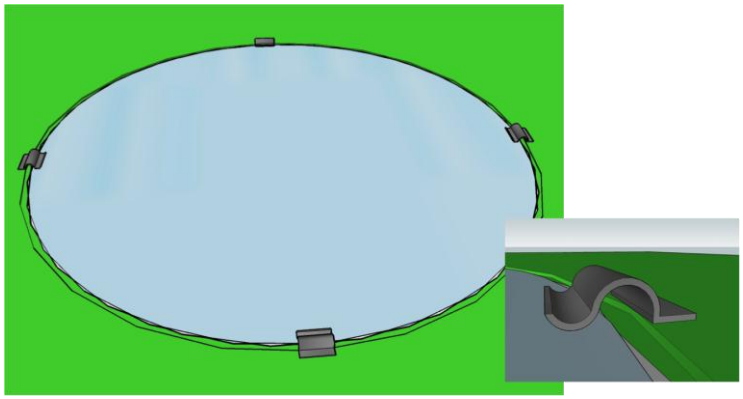

Figure 120 Carbon fiber clips.

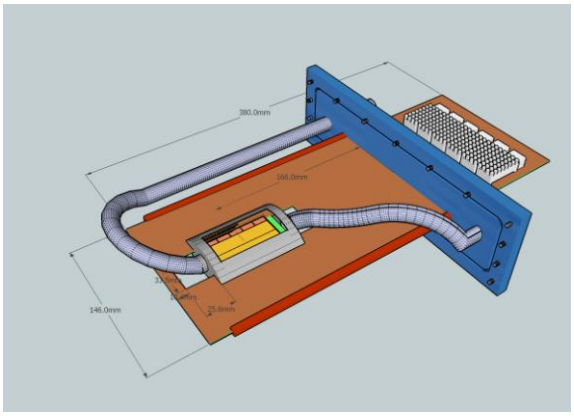

Figure 121 Gas cooling integration.

2.3.10 GTK Production, Test & Assembly

In the context of the GTK project several tasks are challenging and need dedicated research for production and assembly optimisation. In short these tasks are:

- Design and test of the GTK ASIC
- Assembly bump bonding and thinning.
- Design, optimisation, production and test of the cooling system
- Integration of GTK module – assembly of sensor-ASIC assembly with cooling elements.

- Integration of GTK module – assembly of silicon assembly with GTK assembly carrier.

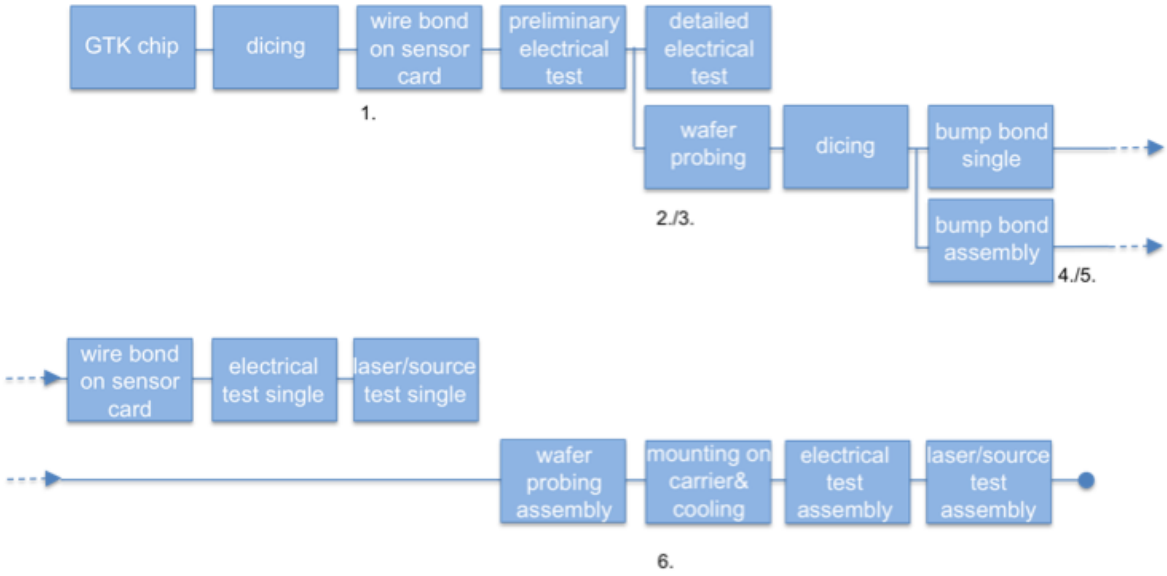

Figure 122 GTK test& production flow chart.

In order to streamline the R&D phase and the production phase in terms of completion time it is mandatory that the tasks listed be tackled in parallel. Consequently for the optimisation process on the integration the final GTK assembly (sensor/ASIC) is not yet available and mechanical and thermal equivalent dummy structures are produced to elaborate a production process. Integrated heating elements mirroring precisely the dissipation on the final ASIC on the dummy structures allows the cooling qualification. The same dummy structures are also used to emulate the electrical behaviour of the ASIC by connecting high speed drivers to it. Thus they also will be used to electrically qualify the GTK assembly carrier PCB. Once the GTK assembly is available the production process is optimized and allows a stream-lined integration.

Figure 122 shows the qualification chain flow chart once the GTK ASIC has been produced. The flow diagram shows several required parallel tasks and assumes the following points (as noted in the figure):

1. Sensor card, where the single chip sensor/ASIC assembly is glued and wire bonded is available when ASIC is ready.
2. Wafer prober has been setup in advance including mechanics, electrical test systems and firmware/software.
3. Wafer probing is run in parallel without the need of the expert performing the detailed tests.
4. Bump bonding & thinning understood in advance using dummy assemblies.
5. Bump bonding of single & assembly is done in parallel.
6. Mechanics assembly & integration with GTK assembly carrier and cooling verified and optimized with dummy assemblies before.

2.3.11 DAQ - Off Detector Readout

2.3.11.1 Introduction

The aim of the off detector readout system is to provide the interface to the on detector electronics, the DAQ system, the trigger system and the NA62/GTK detector control system (DCS). Figure 28 on page 71 shows a full GTK block diagram. The following sections give an introduction to the system and its interface. Section 4.1.8 on page 345 in the Trigger and DAQ chapter give a detailed description of the readout.

2.3.11.1.1 Detector Interface:

In particular the system is connected to the 3 GTK stations via optical fibers. For each of the 10 chips in a station 2 – 4 multi gigabit serial data links totalling a data rate of 6 Gb/s and chip and one slow status data link are foreseen. In the direction to the detector one clock link and one slow serial configuration link per chip is foreseen. The GTK off detector read-out system reads the raw data stream from the high speed serial links, stores the entire data flow until the L0 trigger decision and selects all data in a time window of 75 ns around the event indicated by the L0 decision which arrives with a latency of 1 ms. In an upgrade phase of the trigger system the latency might go up to several ms. The design of the GTK off-detector read-out will take this possibility already into account.

Serial status information is accepted by the off detector read-out for each chip individually. It continuously evaluates and forwards them to the GTK DCS. Serial status information contain temperature values, acknowledge of commands sent on the serial link and read-back of configurations. Furthermore both the data stream and the status data are monitored for its data consistency in the off detector read-out FPGA processor and consequently by the GTK DCS.

The low jitter clock (< 20 ps jitter) is sent to each chip individually on equal length fibers. The serial configurations are sent to each chip individually allowing to address each chip independently. The serial data contains information such as global and local pixel thresholds, programmable read-out parameters and programmable bias circuitry values.

In total the number of links connected to the off-detector read-out is: 60 or 120 gigabit serial links for a total data rate of more then 120 Gb/s, 30 status links, 30 clock links and 30 configuration links.

2.3.11.1.2 Trigger Interface:

The GTK off detector electronics expects a constant latency trigger signal. It is planned to use the TTC protocol to transmit trigger information. As the GTK off-detector system is placed in the control room and thus not subject to radiation the TTCrx chip can be replaced by an FPGA if it turns out to be more practical. In case of a non-constant latency L0 trigger signal a time tag synchronized with the GTK time domain must be attached to the trigger signal allowing extraction of the corresponding time slice from the GTK read-out memory. Two signals are foreseen to communicate from the GTK read-out to the NA62 trigger system. The throttle signal indicates to the NA62 trigger system that the present trigger rate is too high and buffers are going to overflow. The error signal indicates that

data loss already has occurred and the GTK read-out system attempts to recover process integrity and/or needs a reset.

2.3.11.1.3 Detector Control Interface:

The GTK off-detector read-out is connected to the GTK PC based DCS system via a dedicated bus system. The GTK DCS initiates via the GTK off-detector read-out the downloading and verification of configuration as well as online verification of parameters and data flow.

2.3.11.1.4 Safety Interlock Interface:

The GTK off-detector read-out system is directly connected to the safety interlock without passing through a PC. In case of a detection of a temperature deviation outside of the programmable limits the read-out interacts with the power supplies via the safety interlock feature.

2.3.11.1.5 DAQ Interface:

The GTK off-detector read-out system is connected to the DAQ system with Gigabit Ethernet links pushing data upon a L0 trigger decision. Two Gigabit Ethernet links per GTK read-out chip are foreseen resulting in 60 Gigabit Ethernet links for the GTK.

2.3.12 GTK System Components, Auxiliaries

Figure 28 shows the full GTK block diagram.

The three GTK stations are located within few meters from each other and the radiation effects harming commercial electronics is negligible already a few 10 cm from the beam area. Consequently low voltage power supply modules for the read-out chip, the HV modules for the sensors, the safety interlock system and the environmental monitoring will be located in a service area next to the GTK some 10 m distant from the stations. Control busses to the units allow remote configuration and monitoring of voltage, currents and environmental parameters from the control room. The magnetic field in the GTK service area is expected to be negligible.

2.3.12.1 Low Voltage Supply

Each individual read-out chip is supplied individually. The read-out chip requires 2.5V (< 1A) and 1.2V (< 3A). Thus for all three stations 30 2.5V and 30 1.2 V LV power channels are required. One channel is connected to one read-out chip with two remote sense lines connected directly on the module allowing to controlling each read-out chip individually. The common ground connection of one station is established directly in the station. Consequently the power supply modules provide floating outputs. Each channel can be set and monitored with a resolution of 10 mV. The current limit and the current monitoring have a resolution of 10 mA with a sampling period of at least 500 ms. The modules can be disabled by the safety interlock system. In case of over voltage or over current the modules switch off automatically using a hardwired internal interlock system. The modules have a hardware adjustable voltage and current limit, which cannot be overridden by software.

2.3.12.2 High Voltage Supply

Each of the sensors in the three stations is individually supplied with bias voltage. The voltage can be adjusted from 0 V to 700 V with a resolution of 1 V. Current limits are set with a resolution of 100 nA with a maximum of 500 µA per sensor station. The actual sensor leakage current is monitored continuously and included in the GTK interlock system. The leakage current provides a measure of the accumulated radiation damage and the temperature of the sensor. Thus monitoring the current on a long term basis provides feedback on the radiation induced performance degradation and detects a potential need for retuning of the analogue front-end. Monitoring the current on short term provides temperature information, which is fed to the safety interlock system in addition to measurements of on-chip temperature sensors. The current monitoring resolution is around 100 µA with a sampling period of about 500 ms. The modules can be disabled by the safety interlock system. The modules have a hardware adjustable voltage and current limit, which cannot be overridden by software.

2.3.12.3 Safety Interlock

Operation of the stations in a controlled environment is very important. As the modules have little mass they need to be prevented from thermal run-away in case of a cooling failure. This need is accentuated as the modules are operated in vacuum at a temperature of 5 degree Celsius or below and there is no cooling effect via convection. In case the thermal gas cooling option (section 2.3.8.1) is chosen during power supply failure the interlock must act on the cooling as a thermal run-away to cold temperatures can destroy the module as well. Each read-out ASIC contains thermal sensors which are accessible via 2 pads and in addition are read-out via the status link of the ASIC. The direct hard-wired reading of the thermal sensor is also operative when the ASIC is not powered. Humidity sensors are placed close to the stations.

The core of GTK interlock system is a fully hardware based system with a reaction of 1 s and it acts on the power supplies of one chip individually. For each read-out chip it takes the hardwired temperature sensor and the humidity sensor into account on a real time basis. Via the status bus of each chip and the read-out electronics, the temperature values are fed into the read-out processor, which also have one hardwired output for each read-out chip. Each of these wires is also taken into account by the interlock system. This configuration allows a safe and from software process independent operation, as direct hard wired connections to the sensors and hardware allow a fast interlock action.

Although the interlock system operates independently from any computer it is connected to a computer based detector control system (DCS) which logs the environmental parameters, such as temperature, from the hardwired sensors, temperature measurement from the read-out stream, humidity values, low and high voltage power supply voltage and current and cooling system parameters.

2.3.12.4 Detector Control System

The detector control system is a computer based system which:

- sets the low voltage voltages and current limits on each individual chip;
- sets the high voltage, current limits and initiates the voltage ramp up/down;
- configures the read-out ASIC front-end configuration (global discriminator thresholds, individual discriminator thresholds, enables/disables individual pixels, set trim DACs);
- configures programmable parts in the read-out section of the ASIC;
- configures programmable parts in the off-detector read-out electronics;
- monitors the status of the detector elements (environmental parameters, data consistency) regularly and reads back configuration registers;
- configures and sends an input to the interlock system;
- logs permanently the system parameters; and
- issues warnings and alerts.

2.4 CHANTI

2.4.1 CHANTI Detector Requirements

The Charged ANTI (CHANTI) detector is required in order to reduce critical background induced by inelastic interactions of the beam with the collimator and the Gigatracker (GTK) stations as well as to tag beam halo muons in the region immediately close to the beam. The most critical events are the ones in which the inelastic interaction takes place in the last GTK station (GTK-3). In such cases, pions or other particles produced in the interaction, if emitted at low angle, can reach the straw tracker and mimic a K decay in the fiducial region. If no other track is detected, these events can appear like a signal event, i.e. one single π^+ in the final state. A GEANT4 simulation has shown that kaon inelastic interactions with GTK-3 happens in about $1/10^3$ cases, so that the combined rejection factors of the analysis cuts and the CHANTI veto must lead to a remaining inefficiency of 10^{-8}.

The purpose of the CHANTI is to identify inelastic interactions in the GTK by tagging particles at higher angles w.r.t. to the beam. This can be achieved by placing a number of guard counters right immediately after GTK3. At the same time the CHANTI can also veto beam halo (muons) in the region closest to the beam.

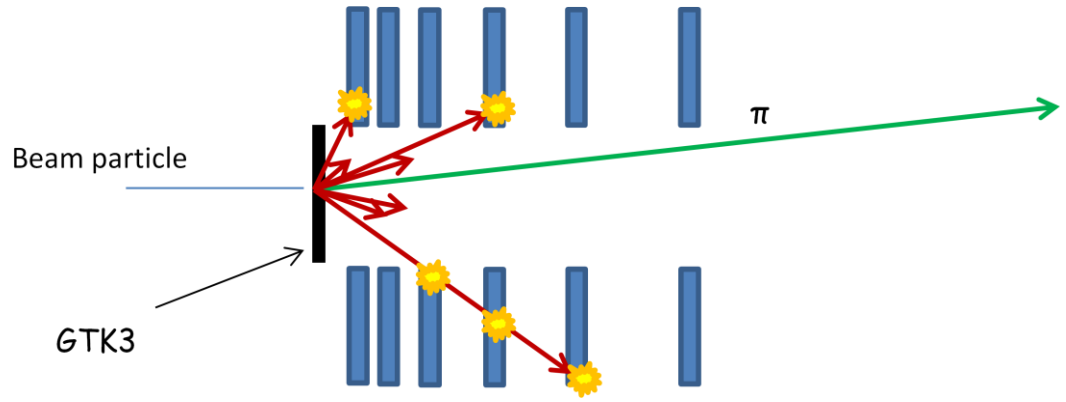

Figure 123 Despite the ulltra thin GTK detectors, beam particles can create accidental background if they undergo inelastic interaction in the GTK material. The scattered particles are revealed in the CHANTI which is placed immediately after the third GTK station.

Finally a veto for charged and neutral particles placed just downstream of the last GTK station, provides additional rejection of the accidental background coming from hadronic interactions of the beam particles in the last GTK station, as previously discussed. This detector, called CHANTI, consists of scintillators assembled in a rectangular shape surrounding the beam.

The CHANTI detector must be able to tag inelastic events with high efficiency. Given that it will be sensitive to the muon halo and to the inelastic interactions the expected rate of particles that release enough energy to be detected will be around 2 MHz. Even if it is not intended as a trigger veto at L0, the CHANTI must have a good time resolution (≤ 2 ns) to keep the random veto rate at an acceptable level: for instance, assuming a 5 sigma (10 ns) time coincidence window with the event fine time at reconstruction level, a 2% inefficiency on the signal would be induced by CHANTI

random vetoes. Although tracking capability may be not mandatory for the system, it can help in distinguishing beam halo events from inelastic interactions and in monitoring the beam halo itself very close to the beam. Last but not least, position sensitivity is useful in improving time resolution without increasing too much the granularity of the detector. Charge measurement is needed only to the extent of obtaining time slewing corrections, and, in the layout proposed, to improve the spatial resolution.

2.4.2 General Layout

The general layout of the system is sketched in Figure 124. It is composed of six stations, placed inside the vacuum tube respectively at 27-77-177-377-777-1577 mm distance from the GTK-3 (see Figure 124).

The rectangular hole inside each station is 50 mm in y and 90 mm in X due to the rectangular shape of the beam. Outer square side length is 300 mm. For particles hitting the GTK-3 center the CHANTI covers hermetically the angular region between 34 mrad and 1.38 rad wrt the beam axis, for particles hitting one of the GTK-3 corners the coverage is hermetic between 50 mrad and 1.16 rad. This must be compared to the highest angle under which a LAV station is possibly able to detect particles produced in the GTK-3 that is 49 mrad for particle produced at GTK-3 corner, so that LAV complements at low angles the information given by CHANTI. The CHANTI, by itself, is able to veto about 95% of all inelastic interactions of K in GTK-3 regardless of the final state. This vetoing efficiency reaches almost 99% if one restricts to potentially signal-like events, namely the ones where the kaon either did not survive the inelastic interaction or did not decay in the fiducial volume, and one track is reconstructed by the straw spectrometer.

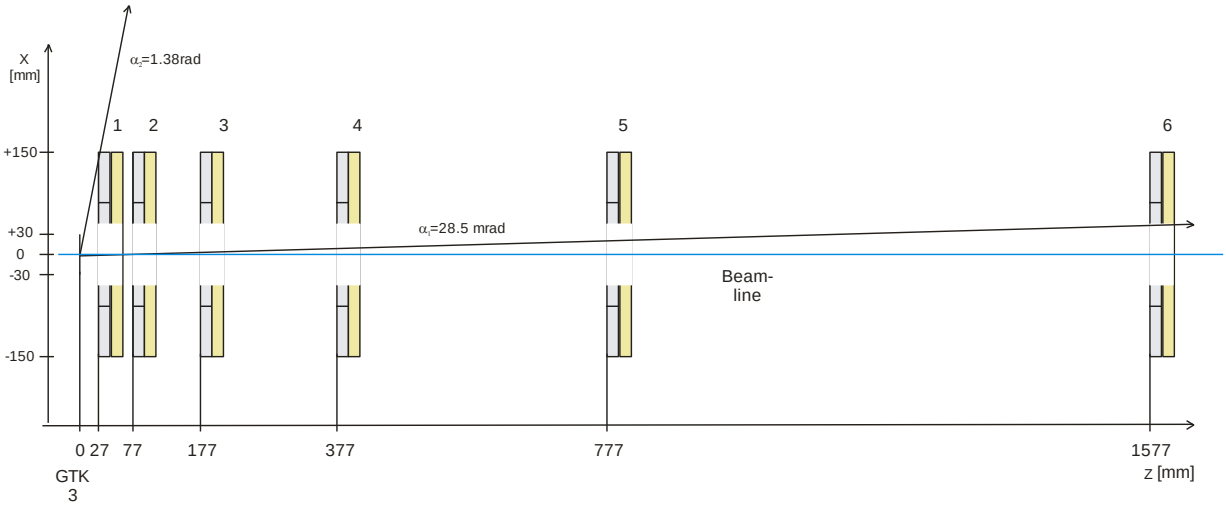

Figure 124 Positioning of the six CHANTI stations on the beam line (top view). The first colored line is GTK-3, the distance between GTK-3 and the first CHANTI station is 27 mm.

Each station is made up of two layers, called layer x and y respectively. A y(x) layer is composed of 22 (24) scintillator bars arranged parallel to the x(y) direction and individually shaped at the appropriate length (see Figure 125). Each layer is in the end composed by two sublayers, made up by 10+12

(10+14) bars, and staggered by half bar. Each bar is triangularly shaped, and two staggered bars face oppositely as shown in Figure 126.

Light is collected by means of one WLS fiber placed inside each bar and read at one side by a photodetector.

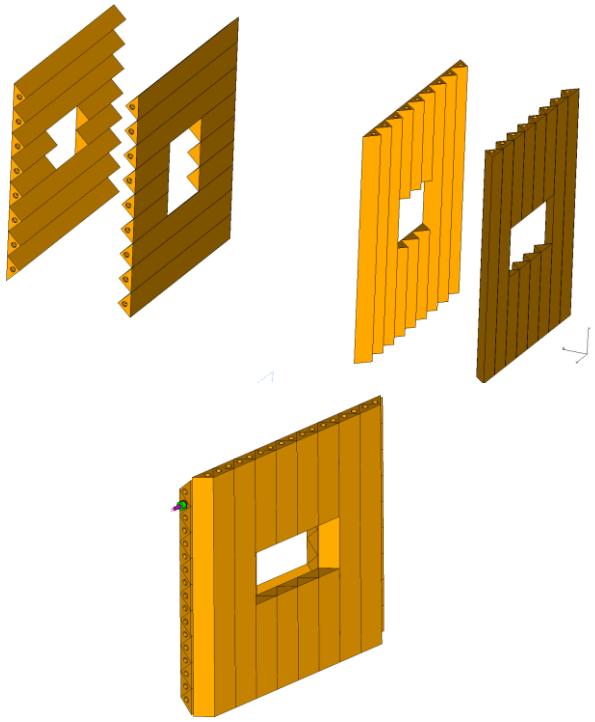

Figure 125 Layout of a complete CHANTI station. Top: exploded view; bottom: assembled station

2.4.3 Basic Scintillator Layout

Figure 126 Staggering of triangular scintillator bars to form a plane

The basic building block of the CHANTI is a scintillator bar in form of a triangular prism similar to the ones used in the D0 preshower detectors [45] and the MINERVA experiment [46] . It is produced at the NICADD-FNAL extruded scintillator facility [47] and consists of an extruded polystyrene core (Dow Styron 663 W) doped with blue-emitting fluorescent compounds (PPO 1% by weight and POPOP 0.03% by weight) and a co-extruded TiO_2 coating (0.25 mm thick) for reflectivity. The cross-section of the bar is a isosceles triangle with a base 33 mm and height of 17 mm, with a hole placed at 8.5 mm from the base . The hole has 1.7 mm diameter to host a WLS fiber. Optical glue ensures the coupling between the fiber and the scintillator. The main characteristics of the scintillator are:

- good LY (100% of Kuraray SCSN-81)
- radiation hardness (5% degradation observed after 1Mrad γ irradiation)
- Low cost
- fast response (τ few ns)

The WLS fiber used is Bicron BCF92 [48]. The fiber is read only at one side, and in order to recover light emitted in the direction opposite to the photodetector it has been mirrored at one end, by means of Al sputtering in vacuum, using the same technique developed for the fibers of the ALICE EmC calorimeter [49]. BCF92 has a very short fluorescence time (2.7 ns) and a light output and attenuation length comparable to the one of Kuraray Y11 .The WLS fiber is coupled to a SiPM device for photon detection.

The triangular shape allows a gap-free assembly when two bars are put one facing the other, in an almost self-sustaining shape. Moreover, the amount of light shared between two adjacent bars depends on the position of the impact point of the particle w.r.t. to the triangle centres (i.e. the position of the WLS fibers) This allows to determine the impact position in the direction orthogonal to the fiber with about 3mm resolution, much better than the one expected for rectangular shaped bars (roughly 10 mm) given the 33 mm spacing of the fibers .

2.4.4 Mechanics and Supports

The whole CHANTI detector consists of six stations which are placed inside a single vacuum tight vessel, together with the GTK-3 station. A possible solution using a rectangular vessel with a vacuum tight removable cover, is sketched in Figure 127 where the last GTK station and the CHANTI stations are visible. Mechanics and supports for GTK-3 will be the same as for the other two stations.

Each station is composed of scintillator bars glued together, so that they form a self-sustaining structure. A light Al frame is supporting mechanically each station to avoid any risk of ungluing and to fix the station to the outer vessel (see Figure 128).

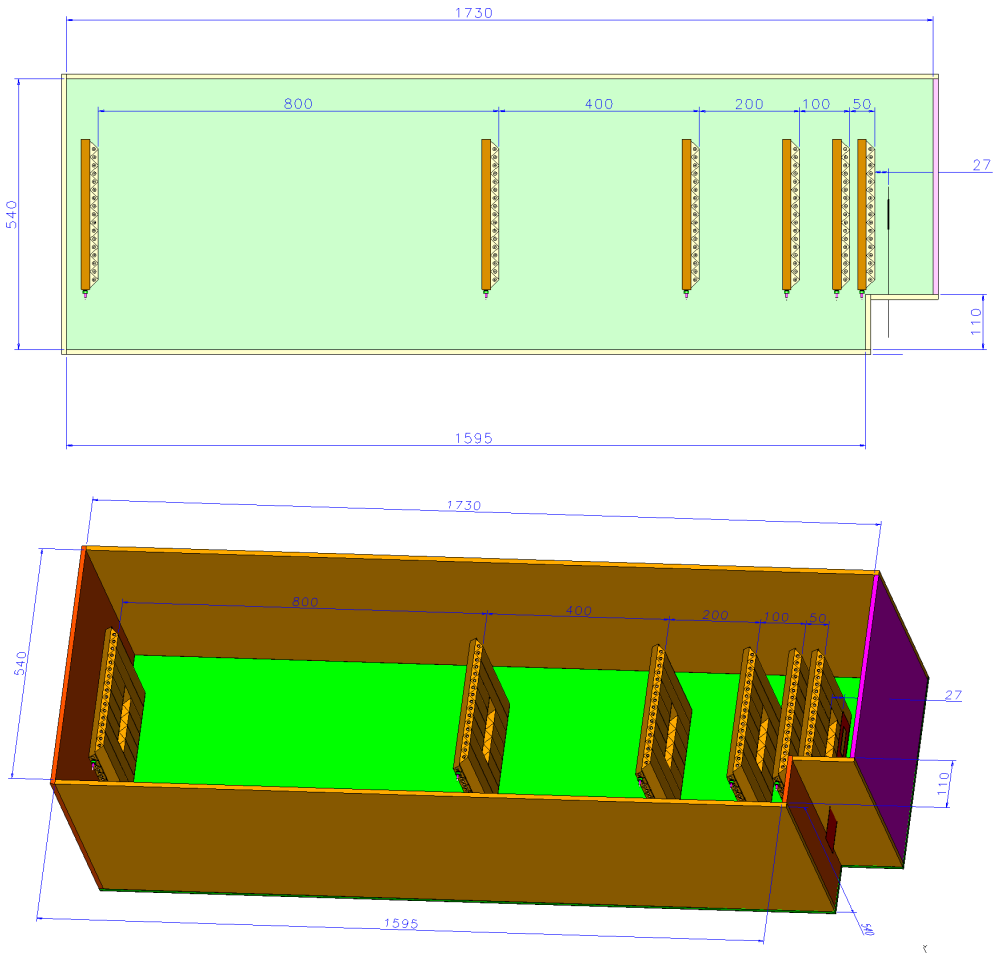

Figure 127 Layout for GTK-3 - CHANTI vessel. Vacuum flanges are not shown.

One station is composed of 46 bars with different length. The bars outside the beam gap are all of the same length, the so-called long bars (L). Since the gap is of rectangular shape, the bars in the central parts are of two different lengths, depending whether they are in the horizontal layer or in the vertical one. There are so-called middle (M) bars and short (S) bars respectively. The full lengths is 300mm (L), 117.5mm the middle type (M) and 102.5mm the short type (S). Table 24 summarizes the composition of one station.

Table 24 Lengths and number of long (L), middle (M) and short (S) type scintillators.

Bar type/Layer	*L*	*M*	*S*	**Total**
Layer Y	10	-	12	22
Layer X	10	14	-	24
Full station	20	14	12	46

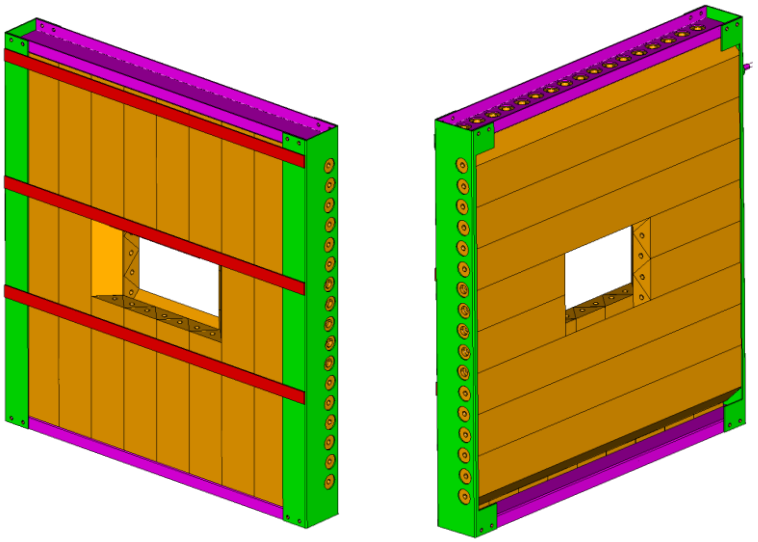

Figure 128 A sketch of a single CHANTI station, with the elements of the supporting frame (in green, violet and red). Left: beam view; right: rear view

Every scintillator bar is provided with a custom designed connector (see Figure 133) which is inserted in a precision hole (Ø 1.02 mm) and coupled to the photodetector. A precisely machined screw cap holds the photodetector in the right position. Alignment between the fiber center and the photodetector sensitive area is guaranteed at the level of 50 μm and is dominated by the tolerance in the position of the photodetector area w.r.t. its frame as provided by Hamamatsu.

It is worth to note that it is possible to substitute a photodetector by just unscrewing the connector. Moreover, quality tests can be easily arranged by using the same photodetector for all bars.

The precise definition of the position of the bars is given by a custom machined reference jig, (see Figure 138) on which each bar is placed “cusp down” to form a planar surface with the nominal dimensions. Bars in the orthogonal direction are then glued to these ones to form a “waffel – like” structure . The complementary bars are finally added in both directions to form the final planar station. See also 2.4.6, Figure 139 and Figure 140 for details.

Figure 140-e shows the final geometry with the readout connectors on the four sides of the station.

2.4.5 Photodetectors

The use of Silicon Photomultiplier (SiPM) devices as alternative to traditional photomultipliers tube (PMT) is becoming a widespread solution in particle detectors when high number of channels or high level of integration is needed.

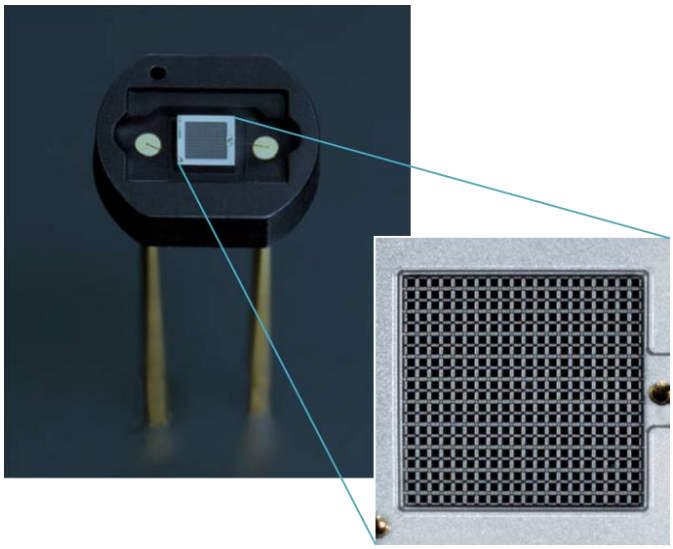

Figure 129 Hamatsu SiPM in ceramic packing

SiPM's are essentially [50] [51] [52] arrays of O(100) avalanche photodiodes (APD) (typical pixel size are 50x50 or 100x100 μm^2), see Figure 129, on one chip with shared readout in which each pixel is working in (quenched) Geiger mode (i.e. in a digital on/off configuration) but the number of firing pixels (and thus the overall response of the device) is proportional to the amount of incoming light (in the hypothesis Nphotons<<Npixels). Typical photon detection efficiencies (PDE) are highly dependent to SiPM model, they spread from 20-25% and can reach also 60-65% and gains are at level 10^5-10^6 (typical APDs gain are 10^2). This gain can be reached with no need for an HV system since the operation voltage is typically between 30 and 70 V and the leakage current of order of few nA. They can sustain very high rate, O(10 MHz) without problems. The dark rate (i.e. the rate observed with a counter at 0.5 pe threshold) , depending on manufacturer, on the pixel size and on the overvoltage, can vary from 100 kHz and few MHz at room temperature. It is not a concern if the number of expected photoelectrons is reasonably high, since it scales roughly by an order of magnitude per photoelectron as far as the threshold is increased to 1.5, 2.5, 3.5 pe and so on. Working with, for example, 3.5 pe threshold reveals a typical dark noise of about 100 Hz/channel.

SiPMs are intrinsically radiation hard devices. There is however, as for all semiconductor detectors, a known issue with their behaviour after intense neutron flux irradiations. The ASTM E722-93 standard practice allows to compare damage on silicon devices from different neutron sources by normalizing it to the damage induced by monoenergetic 1 MeV neutrons. In this context the figure of merit of an environment for SiPM devices is the number of equivalent 1 MeV neutron /cm^2 crossing the detector. It is known from literature [53] that a neutron irradiation corresponding to 4* 10^8 /cm^2 1 MeV equivalent neutron or less gives no visible effect on SiPMs, while increasing further the irradiation the dark noise starts increasing, reaching about 10 times its initial value at about 2-3 10^9 /cm^2 neutron fluence. Even if a 10 times larger noise w.r.t. to the standard one could still be manageable by increasing the threshold by 1-2 pe, we have checked that the radiation should be below this level at the CHANTI for at least two years of operation. Two different simulation have been carried out: the first, using FLUKA, has shown that the neutron dose at the CHANTI, generated by the primary proton beam on target should be of the order of 0.04 Gy/y which corresponds roughly to 10^8 n_{eq}/cm^2 /y. The second simulation has been done using GEANT4, to understand the

contribution to the neutron flux from the inelastic scattering of the beam in GTK, which were not considered in the first simulation. It has shown that the expected neutron fluence from GTK corresponds to less than 10^8 n_{eq}/cm^2 /y .Combining the two result we have that in two years running of NA62 the CHANTI SiPMs should integrate no more than 4 *10^8 n_{eq}/cm^2 i.e. should operate in safe conditions.

In order to choose the best solution for our photodetectors we tested some Hamamatsu SiPM series 13-50, 11-50 and 11-100 (first number 13 or 11 is the SiPM dimensions in tenths of millimeter, second number 50 or 100 is the pixel size expressed in microns). Hamamatsu provides, for each SiPM, specifications and working parameters as bias voltage, gain and dark rates, all of them measured at 25°C. In order to reproduce climatic conditions we used a thermostatic chamber that could fix temperature better than 0.1°C. A collimated Sr^{90} source has been used to compare different devices coupled to the same scintillator bar. Signals are amplified using a fast 10X amplifier. Data is collected using a Tektronix TDS5054 5GS/s oscilloscope via GPIB connection and a custom LabView program. Oscilloscope bandwidth is 500MHz, enough to follow few nanoseconds signal rise time.

Relative comparison in terms of number of photoelectrons collected is our figure of merit. First of all we measured for each SiPM single photoelectron response (SPR), this operation has been done using thermal-generated signals (dark noise) (see Figure 130). Once obtained the single photoelectron normalization factor the ratio signal/SPR for the irradiated bar could be used to compare different devices.

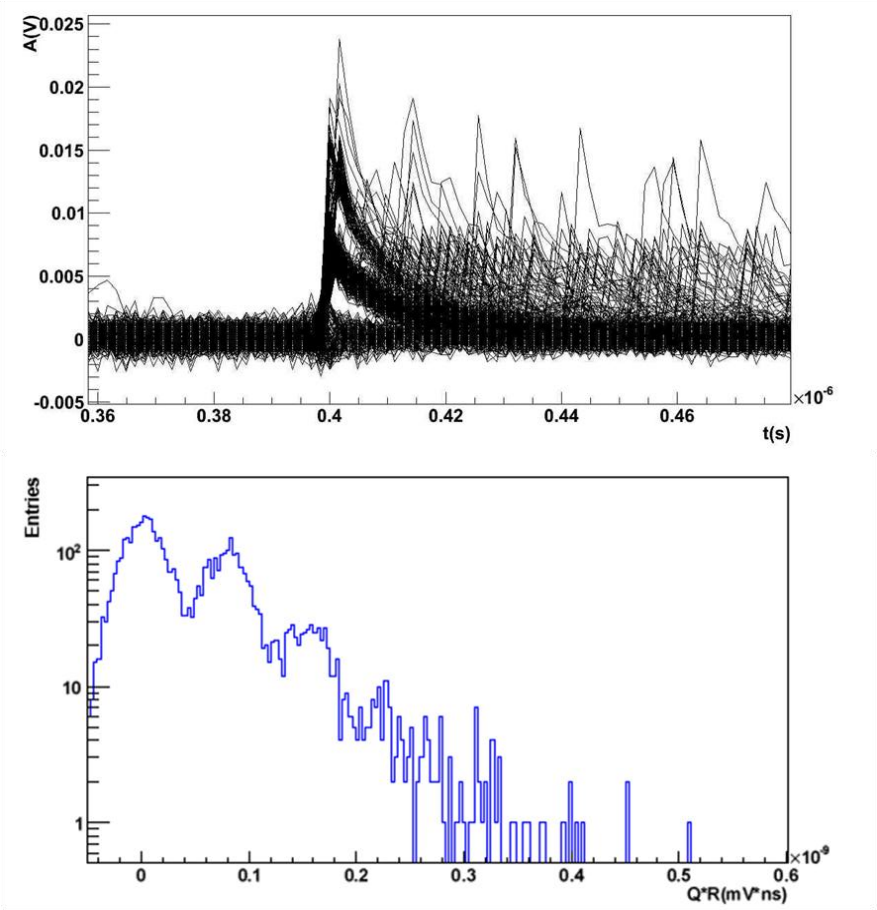

Figure 130 a) Thermal generated waveforms at 25°C. b) Dark noise spectrum obtained integrating waveforms shown on top.

Measured light yield for each SiPM is reported in Figure 131. We can conclude that SiPM serie 11-50 has a significantly lower photoelectron yield, while both 13-50 and 11-100 series seem viable solutions for our purposes.

In order to obtain an absolute scale of the expected number of photoelectrons (pe) the detector response was measured using cosmic rays. These tests show that for 300 mm long scintillator bars, equipped with (mirrored at one side) BCF92 fibres and Hamamatsu SiPMs (Type: S10362-13-050-c Ø=1.3 mm, with 50 X 50 μm^2 pixels) more than 100 pe are generated per detector plane, if crossed by straight MIPs i.e. more than 50 pe per channel/ MIP.

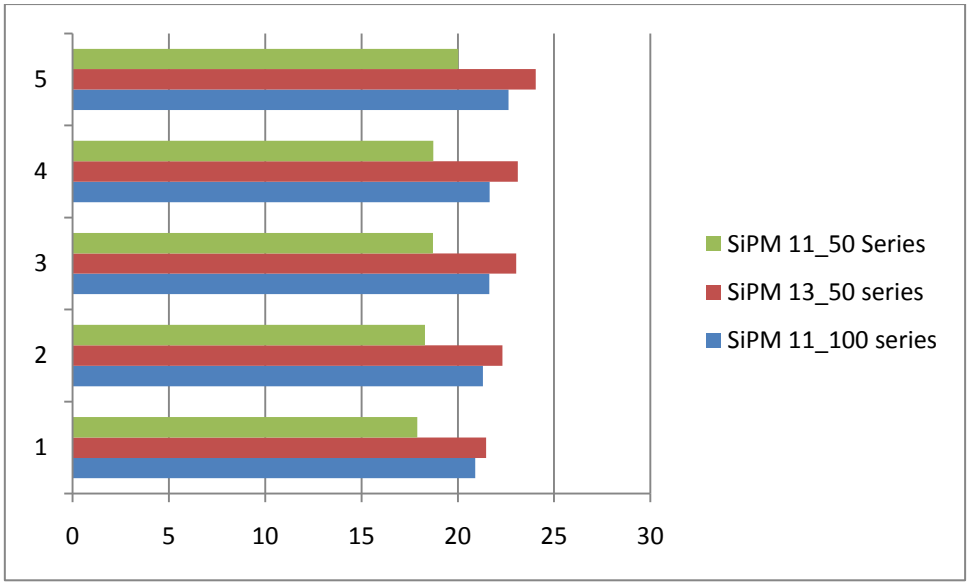

Figure 131 Photoelectron yield for photodetectors of three different Hamamatsu SiPM series. Five photodetectors have been measured for each series using the Sr^{90} source.

This allows to operate the device with a threshold in the range of 3-4 pe per channel or more with negligible efficiency losses.

Time resolution using the fast fluorescent Bicron BCF92 fiber has been measured (see Figure 132) to be about 800 ps for MIPS crossing the detector at a fixed distance from the SiPM: for an event with full x-y reconstruction this is thus the expected time resolution, well in line with the CHANTI specifications.

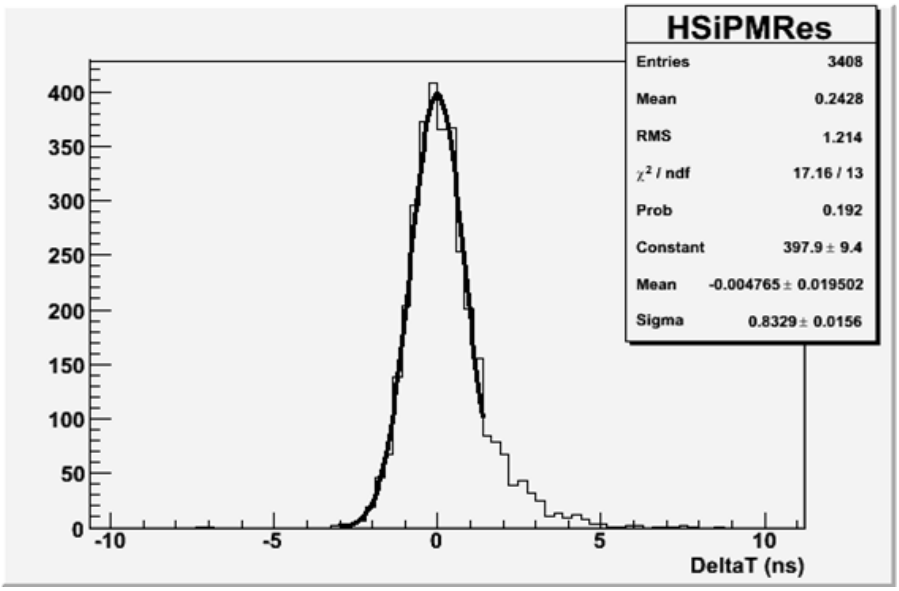

Figure 132 Time resolution measured for a triangular bar equipped with a Bicron BCF-92 fiber readout using a FBK-IRST SiPM.

2.4.6 Prototype construction

CHANTI prototype has been assembled in Napoli at the end of July 2010. It is a full dimension prototype of a X-Y station. Scintillator bars for the prototype were obtained courtesy of FNAL and Al sputtering of the fibers was performed at LNF.

The construction procedure adopted is hereby briefly described.

First, some custom tools to simplify the mechanical assembly were developed.

The assembly took about 20 days, and can be divided in three main parts:

1) Gluing fibers into bars,

2) Test of the bars,

3) Assembly all test-passed bars into final X-Y station.

A single X-Y station contains three different types of bars: 300mm (L), 117.5mm (M) and 102.5mm (S).

All operations have been done in a class 100 clean room environment and all components have been accurately washed using an ultrasound cleaning before manipulation.

2.4.6.1 Gluing of the Fibers into the Scintillator Bars.

This operation is itself done in two steps. First mirrored fibers and connectors are glued together and finally this ensemble is glued into a bar.

2.4.6.1.1 Fibers-Connectors Gluing

Each connector is provided together with a Teflon© cap. It is a multi-purpose tool (Figure 133) to define a reference plane for the fibers as well as to protect the polished side of the fiber during transport and handling.

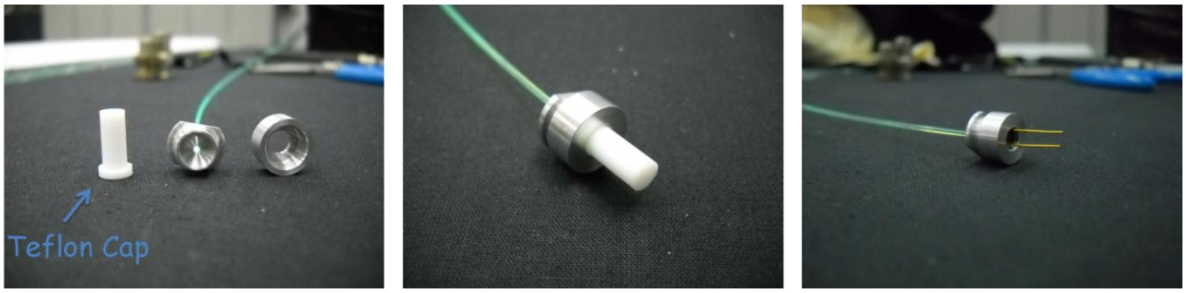

Figure 133 a) Connector in its parts b) Assembled connector c) Connector with SiPM

During the prototyping phase different glues were tested and the epoxy ARALDITE 2011 was found the best choice for this application. It has a high viscosity which helps to prevent glue leaking into the wrong connector side. This glue is also solvent-free avoiding cladding damages. Optical properties are not important at this level, because no coupling is required.

Fibers are plugged in their final position into connector, being careful they reach the Teflon cap. A special support (Figure 134) has been developed in order to parallelize this operation. It is able to carry 24 fiber-connector couples. When the fiber is in its correct position glue is put using a syringe (with a Ø 1.3mm needle).

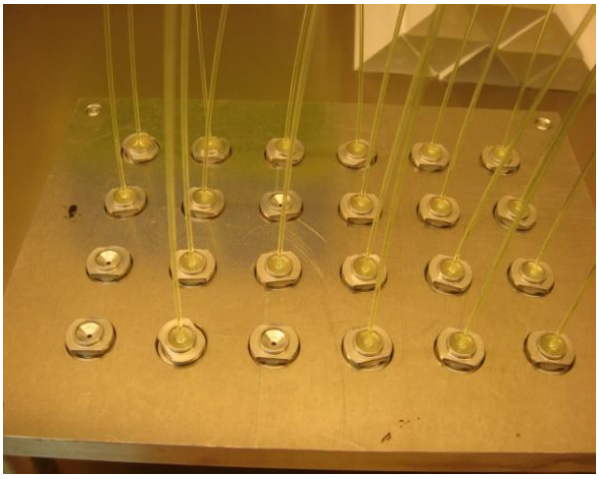

Figure 134 Tool for connector-fiber gluing.

Figure 135 The fiber-bar gluing tool in action. Its back side is identical to the one shown. Glue is injected from bottom using syringer.

2.4.6.1.2 Fibers-Bar Gluing.

Once obtained the fiber-connector ensemble they were glued into the bars. A custom tool (Figure 135) allowing to fix the bars in vertical position was developed. It can host up to 10 bars, and was used to hold the bars during the hardening of the glue. Five days are necessary to glue all the bars for one station.

Glue used is a SCIONIX Silicon Rubber Compound, that guarantees a good fiber-scintillator optical coupling and is known from NASA database [54] to be low outgassing.

The glue is injected from the bottom using again a syringe, this method was found to reduce the risk of trapping air bubbles in the glue. The required glue quantities are adapted for each bar length (L -> 2.1ml, M -> 1.0ml, S -> 0.9ml). This is important in order to avoid leaking at the top of the bar. The whole prototype contains about 65ml of glue.

2.4.6.2 Scintillator Bar Test.

Since after complete assembly any bar substitution is impossible a quality test before assembly is needed. For each bar the response to a Sr^{90} collimated beta-source is measured. Measurements have been carried out with the same setup described in section 2.4.5 for the SiPM comparison, in a controlled temperature environment using the same photodetector (an Hamamatsu 13-50 type) coupled each time to a different bar. Experimental apparatus is in Figure 136.

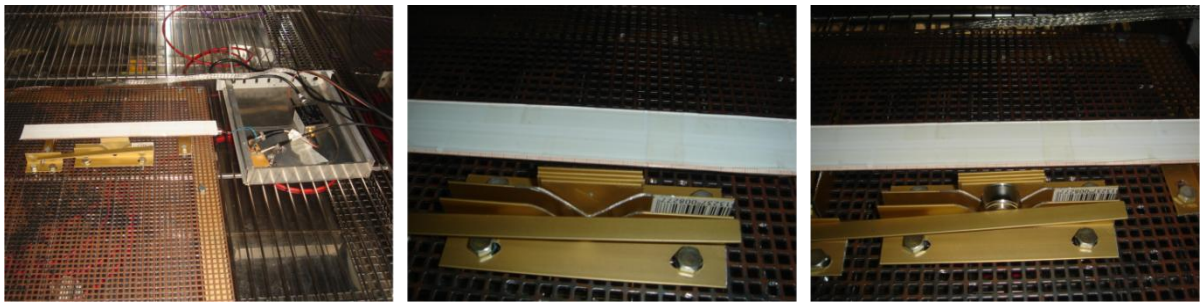

Figure 136 Experimental apparatus for bar test. a) Setup overview on the right the readout electronics and amplifier. b) Collimator detail. c) Source in finial position.

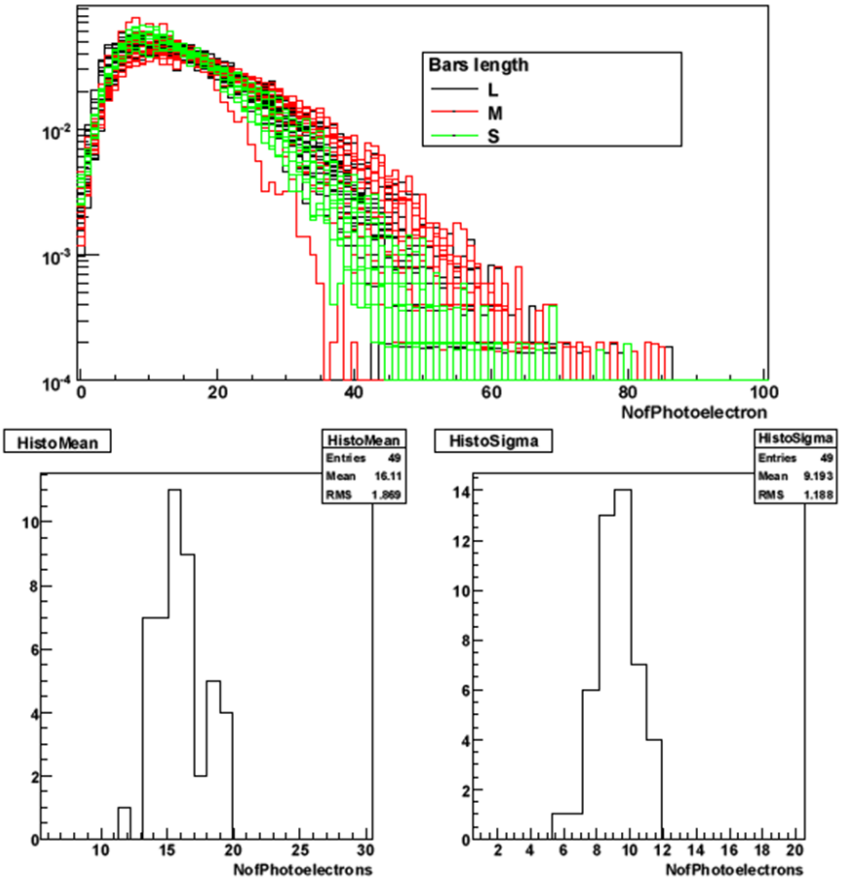

Figure 137 Bars quality test: a) Measured Spectrum for the different scintillators, b) Measured spectrum mean value, c) Measured spectrum RMS

Relative comparison among different bars in terms of number of photoelectrons collected is again our figure of merit.

Figure 137a shows all bars normalized source-response in terms of number of photoelectrons. Mean values for each curve are given in Figure 137b : one can conclude that the bars quality is very uniform. Only one bar showed a significant difference, the cause of low light response was understood, after inspection of the bar: the fiber edge was found slightly backward with respect to its reference plane facing the SiPM sensitive surface, leading to lower light collection efficiency.

2.4.6.3 Module Assembly

Once all the material was ready, a prototype could be assembled in 2 steps. Each step takes one day.

During the first day a half-layer X and Y are glued together. First of all bars are arrange on a jig (Figure 138), afterwards glue points are defined using a Teflon mask (Figure 139). Then the bars of the other half-layer were aligned on top. A second jig is put on top of the assembly to align the last half-layer. Pressure is added to increase the glue uniformity. Each glue spot contains 0.1ml of glue.

Figure 138 Jig used to align bars during gluing, on top Teflon mask to distribute glue spot

On the second assembly day the prototype is completed. Figure 140 shows all the steps of the procedure. First of all bars of a half-layer are aligned on the jig and glue spots are applied (3 for L-type bars and 2 for M and S type bars). Then the complementary half-layer is put on top. Again glue spots are placed and the last half-layer placed. A weight has been applied in order to uniformly distribute glue. For the whole module about 9.5 ml of structural low outgassing epoxy glue (3M DP490), identical to the one used for the reinforcement of the LAV blocks, were used. .

Figure 139 Assembling "First step": a) A half-layers is arranged on the jig, b)Teflon mask is aligned, c) Glue spots are applied, d) Mask is removed, e) f) Bars of the other half-layer are placed, g) h) second jig is applied on top to align last placed half-layer and to redistribute pressure, i) Half-module as appear the day after gluing

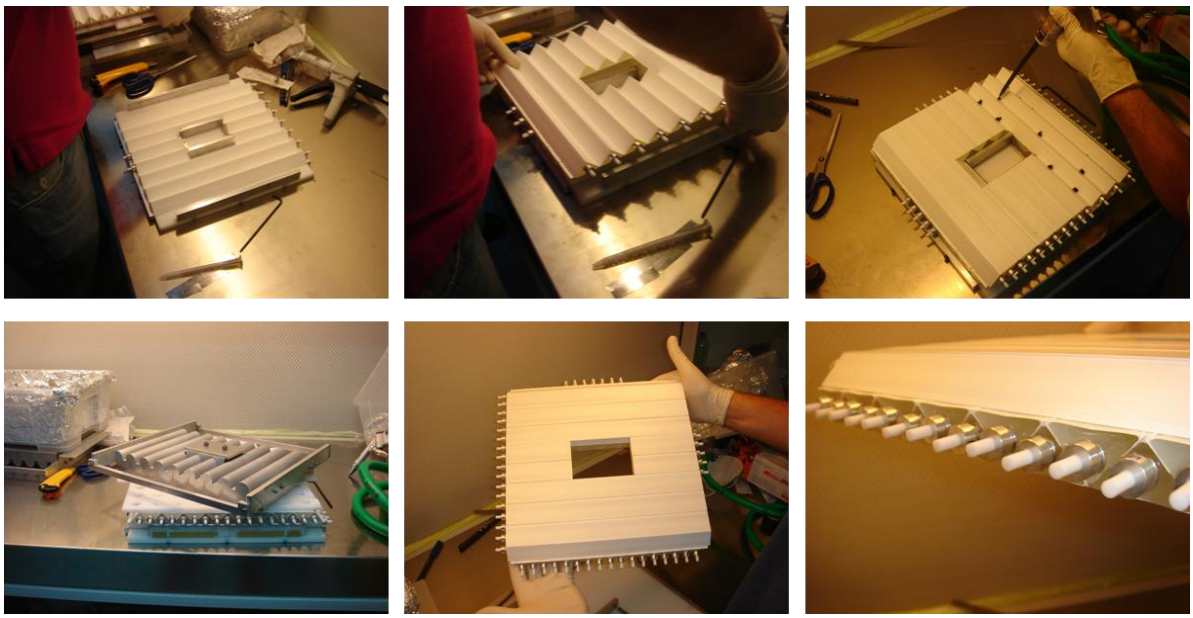

Figure 140 Assembling "Second step": a) Half-layer bars ready to be glued to half-module, b) Gluing half-module, c)Last half-layer gluing, d) Weight applied, e) Prototype as appears after one day, f) Connectors detail

2.4.7 Readout

The conceptual scheme of the CHANTI readout is shown in Figure 141.

As previously sketched each scintillator bar is coupled individually to a SiPM which converts light collected by the fiber into electrical signals. Each SIPM has two pins which are used both to polarize it and to read these signals.

The bias voltage is brought to the SiPMs inside and the signal are carried out from the vacuum tube using appropriate vacuum tight flanges as done for the LAV signals and HV. Both twisted pair and coaxial cables are being investigated as possible solution for the internal cabling , the final choice being the one which optimizes cost/benefits ratio.

A small amplification board is placed just outside vacuum in order to be able to transfer signals from the detector to the FEE crate with an acceptable S/N level.

Typical signals are expected of order few mV (on 50 Ω impedance) with a fast rise time (1 ns) and a somewhat long decay time (in the range 10-100 ns). The maximum expected rate in input to the FEE will be, for the inner bars, of the order of about 1 MHz per bar, as shown by Geant4 simulations, plus the dark rate (for Hamamatsu SiPMs, some 100 kHz). In order to keep some safety factor, the electronics will be designed to cope with a 5 MHz rate.

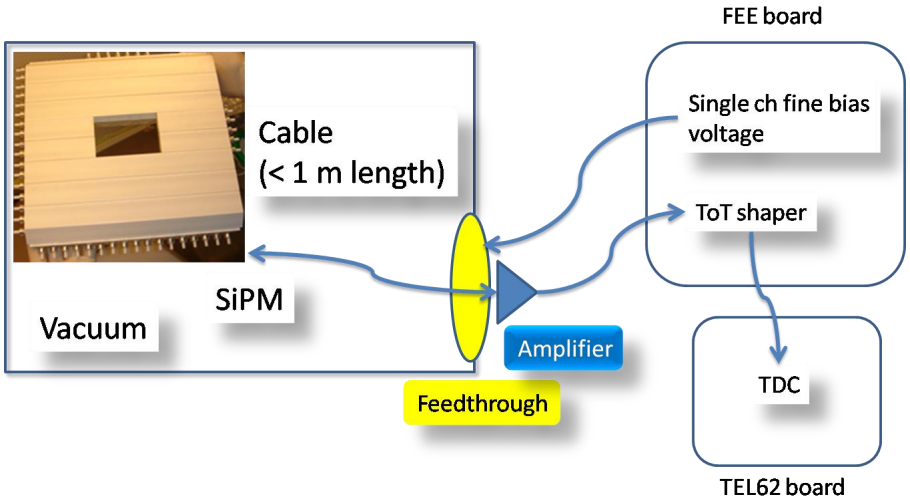

Figure 141 Conceptual scheme of the CHANTI readout.

2.4.7.1 FEE

The FEE boards provide for each channel:

- a way to control the V_{bias} with O(10 mV) accuracy
- a fast, DC coupled, conversion to a Time Over Threshold-LVDS signal output
- a temperature and/or a dark current (with nA resolution) monitor for slow control adjustment of the V_{bias}

Thresholds and V_{bias} settings will be controlled using the CANOpen standard.

The LVDS output will be directly sent to a TEL62 board equipped with HPTDC for both leading and trailing edge measurement. The total number of channels needed is 46 X 6 = 276. One TEL62 board equipped with three 128 ch TDC boards will be able to readout the whole system and provide also a large number of spare channels.

The TOT technique will allow also approximate charge measurement to improve the spatial resolution of the system.

2.4.7.2 Integration with Trigger/DAQ

The CHANTI is not supposed to give primitives for the L0 trigger. It uses the standard TEL62/HPTDC structure and obtains the clock via the LTU+TTCex system.

An evaluation of the maximum data rate can be done as follows. The highest multiplicity in the CHANTI is expected from the beam halo events, who cross all six stations. If fully efficient on these events the system will give at most 4 X 6 = 24 hits or 192 bytes per event. These events occur at 1 MHz rate. Inelastic interactions in the GTK will be detected at approximately the same rate, with comparable hit multiplicity. One can safely estimate a maximum of 200 bytes at 2 MHz. Assuming a 1 MHz trigger rate and a readout time window of O(100 ns) one expects O(200 kHZ) of such events in coincidence with a trigger. This would generate a data rate of about 40 MB/s well below the TEL62 specifications.

2.4.7.3 External Services

The external services needed for the CHANTI are power for the crate and for the common Vbias generator, and the CANOpen system for the DCS.

3 Decay Region and Downstream Detectors

3.1 Vessel, Vacuum System and Magnet

3.1.1 The Vacuum Vessel

The beam & decay vacuum vessel for the NA62 experiment, see Figure 142, is mostly made up of vacuum vessel parts from the NA48 experiment. In the NA48 configuration, the vacuum vessel parts referred to as the 'blue tube' were numbered from 1 to 22. Typically, 6m long cylinders were connected together into couples of 12m length. In the configuration for the NA62 experiment, the cylindrical parts start with number 7 which is positioned 105.595m from the target, and end with the connection window to the new LAV11 at 219.65m, a total length of some 114m.

Figure 142: The NA62 beam & decay vacuum vessel

The vacuum vessel starts in the experimental cavern TCC8 and continues into the cavern ECN3 where the floor is 1.66m lower. The beam axis in TCC8 is 1.2m from the floor whilst in ECN3 it is 2.86m. The bottom of the vacuum vessel in TCC8 is very close to the floor so that the supports of the vessel connect directly to the floor. In the ECN3 cavern, radiation shielding blocks will be used as a means of intermediate support from the floor to the supports of the vacuum vessel, as they were for NA48.

For the purpose of this description, the vacuum vessel can be described in three parts; the sections that are situated in TCN8 with a diameter of 1.92m, the sections situated in ECN3 with a diameter of 2.4m and the sections in ECN3 with a larger diameter 2.8m (the old NA48 helium vessel).

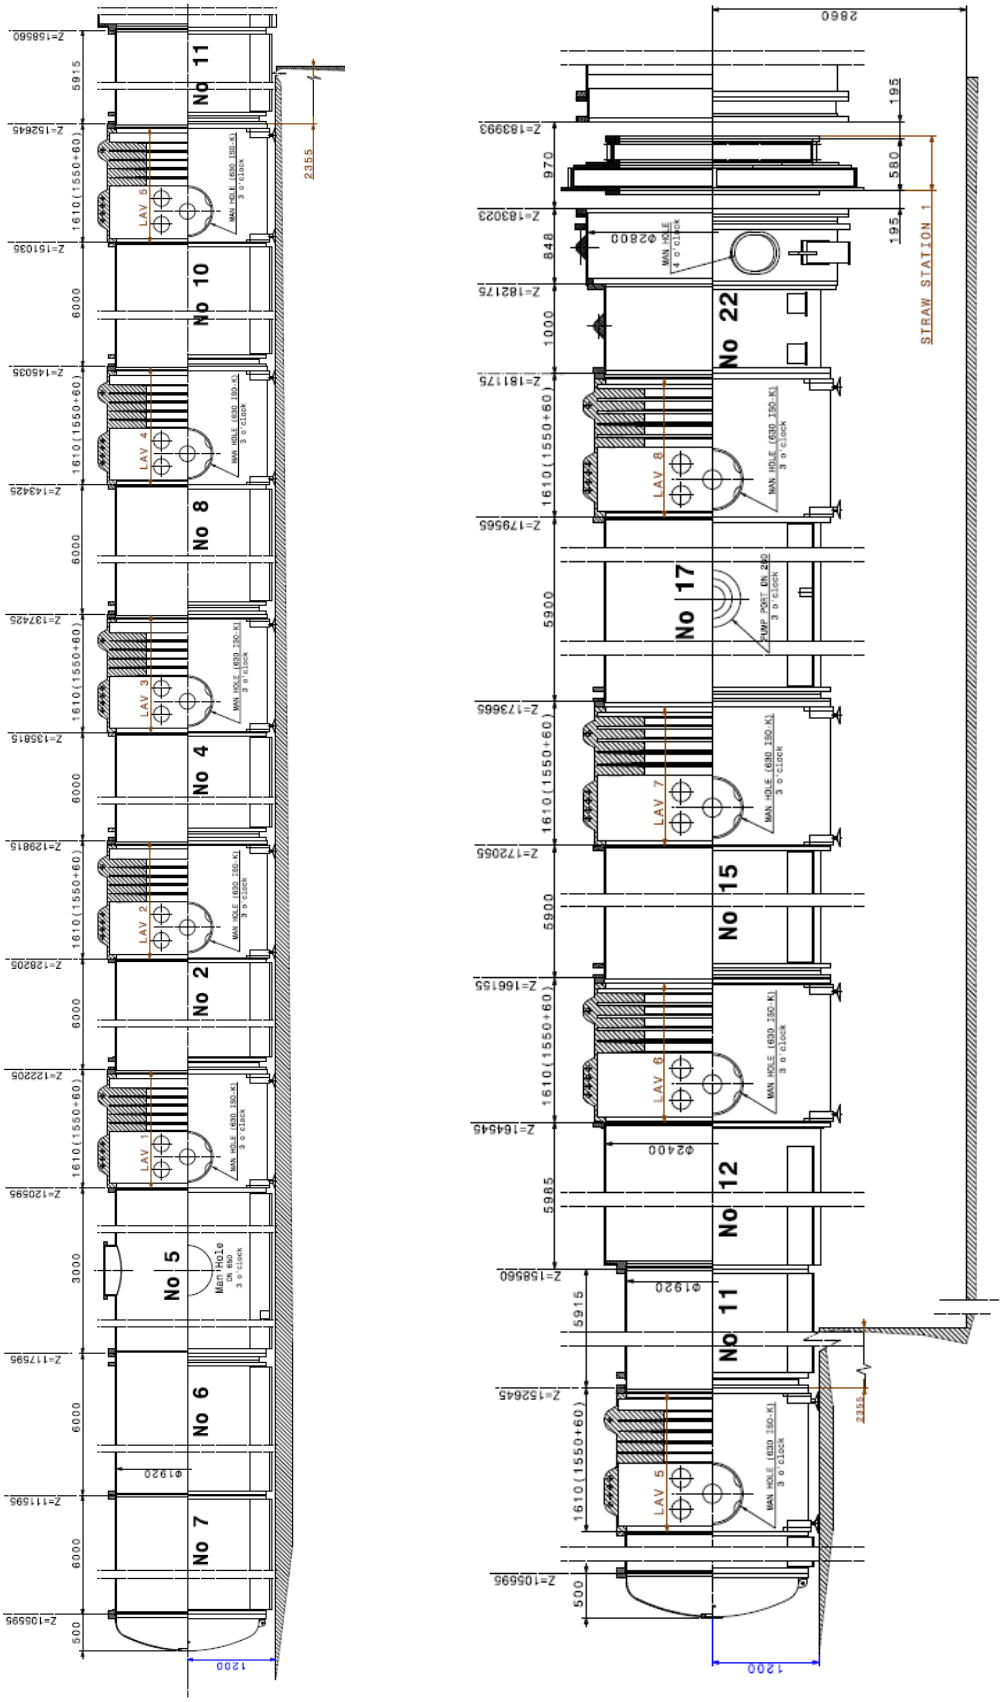

Figure 143 Vacuum vessel sections in TCC8 (parts 7 to 11) and in ECN3 (parts 11 to 22).

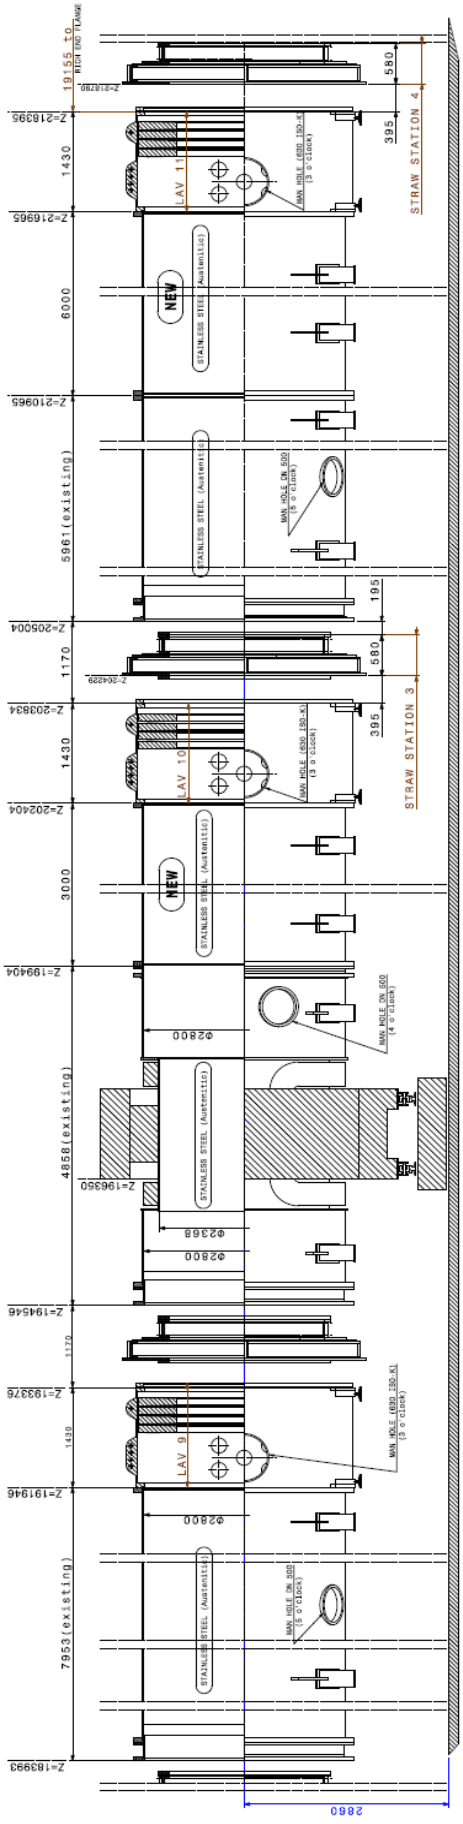

Figure 144 : Vacuum vessel sections in ECN3 of 2.8m diameter that include the MNP33 magnet together with the straw chambers and LAV stations.

The vacuum vessel in TCN8 is made up of the following parts:

- Part 7, 6m length, 1.92m inner diameter
- Part 6, 6m length, 1.92m diameter
- Part 5, 3m length, 1.92m diameter
- Large Angle photon Veto 1 (LAV1), 1.92m diameter (at the connection flange)
- Part 2, 6m length, 1.92m diameter
- LAV 2, 1.92m diameter
- Part 4, 6m length, 1.92m diameter
- LAV 3, 1.92m diameter
- Part 8, 6m length, 1.92m diameter
- LAV 4, 1.92m diameter
- Part 10, 6m length, 1.92m diameter
- LAV 5, 1.92m diameter
- Part 11, 5.915m length, 1.92m diameter (straddles TCN8 and ECN3)

From the above list, and the accompanying Figure 143, it can be seen that: the common diameter is 1.92m, the LAV stations 1-5 are all identical and are spaced apart by 6m long sections of existing 'blue tube'.

The vacuum vessel in ECN3 is made up of the following parts:

- Part 12, 5.985m length, transition piece with diameter increase from 1.92m to 2.4m
- LAV 6, 2.4m diameter
- Part 15, 5.9m length, 2.4m diameter
- LAV 7, 2.4m diameter
- Part 17, 5.9m length, 2.4m diameter
- LAV 8, 2.4m diameter
- Part 22, 1m length, transition piece with diameter increase from 2.4m to 2.8m

From the above list, and the accompanying Figure 144, it can be seen that: the common diameter is 2.4m, the LAV stations 6-8 are all identical and are spaced apart by 5.9m long sections of existing 'blue tube'.

The vacuum vessel containing the magnet spectrometer, the new straw chambers and the LAV's 9, 10 & 11 is made up of the following parts:

- Transition piece (diameter change) 0.848m length, 2.8m diameter
- Straw Chamber 1, 0.970m length (including interface parts)
- 7.953m length, 2.8m diameter, stainless steel
- LAV 9, 2.8m diameter
- Straw Chamber 2, 1.170m length (including interface parts)
- 4.858m length, 2.8m to 2.4m diameter reduction for the MNP33 magnet, stainless steel
- 3m length, new section, 2.8m diameter, stainless steel
- LAV 10, 2.8m diameter
- Straw Chamber 3, 1.170m length (including interface parts)

- 5.961m length, 2.8m diameter
- 6m length, new section, 2.8m diameter, normal steel.
- LAV station 11
- Straw Chamber 4, with connections to LAV 11 and RICH.

From the above list, and the accompanying drawing Figure 144, it can be seen that: the common diameter is 2.8m, the LAV stations 9-11 are all identical, the existing MNP33 magnet is re-used for NA62 and the four straw stations are all in this region. There are also two new vacuum vessel sections of 2.8m diameter required, one 3m long in stainless steel and one 6m long in normal steel.

At the end of this vacuum vessel, and after straw chamber 4, the beam vacuum continues through the RICH detector and the Liquid Krypton calorimeter by means of a beam pipe.

3.1.2 Connections and Adjustment Possibilities

For the NA62 configuration it is common to use 6m long sections of vacuum vessel connecting between LAV stations. Practically every second component (LAV detector or vacuum vessel) has an extendable flange (telescopic extension of the tube) with a sliding vacuum seal. These flanges provide a longitudinal opening of approx. 60mm allowing for lateral extraction e.g. a LAV module. These connections include 'O' ring seals to ensure the leak tightness and fixation bolts to connect the parts together.

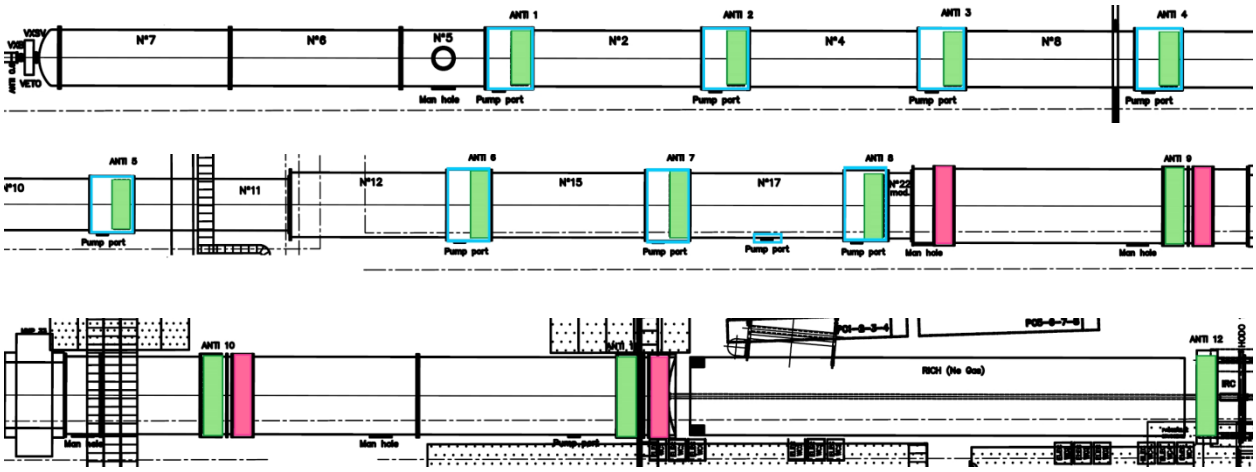

Figure 145: Layout of the decay tube vacuum tank with LAV stations and Straw detectors (in green and pink respectively) Vacuum Requirements

3.2 The Vacuum Pumping System

The beam particles that exit from the beam line at the level of the final collimator are transported towards the main detectors through vacuum. The physics events of interest are kaon decays in the first ~65 metres following the final collimator. Beam particle interactions with the rest gas could cause background events faking the signal of interest. The pressure in the tank must therefore be minimized accordingly. The decay products diverge from the beam axis, the diameter of the decay volume must therefore be sufficiently large to contain the decay products within the detector

acceptance, as well as a number of veto detectors (LAV or Large Angle photon Veto stations). The same beam & decay vacuum tank also houses the four Straw detectors.

3.2.1 Layout

Figure 145 shows the layout of the 480m^3 decay tube volume, where the main new sources of out-gassing are integrated: 12 LAV stations (green) and 4 Straw detectors (pink). The expected gas load is discussed in paragraph 3.2.3. The vacuum tank is made of steel (or stainless steel for the elements around the magnet), painted on the inside, and has a total volume of ~500 m^3.

3.2.2 Vacuum Requirements

A FLUKA simulation has allowed quantifying the effect of interactions of pions, kaons and protons with the residual gas in the decay vacuum tank. The generated beam particle interaction data were passed to a detector simulation where the probability that such an interaction can cause fake triggers, and consequent mis-tagging, was computed. The conclusion of this study is that vacuum in the upstream part of the vacuum tank (that is up to the 1st Straw chamber or a tube length of about 80 m) should be better than 6.10^{-8} mbar to keep the background to less than one fake event per year. With a CEDAR counter positively tagging the ≈6% of kaons in the beam the accidental mis-tagging is minimized as pion and proton interactions with the rest gas are then mostly rejected. Then the vacuum requirement for the first 80 metres can be reduced to a vacuum of 10^{-6} mbar or better.

The vacuum in the downstream parts of the vacuum tank (spectrometer area) must still be 10^{-5} mbar or better to limit Coulomb scattering between beam particles and the particles present in the blue tube.

In order to design a vacuum system capable of achieving the required vacuum levels we must first evaluate the gas load to the system and then take into account available conductance's for the vacuum tube and the pumping ports.

3.2.3 Gas Load (Out-Gassing and Permeation)

Estimates exist for the three main gas loads based on different sets of measurements that were performed between 2006 and 2010. The two main sources of out-gassing and their estimated out-gassing rates are as follows:

- The vacuum vessel [25] itself (blue tube, painted steel): 10^{-2} mbar l/s
- Set of 12 LAV stations [55]: 2-5 10^{-2} mbar l/s

The Straw detectors contribute to the overall gas load of the vacuum tank through permeation of the detector gas through the surface of the Straws. The Straws will either be filled with a CO_2 90%, iC_4H_{10} 5%, CF_4 5% (slow gas) or Ar 70%, CO_2 30% (fast gas) mixture. The proposed vacuum system must be able to meet the vacuum requirements for both Straw mixtures, keeping in mind that the permeation for Argon through the straws is about 1/16 of that for CO_2. As CO_2 has the highest permeation rate of the gases mentioned above, the "slow gas" mixture yields the highest release of

[25] Result from measurements performed in 2006.

CO_2 into the vacuum tank. The Straw detector gas load from permeation (for the "slow gas" mixture) is measured to be 10^{-1} mbar l/s.

Table 25 shows the gas load contributions of the four elements discussed. The last two columns show the fractions of gases which are condensable (non-condensable) at 77K.

Table 25: The gas load, condensable and non-condensable.

	Total gas load [mbar·l·s^{-1}]	Condensable gas load [mbar·l·s^{-1}]	Non condensable gas load [mbar·l·s^{-1}]
Blue tube	10^{-2}	$8 \cdot 10^{-3}$	$2 \cdot 10^{-3}$
LAV	$3 \cdot 10^{-2}$	$2.4 \cdot 10^{-2}$	$6 \cdot 10^{-3}$
Straw (90% CO_2)	$1 \cdot 10^{-1}$	$1 \cdot 10^{-1}$	$3 \cdot 10^{-3}$
Straw (30% CO_2)	$3.4 \cdot 10^{-2}$	$3 \cdot 10^{-2}$	$4 \cdot 10^{-3}$

The gas loads listed above show that Straw trackers dominate the gas load and therefore define the pumping speed needed to match the vacuum requirements. The required pumping speed is then of the order of 150000 l/s.

3.2.4 Conductances

The large cross section of the blue tube, 2.9m^2, or more, gives it nearly zero hydraulic impedance. LAV detectors represent a restriction to a surface of 0.9m^2 or larger. It is assumed that each Straw detector presents an impedance equivalent to 0.2m^2 of free surface. LAVs and Straws can therefore help in establishing a structured pressure profile along the blue tube and in particular, to maintain a pressure difference between the decay region and the Straw region. This throttling may be exploited to restrict the achievement of the more stringent vacuum requirement of 10^{-6}mbar to the decay region instead than to the whole blue tube.

There are twelve manholes of 630mm diameter along the blue tube, corresponding to each LAV detector. Three further pumping ports of 280mm diameter are located at the downstream extremity, in the Straw region. Under the condition that the manholes are easily made available for access and maintenance, all these apertures can be used to install pumping groups.

3.2.5 Technical Pumping Solutions

Two approaches to reach the required vacuum levels are considered, both based on a mixed scheme, with dry pumps and roots pumps for the first level of pumping and then diffusion pumps and/or cryo-pumping for the second level; one solution implements industrial cryo-pumps, while the other foresees diffusion pumps with a custom made cold-trap at liquid nitrogen temperature. It is worth noting that, because of electromagnetic interference with the different tracker detectors in the region, turbo-molecular pumps cannot be used near the blue tube while the detectors are operating.

The main technical principles behind these two cryo-pumping solutions are described here and in the next paragraphs. In both cases, a set of primary dry pumps are used to reach a rough vacuum of ~10^{-3} mbar in the 480m^3 volume. Additional oil diffusion pumps provide part of the required pumping speed. The remainder is ensured by cryo-pumping. In one case, one would use off-the-shelf industrial cryo-pumps in the other cryo-panels, cooled with liquid nitrogen.

3.2.5.1 Industrial Solution

In the first case, a consistent proportion of the pumping speed is provided by off-the-shelf two-stage cryo-pumps, based on the Gifford-MacMahon closed cycle for refrigeration. Helium gas is compressed at warm and expands against a moving piston at cold, providing refrigeration power to a cold-head.

Operating at a temperature below 20K, cryo-pumps feature the highest pumping speed of all industrial vacuum pumps for a given size of pumping port; besides cryo-condensation, acting typically on H_2O, N_2, CO_2 and Ar, cryo-sorption on cold active charcoal provides pumping speed for H_2 and Ne.

Fully automatic operation, including regular regeneration, is ensured by a control unit. A yearly maintenance of the cold-head by specialized personnel is required.

Equipping each cryo-pump with a remotely controlled gate valve allows for the isolation of the pumps from the main vacuum volume for their regeneration during beam time. However, the inclusion of the gate valves will result in a loss of ~30% of pumping efficiency.

3.2.5.2 Custom-Made Solution

At 77K, the liquid nitrogen temperature at normal pressure, CO_2 exhibits a saturation vapour pressure of $2 \cdot 10^{-8}$mbar. With the required vacuum level, it is interesting and cost-effective to apply cryo-pumping only to CO_2 using a large cryo-panel supplied with liquid nitrogen via a small pumping port. Pumping speed is proportional to the cold surface area and can, in principle, be increased at will, limited only by the liquid nitrogen flow needed to maintain a larger panel at the required temperature. In practice, the size of the cold surface is limited by the requirement not to decrease the temperature of the neighbouring LAV detector by more than 20K.

Total capacity is determined by the temperature of the upmost layer of condensed gas, and is therefore limited by the thermal conductance in the condensed layer and the heat load to it. Liquid nitrogen flow, from a Dewar installed in the cavern, would ensure stable temperature conditions by pressure drop, limiting control devices to a flow-regulated warm valve and some temperature sensors.

Laboratory tests on a prototype, scaled down to the NA62 operating conditions, have validated the principle of a liquid nitrogen cooled cryo-panel for CO_2 pumping, confirming that pumping speed, heat losses, capacity and the stability of the system, are sufficient to cope with the required vacuum levels and operation scheme of NA62 [56].

The final choice will strongly depend on the design retained for the Straws and will result from a trade-off between reliability and cost, the latter including equipment, maintenance and operation. A detailed analysis of these aspects is under way. Some preliminary technical considerations must be taken into account to complete this forthcoming analysis.

The throttling effect by the Straws, separating a 10^{-6}mbar region from a 10^{-5}mbar one, is essential in limiting the total required pumping speed and should be validated by numerical calculation or test.

An apparent advantage for the cryo-panel solution, the availability of liquid nitrogen from the neighbouring LKr calorimeter appears, *a posteriori*, impossible to exploit, considering the age and

fragility of the liquid nitrogen feeding system. The purchase and installation of a separate liquid nitrogen storage and transfer system has therefore to be included in the cost estimation. An extended cryo-panel, cooled by a closed-cycle refrigerated cold-head, may represent a valuable alternative to liquid-nitrogen flow. Pumping speed can in principle be increased at will, but it is restricted to a 77K condensable gas species.

In spite of the simplicity of operation, often presented by industry, of off-the-shelf cryo-pumps based on closed-cycle refrigeration, they require regular, costly, maintenance and present an intrinsic fragility, requiring a consistent proportion of installed pumping contingency. The counterpart is a huge, wide-spectrum pumping speed, including air components, where a liquid-nitrogen cooled cryo-panel would ensure a comparably high pumping speed only for CO_2, water and hydrocarbons. In case of an air-leak to the vacuum system, to counterbalance the safety margin ensured by cryo-pumps, an equivalent over-dimensioning of the diffusion-pump system would be necessary in the cryo-panel version. Following this line of reasoning, the cryo-panel solution is particularly adapted to the case of a 90% CO2-filling of permeable Straws, but this may not be optimized to cope with a different filling scheme or a higher permeation by Ar.

There is also the question of contamination, or even saturation, of the cryo-pumping cold surfaces by oil from the diffusion pumps. It is, however, not excluded that special oils with sufficiently low condensation pressures would only have a minimum effect on the cryo-pumping efficiency.

3.3 The MNP33 Magnet

The NA62 experiment uses the same magnet as previously the NA48 experiment, the MNP33. The various parts that make up the magnet have been removed from the beam line and were put to one side. These parts have now been re-assembled in their nominal NA62 position. Some parts of the magnet system – flexible cooling lines for example – will need to be replaced.

This magnet had new coils added at the time of the NA48 experiment. The magnet cross-sectional dimension is 4.4m wide by 4m high and it is 1.3m long. The overall weight of the magnet is 105 tonnes. The main field component goes from top to bottom. Most of the details of the magnet in this section have been extracted from [57].

The iron yoke is made of forged steel. Additional return yokes were added for NA48 with dimensions 3.2m x 0.6m x 1.3m to provide a gap height of 2.40m. The aperture is 3.2m wide.

The helium tank magnet section of NA48 is used again, in a new position for NA62, Figure 146, to cross the magnetic gap and leave an active fiducial region of 2.37m diameter, corresponding to the inner diameter of the old helium tank.

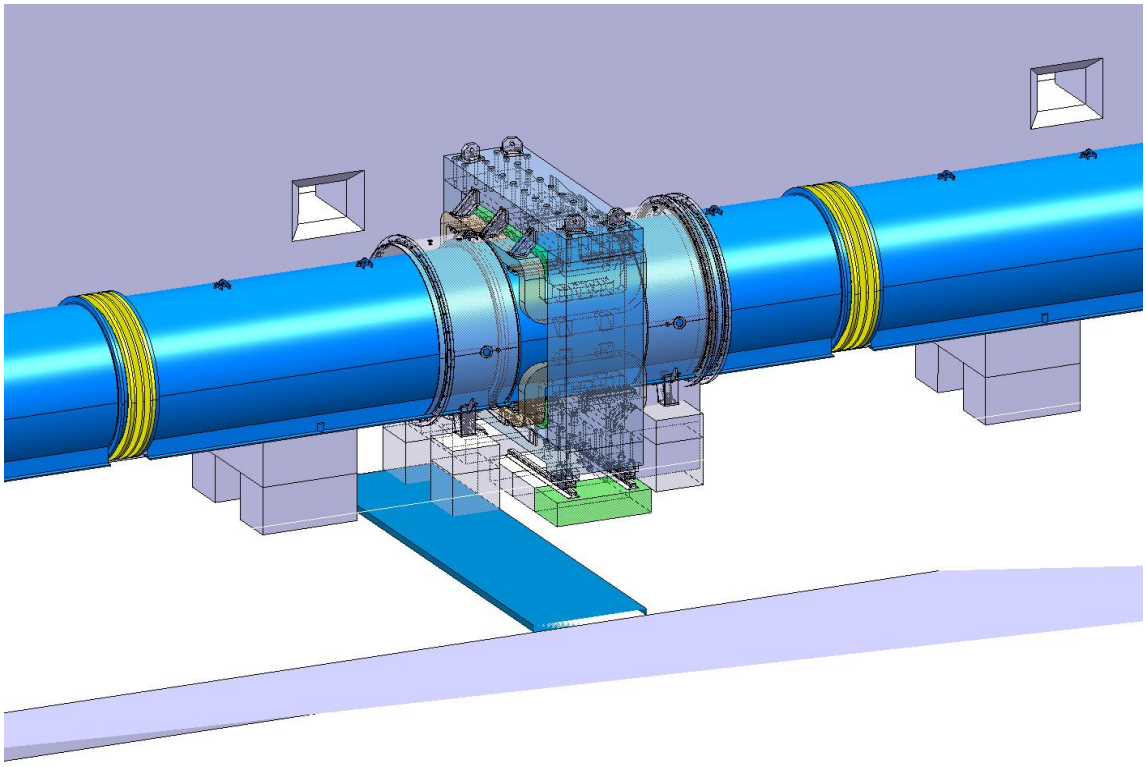

Figure 146: The MNP33 magnet, together with the corresponding section of vacuum vessel (old Helium tank), in its new NA62 position.

3.3.1 The Magnet Coils

The coils are wound with hollow copper conductors of square cross section where cooling water passes through the hollow part. The coils provide an EMF of 0.98 x 10^6 Ampere-turns of which 0.40 x 10^6 are supplied by the original pair and 0.58 x 10^6 by the pair added for NA48. The dissipated power is 3.1 MW. The main parameters of the coils is shown in Table 26.

Table 26: Main parameters of the magnet coils

Coils	Original	Added for NA48	Total
N° of turns	2 x 168	2 x (2 x 120)	816
Nominal current (A)	1200	1200	-
N° of ampere turns (A)	0.4 X 10^6	0.58 X 10^6	0.98 X 10^6
Dissipated power (MW)	1.3	1.8	3.1

Table 27: Main parameters of the conductors

Conductors	Original	Added for NA48
Outer dimensions (mm^2)	14 X 14	15 X 15
Hole diameter (mm)	10	8.5
Net cross-section (mm^2)	117	168
Current density (A/mm^2)	10.2	7.1

3.3.2 The Magnet Power Supply

There are two power supplies that act as master and slave, one for the original and one for the NA48 added coil with a maximum rated current of 1250 and 2500 A respectively. All the coils consist of two parts, upper and lower, electrically connected in series. The newer coil is split into two sections electrically connected in parallel. The main parameters of the conductors is shown in Table 27.

3.3.3 Measured Magnetic Field Parameters

Field calculations and measurements were made in 1995 of the then new magnet configuration. A full field map was measured at the nominal current of 1200A. The measured parameters of the magnet for the nominal operating current are shown in Table 28.

Table 28: Measured parameters of the magnet for the nominal operating current of 1200 A.

Measured field at the centre of the magnet (T)	0.3712
Integrated field (Tm)	0.858
Maximum field close to the poles (T)	<0.9
Maximum field gradient (Tm)	0.9
Transverse momentum (MeV/c)	257.4
Effective magnetic length (m)	2.31

3.3.4 Dismounting and Re-assembly of the Magnet

The magnet has now been disassembled from its previous position in NA48. The following drawing, Figure 147, was made to assist with the disassembly

The magnet has now been reassembled in its new position for NA62. A drawing, Figure 148, was made to show the new position.

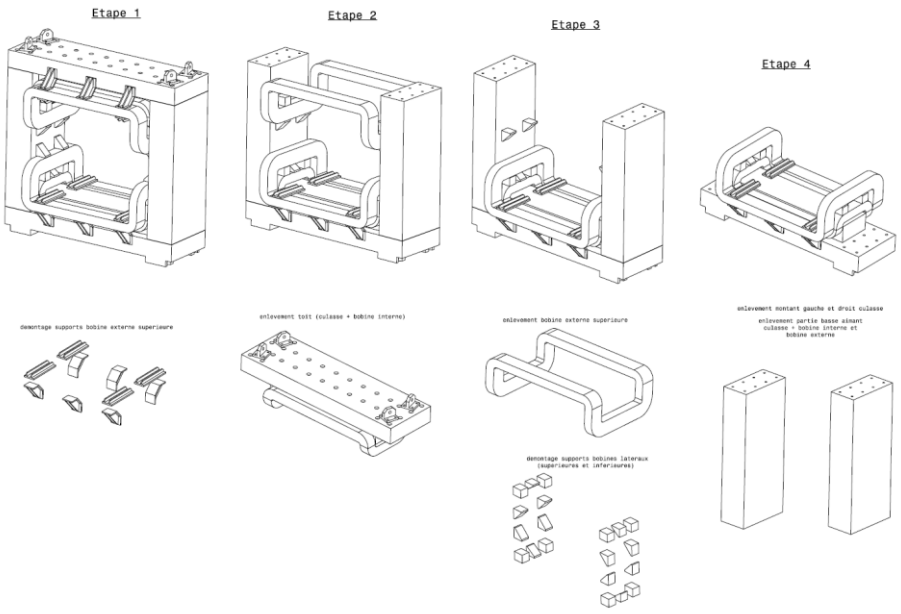

Figure 147: The magnet broken down into its parts

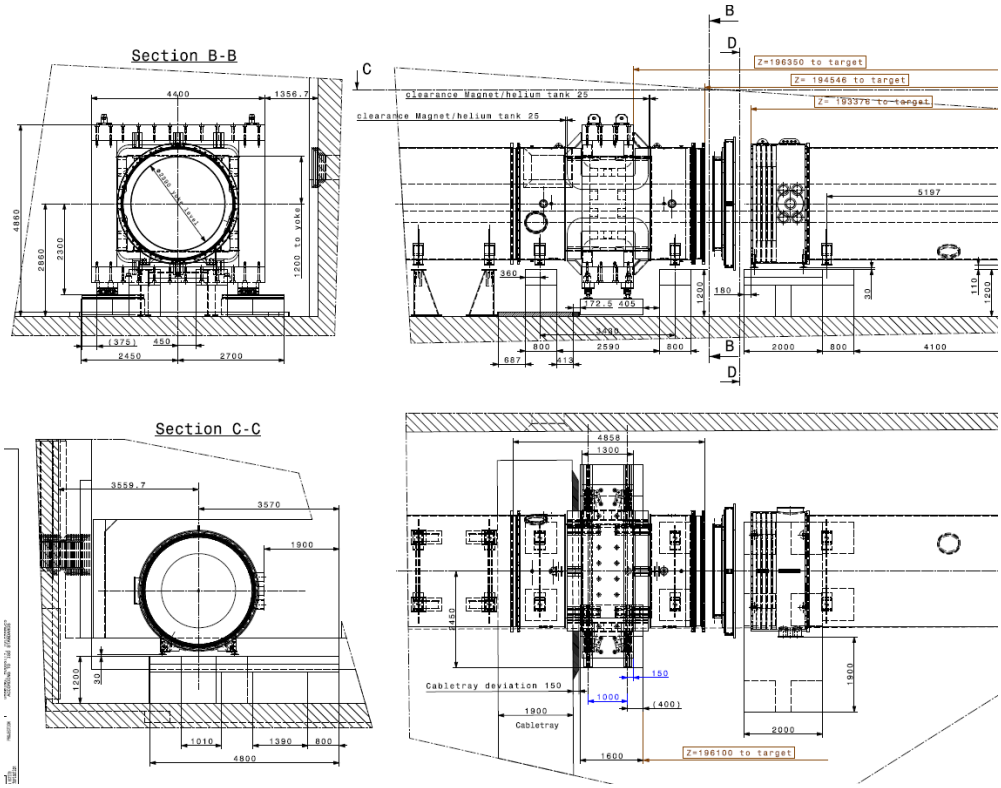

Figure 148: Drawing of the new position of the MNP33 magnet for NA62.

3.4 The Photon Veto Detectors

3.4.1 Physics Requirements for Photon Veto System

Photon vetoes are required to suppress the dominant background originating from the decay $K^+ \to \pi^+\pi^0$ (BR=20.7%) to the specified level. The average inefficiency for the rejection of the π^0 should be smaller than 10^{-8}. The photon vetoes need to have hermetic geometrical coverage up to 50 mr for the photons originating from the kaon decays occurring in the decay region (from 5 to 65m after the final collimator). With such a configuration, only about 0.2 % of the $K^+ \to \pi^+\pi^0$ events have one photon from the π^0 left undetected.

The geometry of the experiment suggests partitioning of the detector into three different angular regions, each instrumented by three different detector technologies:

- Large Angle Vetoes (LAV), covering the angular region between 8.5 mr and 50 mr, distributed along the decay volume and spaced by 6m in the upstream region and by 12m downstream, according to the layout in Figure 149.
- The NA48 Liquid krypton calorimeter (LKR), covering angles between 1 and 8.5 mr (see section 3.4.4 on page 225)
- Small angle vetoes covering the region down to zero degrees (SAC) and the zone around the inner radius of the LKR (IRC) calorimeter. These will have suitable overlap in the angular acceptance to cover the beam pipe and an inner radius smaller than that of the beam pipe (see section 3.4.5 on page 228).

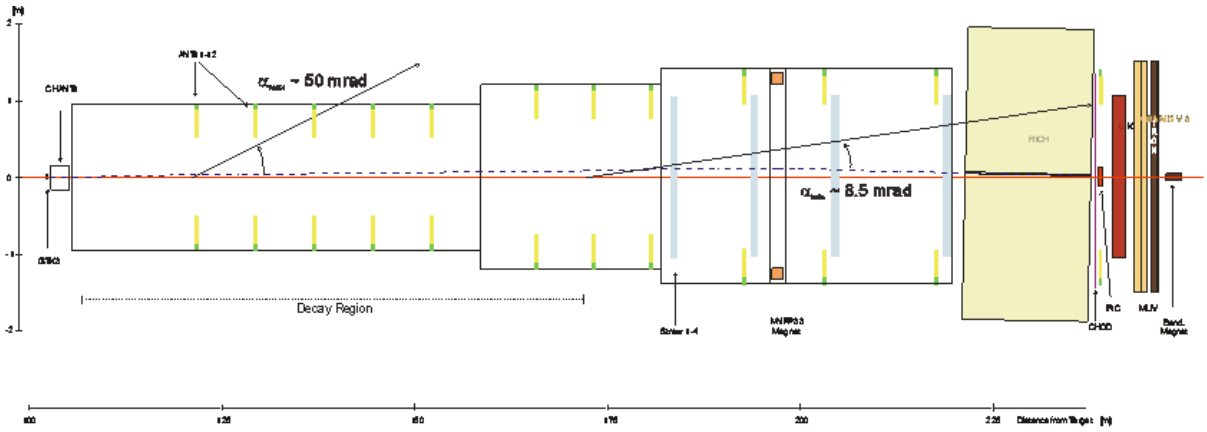

Figure 149 Longitudinal View of the LAV detector (yellow).

The kinematics of $K^+ \to \pi^+\pi^0$ decay in the NA62 decay volume are such that, with a cut on the momentum of the charged pion, only three possible configurations are present:

- both photons from the π^0 hitting the forward calorimeters with a total energy larger than 20 GeV;
- one photon in the forward calorimeters and the other one in the Large Angle Vetoes;
- one photon in the forward calorimeters and the other one lost at angles larger than 50 mr.

This last combination occurs only in 0.2 % of the decays. In order to achieve the required π^0 rejection, all photon veto detectors must have an inefficiency lower than 10^{-4}. With this requirement, the major contribution to the global inefficiency comes from the 0.2 % of events where only one photon is detected.

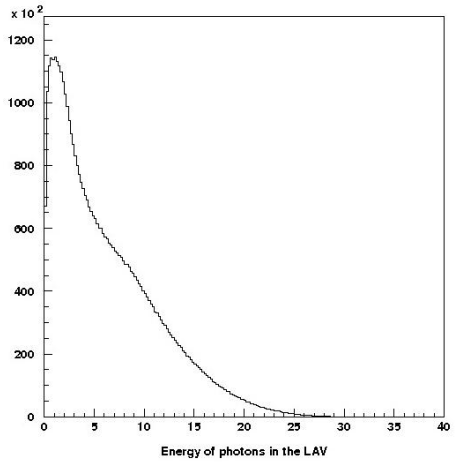

Figure 150 Photon energy distribution in the LAV.

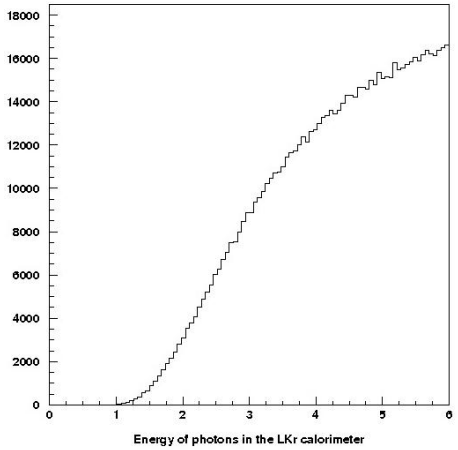

Figure 152 Minimum photon energy in the LKR.

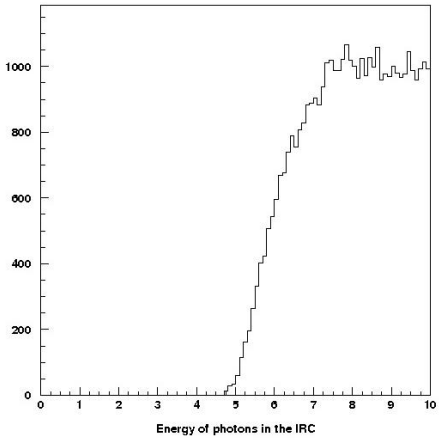

Figure 151 Minimum photon energy in the IRC

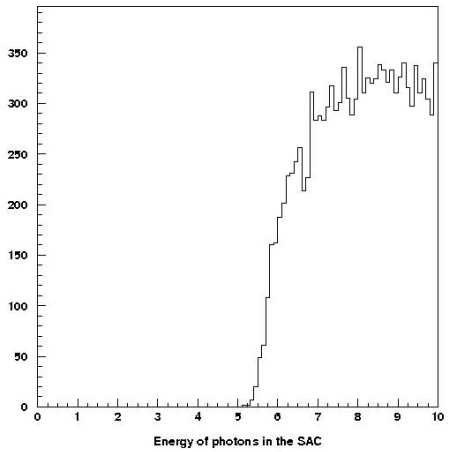

Figure 153 Minimum photon energy in the SAC

A simulation of the decay kinematics and of the geometrical acceptance has been made to compute the average photon detection inefficiency. It was done using an estimation of the inefficiency of each detector based on experimental measurements from test beams and with the NA48 detector.

The average π^0 inefficiency is computed weighting each photon with the inefficiency of the hit detector at the photon energy.

The results shown refer to 10^7 $K^+ \rightarrow \pi^+\pi^0$ decays generated with a cut on the charged pion momentum between 15 and 35 GeV. The photon energy in the Large Angle Vetoes is distributed between below 10 MeV and 30 GeV (Figure 150). The minimum photon energy in the LKR calorimeter is 1 GeV (Figure 152), while in the IRC and SAC it is around GeV (Figure 151 and Figure 153).

The average π^0 inefficiency is reported in Table 29, with and without imposing a cut on the π^+ momentum between 15 and 35 GeV. The fraction of events which mainly contributes to the average inefficiency is also shown in Table 29, imposing the cut on the π^+ momentum. This fraction is defined as the ratio between the number of the events contributing to the inefficiency and the total number of the accepted events (in this example 2747513).

Table 29 Fraction of inefficient events

Destination of the two photons	Number of events contributing to the inefficiency	Fraction (%)
Both photons undetected	0	0.00
One photon undetected, the other in the LAV	0	0.00
One photon undetected, the other in LKR,IRC,SAC	5312	0.19
One photon in the LAV, the other in LKR,IRC,SAC	510534	18.58
Both photons in LKR, IRC, SAC	2231667	81.22
Average inefficiency without π^+ momentum cut	$8.4*10^{-8}$	
Average inefficiency with π^+ momentum cut	$1.6*10^{-8}$	

An in-depth look at inefficient events shows that they are characterized by one low energy photon that is either outside the acceptance of the LAV or inside it, but with very low energy and high inefficiency. The other photon is mainly hitting the LKR calorimeter, with an energy larger than 35 GeV, or the first IRC, with an energy larger than 60 GeV. It is then mandatory to have an inefficiency of the LKR lower than 10^{-5} at those energies to reduce the contribution from these events to the average π^0 rejection inefficiency.

3.4.2 Photon-Veto detector Efficiency Studies

For the large-angle photon veto detectors, a comprehensive research and development program was carried out in Frascati in 2007 in order to guide the choice of technologies to be used. Three possible technologies for the LAV system were considered. The first design, originally proposed for use in the (now cancelled) CKM experiment at Fermilab, featured a modular structure consisting of alternate layers of 1-mm thick lead plates and 5-mm thick scintillating tiles, readout by wavelength-shifting fibers. A second option, adopted for the construction of the electromagnetic calorimeter for the KLOE experiment, featured a structure consisting of 1-mm diameter scintillating fibers embedded between 0.5-mm thick lead foils. The third solution made use of lead-glass crystals

obtained from the dismantled electromagnetic calorimeter of the OPAL experiment [58]. Prototype detectors based on each of the three technologies have been obtained or constructed. Experimental tests conducted with the electron beam at the Frascati BTF demonstrated that all three technologies are suitable for use in the experiment [59]. The results obtained for the inefficiency are summarized in Figure 154. The efficiency for detection of low-energy electrons is similar for all three technologies. For these initial tests at the BTF, the results for the lead-glass solution were obtained with a stack of four OPAL modules. Since there is a significant practical advantage to adapting existing hardware for the NA62 LAV system, the performance of the OPAL modules was further investigated. In early 2008, a 25-block staggered array of lead glass crystals was tested at the Frascati BTF. In particular, the efficiency of a lead-glass array for the detection of 470-MeV electrons was studied as a function of the impact position. These studies essentially confirmed the findings presented in Figure 154 [60]. We therefore decided to base the LAV design on the OPAL detectors, as described in the following sections.

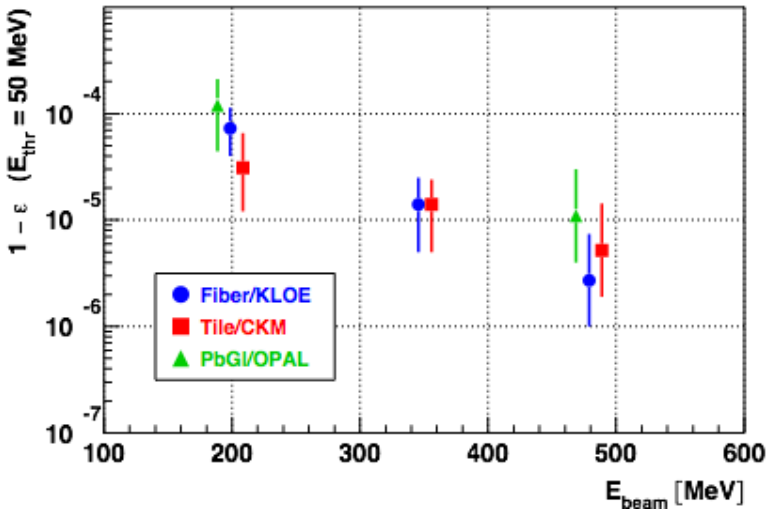

Figure 154 Measured efficiencies vs energy for three types of photon veto detector.

Dedicated studies of the photon detection efficiency of the NA48 LKR calorimeter were performed using two data samples from the experiment [12], as noted below.

In 2004, $K \rightarrow \pi\pi^0$ decay events were collected during a dedicated run with the NA48 apparatus, using as trigger criteria a drift chamber multiplicity and hodoscope signal requirement to select single track events, as well as a muon veto. $K \rightarrow \pi\pi^0$ decays were identified offline using missing mass cuts. Various other cuts were made to eliminate backgrounds. Among the clusters not associated with the track, the one with the least energy was assumed to be the lower energy photon from the π^0 and used to tag the other photon. Among 730k events, the other photon was not detected in five cases. Allowing for uncertainties in the background rate and other systematics, the inefficiency of the LKR for the detection of photons with $E > 10$ GeV was found to be less than 9×10^{-6} at 90% CL.

An alternative measurement of the LKR calorimeter efficiency was performed during a dedicated run in 2006. In this run, a 25-GeV/c secondary electron beam was produced using protons from the SPS,

and traversed the NA48 apparatus, producing bremsstrahlung photons. For events in which the photon was emitted from an interaction with the 1.3% X_0 of material upstream of the spectrometer magnet, the initial beam direction defined the impact point of the photon on the calorimeter, while the beam itself was deflected horizontally by the magnet, so that it arrived at the calorimeter with a displacement of about 12 cm from the impact point of the photon. The detection efficiency measured for photons with E > 10 GeV was consistent with the result of the 2004 study using $K \rightarrow \pi\pi^0$ decays.

The baseline design calls for the small-angle calorimeters to be constructed with shashlyk technology. During the 2006 test run, a prototype shashlyk SAC module was constructed and tested using the 25-GeV electron beam [12]. The active part of the detector was assembled from layers of 20.5 × 20.5 cm^2, 1.5-mm thick lead plates and 1.5-mm thick scintillator plates, for a total of 16 radiation lengths. The light was read out with wavelength shifting fibers coupled to four photomultiplier tubes. During the tests the detector was placed between the two planes of the NA48 scintillator hodoscope, which helped define the trigger. The electron beam was deflected to point at the SAC module; the beam position was varied to collect information on the variation of response as a function of impact position. The offline selection also required a full-momentum track in the NA48 spectrometer. Of 36k total tracks passing the offline cuts, none were found without an associated signal in the SAC prototype. The corresponding upper limit on the inefficiency for 25-GeV electrons is 6.4×10^{-5} at 90% CL. A refined analysis with the whole statistics gave a value for the inefficiency of $(2.9\pm0.3)\times10^{-5}$.

The results from these tests have been used to compute the average inefficiency as described in 3.4.1.

3.4.3 The Large Angle Detectors (LAV)

3.4.3.1 General Layout of the LAV

The Large Angle Veto is to be composed of 12 stations situated between 120 and 240 m along Z (Figure 149). The first eleven stations are part of the vacuum decay tube, while the last one is located outside the vacuum tank.

The basic building blocks of these detectors are lead glass crystals with attached photomultipliers (PMT) from the former OPAL electromagnetic calorimeter (see 3.4.3.2). Four crystal detectors (lead glass crystals + PMTs) are mounted on a common support structure forming an azimuth segment. Inside the vacuum tube the azimuth segments are assembled to form a complete ring of lead glass blocks. Each LAV stations is made up of 4 or 5 rings, which are staggered in azimuth providing complete hermeticity of at least three blocks in the longitudinal direction. All the counters in a complete ring of lead glass lie in one plane that is perpendicular to the beam line, with all the PMTs on the outside of the ring.

3.4.3.2 The Basic Crystal Detectors

3.4.3.2.1 Mechanical Assembly of the Crystals

Each lead glass block from the former OPAL experiment has the shape of a truncated prism of Schott SF57 lead glass [58], whose properties are shown in Table 30. The blocks are available with different shapes and dimensions (with minimal variations between different types). Figure 156 and Table 31 describe these differences and all dimensions which are relevant for the design of the modules. The block length (BL in the figure) is always 370 mm. One of the (almost) square faces of the lead glass has a 1 cm-thick steel flange glued to it (Figure 155). This flange has four threaded holes for fixing the counter to the support bracket, one for the connection of a calibration optical fibre, and a central large hole for the passage of a cylindrical light guide for light collection. The light guide is a cylinder of SF57 lead glass with a diameter of 73 mm and a height of 60 mm. It is glued to the lead glass block and, at the other end, to a Hamamatsu R2238 photomultiplier. An external mu-metal shield, enclosing the guide and the PM, is glued to the steel flange.

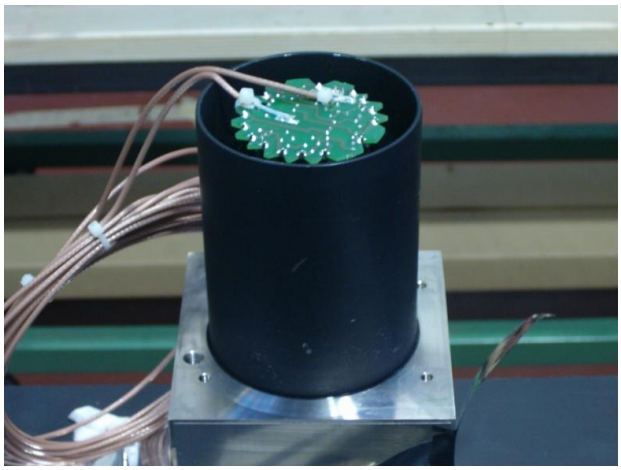

Figure 155 The steel flange

Table 30 Physical characteristics of the SF57 lead glass

Density	5.51 g/cm^3
Radiation length	1.54 cm - 8.49 g/cm^2
Moliere radius	2.61 cm
Interaction length	20.6 cm - 115 g/cm^2

3.4.3.2.2 Photon Detector Assembly

The photomultiplier pins are directly soldered to the standard Hamamatsu divider for the R2238 tube mounted on a round printed circuit board. HV and signal cables are made of Teflon insulated RG316 cable, with a length varying from 4.5 to 6.2 m (depending on the ring to be built). The HV ground connection on the divider is decoupled from the signal ground by a 10 Kohm resistor in series. On the side opposite to the PMT, the signal cable has a mini coax connector (SOURIAU type RMDXK10D28), while the HV cable has a Burndy pin (type RM24M-9K) on the voltage wire and a faston connection on the ground.

Block type	W2 (mm)	mumetal length (mm)
15	95.461	110
14	95.598	110
13	95.808	110
12	96.086	110
11	96.426	110
10	96.820	115
9	97.261	125
8	97.740	145
7	98.250	145

Table 31 Types and dimensions of the blocks

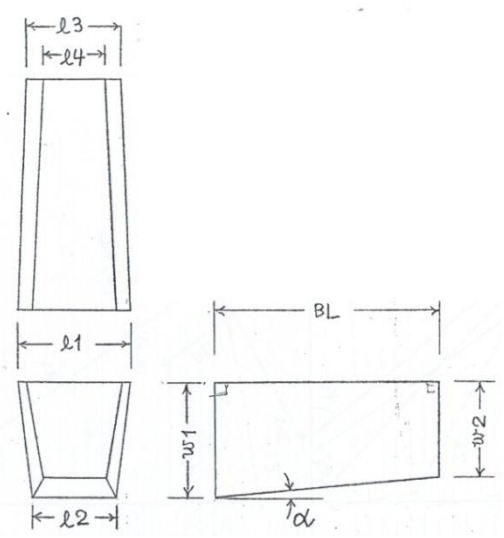

Figure 156 Shape and dimensions of the block

3.4.3.2.3 Validation and Selection of Existing Crystal Detectors

The final number of available OPAL lead-glass blocks after the validation test (see 3.4.3.3) is expected to be close to the quantity of blocks (2496) required for NA62. A possible optimisation of the assignment of blocks to the different rings is at the moment under study. Table 32 lists the provisional assignments of block types to each detector station.

Table 32 Provisional assignments of block types to the detector modules and layers.

Module	Number of Layers	Crystals/layer	Block types used
LAV 1	5	32	15
LAV 2	5	32	14
LAV 3	5	32	4*11 + 1*14
LAV 4	5	32	4*13 + 1*14
LAV 5	5	32	15
LAV 6	5	48	13
LAV 7	5	48	14
LAV 8	5	48	2*11 + 3*15
LAV 9	4	60	12
LAV 10	4	60	10
LAV 11	4	60	9
LAV 12	4	64	2*7+2*8

3.4.3.3 Testing and Characterization of the Crystal Detector Blocks

Before re-using the OPAL crystal detectors , a number of tests are performed to validate each crystal detector block :

- Operational checks of the PMT (HV, signal, and light tightness)
- A measurement of the PMT amplification (gain) as function of anode voltage
- A measurement of the photoelectron yield in the crystal with the connected PMT

3.4.3.3.1 The Lead Glass Test Station

The gain and the photoelectron yield of all the PMTs will be measured in conjunction with their lead-glass crystals using the Lead Glass Test Station (LGTS). The LGTS (see Figure 157) consists of a light-tight box that can accommodate twenty lead glass detectors (LG), twelve for testing and eight for triggering on cosmic ray particles. A bundle of 12 (+2 spares) optical fibres distributes light pulses from a LED. The whole system is remotely controlled by a personal computer and signals are acquired by a VME ADC module (CAEN V792, 12 bit 0-400 pC).

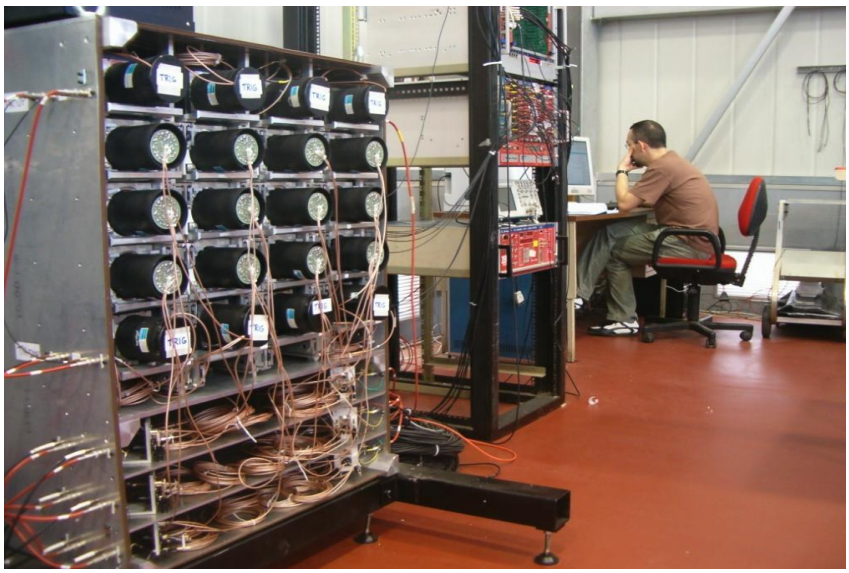

Figure 157 The Lead Glass Test Station.

Optical fibres can inject light pulses using a small inspection window present in the wrapping. The external lateral sides house the High Voltage (HV) and signal connectors. An LED system has been designed for the LGTS to distribute the light generated by one LED to all 12 blocks under test. The temperature of the LED is stabilized at a level of about 0.1 °C by means of a Peltier cell. The amplitude of the light pulses can be varied by more than a factor of 10 and is software controlled via an RS-232 interface. Pulses can be triggered either by an internal clock or by an external source, the typical trigger rate is 100 Hz. The LGTS is used to measure the gain curve of each PMT and to

equalize the responses. Beforehand, checks are made to see that the block can stand high voltage without discharging and that the current from the PMT, at a fixed HV value, is within specifications.

3.4.3.3.2 Gain Measurement

The gain G of a PMT can be measured using several light pulses having the same amplitude. If R is the measured response, then the mean value <R> and the standard deviation σ_R are related to the gain by the relation:

$$\sigma_R^2 = \langle R \rangle \cdot G \cdot \left(1 + \delta_{SER}^2\right) \qquad (1)$$

where δ_{SER} is the single electron response (SER) fluctuation contribution. This term depends on the gain value and the PMT structure, and can be non negligible, of order 20-30%. A good approximation of its value is obtained by the relation:

$$\delta_{SER} = G^{-1/N} \cdot \left[\frac{\left(\prod k_i\right)^{1/N}}{k_1} \right]^{\alpha}$$

where

N is the number of dynodes

$k_1 ... k_N$ is the fraction of the HV applied to the i-th dynode

α is a parameter usually comprised between 0.6 and 0.7, which is kept here at 0.6.

After fixing a HV value, <R> and σ_R are measured for a sample of typically 10^4 light pulses generated with the same amplitude over a relatively short time period (100 s) . Then the measurement is repeated for a different value of the light pulse amplitude. G is extracted by applying the obtained data with a least square fit to equation (1).

The whole procedure is then repeated varying the HV between 1150 V and 1350 V (in 50 V steps). The gain curve G(V) obtained in this way is then obtained by analytical interpolation of the measured points.

3.4.3.3.3 Photo-electron yield

The photoelectron yield (PeY) for each block is defined as the number of photoelectrons released by the detector per MeV deposited by a MIP. It depends mainly on the photocathode quantum efficiency, on the photon collection efficiency (crystal to PMT) and on the photon yield of the crystal. Since the crystal and the PMT are glued together, this number is the main figure of merit for a detector block. At the beginning of 2008 the PeY was measured for 40 randomly chosen detector blocks. Using a tracker it was possible to determine the detector response as a function of distance from the photocathode of a straight MIP crossing the crystal, and it was shown to be flat.

The total charge measured when a MIP crosses the detector is :

$$Q = G \cdot N_{pe} = G \cdot E \cdot PeY$$

where G is the PMT gain and E is the energy released by the MIP. One can evaluate for each block the PeY from the measured response Q when a cosmic ray crosses the detector and the known gain G given by the gain curve G(V). The average energy E released by a MIP impinging orthogonally on the crystal is about 77 MeV. From these data one can evaluate for each crystal the voltage V to get a gain $G=10^6$. Then a cosmic run is performed, by acquiring typically 10^4 events at 0.5 Hz trigger rate. An automatic procedure fits the cosmic peak to find Q and extract the PeY. The whole procedure is repeated then with a different value of the gain ($9x10^5$) as an independent cross check of the PeY value, which, of course, does not depend on the gain.

3.4.3.3.4 Response Equalization

Given the gain curve G(V) and the PeY for a block one can estimate for each lead-glass counter the working point V_{eq} , i.e. the voltage needed to get equal response for MIP's. After this first estimate the response equalization is cross-checked and fine tuned in a new cosmic run by fixing an arbitrary value of Q = 4.5 pC and adjusting the PMT voltage according to the estimation. If the equalization is not satisfactory a small change along the G(V) curve is done to improve it and an additional run is done to check if good equalization is achieved. Results for PeY's of about 200 blocks measured for the LAV-1 and for the spread of the equalized responses are shown in Figure 158 and Figure 159.

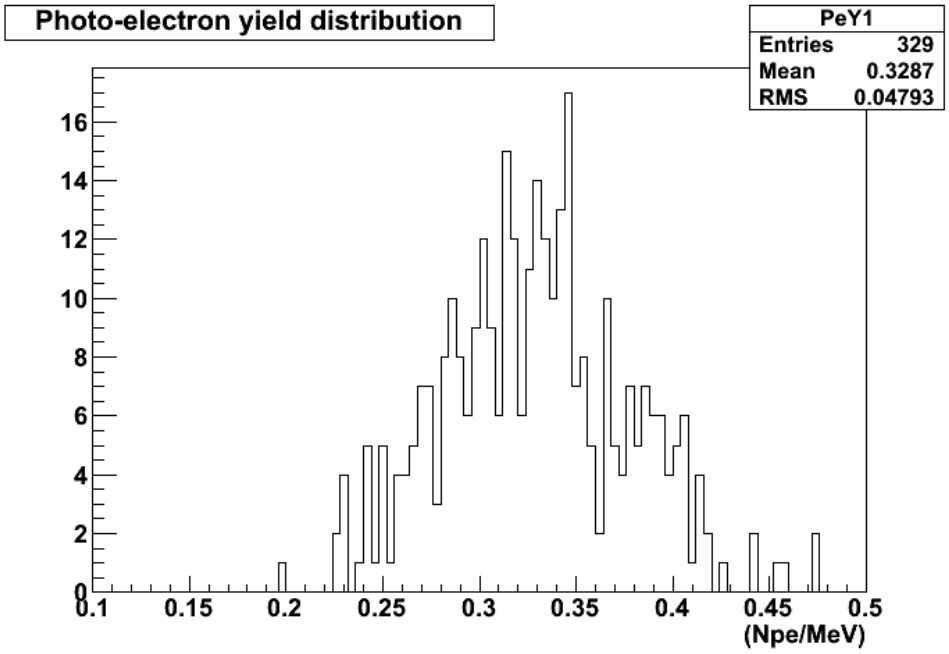

Figure 158 Photoelectron Yield (PeY) for a typical set of blocks

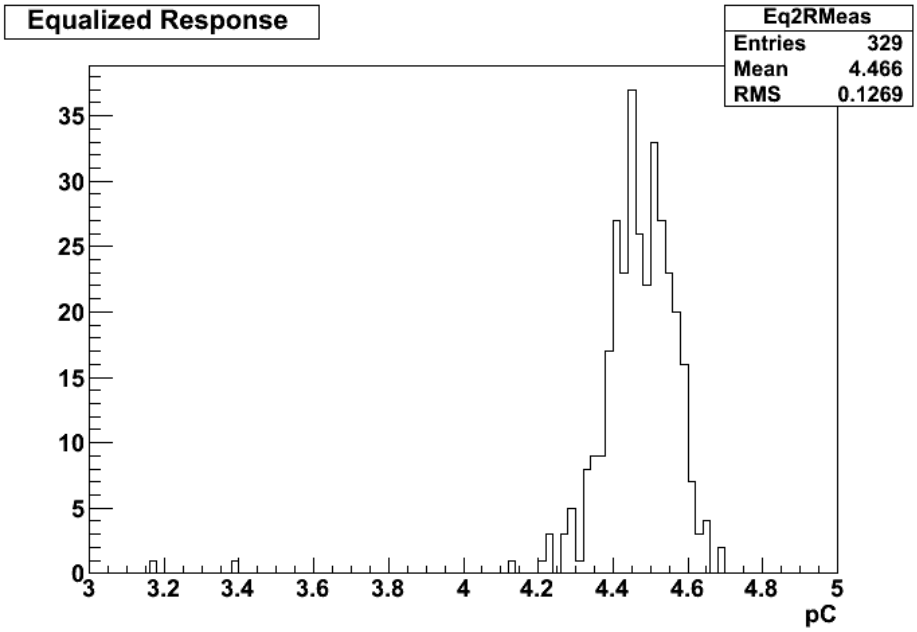

Figure 159 Spread of the equalized response for a typical set of blocks

3.4.3.4 Operation in Vacuum

The interaction of the beam with residual gas in the decay tube can produce a photon-free background to $K^+ \rightarrow \pi^+ \nu\nu$ that will become important if the vacuum is worse than 10^{-6} mbar in the fiducial region and worse than 10^{-5} mbar the downstream region of the spectrometer. The Large Angle Veto system must be able to operate inside the vacuum environment. The choice of having both the lead glass blocks and the PMTs inside the vacuum simplifies the design of the mechanics. In this context, it is important to note that the PMTs when used in OPAL were never operated in vacuum.

To validate the technological choices and the materials used a series of outgassing measurements were done. Heating tests for the voltage divider have been made to confirm the vacuum operation of the PMTs.

3.4.3.4.1 PMT Heating Test and Measurements

Two independent checks were done in Frascati and Pisa using vacuum vessels with infrared-transparent windows and a thermocamera. The results are described in the following.

The Hamamatsu R2238 PMT voltage divider has a maximum resistance R1=480 KΩ between the cathode and the first dynode which dissipates ~100 mW during operation. When the PMT is operated in air the power is dissipated through convection mechanism using air as medium. Once in vacuum the only possibility is through radiation and conduction. Figure 160 shows a picture of the PMT base made at a pressure of few 10^{-4} mbar using an infrared termocamera.

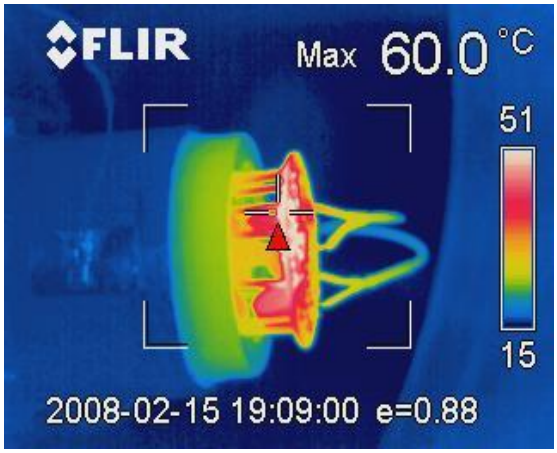

Figure 160 Measurement in vacuum of the temperature of the divider connected to the PM

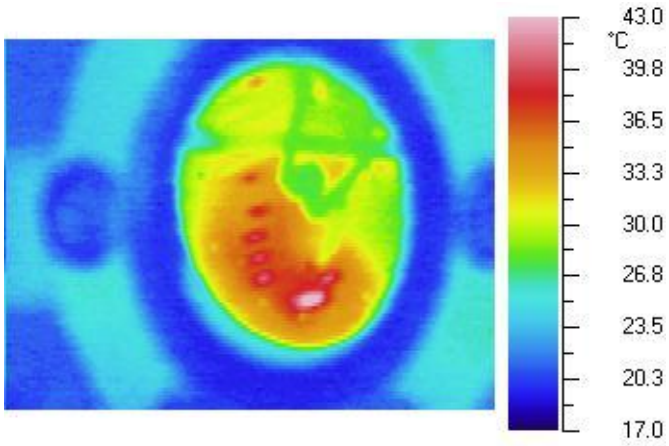

Figure 161 Measurement of the divider temperature in vacuum

The base reaches a maximum temperature of 60° C at the R1 resistor after more than 4 hours of PMT operation in vacuum. The picture shows that in addition to the radiative cooling, a conduction mechanism is active through the base connections to the PMT.

A similar test was performed in Pisa using the divider alone (without connections to the PMT) with an applied voltage of 1500 V in vacuum. Figure 161 and shows the results: the thermo camera was looking at the sample through a CaF_2 window and a calibration correction should be applied to the displayed temperature. With this correction, the temperature of the hottest point was found to be 70 ^{0}C, in agreement with the result from the test in Frascati.

To monitor the divider temperatures on a long time scale, PT100 temperature probes have been installed on R1 in more than 30 PMT dividers in the LAV 1 module. The temperature will be monitored during months of operation to spot unexpected behaviors.

3.4.3.4.2 Outgassing Measurements

Two methods were used to assess the outgassing rates:

- measure the vacuum pressure limit at known pumping speed
- measure the pressure rise time without pumping.

Details on how the measurements have been performed can be found in [61].

The use of two redundant measurements has allowed us to cross-check the results and to evaluate some systematic effects. The following notations will be used:

- Q_S: outgassing rate of the sample under investigation (mbar*l/s);
- Q: specific outgassing rate, i.e., the outgassing per cm^2 ($Q=Q_S=A$).

The outgassing rates for individual parts of the lead-glass detectors have been measured in a small vacuum chamber. The samples were pumped until the pressure became stable. After that the limit

pressure and the pressure rise time were measured. Details on the experimental setup can be found in [15]. There is no evidence of significant contributions to the outgassing from the wrapping, the glue and the PMTs. The observed outgassing seems to arise mainly from water and air absorbed by the crystals. A conservative estimate of the single block outgassing rate is $Q_{BLOCK} = (7\pm10)\cdot10^{-6}$ mbar l/s after two weeks of pumping time.

The vessel of the first LAV module was tested for vacuum leaks using a helium leak detector with a limit sensitivity of $5\cdot10^{-11}$ mbar l/s. A limit of $< 1\cdot10^{-10}$ mbar l/s can be set on the leak rate for the vessel, 4 orders of magnitude lower with respect to normal operation pressure. To measure the outgassing rate the vessel was pumped for 96 hours reaching an ultimate pressure of $P_{lim} = (1.4\pm0.3)\cdot10^{-6}$ mbar. The result yields a total outgassing rate for the LAV 1 module of $Q_S = (1.1 \pm 0.2)\cdot10^{-5}$ mbar l/s, from which the specific outgassing rate of the vessel can be computed to be $Q = (0.9\pm0.2)\cdot10^{-10}$ mbar l/s/cm^2.

The outgassing of the LAV 1 vessel can be compared with preliminary estimates obtained at CERN for the NA48 vacuum tank, which has a specific outgassing rate $Q \sim 1$ to $3\cdot10^{-9}$ mbar l/s·cm^2 after 15 days of pumping time, while the LAV 1 vessel reaches $Q \sim 1\cdot10^{-10}$ mbar*l/s·cm^2 after only 4 days. The difference is most probably due to the fact that the inside walls of the NA48 vacuum tank are painted.

As a final consideration, the contributions to the overall outgassing rate from each of the components of the LAV 1 prototype are listed in Table 33. The estimates are obtained by scaling the measurements for each component to the amount used in the assembly of the entire LAV 1. For the outgassing of the blocks, both the worst case ($Q_S = (2\cdot10^{-5})$mbar l/s) and the average ($Q_S = (1\cdot10^{-5})$ mbar l/s) outgassing rates are listed.

Table 33 Estimated outgassing rate (mbar l/s) of LAV 1 with 160 lead-glass blocks

	Outgassing rate	*Days of pumping*
LG detectors (worst case)	$(2.9\pm2.0)\cdot10^{-3}$	15
LG detectors (average)	$(1.5\pm1.0)\cdot10^{-3}$	15
Monitoring system	$(1.0\pm0.2)\cdot10^{-5}$	15
Tyvek wrapping	$(3.1\pm0.6)\cdot10^{-5}$	5
LAV 1 vessel	$(1.1\pm0.2)\cdot10^{-5}$	4
Total LAV 1 (average)	$(1.5\pm1.0)\cdot10^{-3}$	15

Table 33 shows that the dominant contribution to the LAV 1 outgassing rate comes from the blocks. The contribution to the outgassing rate from the hardware and aluminum plates used to fix the blocks to the wall of the vessel is not included in the computation because it is considered to be small compared to the contribution from the blocks. The estimated outgassing rate of LAV 1 is $(1\text{-}2)\cdot10^{-3}$mbar l/s. A first measurement of the outgassing rate of the fully assembled LAV 1, based only on 7 days of pumping time, produced a value of $< 4\cdot10^{-3}$mbar l/s in reasonable agreement with the prediction based on single block measurements.

Using LAV 1 value a preliminary estimate of the outgassing rate for all the Large Angle Veto stations in vacuum, consisting of 11 rings for a total of about 2250 lead-glass blocks, is $(2\text{-}5)\cdot10^{-2}$ mbar l/s. This value is very similar to that obtained at CERN for the entire NA48 vacuum tank, $(1\text{-}5)\cdot10^{-2}$ mbar l/s.

3.4.3.5 Light Calibration

The LAV system must be accurately monitored given the demanding efficiency requirements. For instance, a 5-10% drop in efficiency in 1% of the blocks could produce effects of the same order of magnitude as the maximum allowed inefficiency of the system. In practice one will use software and hardware thresholds to flag an event, and thus a precise monitoring of the detector response is needed. The light pulse calibration system will allow checking of the response of each block to light pulses at any time in a simple and effective way. It will also allow checking of the stability of the PMT gain by repeating the procedure described in 3.4.3.3. Furthermore, when combined with the measurement of the response to beam halo muons, it will check the stability of the PeY (as defined in 3.4.3.3.3). It will also be useful to inter-calibrate the timings of the blocks.

The layout of the system is shown in Figure 162. It is composed of a light distribution system, based on silica fibres with different core diameters (see 3.4.3.6.7) and a remotely controlled LED pulser.

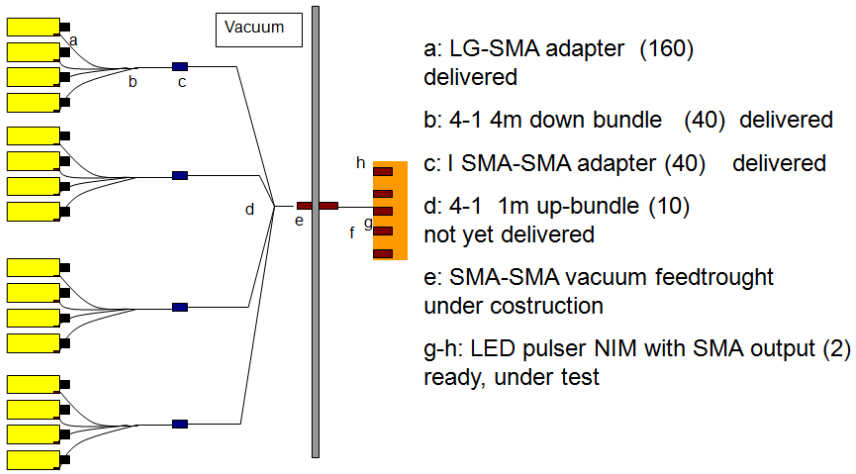

Figure 162 Layout of the fiber calibration system

Each LED pulser module can pulse independently 5 LEDs. Each LED is independently temperature stabilized at the level of 0.1 °C by means of a Peltier cell. The working condition can be software adjusted. The module has TCP/IP communication protocol and has both Ethernet and RS-232 connections. The pulser modules have been engineered both as standard NIM crate modules and as standalone devices.

Pulses can be generated via internal or external triggers at a maximum rate of about 500 Hz. They can be externally triggered in the inter bunch time (as done for the LKR system) for fast calibration purposes, and/or gain measurements.

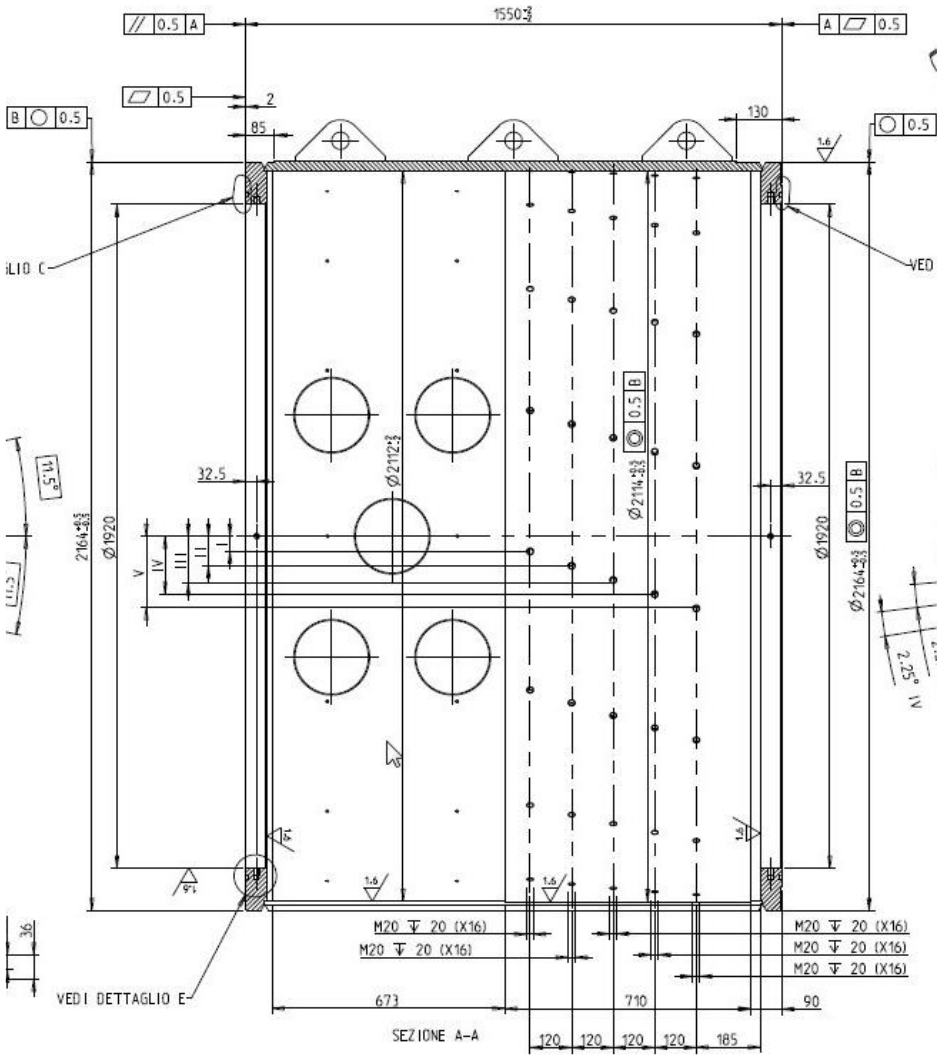

Figure 163 Lateral view of the LAV 1 vacuum vessel

3.4.3.6 Mechanics

3.4.3.6.1 Support Tubes

The detector crystals of the NA62 Large Angle Vetoes are contained inside a flanged vacuum vessel that supports the full weight. The overall dimensions of the different LAV vessels are listed in Table 34 and their longitudinal positions (the downstream ends) are listed in Table 35.

Table 34 Overall dimensions of the LAV modules

Station	Ext. Vessel Ø [mm]	Int. Vessel Ø [mm]	Int. Lead Glass Ø [mm]	Length [mm]	Total weight [kg]
1-5	2164	2114	1064	1550	~10000
6-8	2650	2590	1544	1550	~15000
9-11	TBD	TBD	1944	1430	TBD
12	TBD	TBD	2144	800	TBD

Table 35 Longitudinal positions of the LAV modules

Module	1	2	3	4	5	6
Longitudinal position [m]	122.220	129.830	137.440	145.050	152.660	166.170
Module	7	8	9	10	11	12
Longitudinal position [m]	173.680	181.190	193.391	203.909	218.530	239.645

There are several requirements that must be taken into account for the design of the vessel: the accuracy of the crystal position; the compatibility with the existing vacuum tubes; the vacuum tightness, cleanliness, outgassing, and the pumping & vacuum instrumentation; the cabling & feed-through flanges; the handling attachments; the mechanical strength; the tube attachments; the accessibility of the equipment inside.

The required accuracy for the positioning of the crystals is 1 mm. The inner diameter is accurate to +/-0.5 5mm with respect to the reference diameters listed in Table 34 and it is concentric to the two external flanges to 0.5 mm. Measuring the position of the externally machined flanges allows the determination of the crystal position. Eighty blind M20 holes are used to fix the crystal azimuth elements to the internal surface of the vessel (Figure 163). The compatibility with the existing vacuum tubes requires the same diameter for the bolting circle and the same diameter for the o-ring contact region. The height from the floor to the axis of the vacuum tube must also be respected. In order to have the correct mounting circle for the crystal, a reverse flange design was adopted (Figure 163).

It is important to ensure a good leak tightness on the interface flanges between the LAV vessel and the vacuum tank whilst respecting the mechanical tolerances of the flanges and the O-ring (Viton) grooves.

Five vacuum flanges, K200, for cable feed-throughs are positioned on one side of the vessels. A man hole (flange K630) is positioned either on the opposite side (stations 1-5) or on below the horizontal plane at an angle of about 45 degrees (stations 5-11) of the vessel (see Figure 164).

The manhole gives the necessary access to the inside and it allows to branch a large diameter pumping port. The inside surface of the vessel is equipped with a stainless steel grid to which the cables can be attached (see Figure 164). The handling of the vessel is assured by three lug attachments on the top. The total weight of the detector stations ranges between 9 and 10 tonnes.

It has been verified that the vessel, if open and with no cover, does not deform from its circular shape (ΔR <2.5 mm). The stress is also kept to a reasonable level for the material used, carbon structural steel 1.0044 EN10028-7 EN10025 (1993) S275JR Fe430B EN10025 (1990).

The construction and welding rules were reported in the PED code and must be followed. The design has been carried out using the standard PED EN13445. The qualification test and control will also be made using this standard. The design external pressure is 1 bar at a temperature of 20^0C. For this reason it is necessary to provide covers for all openings in order to perform vacuum tests. During construction vacuum tests were made after the tube was welded and prior to the final machining.

The test was then repeated after machining. The vessel is supported by four feet located at 30 degrees from the vertical axis of the cylinder (Figure 164). The tube, when supported by the four feet and with the detector weight, has a sag of 0.2mm.

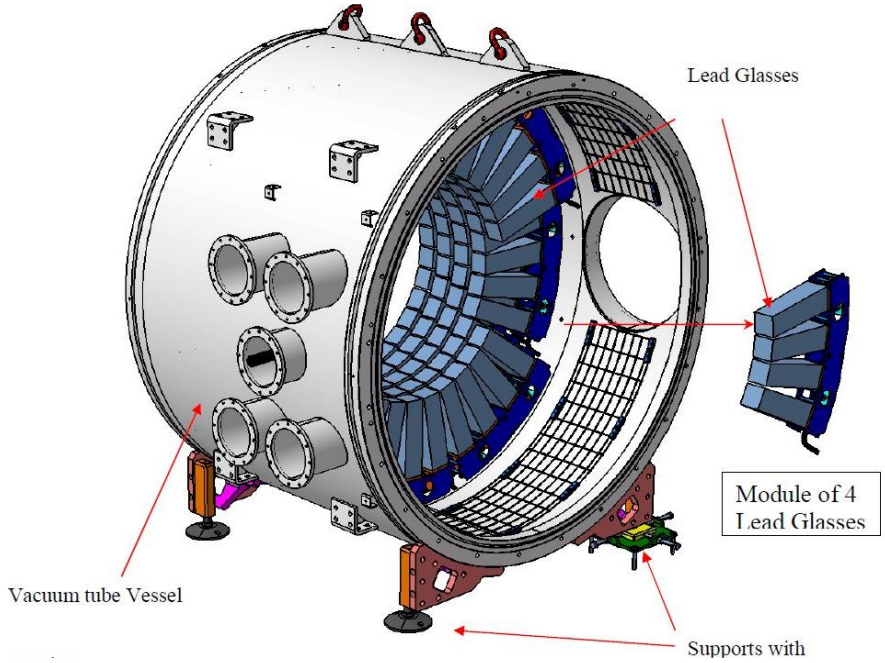

Figure 164 Complete station layout.

3.4.3.6.2 Support Structure for an Azimuth Crystal Segment

The structure supporting the crystals is made mainly by two aluminum plates connected together by spacers (Figure 165).

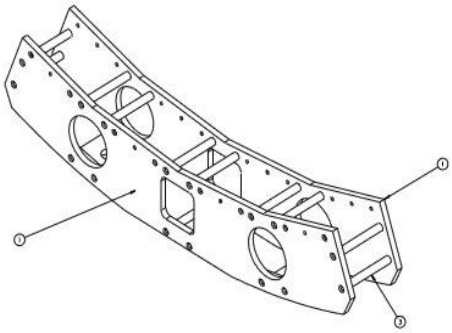

Figure 165 The support structure for one azimuth segment holding four crystals.

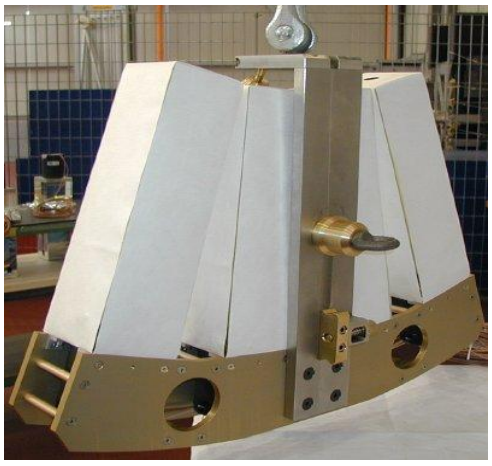

Figure 166 A complete azimuth segment hanging vertically

Four crystals are mounted to this aluminum structure by fixing the 1 cm stainless steel plates mounted on the crystal to the aluminum. The connection between the crystals and the stainless steel plates has been reinforced by gluing four stainless steel strips (106x26 mm, 0.3 mm thick) across the junction between the two elements, on the outside of the counter assembly (Figure 167).

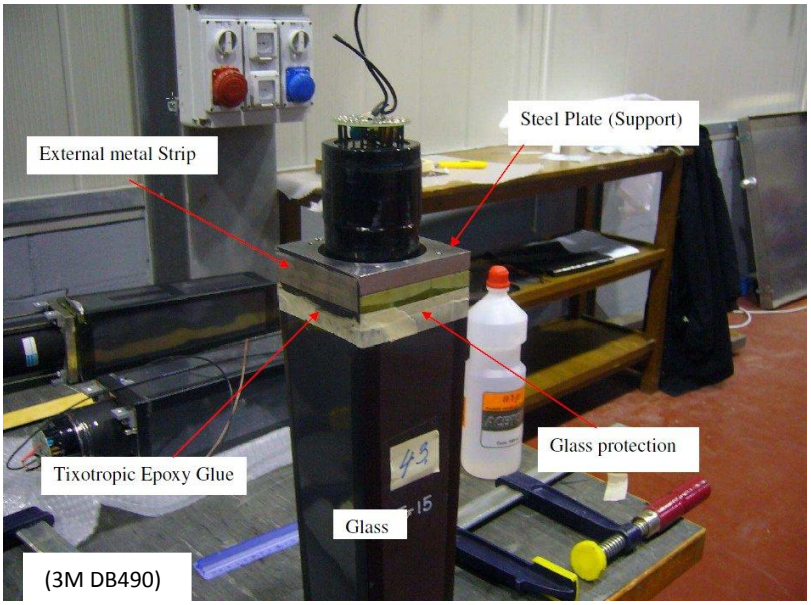

Figure 167 Reinforcement plates on a lead glass block

The two lateral plates contain two circular holes on the side; these give access to tightening the screws to attach it to the tube. The third square hole in the centre is used to handle the four crystals for the installation (Figure 166 and Figure 168).

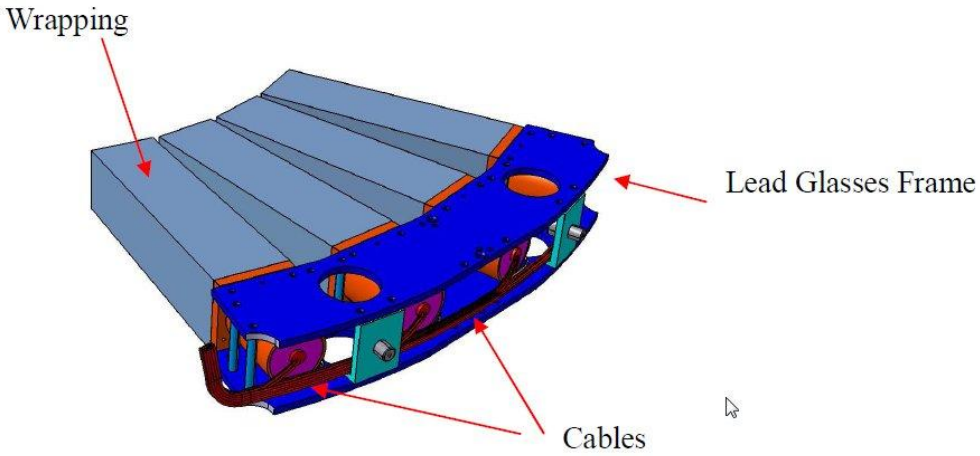

Figure 168 Schematics of the four-crystal assembly

3.4.3.6.3 Flanges

The connection between the internal signal cabling and the external electronics is to be made using 10 37/-pin, D-type CANNON connectors soldered onto a K200 flange. The signal ground is kept separate for each channel. This flange has been built with a custom configuration.

The connection between the internal HV cabling and the external high voltage power supply is to be made using eight 32-pin, MIL C-26482 type, round connectors soldered onto another K200 flange. It should be noted that the 32 pins are connected to 32 individual HV distributor channels. The HV cables will be grounded at the front end electronics crate. The ground connections for all channels will be tied together and connected to a dedicated feedthrough (one per layer).

The flange for the calibration fibres allows the coupling of the external optical fibres (600 μm core) to the internal bundles using SMA connectors. Both fibres face a polycarbonate window and both the internal and external SMA plugs can be finely positioned using small screws to maximize the light transmission. An appropriate O-ring ensures vacuum tightness.

3.4.3.6.4 Cable Routing

Inside a module, cables are routed first along each support structure to one of the short sides. Then the cable bundle (HV, signal and fibre) is laid along the space left between supports up to the section of the module where the grid for the cables has been mounted. The bundles are fastened to the grid in routings that permit each bundle to be properly terminate at its flange cable feed-through. An extra 20 cm of cable is left in the pipe to allow access to the connectors when dismounting the flange from outside.

3.4.3.6.5 Signals

As described in 3.4.3.2.2, the anode of the PMT is connected to a cable (whose length is a function of the radius of the ring) which has on the other side a crimped pin for the signal and the relative ground; these pins are inserted in a DB/37 connector that is connected directly to the flange.

3.4.3.6.6 High Voltage

To minimize material inside the vessel, the negative HV terminal of the PM distributor cable will be cut to length and terminated on the other side with a single pin to be inserted in a 32/pin MIL/C26482 connector, directly connected to the flange. The ground for each layer is grouped together and connected to a dedicated feedthrough.

3.4.3.6.7 Calibration Fibers

In each lead-glass detector the steel flange is equipped with a small hole, used in OPAL to inject light

for calibration and monitoring. The same hole is used for monitoring and calibration purposes. With respect to the one used by OPAL the connector has been changed to a standard commercial SMA connector. The scheme of the calibration system is sketched in Figure 162.

Light is injected in each block by a 200 μm core silica fibre. Fibres are grouped in bundles of four ("downstream bundles") supplying one azimuth segment, each bundle ending with a single SMA connector on one side and with four SMA connector on the blocks side. Each bundle is covered using Teflon, and the connectors are glued using EpoTeK 353 ND to reduce outgassing. For the small rings (ANTI 1-5) each downstream bundle is 4m long. Each downstream bundle is connected to one "upstream" fiber (400 μm core) using an I–shaped SMA connector, which is equally shaped and treated. They end with a with a single SMA connector on the vacuum flange. For the small rings (LAV 1-5) the upstream bundles are each 1m long. The routing of the fibres to the vacuum flange is the same as described for the HV and signal cables.

Externally the light is distributed by a 600 μm core silica fibre which is connected to the LED module. With this scheme, light generated by one LED is distributed to 16 blocks. The intensity of the LED pulser and the quality of the light couplings gives enough light output for our purposes at the end of the light distribution chain.

3.4.3.6.8 Temperature Probes

About 30 PT100 temperature probes are installed inside LAV 1 module, the major part on the hottest resistor in the divider and few on the plates glued to the glass. The purpose of this installation is to monitor the temperature rise in the divider and in the steel-glass junction (known to be delicate and sensible to temperature differences of some tens of degrees) and to possibly monitor the gain variation of the PMTs as a function of the temperature.

On the LAV 1 the connections to the outside readout have been done using spare signal pins (160 were used out of 256). For the other stations, one PT100 for each layer will be installed.

3.4.3.7 High Voltage System

The high voltage system is almost entirely based on the former OPAL HV distributors, powered by new HV power supplies. A specific controller board has been developed to integrate the control of the distributor in the CERN Detector Control System.

3.4.3.7.1 Power Supplies

One unit of the former OPAL HV distributors is a 19" chassis containing up to 15 HV distributor boards plus a controller one. Each board has 16 channels with individual DAC for voltage setting and individual series regulator for HV. The output setting of each channel can be read, after proper selection, through a specific HV connector on the front of the chassis. Each individual PMT high voltage can be set and regulate within a nominal voltage window of 510 V.

Each chassis can power up to 240 channels via output connectors on the back (55-pin round MIL C26482). One chassis will be used for each station, except for the last, where two chassis are needed.

The source power to each chassis will be provided by an ISEG HPn 20 757 power supply (2KV, 750mA max) controlled remotely by the experiment DCS.

3.4.3.8 Readout Electronics

3.4.3.8.1 Requirements

The LAV system will mainly detect photons from kaon decays, as well as muons and pions in the beam halo. For each incoming particles the veto detectors are expected to provide a time measurement, with a ~1 ns resolution, and an energy measurement with a moderate precision of order 10%. The system should be able to operate with very low threshold, well below one minimum ionizing particle (MIP), in order to keep the detection efficiency for muons and low energy photons as high as possible.

The simplest approach would be to split the signal from the PMTs of the lead-glasses in order to have:

- A fast measurement of the charge of the signal, using Flash ADC's or wave-form sampling;
- A fast digitization of the signal, using a discriminator and a TDC for time measurement.

The intrinsic time resolution of the lead-glass blocks (<1ns) and the rise time of the Hamamatsu R2238 PMT (~5 ns) do not put stringent requirements on the time measurement accuracy. On the other hand, the expected energy deposit in the LAV stations from photons coming from π^0 decays covers a very wide range, from ~10 MeV up to 30 GeV as shown in 3.4.1. As a reference, the PMT signal for a cosmic ray muon is shown in Figure 169.

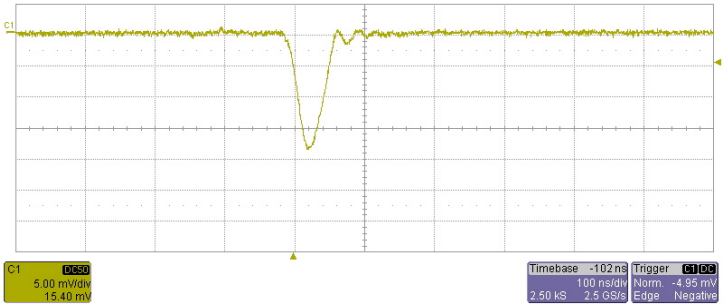

Figure 169 Typical signal from cosmic ray muons

This wide photon energy range, together with the necessity of detecting muon signals, requires operating the R2238 PMTs at a gain of order 10^6. Using the measured average photoelectron yield of 0.3 p.e./MeV, one expects a ~4.5 pC charge for a MIP, corresponding to a signal amplitude of ~20 mV over 50Ω. On the upper part of the range, signals from 20 GeV showers can reach an amplitude of 10V over 50Ω. Such a signal range would be outside the possibilities of commonly used ADC chips, whose dynamic range is at most 50. On the other hand, the range of a typical amplifier-shaper-discriminator chain is of order 10.

In selecting the readout scheme, the cost and data through-put must also be considered. An approach using Flash ADCs is more expensive as the modularity of channels per readout board is low. In addition, it needs shaping electronics to stretch the signal slightly in order to have a reasonable number of samples at 40 Mhz and it imposes huge constraints on the readout board as far as the data throughput handling is concerned.

We have identified three possible strategies to cope with the wide signal amplitude range, while keeping under control the cost and complexity of the system:

- Use a front end electronics with a clamping stage followed by an amplifier;
- Use of multiple scales with different amplification factors for the same input signal;
- Measure the signal height using the Time over Threshold (ToT) technique.

The requirement of signals as high as few Volts rules out most part of commonly used ASIC produced for HEP applications, such as CARIOCA, NINO, ASDQ, etc. On the other hand, the application does not justify the effort of producing a new dedicated ASIC. The idea is then to use commercial amplifier and comparator circuits, to produce a digital output signal with time duration equal to that of the analog input, exploiting the advantage of having high amplitude signals and a relatively slow rise time. Using this technique, the resulting dynamical range of the input signal is enlarged, due to the typical non-linear relationship between the ToT width and the input signal amplitude.

3.4.3.8.2 The LAV readout chain

The output signals from each station are connected to 32 channels front end electronics (FEE) cards. These FEE cards convert the analog signal from the PMTs to a digital signal of proper width (equal to the ToT duration) using the differential LVDS standard. Then, the digital signal time of arrival and width are measured by TDC daughter-boards (using the HPTDC as described in section 4.1.5 on page 339) mounted on a TEL62 board. The FPGA inside the TEL62 produces the time corrected from the slewing introduced by the discriminator, using a double threshold system (described in the following paragraphs), and the charge for each hit reconstructed from the ToT width and a time-charge calibration. This information is then sent to the following DAQ stages. L0 trigger primitives are calculated inside the TEL62 and sent to the Level 0 supervisor. The general layout is shown in Figure 170.

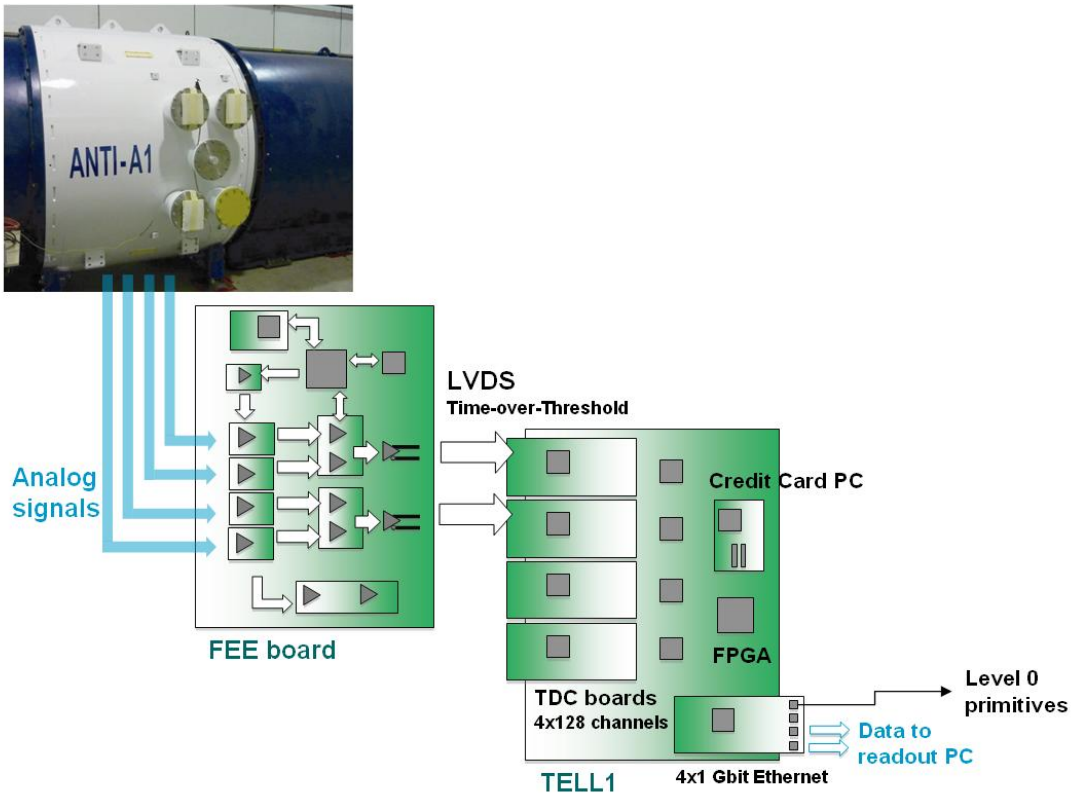

Figure 170 Schematic scheme of the LAV readout (for one station).

3.4.3.8.3 Frontend Electronics

The most specific part of the LAV readout is the dedicated FEE card. This is a custom project designed to convert the analog signal coming from the PMT into an LVDS digital signal of the same width. The analog input signal is clamped to avoid too large input signal to the amplifier stage. The analog signal is then amplified, split and send to two separated comparators (Figure 171). Each of them compare the analog with a programmable threshold and produce an LVDS output signal whose width is equal to the time the analog is above threshold.

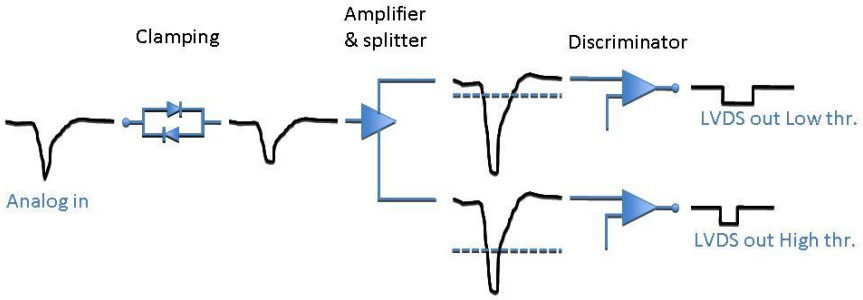

Figure 171 Basic functionality of one channel

The LVDS signal is then transmitted to a TDC to measure its width that is related to the deposited charge. The functional relation between ToT and charge is logarithmic and therefore saturates for

high charge values decreasing the sensitivity of the reconstruction algorithm. The use of a second threshold crossing the top part of the PMT signal, where the slope is steeper, allows to extend the range of sensitivity. A detailed description of the FEE functionality will be given in the following paragraphs. A single channel layout is shown in Figure 172. Each board will house 32 analog input channels delivering 64 LVDS output to the TDC cards. The boards will also provide calibration signals and analog sums useful for the board diagnostic.

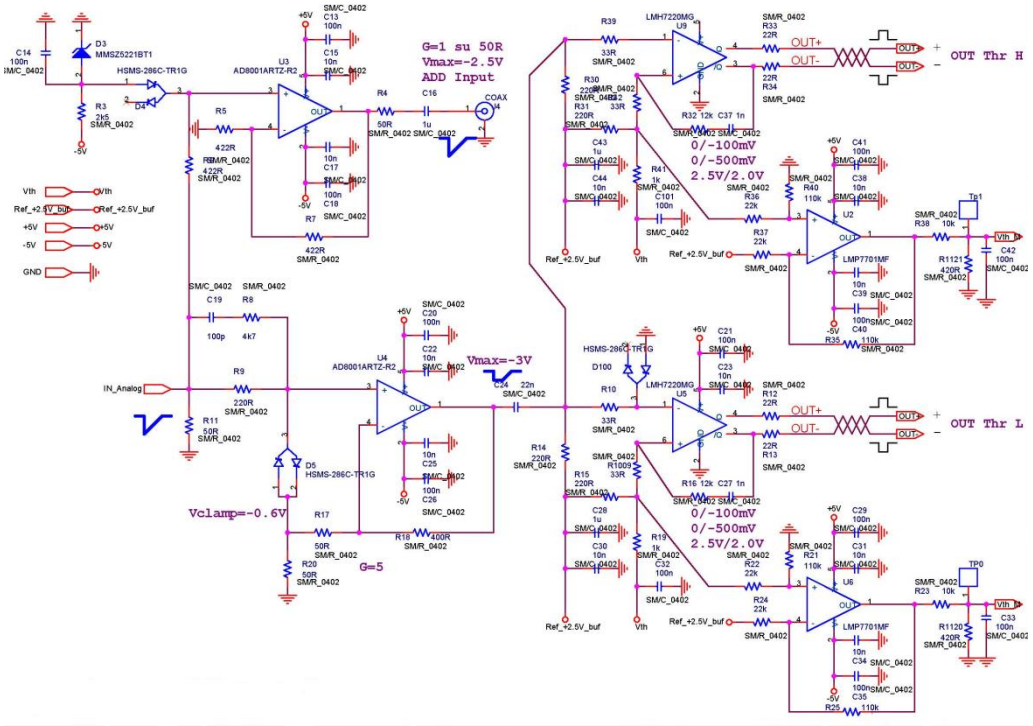

Figure 172 Single channel layout for Time-over-Threshold custom board

3.4.3.8.4 Clamping Stage

In order to protect the amplifier stage from too large signals, the input amplitude must be limited to a maximum of 600 mV in order not to exceed, after amplification, the maximum dynamic range of the comparator (~2V). The clamping circuit is designed to be able to sustain high signal rate up to 12 V, but is able to tolerate an isolated larger signal. To preserve the ToT measurement, the clamp circuit must cut the edge of the signal without changing its time duration. This is achieved by using a couple of very fast low capacitance Schottky diodes (type H5M5-286C-TR1G). The clamping performance has been simulated and the results are shown in Figure 173.

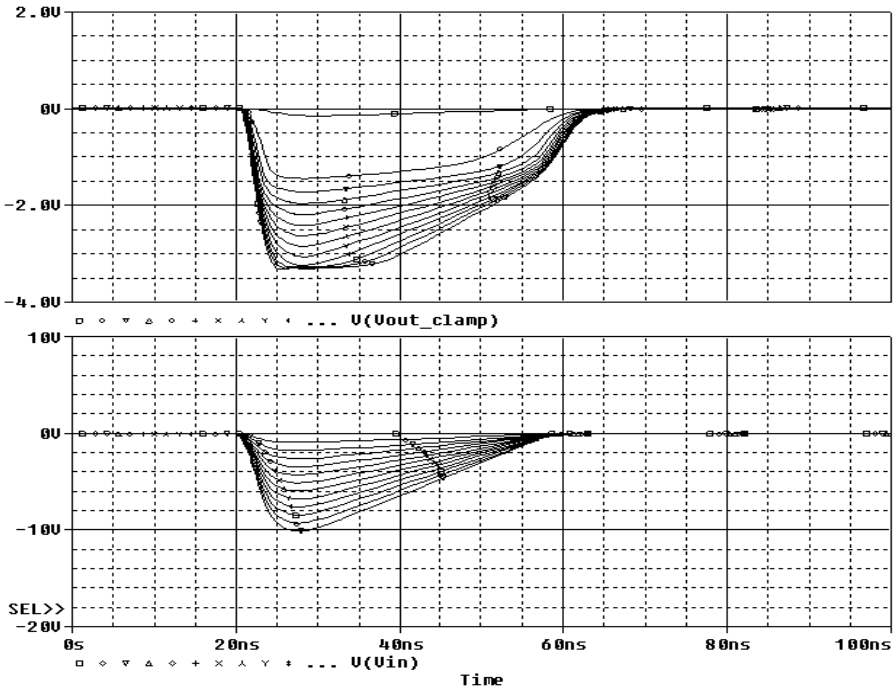

Figure 173 Clamping stage simulation

The figure shows how, when changing the pulse-height of the input signal (bottom) and maintaining the same time basis, the clamping circuit produces a response (top) with clamped output amplitude, but with unchanged time width. To avoid diode breakdown due to excessive heating, appropriate dissipative pads are inserted on the printed circuit board.

3.4.3.8.5 Amplifier

The ToT system should work with a very low effective threshold on the analog signal in order to maximize efficiency. To improve signal to noise separation and reduce the walk dependence on the analog amplitude (overdrive), a moderate amplification is needed (we have chosen a gain of ~3). This also allows compensating the degradation of signal rise time due to the cable between the PMT and the FEE card (4.5 m inside the LAV station vessel plus a 1 to 2 m cable from the flange to the FEE card).

A high bandwidth, high speed Current Feedback Amplifier (type AD8001) is used for this purpose. After the decoupling capacitor C24, to avoid amplifier DC offset with a band of 50 KHz, the output is sent to the comparator input.

3.4.3.8.6 Comparator and LVDS driver

The amplified input signal is divided in two copies and each sent to LMH7220 High Speed Comparator with LVDS Output. This device compares the input with a programmable threshold which can be adjusted in the range 0-250 mV with a 12 bit resolution. To reduce double pulses at the comparator output, due to noise in the input signal, a 10 mV hysteresis is also provided. The comparator produces an output signal starting when the leading edge of the analog signal crosses

the threshold, and stopping when the trailing edge crosses the threshold again. This digital signal is transmitted to the TDC using the LVDS differential standard.

3.4.3.8.7 Analog Sum Circuits

In order to have the possibility to monitor the input analog signals to the FEE board, an analog output is required. Due to mechanical constraints it is impossible to have 32 analog output LEMO connections in the front panel. For that reason the single analog signal are collected in sums of four blocks (one azimuth segment) and then summed again in groups of four to get an half ring sum. The dynamic range of both the inputs and the sums is 2V. The board is equipped with 8 LEMO connection for the sum of 4 blocks, in yellow in the figure below, and two LEMO connection for the sum of a half quadrant, 16 blocks (Figure 174).

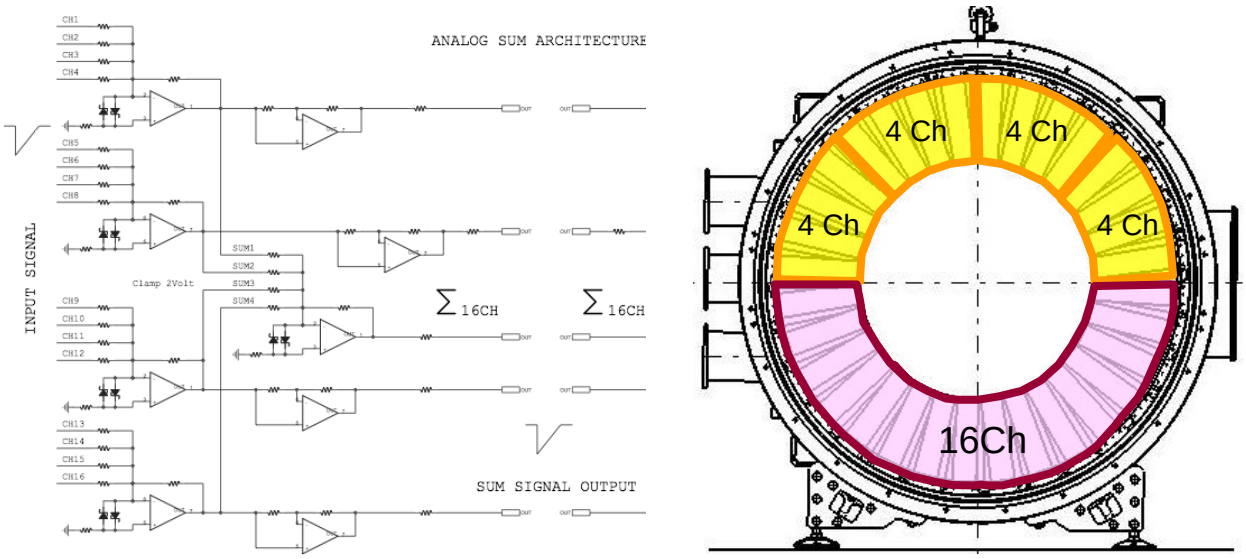

Figure 174 Layout of the analog sums

3.4.3.8.8 Test Pulse Generator

To diagnose electronics functionality, channels integrity and channels calibration, an internal pulse generator has been introduced into the FEE board. The pulse generator it is able to provide pulses with 10-50 ns programmable width, and 10-500 mV programmable amplitude using 12 bit words. The stability of the both amplitude and width is order 1% and the pulse rise time will be order 2 ns. The internal pulser can be triggered both locally, controlled by the on board CPU, or externally, using an external clock. To allow maximum flexibility the pattern of channels to be pulsed can be software programmed.

3.4.3.8.9 Test of the Prototype Boards at CERN

Five prototype boards, with 16 channels each, were produced during spring 2009 with the aim of validating the ToT technique. They have only 1 threshold, adjustable in group of 4 blocks using a trimmer, for each channel without all the control. The boards include an amplification stage with dual output: one copy of the signal is fed to the discriminator, which produces LVDS digital signals with time-over-threshold duration, while a second copy is available for routing to a different readout.

After the assembly of LAV 1 was completed at LNF in July 2009, the module was transported to CERN. It was installed in the existing NA48 vacuum vessel in October 2009.

The five prototype frontend boards were used to fully equip one half of the ring (80 crystals), so that we were able to perform extensive tests with muons and electrons using the SPS beam line. The second copy was used for direct charge measurements for the purposes of comparison.

For the charge measurement, we used commercial 12-bit charge QDCs (CAEN V792), while for the time measurements we used commercial TDCs (CAEN V1190B). The CAEN TDC is based on the CERN HPTDC chip which is the same used in the TDC boards built for the TEL62. In order to provide a gate for QDC integration, a trigger signal was provided by the fast-OR of the 16 digital signals in the upstream ring of lead-glass detectors. A stand-alone DAQ system was also prepared, capable of gating the acquisition with SPS status signals and sustaining a trigger rate of about 1 kHz. We collected large samples of data both with a diffuse halo of muons, thus illuminating all the counters in the ring, and with electron beams of energy 2, 4, and 6 GeV.

3.4.3.8.10 Charge Reconstruction Technique

The reconstruction of the charge of a single hit is based on the correlation between the charge and the Time-over-Threshold produced by our FEE card. This correlation is shown in Figure 175 for events collected during our tests with the cosmic ray setup.

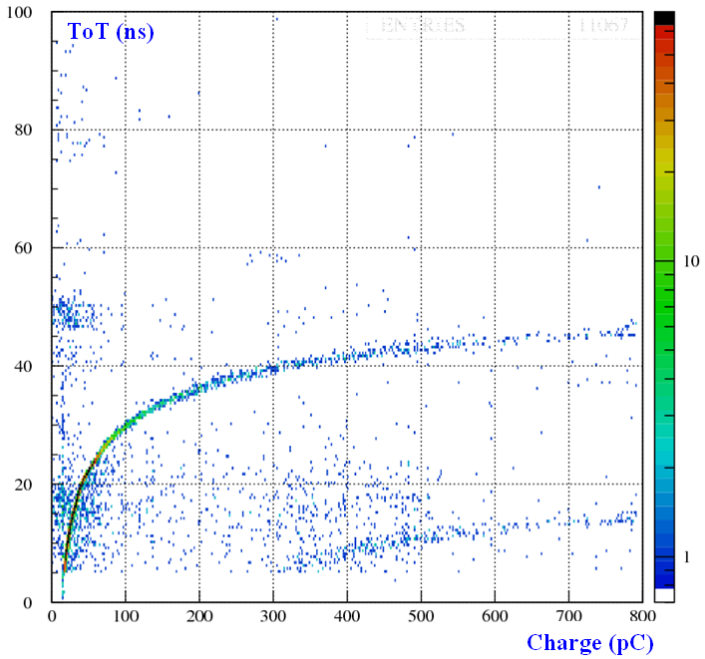

Figure 175 Correlation of ToT measurement and charge for cosmic muon signals

The distribution shows different behaviours in different parts of the spectrum:

- The curve shows signals with a very clean correlation over the entire charge range covered by our QDC (CAEN V792, 400 pC full range, 12 bits resolution)
- The PMT re-firings are identified as high charge and low ToT signals (in the right-bottom part of the plot)
- Noise in the PMT generates a halo of wrong charge events

The charge to ToT correlation also shows saturation at high charges, reducing the sensitivity of the measurement for very high signals. This is not an issue since we would in any case veto events with very high energy deposits in the LAV.

3.4.3.8.11 Charge Reconstruction Performance

The clear relation between the charge and the ToT shown in Figure 175 can be exploited to reconstruct the hit charge from the digital signal time width, namely the difference between leading and trailing edge times. In Figure 176 (left), the charge average value has been represented as function of the time width of the signal. The relation can be parameterized in first approximation with a 4th degree polynomial function:

$$Q_{REC}(pC) = p_1 + p_2\Delta t + p_3\Delta t^2 + p_4\Delta t^3 + p_5\Delta t^4 \qquad (1)$$

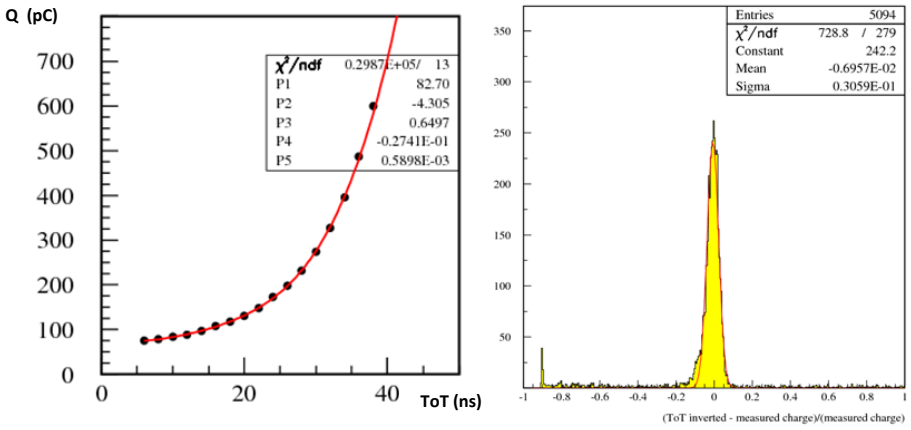

Figure 176 Left: Charge vs ToT. Right: Resolution of charge measurement using ToT

The fit result is superimposed in red. Using (1) and the parameters obtained from the fit, the reconstructed charge Q_{REC} can be calculated. The histogram on the right of Figure 176 represents the relative error on the reconstructed charge obtained with the formula:

$$\frac{\delta Q}{Q} = \frac{Q_{REC} - Q_{ADC}}{Q_{ADC}}$$

The Gaussian fit shows that the core resolution is ~3% and that there are very small tails in the reconstruction. The performances are very good, but we must remember that the charge range is essentially the low part of the spectrum, covering therefore the region of maximum sensitivity for the correlation function.

3.4.3.8.12 Time Resolution

The intrinsic time resolution of the lead-glass blocks has been estimated to be better than 700 ps using electrons during the first LAV prototype test beam in 2007 at the Frascati BTF, due mainly to light collection in the crystal. Nevertheless during operation at NA62 the wide range of signal amplitudes will introduce a considerable slewing, contributing worsening the single block time resolution. Estimates for 10 mV threshold and 5 ns rise-time indicate that the slewing effect can add up to 4 ns for small signals, and at least 2 ns for MIP signals. For this reason a slewing correction is needed in order to reach a good on-line time resolution. The correction can be easily computed once the charge of the hit is known. Due to the absence of direct charge measurement in the present project, the charge should be reconstructed using the ToT signal width, which is of course still affected by the slewing in the leading edge. This problem can be overcome by using a second threshold: the leading edge slope can be easily computed from the two threshold crossing times, and then used to improve the time measurement correcting for the slewing effect:

$$T_0 = T_L - L_{THR} \cdot \frac{T_H - T_L}{H_{THR} - L_{THR}}$$

Using the above formula the slewing corrected time of the hit, T_0, can be computed just using the values of the thresholds in mV, H_{TRH} L_{THR}, and the corresponding crossing time of the leading edge T_H T_L.

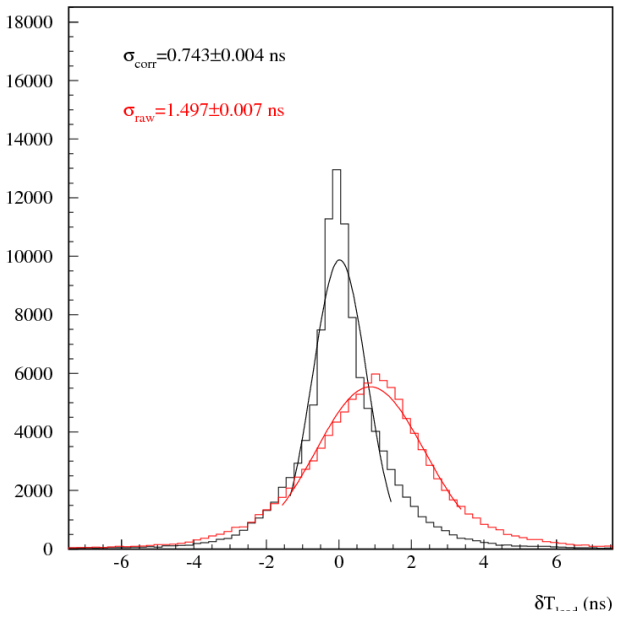

Figure 177 Time resolution with electrons

Figure 177 shows the distribution of hit's times of two blocks produced by 2 GeV electrons during the LAV test beam at CERN. The preliminary result of resolution for slewing corrected distribution in black is 750 ps with an improvement of a factor 2. The core resolution is better than 500 ps for a single block. For each particle crossing the LAV detectors there are many firing blocks allowing further improvements of the time resolution. Averaging all the block measurements a resolution ~300 ps for a MIP and ~150 ps for electrons and photons can be reached in each single station.

3.4.3.8.13 Improving Divider Performances

The correlation ToT vs charge shown in Figure 178 (right) below is obtained using 2 GeV electrons during the 2009 test at CERN of the LAV1 module with SPS beam. We discovered the presence of a discontinuity in the logarithmic relation due to a jump in the distribution of signal ToT, which can cause the charge reconstruction algorithm to fail. Figure 178 (left) shows the analog signal of one Hamamatsu R2238 PMT. The tube has a ringing following the main signal. This behavior has been observed already by Hamamatsu in the PMT datasheet.

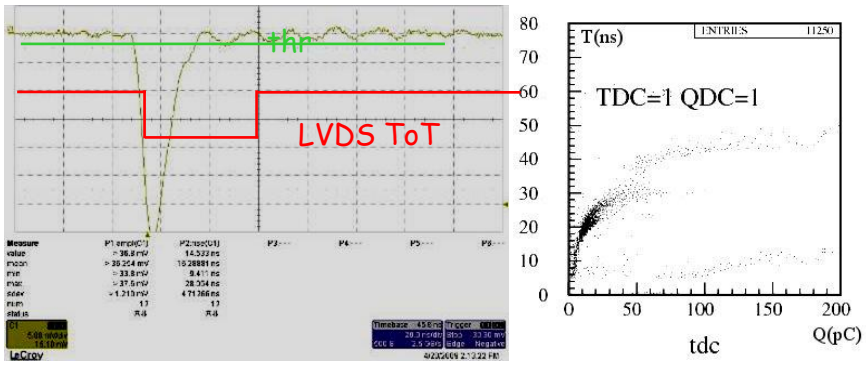

Figure 178 Left: Discriminated pulse – Right: ToT vs charge correlation

The ringing enlarges the LVDS signal, representing the ToT, as soon as it comes above the threshold. In fact, the TDC is unable to separate signals, such as the ToT ringing, that are narrower than 5ns, and therefore the ringing signal is merged with the LVDS of the main signal. After analyzing the PMT in detail we discovered that the ringing was due to parasitic inductance among the last 3 dynodes and the divider of the PMT. To quench the oscillating behavior we projected and built a modified divider including resistors on the last 3 dynodes (Figure 179).

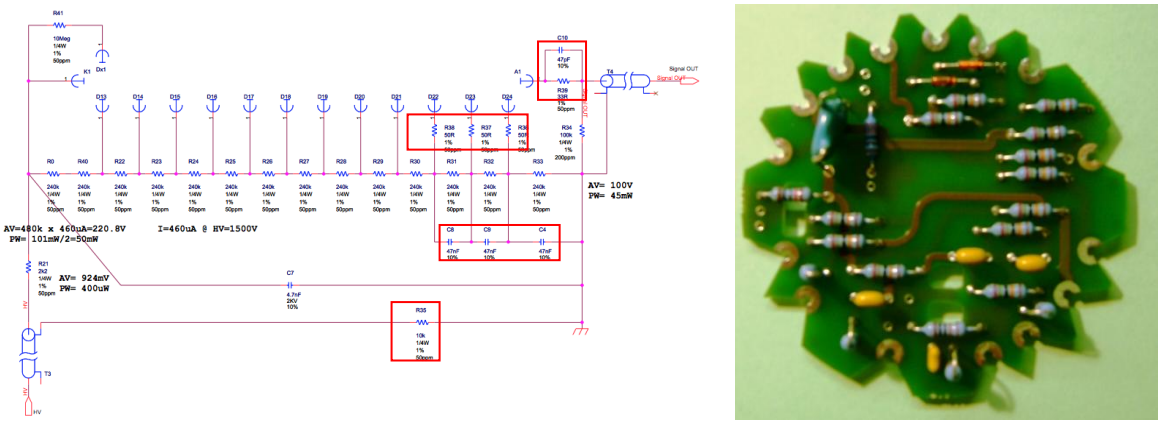

Figure 179 The new PM divider

These resistors are used to reduce the dumping coefficient of the circuit to ~1. The same technique has been also applied to the anode. The red boxes in Figure 179 above show the positions of the four resistors added, as well as of the capacitors in parallel to the last stages and the 10K resistor in the ground line to decouple HV and signal grounds. This solution cured the oscillations in the analog signals allowing us to improve the charge sensitivity of the ToT technique by a factor greater than three. We produced 2500 new dividers which will replace the original Hamamatsu ones. Test of gain and time characteristics of the PMT main signal show no changes with respect to old dividers.

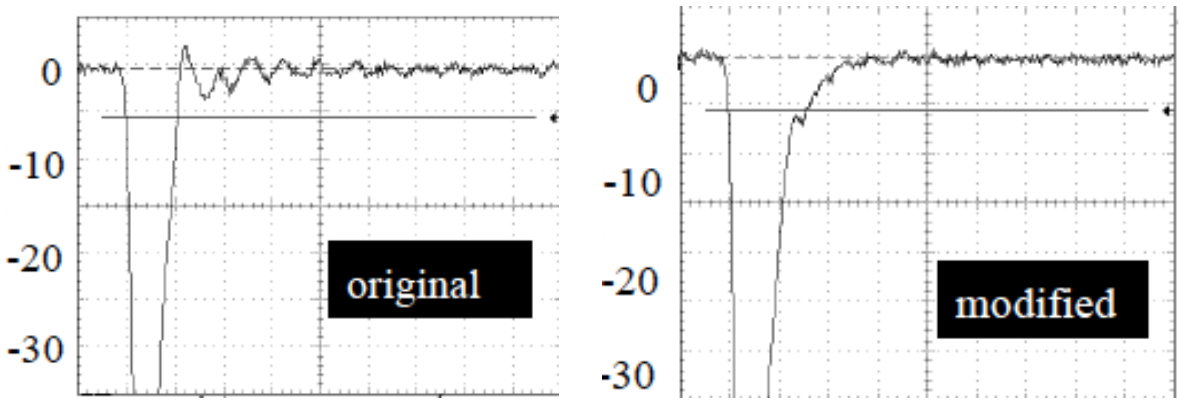

Figure 180 Difference of the signal between the old and new divider

Figure 180 above shows the difference between the original divider and modified divider signals. The ringing in the right part of the left figure is largely suppressed.

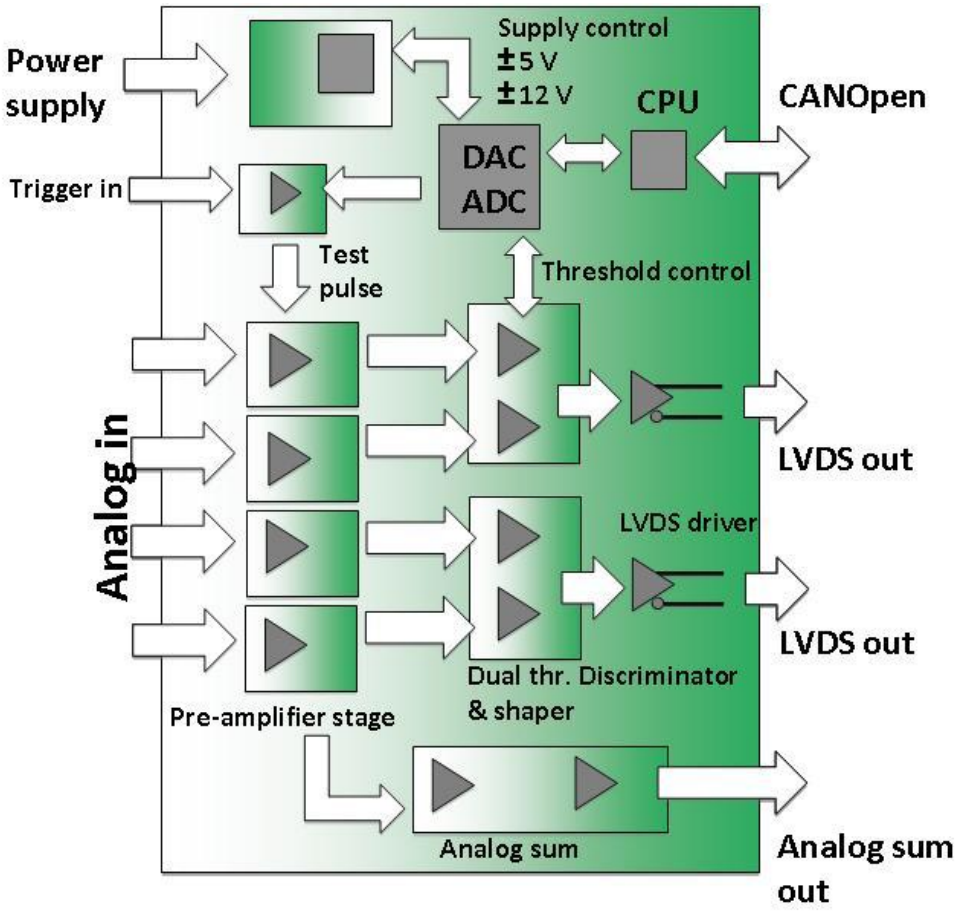

Figure 181 The layout of the frontend card

3.4.3.8.14 LAV Frontend Board Layout

The final LAV FEE card is a VME 9U board able to manage 32 analog input channels.

The input connection is made by two DB37 connectors, in which for each channel both the signal and its ground connection are transmitted. The general layout of the FEE card is shown in Figure 181.

Each board has a two x32 channels output LVDS connectors. The FEE board is connected to the TDC using two 34-Pair, 68-Conductor double-shielded SCSI cables.

The channel multiplicity is very easy to match due to the fact that each TDC daughter-board will allocate 4 connectors filling the 128 channels.

The FEE is made by different modules: the motherboard, which manages all input and output connections and the power supply, 16 daughter board housing 2 channels of the ToT discriminator electronics each, 8 daughter boards for the analog sums, and finally a daughter board housing the DAC and ADC and the slow controls. This design minimizes the cost in term of spare channels allowing to have many spare for the different daughter boards and a few spare for costly motherboards. The boards use standard VME 9U mechanics, 400x400 mm^2, with non-standard power supply ($\pm$7V, +3.3V, $\pm$12 V) and no VME bus. All the slow control and communication functions are implemented using CANOpen standard in agreement with present CERN standards

3.4.3.8.15 Crate for the LAV Frontend

In the present design of the LAV readout system each of the twelve stations will be equipped with a VME 9U non-standard crate housing and, depending on the type of LAV station involved, 1 TEL62 board (9U, 2 slots) and from five to eight front end boards. Moreover there is the possibility that from five to eight monitoring system cards will be placed in the same crate.

The FEE boards will need an extra $\pm$7 V power line and no standard VME bus. For that purpose we are currently considering the possibility to ask for a custom low cost version of the UEV 6021-LHC VME crate to be produced by Wiener. The cost reduction will be based on the substitution of the present 6U VME bus with a 3U J1 power bus only.

3.4.3.8.16 Number of Channels

The present readout scheme is based on three basic elements:

- FEE card (32 analog inputs, 64 LVDS outputs);
- TDC daughter board developed in Pisa (128 channels each), to be hosted on the TEL62;
- TEL62 readout board (can handle four TDC boards for a total of 512 channels)

We assume to use two different thresholds for each physical signal to allow online slewing correction, improve high charge sensitivity, and ensure redundancy in case of TDC broken channels. The presence of a second threshold will not change the number of FEE cards and TEL62 but only the number of TDC boards and cables are affected. In Table 36 the number of channels, the number of FEE cards, TDC cards, and TEL62 board is estimated for the whole veto system.

Table 36 Number of detectors and readout channels

Detector Station	N° of detectors/station	$N_{stations}$	N_{ch}	N_{FEE}	N_{TDC}	N_{TEL62}
1-5	160 (5x32)	5	800 x2	25	15	5
6-8	240 (5x48)	3	720 x2	24	12	3
9-11	240 (4x60)	3	720 x2	24	12	3
12	256 (4x64)	1	256 x2	8	4	1
Total	2496	12	4992	81	43	12

The table shows that we have ~2500 detectors with ~5000 readout channels. To equip the whole system around 85 FEE, 50 TDC boards, and 15 TEL62 are needed.

Each board has two 32 channel output LVDS connectors. The FEE boards are connected to the TDC using flat or round shielded cables with Robinson-Nugent 68 pin connectors. The channel multiplicity is very easy to match due to the fact that each TDC daughter-board will allocate 4 connectors filling the 128 channels. The connections to the TDC system are described in Section 4.1.9 on page number 350.

3.4.3.9 Module Construction Phases

3.4.3.9.1 Assembly

The assembly area at LNF (Figure 182) is located in a newly constructed hangar building, with the following characteristics: the floor is capable of supporting the weight of the apparatus; an area to the right of the hangar doors has been fenced off for use as the NA62 experimental area; a 20-ton

overhead crane has been modified, together with the control unit, to ensure that operations can be carried out safely and at low speed.

Figure 182 The assembly area in Frascati

The assembly operations can be grouped into three phases:

- Preparation and installation of detectors into the cylinder
- Rotation and electrical and vacuum testing
- Packing for shipping

3.4.3.9.2 Preparation and installation of the Detectors into the Cylinder

Preparation of the blocks

The detectors are first removed from the shipping crates, by using a specially designed tool (Figure 183). A single detector weigths about 25 kg.

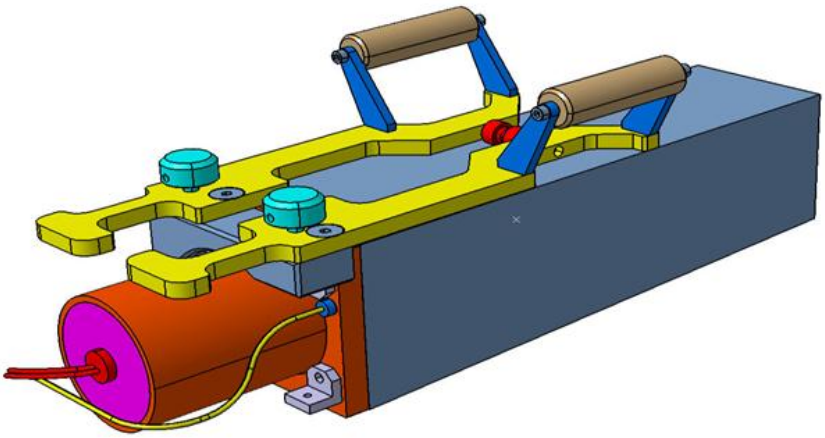

Figure 183 The lifting tool for individual blocks

Reinforcement plates are glued across the junction between the glass and the steel flange on the four crystal sides, to guarantee better mechanical resistance. The glue used is DB490 by 3M, a two-component epoxy resin. The following gluing procedure have been proved to sustain stresses as high as 40 times the crystal weight in the vertical direction. The detectors are positioned inside the holes of a 4.0-m long gluing workbench (Figure 184). They are oriented vertically and a specially designed gluing rig is placed on the top and left in position for five hours for glue curing.

Figure 184 The gluing workbench

After reinforcement, the detectors are wrapped to increase light collection. Wrapping is not required to be light-tight, since detectors will work in complete darkness. The wrapping material consists of laser-precut sheets of DuPont 4173D-Tyvek, chosen because of its high load resistance: with a weight of 75g/cm^2 and a thickness of 205 μm, it can sustain up to 60 N/m^2. The Tyvek is folded over a block mock-up and the folded ends soldered together to give it the correct shape. The formed wrapping has holes to be anchored by screws to the holes on the flange at the top (Figure 185).

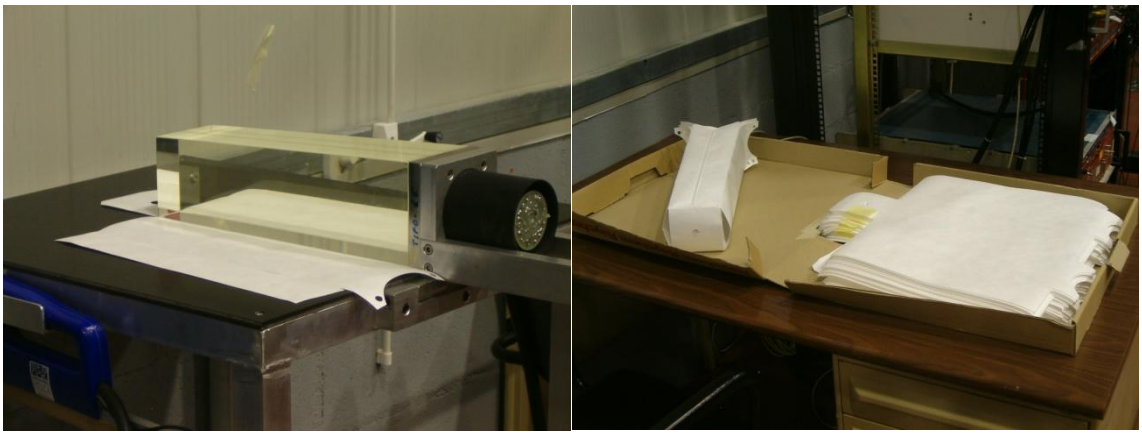

Figure 185 Wrapping preparation

The detectors are gripped by attaching the lifting tools to the aluminum mounting hardware and are then taken to the test area. Here tests of electronics, continuity, etc. are performed continually during the installation period.

In groups of four, the detectors are then mounted in the support brackets, mainly made by two aluminum plates connected to each other by cylindrical and square spacers (Figure 165). The plates have two circular and one square holes per side, guaranteeing access for tightening the support screws to the vessel and for mounting a suitable lifting tool. Crystals are connected to an azimuth segment by stainless steel screws. For each lead glass block the high voltage, the signal cable and optical fiber are routed and fixed into the mechanical structure. At the end of the operation the fibers are checked for light transmission using a LED. A completed segment with four crystals is shown in Figure 186. The assembly is then moved using a cart and a special lifting tool (Figure 187) which uses the central hole of the support structure for the connection and has a plate underneath the crystals as a protection in case of damage during handling.

Figure 186 A complete azimuth segment with four crystals.

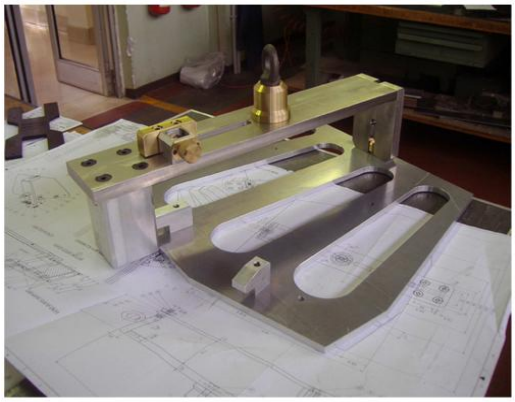

Figure 187 The special handling tool for assembly structures

Since the entire detector structure must operate in vacuum, all the screws used in the mounting and installation either are traversed by an axial hole or they are milled on the side to avoid virtual leaks during pumping; moreover all the screws and the support structure components are cleaned in an ultrasound bath.

3.4.3.9.3 Installation of the blocks into the cylinder

During the installation the cylinder is placed with its axis in the vertical direction (Figure 188), on three concrete blocks with reinforcement frame, the cable mounting grid is installed (

Figure 189) and a first protection ring is lowered into the cylinder, for the safety of the personnel working inside the cylinder in the event of a mechanical failure (Figure 190).

Figure 188 Installation of the vessel for mounting.

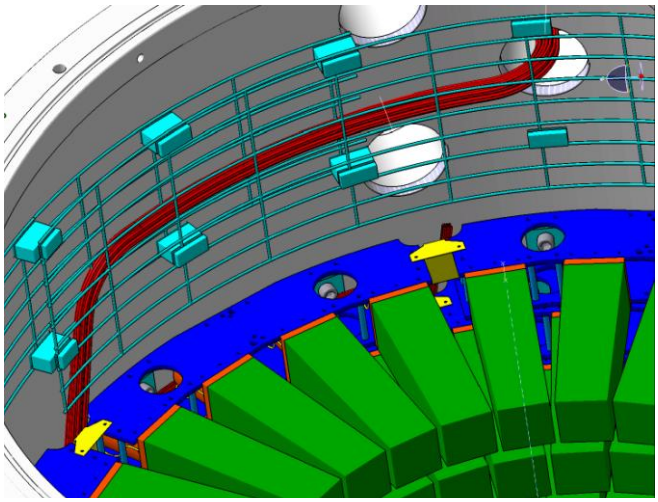

Figure 189 Detail of the cable grid.

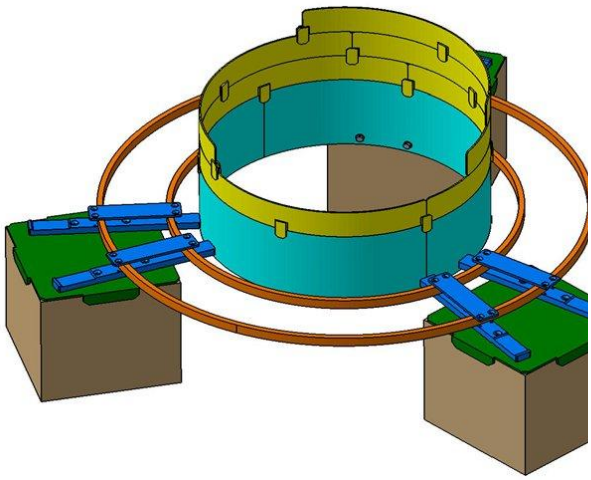

Figure 190 The protection structure.

One azimuth segment, with the lifting bar and safety plate attached, is lowered into the cylinder. The crane movements are performed by an operator on an elevated stand with an unobstructed view of the operation field. The operator lowers the detector segment into the cylinder from above, while the operator inside the cylinder anchors the piece to the cylinder wall and removes the lifting bar (Figure 191). Immediately afterwards, the cables are fixed to the grid and routed towards the portholes for cable exit. This operation is carried out repeatedly during the assembly phase.

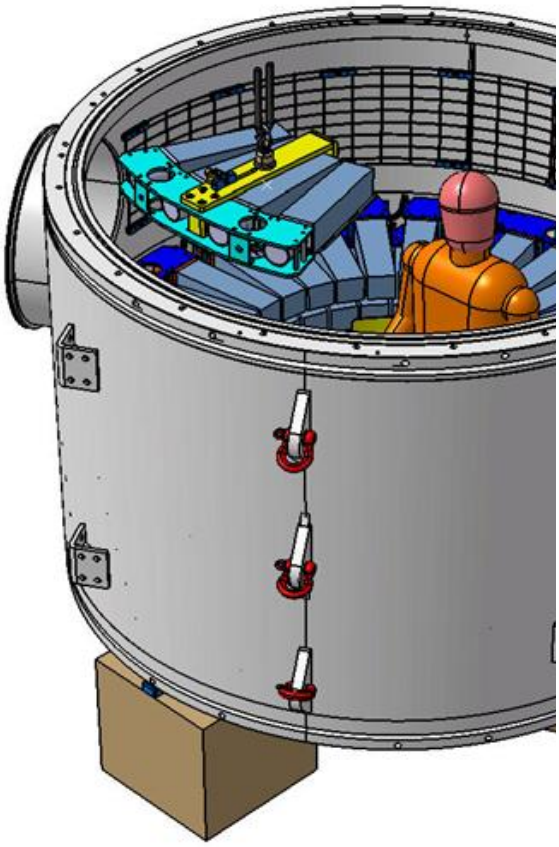

Figure 191 Mounting of a crystals.

One day is needed to mount and cable one layer of crystals. Before mounting the subsequent layer, the cosmic ray signal of each lead glass at a voltage of 1.2 KV is checked using an oscilloscope.

The previous operations are repeated for each layer; once the mounting and cabling is completed, additional checks of the signal and of the optical cabling are performed using the final HV, signal and optical fiber feed-through connections. The final module is shown in Figure 192.

Figure 192 Inside view of one completed LAV module.

3.4.3.9.4 Rotation, Electronics and vacuum testing

To prepare for the vacuum test, the cylinder is rotated by means of an appropriate tool (Figure 193) and it is placed on its feet in the horizontal position and leveled by adjusting the leveling screws. The O-rings are lubricated with vacuum grease and seated in the flanges at both ends of the cylinder. The two yellow cover flanges are lifted and mounted on the flanges of the cylinder. The 24+24 class 8.8 M20 bolts on the two flanges are then tightened to a torque of 250 N m. A 100-mm port on the manhole allows the installation of a vacuum valve and a turbomolecular pump. A disk with reduced aperture of known conductance is installed inside for the measurement of the outgassing rate.

Figure 193 The rotation stand

The pumping system consists of a scroll pump for the roughing stage, of a turbomolecular pump installed on one of the endcap flanges, and of another turbomolecular pump installed on the manhole flange for outgassing measurements (Figure 194 and Figure 195).

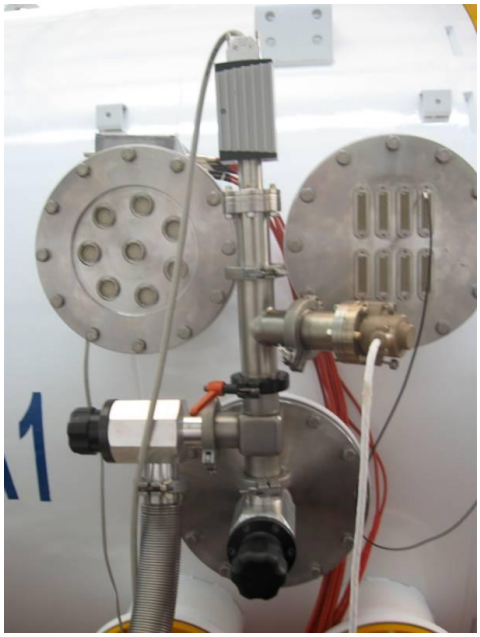

Figure 194 Pumping connections - Small flange side.

Figure 195 Pumping connections - Manhole side.

After the vacuum test is completed, one of the turbomolecular pumps is left on for several days to allow measurement of the specific outgassing rate (see 3.4.3.4.2).

At the pressure of 10^{-5} mbar, a voltage of 1200 V is applied to all photomultipliers and an additional check of the signals is performed.

3.4.3.10 Transport Procedure

The completed stations are to be shipped approximately 940 km from Frascati to the CERN Prévessin site. The first LAV station was shipped from Frascati at 18:30 on 21 July 2009 and arrived at the Prévessin site the next morning at 09:00. The load included, in addition to the first LAV station, the equipment needed for its rotation and handling, for a total weight in excess of 15 tons. The speed of the transport vehicle, a canopy truck, was limited to 90 km/h. Two technicians involved with the logistical planning followed the vehicle. The transport went smoothly, so that we can extrapolate from this experience with confidence. The primary concerns in planning for the shipping were to minimize the exposure of the station to thermal stresses and mechanical shocks.

To minimize thermal stresses, the station was transported during the night. It was not possible to load and unload the station from a thermostatically controlled vehicle. While canopied truckbeds provide easy overhead crane access, loading and unloading from a vehicle with insulated walls would have required a 10-ton forklift, which was not readily available. The construction of a thermostatically controlled shipping container was likewise judged to be unwarranted. Shipping of the module at night eliminated the problem of excessive heating in the summer sun and optimized travel time during which the module was in an uncontrolled environment. There are legal restrictions on the number of hours a single truck driver may operate a vehicle without a suitable resting period. Given the 90 km/h speed limit for the transport, two drivers are needed if the transport is to be completed without stopping for a substantial rest period during the day.

Once on the truckbed, the station was draped with quilts of fiberglass insulation heat sealed between sheets of Propacktene polyester/aluminum laminate. The station was also fitted with two max/min glass-bulb thermometers, as well as with two continuously recording temperature sensors in the shock recorders (see below). Analysis of the recorded data showed that the maximum temperature of 30 ^{0}C corresponded to the temperature of the assembly hall at departure, and that the minimum temperature of 25 ^{0}C was reached during the early morning hours and remained stable until arrival. These results demonstrate the adequacy of the measures taken to prevent excessive thermal stress for summertime transportation. Since heating of the station in direct sunlight is considered to be the primary thermal risk, summertime transport is expected to be more challenging than wintertime transport.

The following measures were taken against mechanical shock:

- The module was transported on end, i.e., with the axis of the crystals in the horizontal plane. This configuration was thought to be the most stable one; in addition, the stress on all crystals due to vibrational shock in the vertical direction is the same.
- Open-cell polyurethane foam wedges were tightly fitted in the horizontal planes between the crystals on different layers. The modules on the bottom layer were supported on foam blocks. An aluminum cylinder slightly smaller than the inner diameter defined by the crystals was placed on the central axis of the station to hold the foam wedges and blocks in place.
- When loaded onto the truckbed, the station was placed on a Sylomer R (Angst+Pfister) closed-

cell polyurethane foam mattress about 20 cm thick. The mattress was designed to reduce the vertical acceleration at 25 Hz (the fundamental resonance frequency of the crystals) by a factor of 20.

- The 90-km/h limit on the speed of the truck reduced the intensity of vibrations and shocks to the station.
- The shock absorbers of the trailer contributed some protection against major vertical shocks. The LAV station was positioned on the truckbed at the point where vibration damping by the suspension was most effective.

Two ShockLog RD298 shock recorders were rigidly fixed to the hull of the vessel for transportation. The acceleration threshold at which the recorders started to register detailed event data was set to 1 g on each of the three axes. No events were recorded on either of the transverse axes; accelerations in the transverse direction remained below the level of 0.2 g. There were two events in which the vertical acceleration was greater than 1 g and thus recorded in detail. In these events, the acceleration oscillates sinusoidally with a frequency just under 1 Hz. The largest momentary acceleration reached was 1.5 g, well within the alarm threshold of 3 g.

3.4.3.10.1 Installation Procedure

All the modules, except the LAV 12, are connected to the vacuum of the main decay tank (blue tubes). The vacuum should be around 10^{-6} mbar and it is important that the flanges are flat and well aligned to provide good tightness. The gaskets providing the seal for the vacuum are "O-Ring" type "Viton" material with hardness of 75 shores and 7mm core diameter. The flanges will be tightened with 24 M20 x 90mm screws in high-tensile steel, type 10.9, with 90kg/mm^2 of tensile resistance.

The installation sequence will be mostly identical for 11 modules. The last module (LAV 12) will be an exception as is not connected to the vacuum tank.

Preliminary operations:

- The module arrives, closed up with cover flanges, in the horizontal position to building 911 at CERN. Here it is unloaded and lowered with the overhead crane (40T) to the ECN3 pit where it is then transported by trolley to the TCC8 area.
- In the TCC8 area the module must be rotated to the vertical position using the same custom made rotation tool used during the mounting phase (Figure 196) . At the same time it is equipped with its four feet for support and adjustment.

Figure 196 Performing module rotation at CERN

Figure 197 Vacuum test before the installation

- Following the rotation, one flange is opened to inspect the inside and to remove the transport protection from the crystals . If no damage is found, the cover flange is mounted again and the module can be dressed with all the components needed for a vacuum test (Figure 197).
- After the vacuum test the two cover flanges are removed and the module is prepared for connection to the decay vessel.
- At the same time the circularity and the planarity of the two flanges of the blue tubes are measured, upstream and downstream of the module. If the divergence, particularly of the circularity, is more than 3 mm then a special tool will be used to restore the circularity inside the prescribed tolerance.

Installation phase:

All LAV modules (except LAV12) are connected to the vacuum vessel (Figure 198) with a fixed flange connection on one side and the sliding ring connection on the other (Figure 199).

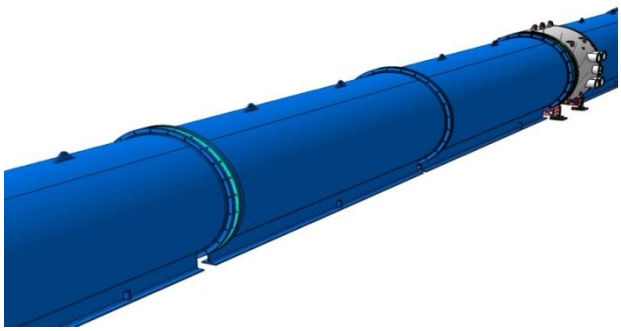

Figure 198 Module between the blue tubes

The connection with a sliding ring allows, in case of necessity, to remove the modules from the decay tank without dismounting other tank elements.

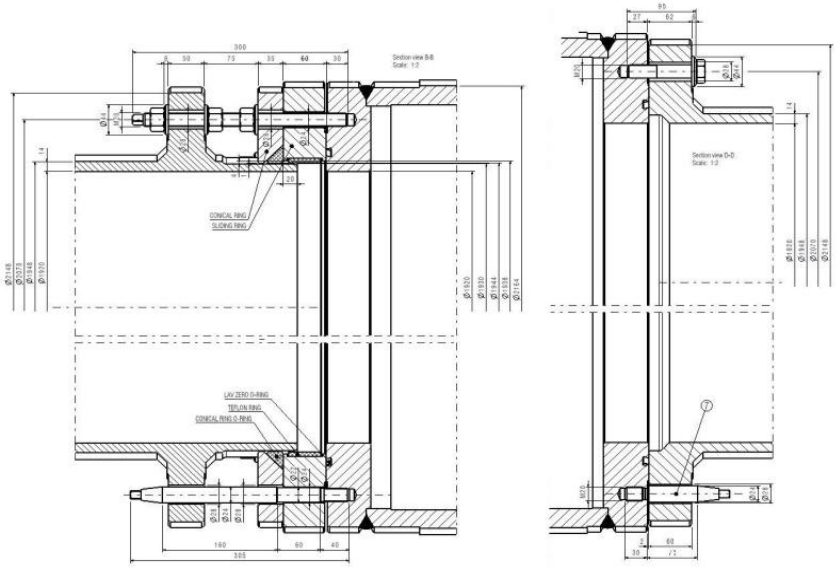

Figure 199 Left: Sliding ring connection – Right: Fixed flange connection

First stage:

It is advisable to check the circular shape and the planarity of the two connecting flanges from the vacuum tank before connecting the LAV module (Figure 200). If the variations of the vacuum flange diameter exceed 2.5mm a special tool is used to restore the circularity of the flange.

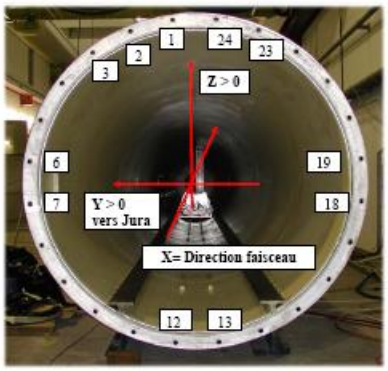

Figure 200 Survey of the blue tube flanges

Second stage:

The LAV module is brought in position using the overhead crane, the fine positioning is made using adjustable ropes (Figure 201). The connection to the vacuum tank starts on the side with the fixed flange where two conical rods, fixed on the module, guide the module to the contact with the flange. At this point 24 M20 bolts are inserted and the flanges are tightened.

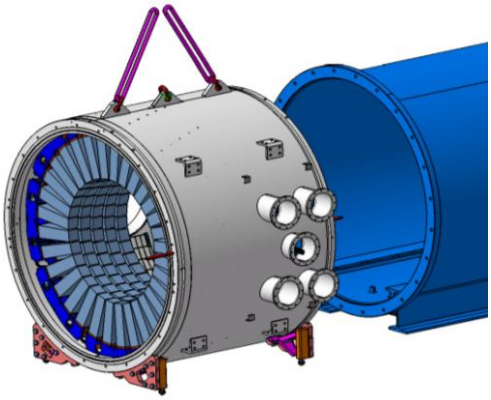

Figure 201 Module hanging from the crane

Third stage:
After checking the axis alignment the four adjustment screws in the feet are lowered and the module is disconnected from the overhead crane (Figure 202).

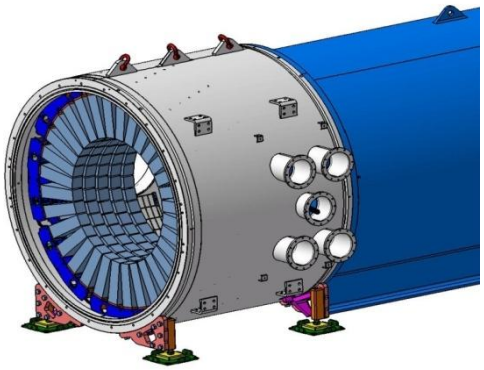

Figure 202 Module in final position

Fourth stage:
The sliding ring is placed on the two long conical rods on the opposite side of the module (Figure 203).

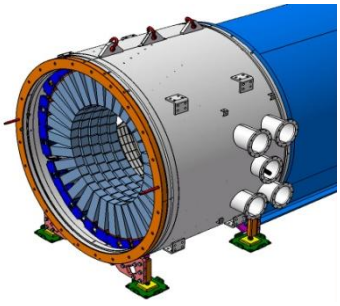

Figure 203 Sliding ring pre assembled

The overhead crane is used to bring the vacuum tank element in front of the LAV module. The two conical rods serve again to guide the vacuum tank onto the flange of the LAV module (Figure 204).

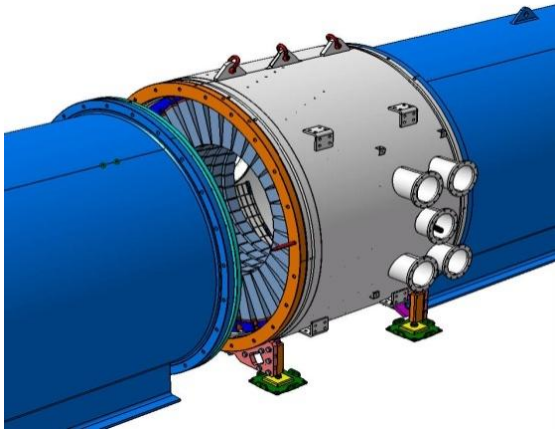

Figure 204 Approach of the blue tube to the module

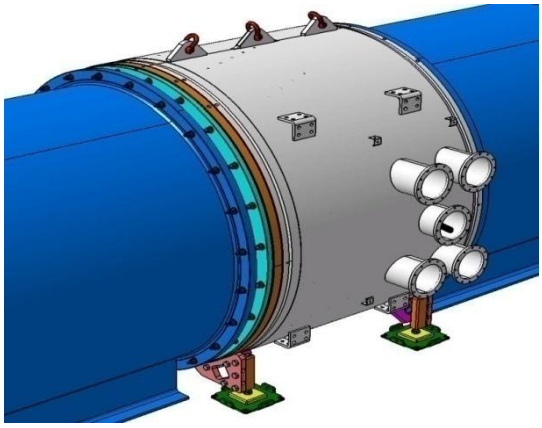

Figure 205 Module installed on the decay channel

After the connection to the LAV module the opposite flange of the vacuum tank element can be aligned to the beam axis. (Figure 205) shows an estimate of the time needed for the installation of one LAV module.

Table 37 Installation time estimate for one LAV module

Operation	Time (h)
Module arrives to building 911, unloaded and lowered by crane to the pit ECN3, transported by trolley to TCC8	4
Module rotation	2
Dismount flanges, remove inner protection, re-assemble flanges for the vacuum test.	3
Vacuum test	48
Dismount flanges, installation on the channel, insertion and tightening of the 24 screws.	4
Survey and align the module	4
Preassemble sliding ring	1
Installation of the blue tube	4
Total	74

3.4.4 The Liquid Krypton Calorimeter (LKR)

3.4.4.1 Calorimeter Requirements

As it has been explained in section 3.4.1, the Liquid Krypton Calorimeter is a key element for vetoing photons from K decays, with the requirement to have a photon detection inefficiency of better than 10^{-5} for energies larger than 35 GeV. In addition the calorimeter is supposed to provide trigger signals based on energy deposition to contribute reducing the L0 trigger rate.

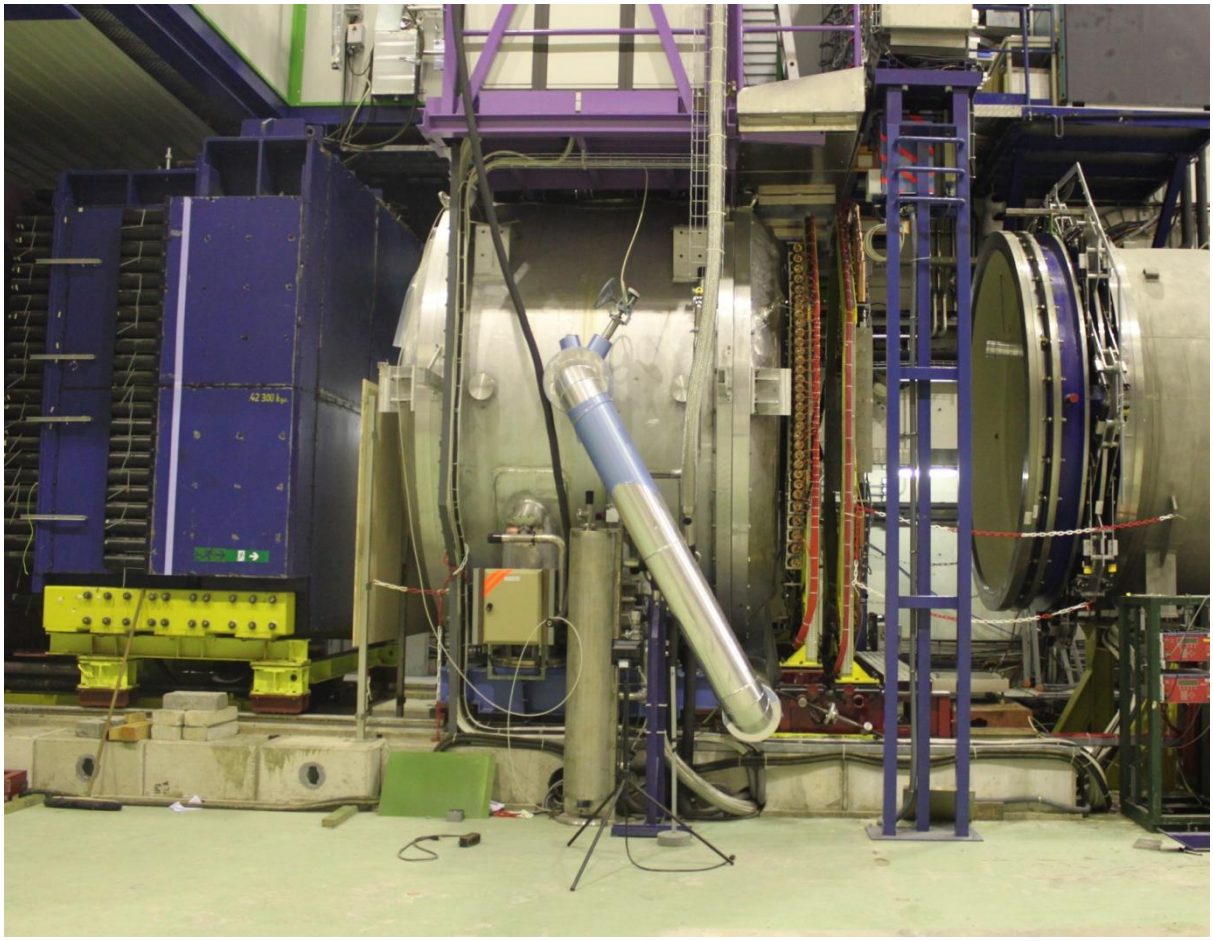

Figure 206 The LKR calorimeter during the dismantling of the NA48 experiment.

Another important requirement imposed to the readout is the capability of reading interesting events with a basic hardware zero suppression (or even without), leaving to offline algorithms the optimization of this job: for example, disentangling a photon shower from a nearby charged pion shower could be tricky if a unique hardware zero suppression algorithm is implemented.

The performance of the calorimeter as used in NA48 has been checked to be good for NA62. Since the L0 rate is expected to be two order of magnitude larger, the old readout system will no longer be usable, as it was limited to 10 KHz (and many of its components are obsolete already).

In addition, other subsystems (e.g. the calibration system logic and software) will need updating. The calorimeter is fully described in [62], and because of that, it will not be described in detail here.

3.4.4.2 The new Readout of the Calorimeter

The readout system of the calorimeter should be updated. Indeed the old CPD-based system has a pipeline memory which can keep the data for 200 μs and the readout of the 64 channels in each individual CPD modules is done serially, limiting the readout rate to 10 KHz. The expected rate of L0 triggers in NA62 is of the order of 1 MHz and the latency of the L1 is of several seconds. This will rule out the use of the CPD pipeline memory. The idea is then (see a complete description in section **4.1.15** on page number 368 to have FADCs feeding a circular buffer with a depth of few seconds, a logic to allow events flagged by L0 to be transferred to a linear buffer of a suitable dimension and a readout procedure which will send L1-selected events to event building PCs using Ethernet. The new system will use commercial-of-the-shelf (COTS) components such as standard DDR2 memories, Gigabit Ethernet (GbE), network switches and standard PCs (processor nodes).

The system will also provide sums of the input signals to the trigger system.

3.4.4.3 The Calibration System

The calibration system is used both to monitor and calibrate the cell response and to spot possible problems in the readout chain. Calibration data are taken in the so-called “out of burst” period, namely 0.5 s, in order to not upset the neutral trigger response during the burst.

The calibration system hardware consists of:

- A calibration pulse generator, which produces signals to the calibration pulser and which has some programmable characteristics, like in-burst (if needed)/out-burst frequency, synchronism with the clock frequency, delay w.r.t. the clock system, etc
- A set of 30 pulser boards which receive the above pulses and which can be programmed to pulse any combination of their 8 outputs
- A 16-bit DAC which provides a voltage reference for the calibrators
- A multiplexer which distributes the DAC voltage to 8 channels connected to specific calibration patterns

It is planned to have a complete renovation of the calibration system, including an upgrade of the logic used to generate the pulses and control the system, moving from old NIM and CAMAC modules to a simple programmable VME system.

3.4.4.4 The Neutral Hodoscope

The neutral hodoscope is a detector made up of scintillating fibers which is installed inside the calorimeter at the depth where the shower development is at its maximum [15]. The photomultiplier

gain has been once equalized and the HV channels are grouped in two HV connectors on the outside HV flanges.

Signal cables are going from one of the brass boxes on top of the calorimeter to the Technical Gallery. The original cabling and amplifier connections has been left in place and it will be relocated in due time.

In case of use, the neutral hodoscope readout should be a new one, because, for past experiments, it was read by the old PMB system (now dismantled). It needs anyway an amplifier system. The readout could be done for example either by the LAV or by the RICH electronics.

3.4.4.5 The Cryogenic System

The complex cryogenics system of the calorimeter has been working for more than twelve years to maintain the calorimeter in a fully operational state at the temperature of the liquid krypton. An upgrade of the control system has been done recently. In order preserve the cold parts of the calorimeter in optimal conditions the collaboration has decided not to empty the calorimeter during the installation period. The main issue here is that repeated temperature cycles may increase the failure probability of the cold electronics (preamplifiers) inside the calorimeter.

However, a list of items that need consolidation on the cryogenic system has been established, and it is considered that the repair work can be carried out while keeping the liquid Krypton inside the calorimeter.

The following main interventions have been planned:

- Replacement of the pump of the Ar cooler vacuum system and its integration in the control system
- Repair of a blocked valve on the Liq. N2 line of the emergency cooler
- Control and replacement of leaky HV feed-troughs
- Adding another pump on the Monotube[26] vacuum system as a backup and its integration in the control system
- Cleaning of the Monotube
- Replacement or improvement of the vacuum valves on the Monotube (Guillotine Valves)
- Various maintenance operations (like change of the thermal insulation)

Moreover, dissemination of the information about the operation of the cryogenics system is under way, as the people with experience of the system are retiring and new people will follow the operations. A formal agreement with the CERN cryogenics support group for the sharing of responsibility is under preparation.

[26] The term Monotube refers to the aluminum beam pipe inside the calorimeter.

3.4.5 The Small Angle Vetoes (SAV)

3.4.5.1 General requirements

In order to provide hermeticity for photons flying at angles approaching zero with respect to the kaon flight direction, two photon veto detectors are necessary. One of them, the inner ring calorimeter (IRC), is to be placed in front of the LKR, and the other, the small angle calorimeter (SAC), at the end of the experimental setup. The spectrum for photons from $K^+ \to \pi^+\pi^0$ decays in the fiducial region for $K^+ \to \pi^+ \nu\,\nu$ decay is shown in Figure 207. Cuts on the kaon decay vertex, on the pion momentum (15-35 GeV) and on the LKr acceptance have been applied. The total energy spectrum is different from that one due to additional contribution of kaon decays further downstream and muon halo of the beam.

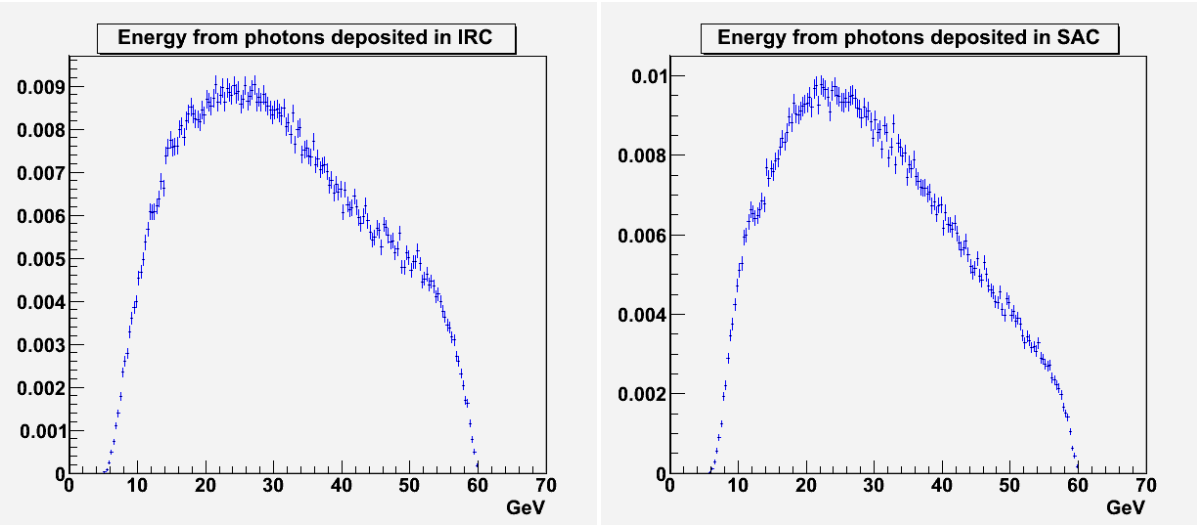

Figure 207 Photon energy in the IRC and in the SAC

Both the SAC and the IRC are exposed to photons with energies higher than 5 GeV. In order to achieve the necessary suppression of the background from $K^+ \to \pi^+\pi^0$ decays, both the SAC and the IRC have to provide detection inefficiency better than 10^{-4}.

The IRC detector will be situated around the beam pipe and its active volume should start as close as possible around the non-decayed kaon beam, and yet far enough away so that the beam halo doesn't generate extremely high rates. The SAC is at the end of the experiment with the only geometrical requirement to cover the region not covered by the IRC and LKR, and to have dimensions small enough not to interact with the deflected non-decayed beam.

The SAC and the IRC are exposed both to kaon decay products as well as the muon halo of the beam. The expected rates are few hundreds of kilohertz during the burst time for photons from $K^+ \to \pi^+\pi^0$ decay. These photons form almost the total rate in SAC while the majority of the events in the IRC will be from the muon halo (rate of 6 MHz). Since the SAC and the IRC are basically a single channel detector, the readout electronics should be able to sustain such rates and to provide good time resolution in order to have a low probability for random coincidences with the events of interest.

3.4.5.2 Description of the Shashlyk Technology

A detector of “shashlyk” type is based on consecutive lead and plastic scintillator plates. The first such detector was suggested in 1991 by Atoyan et al. as an electromagnetic calorimeter for the E-865 BNL experiment devoted to the search for lepton flavour violating decay $K^+ \rightarrow \pi^+\mu^+e^-$ [63]. The incoming electron or photon interacts with the lead and develops an electromagnetic shower. The charged products of the shower produce scintillation light inside the plastic material which then could be absorbed and re-emitted to longer wavelengths by fluorescenting additions. At these longer wavelengths, the attenuation length of the plastic is lengthened considerably. The light is taken out by wavelengths shifting fibres (WLS) where a second wavelength shifting takes place usually to the green part of the spectrum to a photodetector (most commonly photomultiplier). The fibres pass through the plastic scintillator and lead plates via holes in the plates. A schematic view of the layout is shown in Figure 208.

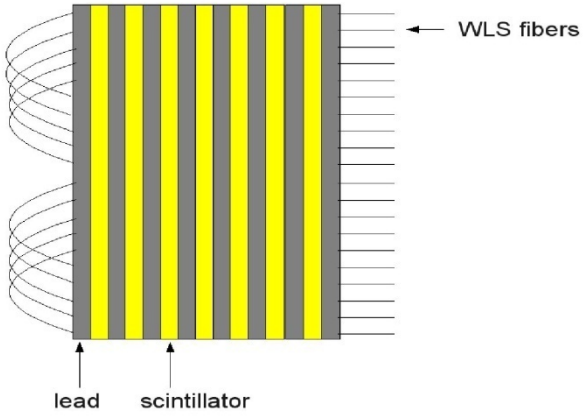

Figure 208 Shashlyk technology

The choice of lead for radiator and photon converter is motivated by the fact that lead has the maximal radiation to interaction length ratio. For a full general description of this type of calorimeter see [64].

A single module of shashlyk type calorimeter is also a single channel detector. The attenuation length of the emitted scintillation light in the plastic scintillator is much longer than the actual transverse size of the module which leads to light in all the fibres. It is important to note that splitting of the total number of fibres into bunches to be readout by different photodetectors doesn't diminish the single channel rate but only matches the geometry and the surface of the active photocathode area to the total surface of WLS fibres.

The shashlyk module has an internal particle detection inefficiency connected with the presence of the holes for the WLS fibres. Concerning the present application of such detectors this inefficiency is not an issue since if can be recovered by tilting or rotating the single shashlyk module as will be discussed in Section 3.4.5.5 for the SAC and section 3.4.5.6.1 for the IRC.

3.4.5.3 Results from the Prototype Tests

A Small Angle Calorimeter (or SAC) prototype was constructed and tested during a 2006 test run. The active part of the constructed prototype was assembled from 70 lead plates with thickness of 1.5mm

with 1.5mm scintillator plates in between, for a total of 21cm (~17 radiation lengths) in order to keep the punch through probability below 10^{-5}. The transversal dimensions were 20.5x20.5 cm^2. The light was taken out with 480 wave-length shifting fibers 1 mm Y-11(250)MSJ (Kuraray) which are separated into four bunches (four channels). The pitch between the holes is 9.5 mm. The assembly was done as follows: first the lead and the scintillator plates were put together to form a sandwich construction supported by four rods at the four edges of the plates. Each scintillator was wrapped in a layer of white paper. When half of the construction was assembled the WLS fibres were guided through the holes of the plates. This way of constructing the detector requires high precision of the holes coordinates with respect to the plates dimensions and with respect to the four guiding holes for the support rods. The WLS fibres were curved at the front so that both ends were readout. The rest of the lead, paper sheet and scintillator plates were put in groups of 5 layers and the WLS were guided through them in steps. When the active part of the detector was ready the WLS fibres were divided into bunches to ease their passage through the PMT supporting aluminium plate.

The detector was equipped with four conventional FEU-84 photomultipliers. The complete detector was placed inside a black-painted metal box hermetic for external light. The assembled prototype just before being closed is shown in Figure 209.

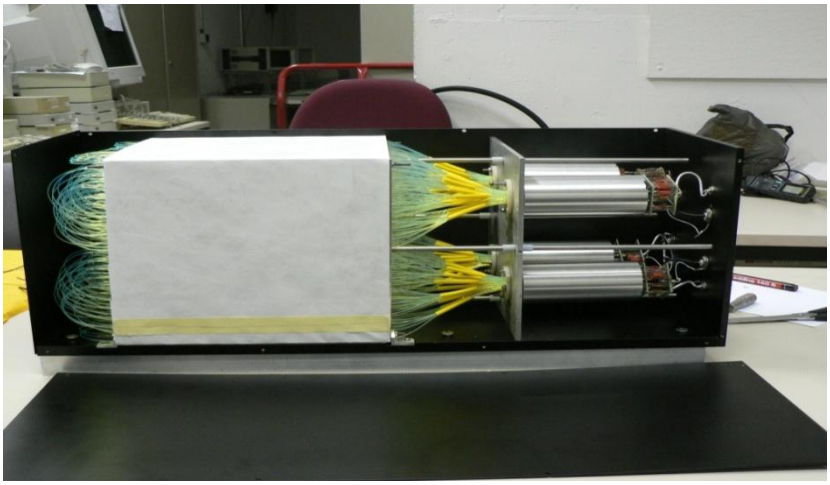

Figure 209 The SAC prototype, with the active part wrapped with white paper

The complete detector was first tested at the assembly room before being installed at the NA48 setup. Its response to cosmic muons was checked. All the channels were verified to provide signal and the time width of the signal was measured. It is a limitation factor for the maximal rate the detector could sustain. The pulse shape is shown Figure 210. Approximately 30ns FWHM was obtained. The detector was tested in 2006 with 25 GeV electron beam, using partially the setup of the NA48 experiment. The charged hodoscope was used for triggering and the magnetic spectrometer to measure the coordinates of the electrons at the SAC prototype position.

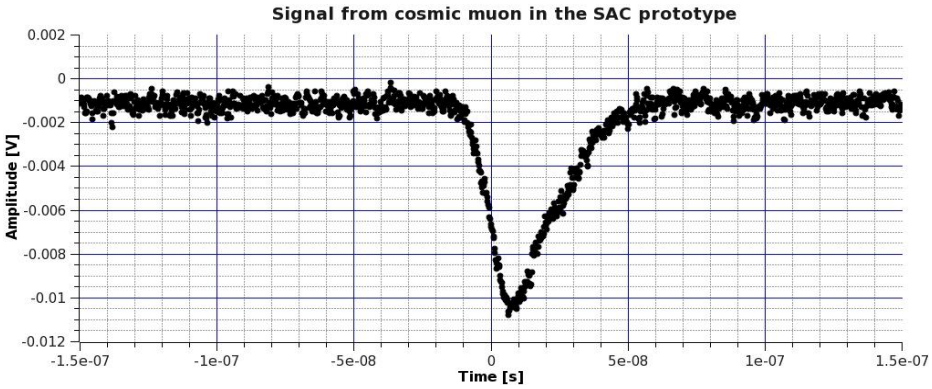

Figure 210 Oscilloscope signal from one of the PMTs of the SAC prototype for a cosmic muon

During the tests SAC was placed between the two planes of the scintillator hodoscope of the NA48 experiment. The signals were sent to an existing CPD based [65] readout of the muon veto. The CPD is based on 10-bit, 14-bit dynamic range FADCs which digitize the amplitude of the signals every 25 ns. The 25 GeV electron beam was deflected to point directly at the SAC. Several runs were taken with the beam pointing to different spots of the detector which allows to study the inefficiency as a function of the electron entry point. A total of 10^7 events were recorded.

The following preliminary results are based on 462 bursts analysed for a total of 2410184 reconstructed events. For each of the four channels, the signal is defined as the measured amplitude minus the pedestal of the channel. The total signal from SAC is

$$S = \sum(A_i - P_i)$$

where A_i is the amplitude in channel i , and P_i is the pedestal for channel i (i=1,2,3,4) in ADC counts. In order to determine the pedestals, only the events for which there were no associated tracks in the Drift Chambers were counted. Since for these events there is no particle, the signal in the SAC is just the pedestal. The pedestals were almost constant for a particular channel during different bursts: the fluctuations observed in the first bursts while tuning was done are less than 0.1%.

The sidemost parts of the detector are expected to be less efficient. The inner area that is effective enough should be determined, i.e. the “effective” size of the SAC. The effective size of the detector is obtained by considering the tracks, extrapolated from the Drift Chambers to the Z-coordinate of SAC (i.e. expected number of tracks), and the tracks for which there was response in the SAC (i.e. detected events). The result for the active area (Figure 211) is obtained using the ratio of these two kinds of tracks (detected/extrapolated).

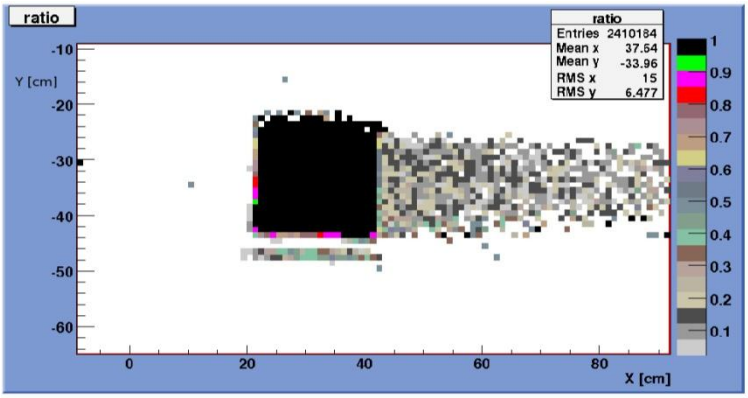

Figure 211 Measured efficiency of the SAC as a function of the photon impact point .

The effective dimensions of SAC can be considered as: X: 22÷42cm, Y: -43÷-23 cm from which the active area is 20x20 cm2. The actual coordinates of the SAC in between the two layers of the charged hodoscope were 22 < X < 43.5 cm, -44.5 < Y < -22 cm meaning that the effective size of SAC was very close to the actual physical size of the detector.

Considering the effective size of SAC, the efficiency map is studied within this area. Figure 212 shows that in general the efficiency is uniform; there are some areas that are less efficient, but the difference is approximately 10^{-4}.

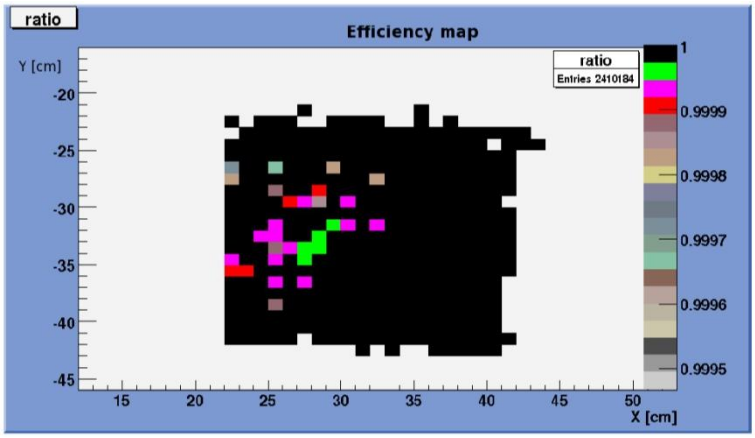

Figure 212 Map of the SAC prototype efficiency

The inefficiency for the effective size of SAC is defined as

$$\varepsilon = \frac{N_{missed}}{N_{extrapolated}}$$

where $N_{extrapolated}$ is the number of electrons hitting the SAC within its effective size and N_{missed} is the number of electrons that were not detected by SAC (below a chosen threshold). On the basis of the analysed amount of data the result is

$$\varepsilon = (2.9 \pm 0.3_{(stat)}) \times 10^{-5}$$

This result successfully conforms to the requirements of the NA62 experiment for inefficiency less than 10^{-4}. Another characteristic of the detector is the random signals determined by the events for which there was signal in the SAC when there was no event, i.e. the track is outside of the detector and the signal in the SAC was above the chosen threshold. For this given threshold the random signals (mostly due to electronic noise) were less than 1% for the assumed event readout time window of 300 nsa nd for cut on the Drift Chambers of 25 GeV.

The Small Angle Calorimeter prototype was constructed with the financial support of the Ministry of Education and Science (Bulgaria) under contract VUF-04/05.

3.4.5.4 Materials

3.4.5.4.1 Lead plates

Since the pure lead is very soft and is difficult to work, tin enriched lead will be used for the plates production. In order to achieve the desired thickness of 1.5 mm the lead-tin alloy (96% lead and 4% tin) will be extruded. The extrusion precision on the thickness achieved for the SAC prototype was 0.15 mm per 205 x 205 mm lead plate.

3.4.5.4.2 Scintillator

The base material for the scintillation tiles is chosen to be polystyrene produced by the Amcrys-H company - the commercial branch of The Scientific and Technical Concern "Institute for Single Crystals" (Kharkov, Ukraine). Industrial granulated polystyrene will be used with PTP (para-terphenyl, 1.5% by weight) and POPOP (0.02% by weight) as dopants for the tiles production by casting technology. This type of scintillator, namely UPS-923A is developed as analogue of the BCN-400 (Bicron) and NE102 (Kuraray) scintillators. The initial excitations of the bulk polystyrene due to the passage of charged particles is transferred to the PTP (with maximal absorption at 275 nm and emission length of 340 nm) via Foersters resonant dipole-dipole interactions. The light emitted by the PTP is reabsorbed by the POPOP (abs=365 nm) and re-emitted in the blue region of wavelengths (maximal emission at 420 nm). The light yield of such scintillator is ~60% compared to anthracene and the attenuation length is about 0.5 m (depending on the coverage – aluminized mylar, tyvek). The rise time is 1 ns and the decay time of the chosen scintillator is about 3.5 ns which is enough for the expected rate of particles passing through the detector [66].

In order to provide better reflection from the scintillator surface, the plastic tiles will be wrapped in Tyvek. This will provide an increased light yield and an improvement of the uniformity of the response.

3.4.5.4.3 WLS fibers

WLS fibres are used to take the signal to the photomultiplier. The chosen type of fiber was multi-cladding Y-11(250)MSJ of KURARAY with 1.2 mm in diameter. The attenuation length inside the fibres is higher than one meter and the usual decay time of the light in the fibre is about 9 ns. The fibres are the major contributor to the length of the signal and thus to the two consequent signal separation capability of the detector. If possible an alternative fibre will be used, BCF-92 produced by BICRON, which shows the shortest decay time (about 3ns) with respect to other fibres [67]. The light emitted by these fibres is with a maximum at approximately 490 nm (green light).

3.4.5.4.4 Photodetectors

Photomultipliers were chosen as photodetectors for the light from the WLS fibers. They provide gain up to 10^6. A FEU-115M-10 PMT was chosen. It has a 10 stage Al-Mg alloy linear focused dynode system. The photocathode is made from Sb-K-Na-Cs and is semitransparent. The average quantum efficiency for wavelengths in the interval 400-500 nm is higher than 12%. This PMT has been tested for LHCb ECAL and HCAL [68]. The gain is between 10^3 (for 1250 V applied high voltage) up to 10^6 for 2000 V. The transit time varies from 36 to 27 ns depending again on the high voltage (1250V to 2000V). An important characteristics is the rise-time which for this PMT doesn't depend on the HV and is less than 5 ns. It also provides linearity better than 1% for anode currents up to 50mA (and even higher).

The phototube has a diameter of 30 mm with 24 mm useful photocathode diameter. The length is 72 mm and a typical base is of length 30 mm for a total of 102 mm long tube. It should be placed in a μ-metal cylinder for magnetic field shielding. The alternative phototube considered is HAMAMATSU R5800 which provides better photocathode uniformity and is operated at lower high voltage.

3.4.5.4.5 Outgassing

Since the SAC is planned to operate in vacuum the outgassing properties of the materials are important parameters. No dedicated study on the outgassing of the SAC prototype was performed. However the outgassing of scintillator, lead and Tyvek paper assembled at Protvino as a potential prototype module for the MUV has been tested. The obtained results both from the CKM collaboration during the tests of the Protvino module [69] show that the most leaking material is the scintillator. The rate was found to be $1.2*10^{-7}$ mbar.l/(s.cm^2) after five days of pumping. The total expected outgassing from the SAC was calculated assuming the total scintillator surface and was found to be approximately $6.1*10^{-3}$ mbar.l/s. Taking only into account the dimensions of the tube where the SAC will be located (rough estimate of ~50 cm diameter, ~7 m long), this turns into $1.1*10^{-6}$ mbar/s.

3.4.5.5 Mechanical Design of the Small Angle Calorimeter (SAC)

The SAC will operate in vacuum and will be placed on a rail in an approximately 7 m long 100 cm diameter vacuum tube. The schematics of the placement of the SAC in the NA62 complex is shown in Figure 213.

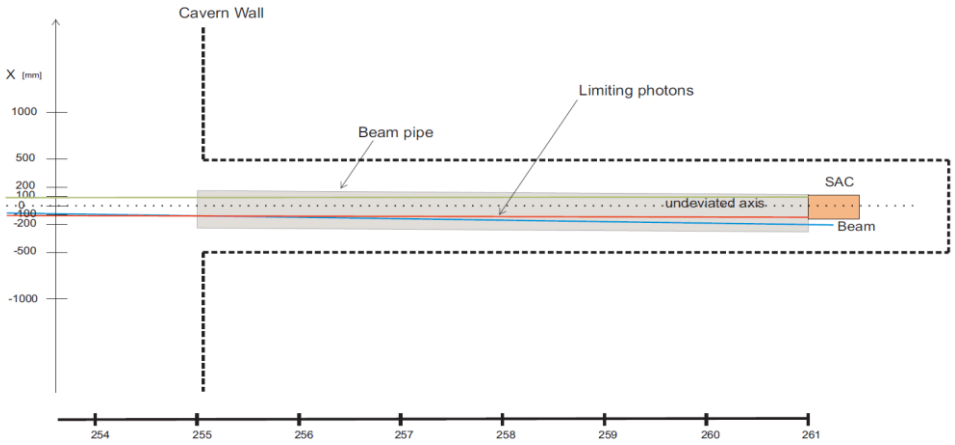

Figure 213 Placement of SAC inside a vacuum tube. The deviated beam is also shown

The detector itself will stay on a table which will be supported from below on two rails which will be welded to the inside of the vacuum tube. The table will have ability for height adjustment (Y axis) and rotation in the XZ plane (along Y axis). The detector will be put at a small angle with respect to the X axis in order to avoid the possibility of a photon originating from the fiducial region to pass completely along a scintillating fibre. The studies show that an angle of few mrads is enough to recover the high efficiency for photon registration (Figure 214). This angle can also be obtained by just requiring that a photon parallel to the non-deviated axis passes at least through part of the converter material (taking 1.5 mm the hole diameter per 210 mm active detector leads to an angle of 7 mrads).

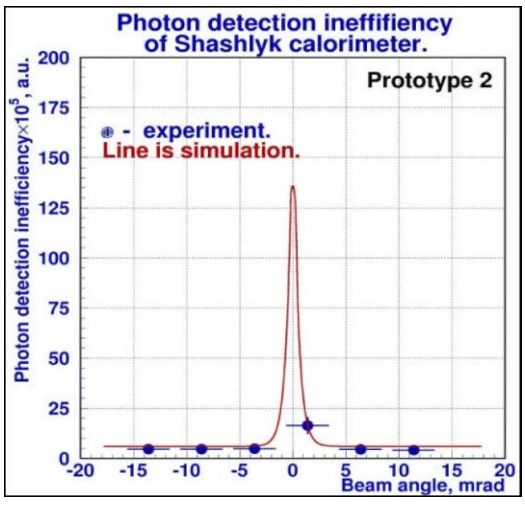

Figure 214 Shashlyk calorimeter inefficiency as a function of the beam angle [70].

The lead plates will be manufactured by the Laboratory for Technical Development in Physics, Bulgarian Academy of Science (as a subcontractor, resources from project DID02-22/17.12.2009, Ministry of Education, Youth, and Science, Bulgaria). A total amount of 80 square plates will be produced with dimensions (205 ± 1) mm times (205 ± 1) mm and thickness (1.5 ± 0.15) mm (

Figure 215). Symmetrically with respect to the center of the plates 484 holes with diameter (1.5 ± 0.05) mm will be made in 22 columns each with 22 holes. The distance between the centres of the holes in each row and between adjacent rows will be (9.5 ± 0.05) mm. The distance between the centres of the outermost holes will be 195.5 mm so leaving 1mm distance from the edge of the hole to the edge of the plate.

In the four corners of the plates 4 bigger holes will be made for the support rods – again symmetric with respect to the central axis of the plate, the centres of which form a square with dimensions 190mm x 190mm. 70 out of the 80 plates are necessary for the production of the SAC, 10 more are intended to be used in case of obstructed holes. During the prototype construction only one out of 33880 holes in the lead was blocked (probability ~$3*10^{-5}$). This could happen since during the drilling process there is a chance that the remnants of the lead fall into the hole and remain there due to high lead plasticity. As in the prototype, the SAC will be 21 cm (~17 radiation lengths) thick.

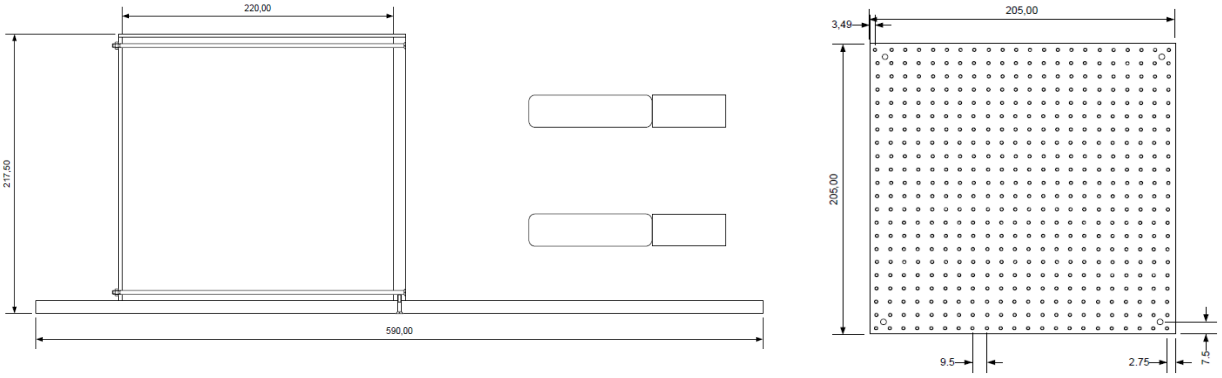

Figure 215 SAC general layout and a single lead plate

A special matrix for drilling holes in the lead will be used. The same matrix afterwards will be used to drill holes in the scintillators. The scintillating tiles will be the same size as the lead ones. The holes will have (1.4 ± 0.15) mm diameter and again the pitch between them is (9.5 ± 0.05) mm.

Tyvek paper will be put between the scintillator and lead plates. Also the whole detector will be wrapped in TYVEK. The assembly of will be done with the detector lying on its supporting main plate (shown below, also used as the table for the support on rails). Lead, paper and scintillator will be put vertically and the fibres will be strung through their holes, layer by layer. 1 meter WLS fibres read from both ends and curved at the front will be used – long enough to assure a curvature diameter of 5cm (shorter radius leads to light loss) in the front of the SAC and to operate freely when combining them at the back.

In order to support the lead and the scintillator tiles a front plate of 4mm thickness made of aluminium is foreseen. It will have the same distribution of holes like the converter plates with the only difference of the diameter which will be 3mm. The packed active SAC volume will be tightened with screws on the supporting rods at the end of the assembly, behind another aluminium supporting plate (again with 484 3mm holes). Yet another aluminium plate will be placed 10 cm after the end of the lead-scintillator sandwich where the PMTs will be mounted (Figure 216). The 10 cm are necessary in order to have enough space for WLS fibres grouping before they go the photomultipliers.

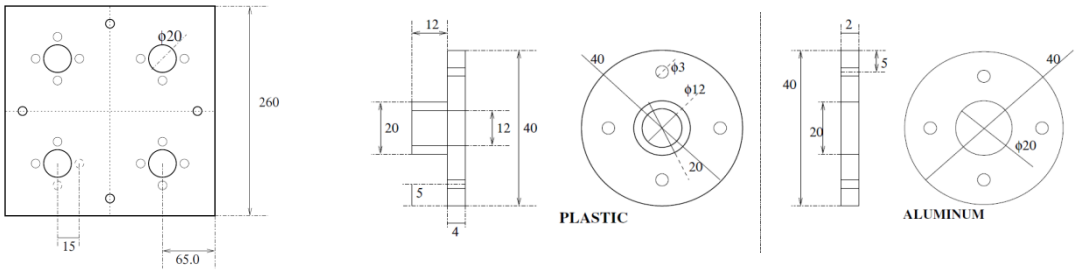

Figure 216 PMTs supporting frame at the end of the active detector volume. The WLS fibers through the hole in the plastic to the PMTs faced at them.

3.4.5.6 Mechanical Design of the Intermediate Ring Calorimeter (IRC)

3.4.5.6.1 Detector design

The detector will be made as a cylindrical tube (ring) with two active regions. The first half will be with 25 layers of lead and scintillator while the second half will have 45 layers, for a total of 70 layer sandwich. The matrix for holes drilling for both halves will be rotated at a small in order to prevent a photon from following a wavelength-shifting fibre completely through the detector, a path which would lead to a loss in detection efficiency. The outer radius of the IRC will be centred at the Z-axis of the experiment while the inner radius will be centred at the beam axis which, at the IRC coordinates, is displaced 1cm away in X-axis. In this way the detector will not have a rotational symmetry. The front half will have radius of the inner region 60 mm and an outer radius of 145 mm, while the second half will have inner and outer radii of 61 and 145 respectively.

The second half of the detector will have an inner radius very slightly larger than that of the first half in order to prevent photons hitting the detector edges too far downstream and thus escaping detection. The matrix for holes drilling will be rotated 40mrad with respect to the first half along Z axis. The small rotation angle in the XY plane of the experiment will provide the necessary coverage for the whole geometrical region.

The lead plates will be produced by the Laboratory for Technical Developments in Physics, BAS-BG. They will be made as a single ring converter. The same identical hole matrix as for the SAC with an

additional tool to fix the ring lead plate at the centre of it will be used for the holes drilling - this means a pitch between the holes of (9.5 ± 0.05) mm and a hole of (1.5 ± 0.05) mm. The distribution of the holes on the lead plate is shown in Figure 217 .

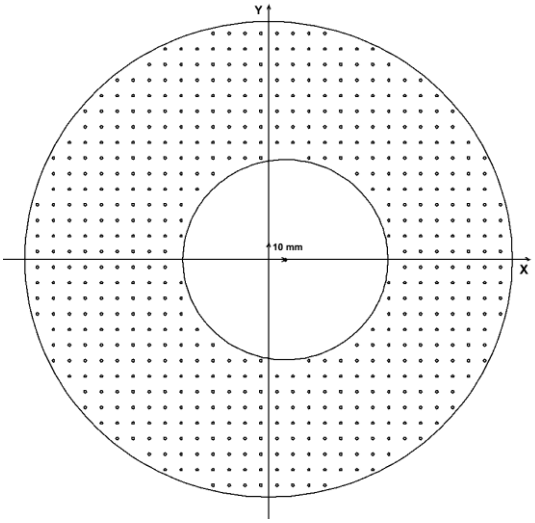

Figure 217 Layout of the holes in the IRC plates – XY plane

The high rate in the IRC could lead to event overlapping as described in the earlier sections. The width of the avalanche from electrons with 25 GeV energy in the shashlyk prototype during the tests was estimated to be about 5cm. This allows to make the inner ring calorimeter segmented into four parts, each representing 90 degrees sector (centred at the Z-axis). The segmentation will be achieved by dividing the scintillating tiles into four sectors. Each tile will have its edges aluminized in order to avoid cross talk between the four channels. The total width of the first half will be approximately 80 mm while the second will have width 140 mm (including the Tyvek, without the enclosing aluminium plates, each 5mm thick) (see Figure 218).

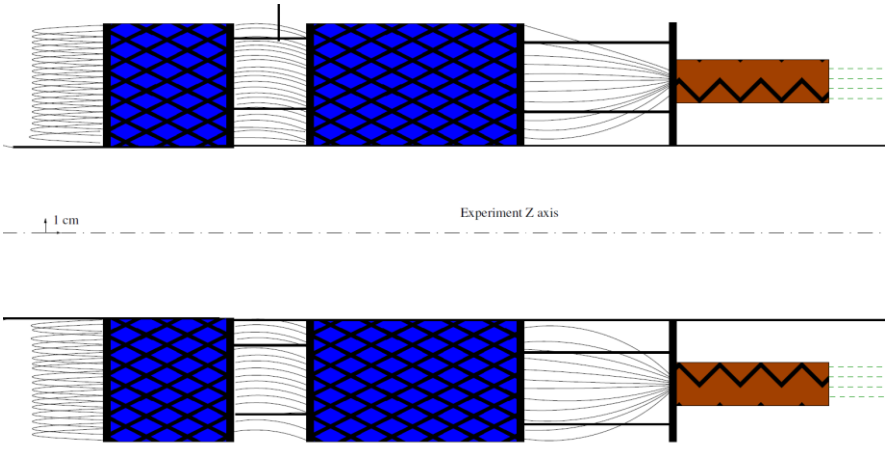

Figure 218 IRC design in the form of two halves in the YZ plane. Inner hole of the first half is ~59mm.

The IRC will be equipped with four photomultipliers with additional high voltage stabilizer in order to assure stable operation in high rate regime. The last four dynodes of each PMT will be provided with

additional high voltage supply. PMTs will be mounted on a aluminium ring located 10 cm after the closing plate of the second half of the detector. This will provide enough room for curving the fibres.

3.4.5.6.2 Support Structure and Integration

Since the IRC is a monolith detector as the lead plates are in a ring shape, it necessarily incorporates part of the beam tube. The beam tube passing through the IRC will also be made in two welded parts of 1mm stainless steel with outer diameter of 120 mm in the front and 122 mm downstream. The upstream part of the IRC beam pipe will be attached to the aluminium beam pipe passing through RICH by an elastomer O-ring made from fluorocarbon (Viton® by DuPont). It is good enough to provide small permeability for a vacuum of 10^{-4} mbar as expected in the beam tube. The downstream tube will be welded to a steel bellow which will be then connected to a CF flange, connected to the LKR guillotine valve (Figure 219).

The total IRC assembly will weight about 70 kilograms and cannot be left only on the beam pipe. The detector will hang on a string made of carbon fibre or stainless steel from a rod attached at the top of the LKR. The string will be attached to a rod located between the two halves of the IRC. The rod will be attached to the two aluminium disks enclosing the active detector region.

3.4.5.7 Readout Electronics and DAQ integration

The SAC and the IRC will use a TEL62 based readout with high-speed Flash ADC modules to be used waveform digitizers, because of the high rate. The description of the readout scheme is done in section 4.1.16 on page 384.

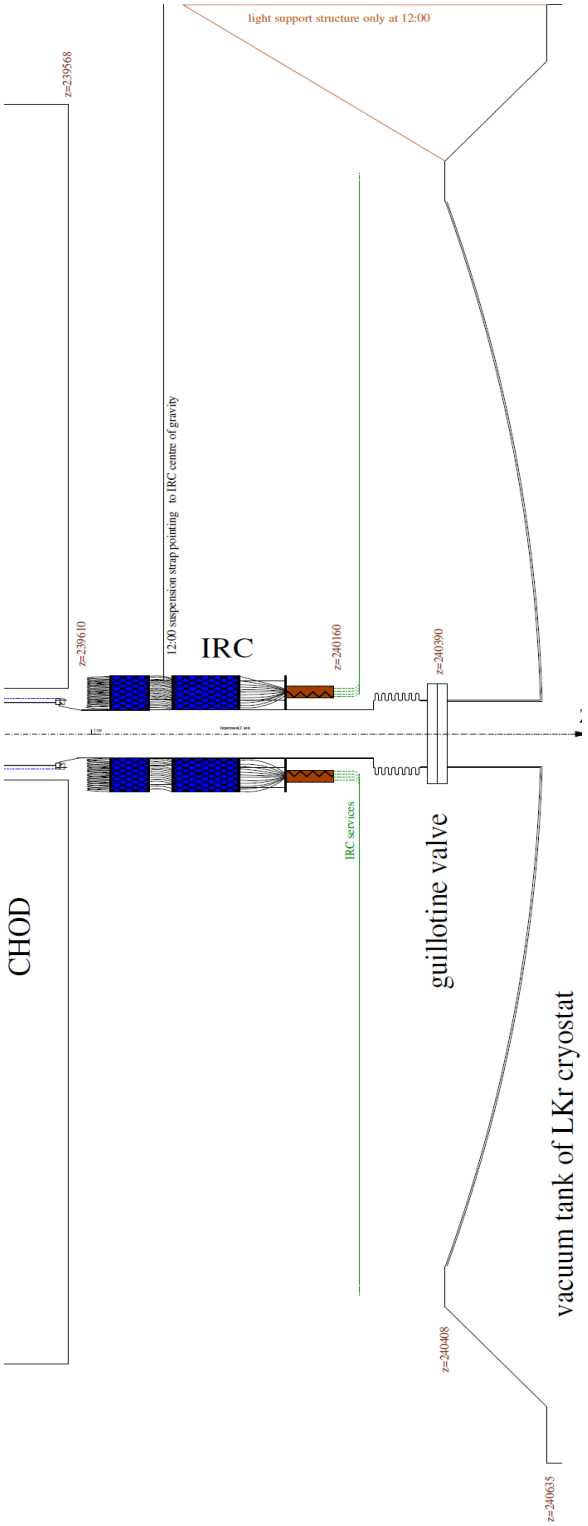

Figure 219 Integration of the IRC in the NA62 setup.

3.5 Straw Tracker

3.5.1 Introduction

The purpose of the STRAW Tracker is to measure with good accuracy the direction and the momentum of secondary charged particles originating from the decay region. The spectrometer (see Figure 220) consists of four chambers intercepted in the middle by a high aperture dipole magnet providing a vertical B-field of 0.36T. Each chamber is equipped with 1'792 straw tubes, which are positioned in four "Views" providing measurements of four coordinates (see Figure 221).

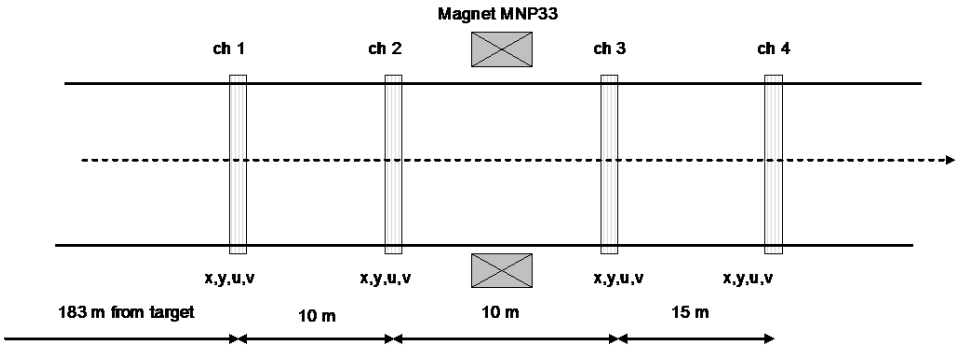

Figure 220. Schematic view of the magnetic spectrometer.

The main building block of the detector is an ultra-light straw tube which is 2.1m long and 9.8 mm in diameter. The tubes are manufactured from 36 μm thin PET[27] foils, coated –on the inside of the tube- with two thin metal layers (0.05 μm of Cu and 0.02 μm of Au) to provide electrical conductance on the cathode. The anode wire (Ø=30 μm) is gold-plated tungsten.

3.5.1.1 Straw Tracker Detector Requirements

As illustrated in the proposal [12], kinematical constrains in the events allow the rejection of the majority of the background provided the kaon and pion tracks are reconstructed with good accuracy. Two principal performance requirements –for secondary particles- follow from this:

$$\frac{\Delta P}{P} \leq 1\%$$

and

$$\Delta\theta_{K\pi} \leq 60\ \mu rad$$

[27] PET = polyethylene terephthalate

In addition to the above, it is important to stress that the overall physics performance of NA62 depends on a number of experimental necessities for the Straw Tracker:

- Use of ultra-light material along the particle trajectory in order to minimize multiple Coulomb scattering, in particular, near the first chamber.
- Integration of the tracker inside the vacuum tank.
- An intrinsic spatial resolution that allows a precise reconstruction of the intersection point between the decay and parent particle.
- Average track efficiency near 100%.
- Capability to veto events with multiple charged particles
- Sufficient lever arm between the four chambers allowing to re-use the exiting dipole magnet.

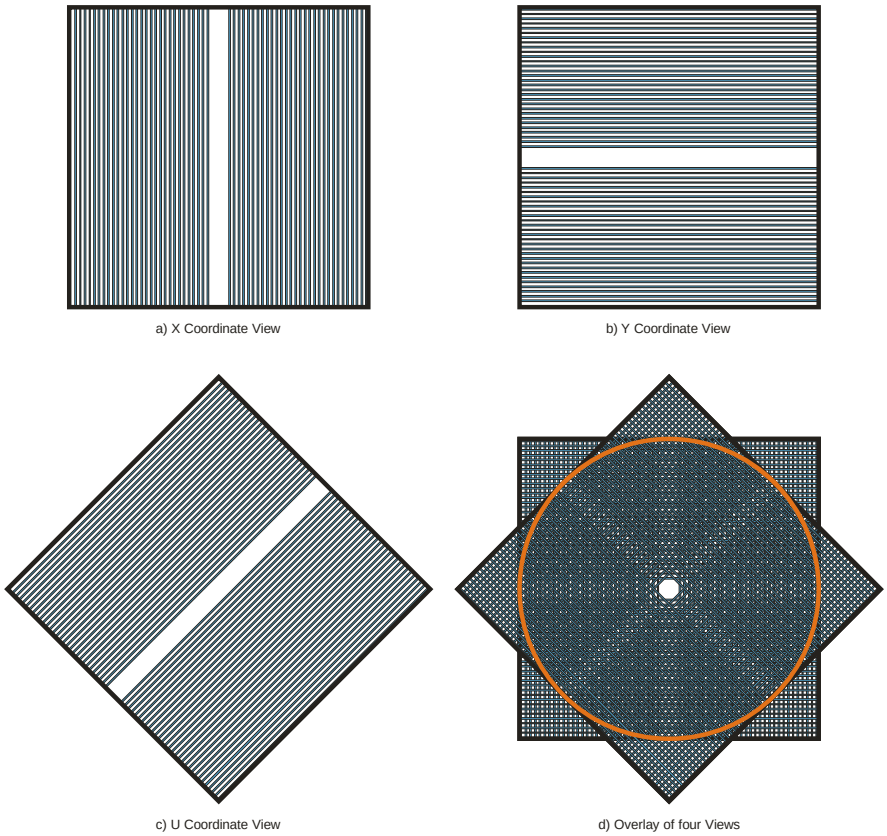

Figure 221. Schematic drawing of the four "Views" that compose each straw chamber. a) the x-coordinate view with vertical straws, b) Y-coordinate View with horizontal straws, c) the U-coordinate view (the V-coordinate view is rotate by 90 degrees compared U-Coordinate), d) A full chambers consisting of the X,Y,U and V Views; the active area of the chamber covers a diameter of 2.1m. The gap near the middle of each layer is kept free for the beam passage.

From these constrains follow the main requirements of the detector:

- Spatial resolution ≤ 130 μm per coordinate and ≤ 80μm per space point
- ≤ 0.5% of a radiation length (X_0) for each chamber
- Installation inside the vacuum tank (P = 10^{-5} mbar) with minimum gas load for the vacuum system $(Leakrate \leq 10^{-1}\ mbar \cdot l/s)$

- For straws near the beam, operation in a high rate environment (up to 40 kHz/cm, and up to 500 kHz/Straw)
- Contribution as multiplicity veto to the trigger

The most important parameters for the straw tracker of NA62 are summarized in Table 38.

Table 38. The main parameters of the straw tracker.

Chamber layout	**Value**			**Comment**
Number of chambers	4			
Number of views/chamber	4			
Number of straw layers/view	4			
Number of straws/view	448			
Central gap	~12 cm			
Central gap off-set vs. beam axis				Beam shifted towards the Jura side
	X	Y	U=V	
Chamber 1	101.2 mm	0	70.4 mm	
Chamber 2	114.4 mm	0	79.2 mm	
Chamber 3	92.4 mm	0	66 mm	
Chamber 4	52.8 mm	0	39.6 mm	
Track angle coverage	± 3º			
Straw length	210			Active length
Straw position accuracy	±0.1mm			
Wire	30 μm gold-plated Tungsten			Toschiba
Straw inner diameter	9.75 ± 0.05 mm			
Straw straightness	±0.1mm			
Maximum wire off-set	0.2 mm			
Gas volume in one straw	160 cm^3			
Straw material (option 1)				
Mylar film	36 μm			Hostaphan RNK 2600
Density	1.39 g/cm^2			
Copper layer	500 Å			
Gold layer	200 Å			
Material budget (1 view)	**Radiation length in %**			
Gas	0.010			
Straw wall	0.099			
Wires	0.0046			
Total	0.1136			
Straw operating conditions				Inner most straws
Wire tension	(90 ± 10) g			
Gas	70% Argon ,30%CO_2			Option: 90%CO_2,5%Isobutan,5%CF_4
High Voltage	1.75 kV			Option: 2.5kV
Gain	$1*10^5$			
Cathode resistivity	~70 ohm/			
Max Counting rate/straw	0.5 MHz			
High rate per unit area	40 kHz/cm^2			Straws close to center gap
Accumulated charge	0.015 C/cm/year			50 000 bursts of 3s
Maximum current/cm	64 nA/cm			
Gas flow per straw / per view	160 cm^3/h / 70 l/h			High flow straws
Gas pressure (absolute)	1.02-1.04 bar			
Nominal electron drift time	140 ns			
Nominal ion drift time	100 μs			
Effective radius	4.8 mm			>95% efficiency
Cross-talk	3%			Estimated
Nominal threshold	3 fC			5 fC
Termination resistor (far end)	330 Ohm			Estimated

3.5.1.2 Simulation of Straw Performance

3.5.1.2.1 Track Multiplicity Veto in High Rate Environment

To verify if the straw inefficiency is acceptable for the detector veto function, a fast and simplified Monte Carlo simulation of the possible background decays has been carried out. It was found that an overall inefficiency level of 5% does not lead to a significant change of the three-track veto efficiency with the current detector geometry. Decays like $K\rightarrow 3\pi^{+-}$ and $K\rightarrow\pi^{+}\pi^{-}\mu^{+}\nu$ with rather symmetric track distribution (caused by a small amount of kinetic energy), can be easily suppressed by the spectrometer multiplicity veto down to the level of 10^{-8} (limited by the simulated statistics of 100 million decays per mode).

In this respect the decay $K \longrightarrow \pi^{+}\pi^{-}e\,\nu$ (called Ke4) represents one of the most problematic backgrounds. With a low-energy electron and a negative pion that misses the first three chambers -due to the Straw Tracker beam gap- the veto relies on the measurement in the last Straw Chamber. If the pion is very close to the beam, it can occur that the measurement is done in only one single layer of the last chamber. From fast Monte Carlo simulation one can conclude that if the fourth chamber measures the negative pion from Ke4, with a straw inefficiency of better than 5%, it will suppress this background to an insignificant level.

The biggest track rates are expected on straws located near the beam gaps (Figure 221). Simulations show that the highest peak rates are in chamber 4 on the most inner straws on the Saleve side (positive X). However, more important is the efficiency on the inner straws on the Jura side, where Ke4 negative pions must be detected. Halo simulations, presented in Section 2.1.2.9 on page 39, predict a rate of 0.3 MHz due to kaons and pions inside and outside the fiducial region. Simulations of the other five dominating modes of kaon decays in the beam add another ≈0.08 MHz. Therefore, one can conservatively estimate the maximum rate for the important straws to be at the level of 0.5 MHz, which gives the final requirement for the straw rate capability. In these conditions the straw efficiency at the straw centre should remain at a level of 95%.

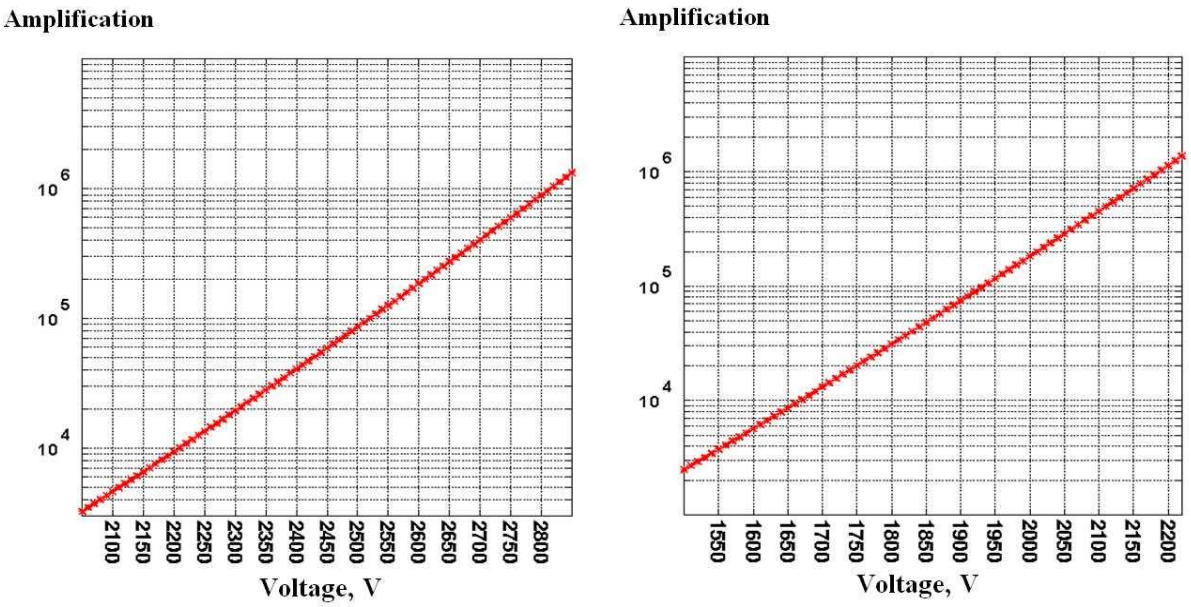

Figure 222. The simulated dependence of the gas gain on the high voltage for gas mixtures CO_2 90%, iC_4H_{10} 5%, CF_4 5% (left) and Ar 70%, CO_2 30% (right).

3.5.1.2.2 Spatial Resolution and Efficiency

GARFIELD was used to model a single straw tube [71]. The gas mixtures were CO_2 90%, iC_4H_{10} 5%, CF_4 5% (slow gas) and Ar 70%, CO_2 30% (fast gas). The dependences of the gas gain on high voltage are shown in Figure 222.

The dependences of spatial resolution and efficiency on the high voltage and thresholds are shown in Figure 223. The threshold for the accumulated charge from the anode wire was set to 4 fC, 10 fC, 15 fC and the resolution was calculated for a distance to the straw wire of 3 mm. The working points (high voltage) were defined on the basis of two criteria; the spatial resolution has not to exceed 40 µm for the gas mixture CO_2 90%, iC_4H_{10} 5%, CF_4 5% and 80 µm for the gas mixture Ar 70%, CO_2 30% and the straw efficiency should not be less than 99% for both gas mixtures. For the slow gas mixture, the working point is 2300 V for the threshold of 4 fC, 2400 V for 10 fC and 2500 V for 15 fC. For the fast gas mixture, it was 1850 V for 4 fC, 1950 V for 10 fC and 2050 V for 15 fC. The dependencies of the spatial resolution and the efficiency on the distance to the wire for both gas mixture and threshold 4 fC are shown in Figure 224. The resolutions and efficiencies almost do not depend on the thresholds, if the corresponding optimum working points are used. The r-t dependency is simulated and shown in Figure 225.

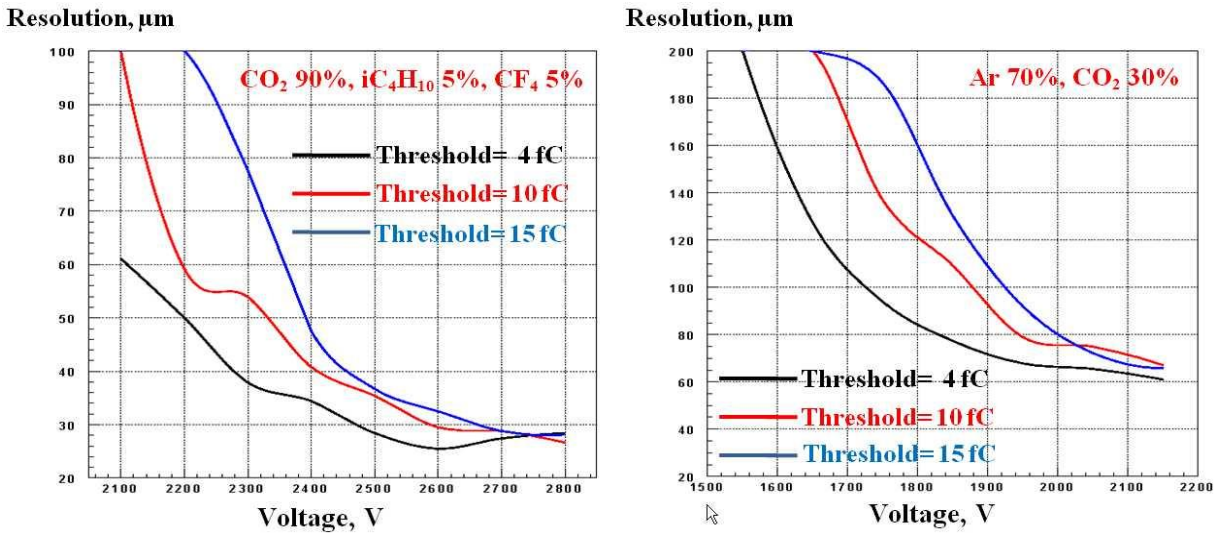

a)

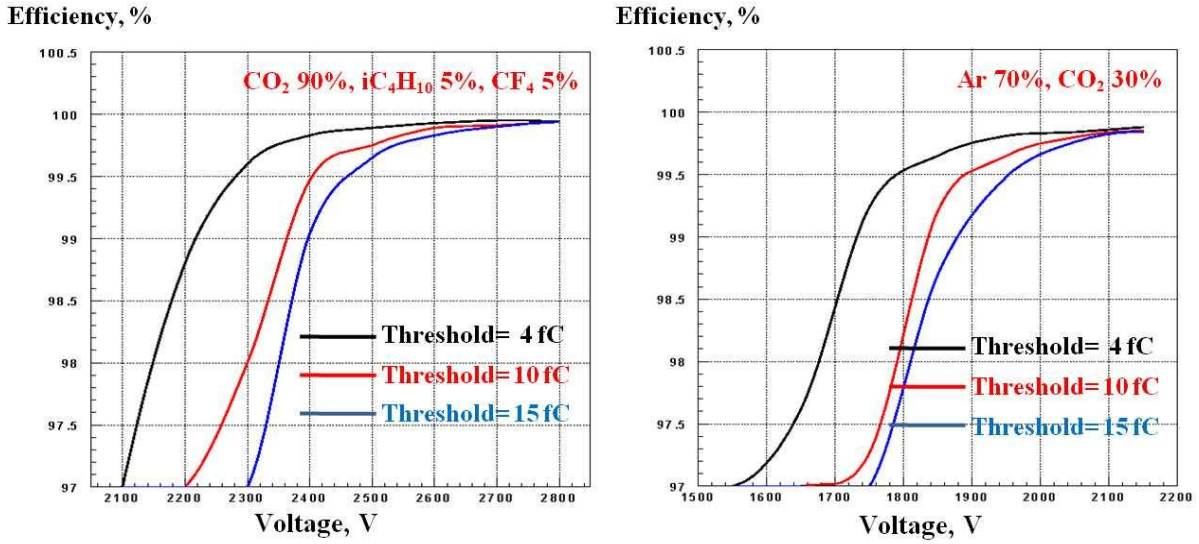

b)

Figure 223. The dependence of the spatial resolution (a) and the straw efficiency (b) on the high voltage and on the threshold for the gas mixtures CO_2 90%, iC_4H_{10} 5%, CF_4 5% and Ar 70%, CO_2 30%. The spatial resolution was calculated for a distance to the straw wire of 3 mm.

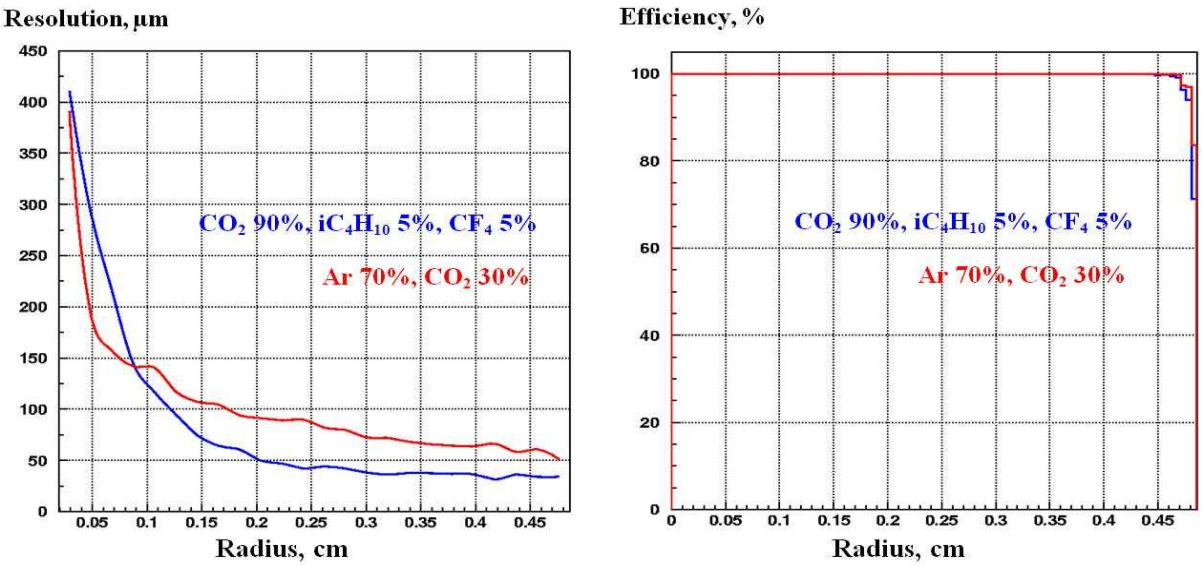

Figure 224. The dependence of the spatial resolution and the straw efficiency on the distance to the wire for gas mixtures CO_2 90%, iC_4H_{10} 5%, CF_4 5% and Ar 70%, CO_2 30%.

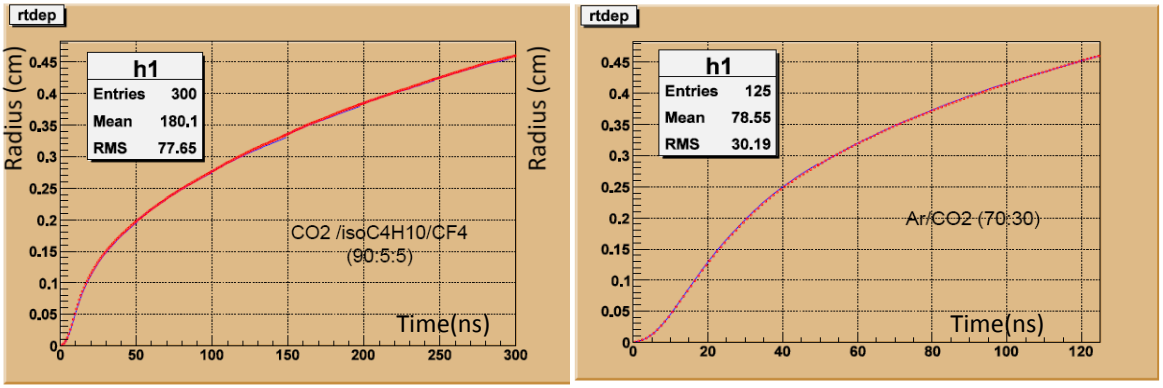

Figure 225. The r-t dependency for slow (left) and fast gas (right).

3.5.1.3 Conclusions from 48-Straw Prototype

To check experimentally the straw resolution and efficiency, a 48-straw prototype was designed and built. The 48 straws were arranged in 8 columns and 6 rows. This prototype has been successfully tested at CERN test beams in 2007, 2008 and 2009 runs with muon and pion beams. A dependency of residuals (that approximately represent the resolution) from track-wire distance for the gas mixture, was obtained in 2009 test run in muon beam as shown in Figure 226 and is qualitatively in agreement with the results from simulations [72].

Ar/CO2(70:30)

CO2/isoC4H10/CF4(90:5:5)

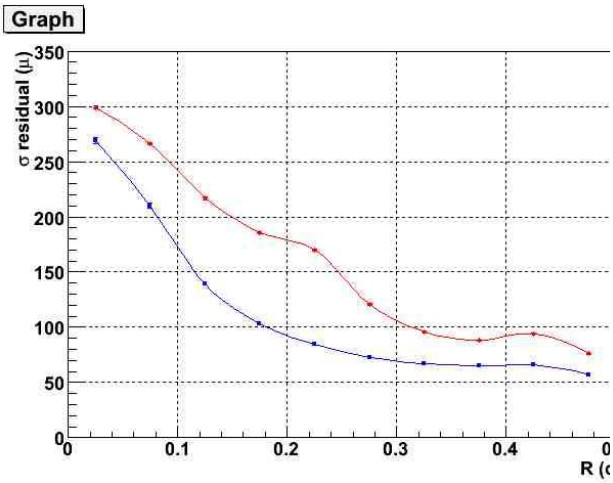

Figure 226. Dependences of the residuals from track-wire distance.

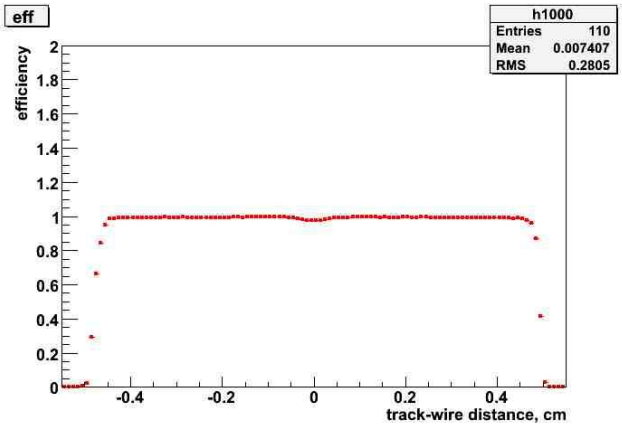

Figure 227. The measured of the efficiency in a single straw as function of the distance between track and wire.

The dependence of the efficiency for a single straw as a function of track-wire distance has been investigated in a pion beam using a CO2/isoC4H10/CF4 gas mixture. Due to the drop of the straw efficiency near the straw wall, the experimental resolution of the coordinate measurement affects the efficiency. The efficiency drop is smoother than in the results from the simulation, which makes difficult to measure precisely the effective straw radius. Nevertheless, the experimental results confirm the simulation and at the 50% level, the radius (almost unchanged by the coordinate smearing) corresponds to the same radius as in simulation (close to the value of the straw physical radius). The result from measurements is shown in Figure 227.

The dependence of straw efficiency in the straw center on the beam intensity has been investigated in pion beam, using CO2/isoC4H10/CF4 gas mixture. The result is shown in Figure 228. The efficiency is changed from 99.5 % at 4 kHz down to 97.4 % at 520 kHz. The exact reason for this drop in the efficiency is under study.

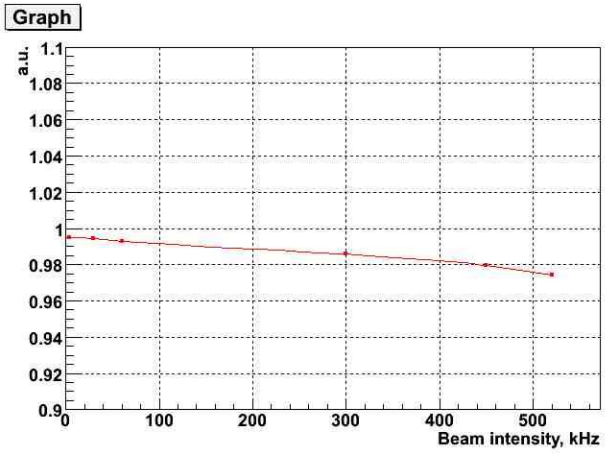

Figure 228. Straw efficiency versus particle rate in a straw.

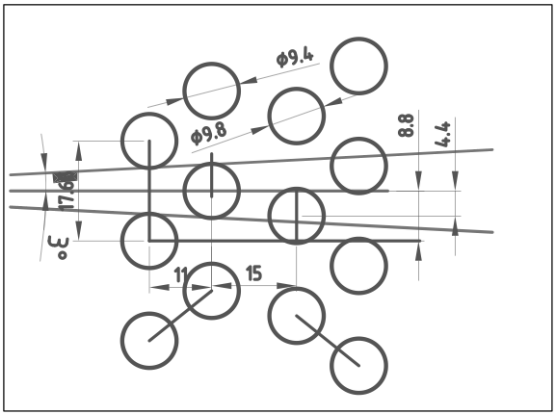

Figure 229. Straw layout in one view (beam direction from left to right). The distance between the straws in one layer is 17.6mm.

3.5.1.4 Straw Chamber Layout

The straw layout in a section of one view is shown in Figure 229, to solve the right-left ambiguity, each view has four layers. There are 112 straws in each layer and the distance between the straws in one layer is 17.6 mm. It is enough to guarantee at least two hits per view and be efficient for track angles from -3 to 3 degrees.

3.5.2 Wire Centering and Wire Off-Set

Due to the length of the straw and the high voltage on the wire, a calculation of the effect from the electrostatic forces on the wire offset was carried out. The wire offset was calculated using the formula:

$$y = \frac{C}{k^2}\left(\frac{1}{\cos(kL/2)} - 1\right) \tag{1}$$

where C is:

$$C = \rho g \sigma / T \tag{2}$$

and k^2 calculated from the expression:

$$k^2 = 2\pi\varepsilon_0 E_0^2(b)/T \tag{3}$$

[73] where

ρ =: is the density of the wire
σ =: the cross-section of the wire
T =: is the straw tension
E_0 =: the field at the straw wall
L =: the length of the straw
ε_0 =: has the value 8.854×10^{-12} As/Vm

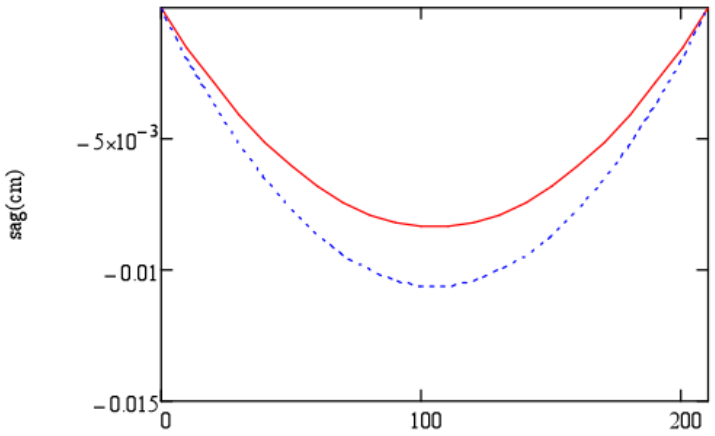

Figure 230. The wire deflection for a 2.1m long straw with (blue) and without (red) electrostatic deflection at 2.5 kV.

The results plotted in Figure 230 shows an amplification of the wire sag of ~27%, due to the electrostatic forces. The voltage was set to 2.5 kV (slow gas) and the tension in the wire to 90 g. For the fast gas with a high voltage of 1.8 kV, the corresponding increase in the wire sag is 9%.

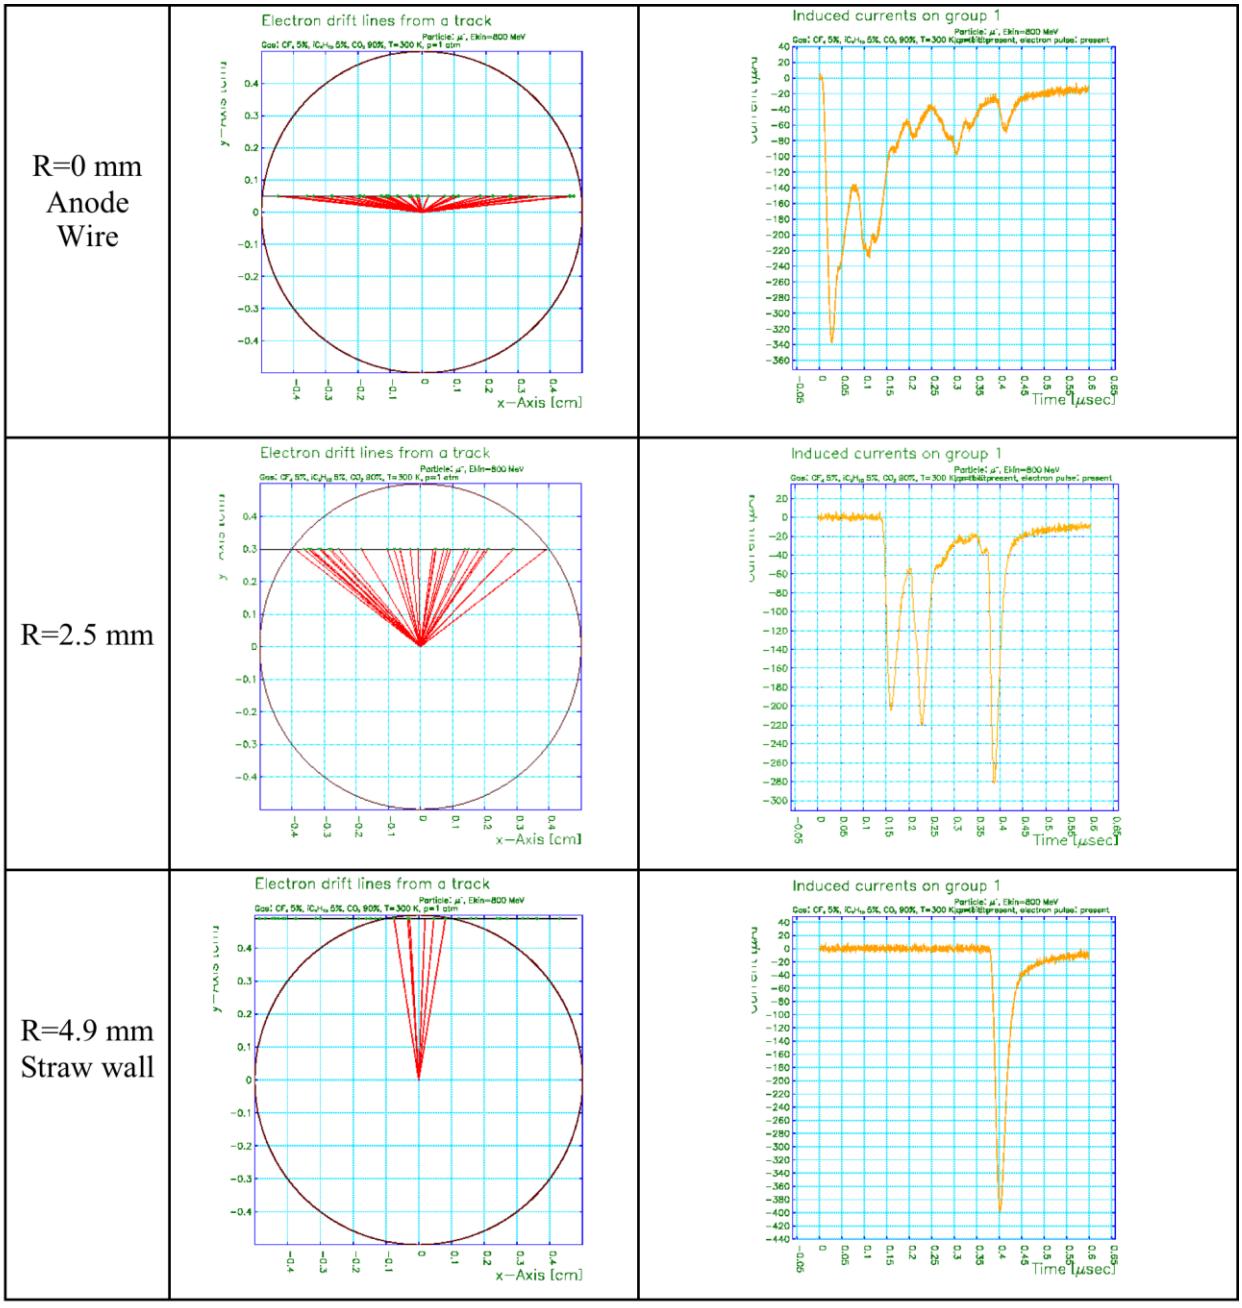

Figure 231. Signal shape for muons for different distances between tracks and straw centre (simulation using Garfield).

3.5.2.1 Straw Signal

The straw signal shapes for muons were simulated with Garfield and a few examples are shown in Figure 231. The straw diameter is 10 mm and the anode wire diameter 30 µm. The gas mixture is CO_2/isoC$_4$H$_{10}$/CF_4 (90:5:5) with a gas gain of about 10^5. A track contains a number of clusters with 1-3 primary electrons in each. The average number of clusters for a 1 cm long track is about 75.

The characteristic impedance of a straw as transmission line and as a function of frequency, is calculated using the following formula (see Figure 232).

$$Z(\omega) = \sqrt{\frac{R + i \cdot \omega \cdot L}{G + i \cdot \omega \cdot C}}$$

With $R = R_{wire} + R_{straw}$, , G = conductance between the straw and the wire , $L = 1.16\,\mu H$ and $C = 9.6\,pF/m$. At high frequencies the impedance can be expressed as $\sqrt{L/C} \approx 350\ \Omega$.

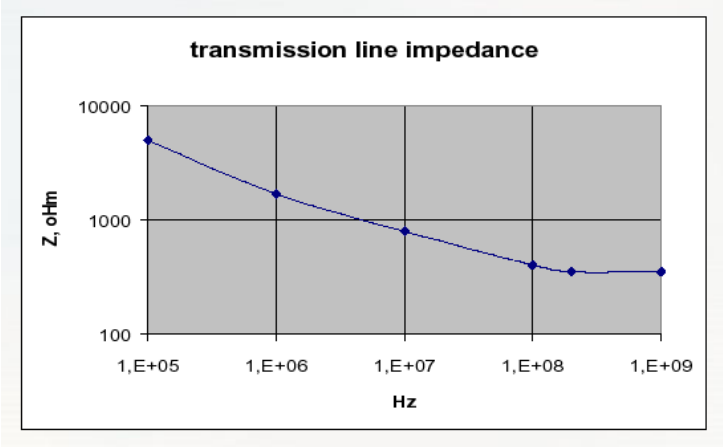

Figure 232. The dependence of the straw impedance from the frequency.

An example of the signal reflection study by the Spice program is shown on Figure 233. The current division and the signal attenuation are clearly visible. The signal reflection phase depends on the value of R1 and Rs+Rin.

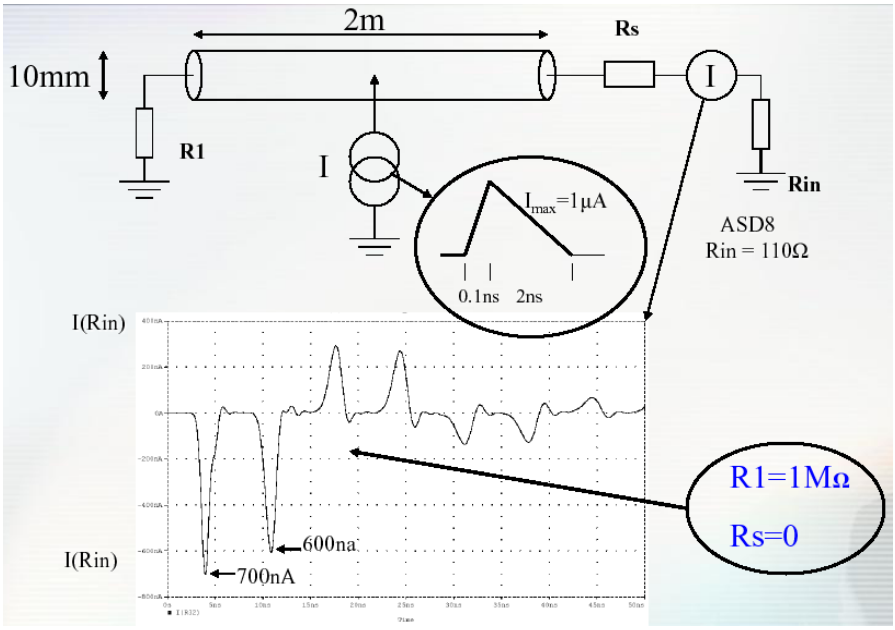

Figure 233. Signal reflection simulation with SPACE.

The result of the signal termination optimization is presented in Figure 234. A resistor R1 of 350 Ω cancels the reflection at the "far" straw end. With Rs=0 Ω the signal amplitude is higher by factor 1.5, compared to a termination of 240 Ω.

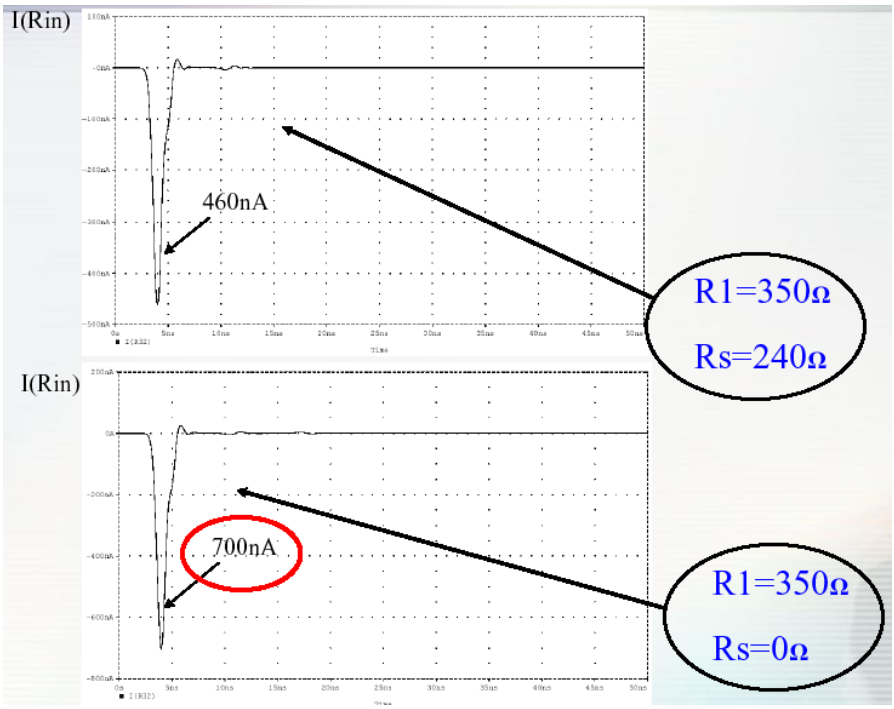

Figure 234. Optimal signal termination for the straw.

The signal reflections in the straw have been studied experimentally. A schematic of the set-up is presented in Figure 235. The straw diameter is about 10 mm, the anode wire diameter 30 μm, the straw length 210 cm, the straw capacitance 20 pF and the straw inductance 2.5 μH. Furthermore, the anode wire resistance is 250 Ω and the cathode resistance ~70 Ω. A fast current

amplifier (rise time about 10 ns) with a gain of 25 mV/μA and input impedance of 120 Ω was used. The straw termination resistor had the following values: 0 Ω, 250 Ω and 1 MΩ.

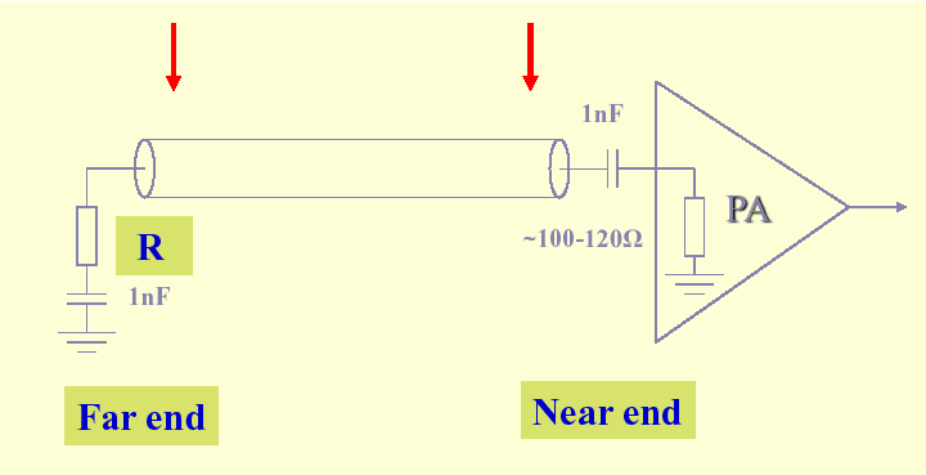

Figure 235. Set-up for straw termination experimental study.

A radioactive source Fe-55 (gamma, E=5.9 Kev) was used. The influence by each option of straw termination on the amplifier signal shape and its value is presented in Figure 236. The optimal straw termination (no signal reflection) is for R=250 Ω (C blocking=1000 pF). The signal attenuation value from the "far" straw end is about 55%.

The straw acts like a transmission line with the characteristic impedance of 350 Ω at high frequency. The straw "dead" time increases (reflections are up to 100 ns long) and at high rates, the reflections may deteriorate the signal. Experimentally it was shown that the "far" end should be terminated with a resistance of 250 Ω, if the capacitance of the blocking capacitor is 1000 pF and at the "near" end, there is no termination resistor (Rs=0 Ω). The signal attenuation along the straws is 55%. However removing the signal reflections with the termination also decreases the signal amplitude. The final choice on the straw termination will be done once the data from the 2010 test beam is analyzed (see also 3.5.110).

3.5.3 The Straw

The most important constituents of this detector are 7'168 Straws. Each straw is the cathode of a drift tube that is capable to operate inside vacuum. A straw tube is manufactured from thin polyethylene terephthalate (PET) foil that is longitudinal welded (ultra-sonic weld) to form a tube. The straw has an active length of 2.1m long and a nominal diameter of 9.8mm[28]. It's inner surface has a metal coating (Cu/Au) to provide electrical conductance on the cathode.

[28] The exact diameter of the straw depends on its overpressure. 9.8 mm corresponds to the diameter without any overpressure (manufacturing conditions).

R=250 Ω	Amplifier end 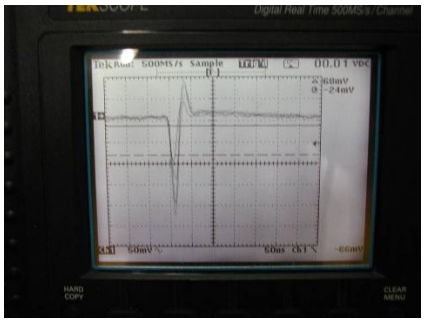	“Far” straw end 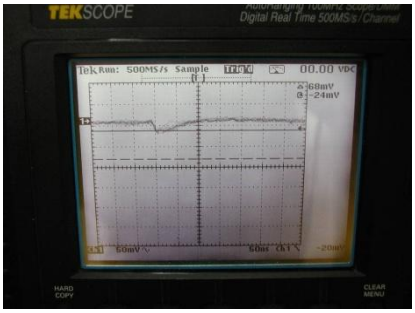
R=1 MΩ	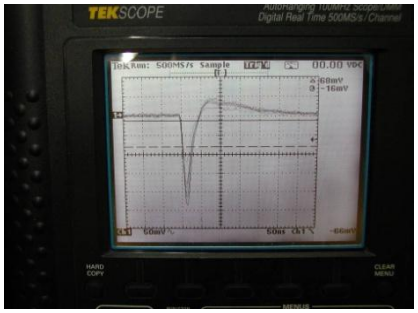	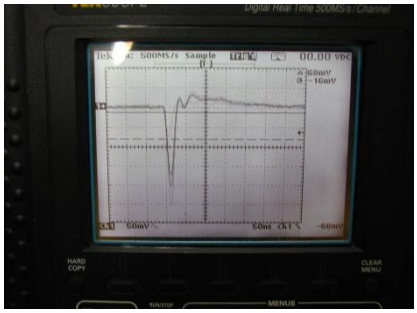
R=0 Ω	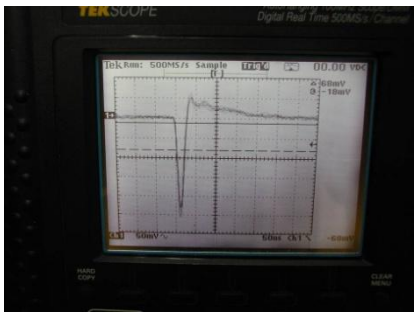	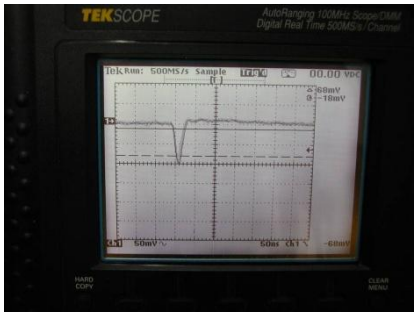

Figure 236. The signal shape for different straw terminations.

3.5.3.1 Material

The choice of the straw material is a compromise between many different requirements, i.e. radiation length, permeation of gases, mechanical properties, adhesion to metal coating, bonding with epoxy and the ability to be ultrasonically welded into a tube. A summary of the specifications is given in Table 39.

Table 39 Summary of coated material specifications

Description	Specifications
PET film	polyethylene terephthalate type Hostaphan RNK 2600 of (36 ± 2.0) μm thickness
Density	1.4 g/cm³
Copper layer thickness	50 nm
Gold layer thickness	20 nm
Resistivity of finished straw (2.1m)	~70 Ω
Permeation of naked film (10^{-12} Torr·l·cm/s·cm^2·Torr)	6 for He, 0.06 for Ar, 1 for CO_2

The straw tubes for the tracking detector in the NA62 experiment are manufactured from a biaxially oriented coextruded film made of polyethylene terephthalate (PET). One side of the film is chemically pre-treated to improve adhesion; this side was chosen for the epoxy bonding between the straw tube and the (PEI) polyetherimide straw fixation plug. The thickness of the film is 36 μm, and the non-treated side is coated by a conductive layer of 50 nm of Copper (Cu) followed by a protective layer of 20 nm of gold (Au). Once the tube is manufactured, the conductive layer, is found on the inner diameter of the tube, while the outer diameter remains uncoated, but chemically pre-treated for gluing.

3.5.3.2 Long Term Tests

Long-term tests were carried out with straws mounted under tension with a load of 15N applied before the extremities of the tubes are bonded to the PEI supporting parts. The sag and the tension of the tubes were measured during a long-term test (18 months). Values are recorded with straws under pressure (in order to simulate the final experimental set-up conditions in vacuum) and compared to a reference straw at atmospheric pressure.

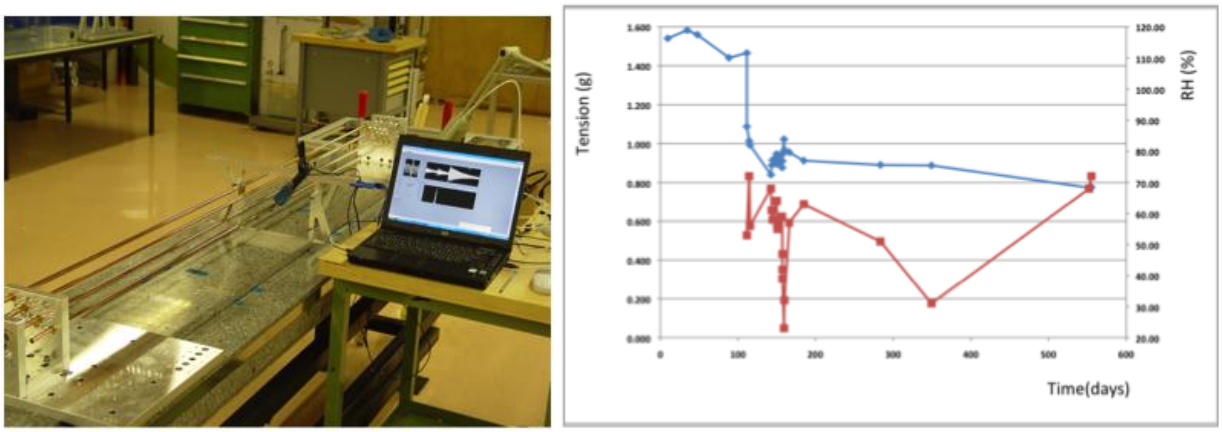

Figure 237. A view of the long-term test apparatus (left) and the straw tension (blue curve) over time. The red curve indicates the relative humidity.

3.5.3.3 Tensile Test of PET samples:

The base material for the straw tubes was tested on CERN's Instron tensile test apparatus. The values observed for welded samples are typically 25 % lower than those of the coated un-welded PET. Un-welded PET has a tensile strength close to 130 MPa whereas welded and coated samples fail at ~100 MPa. The stress in the straws due to the pressure difference (~1 bar) is 13 MPa which gives sufficient safety margin.

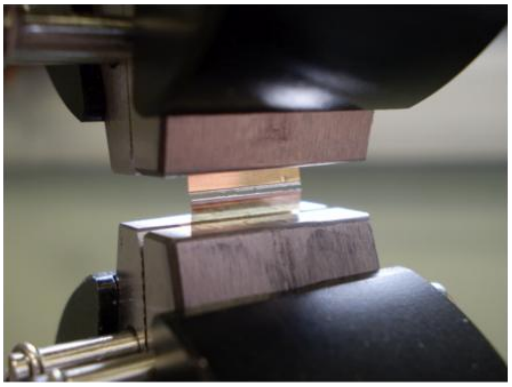

Figure 238. Close-up view of a tested sample (the weld in the center).

3.5.3.4 Coatings and Permeation:

The metal coatings are accomplished by a sputtering process; the metallisation provides electrical conductivity as well as an improvement in the permeation properties of the PET substrate. Measured values of permeation through a 50 nm Cu coated membrane show a BIF (barrier improvement factor) of 8.5 for He (helium); 9 for CO_2 (carbon dioxide) and 5 for Ar (argon). The quality of the sputtered coating is measured by a peel test using adhesive tape.

3.5.3.5 Scanning Electron Microscope Investigations

The coated PET was observed using scanning electron microscope (SEM); this showed the quality of the coating and sputtering on the substrate. The SEM was also used to verify the metal coating thickness, not as a direct measurement, but by using the EDX (energy-dispersive x-ray spectroscopy) facility and a third party software.

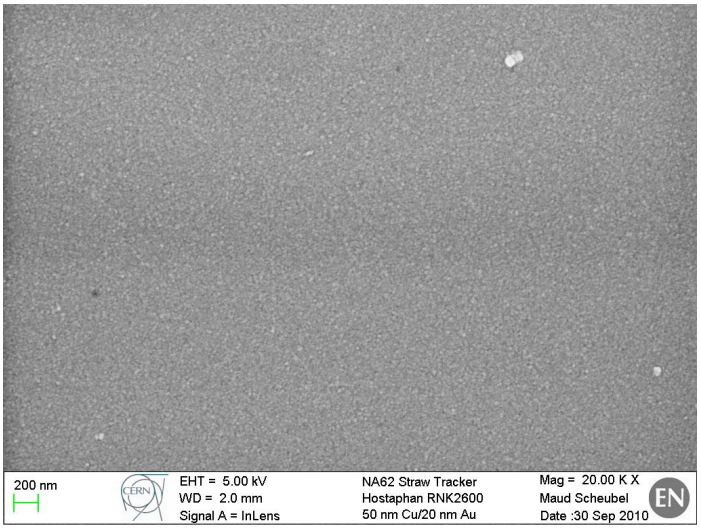

Figure 239. A typical view of coated PET with a magnification of 20 000 (RNK 2600).

Metal coating and weld quality of the seam was checked using both SEM and traditional optical microscopy (see Figure 239 and Figure 240)

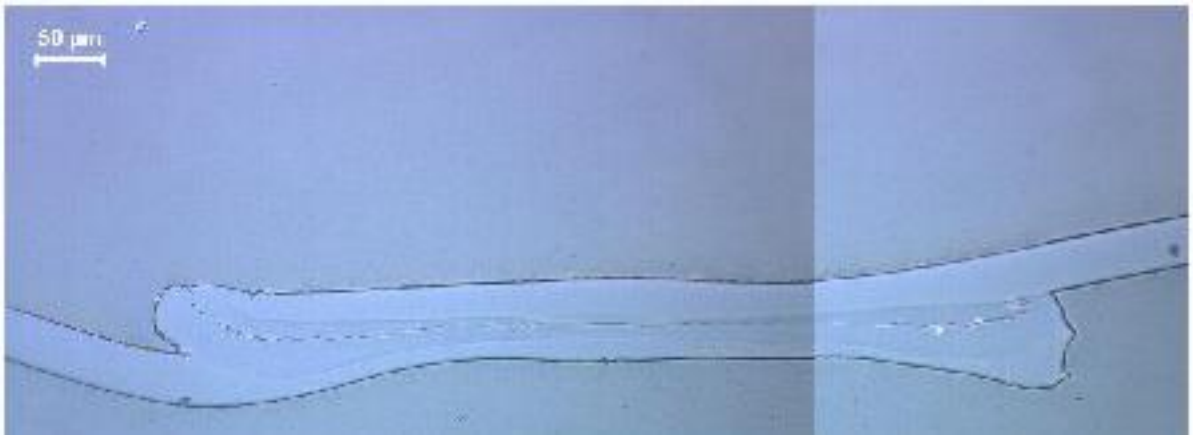

Figure 240. Cross section of a welded seam where the two sides of the strips are fused.

3.5.3.6 Glue Bonding Test

The adhesion between the outer diameter of the straw tube and the PEI straw support is critical. The selected epoxy resin for this operation was TRA-CON's TraBond 2115 who's out gassing properties and radiation hardness were studied at CERN. This epoxy has a good resistance to shocks and temperature cycling. Its pot life is close to two hours and curing in 24 hours. Samples of PET were peel tested after being bonded to an aluminium plate. The bond of the selected candidate was greater than the natural resistance of the material. The selected material was also analysed by X-ray photoelectron spectroscopy. This allowed us to better understand the chemical treatment made by the PET suppliers and its effects on bonding.

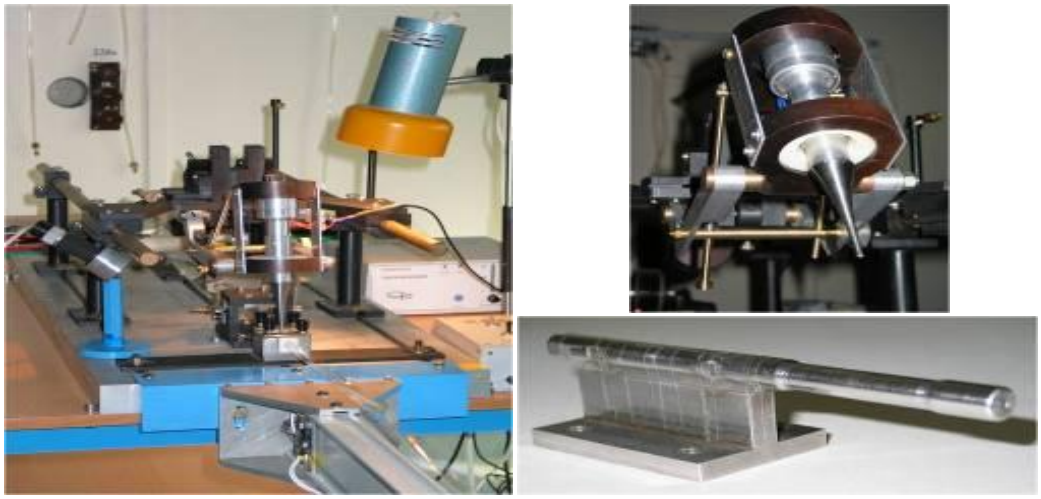

Figure 241. A new machine for straw production (left), new ultrasound head (right top) and anvil (right bottom).

3.5.3.7 Straw Manufacturing and Welding

A dedicated machine for straw production has been designed and constructed to manufacture the 7168 straws (see Figure 241, left). An ultrasound head with movable fixation at the frame (the Figure 241, right top) and anvil for straw welding (see Figure 241, right bottom) were designed and constructed to obtain a good quality of weld seam (see Figure 242) and to minimize differences in straw diameter. The distribution of the straw seams widths and tube diameters obtained during the production of a large set of straws are shown in Figure 243 and Figure 244. After manufacturing, each straw is equipped with glued end plugs and leak tested. All measurements are individually recorded in a logbook where all straw characteristics are stored like material, manufacturing date, inner diameter, etc. A unique serial number is allocated to every straw tube and a complete feedback is readable at any moment of the detector assembly.

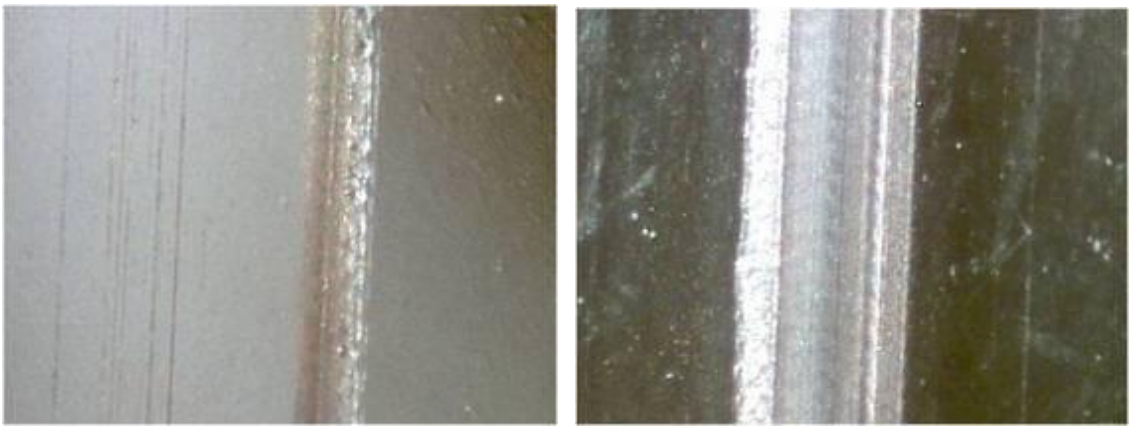

Figure 242. Quality of the weld seams (0.4 mm – left picture, 0.85 mm – right picture); 60 times magnification.

3.5.3.8 Straw Conditioning

Once the ultrasonic welding apparatus produces the straws, they need to be carefully stored, prior to installation into the detector. The straws need to be stored during a period of time that varies from several weeks to several months. The tube storage has several constraints that need to be addressed:

The inner diameter of the tube must remain free of dust, and of foreign particles.

The tube must also have the ability to be manipulated by technicians without buckling.

With the above constraints in mind, it was decided that the optimal way to store the tubes would be under pressure. This would also have the added value that the changes of pressure in a tube with respect to time could be observed. Straws that show gas leakage and have poor weld quality can be removed. Taking into account further quality control tests that are performed on the tubes before installation, and the manipulation of the tube itself during its installation into the detector, a dedicated system was developed, in the form of valves as shown in Figure 245.

For storage, the straws are gas filled with an approximate overpressure of 1 bar. The loss of pressure over time is measured indirectly by measuring the local straw deformation under a 300g weight and comparing with the day it was first pressurized. The first valve must have the ability to let gas into the straw, then act as a non-return valve, thus keeping the straw under pressure. The valve design is loosely based on the “Dunlop” / “Woods” valve design, with an additional feature. To prevent leakage through the non-return system, an O ring and threaded screw have been added. This is also used to fit the gas connector.

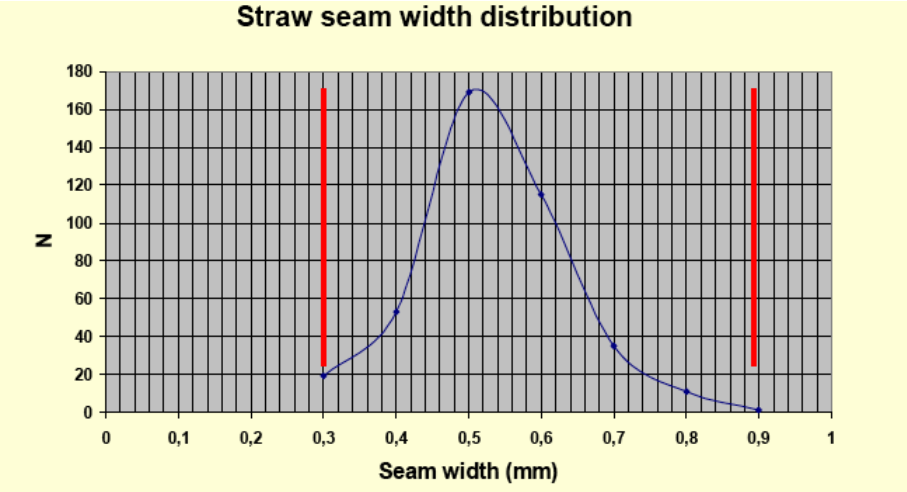

Figure 243. Weld seam widths distribution.

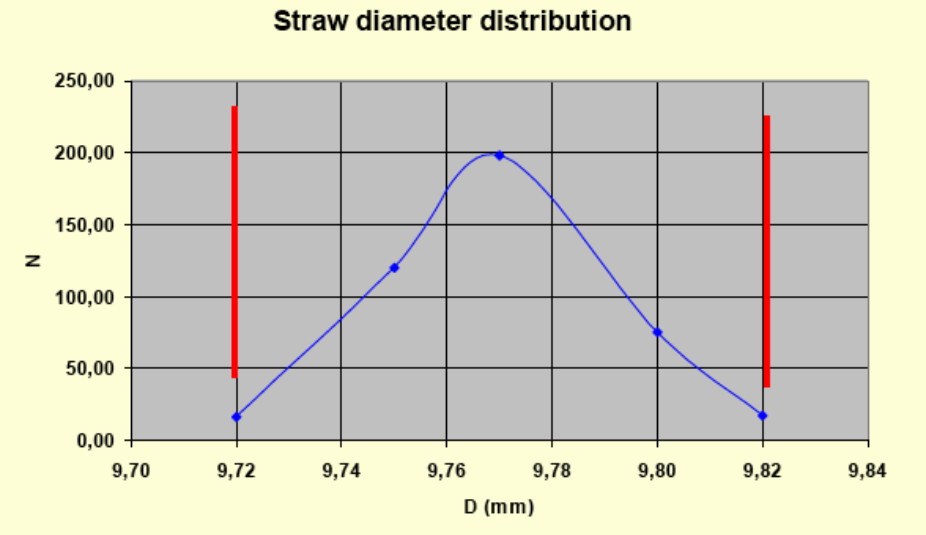

Figure 244. Straw diameters distribution

Figure 245: The PEI non-return valve (left) and the opposite side with O ring and an M4 screw.

The second valve, used on the other extremity of the straw tube, it is a simplified version without the non-return mechanism. To run a leak test, we simply connect the test system to an M4 thread inside the valve and pressurise the straw. To simplify Sagitta measurements of the simply supported tube both valves are of the same weight. In addition, the ends have the ability to be connected to standard M4 threads. Furthermore, both valves have an outer diameter smaller than the straw tube outer diameter, which allow the straw tube to be installed under pressure. This facilitates the manipulation of the straws and limits the risk of buckling the straw wall and the welding seam.

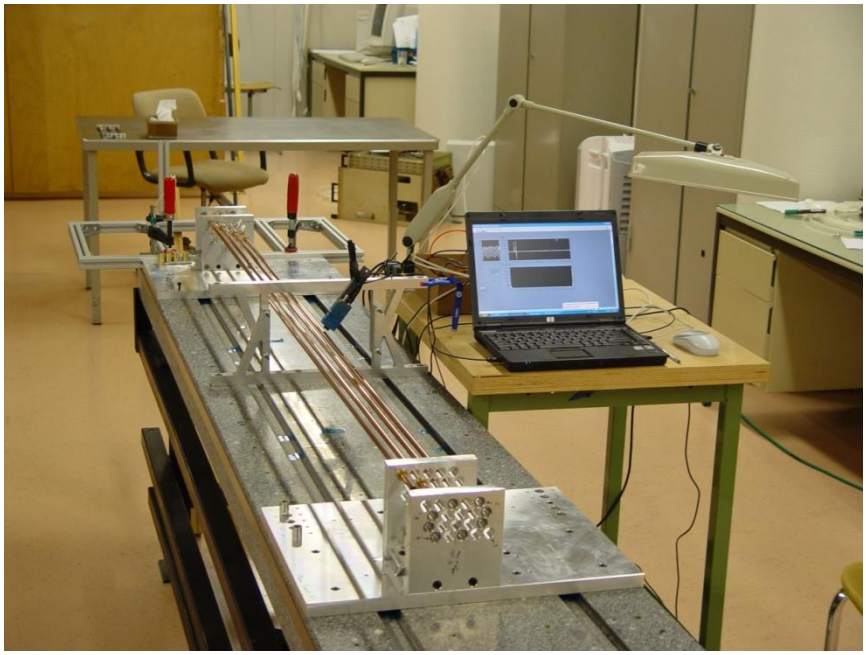

Figure 246. Set-up to measure the straw sag and frequency as a function of different pre-tension and pressure. The computer is connected to an IR emitter and receiver (OPB732) to measure the frequency of the vibrating straw.

3.5.3.9 Mechanical Properties and Pre-Tension of the Straw

3.5.3.9.1 The Measurement of the Straw- Sag

In order to determine the necessary straw pretension, a dedicated set-up was designed to measure the straw sag and the straw frequency, as a function of pressure (see Figure 246). In order to simulate the effect of the vacuum, the straw was connected to a gas bottle and the absolute pressure inside the straws was set to 2 bar. Four straws were glued horizontally into two plates. One straw was equipped with a sliding fixation on one end to allow for variable straw pre-tension. However, the sliding end of the straw was blocked (clamed ends) before measuring the deformation (sag) and frequency. The results are shown in Figure 247. The increase of the sag, as the pressure inside the straws increases, is clearly shown.

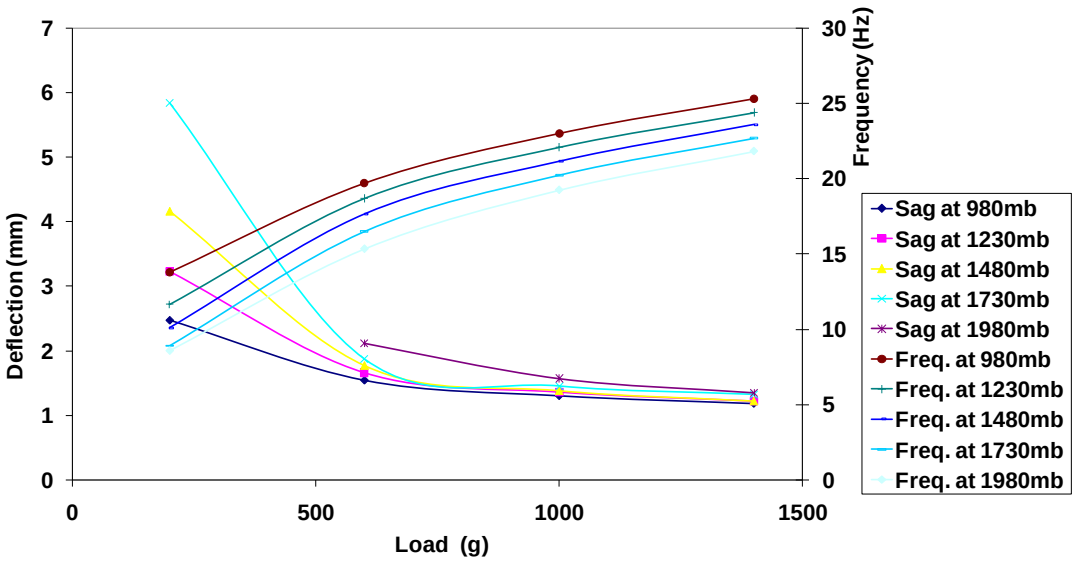

Figure 247. Measured sag and frequency of the straw as function of pre-tension at different over pressures.

3.5.3.9.2 Pressure Influence

A study of the straw sag as function of overpressure at different pre-tension was carried out to determine the required straw tension. The results are shown in Figure 248. A minimum tension of 10N is necessary to obtain an acceptable straw straightness. We have decided to apply a force of 15N during installation in order to allow for some loss of tension over time [74].

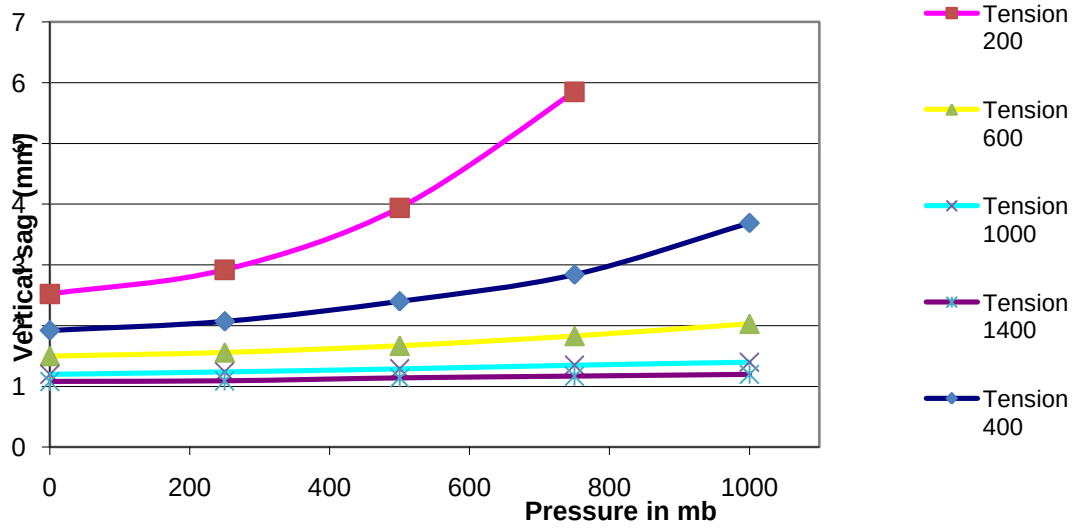

Figure 248. Straightness as function of pressure for a 1.85 m long horizontal straw for different values on the pre-tension (g).

3.5.4 Chamber Design, Construction and Installation

3.5.4.1 Detector Geometry and Layout

The detector consists of four stations and each station has four views and gives four measuring points (x,y,u,v) as shown in Figure 249. The mechanical structure of the chambers requires good dimensional stability, rigidity and strength. The active volume of the detector is 2.1 m x 2.1 m and each straw plane has a 12 cm wide region for the beam in the middle of the frame without straws. A 3D drawing of a chamber is shown in Figure 250.

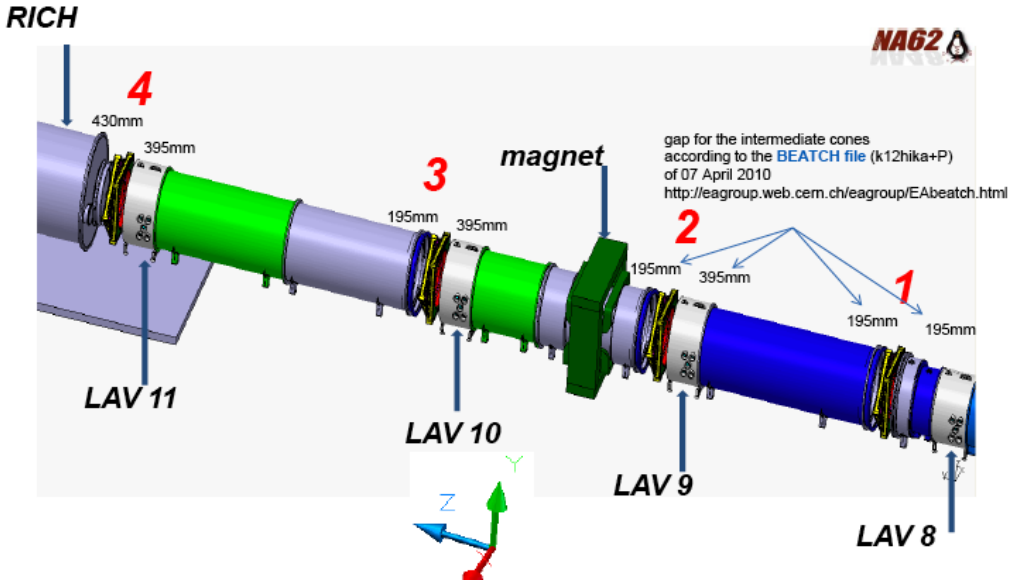

Figure 249. Layout of the straw spectrometer in ECN3 with the straw chambers marked 1,2,3 and 4.

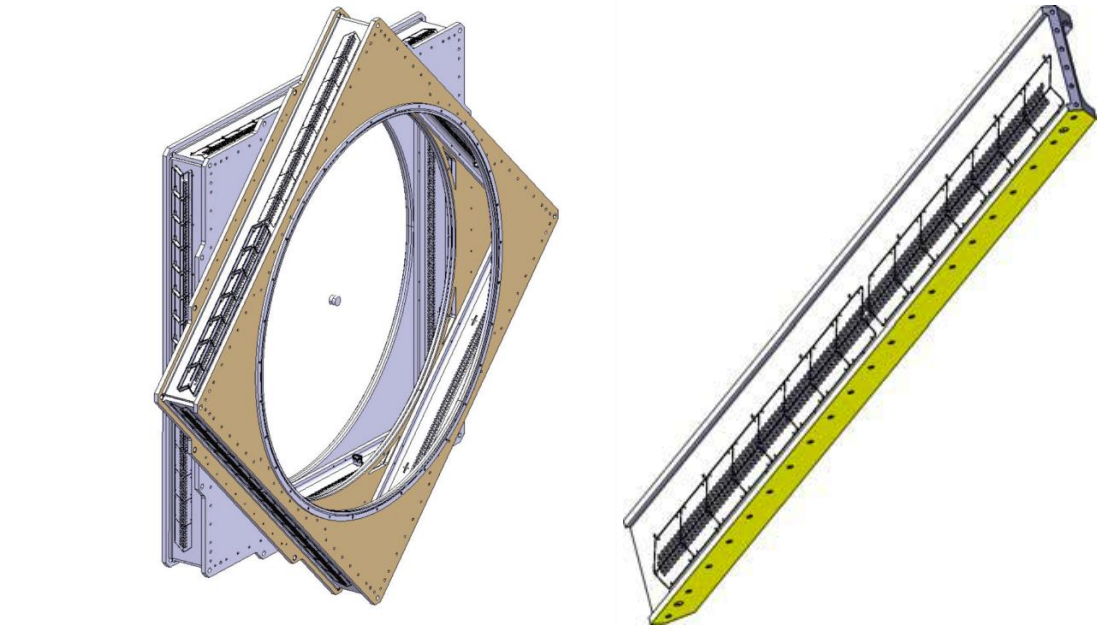

Figure 250. One chamber consisting of two modules (left). Detail of the straw fixation beam (right).

The chambers consist of two modules and each module is built from six parts; four beams for the fixation of the straws and two flanges. The basic element during the detector assembly and testing is the module.

3.5.4.2 Mechanical Validation of the Design

In order to dimension the mechanical structure of the detector, a finite element model (FEM) of the chamber and of the interface parts (between chamber and vacuum tube) was built. The aim is to calculate the deformation of the chambers and to verify that the stress levels are acceptable in the different parts. In addition, the model is used to determine the dimensions of the fixation bolts.

The following loads have been applied to the chambers and interface parts for the calculations:

- A straw tension of 14.7 N/straw
- A pressure difference of 1 bar between the outside and the inside (straw volume) of the chamber
- An axial force of 30 t on the interface parts, originating from the atmospheric pressure pushing on the end cap of the decay tube (a plate of ~ 2m in diameter)
- The gravitation
- Note: The wire tension of 0.9N/wire was neglected.

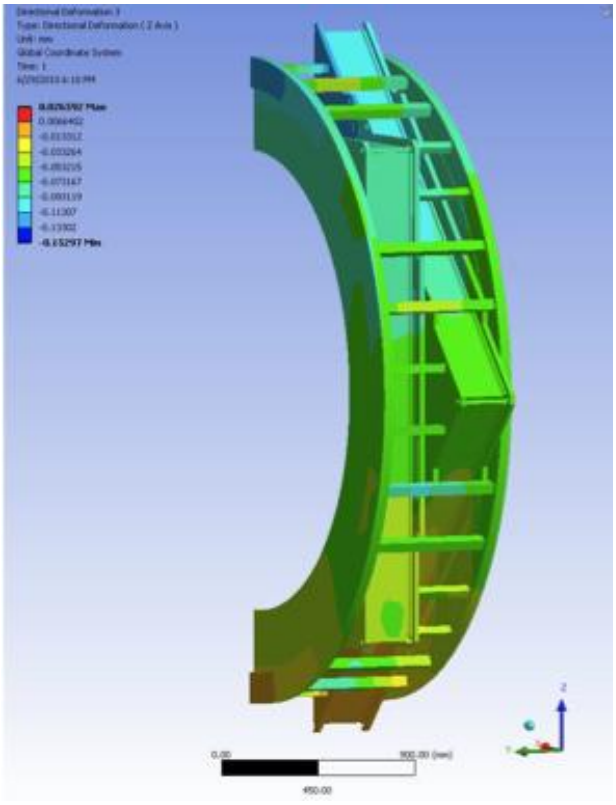

Figure 251. Calculation of the chamber deformation.

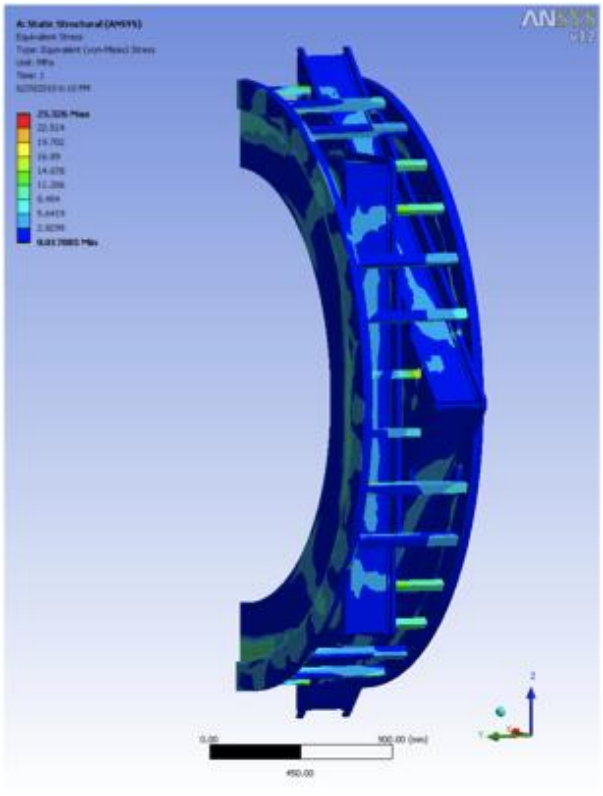

Figure 252. Calculation of the equivalent stress levels.

The resulting deformations and the stress levels are shown Figure 251 and Figure 252. The maximum deformation is 0.2mm and the highest stress level is calculated to < 10 MPa. For comparison it can be

mentioned that the elastic limit is around 145 MPa. In this context it should be pointed out that some safety was privileged at the cost of a somewhat higher weight. The total weight of a chamber is 1200kg (2x 600 kg).

3.5.4.3 Structure Assembly and Tightness

One of the challenges in the NA62 tracker construction is the structure gas tightness. The chamber will be built using four "H" beams linked to two octagonal flanges, as shown above in 3.5.4.1 The whole chamber will be bolted together and then sealed by a two-component epoxy resin in the joining surfaces. A prototype of the corner was built in order to validate the assembly technique (see Figure 253).

Figure 253. Prototype detector's inner corner (left) and a set-up verify leak tightness (right).

3.5.5 Detector Components and Assembly Principles

3.5.5.1.1 Active Web and Wire Connection

A cross-section of the mechanical structure holding the straws is shown in Figure 254. An Aluminum frame has precisely drilled holes for precise positioning of the straws and wires. The modularity in terms HV, gas supply and readout is 16 straws. Removable partitions for the gas manifold and a dedicated polyurethane joint for gas tightness was developed. The partitions also serve as supports for the so-called cover, which contains the front-end electronics (FE), high-voltage and gas connections. Between the straws and the frame are glued insulating sockets in polyetherimide (PEI) (see Figure 255).

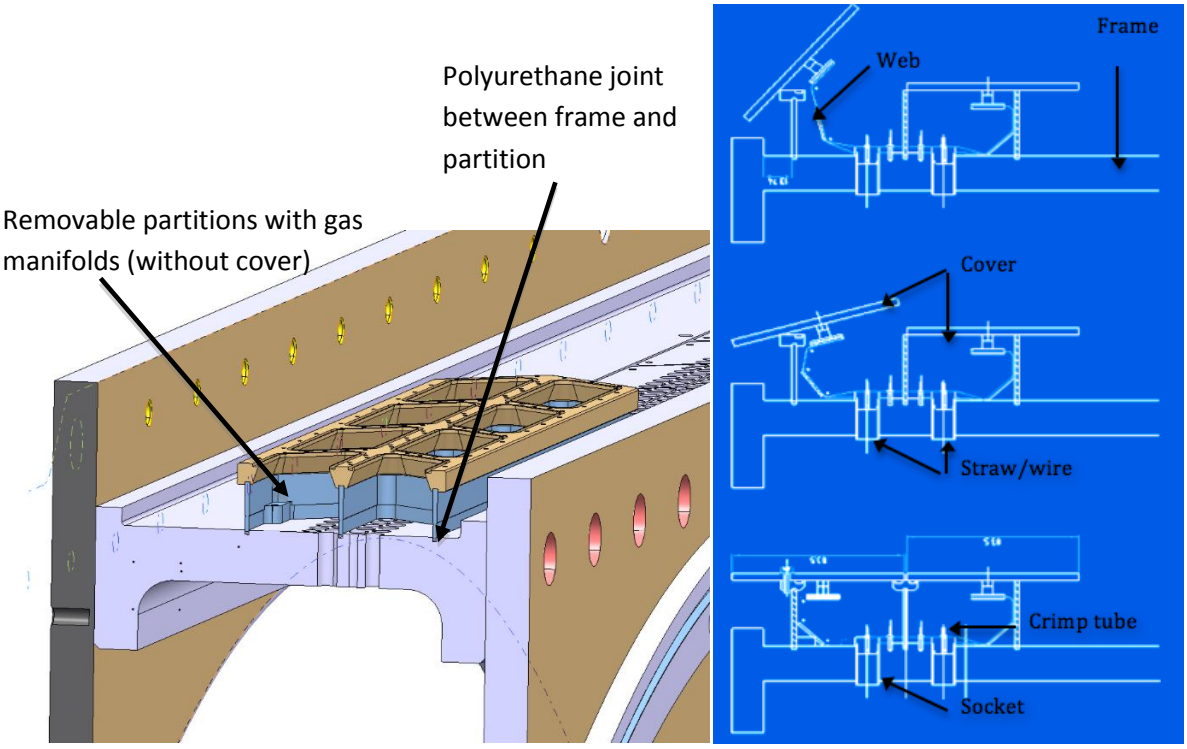

Figure 254 Principle cross-section showing detector frame with gas manifold and cover (left). A schematic view showing how the web connects to the straws and wires to the inside of the cover (right).

One of the key components of the chamber is the so-called active web, shown in Figure 256. The active web is a multi-layer flex-rigid printed circuit board, which carries the high voltage to the wire and transmits the signal back to the front-end electronics. The web has two connectors; one for HV and one for the signals. The web connects to the backside of the cover, which is shown in Figure 257. The cover does not only provide the leak tightness of the manifold, it also comprises the front-end electronics, high voltage and gas connections. The front-end chip is the 8-channel CARIOCA chip which was developed for LHCb [75], [76] .

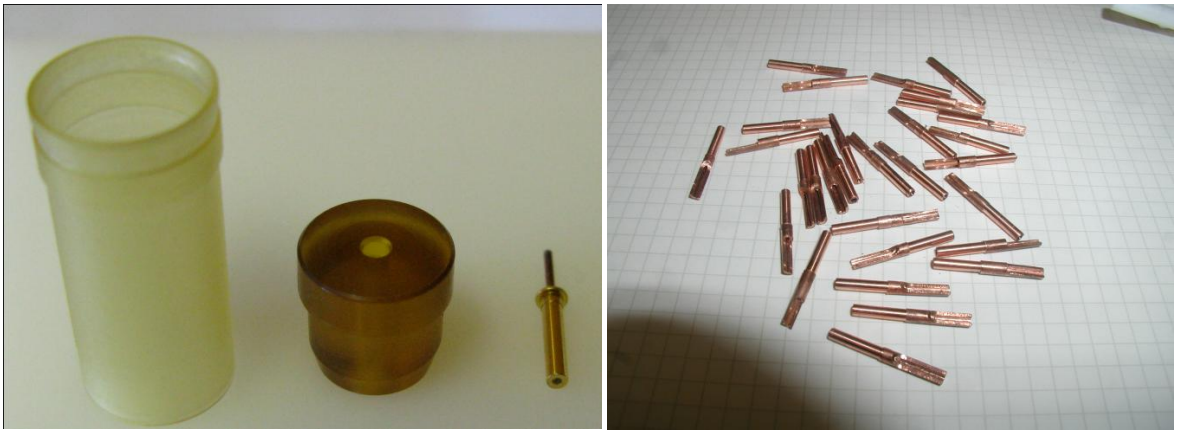

Figure 255. Details of the straw end-pieces are shown. The socket is glued between the straw and the structure (left). The connection plug pushes the petals of the web to the straw wall. The crimp-tubes connect the wire to the Kapton® layer of the web circuit by crimping (left picture) or by soldering (right picture).

The wire is tensioned and positioned using a copper tube inserted in the so-called connection plug (see Figure 255). The copper tube connects to the second layer of the web, which brings the high voltage to the wire and the signal back to the front-end electronics. At present two options for stringing the wires are studied; one based on crimping and one where the wire is soldered to a copper tube (see details in 3.5.11 on page 282).

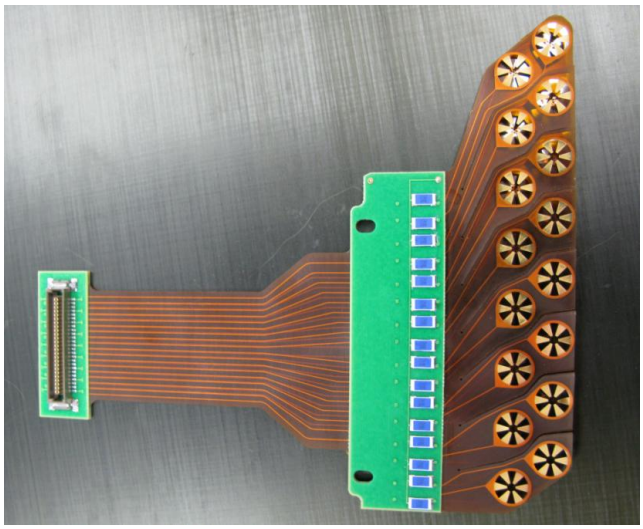

Figure 256. Detail of the flex-rigid circuits board (web) that connects to the straws (ground) and the wire (high voltage and signal).

Figure 257. Front-End board (cover) with the two CARIOCA chips (front side). The cover is mounted on partition serves as a leak-tight roof for the gas manifold and feed through for gas, high-voltage and signal connections to the straws (back side).

3.5.5.1.2 Alignment Tool for the Straw Insertion and Alignment

In order to keep the straws straight the straws are pre-tensioned with a force of 14.7 N. Mechanically, the straw under tension behaves like a string joining the two sides of the mechanical frame. By using spherical bearings, the straw is brought under tension while it is in vertical position (see Figure 258). The spring keeps the pre-tension on the bottom center pin during the gluing process. The pre-tension in combination with the bearings self-aligns the two pins, which hold the straw during the gluing, and

guarantees the straw straightness. Using a digital microscope fixed on a linear rail we measured the straw straightness under a tension. The measured straightness was better than 100µ. This measurement validated the assembly principle. The influence of gravity on a 10 N tensioned straw, tilted at 45˚, was measured and showed a deformation of ~300µ. This sag is too large for safe operation and imposed a spacer support development (see 3.5.5.1.3).

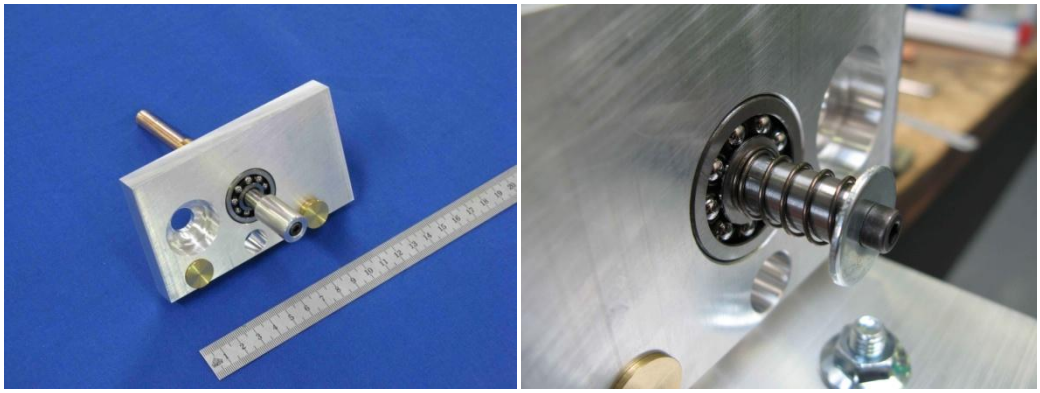

Figure 258. Alignment support with spherical bearings and (left) and the spring-loaded tension tool (right).

3.5.5.1.3 The Spacer

In order to support a straw along its length each straw is held by two "spacers" (see Figure 260). In this configuration the maximum unsupported straw length is ≈ 70 cm. In order to minimize the amount of material, the spacers consist of polyetherimide (PEI) rings surrounding the straw (Φ_{inner} = 9.92 mm, Φ_{outer} = 10.40mm, H= 3mm). The weight of one ring is 0.018g. They are glued at a distance of 17.6 mm from each other between two CuBe wires with a diameter of 0.1mm. The wires are given a tension of 500g to guarantee the straightness at all times. A shorter version of the spacer for the 64-straw prototype (see Figure 269) is shown in Figure 259.

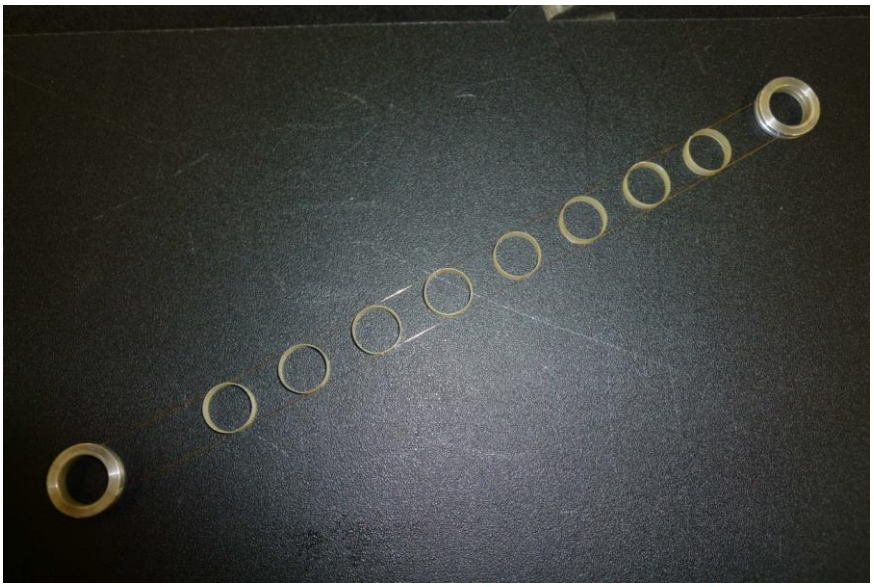

Figure 259: The spacer made for eight straws with a metal ring for alignment and tensioning at both ends. The PEI rings are seen in the center.

The design and assembly procedure was validated in a full-length prototype and measurements showed a maximum difference between the measured and theoretical ring position (ring centre) of 10 μm along the 2 m long spacer.

3.5.5.1.4 Straw Tension Measurement

Straw tension is measured with a long distance infrared reflective switch (OPB732). It measures the frequency of the straw vibration. The output of the device is connected to the microphone input of a laptop equipped with a LabVIEW® program. It displays the natural frequency of the straw and calculates the straw tension. The principle is similar to the wire tension measurement except a smooth rod is used to excite the straw instead of a current pulse in a magnetic field as in case of the wires. A schematic view of the set-up and the LabVIEW® panel is shown in Figure 260. Note the straw spacers limit the movement of the straws and create two nodes along the straws.

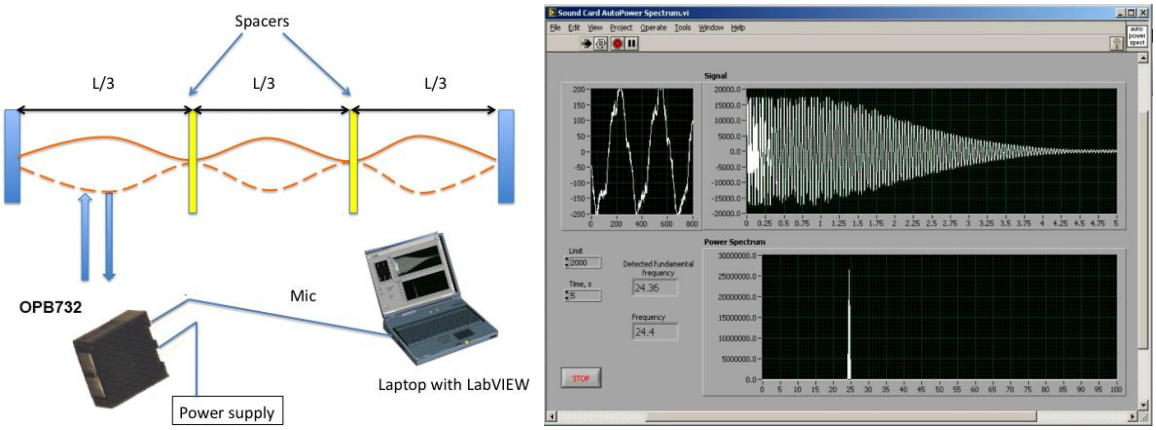

Figure 260. Set-up for straw tension measurement (left) and the LabVIEW control panel with the input signal from the OPB732 and the Fourier spectra (right).

3.5.5.1.5 The Measurement of the Straw Straightness

Once the straws are fixed at both ends in the detector frame, the straightness is measured layer by layer with two light sources and a CCD camera/microscope. The two laser diodes illuminate the straw and the camera/microscope takes a photograph of the straw section. The system is calibrated using a wire stretched from top to bottom parallel to the straws. A stepper motor moves the camera and a picture is taken at 10 equidistant positions along the straw. A dedicated LabVIEW® program transforms the pictures into a black and white image and finds the two straw edges automatically. The program then calculates the positions of the straw centre (straightness) in the xy-plane and the diameter of the straw.

3.5.5.1.6 Wiring

A gold-plated Tungsten wire from Toshiba with a diameter of 30μm is chosen [77]. The limit of elastic deformation has been measured to 150 g and the rupture occurs around 220 g. The nominal wire tension is set to 90 g (see 3.5.5.1.7). The detector is placed in horizontal position and the straw resistivity is measured. All straw diameters are measured and recorded in the logbook. The webs are

then installed and the ground circuit is fixed to the inside of the straw with the help of the connection plug. A tight fit is necessary to ensure a good electrical contact. The ground petals are formed with a special tool to facilitate the insertion of the PEI plugs. A 0.1mm diameter wire ("needle" wire) is blown through the straw. The 30 µm wire is fastened to the needle wire and kept under a small tension at all times during installation to avoid kinks. The wire is gently pulled through the straw. From both sides, the 30µm wire is inserted in the copper tubes and the copper tubes are then inserted through the electronic circuit into the contact plugs. Once the wire is fixed to the web side, a 90g weight is suspended to the wire on opposite side and fixed. The electrical continuity between the two pins is checked as well as insulation between wire and straw (broken wire). After wiring one layer the wire tension is measured and a HV test is performed to measure the leak current, which should not exceed 1 or 2nA per high voltage group at 1600V.

3.5.5.1.7 Measurement of Wire Tension

The full Straw Tracker contains a total of 7168 anode wires and the wire tension need to be verified during and after installation. The wire resonance frequency is close to 50 Hz for a tension 80 g and a wire length of 2200 mm. 50 Hz is close to the 220 V voltage supply and in order to avoid this frequency, the nominal wire tension was set to 90 g. The acceptable values of wire tension during the module production are between 85 and 95 g. The upper limit of 95 g is considered to be comfortably below the elastic limit of 150 g, while the lower limit is high enough to allow good operation of the straw. Electrical instability of the wire was found for tensions below 35 g. However, to minimize electrostatic deflections on the wire the lower limit is set to 85 g. The wire tension measurement principle is shown in Figure 261: wire oscillations are stimulated in an external magnetic field with the help of a current generator; the resonant frequency of the wire is measured [78].

The current induced on the wires will be amplified with a high-input impedance amplifiers and an input signal threshold (typically of about 1.5mV). Signals will be digitized by a PC based A/D board and the measured frequency will be translated into applied wire tension *T* (in grams) following the formula:

$$T = \frac{4 \times \mu \times L^2 \times \nu^2}{a}$$

where *L* is the free wire length in cm, ν is frequency in Hz, μ is the mass per unit length in g/cm and *a*, the gravitational acceleration. Uncertainties in the measured tension arise from variations in the wire diameter and the length of the vibrating wire. The minimum uncertainty can be estimated to ~1%. The readout electronics is connected to a cell of 16 straws and measures the main resonance vibration frequency of each wire. The magnetic field near the tested wire must be at least 100 Gauss. The operator can modify the parameters of the pulse generator and Fourier analysis on the Labview based computer panel. The tension of the wires will be measured during the stringing process itself and a final global wire-tension measurement of all wires in the module will be carried out before the next assembly step. The results of the measurement will be written to the production database. Loss of wire tension between module assembly at JINR and arrival to CERN can happen and hence the wire tension will be measured again after delivery to CERN. Major attention will be paid to wires, which show a tension loss of more than 10 g. These wires will be investigated carefully and replaced if necessary .

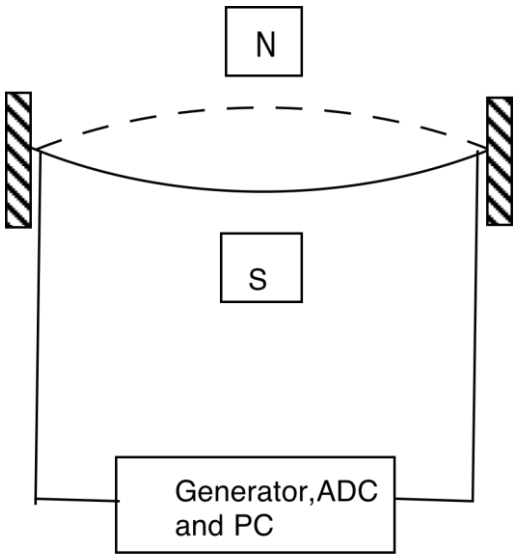

Figure 261. Schematic view of the wire tension measurement device.

3.5.5.1.8 Gas Tightness Tests

The next stage will be to verify the glued joints between straws and the frame. It is important to perform this test at this stage because the straws and their glued joints to the straw support frame are still easily accessible and repairable in case of problem. A dedicated gas-tightness set-up is to be used to verify the quality of the about 2000 glue joints (per module) between the straws and their end plugs and between the end plugs and the straw support frame. The set-up used for the leak test consists of a flat steel table and a cover (see Figure 262). The system contains a temperature sensor, pressure gauges, vacuum stand and a system of valves and pipes. The inner module flange with a groove for the O-ring will be put on the flat table. The cover must be centred with respect to groove with the O-ring on the outer module flange. The module gas tightness will be achieved by the vacuum inside the module and with the help of 12 bolts that clamp the table and cover.

Straws and gas manifolds will be filled with Ar/CO_2 from a bottle with 20 mbar over atmosphere pressure. The vacuum volume of the module to be tested is about 1 m^3. The leak rate (mbar•l/sec) will be evaluated by measuring the pressure increase of the vacuum around the straws as a function of time, once the pump has been stopped. A module will be approved for subsequent assembly whenever the measured leak rate does not exceed 10^{-2} mbar•l/sec. Nevertheless, all effort will be made to get a value as low as possible [78].

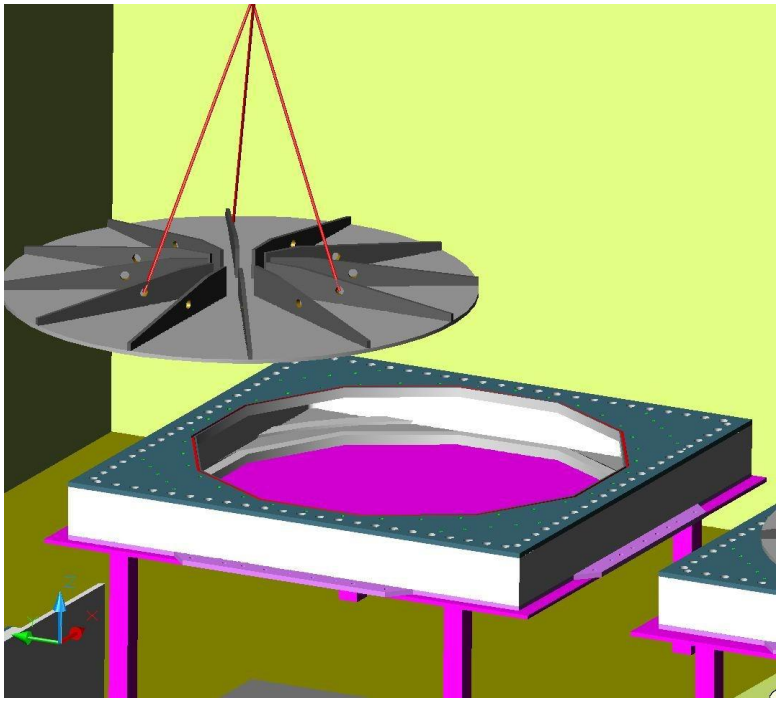

Figure 262. Flat steel table and a cover for module vacuum test.

3.5.5.1.9 Gas tightness of the Module

The module includes about 500 straws and gas manifolds (see 3.5.5.1.1) with total gas volume about 110 l. The gas tightness test of assembled module will be carried out by measuring the pressure drop in the closed detector gas system filled with Ar/CO_2 at the over pressure about 1 bar. The minimum test duration will be not less than 20 hours. Internal pressure, atmospheric pressure and temperature will be recorded during this test. The recorded data (temperature and atmospheric pressure) will be used to correct the module gas system pressure drop. The leak rate (mbar/l/s) is calculated from the measured pressure drop over time. A module will be considered valid for installation whenever the measured leak rate of the module gas system does not exceed 100 mbar per 24 h at P=2 bar (about 4 cm^3/min). The expected gas leak of 500 straws will be about 10 times lower.

3.5.6 Frontend Electronics

The NA62 straw detector is a gaseous detector consisting of 7‘168 drift tubes, 2.1 meter long, organized in 4 chambers, each station having 4 views. It aims at providing tracking information with a good resolution and a charged particle veto signal. The drift time measurement performed on the front-end electronics is indispensable for good tracking resolution. After analog processing of the straw signal, only timing information (start and end of the signal) is measured. From an electronics point of view, the following requirement can be defined:

- 130µm resolution for a single straw, 80µm for the detector
- Efficiency close to 100% (95% in high rate environment)
- Low noise
- Mean particle rate 40kHz, for straws close to beam up to 0.5MHz

3.5.6.1 Signal Parameters, Processing and Straw Properties

The basic detecting element is a straw drift tube of ~10mm diameter, 2.1 meter long. The cathode is formed by a very thin layer of copper and gold, few hundred atomic layers thick, so the electrical properties are determined by surface effects (one does not need to consider the skin effect). Indeed, the measured DC resistivity of the cathode is ~70 Ohms (2.1 m) entirely determined by surface effects. The anode is 30µm diameter gold plated tungsten wire. The straw can be considered as a very lossy transmission line and termination effects on both ends should be evaluated. With the gas mixture Ar/CO_2(70/30), the electron total drift time is ~150ns.

As a response to each cluster of primary and secondary electrons created by the passage of particle, straw outputs a current signal with approximately hyperbolical shape and few microamperes amplitude. The shape differs from hyperbolical curve for the first few ns due to variable ion mobility. Also, when the selected gas would contain electronegative components, like CF_4[29], the signal is distorted and of smaller amplitude due to electron attachment. The leading edge of the signal is very fast, of the order of 1ns, the trailing edge time depends on a number of factors like geometry, gas and voltage applied. The ion tail lasts up to few 100µs of microseconds. The electrical properties of the straws are the following:

- Characteristic impedance 350 Ohms (calculated) at 20MHz, 1000 Ohms at 1MHz
- Capacitance 23 pF
- Attenuation 2.3 at 20 MHz

The on-detector electronics consists of an 8-channel analogue front-end chip containing a fast preamplifier, semi-gaussian shaper, a tail cancellation circuitry, base line restorer and a discriminator. Input impedance of the preamplifier is a compromise between good straw impedance matching and a small crosstalk. The lower value is advantageous as it guarantees lower crosstalk and signal enhancement due to current increase during signal reflection on low termination impedance. The reflected signal, which returns after travelling to the far end of the straw and being reflected there, is already strongly attenuated and does not contribute significantly to the output. The shaping time should be short in order to get a response from the first primary cluster and thus better time resolution, but should be also sufficiently long to integrate further clusters so that there is only a single output pulse per particle crossing straw. Assuming 0.3mm average spacing between primary clusters, the shaping time should be much longer then 15ns for slow gas and 6ns for fast gas.

The ion tail cancellation is of upmost importance for straws with high rate of particles. Pile-up at high particle rate would cause loss of both efficiency and time resolution.

As baseline for the frontend analog electronics it is proposed to use the CARIOCA chip developed for the LHCb muon chambers. However, given the relatively small number of channels, electronics built from discrete components could be used if none of these chips could provide optimal performance. The CARIOCA chip has a shaping time constant of ≈15ns, so its operation with NA62 straw detector must be verified. This will be done in the 2010 test beam with the 64-straw prototype.

[29] CF_4 is considered in the back-up gas mixture, i.e. $CO_2/iC_4H_{10}/CF_4$ (80%/10%/10%)

3.5.6.2 Modularity and Readout

The frontend modularity is following straw smallest unit. There are 16 straws in the basic unit of gas, high voltage and readout. The outputs from discriminators are either directly fed to the on-board TDC or transmitted to external module with TDC. Study and on-going design confirms the possibility to integrate 16 TDCs with 1ns RMS resolution inside Field Programmable Gate Array (FPGA) chip together with time-stamping, serialization and readout protocol using high speed serial links. Both starts and ends of signal must be recorded. The falling edge (end) of signal is the same for all straws seeing the same particle (Figure 263) and can be used for time measurement validation or even as a crude time measurement. It is also ideal for building fast hardware trigger or veto for multiple charged tracks. The rising edge (start) of the signal is used for precise drift time measurement and the track position is calculated through known r-t dependence.

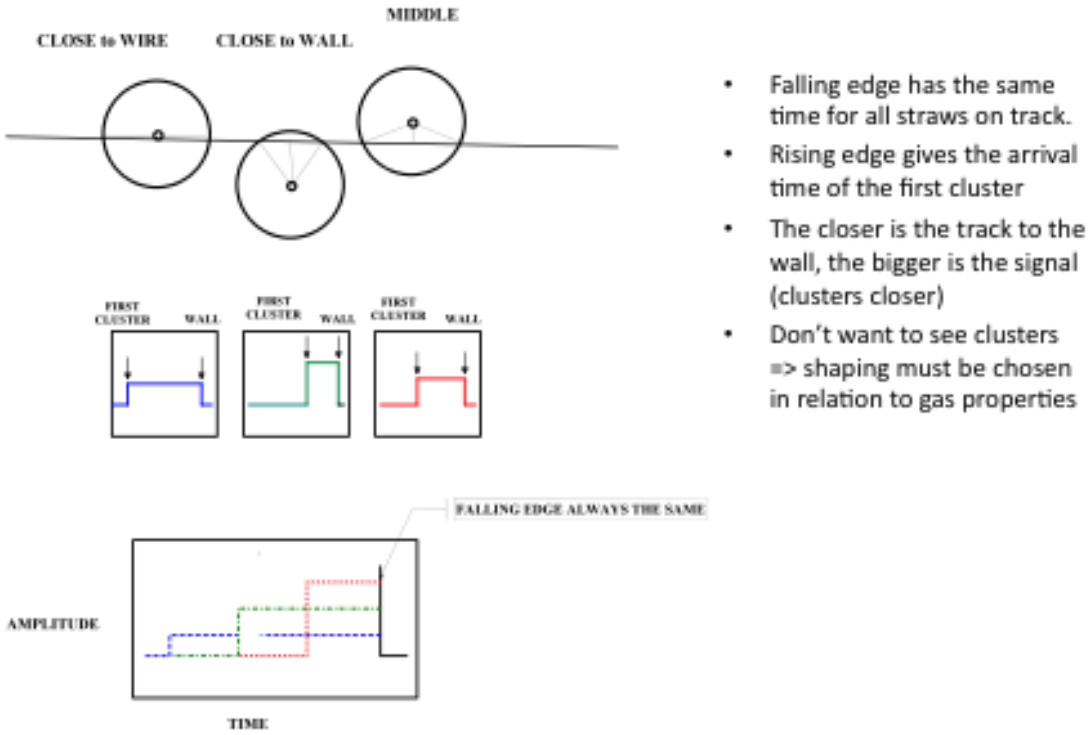

Figure 263: Timing of the straw detector.

3.5.6.3 Noise and Internal Crosstalk, Straw Termination

The threshold setting is a compromise between time resolution and the rate of noise hits. The lower the threshold, the better should be a time resolution and efficiency but the higher is the rate of noise hits. As the straw detector provides also multiple charged particles veto signal, the noise must be as small as possible. Crosstalk can also cause fake hits; multiple measures have been taken to reduce it as much as possible. For conductive crosstalk, the return paths of straw signals are separated already at the level of web connection and kept until frontend board input connector. Capacitive coupling effect can be mitigated by lowering input impedance of readout electronics and eventually by individual shielding of straw cathodes. Input impedance is the sum of preamplifier input impedance and input discharge protection resistor. The input impedance of the CARIOCA is about 50 Ohms at the working

frequency; the value of the protection resistor is of the order of 100 Ohms. The total load impedance for straw is thus ~150 Ohms, to be compared with 350 Ohms straw characteristic impedance. The lower load impedance will cause signal reflections, but in a useful manner; it will increase the input current, thus enhancing input signal. As the straw is a transmission line with very high loss, the signal which returns after reflecting from the far end of straw, is already negligible. The other possibility considered for lowering capacitive crosstalk is covering straws from outside by very thin metallic layer connected separately to ground of the frontend board.

The straw far end is left open. This should enable for collecting as much charge as possible as the straw has a very high attenuation (2.3) at working frequency.

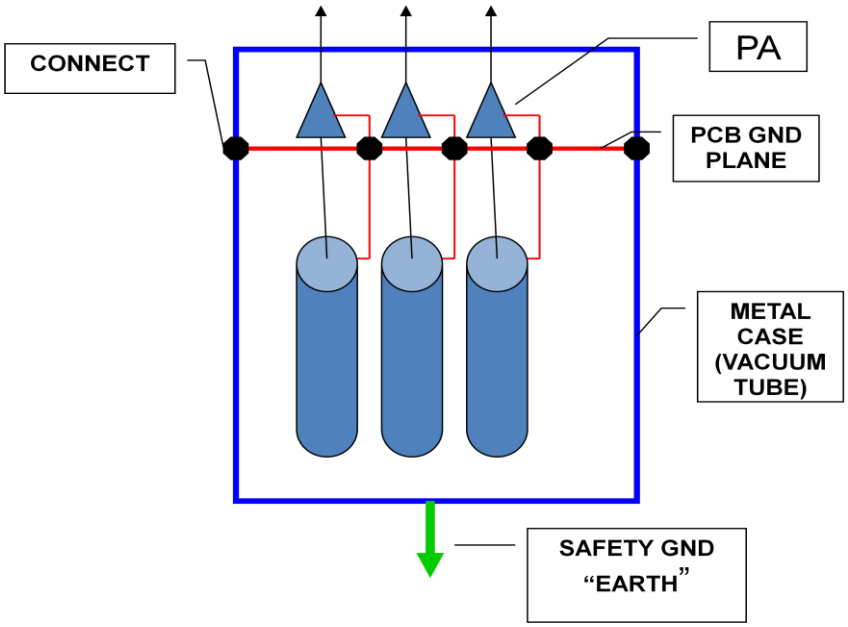

Figure 264: Grounding and shielding scheme

3.5.6.4 External Crosstalk, Grounding and Shielding

External crosstalk is an interference with external objects surrounding the detector, like cranes, magnets, power supplies, wireless connections, etc. Another source can be caused by voltage potential difference between front and back end. One must also ensure connection of all conductive parts to safety ground. The proposed grounding and shielding scheme in Figure 264 ensures that the straw signal returns are separated and connected close to preamplifier inputs at frontend board level. This scheme was implemented in the 64-straw prototype.

The ground of the frontend board is connected around its perimeter to metal case, which is surrounding the whole detector, thus forming tight electromagnetic shield. Straws are connected to signal returns only on one end to prevent external currents flowing through the cathodes. All cable shields should be connected on both ends; one side directly to frontend board, backend side through damping impedance. If there are metallic pipes servicing gas or cooling, they should be electrically connected to metal frame as well.

The metal case of the station should be connected to safety earth at one point. The backend electronics (VME crates) must be connected to the same earth, thus providing both required safety connection and eliminates potential difference between frontend and backend electronics.

3.5.7 The Gas System

The gas system consists of two parts; first, the mixer system which delivers quantity, mixing ratio and pressure conditioning to downstream elements and secondly, the distribution system, which delivers the gas in well defined quantities to the individual detector components.

3.5.7.1 Gas System Requirements

The total gas flow chosen correspond to a normal gas flow rate of ~2-4 cm^3/min per straw. The gas modularity is optimised in order to minimise the number of lines between the detector and the distribution racks, and, on the other hand, to minimise the loss of performance in the case of a accidental leak in any module.

3.5.7.2 The Mixer

The detector shall be supplied with a constant gas mixture of Ar/CO_2 (70%/30%) with a precision better than 1%. The total flow of the mixer will be 1500 l/h. Each primary gas line is equipped with a Digital Mass Flow controller to measure the component flow with appropriate accuracy. An output pressure regulator adjusts the downstream pressure from 0.2 to 2 bar. The control system provides a flow independent mixing ratio and adequate error handling. The mixer has to automatically follow the demand of the distribution system. The mixer system is in fact a standard mixing station of the LHC gas system project: a none-ATEX, two-gas mixer for a purged gas system with output pressure regulation [79].

3.5.7.3 The Distribution

The gas from the mixer is distributed to four gas distribution racks, one for each chamber. The major design criteria of the distribution system is the uniform gas supply to each cell with adequate separation capabilities in case of pressure loss due to a leaking straw.

Figure 265 shows an overview of the entire gas system and Figure 266 the gas channel allocation of one chamber. Each chamber has four views, each containing 448 straws or 30 gas channels (Each view consists of 2x15 gas channels[30] and each channels supplies 16 straws). The gas distribution of one layer is split into 13 “normally flushed cells” and two “highly flushed cells” for the centre beam region. The 13 low flow cells per layer are grouped into one gas channel. The remaining two high-flow cells in the centre of the layer are grouped with similar cells of all layers to one gas channel (2 cells x 4 layers). In this way, one chamber has eight gas channels with low flow rate and two gas channels with high flow rate. The gas volume of one cell is about 3 litres, which results in about 20 l/h (1/2 Volume exchange per hour) and 24 l/h (1 volume exchange per hour) for low and high flow respectively. The following main elements in the gas system layout have been identified (see Figure 265):

[30] Each view consists of 2x15 gas compartments, however due to the beam window the number of straws near the beam can be smaller.

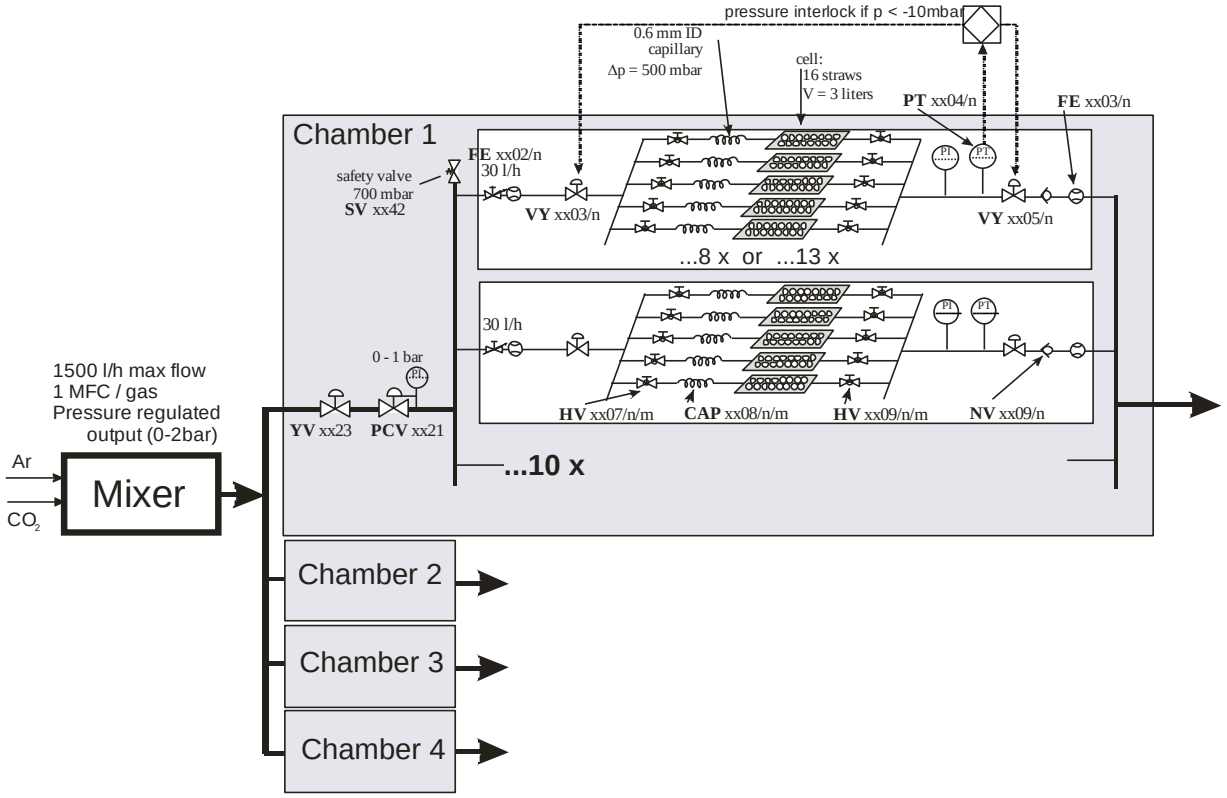

Figure 265. NA62 Straw Gas System overview.

The pressure regulator PCVxx21 (xx=61,62,63,64 for chamber 1,2,3,4) provides a constant relative pressure of 500 mbar to a gas manifold of 10 gas channels. The safety valve SVxx42 protects this manifold to about 700 mbar of over pressure. Each of the 10 gas channels is equipped with the following elements:

1. Input flow measurement FExx02/n (n=1,..,10)
2. Separation valves VYxx03/n and YVxx05/n
3. A gas distributor to 8 or 13 cells
4. Cell separation input valve HVxx07/n/m (m=1,...,8 or 13)
5. A capillary, CAPxx08/n/m, of 0.6mm inner diameter and a pressure drop of 500 mbar. The length of this capillary determines high or low flow rate for the cells.
6. Cell separation output valve HVxx09/n/m
7. Pressure transmitter PTxx04/n with interlock to VYxx03/n and VYxx05/n (fires in case less than 10 mbar under pressure is detected).
8. A none return valve, NVxx09/n, avoids backflow and therefore air contamination in case of pressure loss in a cell
9. Output flow measurement FExx03/n

The full functionality of the distribution system is covered by the standard LHC gas system control project.

As the straws will work in vacuum a protection system is required against a sudden leak in a straw. As shown above the modularity of the system (HV and gas) is 16 straws. PTxx04/n in combination with inlet and outlet shut-off valves VYxx03/n and YVxx05/n, assures a fast closure of the gas supply to one gas channel in case an under pressure is detected (e.g. broken straw). The contamination of the vacuum surrounding the straw tubes is therefore kept at a minimum. The none-return valve at the outlet of each gas channel will cause a slightly elevated straw pressure of 10 to 15 mbar above atmosphere.

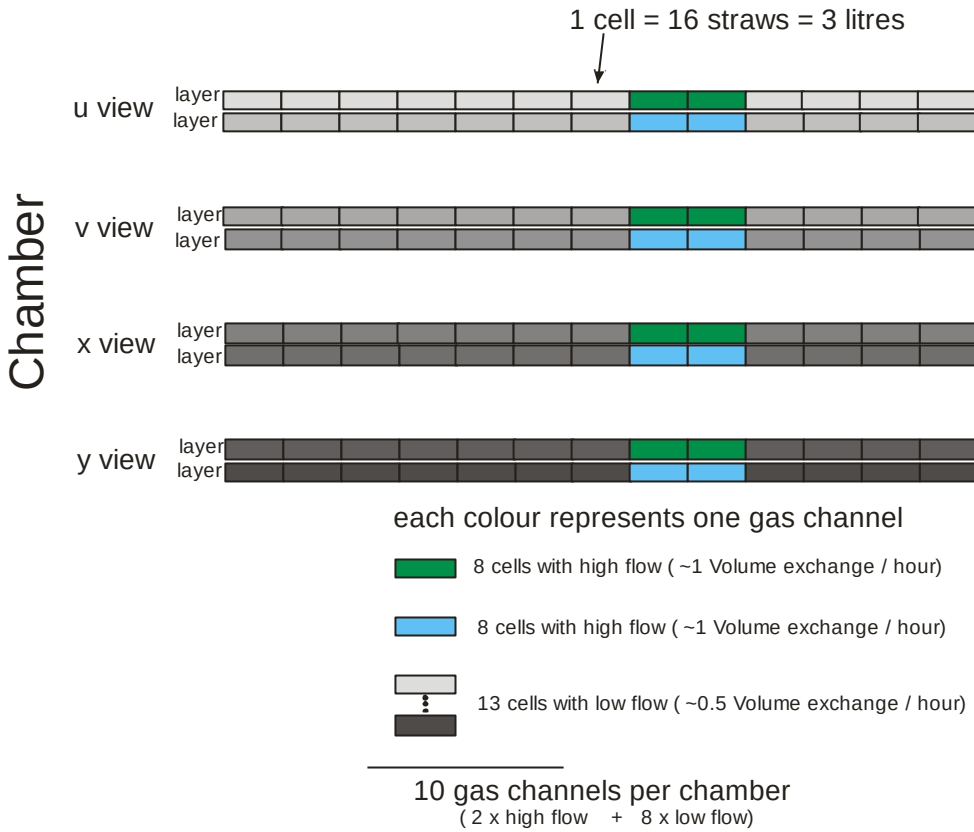

Figure 266. NA62 Straw Gas channel allocation.

3.5.8 Straw Tracker DCS

3.5.8.1 DCS Architecture

The detector control system (DCS) should provide control and monitoring of the detector hardware. In addition, it should perform archiving of the hardware parameters to the database in order to provide access to this information during the off-line data analysis.

It has been decided to use the common DCS system, now under development at CERN, which is based on the PVSS II SCADA toolkit and the JCOP framework.

The Tracker controlled/monitored equipment includes the following items:

- Low voltage power supplies.

- High voltage power supplies.
- Gas mixing and distributing systems.
- Low voltage and temperature monitoring system.

3.5.8.2 Low Voltage (LV) System

The LV system should provide power to the front-end (FE) boards placed on the four straw tracker chambers. Each chamber contains four views and each view contains 30 FE boards. In the present design, one FE board consumes ~ 1 A at 5 volts and therefore one module requires the LV power supply to provide a current of about 30A at 5 V. These requirements could be fulfilled by eight Wiener power supplies MPV8008, which have eight channels of 10 A output current each. One view requires four channels and therefore, two power supplies will provide the required power for one chamber. The MPV8008 could be controlled in a similar way to the solution used in the LHC experiments. It includes a CAN-bus interface card and the OPC-server providing the interface to PVSS. A second possibility to control the MPV8008 is to use the TCP/IP protocol together with the corresponding TCP/IP PVSS driver which is part of the PVSS driver package. The final choice of the FE electronics and corresponding LV power supplies will be done once the results from the 2010 test beam with the 64-straw prototype is known.

3.5.8.3 High Voltage System

The Straw tracker requires a high-voltage (HV) source with voltage below 2 kV and low current. The present design assumes two HV channels per view which gives 32 channels in total. The CAEN HV power supply board A1535, containing 24 channels with 3.5KV/3mA output, should be sufficient. To have 32 HV channels, two boards of this type housed in CAEN SY2527LC mainframe are needed. The mainframe should be controlled via Ethernet line with TCP/IP protocol and corresponding CAEN OPC server. The JCOP framework is suitable for the board A1535 and this component exists as it was developed for the CMS ME1/1 muon chambers.

3.5.8.4 Gas System Controls

The gas system for the straw tracker should provide and distribute Ar/CO_2 gas mixture to individual straw cells (16 straws). The plan is to use the CERN standard gas mixing/distribution system based on PVSS with the possibility to control and monitor the gas mixing and gas flow values from the DCS PC. For this purpose PVSS distribution manager will be used. Both software and hardware interlock signals to prevent HV turning on in case of missing gas flow or incorrect gas mixture will be implemented.

Due to the fact that the straws are placed in vacuum, the gas system should rapidly close the gas line to a cell in case of a leak in a straw. Pressure sensors will detect a sudden drop in pressure on the supply lines. In this case, the gas system should also send a signal to the Straw Tracker DCS to switch off the corresponding HV channels or the whole HV system.

3.5.8.5 Thermometry and FE monitoring

In order to get information about the cover (front-end) temperature at the straw group inlet, one thermo-sensor per FE cover will be mounted. Two voltages one the FE boards and the current consumed by a board will be measured using Embedded Local Monitoring Boards (ELMBs) [80].

A chamber view has 30 boards and four thermo-sensors will be mounted diractly on the mechanics structure. In total, one view needs to measure 34 temperature values and 90 voltages. Therefore, to

read out all monitoring information from one chamber 10 ELMBs are needed. The readout should be performed via CAN-bus and OPC server, which is in a way similar to the monitoring of the LV power supply. It seems reasonably to have one CAN-bus branch per chamber, which gives in total four CAN-bus branches.

In case the LV power supplies control protocol is compatible with the ELMB control protocol, the LV power supplies will be connected to the chamber CAN-bus branches. The final decision depends on the final choice of the FE electronics.

3.5.8.6 Logical Trees in DCS and FSM

To get a convenient way of navigation through the detector elements, the DCS logical tree should have a structure including both tracker nodes and gas system nodes. Due to the fact that the detector has only a few LV and HV channels, the bottom node of the tree should be linked to the corresponding channel of the LV or HV power supply in the hardware three. The FSM tree should correspond to the Straw Tracker logical tree. The states of FSM should correspond to the states of the whole detector.

3.5.8.7 DCS Development and Maintenance

The plan is to use our own DCS PC running tracker PVSS system and all detectors DCS OPC servers in one Windows-based computer. This PC should have CAN-bus adapter cards installed. After the development is finished and to simplify the maintenance of the detector DCS, the plan is to create a JCOP FW component containing all Tracker PVSS panels, scripts and port it to the NA62 Central DCS computer. A copy of this component should be stored in the NA62 DCS repository. The Tracker DCS PC should be used to run the DCS servers and to house the CAN-bus adapters. To simplify the control of the detector during further hardware development, debugging and maintenance, the PC to access PVSS panes located at the Central DCS PC, will be used. For this purpose the DCS PC should have a PVSS User Interface software installed.

3.5.9 Alignment and Installation

3.5.9.1 Requirements

Alignment is the procedure in which the positions of the detector elements are determined. The residual uncertainties on these positions after the detector alignment has been performed are defined as the detector misalignments. Any misalignment of the elements of the detector will degrade the resolutions and the overall detector performance. The chambers will be equipped with eight reference points visible from the outside as shown in Figure 267. The reference points are positioned with respect to the straws with a precision of about 150 μm and will be measured before the straws are inserted. The real position of the reference points will be measured after the assembly of the module structure, but before installation in the experiment. The main principles of the alignment strategy are:

•The holes for the straws are drilled with a precision of ~0.05 mm

•The beams fixing the straws are positioned on the center plate (#1) of ~0.1mm

•The straws are assembled with a precision better then 0.1mm over their full length

•The positions of the straw holes are measured relative to the “targets" (red dots) with a precision better then 0.05 mm.

•Alignment with tracks: procedure to be defined

•Adjustable supports under each chamber for installation and alignment of the chambers.

3.5.10 Detector installation

In order to install the chambers, an interface part is connected on each side of the chamber. 16 bars interlink the two interface parts and will protect the chamber from axial forces as shown in Figure 267. Finite element calculations will determine the exact number and dimensions of theses bars. Before the chambers are installed, the two modules are joined together. This operation is carried out after the two modules have passed the quality control at CERN.

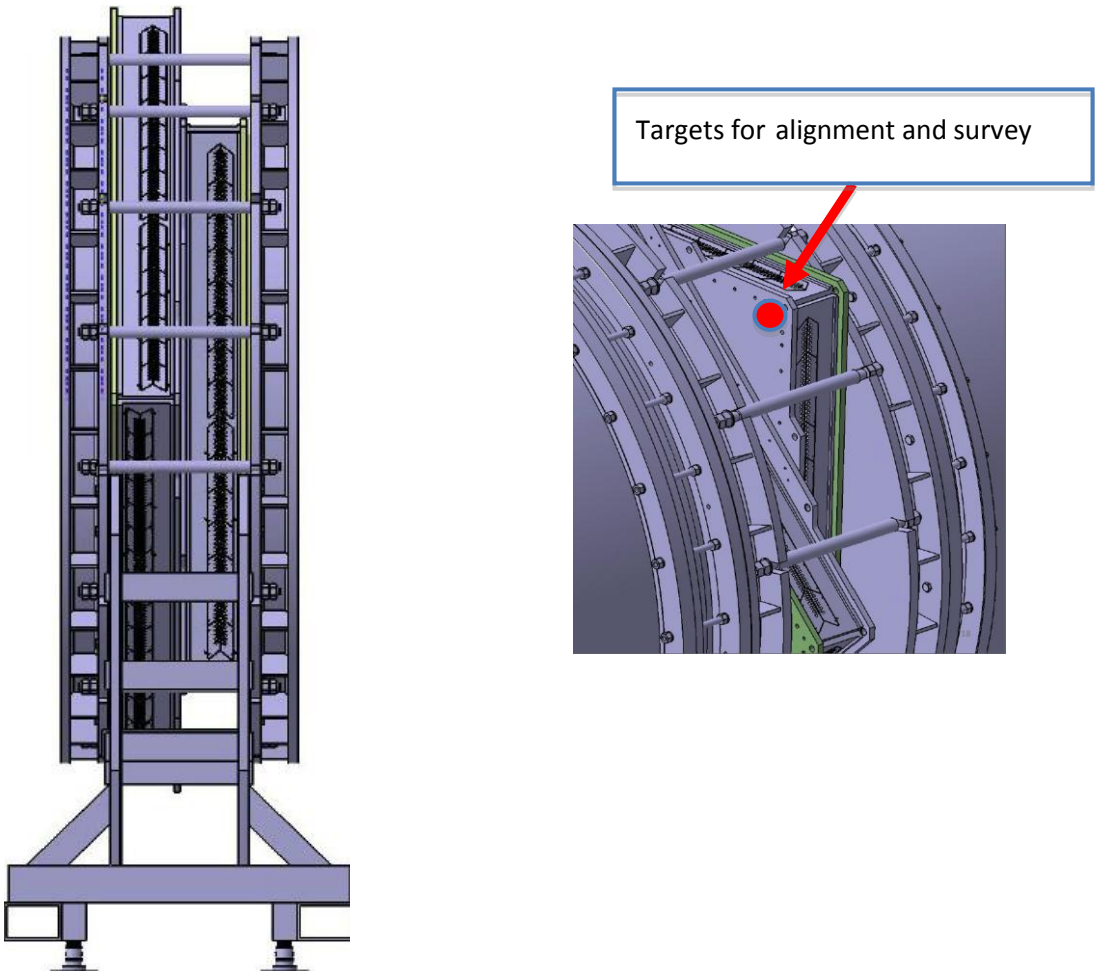

Figure 267. 16 bars parallel to the beam connects the two interface parts and protect the chamber from axial forces (left) and detail of the connection to the blue tube. Eight reference points per chamber are for survey and alignment.

3.5.11 The Assembly and Testing of a 64-straw Prototype

In order to validate the final design and assembly procedures, a full-length 64-straw prototype was built [81]. The straw tracker prototype support was assembled and measured on a precise flat table using a height gauge with a precision of +/-10 microns. All reference holes were measured. A complete mapping of the prototype was recorded and used for the prototype realignment during all assembly steps (see Figure 268).

Figure 268. View of the prototype structure (left) and the measurement of the frame on the granite table.

3.5.11.1 Installation and Alignment of the Straw Spacer

A full-scale low-density spacer for straws support was manufactured according to the principles explained in 3.5.5.1.3. Their position was adjusted on the prototype using the granite flat table taking into account their final position and including the deformation generated by gravity acting on the aluminium structure. The spacer in its final position with its dedicated adjustment tools is shown in Figure 269.

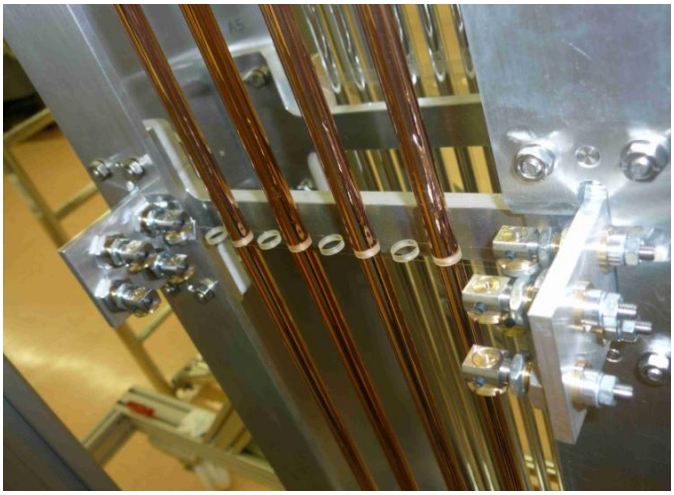

Figure 269. A view of four inserted straws supported by a spacer.

3.5.11.2 Assembly Tooling

A dedicated tool was designed for the prototype assembly. This support allows to work in two different positions: horizontal for the straw insertion and straightness adjustment shown in Figure 270 and vertical for the final gluing under tension.

Figure 270. The prototype in the assembly tooling in horizontal position (left) and the straw fixation tooling (right).

3.5.11.3 **Straw Insertion and Gluing**

The PEI end plugs are then introduced and the alignment tooling support installed. Every straw is fixed with a dedicated pin inserted through a spherical bearing at both ends. The straw fixation tooling blocks the straw and guarantees the 1.5 kg pretension to the straws as shown in Figure 270 (right). The straw is free to move in the straw fixation plate, which guarantees the straightness of the straw after the curing of the glue. On the upper side, the linear translation is blocked, while, on the bottom side, the pin is allowed to slide through the bearing. A tension of 15N is put on the straws using a spring-loaded tool. On the bottom side, the end-plug groove is protected with an O-ring joint and a soft spring is inserted to maintain the end plug in the correct position during the gluing process as shown Figure 271 (left). Tra-bond 2115 epoxy resin is used to glue the straws and as its viscosity is very low the, glue flows by capillarity. On the bottom side, the glue is applied to the outer diameter of the end plug before it is pushed gently vertically until it touches the aluminium plate. The soft spring maintain it in position. The glue is injected between the straw and the end-plug. On the top side, glue is applied to outer diameter of the end plug and it is gently slid vertically downwards till the aluminium end plate. The glue is then injected between the straw and the end plug as shown in see Figure 271 (right). After 24 hours of polymerisation, the alignment pins are removed and the alignment tool transferred to the next straw gluing position. The end plugs are then cut to a final length using a dedicated cutting device.

Figure 271. Bottom side gluing (left) and the gluing of the tope side (right).

3.5.11.4 **Straightness Measurement**

After the assembly of the first layer, every straw tube position is measured with a CCD camera fixed to a height gauge. Seven positions are recorded for each tube, measured under a 2 bar absolute pressure As final result, the straightness obtained on the fourth layer is shown in Figure 272. The precision in the spacer assembly is about ±35µm and the straightness is within ±100µm for each straw.

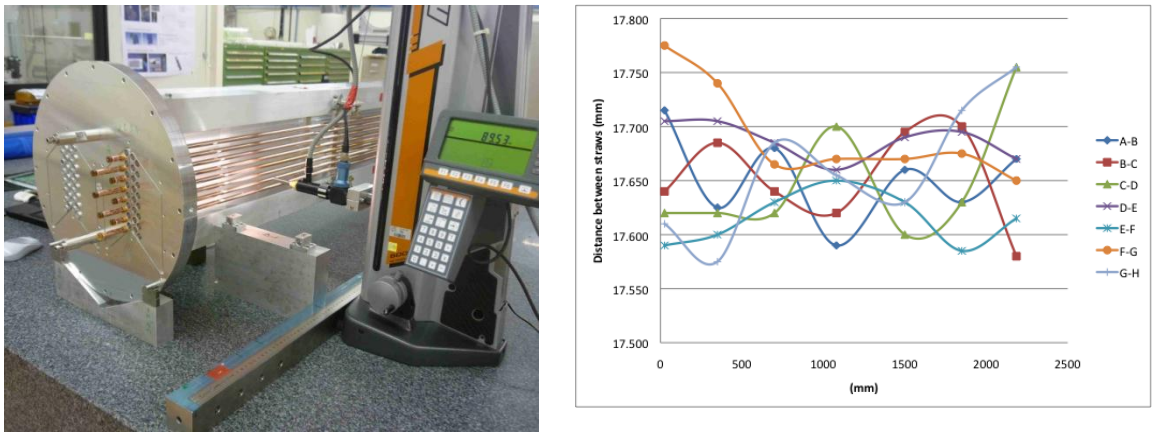

Figure 272. Measurement set up (left) and the distance between neighboring straws (right). The straws are kept at 2 bar absolute pressure [31].

3.5.11.5 **Straw Tension Measurement**

As explained in 3.5.5.1.4, the straw tension is measured with a long distance reflective switch (OPB732). It measures the frequency once a smooth rod gently excites the straw. This method allows for a fast and accurate measurement of the resulting pretension in the straws after straw installation. The results for the first 32 straws are shown in Figure 273.

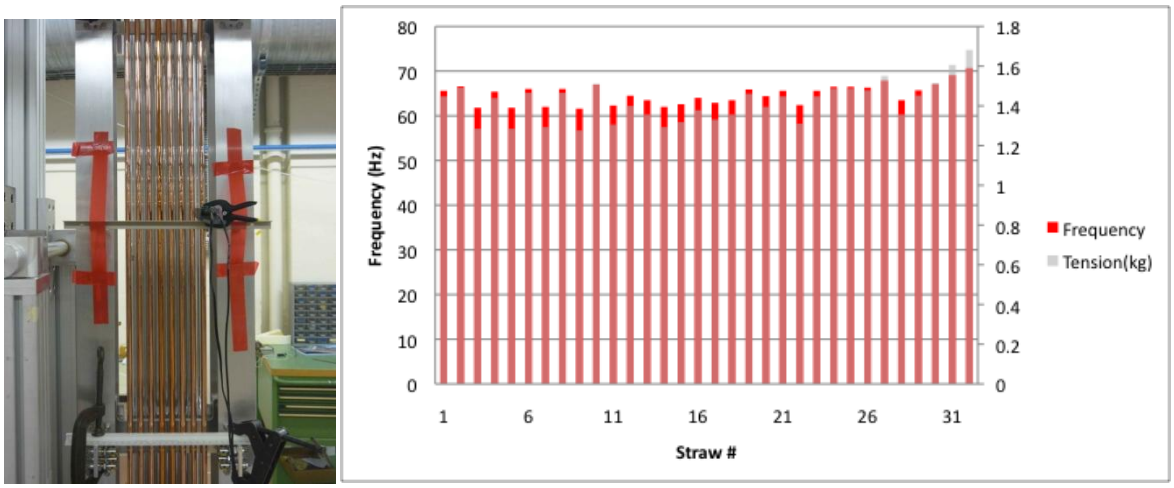

Figure 273.Set up in vertical position for straw tension measurements (left) and the measured tension (right).

[31] The distance between straws is 17.65 mm in the prototype. For the final chamber design it was changed to 17.6 mm.

3.5.11.6 Wiring Fixation option #1 (Crimping)

The detector is placed in horizontal position and the straw resistivity is measured. All straw diameters are measured and recorded in the logbook. The circuits (webs) are then installed and fixed. The ground petals(connects to the inside of the straw) are formed with a special tool to facilitate the insertion of PEI plugs. The wiring tool connected to a nitrogen bottle and a 0.1mm diameter needle wire is installed. The wire is blown through the straw as shown in Figure 274 (left). The 30µm anode wire is fixed to the needle wire. From this moment on, the 30µ wire must be kept under a small tension to avoid any kinks on the wire. The wire is gently pulled through the straw. From both sides, the anode wire is inserted in the copper crimp tubes. The crimp tubes are then inserted through the electronic circuit into the HV plugs. Once the wire is crimped on the web side with a special crimping tool, a 90g weight is attached to the wire on the opposite side before crimping the wire as shown in Figure 274 (right). The electrical continuity between the two pins is checked as well as the electrical insulation between wire and straw. Then, a complete high-voltage test is performed to measure the leak current, which should not exceed 1 or 2nA per high voltage group at 1600 V.

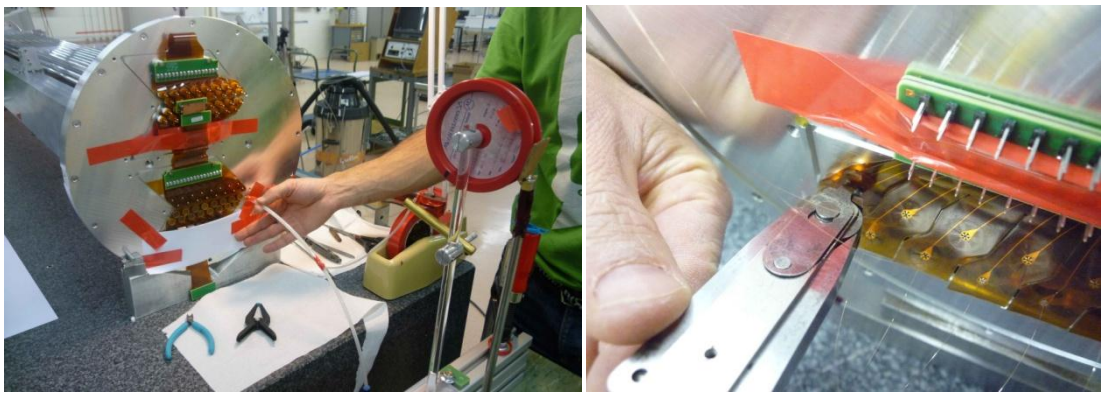

Figure 274. Blowing the wires (left) and crimping of the wire (right)

3.5.11.7 Wire Fixation option #2 (soldering).

A second method of positioning the wires was tested in the prototype. The principle of wire position is shown in Figure 275. An external tools was fixed on the flange reference holes by dowel pins. The accuracy of wire spacing with this method is +/- 5 µm and given by the precision in tooling. Each wire is blocked between 3 pins.

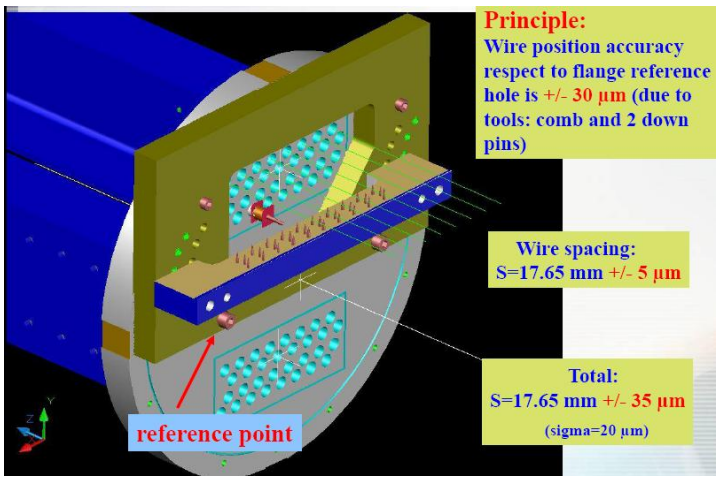

Figure 275. Wire positioning method

The method to verify the wire position at the straw ends during the assembly of the wires is shown in Figure 276.

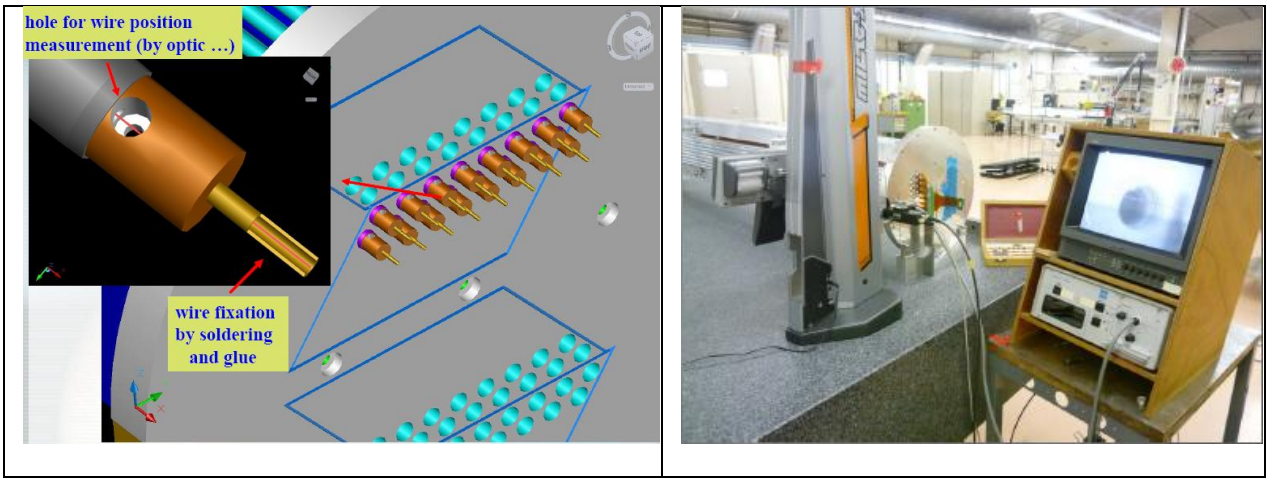

Figure 276. Verification scheme and measurement procedure

The measurements of the relative wire positions are presented in Figure 277. The values (distance between two neighbouring wires) showed that at least one of the two wires was badly positioned and it was replaced. The nominal value of the wire spacing is 17.65 mm. The wire positioning precision is +/- 35 µm, which includes a measurement accuracy of +/- 10 µm. σ of this distribution is about 10 µm and is negligible compared to the requirements (the straw spatial resolution should be better than σ=130 µm).

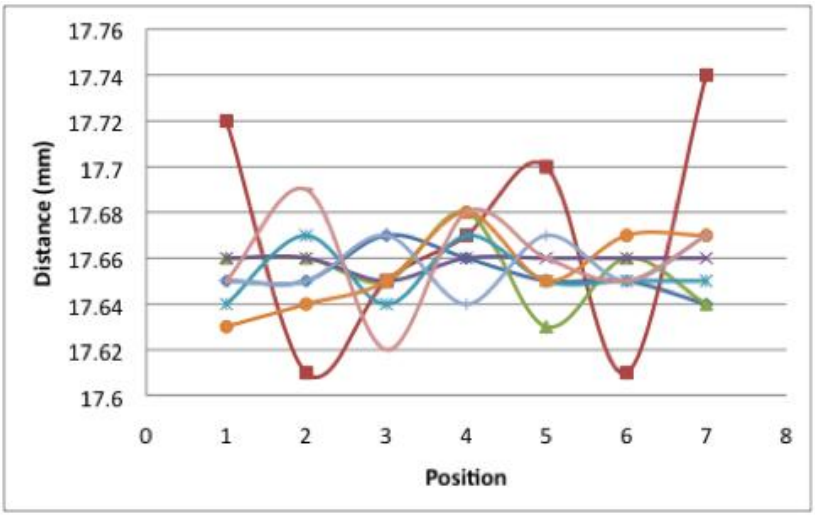

Figure 277. Measured wire positions for the soldered wires.

3.5.11.8 The Partitions

The dismountable partitions are mounted after the wiring and HV check. Thanks to the dry mounting, they can be removed if necessary, e.g. to replace a wire or access a straw (see Figure 278). After the assembly of the partitions, the first leak tightness test takes place.

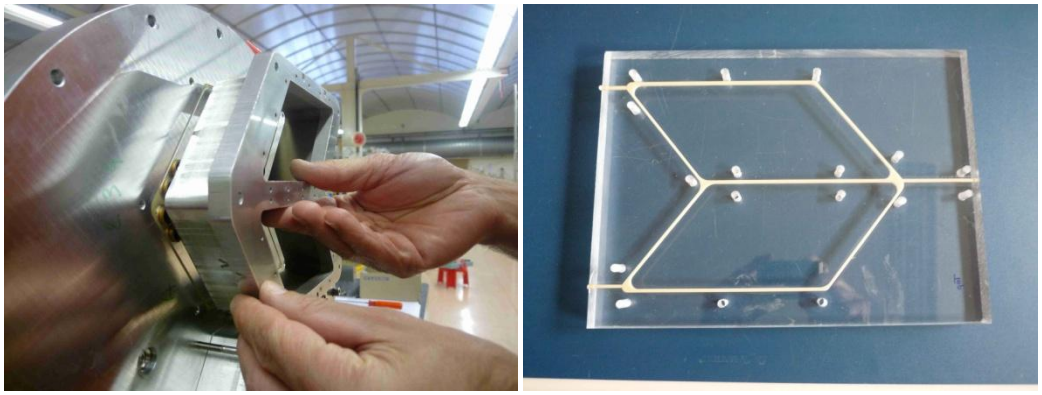

Figure 278. View of the cells partition during assembly (left) and polyurethane joint in the mould (right).

3.5.11.9 Measurement of the Leak Rate

In order to test the straws in real working conditions a dedicated vacuum vessel was built. The prototype is directly fixed on the first edge, and fastened through a metallic bellows to the second end (see Figure 279). With a turbo pump we obtained a vacuum of 2.5×10^{-4}mbar . Through a window on the side of the tube, the centre of the first layer could be observed. No deflection or movement of the straws was noticed during this operation. The permeation tests, using Argon (70%) and CO_2 (30%), showed a total leak rate of 9.3×10^{-2} mbar l/s while the specification is 10×10^{-2} mbar l/s. The requirements and measurements are summarized in Table 40.

Figure 279. View of the far end with the metallic bellows for the vacuum and leak test. The two gas cells (without electronics) are visible in the center.

Table 40. NA62 gas load

Straw detector at NPT (measurement)		LAV (estimate)	Blue tube (estimate)
Single straw	Full detector (7168 straws)		
0.82 mbar cm^3/min	9.8×10^{-2} mbar l/s	3×10^{-2} mbar l/s	2×10^{-2} mbar l/s

3.5.11.10 Aging studies

In order to validate the different components and material in the detector, a dedicated prototype was built. The prototype contains two straws; one with Aluminium coating and with Copper/Gold (see 3.5.3.1). The straw end-pieces were made from ULTEM and the connection to the straws was made with a section of the web (see 3.5.5.1.1). The wire is the 30 μm gold plated Tungsten from Toshiba foreseen for the chamber production. Final crimp tubes were used for electrical and mechanical connections of the wire to the web. The prototype was mounted in the CERN aging test facility. A gas mixture of 70% Argon and 30% was used. The parameters during the aging test are summarized in Table 41.

Table 41: Parameters for the aging test-up

Gas mixture	Gas flow	Current	Total charge	Irradiation area	High voltage	Fe^{55}-scan slit size
Ar(70%) CO_2 (30%)	0.5 cm^3/min	280 nA	0.27 C/cm (500h)	10 mm	1700 V	1 mm

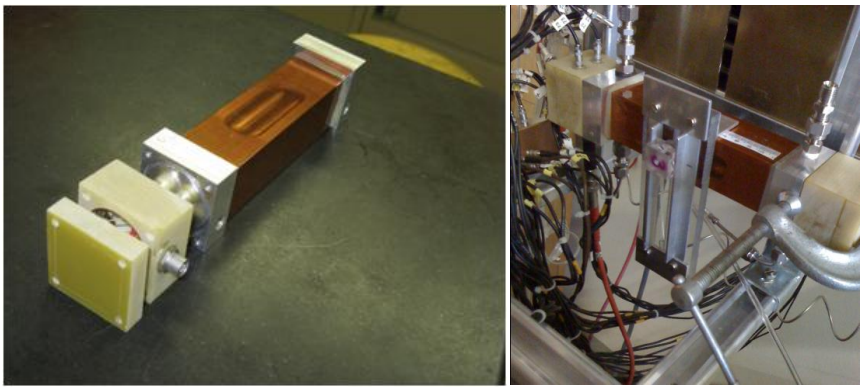

Figure 280. Straw prototype for material validation (left) and mounted in the set-up (right)

A scan with a Fe^{55} source is made at intermediate irradiation levels over the window from 10 to 75 mm [82]. The results are shown in Figure 281 and no change in amplitude can be seen up to a total charge of 0.27 C/cm. The estimated accumulated charge for the hottest straws in the experiment is 0.04 C/cm. After the initial run with Ar/CO_2 it was decided to dismount the prototype and investigate the wire and the inside of straw surface around the irradiated region.

Figure 281. Amplitude scan. for the Cu/Au straw and Aluminium straw respectively.

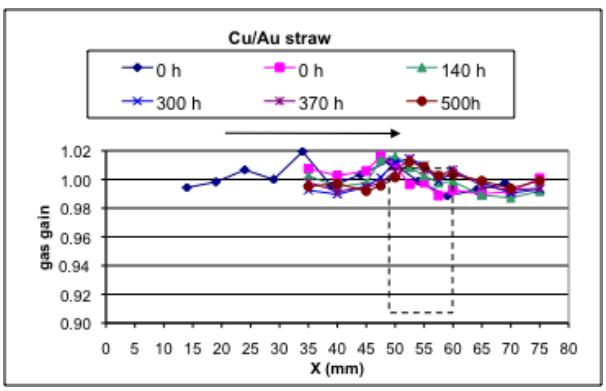

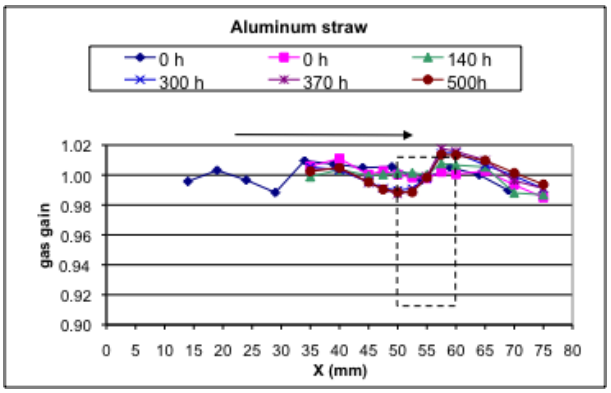

3.6 The RICH

A detector is needed to fulfill the following tasks:

- Separate π^+ from μ^+ between 15 and 35 GeV/c momentum providing a muon suppression factor of at least 10^{-2};
- Measure the pion crossing time with a resolution of about 100 ps or better;
- Produce the L0 trigger for a charged track.

A Ring Imaging Cherenkov (RICH) detector is the natural candidate for particle identification. The first requirement leads to the choice of Neon gas at about atmospheric pressure as the Cherenkov radiating medium; a reasonable compromise between the number of produced photoelectrons [11] (roughly linear with the length of the radiating medium) and the available space in the NA62 layout between the last straw chamber and the LKR calorimeter is achieved with a gas container not longer than 18 m in the beam direction. The second requirement leads to the choice of fast single anode photomultipliers, while the first one would point to PM as small as possible and packed as close as possible. A brief explanation of the choices is given in the following.

When a particle of velocity β (in units of speed of light c) crosses a medium with index of refraction n, it emits a radiation at an angle θ_C with respect to the particle trajectory such that:

$$\cos\theta_C = \frac{1}{n\beta}$$

forming a Cherenkov cone of half angle θ_C . The threshold velocity β_t, below which no radiation is emitted, is given by $(\theta_C = 0)$ yielding $\beta_t = \frac{1}{n}$, while the maximum angle of emission θ_{max} is reached for $\beta = 1$.

$$\cos\theta_{max} = \frac{1}{n}$$

The momentum threshold p_t for a particle of mass m to emit Cherenkov radiation is:

$$p_t = \frac{m}{\sqrt{n^2 - 1}}$$

In order to have full efficiency for a 15 GeV/c momentum pion, the threshold should be about 20% smaller or 12.5 GeV/c, corresponding to $(n-1)=62\times10^{-6}$ which corresponds almost exactly to Neon at atmospheric pressure. In a RICH detector, the Cherenkov cone is reflected by a spherical mirror, placed at the downstream end of the radiating volume, onto the mirror focal plane (which is located at the upstream end of the gas container). The cone image is a ring of radius

$$r_c = f\tan\theta_c$$

Where f is the mirror focal length.

The number of emitted photons N through a radiator thickness dx per unit of photon energy E is predicted by the Franck-Tamm equation as a function of the Cherenkov angle (assuming a charged particle of unit charge):

$$\frac{d^2N}{dEdx} = \frac{2\pi\alpha}{hc} sin^2\theta_c = 370\ sin^2\theta_c\ eV^{-1}cm^{-1}$$

The actual number of photoelectrons $N_{p.e.}$ (i.e. electron produced by photons impinging on the PM photocathode) is a convolution of the previous equation, the spectral response of the PM, the reflectivity of the mirror, the transparency of the gas and of any medium in front of the PM, the geometrical acceptance of the PM, multiplied by the radiator length *L*, i.e. the mirror focal length. Each of these quantity depends on the photon energy. Even if the index of refraction depends on it, it is traditional to summarize the performances of a RICH through a quality factor N_0:

$$N_{p.e.} = N_0\ L\ sin^2\theta_c$$

where a good detector can have N_0 around 100 cm^{-1}. If the probability to produce more than one photoelectron in the same PM is not negligible (as in our case) what is really meaningful is the number of fired PM N_{hit} which will be smaller than $N_{p.e.}$.

3.6.1 Expected Performances

The RICH performance expectations have been validated testing a full length prototype with particle beams. These results are described in section 3.6.6.

3.6.2 The Vessel

The RICH vessel must contain the Neon gas without leak and will be operated in slight overpressure w.r.t. the external atmosphere. The gas density must remain constant (within 1%) over long time. The PM will be located at the upstream end of the vessel: together with their support structure they must be outside the fiducial acceptance of the apparatus, bringing the diameter of the vessel to 3.8 m.

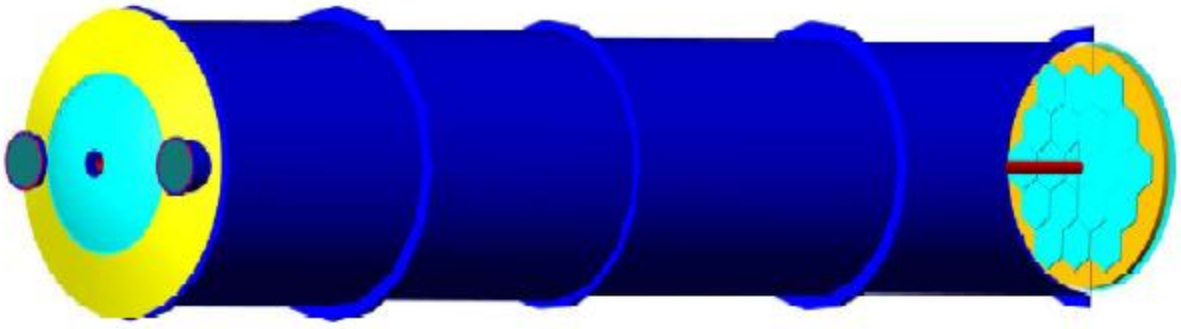

Figure 282 Schematic drawing of the RICH detector; the downstream section of the vessel is cut to show the mirrors and the beam pipe.

3.6.2.1 The Vessel Functional Requirements (FR)

FR.1. Develop a rational installation strategy with in-situ part assembly and alignment possibilities for this large size vessel.

FR.2. Provide a tight, clean and non-reflective containment to the radiator gas. Provide a stiff gas containment, keeping in mind possible pressure variations, between 0 and 150mbar overpressure.

FR.3. Provide an interface to the two photomultiplier arrays at the upstream end near 3 and 9 o'clock, keeping in mind gas and light tightness and space requests for services and access.

FR.4. Hold the mirror support panel at the downstream end; provide possibilities for fine-adjustment.

FR.5. Support a beam pipe.

FR.6. Provide minimal-material budget, in the form of "entrance window" and "exit window" respectively, of the frontal surfaces at the upstream and downstream vessel ends. The minimal-material outer diameters are dictated by criteria of downstream photon acceptance. Furthermore, minimal material near the beam axis is especially important.

FR.7. Provide as rational as possible an interface between the decay vacuum volume (the downstream end of which is materialized as STRAW 4, see chapter 3.1) and the radiator gas volume.

FR.8. Incorporate strategic flexibility, such that the important investment in hardware could be re-used with ease in a possible, second-generation Experiment with neutral, undeflected beam.

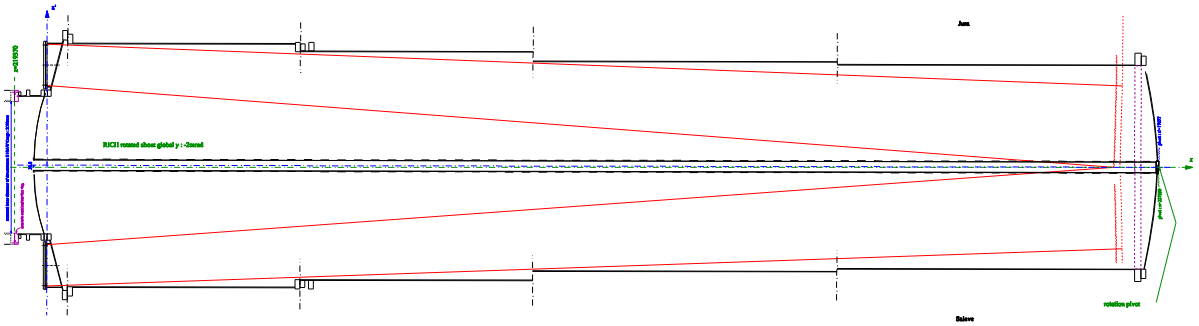

Figure 283 Main parameters of the integration of the rotated vessel in NA62; section along y=0 plane.

3.6.2.2 Proposed Solution

From the combination of FR.5 , FR.6 and FR.7 , it is logical to aim at a vessel design with centered beam pipe, and to rotate the entire vessel (by 2 mrad) with embedded beam pipe about a vertical axis, such that it be aligned along the deflected, secondary NA62 beam (see Figure 284).

Then, it follows that the minimal-material requirement extends up to r = 1.1 m at the entrance, and 1.4 m at the exit[32]. The two photomultiplier arrays count up for a diameter of about 0.7 m each. Adding up necessary overhead such as flanges, it results in a vessel about 4.2 m wide at the entrance.

It can be argued that the vessel does not have to be as high as that. Such remark would hint at making a vessel of e.g. race-track shape, resulting in a smaller radiator gas volume. However, at these scales, it

[32] Due to the photon geometrical coverage of LAV11 and LAV12, respectively.

is mandatory to aim at constructions providing stiffness by *traction*[33]. This argument drives the following shaping of the vessel parts:

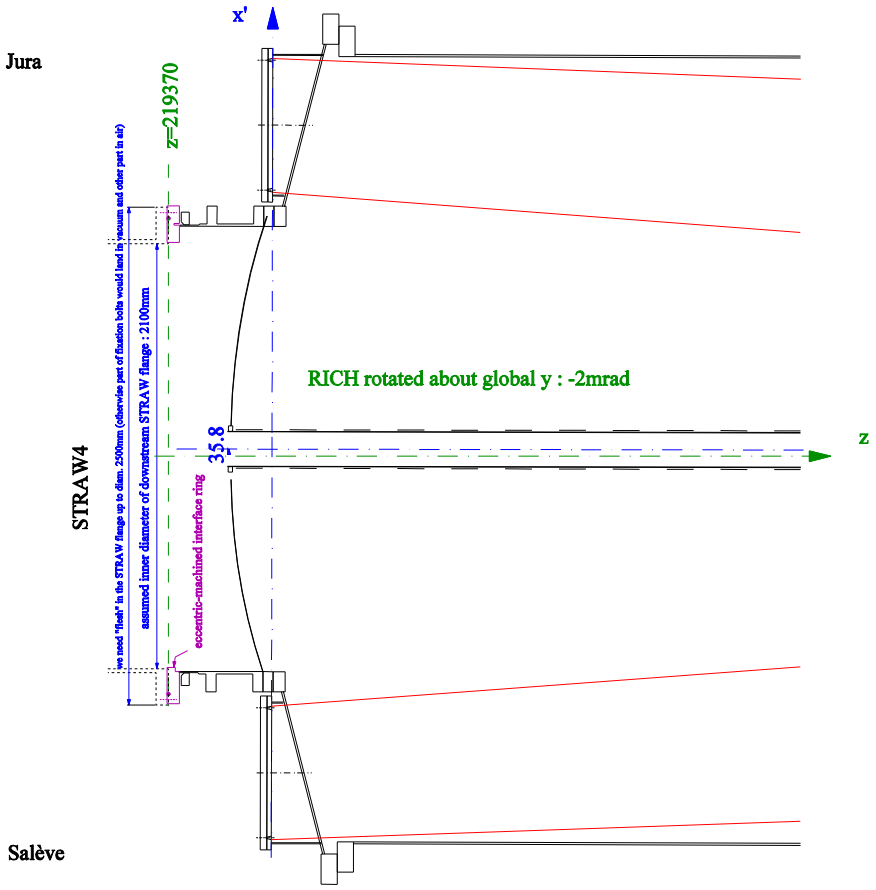

Figure 284 The RICH upstream end is offset in +X direction by 36mm to accommodate the deviated beam inside the vacuum tube. This lateral displacement with respect to STRAW4 is compensated by an eccentric interface ring (in red). The downstream end of the RICH is centered on the experimental axis, so that the whole vessel is in fact rotated by 2mrad.

- circular, rather than elliptical or race-track, outer cylinder. Outer diameter >4m at entrance;

- windows as membranes rather than bending-stiff panels ;

- bridging of the diameter gap between entrance window and outer cylinder: *conical* sheet rather than thick, flat plate perpendicular to axis. Protrude the cone by appropriate, short cylinders to receive the flanges of the photomultiplier array.

Technological difficulties with non-circular flanges are another argument in favour of a circular concept.

[33] For a load of pressure type, deflections scale with the *fourth* power of the relevant linear dimension in case the stiffening comes from *bending*, and with the *second* power in case of *traction*.

The windows would be sheets welded in inner and outer flanges, all in aluminum alloy 5083[34]. The welded assembly would be hydro-formed, to give the membrane sheet the optimal shape, and at the same time induce the optimal strain hardening in the sheet material (see Figure 285). Although detailed dimensioning has yet to be done, it is believed that the sheets would be 4 and 3mm thick, for entrance and exit window respectively.

All other components would be in ferro-pearlitic structural steels like S235 or S355. Indeed, there do not seem to exist (non-)magnetic requirements on equipment, so far downstream of the NA62 spectrometer magnet. The integrity and cleanliness of the inner surfaces would be achieved by an epoxy painting. One could think of either a two-component, room-temperature cure system, or, if possible for our dimensions, a well-known "epoxy coating", being a spraying of a mono-component epoxy powder, sticking by electrostatic forces, then cured in an oven. One should select a process resulting in a "mat", non-reflective surface. The steel sheet thickness will be driven by technological issues, and is believed to range between 5 and 10mm[35].

The cylindrical container shall be made in several parts, to be bolted together in-situ. The parts can be of gradually somewhat diminishing diameter. The diameter of the last part is driven by the size of the mirror support panel[36]. Intermediate parts shall not hinder the reflected light. The concept of interface between adjacent cylindrical vessel parts is a copy of that of the vacuum and helium tanks of NA31/NA48. It allows small adjustments in rotation about all three principal axes, prior to compressing the O-ring seal.

The total cylinder length is about 17m, driven by the mirror focal length.

There is a cast-in-concrete bottle neck of passage at the downside of the main shaft to the ECN3 cavern, being about 3.87 m wide. Thus, we have set a limit of 3.83 m on the dimension of each vessel part. Parts of greater diameter have to be kept below limit length, and vice-versa, long parts shall not exceed the limit diameter. We have four parts in the current design[37].

The remaining main steel part in the vessel assembly is the "cone" in front. Even though short in *z*, it has the greatest diameter of all parts, and is quite complicated.

Upstream of the entrance window, there is an intermediate vessel piece bridging the gap to STRAW 4 (see Figure 284 and Figure 285). The RICH entrance window is thus the unique separator between the decay vacuum volume and the RICH radiator gas volume. The intermediate piece would also be in structural steel, and is a vacuum vessel. The entrance window has to stand up to 1.15 bar. The rest of the vessel, including exit window, has to stand up to 0.15 bar. A special interface ring will allow the

[34] 4.5% Mg, 1% Mn, mass fractions.

[35] Indeed, 1mm or even less would already be sufficient to stand the pressure.

[36] Currently, the support panel rather than the minimal-material aperture requirement, drive the outer diameter of the exit window.

[37] It turns out that ECN3 passage, rather than transport, sets the stricter requirement.

intermediate vessel piece to be connected to STRAW 4. For the initial Experiment with charged beam, this ring will have to be machined with eccentricity, in order to cope with the transverse offset of the RICH, about 43 mm. For a future neutral beam run, a new, concentric, ring would be the only necessary piece of new equipment.

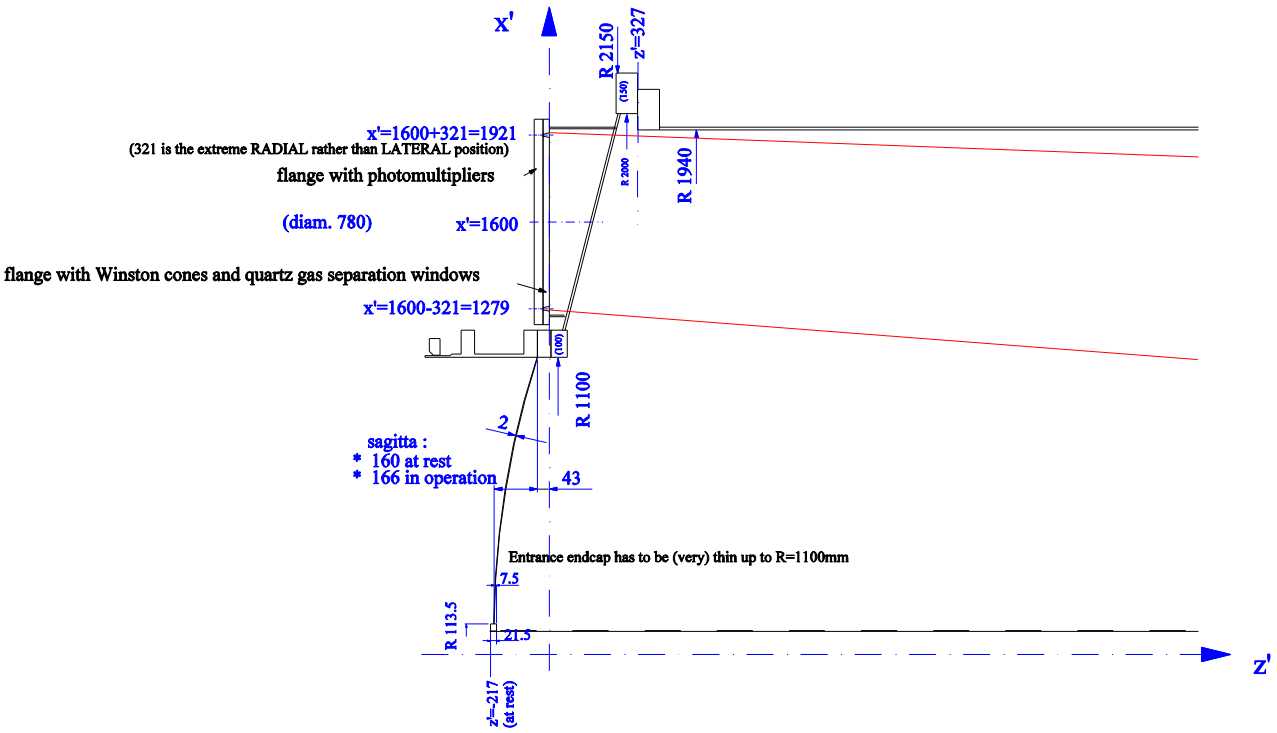

Figure 285 Upstream extremity of the RICH showing the PMT flange and the aluminum beam window (section along y=0 plane). The beam pipe diameter near the upstream end of the RICH has been designed to be Ø ≥ 170 mm in order to prevent any material in front of the SAC photon acceptance.

In order to minimize material, especially at the upstream side, the beam pipe "starts" at the inner flange of the entrance window, with a minimal-engineering connection, thus simply flanged. Longitudinal decoupling is sought far downstream, in front of the LKR cryostat. Consequently, the beam pipe assembly in the RICH and even further downstream, is allowed to move, in z, according to any displacement of the entrance window. Inside the vessel, pipe weight will be supported, and transverse movement blocked, by thin wires. The exit window is decoupled from the vacuum system. The beam pipe continues through the window's inner hole. The window is connected to the beam pipe with a special bellows, providing radiator gas tightness. The window could be taken off from the vessel without disturbing the vacuum (see Figure 286).

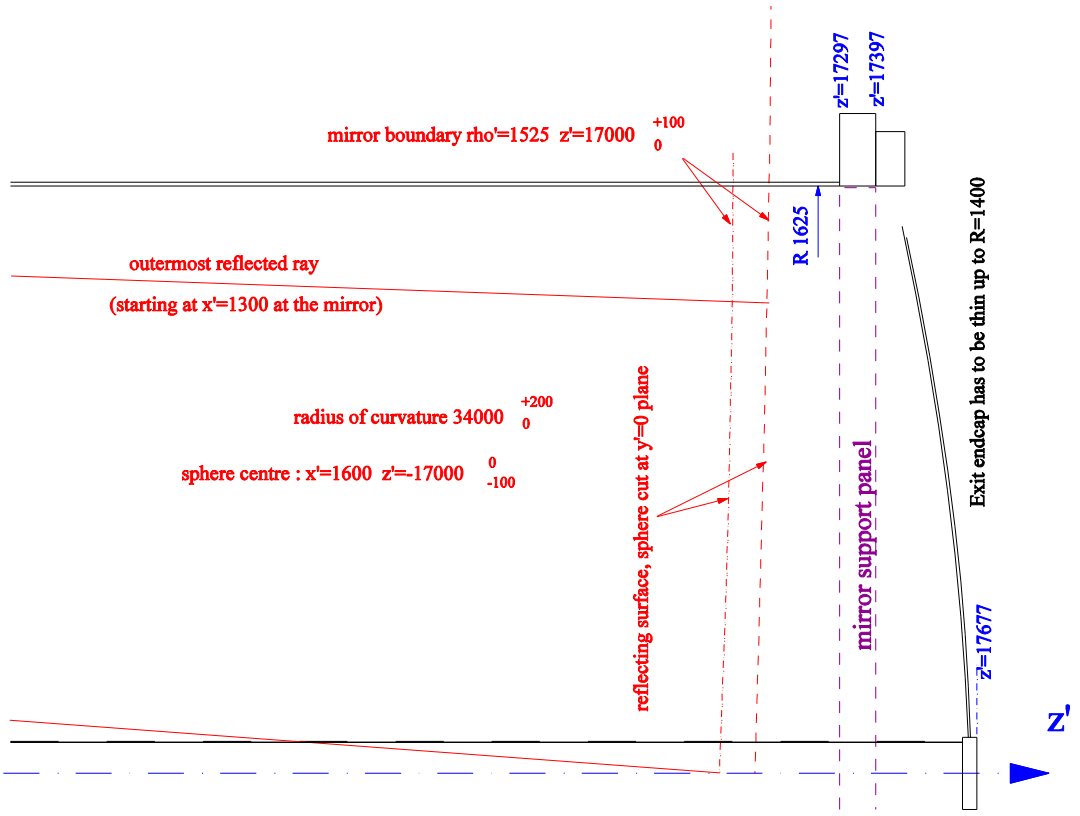

Figure 286 Main parameters of the vessel near the downstream end; section along y=0 plane.

3.6.3 The Gas System

There are five important requirements that need to be met by a suitable radiator gas for the NA62 RICH detector.

- Appropriate refractive index at atmospheric pressure
- Good light transparency in the visible and near UV
- Low chromatic dispersion
- Low atomic weight to minimize radiation length
- Non-flammable

Looking at the above criteria it is relatively straightforward to identify Neon as the most appropriate choice. Besides its suitable refractive index, Neon has a small chromatic dispersion combined with a low atomic mass.

The RICH operating range starts at wavelength > 190nm which makes its performance quite insensitive to impurities in the gas. For example, the light absorption for trace pollutants of Oxygen and H_2O become noticeable only at wavelength below 190 nm. However, residual gases affect also the refractive index, which deteriorates the Cherenkov angle resolution, if their concentrations vary in time.

CO_2 is used as transition gas during the filling and emptying of the radiator and residual quantitative could therefore be present in the gas. Figure 287 shows the light transparency of CO_2 for an average light path length of 25.5m and different CO_2 concentrations. The plot indicates that for low CO_2 concentrations the transparency is near 100% down to wavelength of 195 nm. Test beam measurements confirm that the signal loss due to a residual CO_2 concentration of 1% is smaller than 0.3%.

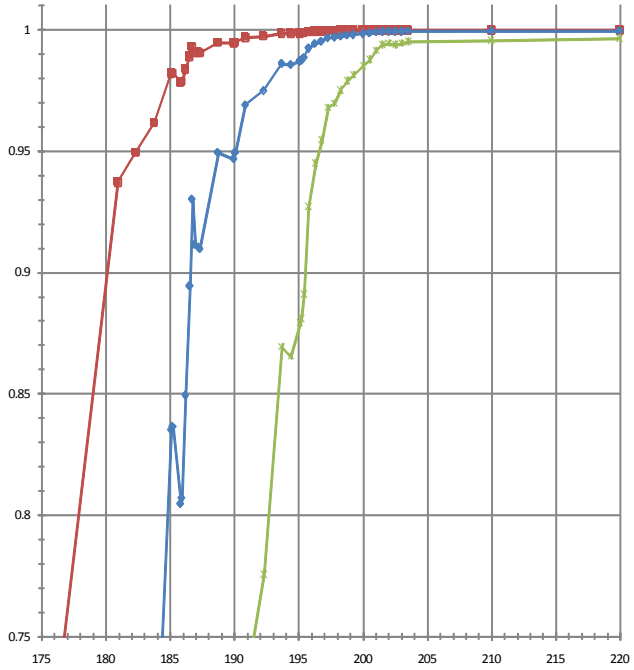

Figure 287 Relative light transparency in residual CO_2 gas for a light path length of 25.5 m, for 0.1% (red), 1% (blue) and 10% of CO_2 (green) [83].

A number of relevant detector and gas system parameters are summarized in Table 42.

Table 42 Relevant detector and gas system parameter.

RICH radiator volume	≈ 200 m^3
Operating gas	100% Ne
Residual trace gas concentrations: CO_2 O_2 H_2O N_2	 < 0.1% < 0.1% < 0.1% < 0.5 %
Operating pressure	Between 990 and 1010mbar (absolute)
Max. over/under pressure (rel. to atmosphere)	+ 150 mbar - 0 mbar
Pressure regulation accuracy	± 1 mbar
Flushing and purging flow rates	3 to 10 m^3/h
Gas density stability	< 1%
Detector leak rate	< 1.0 * 10^{-2} Std. cc/s

3.6.3.1 Outline of the Gas Modules

The Rich gas radiator is a cylindrical vessel with an overall volume of 200m^3. For mechanical reasons the gas pressure needs to be kept above atmospheric pressure with a maximum overpressure of 150 mbar.

To reduce design overheads and optimize long-term maintenance cost the gas system is build-up of functional modules (Table 43), which are technically very similar to the ones used in the LHC experiments.

Table 43 Gas system modules and their locations.

Module	**Situated in**
Neon Supplies	Gas Building B920
Distribution Module	Experimental Area ECN3
Circulation and Pump Module	Build. 883
CO_2 Absorber Module	
Buffer Volume (size 1 m^3)	
Purifier (optional)	

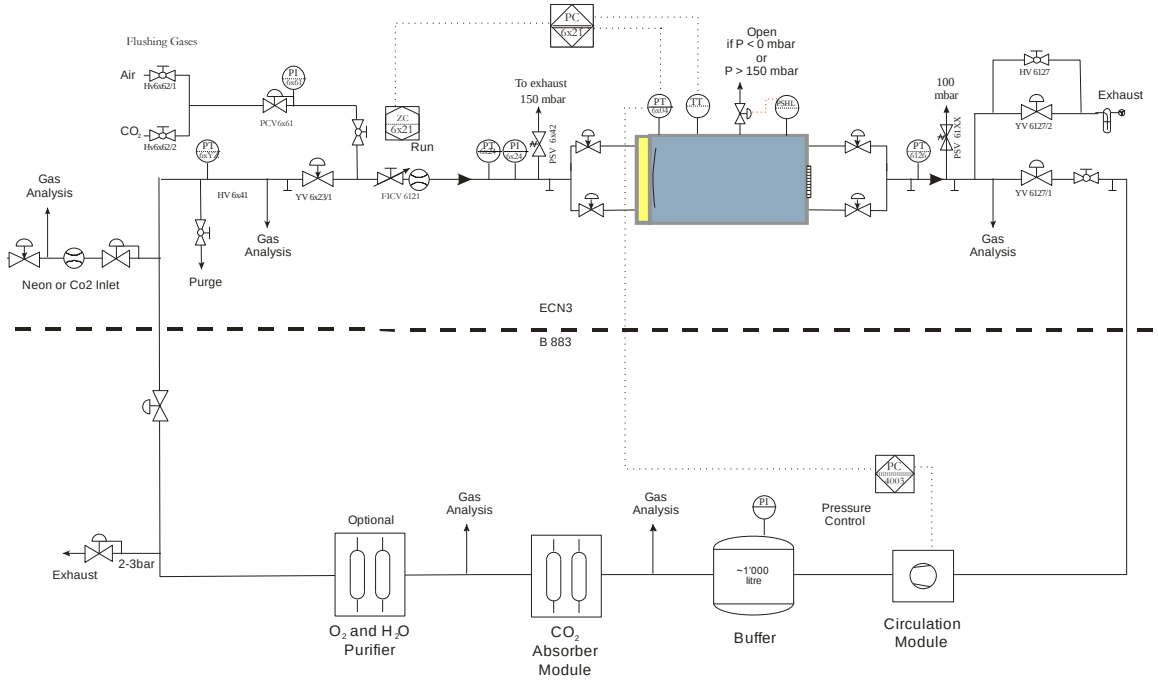

Figure 288 Schematic layout of the RICH gas system (blue = RICH radiator volume). The upper part shows the distribution located in the experimental hall. The lower part shows schematically the modules situated in B883.

A schematic drawing of the gas system is shown in Figure 288. The large volume combined with a relatively expensive gas (135 CHF/m^3) requires special precautions for a cost effective filling of the detector. As the vessel is unable withstand vacuum, the radiator volume is first filled with carbon-dioxide to replace air. The gas is then circulated in closed loop, and the Neon is introduced while absorbing the CO_2 in a molecular sieve filter.

All remote valves, instruments and sensors (pressure, temperatures and flows) are read-out by a programmable control computer (PLC) allowing remote filling or emptying of the detector.

In the following a short technical description of the modules is given.

3.6.3.1.1 The Distribution Module (DM)

The DM is located in the experimental cavern near the detector itself; a schematic layout of the module is shown in the upper part of Figure 288. Two gas inlets and two gas outlets, one of each at the bottom and on top of the detector, connect the distribution rack to the radiator, thus the higher density CO_2 gas can be filled from the bottom and the lower density Neon from the top. The inlet flow is metered using a variable area flow meter; pressure transmitters measure the pressures at the inlet, inside the radiator and at the outlet.

All four gas lines have pneumatic shut-off valves, so that the detector can be hermetically isolated from the system. The DM contains a passive and pneumatically driven pressure protection, which open the radiator volume to atmosphere if the pressure exceed the mechanically allowed limits ($0 < P_{vessel} < 150$ mbar).

From the DM the radiator can be purged manually with CO_2 or air in single pass mode.

3.6.3.1.2 The Circulation Module (CM)

When running in closed loop mode the outlet gas from the radiator must be compressed for return to the surface and re-cycling through the CO_2 absorber. Using a dry membrane compressor the gas is compressed by about 0.5 to 1.0 bar.

The CM is shown in Figure 289. The return gas from the radiator is entering the circulation module and the pressure at this point is regulated by means of a remotely controlled bypass valve around the compressor.

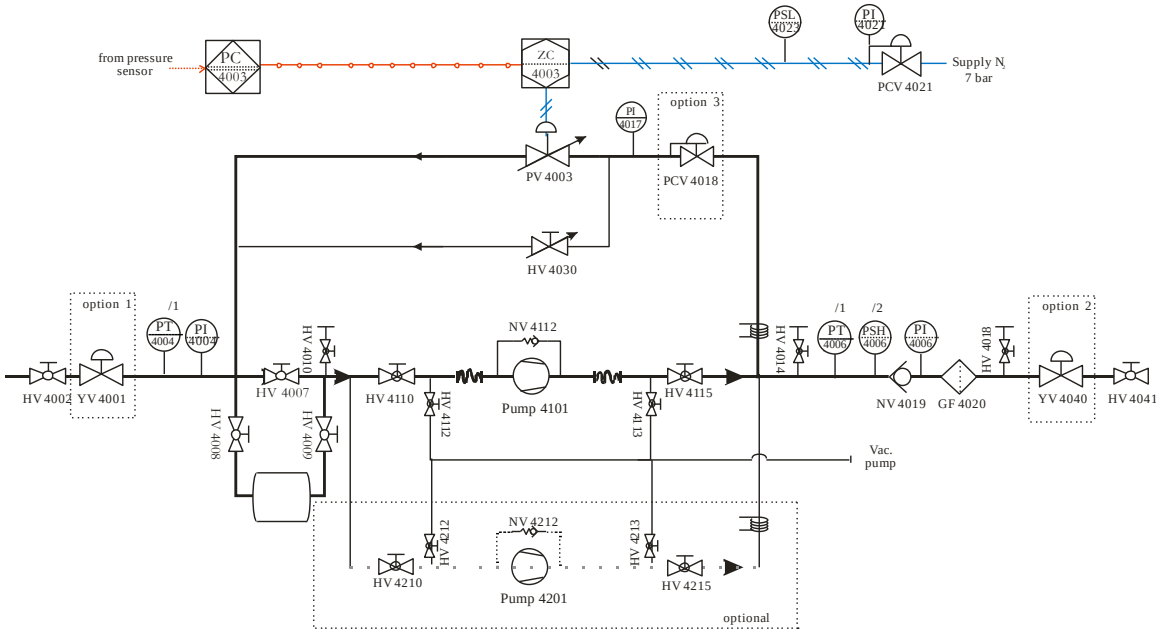

Figure 289 Component layout of the pump or Circulation Module.

3.6.3.1.3 The CO2 Absorber Module (AM)

The motivation for the CO_2 AM is a cost effective filling of the large radiator volume in closed loop circulation with a minimal loss of Neon (see section 3.6.3.2).

The AM (Figure 290) consists of two absorption cartridges of identical size (24 litre volume), which are installed in parallel. One of the cartridges is in operation while the second one can be regenerated. To allow a fast and efficient regeneration the cartridges are equipped with heating jackets (+180°C) and active cooling.

Each cylinder contains approximately 15 kg Molecular Sieve 13X, and can absorb up to 3.5 Nm^3 of CO_2 gas. After saturation with CO_2 the cartridges can be regenerated by either:

- lowering the pressure to a few mbar (absolute) at room temperature, or by
- lowering the pressure and heating up to 180° C

The first method is very fast (regeneration time between 10 and 20 min.), but yields a lower absorption capacity. The second method combines low pressure and high temperature prolonging the cycle time to approximately 5 hours, but it provides a higher absorption capacity and an outlet purities of a few ppm residual CO_2 concentration.

To absorb 200 m^3 of CO_2 between 100 and 150 cycles are required if both methods are used. In the beginning the returned CO_2 concentrations is near 100% and one cartridge will saturate in 10 to 20 minutes. As the CO_2 concentration in the loop diminishes the saturation time of the cartridges prolongs gradually. When the CO_2 fraction reaches a few %, the 2nd regeneration method can be used, to reduce the residual CO_2 fraction after the absorber module to less than 5 ppm.

The CO_2 concentrations in the gas are monitored at the inlet and outlet of the module.

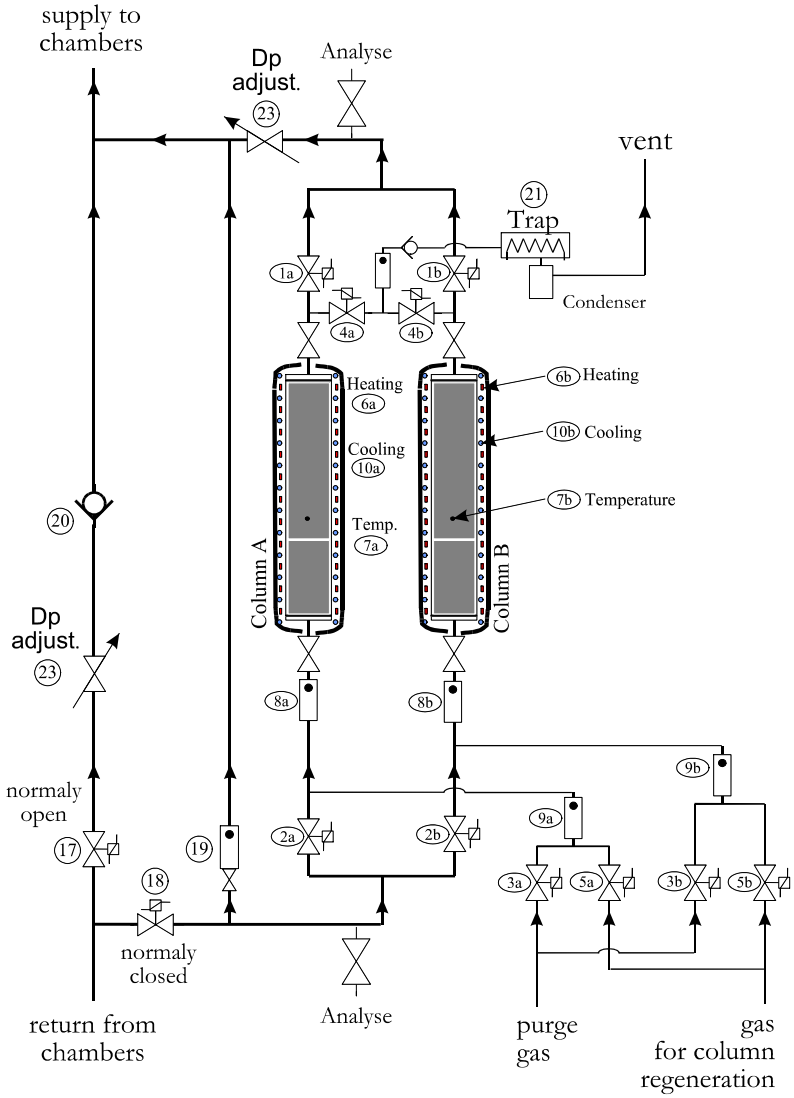

Figure 290 Schematic layout of the CO2 Absorber Module.

3.6.3.2 Filling and Emptying of the Detector

Before introducing the Neon, the entire system (gas system and radiator) is filled with clean CO_2 to remove air from the system. When the residual air in the radiator is at the required level the inlet flow is switched to Neon. As soon as the Neon fraction in the radiator outlet flux exceeds a few % the system is changed to closed loop circulation and the CO_2 absorber is switched on. The circulation flow rates can range from 3 and 10 m^3/h, so that in the beginning the capacity of one absorber cartridge will be used up quickly, and if high flow rates are used the absorption may be interrupted by the regeneration cycles of the CO_2 absorber. The Neon input flow is regulated by the pressure in the system and will, therefore, automatically match the quantity of CO_2 absorbed. When the residual CO_2 concentration in the radiator is at the required level, the circulation flow can be stopped, and the radiator pressure is ramped up (or down) to the desired value.

To empty the Neon, the vessel can be purged in single pass mode with CO2 or Air.

3.6.3.3 Running of the Detector

The neon density influences the refractive index n following the relation

$$n = 1 + (n_0 - 1)\,\frac{\rho}{\rho_0}$$

where n_0 is neon refractive index at NTP (1.000067), ρ_0 is the neon density at NTP (0.9001 kg/m^3), and ρ the density at operating conditions (≈ 0.814 kg/m^3 for T= 25°C and P= 1 bar).

Since the detector performance is rather immune to impurities one can envisage operating the detector with a sealed gas volume without continuous gas renewal. This concept is advantageous because temperature variations will in this case not influence the gas density.

At the beginning of a physics run the radiator pressure will be set accurately to the desired operation pressure, e.g. 1000 mbar (absolute). The gas density is then left constant by closing all valves to the radiator volume.

To guaranty sufficient stability for long physics runs the leak rate of the detector should be smaller than 1.0 * 10^{-2} Std. cc/s. In case the leak tightness or the gas purity requirements cannot be met it is possible to upgrade the system with a density feed-back loop and a purifier module, so that the gas system can be operated in closed loop with in-line purification. In this case the radiator pressure P_{Set} can be regulated such that the gas density $\rho_{const.}$ remains unchanged following the average Temperature T_{Vessel} of the vessel:

$$P_{set} = \frac{\rho_{const}}{M_{Ne}} * R\,(273.15 + T_{vessel})$$

R is the universal gas constant (8.314472 J/K mol) and M_{Ne} the atomic mass of Neon (20.179 gr/mol).

3.6.4 The Mirrors

A mosaic of spherical mirrors is used to image the Cherenkov cone into a ring on its focal plane. To avoid absorption of reflected light on the beam pipe the mirrors are divided into two spherical surfaces: one with the centre of curvature to the left and one to the right of the beam pipe. The total reflective surface exceeds 6 m^2, so that it is practical to use a matrix of 20 mirrors (see Figure 293 and Figure 294).

3.6.4.1 Mirror Layout and Quality

The RICH detector uses spherical mirrors with a nominal curvature radius of 34m and hence a focal length of 17m. Twenty mirrors will be needed, 18 of hexagonal shape (35 cm side) and two of semi-hexagonal shape (see Figure 291 and Figure 292). The two latter ones are used in the centre and have a circular opening to accommodate the beam pipe. The mirrors are made from 25mm thick glass substrate coated with aluminium. A thin dielectric film is added for protection and to improve the reflectivity.

There are three main optical parameters that must be met by each mirror:

- A D_0 not larger than 4 mm (the D_0 is the minimum diameter of a circle which collects 95% of the light of a point–like source placed in the curvature centre)
- A radius of curvature within ±20 cm from the nominal one.
- In the wavelength range between 195 and 650 nm the average reflectivity should be better than 90%.

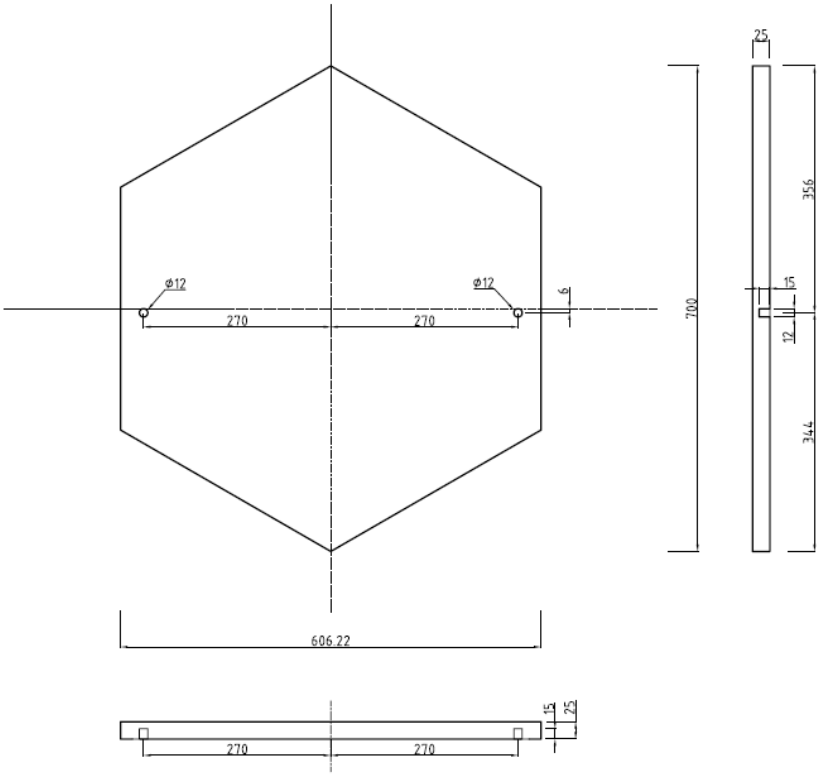

Figure 291 Dimensions of an hexagonal mirror. The positions of the two holes to support the mirror are indicated.

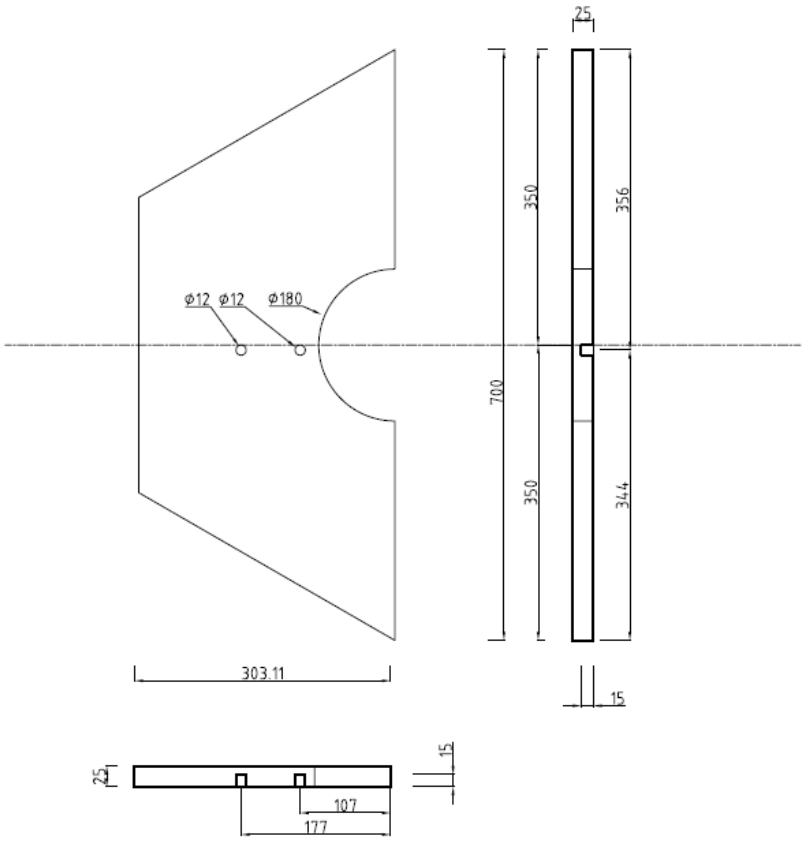

Figure 292 Dimensions of a semi-hexagonal mirror. The positions of the two holes to support the mirror are indicated.

3.6.4.2 The Mirror Supports

The mirror support structure has to sustain a total weight of approximately 400 kg and must guarantee long-term stability of the mirror positions. On the other hand it is important not to add too much material to minimize photo-nuclear reactions which can deteriorate the veto capabilities of the subsequent detectors. In that context, a 10 cm thick carbon-fiber or aluminium honeycomb structure has been chosen as mirror support structure (see Figure 294).

Each mirror must be supported and adjustable for alignment. Two holes, 12 mm wide, will be drilled on the not-reflecting surface of the mirror leaving the barycentre in the middle of them, with an inter-distance of 540 mm for the hexagonal mirrors and of 70 mm for the semi-hexagonal ones. A dowel with a spherical head will be inserted on each mirror hole and will be used to sustain the mirror. Piezo-electric actuators(see Figure 295) will allow a remotely controlled orientation of each hexagonal mirror. The two half-hexagonal mirrors in the center are fix.

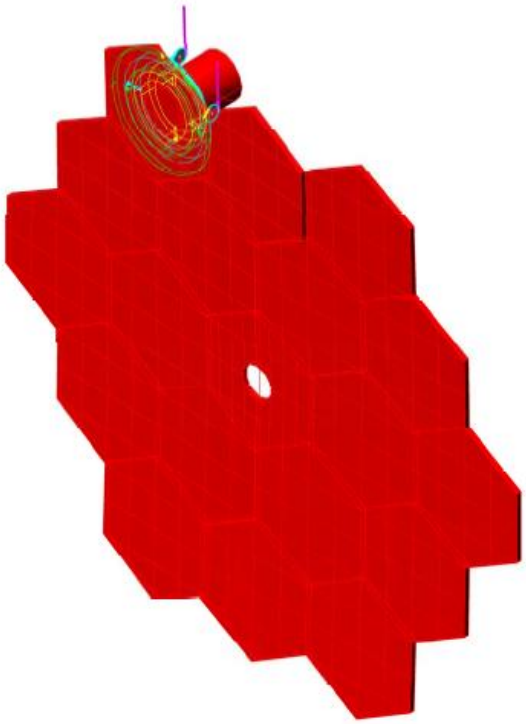

Figure 293 The mirrors layout with the detail of the support for one mirror.

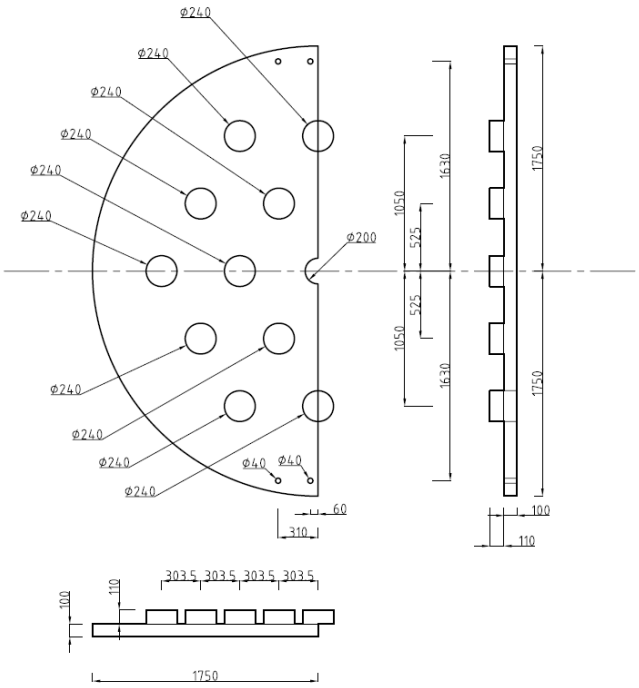

Figure 294 Half of the Mirror support wall.

The Piezoelectric actuators[38] have a ±35 mm travel range with 1 nm resolution and can produce up to 20 N of push/pull force. These devices are self-locking and remain in the same position if the supply voltage is turned off. The piezoelectric actuator size is 22×10.8×21 mm.

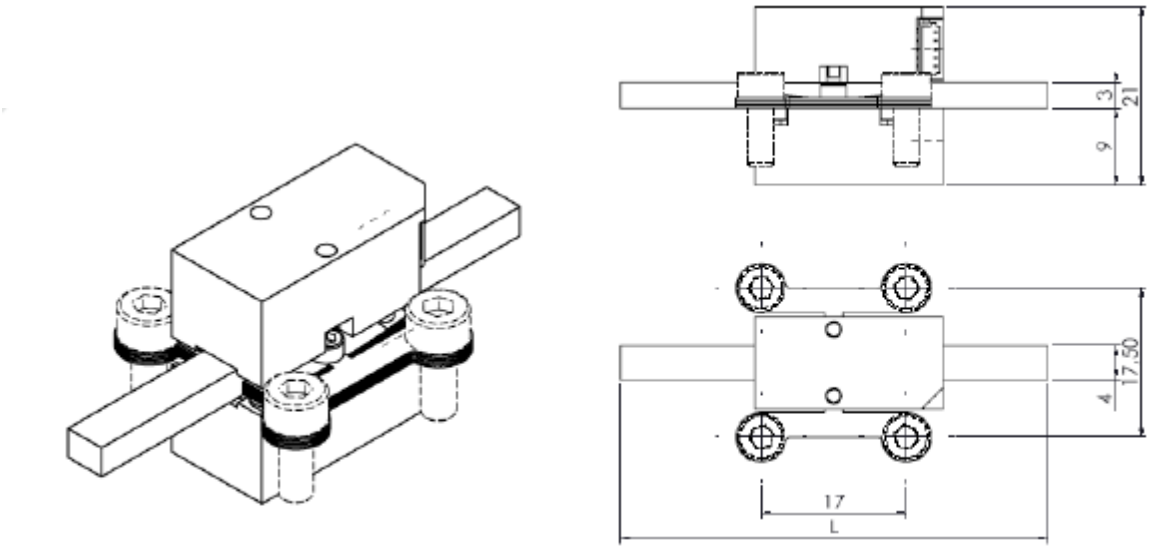

Figure 295 The piezoelectric actuator LEGS-LT02SV-10 from PiezoMotor.

[38] type LEGS-LT02SV-10 produced by PiezoMotor Uppsala, Sweden.

3.6.4.3 The Mirror Alignment

The mirror alignment will be routinely checked during data taking selecting particles whose Cherenkov rings are completely contained in a single mirror and with a track angle measured by the magnetic spectrometer. However, the first time and perhaps later during the periodical shut-down a laser alignment will be done. The laser alignment starts with one mirror that will be chosen as reference. A laser beam, originating in the focal plane, will be directed to that mirror and oriented so that the reflected beam impinges on a selected point in the focal plane. The laser beam will be translated, parallel to itself, to a neighbour mirror by means of two high quality retro-reflectors: this mirror can then be aligned in such a way that the reflected laser beam impinges on the same point as for the reference mirror. The laser beam needs to be adjusted close to the borders of each mirror, because an accurate translation can be done only over a few centimetres. This procedure will be iterated to cover all the mirrors. Two reference mirrors will be chosen, one for the mirrors pointing to the left and one for those pointing to the right of the beam pipe. The laser beam will be prepared as parallel as possible using a telescopic optical device.

3.6.5 The Photon Detection System

The granularity of the photon detection is an essential parameter to optimize the angular resolution of the detector. On the other hand the number of photo detectors has large impact on the cost of the RICH apparatus. A reasonable compromise between the number of sensors, the photon acceptance and the sensor dimensions must be found. According to the experiment simulation, a total of 2000 photomultipliers is sufficient to match the detector requirement. This arrangement leads 1000 PMT's on each side and to a photo-sensor pitch of 18mm.

3.6.5.1 The Photomultipliers (PMT)

Several photomultiplier options have been investigated and the best compromise between fast response, small dimensions and cost seems to be the Hamamatsu R7400 series (see Table 44). These metal cover single anode photomultipliers exist in several versions and three of them have been tested in the laboratory and in the beam.

- The U-04 was considered for its spectral response ranging from the red to close UV but was discarded because of its low efficiency compared to other versions.
- The U-06 version has a quartz entrance window which allows to be sensitive up to 165 nm.
- TheU-03 model has a UV glass entrance window which cuts the sensitivity at 185 nm.

Due to other factors (mirror reflectivity, neon transparency, separation windows in the vessel, etc.) no major difference has been seen between the two latter types of PMT's in a test beam held in 2007, therefore the cheaper U-03 has been selected. In Table 44 the main properties of the PM (at 800 V) are summarized.

Figure 296 and Figure 297 show the dimensions and the spectral sensitivity of the R7400U-03 PMT . The selected PMT's have a minimum gain of 5×10^5 at 800 V; the first delivered 328 PMT's have an average gain of $(8\pm2)\times10^5$ at 800 V.

Table 44 Main characteristics of Hamamatsu R7400U-03 photomultiplier at 800 V

PM	R7400U-03
Type	Head on; Metal PackageTube for UV-visible range
Size	16 mm
Active diameter	8 mm
Min λ	185 nm
Max λ	650 nm
Peak sensitivity	420 nm
Cathode Radiant Sensitivity	62 mA/W
Window	UV glass
Cathode Type	Bialkali
Cathode Luminous Sensitivity	70 mA/lm
Anode Luminous Sensitivity	50 A/lm
Gain	7.0×10^5
Dark current after 30 min.	0.2 nA
Rise Time	0.78 ns
Transit Time	5.4 ns
Transit Time Spread	0.28 ns
Number of Dynodes	8
Applied Voltage	800 V

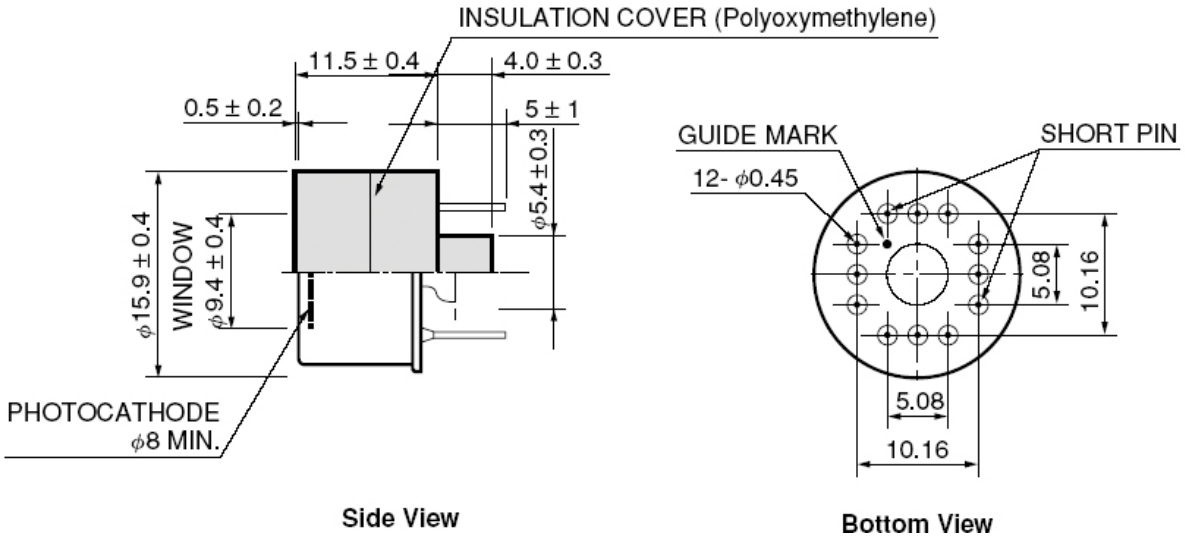

Figure 296 Dimensions of the Hamamatsu R7400U-03 photomultiplier: side view (left), rear view (right).

The 8 dynode voltages are provided by a Hamamatsu E5780 HV divider (2.8 MΩ total resistance) with cylindrical shape with 17.0±0.2 mm diameter and 15.0÷0.5 mm length. It has three connectors: a RG-174/U cable for signal output and two AWG22 cables for the high (negative) voltage supply and for grounding.

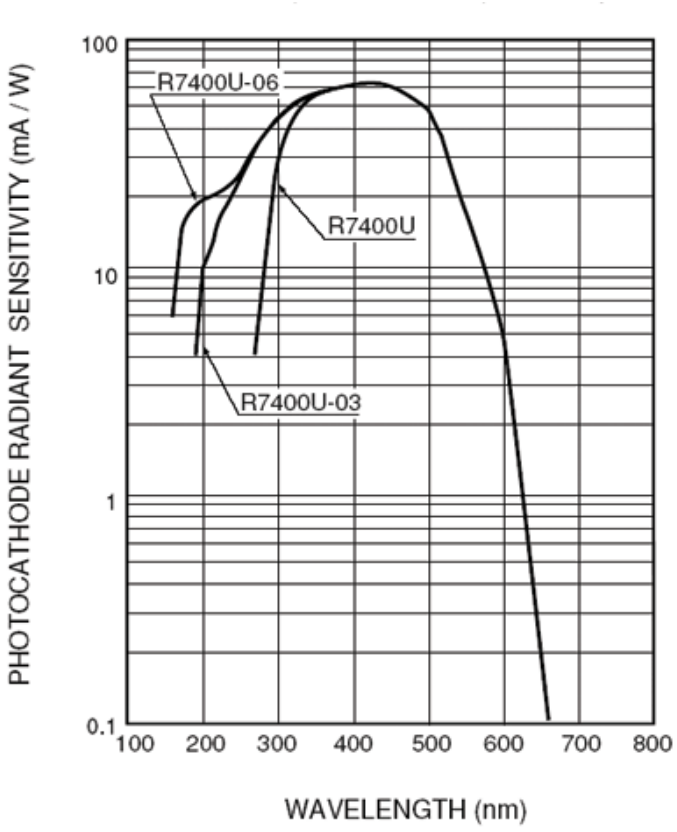

Figure 297 Spectral sensitivity for different Hamamatsu R7400 photomultipliers.

The PMT time response is of great importance for the RICH and has been investigated in detail. The Transit Time Spread given by the supplier is expressed in FWHM because the response has important tails and is far from a Gaussian shape. Several systematic contributions have been identified in the time distribution:

- a small peak before the main peak, attributed to the electron extraction from the first dynode instead from the photocathode;
- a second small peak about 1.2 ns after the main peak is attributed to electrons extracted from the photocathode, after reflection from the first dynode towards the cathode and coming back to the same dynode.

3.6.5.2 The PMT Power Supply

The PMTs need 800-1000 V negative voltage (we assume 900 V in the following) supplied with low noise. To reduce cost, four PMTs are supplied by a single HV channel, so that about 500 HV channels are needed. CAEN 12 or 24-channels HV boards[39] fulfil the requirements and they can be inserted in appropriate standard CAEN crates[40]. The HV system will be remotely controlled by the DCS system. The 12-channel s A1733N board has a maximum of 4(3) kV output voltage (dual range) with 2(3) mA

[39] CAEN 12-channels A1733N or 24-channels A1535S boards

[40] CAEN SY1527 (16 slots) or SY2527 (6 slot)

maximum output current. The voltage and current resolutions are 250 mV and 200 nA, respectively; the voltage ripple is smaller than 30 mV pp. The 24-channels A1535S has a maximum 3.5 kV output voltage, with 500 mV resolution, and 3 mA output current, with 500 nA resolution; the typical voltage ripple is smaller than 20 mV pp; all HV channels have a common floating return.

3.6.5.3 The Frontend Electronics

The PMT output signal has a roughly triangular shape with the same rise time as the PMT (0.78 ns on average) and a fall time about twice the rise time. At 900 V PMT supply voltage (average gain of 1.5×10^6) the output charge is about 240 fC corresponding to a peak current of 200 µA and to a negative peak voltage of 10 mV over 50 Ω. There is also a large variation in gain performances among the PMT's. In order to profit of the fast PMT response, the 8-channel NINO ASIC [84] was chosen as discriminator. This chip has an intrinsic resolution of 50 ps and was developed for the output signal of multigap resistive plate chambers. To match the optimal NINO performance region, the PMT output is sent to a current amplifier with differential output: a 32-channel customized printed circuit will be used for this purpose, sending the output to a board containing 4 NINO ASIC. The NINO chip is operated in time-over-threshold mode and its LVDS output signal is sent to 512-channels TELL1 boards equipped with HPTDC chips with 97.7 ps LSB producing 19 bits long words (corresponding to a maximum of 51 µs). Both the leading and trailing edge of the LVDS signal are recorded providing information on the original signal width, used for the time slewing correction in the analysis.

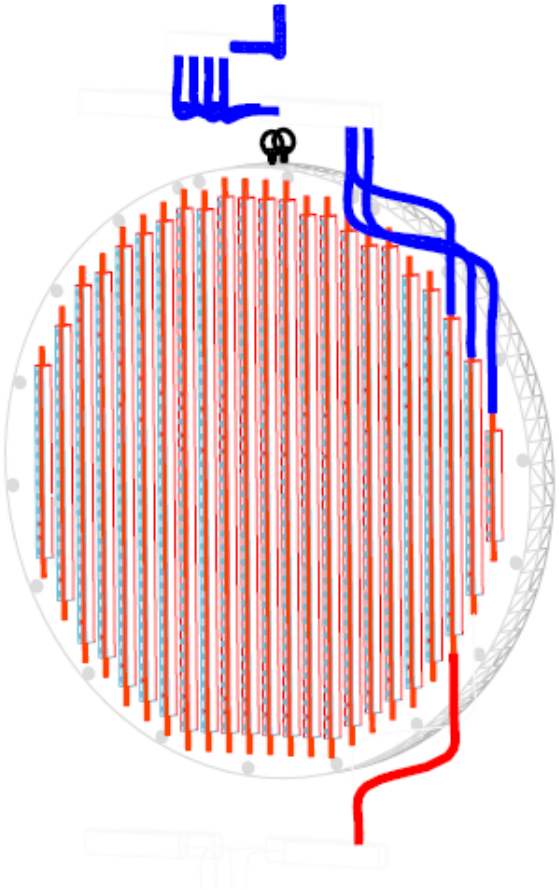

Figure 298 View of one of the flanges with 1000 PMT's. The cooling system is shown.

3.6.5.4 The PMT Support

To simplify the access to the photomultipliers the PMT's are mounted outside the radiator gas volume. This PMT assembly consists of two independent flanges: a radiator flange with quartz windows and an independent aluminium flange holding the PMT's. A schematic cross-section of the arrangement is shown in Figure 300. The aluminium flange holding the PMT's is shown in Figure 299. The radiator flange is shown in Figure 298.

The light exit holes on the radiator flange are covered with a highly reflective aluminized Polyethylenterephthalat (trade name Mylar) foil and they have the shape of a truncated circular parabolic (also called "Winston Cone") to tunnel the light through the window aperture. On the outside of each hole (the PM side) a 1 mm deep, 14 mm wide cylindrical hole allows to position and glue the quartz window (a 0.5 mm deep, 2 mm thick groove accommodates the glue).

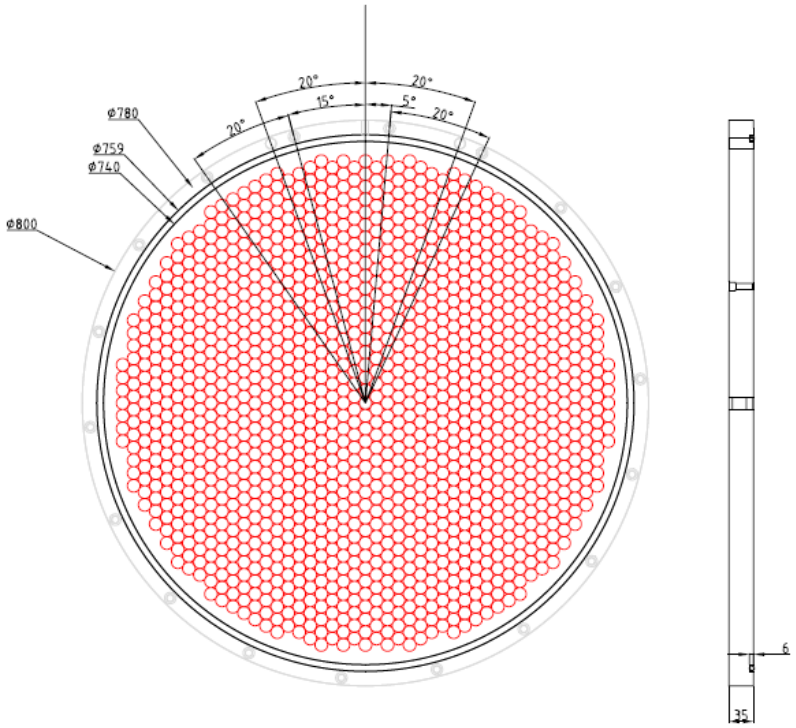

Figure 299 One of the aluminum flanges which holds the PMT's.

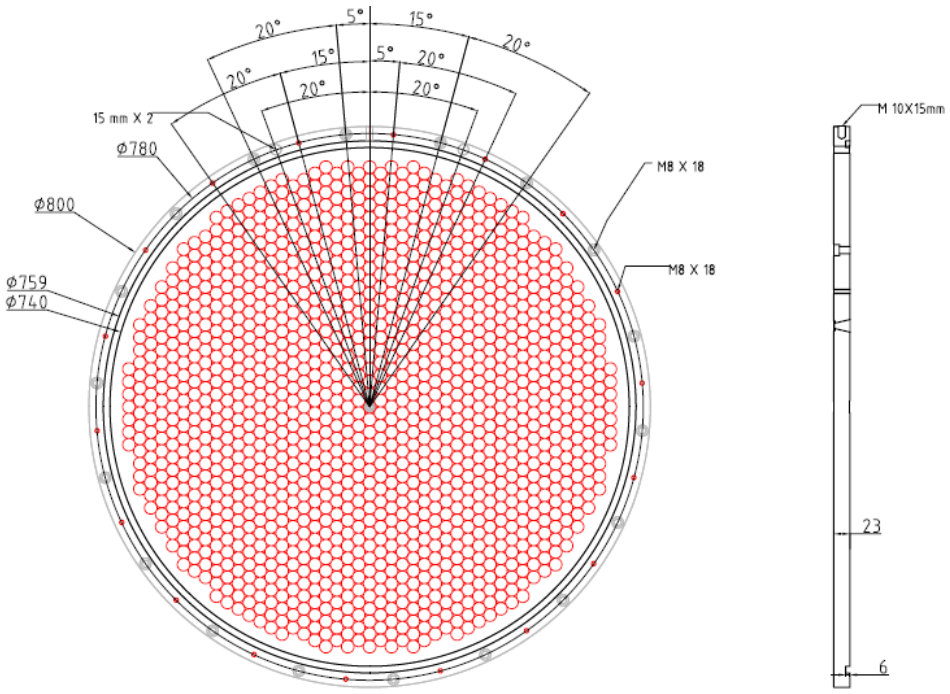

Figure 300 View of one of the two stainless steel flanges which separate the Neon gas from the PMTs.

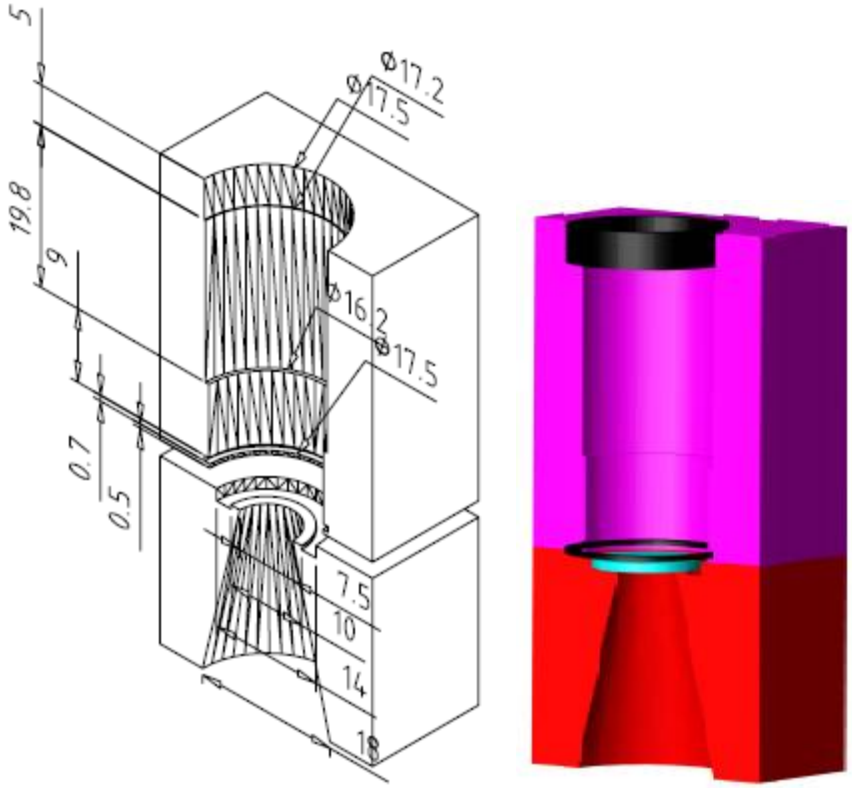

Figure 301 Detail of the PM support structure: the upper part (purple) is the aluminum flange holding the PMTs with the O-ring (black); the bottom part(red) is the flange of the radiator volume with the quartz window in blue.

The PMT's are mounted on the external aluminium flange in front of the quartz windows. They are isolated from environment light and water cooled. A cylindrical hole, 16.5 mm wide and 10 mm high will be drilled in the aluminium flange for each PM, followed by a 17.5 mm wide and 22 mm high hole for the HV divider. A 1 mm thick O-ring (with 17.5 mm outer and 13.5 mm inner diameter) will be placed in front of the PMT (kept in place by a 2 mm thick gorge in the hole placed 0.5 mm above the end) and pressed against the quartz window to avoid external light to reach the PMT. A 5±1 mm thick O-ring (with the same outer and inner diameter of the 1-mm O-ring) will be placed on the back of the PMT (after the end of the HV divider) to close the hole and avoid external light; this O-ring also guarantees a good thermal contact between the PMT and the aluminium flange and absorbs the tolerance in the PMT total length .

3.6.5.5 The Cooling System

The PMT HV divider will dissipate about 0.3 W per tube or 300 W per side (1000 PM) and a cooling system is required to avoid a local heat source on the radiator. A water cooling pipe (ϕ=10mm copper) situated inside an aluminium rod in between the PMT rows is used to remove the heat. An industrial chiller[41] unit is used to cool and pump the water .

3.6.6 Prototype Test Beam Results

A RICH prototype was built to demonstrate the feasibility of the RICH project. In 2007 the prototype vessel was installed in the ECN3 cavern and commissioned. The radiator was cylindrical, 17 m long and 60 cm diameter, filled with Neon gas at atmospheric pressure. A spherical mirror with 17 m focal length was used with PMT's placed in the mirror focal plane.

3.6.6.1 The 2007 Prototype Test Beam

In fall 2007 the RICH prototype was equipped with 96 PMT's of different types and exposed to a 200 GeV/c momentum negative pions beam derived from the CERN SPS. An event time resolution of 65 ps was measured with an average number of 17 hits per event; the Cherenkov angle resolution turned out to be 50 μrad with some bias induced by the limited number of PMT's used in this test. Both Hamamatsu R7400U-03 and U-06 PM were tested and the latter type showed a 30% broader time resolution than the first one without a major improvement in light collection efficiency (the U-06 has a quartz entrance window). This result drove the choice of the U-03 type for the final detector.

The results of the 2007 prototype test beam have been published in [85].

3.6.6.2 The 2009 Prototype Test Beam

In spring 2009 the RICH prototype was equipped with 414 PMT's of Hamamatsu R7400U-03 type to validate the $\pi-\mu$ separation. The prototype was exposed to a positive beam derived from the CERN SPS with adjustable beam momentum. Positrons had an average number of 19 hits per event; the time resolution of the previous test was confirmed. The $\pi-\mu$ separation between 15 and 35 GeV/c was measured at various momenta and the μ suppression factor was estimated to be 1% integrated in the range of interst. In a first part of the test the same mirror as in 2007 was used, while in a second part a

[41] For example Fryka Kältetechnik DLK402 between -10 and +40 C at a maximum power of 380 W.

new mirror was installed; both mirrors have been manufactured by the same supplier, but the former was aluminized and coated with MgF_2 by the supplier while the latter was aluminized and coated with SiO_2 and HfO_2 at CERN in an attempt to improve the total mirror reflectivity. No major difference between the two mirrors was observed in the data.

The results of the 2009 prototype test beam have been published in [86].

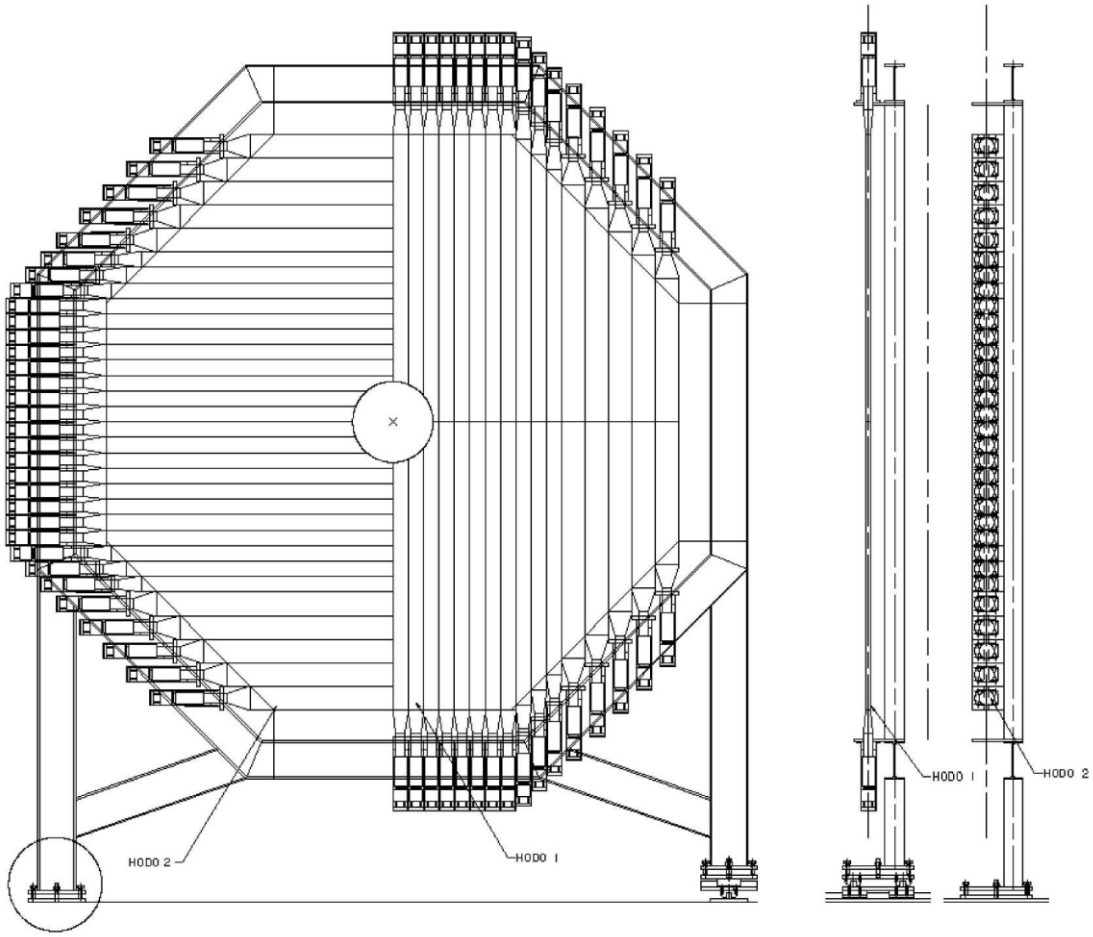

Figure 302 Schematic view of the CHOD detector

3.7 The CHOD

The existing NA48 charged hodoscope is a system of scintillation counters with high granularity and excellent time resolution (200ps) [15]. It will be re-used to detect possible photo-nuclear reaction in the RICH mirror plane and to back-up the RICH in the L0 trigger for charged tracks.

The detector consists of 128 detection channels which are arranged in two planes of 64 horizontal and vertical scintillators. Each plane is divided in four quadrants with 16 counters (see Figure 302), so that the acceptance in the X-Y plane covers a radius of 121 cm. The scintillator dimensions are summarized in Table 45.

The counters are made with BC408 plastic scintillators which have fast light output and good attenuation properties. The scintillation light from each counter is collected via a short fishtail, (made of Plexiglas) light guide, followed by a Photonis XP2262B photomultiplier. Table 45

Table 45 CHOD Scintillator Dimensions

Lengths:	Between 121 and 60 cm
Width:	6 cm
Thickness:	2 cm

The RICH mirror system amounts to about 20% of radiation length and photons from π^0 decays can convert or, even more critical, undergo photonuclear interactions producing low energy hadrons. The Liquid Krypton Calorimeter (LKR) -as subsequent photon detector- has to veto these photons with an inefficiency that is better than 10^{-5}. MC simulations show that photons - which experience photo-nuclear reactions in the RICH - can weaken the photon-veto function of the LKR. In order to re-establish the veto sensitivity to the required level a detector for low momentum charged track after the RICH is needed. This function can be fulfilled by the CHOD.

Another motivation for keeping the present CHOD detector is its timing capabilities, which can be useful in complementing the RICH detector in the L0 trigger selecting charged tracks. The time resolution of an individual plane remains, however, limited by the size and the age of the counters to a level between 3 and 5 ns. This resolution can be improved (drastically) if the two planes are used simultaneously to correct the timing with respect to the crossing point of the track.

The frontend and readout electronics of the CHOD detector have to be entirely re-done in order to cope with the particle flux rate in the CHOD, which is estimated to be around 11 MHz. At this stage several options are under and the choice will depend on the facilities needed for the trigger.

3.8 The Muon Veto Detector (MUV)

3.8.1 Introduction

3.8.1.1 Physics Requirements and General Layout

In addition to the straw chambers and the RICH detector, further muon reduction of the order of 10^{-5} with respect to pions is required and has to be fulfilled by the calorimetric and muon veto systems. The major part of the rejection is achieved by just requiring charged particles not to deposit significant energy in the calorimeters and to traverse a sufficiently thick layer of iron. However, in order to obtain the necessary total rejection power, muons that undergo catastrophic bremsstrahlung or direct pair production and deposit a major fraction of their energy in the calorimeters also have to be suppressed. To reject these rare events, electromagnetic muon showers must be distinguished from hadronic pion showers by measurements of the shower shape, therefore requiring a sufficient segmentation of the calorimetric system.

In order to suppress muon events already at the first trigger level by a factor of about 20, a fast muon veto detector is needed. This sub-detector should have a time resolution of less than 1 ns to reject events with coincident signals in the Giga Tracker and the CEDAR.

3.8.1.2 General Layout

The muon veto system (MUV) consists of three distinct parts, called MUV1, MUV2, and MUV3 in the following, according to their longitudinal position along the beam axis.

The first two modules, MUV1 and MUV2, follow directly the LKR calorimeter and work as hadronic calorimeters for the measurement of deposited energies and shower shapes of incident particles. While MUV2 is the front module of the former NA48 hadron calorimeter (HAC), but turned by 180° (as later explained), the MUV1 module is a newly constructed detector. Both modules are classic iron-scintillator sandwich calorimeters with 24 (MUV1) and 22 (MUV2) layers of scintillator strips. In both modules, the scintillator strips are alternatively oriented in the horizontal and vertical directions. In the MUV1 module, light is collected by wavelength shifting (WLS) fibers, while the MUV2 module routes the scintillator light by light guides directly to photo multiplier tubes (PMTs). The possibility of reusing the complete NA48 HAC (both front and back module) was investigated but rejected for several reasons:

- Required refurbishment of the NA48 HAC back module: The NA48 HAC was originally the NA31 hadron calorimeter and has been reused for the NA48 experiment. In contrast to the front module, the back module had not been equipped with new scintillators and PMTs in the transition from NA31 to NA48, leading to almost blind back module scintillators at the end of NA48 running period. From the experience of the refurbishment of the front module for NA48, the effort required for exchanging scintillators of the HAC back module was estimated to be similar to the effort of building a new module.

- Transversal segmentation: A newly constructed front module allows a finer segmentation in transversal direction to obtain a better distinction between hadronic pion and electromagnetic muon showers.

- Space requirements: Due to the need of a beam deflecting magnet before the end of the beam-line, the longitudinal space available for the HAC modules is less than in NA48. Since light guides and PMTs of the NA48 HAC back module need about 80 cm of longitudinal space in addition to the iron/scintillator layers, a new light collection system had to be built in any case.

The restricted longitudinal space also led to the decision to turn the NA48 HAC front module (MUV2) by 180°. After this rotation, the MUV2 light guides and PMTs point up-stream, surrounding the MUV1 module, leaving the MUV2 downstream space completely free.

After MUV1 and MUV2 and an 80 cm thick iron wall, the MUV3 module, or Fast Muon Veto, has the aim of detecting non-showering muons and acts as muon veto detector at trigger level. To achieve the required time resolution of less than 1 ns at each transversal position, a design is chosen, which employs scintillator tiles arranged to minimize differences in the light path trajectories.

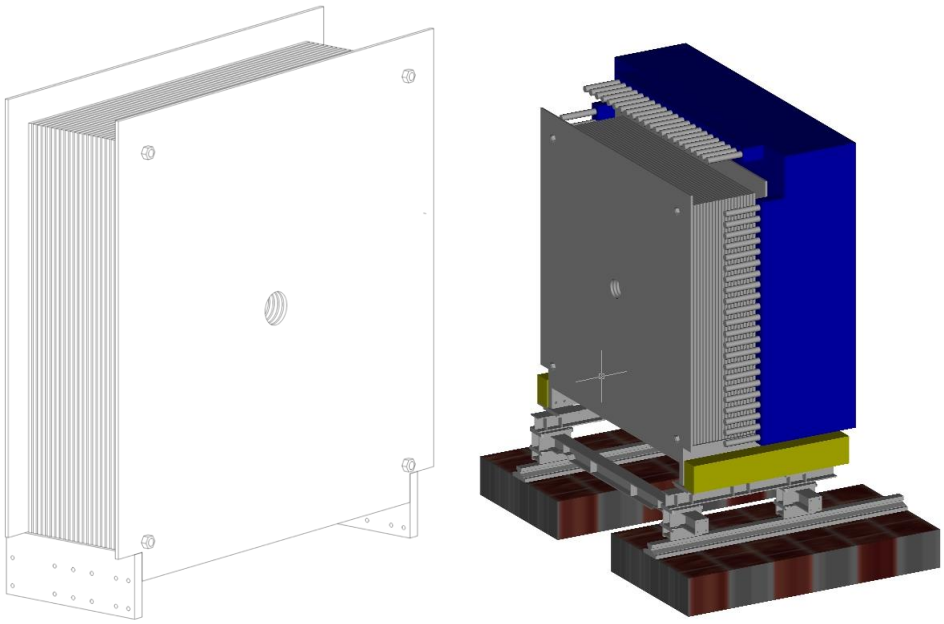

Figure 303 Right: Three-dimensional view of the MUV1 module. Left: View of MUV1 (grey) and MUV2 (blue). The beam is coming from the left.

The number of detection channels is summarized in Table 46.

Table 46 Number of read-out channels of the MUV detector.

Module	**Number of Channels**
MUV1	176
MUV2	88
MUV3 Design A (B)	296 (252)
Total	**560 (516)**

3.8.2 Mechanics and Support Structures

3.8.2.1 MUV1

The MUV1 module consists of 25 layers of SE35 steel. The inner 23 layers have dimensions of 2700 x 2600 x 25 mm^3, while the first and the last layer have the same thickness, but outer dimensions of 3200 x 3200 mm^2. These larger layers serve as support for the whole structure and for the WLS fibers, the photo detectors, and the read-out (Figure 303). The whole iron layer structure is held together by 5 cm diameter steel rods in each corner of the module, maintaining a spacing of 12 mm between the plates. In this way, no welding is necessary, and the MUV1 is constructed by simply stacking alternating iron and scintillator layers onto each other.

Each iron plate contains a central hole of 212 mm diameter for the beam pipe. For additional stabilization during movements and tilts of the MUV1 module, a steel tube of the same diameter can be inserted and fixed inside the central hole.

3.8.2.2 MUV2

The MUV2 module is the old NA48 HAC front module, but turned by 180 degrees, to allow servicing access to PMTs. The welded iron structure consists of 23 steel layers of 2600 x 2600 x 25 mm^3 dimension with 12 mm space between two consecutive iron layers, where a scintillator layer is housed. Each iron plate contains a central hole of 212 mm diameter (Figure 303). The scintillator in the NA48 front module was replaced in 1993/4 when the NA31 (predecessor experiment to NA48) front module was refurbished for use in NA48.

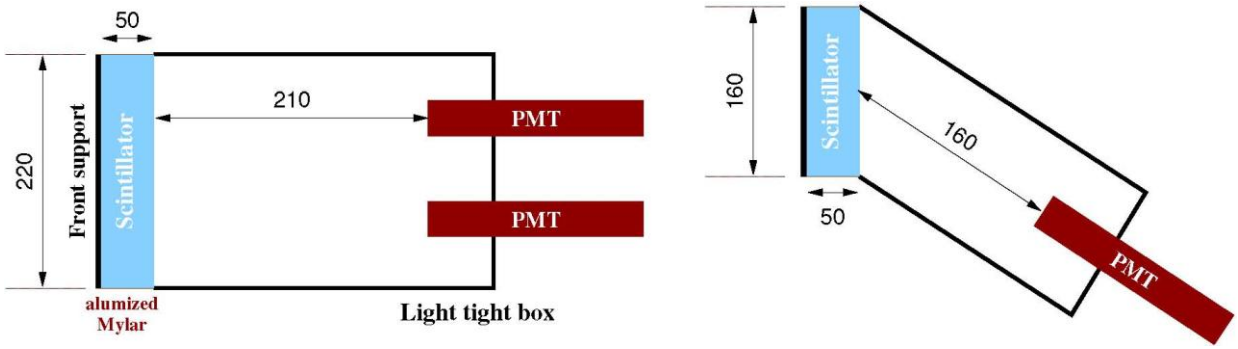

Figure 304 Layout of a single MUV3 tile counter. Left: Drawing A with two PMTs per tile. Right: Drawing B with one tilted PMT per tile. The beam is coming from the left.

3.8.2.3 MUV3

The MUV3 module is located behind an 80 cm thick iron wall filter and serves as a fast muon veto in the lowest trigger level (L0). As a good time resolution is essential, no WLS fibers are used, but instead there are direct optical connections between scintillators and PMTs. The MUV3 will consist of an array of 5 cm thick scintillator tiles. The light produced by traversing charged particles is collected by PMT's positioned about 20 cm down-stream. Due to this geometry, the maximum time jitter between photons from particles hitting different parts of the scintillator tiles is less than 250 ps. However, the

time resolution may be spoiled by Cherenkov photons that are produced by particles traversing the PMT windows. These Cherenkov photons arrive earlier than photons produced in the scintillators, whose typical decay time is about 2 ns, thus shifting the measured arrival time by about –2 ns. To overcome this problem, two different design options are studied and under test (Figure 304). Option A consists of 12 x 12 scintillator tiles of 220 x 220 x 50 mm^3 size, with each tile being read out by two PMT's (Figure 305). The output time of the coincidence of the two PMT signals, corresponds to the time defined by the PMT which is unaffected by the Cherenkov photons. Option B places the PMT behind a neighboring tile, so that possible Cherenkov signals are in coincidence with a scintillator signal from the neighboring tile. In this option only one PMT per tile is needed. Therefore, with the same number of read-out channels and the same MUV3 front face, the tile cross-section may be reduced to 160 x 160 mm^2, arranged in a grid of 16 x 16 tiles. Both design options have been proven, under test, to have sufficient light yield as well as time resolution. The final design decision will be taken after analysis of test beam data, currently being performed.

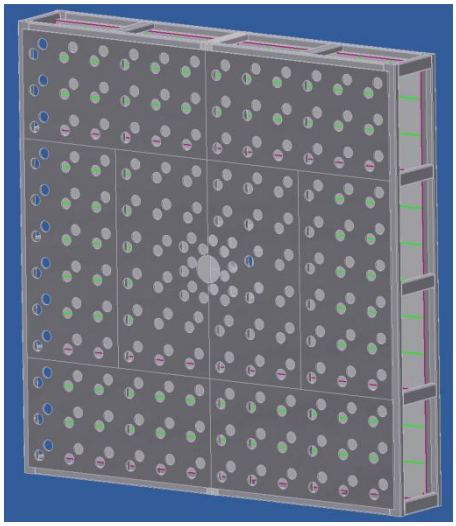

Figure 305 Layout of the complete MUV3 module for option A (without PMT's).

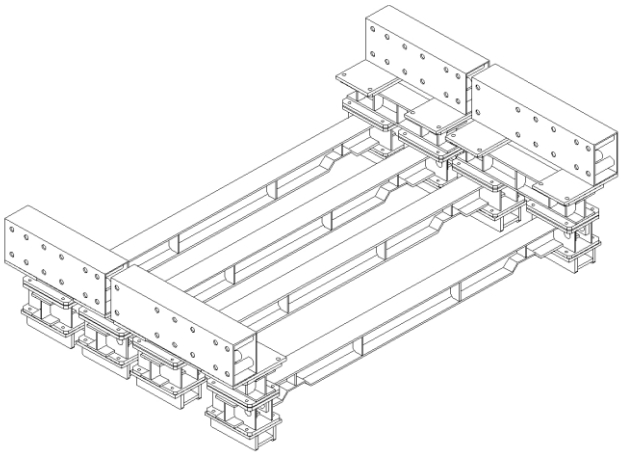

Figure 306 Supporting chariots of MUV1 and MUV2. Most parts of the structure are reused from the NA48 HAC support.

3.8.2.4 Support Structure

The whole muon veto system sits on a pair of rails running in the longitudinal direction. In contrast to the NA48 HAC, where both modules shared a common support structure, each module of the muon veto system (MUV1, MUV2, and MUV3 together with the iron wall) uses its own supporting chariot (see Figure 306 for MUV1 & MUV2). The chariots can be moved independently from each other by the use of winches.

3.8.3 Scintillators and Light Transport

3.8.3.1 MUV1 Scintillators

The MUV1 module houses 2 x 12 layers of scintillators, alternatively oriented in the horizontal and vertical directions. Except for the strips close to the central beam hole and the very outer strips, the size of the scintillator strips is 2616 x 60 x 10 mm^3. They thus cover the whole width of the MUV1 module, allowing a light read-out on both sides. Each scintillator layer consists in total of 48 scintillator strips (Figure 307 left): 34 strips of 2616 mm length, 6 somewhat shorter (with 2496, 2376, and 2256 mm length, respectively) to accommodate the support rods in the corners, and 8 strips of about half length and 54 mm width around the beam hole. The total number of scintillator strips in all 24 layers is 1152.

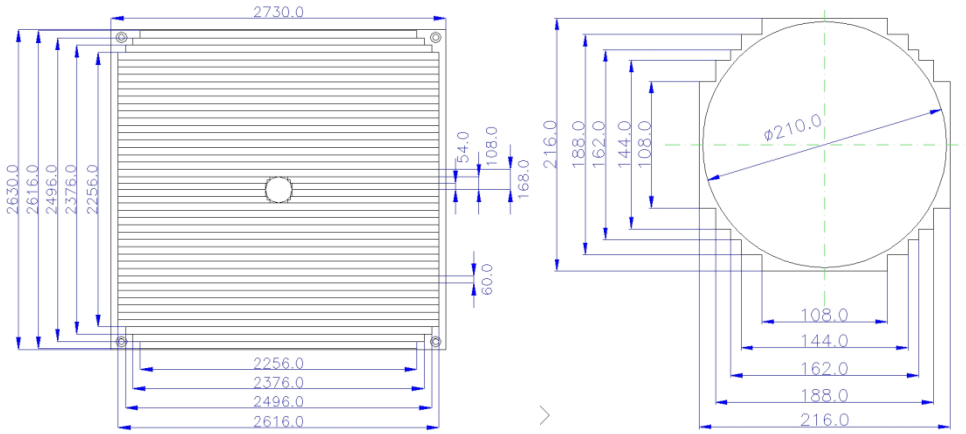

Figure 307 (Left) Layout and sizes of one (horizontal) scintillator layer of the MUV1 module. (Right) Drawing of the spacers in the MUV1 beam pipe region.

In the beam pipe region, the scintillator strips are terminated by a steel spacer ring around the beam pipe. To achieve the maximum acceptance in the region close to the beam pipe, for each orientation the four scintillator strips closer to the beam pipe are cut in two parts and read out at only one end. The ends of four of the eight scintillator strips so obtained are shaped in steps to match the circular shape of the beam pipe as shown in Figure 307 right.

The strip width of 6 cm is a compromise between the need of high granularity and the affordable number of PMTs and read-out channels. Monte Carlo studies showed that a smaller strip width of 4 cm would increase the muon rejection only at a percent level.

The MUV1 scintillators are produced at IHEP in Protvino. They are made of polystyrene (Styron 143E) as carrier substrate with 2% scintillating fluors (p-terphenyl) and 0.05% POPOP. While p-terphenyl emits scintillation light at about 300 – 400 nm in the ultra-violet, POPOP shifts the wavelength to 380 – 500 nm. The procedure used to fabricate the scintillators was newly developed: the mixture of polystyrene pellets, p-terphenyl, and POPOP is melted under a 10^{-4} bar vacuum at about 250 °C. This procedure allows the fabrication of large numbers of long scintillator strips in a relatively short time. One production cycle of heating, melting and cooling needs about 14 hours.

Compared to commercially available scintillators (e.g. Bicron BC 408), the MUV1 scintillators have a shorter attenuation length (< 1 m). However, because of the read-out by WLS fibers (see below), the attenuation length is not an important issue for the detector performance. It was therefore decided to fabricate all MUV1 scintillators at IHEP.

3.8.3.2 MUV1 Light Collection and Transport

For the read-out of MUV1, wavelength-shifting (WLS) fibers are used. This choice was taken to compensate the short attenuation length of the scintillators and also to comply with the space requirements from the surrounding MUV2 PMTs.

Each scintillating strip is read out by two WLS fibers, placed in grooves at 15 and 45 mm along the 60 mm strip width. Several different fiber types from Bicron and Kuraray (Y-11) have been investigated. The preliminary choice is 1.2 mm diameter, multi-cladded fibers BC-91A from Bicron. The fibers shift the scintillator output light to wavelengths between 470 and 570 nm. They are optically connected to the scintillators with optical cement BC-600. It was decided not to use epoxy glue because of possible aging of the fiber cladding.

All fibers have a length of about 500 cm with small variations (< 10 cm), depending on the longitudinal position of the corresponding scintillator. The 12 x 2 fibers of one longitudinal row of scintillators are bundled to direct the light to one single PMT; therefore no longitudinal segmentation exists. The connection to the PMT is made by a matrix, which holds all 24 fibers within the active PMT area of 26 mm diameter.

3.8.3.3 MUV2 Scintillators and Light Transport

The scintillators of the MUV2 module, which is the NA31/NA48 HAC back/front module, were replaced for the start of the NA48 experiment and for NA62 the module is being reused. The scintillators are of type BC-408 from Bicron. Each scintillator plane, inserted between the iron plates, consists of 44 strips. Each strip spans only half the calorimeter so that each plane is made of two half-planes. The two central strips of each half-plane are also three step shaped at one end to wrap around the central hole for the beam pipe (as in MUV1), so they have a width of 108 mm with a length of 1194 mm at the first step and 1243 mm at the third step. All other strips are 1300 mm long and 119 mm wide. The thickness of each scintillator is 4.5 mm (Figure 308).

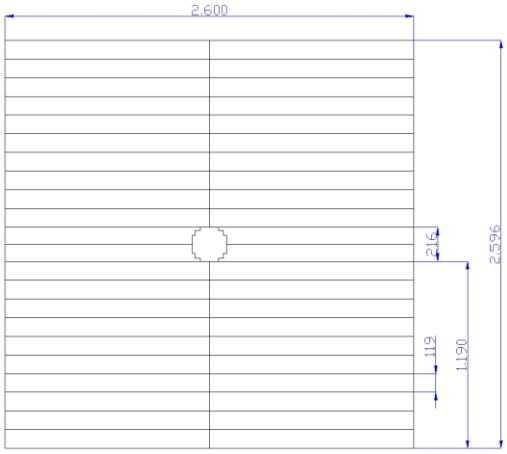

Figure 308 Layout of one (horizontal) scintillator layer of the MUV2 module.

There are in total 24 scintillator planes, so that the MUV2 module – contrary to MUV1 - starts and ends with a scintillator layer. As in the MUV1 module, the strips are alternately aligned in horizontal and vertical directions. Consecutive strips with an identical x/y position and alignment are coupled to the same photomultiplier using *Plexiglas* light-guides "forks".

3.8.3.4 MUV3 Scintillators

The scintillator tiles for the MUV3 module are produced at IHEP in Protvino, using the same mixture of raw materials. However, a different production technique is used: Large blocks of scintillator material are produced by polymerisation with the same mixture of raw materials as for MUV1. The tiles are then cut in the required size from the blocks and polished.

3.8.3.5 Photodetectors

For MUV1, several photomultiplier options have been investigated and tested. The main option is the model Hamamatsu R6095, which has a surface of 5.3 cm^2, enough for the cross section of 24 WLS fibers, and with sufficient quantum efficiency in the green range. The final decision about the MUV1 PMTs will be taken after tests with the full-size scintillator strips.

The MUV2 module reuses the 3'' PMTs (THORN EMI 9265KA) of the NA48 HAC.

For the MUV3 module, the 2'' PMTs (THORN EMI 9814) from the NA48 AKL are reused. Their number is sufficient for about 80% of the MUV3 counters, depending on the chosen design. For the remaining counters, PMTs with similar characteristics will be selected.

3.8.4 Expected Performances

3.8.4.1 Test Results

Scintillator strips of dimension 500 x 40 x 10 mm^3, produced by melting, with Y-11 WLS fiber read-out have been tested with cosmic rays at IHEP and compared to similar strips produced by different techniques. Both the number of photoelectrons (~13 at the fiber far end) and the time resolution (~1.5 ns) did not show significant differences with respect to the other strips.

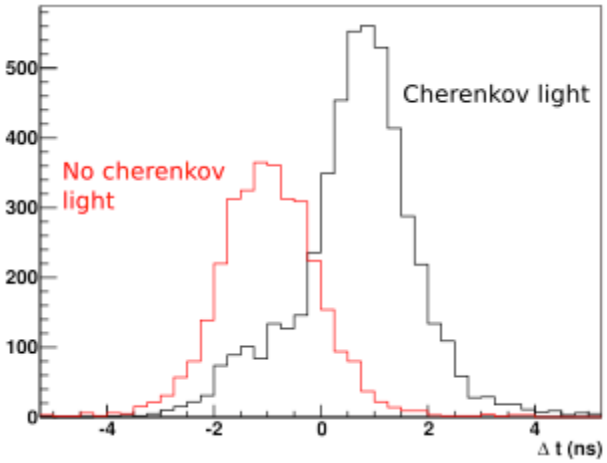

Figure 309 Test beam result for the MUV3 time resolution (layout option A) with the beam passing between the PMTs (red curve) and traversing one of the PMT windows (black).

In addition, several prototype modules of the MUV3 (option A and B) have been tested at a test beam at the CERN PS. The yield of photoelectrons was determined to be 20 – 25 per traversing particle. The time resolution was measured to be 0.5-0.6 ns for particles traversing in between the two PMTs of one module. Most particles passing through one of the PMT windows caused the expected anticipation time shift of about 2 ns (Figure 309). However, when taking into account the low probability of producing Cherenkov photons by a PMT window hit, the overall time resolution of the MUV3 prototype modules was still below 0.6 ns.

3.8.4.2 Simulations

A full GEANT3 simulation of the expected muon suppression of the MUV1 detector together with the LKR calorimeter has been performed. As described in the introduction, the main suppression is achieved by vetoing signals in the MUV3 module and by requiring minimum E/p in the LKR and minimum energy deposition in the MUV1. However, to also reduce background from muon catastrophic energy loss before MUV3, the shower shapes in the LKR and the MUV1 were evaluated. By combining all separation criteria, for momenta 12 < p < 35 GeV/c a total Kμ2 suppression of > 2 x 10^6 at a signal efficiency of 50-60% can be achieved. This figure becomes even better when considering only momenta larger than 25 GeV/c, where the particle ID capability of the RICH detector becomes worse.

4 Detector Read Out and Data Handling

4.1 Trigger and Data Acquisition System

4.1.1 General Overview

The intense flux of a rare decay experiment, such as NA62, necessitates high-performance triggering and data acquisition. These systems must minimize dead time while maximizing data collection reliability. The NA48 trigger and data acquisition systems, designed more than 10 years ago, are unsuited to the task and can no longer be maintained for NA62. A unified trigger and data acquisition (TDAQ) system (1), which, as much as possible, assembles trigger information from readout-ready digitized data, addresses these requirements in a simple cost-effective manner.

The NA62 experiment consist of 12 sub-detector systems and several trigger and control systems, for a total channel count of less than 100 thousand. The GTK has the most channels (54,000), and the Liquid Krypton (LKr) calorimeter shares with it the highest raw data rate (19 GB/s). A summary of the number of channels and typical rates for the primary sub-detectors are shown in Table 47.

Table 47: channel numbers and typical rates of primary sub-detectors.

Sub-detector	Stations	Channels/station	Total channels	Hit rate (MHz)	Raw data rate (GB/s)
CEDAR	1	240	240	50	0.3
GTK	3	18’000	54’000	2’700	2.25
LAV (*)	12	320-512	4’992	11	0.3
CHANTI	1	276	276	2	0.04
STRAW	4	1’792	7’168	240	2.4
RICH	1	1’912	1’912	11	0.09
CHOD	1	128	128	12	0.1
IRC	1	20	20	4.2	0.04
LKr (**)	1	13’248	13’248	40	22
MUV	1	432	432	30	0.6
SAC	1	4	4	2.3	0.02

In the above table: “station” refers to a single unique physical location of electronics (this might include several boards or crates close together); the “Raw data rate” refers to the readout rate from the sub-detector boards to PCs (after a L0 trigger for all sub-detectors except LKr); such rates are quoted as pure payload rates, without any transport overhead being included. (*) Each LAV PMT is read out by two electronics channels; the channel counts in this table are therefore twice the number of PMTs. (**) LKr ADCs are continuously digitized at 40 MHz; the above rate correspond to assuming 8 samples around a trigger are read out at a L1 rate of 100 kHz without zero suppression.

A common coherent clock, with a frequency of approximately 40 MHz[42], generated centrally by a single free-running high-stability oscillator, will be distributed optically to all systems through the Timing, Trigger and Control (TTC) system designed and used for LHC experiments [87]. This "TTC clock" will be the common reference for time measurements[43] (section 4.1.4.1).
TTC optical links will also be used to send to each sub-system:

- a time-synchronous[44] L0 trigger accept pulse (section 4.1.4.2);
- a time-asynchronous L0 trigger type word (section 4.1.4.2);
- a time-synchronous start-of-burst and end-of-burst signals (section 4.1.4.3).

A common time scale is defined by a 32-bit **timestamp** word, with 25 ns LSB and covering the full duration of the interval between two consecutive SPS spills, plus an 8-bit **fine time** word, with 100 ps LSB. While the timestamp will be defined in each system by the phase-coherent distributed clock, each sub-system will locally generate by multiplication a properly locked reference for the fine time.

All clock counters should simultaneously reset at the start of each burst, using an appropriate synchronous command sent to all sub-systems through the TTC link before the actual arrival of the first beam particles. This will also define the origin of the time measurements for the burst. An end-of-burst signal should be sent in the same way some time after the end of the spill, defining the largest timestamp for the current spill. Its value should be recorded by each system and sent to the readout for logging, allowing (online and offline) a consistency check of the number of clock cycles counted by each system during each spill.

For test purposes, each sub-detector readout system should be able to run in a standalone mode, autonomously generating its own TTC signals (including L0 triggers) when not connected to the common system under global experiment control.

In principle, a "triggerless" system under software control, in which sub-detector data are unconditionally readout to PCs, would be the most flexible choice, but the rate and channel count of NA62 make this approach too expensive to implement. NA62 therefore has adopted a hardware lowest-level trigger labelled Level 0 (L0).

[42] Due to TTC system requirements linked to the LHC, the actual clock frequency is not exactly 40 MHz; in this document all references to e.g. "25ns" should be understood as the actual period of the main clock, (close, but not exactly equal, to 25ns).

[43] In a previous proposal for the LKr readout an additional very low-jitter 80 MHz clock, phase-coherent with the TTC clock, was considered to be used as sampling clock for flash ADCs: this could be distributed through the existing NA48 clock distribution system, which remains as an option.

[44] In this context "time-synchronous" denotes a signal occurring in a precisely defined 25ns time-slot with respect to its originating cause.

Following a L0 trigger, most sub-detectors will transfer data for one time-window to dedicated PCs, where a L1 trigger level will be implemented in software, and then to event-builder PCs where a software L2 trigger will be implemented.
The trigger hierarchy is thus made of three *logical* levels:

- a hardware **L0 trigger**, based on the input from a few sub-detectors; after a positive L0 is issued, data is readout from front-end electronics buffers to dedicated PCs (for most sub-detectors);
- a software **L1 trigger**, based on information computed **independently** by each **complete** sub-detector system, using data stored on dedicated PCs;
- a software **L2 trigger**, based on **assembled and (partially) reconstructed** events, in which complex correlations between information from different sub-detectors is possible, using data stored on the event building PC farm.

4.1.2 Trigger Logic and DAQ scheme

4.1.2.1 L0 Hardware Trigger

The hardware L0 trigger will be mainly based on input from the CHOD, the MUV, and the LKr, and optionally the RICH, the LAV, and the STRAW. The default (primary trigger) algorithm will be implemented to collect events with a single track in the CHOD, nothing in the MUV, and no more than one cluster in the LKr. The inclusion of other sub-detector information is possible, both to refine the primary trigger and to implement secondary triggers for control samples and different physics goals: for the main trigger, a multiplicity cut in the RICH and STRAW's may be able to augment the positive CHOD indicator, while the LAV might enhance photon and muon vetoing.

The CHOD will provide positive identification of a charged particle within the detector acceptance, reducing the rate due to *K* decays downstream of the final collimator. Hit multiplicity might also be used to select among multi-track events.

The third ("fast") plane of the MUV (MUV3) will veto muon events, *i.e.* the major background from $K_{\mu 2}$ decays and the muon halo components from decays upstream of the final collimator. This rejection is the single largest rate reduction factor at the trigger level, and, consequently, its efficiency largely determines the L0 trigger rate [88]. The geometrical acceptance of the MUV3 plane must exceed and include that of the positive track-identifying elements in the L0 trigger (e.g. the CHOD), and its online time resolution should be good enough (of order 1 ns), to avoid excessive random vetoing.

The LKr will also be used as veto in the L0 trigger, by allowing, in the primary trigger, no more than a single cluster (compatible with a charged pion EM shower). Simple quadrant energy deposition cuts could sufficiently reduce the $K_{\pi 2}$ rate [88] in the primary trigger line, but online cluster counting with ≈1 ns time resolution may permit a better rejection and a more diverse set of physics triggers.

Hit multiplicity from the RICH might give a further contribution to the reduction of the background rates due to charged particles. Such reduction, however, may be rather limited, as the particle identification capabilities of the RICH cannot be easily exploited without momentum information, which is not available before correlation with the STRAW magnetic spectrometer at L2.

The inclusion of the LAV in the L0 trigger could also contribute to some further reduction of the $K_{\pi 2}$ and muon halo backgrounds. Algorithms differentiating between a minimum ionizing particle and an electromagnetic interacting particle are under consideration and evaluation.

Additional information from other sub-detectors (including sub multiplicities in the STRAW) could increase the flexibility of the trigger, permitting alternative physics triggers. An algorithm requiring the number of MUV hits to be less than the CHOD multiplicity, the total number of clusters (in both LKr and LAV) to be less than the CHOD multiplicity, and no MIPs in the LAV would select $K^+ \to \pi^+ \ell^+ \ell^-$ as well as $K^+ \to \pi^+ \pi \nu\nu$ events.

Trigger primitives from sub-detectors involved in L0 will include both a timestamp and a fine time, in order to allow tight time matching. A L0 Trigger Processor (section 4.1.6) will time-match L0 trigger primitives issued by sub-detector trigger electronics and appropriately generate a trigger signal, which will be dispatched only with timestamps (25 ns time granularity), as it is expected that sub-detectors will readout data corresponding to (programmable) time windows longer than 25 ns.

Data from all sub-detectors will be stored in front-end buffers during L0 trigger evaluation. Upon reception of a positive L0 trigger, most sub-detectors will send their data to dedicated PCs within an adequate time window around a L0 timestamp. At every L0 trigger, each centrally-enabled sub-detector will respond to a L0 trigger by sending a data frame. The type and amount of data it contains may be different for different trigger types; a data frame may even be empty or indicate an error. No sub-detector may ignore a L0 trigger.

A maximum L0 trigger latency requirement implies that if no positive L0-accept signal is received within this period, the data can be discarded by the front-end buffers. A L0 trigger issued after the latency period is an error and should not happen under normal conditions; should it occur, each sub-detector will nevertheless reply with a properly formatted frame (most likely empty or indicating an error).

No data for untriggered events should be sent to PCs: downscaled events for control will be handled centrally by the L0 Trigger Processor, which will accept some events regardless of their failure to satisfy some specific trigger condition (another mechanism by which some sub-detectors might request that some specific event is forcibly collected is by issuing special trigger primitives).

An early and naïve simulation of the detector hit rates and L0 trigger rates when using simple trigger cuts was produced (with an older detector set-up) using the fast FLYO MonteCarlo simulation, and is described in (1).

The maximum L0 trigger rate and latency are discussed in section 4.1.3.2. The possibility of increasing the L0 latency from 1ms to 6 or 7 ms should be foreseen by sub-systems, leaving open the possibility of implementing a smarter and more powerful (albeit possibly slower) L0 trigger scheme (see section 0).

4.1.2.2 L1/L2 Software Triggers

After a positive L0 trigger, all sub-detectors' data (with the exception of the LKr) associated with the L0 trigger timestamp, are moved to PCs for initial processing, which includes, at the very least, quality checks and reconstruction validating, as well as rudimentary pattern recognition.

A L1 trigger will require data quality verification, and then be based on simple correlations between *independently-computed conditions by single sub-detectors*, testing for the presence of a single charged pion in the detector (for the main trigger). A possible algorithm would be an equal odd number of tracks in the front and back of the STRAW spectrometer, no MIP or showering particle through the LAV, a hit multiplicity and pattern consistent with a charged pion in the RICH, and (optionally) a positive indication of at least one in-time kaon in the CEDAR.

Most detectors are expected to provide L1 Trigger primitives, if only in the form of rough data validation. One (single) PC (of perhaps several) associated with each sub-detector will be responsible for dispatching asynchronously the L1 trigger primitives for that sub-detector for *each* L0-triggered event to a central L1 Trigger Processor PC, based on complete sub-detector event data (which may have been merged from independent devices to include the entire sub-detector).

The L1 Trigger Processor will match these primitives and asynchronously issue a L1 decision, at which time the data will be transferred to the event-building farm (in the case of a positive L1) or discarded (in the case of a negative L1 verdict). The L1 decision will be delivered to the master PC of each sub-detector. This PC will dispatch the information to collect event data if they are split between different PCs. In case of positive L1 trigger decision, the information on the farm node to which event fragments have to be dispatched will be also delivered by the L1 Trigger Processor to the (master) sub-detector PC.

All L0-triggered events will get a L1 decision, and no data should be discarded until that decision has been received. The rate of the L1 trigger is not fixed, and there is no strict maximum latency for it, but L1 trigger evaluations are expected to terminate shortly after the end of each spill. No data for untriggered events should be sent to the event-building farm, as downscaled events for control will be handled by the L1 Trigger Processor, which will accept some events regardless of their failure to satisfy some specific trigger condition.

The possibility of running the experiment with more than one L1 Trigger Processor simultaneously active could also be foreseen: exactly one processor will handle one event.

With the data moved to the event-builder farm after a positive L1 trigger, crude full-event reconstruction can be done. Any event with an unaccompanied, single, identified charged pion will be accepted for the main trigger, but particle identification and 4-momentum consistency for other modes could also checked.

A L2 trigger will be based on *correlations between different sub-detectors*. The information upon which these correlations are determined will be provided by event-building PC farms. Most sub-detector activity within an event time window will be at least partially reconstructed in the farm and made available for the L2 trigger decision.

For efficiency reasons the L2 trigger algorithms may be arranged into different hierarchical stages of conditional processing, *e.g.* by delaying time-consuming reconstruction until simpler conditions have been satisfied, or by running multiple reconstruction algorithms with different levels of refinement and complexity as needed. There is no dedicated L2 Trigger Processor, as each event will be dispatched for event-building and L2 triggering to different, dynamically chosen individual nodes in the farm.

All data associated with events satisfying the L2 trigger conditions will be logged to tape. In case L2 trigger conditions are not satisfied for an event, the data will be deleted (a fraction of failed events will be passed for purposes of monitoring and efficiency determination).

The rate of the L2 trigger is not fixed *a priori*, but will be determined by data logging capability. There is no maximum latency: L2 trigger computation can extend into the inter-spill period, but they should terminate before the next spill starts.

4.1.3 Requirements and Specifications

A detailed discussion of design requirements can be found in (1). Among more crucial ones ofr NA62 are reducing to a very low level any undetected partial failure of the readout for vetoing sub-detectors and avoiding any uncontrollable correlations between sub-detector trigger requirements.

Undetected vetoing failure due to data transmission errors can occur through two basic mechanisms:

(a) failure to deliver data from some (part of a) sub-detector going unnoticed, or
(b) time mis-alignment between data sent from one sub-detector and the others.

Once the data is within the processor farm, error checking mechanisms of the networking infrastructure can be exploited to limit the rate of occurrence of such errors, but particular care has to be taken in the first part of the data path, where custom electronics are used and data transfers occur between modules of different types.
Errors of type (a) are partly controlled by requiring that all sub-detectors (and all modules within a sub-detector readout system) always actively respond to a readout request (*i.e.* a L0 trigger), even if they have no data for that particular time region. Periodic DAQ integrity checks must also performed asynchronously in an automatic way.
Errors of type (b) are controlled by continuous burst-level clock alignment checks, and event-by-event timestamp matching checks; the latter are important since data corruption in the timestamp is a single point of failure for vetoing (corruption in other parts of the data might not necessarily result in vetoing failure).

4.1.3.1 Definitions

Unambiguous definitions of some TDAQ terms used throughout the rest of this document are presented here.

Burst or spill: the period of the SPS beam-delivery cycle, it is the basic data-taking time unit (event numbering and timestamping are relative to a burst, and restart from the beginning again at each new burst); the duration[45] is not specified and can usually vary in the range 1-20 s, but is (roughly) constant during each run; each sub-system should be capable of working with any value up to 50s..

Run: an arbitrary but convenient way of grouping a series of bursts taken under uniform data-collection conditions; runs cannot overlap in time.

[45] Includes both the spill time, when beam particles are are hitting the target, and the subsequent inter-spill time, when they are not.

Timestamp: a 32-bit unsigned integer, relative to which all individual channel times are to be interpreted; defined by the L0 Trigger Processor, made available to each sub-system through a TTC receiver, and included in the event structure at all levels of data transport; uniquely related to the event number within a burst; the LSB equals the period of the master clock, roughly 25 ns (1/(40.07897 MHz) ≈ 24.951 ns), and the MSB is reserved, for a time range of about 53.6 s; some lower bits may be ignored when time matching for data extraction.

Fine time: the granularity for detector hit and L0 trigger primitive times; the least significant bit is 1/256 of the main clock period, roughly 100 ps (97.466 ps); not defined for an event, but event fine time will be available within the L0 Trigger Processor data frame.

Event number: the L0 trigger number (together with the burst ID) uniquely identifying a specific event; a 24-bit unsigned integer, corresponding to more than 16 s at the 1 MHz L0 trigger rate; provided by each TTC receiver; a difference in the correspondence between event number and timestamp among different sub-detectors indicates missed triggers (a severe error condition possibly forcing rejection of an entire burst's data). Event building from sub-detector data is based on the event number.

Burst ID: the UNIX time of a conveniently chosen instant in a SPS burst, uniquely identifying it and (together with the timestamp or event number) allows an event; a signed 32-bit number assigned by the farm management system (PC), broadcast to the entire TDAQ farm over the network, and entered into the data stream by sub-detector PCs.

Run number: a 32-bit unsigned integer identifying a continuous data-taking period under roughly homogeneous conditions; a higher-numbered run must contain higher-numbered burst IDs than a lower-numbered one.

Trigger type: a 32-bit unsigned integer indicating the L0 trigger type (the lowest 8 bits, assigned by the L0 Trigger Processor, see Table 55), L1 trigger type (the second 8 bits, assigned by the L1 Trigger Processor), and the L2 trigger type (the third 8, assigned by the Event Builder); the upper 8 bits are reserved (see Table 48).

Table 48: Overall trigger type word format.

Reserved		L2 trigger type		L1 trigger type		L0 trigger type	
31	24	23	16	15	8	7	0

Data block: a series of aligned and possibly padded 32-bit words containing sub-detector specific information, in particular time information with respect to the event time, as defined by the timestamp.

4.1.3.2 Parameters

The values for the most important parameters of the TDAQ system are specified in the following table:

Table 49: Main TDAQ parameters.

Parameter	Value	Description
f(L0) max	1 MHz	Maximum average L0 trigger rate
Δt(L0) min	75 ns	Minimum L0 trigger time separation
T(L0P) max	100 µs	Maximum latency for generation of L0 Trigger primitives
T(L0) max	1 ms (*)	Maximum total L0 trigger latency
f(L1) max	100 kHz	Maximum average L1 trigger rate
T(L1) max	1 s	Maximum total latency for L1 trigger
F(L2) max	O(15 kHz)	Maximum average L2 trigger rate
T(L2) max	Spill period	Maximum total latency for L2 trigger

(*) If possible, designing with a possible upgrade to higher latencies in mind would allow future implementation of different schemes for the L0 trigger generation.

4.1.3.3 Data Format

Each sub-detector and data source is uniquely identified by an 8-bit ID which is used to recognize its data, as indicated in Table 50.

Table 50: data source IDs.

Sub-detector	ID
CEDAR	0x04
GTK	0x08
CHANTI	0x0C
LAV	0x10
STRAW	0x14
CHOD	0x18
RICH	0x1C
IRC	0x20
LKr	0x24
MUV	0x28
SAC	0x2C
L0 Trigger Processor	0x40
L1 Trigger Processor	0x44
L2 event-building node	0x48
Reserved	0x4C
Reserved	0x80-0xFF

Besides detector data, most systems will also produce additional data as trigger primitives, which are identified by the lower bits of the ID, by adding 0x1 for L0 primitives and 0x2 for L1 primitives. For example, the LKr L0 trigger primitives will have ID 0x25, while the LAV L1 trigger primitives (all stations) will have ID 0x12.

The format of the data from each sub-detector is free until the data enter the event-builder farm after a L1 trigger. The format of the data frame (32-bit aligned) will be as follows:

Table 51: Data frame format after L1 trigger.

Word 0	*Reserved*	Total 32-bit word count (N+4)		
Word 1	Data source ID	Event number		
Word 2	Timestamp			
Word 3	*Reserved*			
Word 4 to N+3	Detector data block (N words)			
Bit	31 24	23 16	15 8	7 0

The total word count indicates the size (in 32-bit words) of the entire block (N+4), including headers and trailers; the maximum event data size is 60 MB per sub-system. The detector data block contains sub-detector specific data, 32-bit word aligned and possibly padded.

4.1.3.4 Readout Electronics

Front-end electronics (before digitization) are detector-specific and described in individual sub-detector chapters. Here we summarize general information about readout electronics (after digitization).

Sub-systems will receive a continuous 40-MHz clock, L0 trigger information (L0 accept trigger pulse and trigger type word), and start- and end-of-burst signals through the TTC, driven by the L0 Trigger Processor. These will be dispatched to individual sub-detectors by a TTC transmitter module (section 4.1.4), in principle allowing standalone running (for test or debugging). In practice, some smaller sub-detectors may be grouped to receive clock and L0 triggers from the same TTC module: in this case, they must function together (*e.g.* they will always be both included or excluded from a run).

Each sub-system is responsible for counting the number of clock pulses received between start- and end-of-burst signals, and for transmitting this count (as a data frame) upon request from the L0 Trigger Processor. Each sub-system must respond to every L0 trigger dispatch with a properly formatted data frame. Data can be transferred to respective sub-detector PCs only in response to appropriate trigger signals. Each sub-system must collect (at the end-of-burst, upon request from the L0 Trigger Processor) some status and monitoring information for monitoring and consistency checks, and send it in a special data frame. Finally, each sub-system must be able to drive a pair of lines to indicate that it is overloaded by data or in an error state (section 4.1.4.4): all such lines are merged at each sub-detector level and provided as inputs (one per sub-detector) to the L0 Trigger Processor.

4.1.3.5 L0 Trigger Sub-Detectors

Sub-detectors contributing to the L0 trigger will continuously evaluate their incoming data for the fulfilment of certain conditions (called “primitives” in TDAQ) and associated times. Times include both timestamp and fine time with typically a few nanoseconds intrinsic resolution, so that a number of lower bits of the fine time can be set as desired. Relative time offsets between sub-detectors should be corrected online so that the times of all trigger primitives are consistently aligned. Because each sub-detector has at most one connection to the L0 Trigger Processor (the Processor matches different sub-detector primitives, not primitives from the same sub-detector), data evaluation from multiple front-end cards will be centralized. The trigger primitive data will be packed into 8 bytes for every occurrence of some trigger primitive being satisfied, formatted as follows:

Table 52: L0 trigger primitive message format.

Word 0	Sub-detector ID	*Reserved* (0x0)	Primitive ID	Fine time
Word 1	Timestamp			
Bit	31 24	23 16	15 8	7 0

Inclusion of the sub-detector ID in the data packet offers the possibility of merging in a single link trigger primitives from multiple low-rate sub-detectors.

Each sub-detector can priority-encode primitives, so that when more than one is satisfied at a given time only a single piece of information is sent, but in this case evidence should remain within the primitive ID that such situation occurred[46] (this is required because the sub-detector trigger data for an event should indicate unambiguously ALL trigger primitives which were satisfied at that time). For primitives satisfied over an extended interval of time (of course shorter than the typical detector response time) a representative time should be determined.

Primitives will be sent asynchronously to the L0 Trigger Processor over Gigabit Ethernet (GbE) links; at the expected particle rate of 10 MHz, the maximum single-detector-to-L0 Trigger Processor bandwidth is 80 MB/s, which should be accommodated in a single GbE link. Primitive data will be packed in arbitrarily long lists, which need not be time-ordered, nor is there a fixed minimum time delay for transmission. Trigger primitives may therefore be packed to optimize bandwidth. A watchdog system might be required, however, to guarantee that under low-rate conditions primitives are not sent too late: primitives associated with a given time must reach the L0 Trigger Processor before the maximum L0 latency period T(L0P)max elapses, after which the L0 Trigger Processor is free to make a final L0 trigger decision based on the primitives received so far. Reception of a primitive referring to a time older than T(L0P)max is an error condition.

4.1.3.6 Higher-Level Triggers Sub-Detectors

All sub-detectors should have the capability to check and communicate the consistency of their data with decay mode- and sub-detector-specific trigger requirements, *i.e.* generating L1 primitives. All detectors involved in any L1 trigger decision must send L1 primitives to the L1 Trigger Processor in response to *every* L0 trigger, in the form of a 12-byte data packet shown below:

Table 53: L1 trigger primitive data format.

Word 0	Detector ID	Event number
Word 1	Timestamp	
Word 2	Primitive data	
Bit	31 24	23 16 15 8 7 0

[46] The simplest implementation of this is to associate one bit with a given primitive, but this would restrict the number of primitives to 8; since not all primitives will be independent, better encoding schemes can be used that allow the use of all 256 combinations, while retaining the information on condition overlaps.

Transmission of L1 primitive data can occur any time during the burst, without regard to event ordering. The L1 Trigger Processor will decide whether an event satisfies a L1 trigger only after corresponding L1 trigger primitives from ALL participating sub-detector L1 PCs have been received.

By the end-of-burst (plus some time to finish processing) the L1 Trigger Processor will have evaluated all L0-triggered events in the burst and dispatched L1 trigger messages to all sub-detector PCs, with the L0 and L1 fields of the trigger type word filled. A non-zero L1 trigger type field means the data for this event must be delivered to the event builder indicated in the packet (by its IP number). If the L1 trigger type field is empty, the last word of the packet is missing, and the data for this event should be purged. The format of the L1 Trigger Processor message is shown below:

Table 54: L1 Trigger Processor message format.

Word 0	*Reserved*															Event number															
Word 1	Timestamp																														
Word 2	Trigger type word																														
Word 3	Event builder ID (optional)																														
Bit	31							24	23							16	15							8	7						0

No sub-detector may override L1 Trigger Processor decisions with regard to data handling. Keeping a (standard) fraction of events regardless of trigger decision will be a capability built into, and administered by, the L1 Trigger Processor.

4.1.4 Common Infrastructure

4.1.4.1 Clock Distribution

A single, free-running clock generator will deliver a continuous, high-precision, high-stability ≈40 MHz experiment clock to all sub-detectors[47] through the TTC system [87]. This clock will drive all of NA62 timing systems and will run uninterrupted even when data-taking is not taking place. The TTC system [87] can encode on the same single-mode optical fibre via two time-domain multiplexed channels the clock, synchronous triggers pulses, and asynchronous commands. L0 trigger (from the L0 Trigger Processor) and start- and end-of-burst information (from SPS timing signals or faked by pulsers) will be distributed over the same link.

Twelve master clock partitions (each one dedicated to one or more sub-detectors) are foreseen, at present allocated as follows (the owning group is listed in parenthesis):

[47] If needed, the central clock generator might also drive the old NA48 clock system, to ensure that it maintains a stable relative phase relation with the experiment.

1. CEDAR (Birmingham)
2. GTK (Ferrara)
3. CHANTI (Napoli)
4. LAV (Roma/Frascati)
5. STRAW (CERN)
6. RICH (Perugia)
7. CHOD (Mainz/INR/IHEP)
8. IRC/SAC (Sofia)
9. LKr (CERN)
10. LKr/L0 (Roma Tor Vergata)
11. MUV (Mainz/IHEP/INR)
12. Spare

More than 12 sub-systems can be accommodated, but would share the clock/trigger distribution system with one of the above.

The master clock generator will distribute the 40 MHz clock signal to a fan-out card, which will drive in parallel 12 identical clock/trigger sub-systems housed in two master VME crates. Each clock/trigger sub-system, belonging to the corresponding sub-detector, comprises a modified version of the LTU module [89] designed for the ALICE experiment, and a TTCex [90] module with up to 10 identical optical outputs. Since each clock destination requires an individual fibre, passive optical splitters can be used to serve more destinations (up to 320 per TTCex module with a 1:32 splitter).

The clock from the main clock generator will be fanned out to a set of TTCex modules [90]. Each sub-detector will receive up to 10 identical optical fibres from its TTCex module, and optical splitters can be used to reach more destinations (up to 320 per TTCex module).

Every electronics card requiring reference to the experiment time will be equipped with a TTC optical receiver, a TTCrx chip [91] which will extract the information from the optical signal and provide L0 trigger timestamps, and optionally a QPLL system [92] to reduce clock jitter. These might be integrated into each card, or, equivalently, the CERN-built TTCrq mezzanine board [93] can be used. Each sub-system interfaced to the TTC system should include a 32-bit timestamp counter, counting the number of clock cycles between start- and end-of-burst commands (also distributed by TTC).

4.1.4.2 L0 Trigger Distribution

L0 triggers will be synchronous 25ns pulses[48], and L0 trigger time information will be intrinsically encoded in the pulse occurrence time. The corresponding sub-detector timestamp will be in part provided by the TTCrx chip [91], as the start-of-burst signal will synchronously reset a local timestamp

[48] In TTC jargon, the time-synchronous trigger signal transmitted on “channel A” is called “L1 accept” (L1A); this signal will be used in NA62 to transmit the L0 trigger; the lowest trigger level is called Level 0 in NA62 as it will be the only one performed in hardware, to clearly set it apart from higher trigger levels, L1 and L2, which are performed in software.

counters. Differences in fibre length will cause fixed sub-detector offsets; the TTCrx chip can partially compensate for up to 80 m of fibre length difference (16 timestamp counts), but further adjustments may be necessary in individual sub-detector electronics.

The timestamp generated by the TTCrx for each L0 trigger is a 12-bit word, with 25 ns LSB, and therefore rolls over after 102.4 μs. Coarse time information will be provided locally with a counter incremented by a suitably divided clock, derived from the master clock, and including an offset matching the timestamp offset; appending this coarse time to the TTCrx provided time will yield the complete timestamp.

Sub-detectors issue L0 trigger primitives asynchronously. The L0 Trigger Processor will re-synchronize them before driving the TTC transmitters, and L0 triggers are therefore dispatched in proper time order as synchronous pulses. Consecutive valid L0 triggers will be separated by a minimum of 3 timestamp counts (75 ns). The L0 Trigger Processor will synchronously dispatch trigger and burst information to multiple (NA62 version) LTU modules [89]. Each such module will drive one TTCex optical transmitter module passing the information received from the L0TP, but it will also be able to internally generate triggers in case a sub-detector is running in standalone test mode.

Each central crate will require a VME processor for control and communication. Each processor will run as many DIM server daemons as the number of LTUs in the crate, so each LTU will be controlled by the daemon assigned to it. All daemons will be identical except for the identifier of the LTU crate they serve and their DIM identifier. All the servers in both crates will be approached independently. Each sub-detector group will have access to the processors, in order to be able to control its LTU module at any time, during the “development phase” of the experiment; during the “run phase”, the control of all LTUs will be centralized and handled by a common NA62 run control program. The VME processor configuration used for this “production system” should be considered the “NA62 standard configuration”, supported by the online group.

L0 trigger type information will also be encoded in 6 bits and transmitted through the TTC as an asynchronous command after each trigger pulse (“short B channel broadcast message” in TTC jargon): the two lower bits of the 8-bit word are reserved for start- and end-of-burst encoding (section 4.1.4.3). Five bits are available to encode 32 different types of physics and calibration triggers. The sixth bit is reserved for special commands for certain sub-systems to perform specific tasks and respond with appropriate data frames (remember that all sub-detectors must respond to all L0 triggers, if only with an empty data frame if they cannot handle the indicated trigger). The coding of the L0 trigger type word is shown below:

Table 55: L0 trigger type word encoding.

L0 trigger type	Trigger	Sub-detector action
0b0xxxxx	Physics trigger	Readout data (*)
0b100000	Synchronization	Send special frame
0b100001	*Reserved*	
0b100010	Start of burst	Enable data-taking, send special frame
0b100011	End of burst	Disable data-taking, readout end of burst data
0b100100	Choke on	Send special frame
0b100101	Choke off	Send special frame
0b100110	Error on	Send special frame
0b101111	Error off	Send special frame
0b101000	Monitoring	Readout monitoring data
0b101001	*Reserved*	
0b10101x	*Reserved*	
0b10110x	Random	Readout data (*)
0b10111x	*Reserved*	
0b11xxxx	Calibration	Readout data (*)

(*) The type and amount of data to be read out can be different for different trigger types.

An example of such a special command is the **monitoring trigger**, which requests sub-detectors to send data frames containing monitoring information for inclusion in the data stream.

Start- and **end-of-burst triggers** define the valid data-taking time interval (and must occur after and the start- and end-of-burst hardware signals, respectively). These are the first and last triggers of each burst, and sub-detectors should respond to them by sending data frames as usual: the one for the end-of-burst trigger should contain monitoring data and statistics for the burst, which must include the timestamp count when the hardware end-of-burst signal was received, which will be checked for consistency offline, possibly taking into account relevant offsets).

Other special trigger codes might be defined, *e.g.* to distinguish between physics and calibration data-taking intervals.

Another special command is a **synchronization trigger**, in response to which all sub-detectors send a "sync frame", formatted as a normal event. TDAQ will make use of such frames to monitor the live status of the entire chain through to the offline level, checking that all sub-systems were functional at least immediately before and after the trigger. The average frequency of these triggers will be chosen to allow adequate monitoring of TDAQ, but they will be issued aperiodically so as to avoid masking malfunctions linked to particular timestamp-related bit patterns.

The TTCrx chip will decode the trigger type word and distribute the information, initiating trigger-type-specific responses from front-end systems (*e.g.,* a sub-detector might reduce, or zero-suppress, data for some kinds of triggers but not for others, or issue an empty frame for calibration triggers related to a different sub-detector). A priority scheme will be programmed into the L0 Trigger Processor, so that

distinct, but simultaneously occurring, triggers (*e.g.* a physics trigger and a calibration trigger) can be handled properly.

Higher-level (L1/L2) triggers will define additional trigger types, requiring distinct processing algorithms, and they might also use the L0 trigger type word, *e.g.* for steering different processing algorithms.

4.1.4.3 SPS Interface

In NA62 the distribution of timing and clock will be done with the TTC system. WWE and a delayed EE will define the useful burst interval and will be distributed to all readout systems. The source of the clock and timing distribution, the central TTC crate, will be in the experimental area, and the SPS signals will be brought to this crate. Allowance must be made for user-defined timing sequences, and for part of the existing NA48 infrastructure which can be re-used after refurbishing the NIM logic[49].

The TTC-defined "bunch counter reset" (BCRST) and "event counter reset" (ECRST) signals (the two lowest bits of the short broadcast message) will encode start-of-burst and end-of-burst commands as follows:

Table 56: Start- and end-of-burst TTC encoding.

Command	BCRST	ECRST
Start-of-burst	1	1
End-of-burst	1	0
Reserved	0	1

The start-of-burst command will be dispatched by the L0 Trigger Processor in response to a WWE signal from the SPS, before physics data-taking begins. The end-of-burst command will be dispatched by the L0 Trigger Processor in response to a delayed EE signal, after (possibly quite some time after) the beam extraction has finished. Start- and end-of-burst signals should always occur sequentially and

[49] NA62 re-uses partially the timing system of NA48, briefly reviewed here. The main inputs from the SPS timing distribution -- WWE (Warning of Warning of Ejection, about 1 s before the start-of-burst), WE (Warning of Ejection, just before the start-of-burst), and EE (End of Ejection, just after the end-of-burst) -- were regenerated by NIM logic as custom signals and fanned-out over copper wire by pseudo-differential drivers (CERN design) and audio-video pairs to four separate physical destinations in the ECN3 experimental area: the upstream region of the beam line (the tagger/KABES area), upstream of the detector close to the blue tube (the third drift chamber area), the electronics barrack, and the technical gallery. Each destination was equipped with a receiver module, a fan-out to which users connected, and a transmitter module that issued return signals for monitoring. Custom modules generated delayed copies (referred to as EC and ET) of the EE signal, and the down-counter of a four-fold scaler module generated a CLOCK RESET signal. All these signals were primarily employed to synchronize the readout and single-board VME processors, to create a burst gate for the hardware fast trigger logic, and to generate an in-burst interrupt to the slow control system to read the current in the drift chambers. The NIM logic could also fake a SPS sequence, useful for debugging the readout when the SPS was off.

be present, even in case of failure of the SPS timing signals, so that timestamp counters do not roll over their boundaries. An out-of-sequence command will cause the burst to be lost, and should be reported as a serious error. Each sub-system is responsible for counting the number of clock cycles received between a pair of start- and end-of-burst signals.

One arrangement which was found to be useful in NA48 was to extend the active data-taking time slightly more than the physical burst duration, and to use the additional time at the end (when beam is not present) to have calibration triggers; this required an additional timing signal to mark the end of the physical burst as opposed to the end of the data-taking burst. A similar scheme can be implemented in NA62 by providing some additional timing signal(s) to the L0 Trigger Processor only: sub-detectors need not and should not receive or depend on any other timing signals beyond the start- and end-of-burst ones discussed above.

The information on burst starts and ends might be required by at least a fraction of the experiment's PCs which perform different actions during the spill and inter-spill periods. The central system should take care of distributing such signals to these. Since this is not a time-critical task, such distribution is foreseen to be performed over the network, using the DIM software [94]: one of the VME processors controlling the master clock crates will act as a DIM server for this purpose.

4.1.4.4 Data Flow Choking and Errors

Each sub-system can drive two control lines, each a low-voltage differential signalling (LVDS) pair, transmitted on a single cable (per sub-detector[50]) with RJ45 connectors, informing the L0 Trigger Processor of possible trigger handling problems. Such signals indicate error conditions and *should not be used for data flow control*[51].

The **"choke"** line will be driven high to indicate that a sub-detector is overloaded with data and approaching a situation in which it will not be able to accept further triggers without losing data. The L0 Trigger Processor will respond to choke signals and suspend L0 trigger dispatching as soon as possible. However, all sub-detectors, including the one driving the choke line, are still assumed to handle correctly all delivered triggers even when they have asserted the choke line, so that no data or transmitted L0 triggers are lost. An estimate of the response time of the L0 Trigger Processor to a choke or error conditions cannot be given until the implementation of such a device is known; however, since it is expected that the stopping of triggers will happen at the end of the synchronization board (see section 4.1.6), the response time can be assumed to be no longer than the physical time for transmitting the signals to the L0TP (1 μs might be a safe figure).

[50] Merging signals from *s*ub-cards is the responsibility of individual sub-detectors.

[51] This is the reason for avoiding the name "busy" or "XOFF" for the "choke" line.

If L0 trigger requests for some reason cannot be serviced, or if a L0 trigger timestamp exceeds the L0 trigger latency period and the data has been already purged, or for any other situation which could result in undetected loss of data, sub-detectors should drive high the **"error"** control line[52].

The sub-detector driving the choke and/or error signal should keep it asserted at least until an acknowledgement is received from the L0 Trigger Processor in the form of a corresponding special trigger ("choke on" or "error on"), even if the situation leading to its assertion disappeared in the meantime.

The entire burst may be marked as unusable, but the L0 Trigger Processor will be capable of masking (i.e. ignoring) choke/error lines from individual sub-detectors.

In normal data-taking conditions neither the choke nor the error line should ever be asserted by any sub-detector. The size of the sub-detector buffers should be dimensioned in a way to sustain the maximum average rate, *including the rate fluctuations*.

4.1.5 Common TDC System

An early effort to find common (and possibly existing) solutions to common problems led to the identification of the TELL1 generic readout board [95], developed by EPFL Lausanne for LHCb, as the possible backbone for several applications, thus exploiting an existing product and reducing the amount of new hardware to be developed and maintained. The somewhat dated design was reviewed, and an improved version of the board was designed for NA62, with more powerful computing elements and much enlarged memory buffers.

The new readout and trigger board, called "TEL62" and electrically compatible with the original TELL1, is a 9U-format board which can house 4 independent mezzanine cards, each one served by a dedicated FPGA ("PP-FPGA") with about 6Gb/s input bandwidth and 2GB of DDR2 dynamic RAM. A fifth FPGA ("SL-FPGA") collects data from the four PP-FPGAs and drives another mezzanine card with four 1-Gigabit Ethernet (GbE) links. Other features include an on-board control PC ("CCPC") for slow control (via a dedicated 100 Mbps Ethernet link), and a TTC receiver with a jitter-cleaning quartz-crystal-based phase-lock loop (QPLL), as well as user-definable connections for flow control. The card requires a special backplane with a single VME-like connector, which is only used for power. Apart from the general-purpose Gigabit Ethernet links, no inter-board communication mechanism is present, but a user-defined connector on the mother-board accepts a dedicated daughter-card which can be used for this purpose. Eighty to one hundred TEL62 boards will be produced for NA62.

LHCb developed, for the TELL1, a 16-channel, 40-MHz, 8-bit flash ADC mezzanine card (A-Rx), and a double-sized mezzanine card O-Rx [96] with 8 optical link receivers for CERN's Gigabit optical link transmitter (GOL). Many NA62 sub-detectors require no, or limited resolution on, pulse-height information, and so will use only TDC-based readout systems. Most of these will employ a NA62-built TDC mezzanine card [97].

[52] This includes the situation in which a sub-detector asserted the choke line but after a while, not seeing a pause in the trigger stream, is really no longer able to process additional events.

The NA62 mezzanine TDC board ("TDCB") is equipped with 4 HPTDC chips [98] an FPGA (TDC Controller FPGA, or "TDCC-FPGA"), and 2 MB of static RAM. The board will service 128 input channels (LVDS) over 4 halogen-free SCSI-3 twisted-pair cables, measuring time and time-over-threshold with 100-ps LSB precision. Some capability to process data will be built-in. Up to 4 boards can be housed on the motherboard, for a total of 512 channels per mother-board. When equipped with TDC boards, two slots of crate space per mother-board are required.

Since time measurements are performed by the TDCs with respect to clock edges, several stages of filtering will be present to reduce the time jitter of the reference clock down to the ≈50 ps (RMS) level, without compromising the time resolution. It is possible to monitor one (fixed) channel per TDC chip (1/32) and/or drive it with the TDCC-FPGA for time calibration and debugging. The FPGA can also drive an additional LVDS pair, *e.g.* to trigger front-end board calibration pulses. Additionally, one (fixed) channel per chip (1/32) can receive a NIM signal from a front-panel LEMO connector, rather than from the input connector, for debugging and testing purposes.

The firmware for the TDC system will be split among the four TDCC-FPGAs (on the daughter-boards), the four PP-FPGAs, and the single SL-FPGA (on the mother-board). The code is loaded from on-board EPROMs at reset, and can be modified by accessing the board with a JTAG programming cable or with software via the on-board PC.

The **TDCC-FPGA** should:

- communicate with the CCPC by acting as an I2C slave, for receiving commands and configuration instructions;
- configure the four TDCs by acting as a JTAG master, according to data provided by the CCPC;
- control the TDCs at runtime, reading out data words when available and associating timestamps with them;
- monitor TDCs status, recording the occurrence of fatal errors (and their time of occurrence) in registers which can be read at the end of each burst;
- optionally pre-process TDC data words in local memory;
- provide data words on parallel and independent buses to the PP-FPGAs.

While HPTDCs have extensive multi-hit capability and internal trigger matching capabilities, their buffers are insufficient to store hits for the full latency of the L0 trigger. Data will therefore be read out continuously from the TDCs and buffered in the large RAM on the mother-board. Data time stamped within a programmable time window will be extracted after each L0 trigger.

When configured to run with 100-ps LSB, the TDC timestamp rolls over after 51.2 μs (11 bits) and higher timestamp bits must be added. Although it would suffice to add the minimum number of bits to cover the maximum L0 latency, it is simpler to employ a full 31-bit timestamp word, thus requiring 20 additional (higher) bits. The time-matching feature of the HPTDC chips will be used by the controller FPGA to extract all hits belonging to a given time frame. By periodically triggering each TDC chip, and by setting the time window to the period between triggers, all hits are read out, automatically arranged in time-ordered frames. Each TDCC-FPGA will communicate independently with four TDC chips, triggering them all simultaneously, reading them out, adding upper timestamp bits, and storing all of this in associated FIFO buffers, from which the PP-FPGA will read using the same communication protocol. The TDCC-FPGA should monitor the filling of these FIFOs, and record instances of full

capacity, when possibly data were lost. Should there be no TDC data for a particular trigger (as will often be the case), no data frame will be generated.

For the TDC application, the **PP-FPGA** should:

- independently read TDC words from the four TDC boards and merge them into a single data stream;
- pre-process TDC words, in both a common and a sub-detector dependent way;
- split the TDC-generated time frames into convenient sub-frames (*e.g.* 25ns wide) and store TDC words in (off-chip) memory during the L0 trigger latency;
- (for sub-detector contributing to the L0 trigger) evaluate L0 trigger primitives and send them to the SL-FPGA;
- retrieve from memory a programmable number of data frames in response to a L0 trigger request from the SL-FPGA, and send them to the SL-FPGA.

For the TDC application, the **SL-FPGA** should:

- react to start- and end-of-burst commands, and dispatch time-critical commands to the PP-FPGAs;
- collect L0 triggers coming through the TTC;
- dispatch time stamped (and qualified) L0 trigger requests to PP-FPGAs;
- collect, merge, and format data packets from the four PP-FPGAs;
- prepare and send formatted multi-event packets (MEPs), together with dispatch control data, to the Ethernet link card for transmission to PCs;
- (for sub-detectors contributing to the L0 trigger) collect L0 primitives from the PP-FPGAs, merge them (perhaps with those from other boards), and send them to the L0 Trigger Processor or to another board.

The PP-FPGA will continuously move hits, timestamps, and errors independently from the four TDCC-FPGA FIFOs into other internal FIFOs, and then merge the hit frames into a common “input buffer”, making efficient use of available bandwidth. As the data are being stored, they can be monitored (*e.g.,* saving histograms of hit and error counts per TDC channel, to be read at the end of bursts), reformatted in smaller time frames, and processed (*e.g.*, pedestal can be subtracted, gains computed, etc.). For sub-detectors contributing to the L0 trigger, trigger primitives (*e.g.,* hit multiplicity in small fine-time bins) can be evaluated and temporarily stored in fine-time addressed buffers for transmission to the SL-FPGA.

To ease data retrieval, timestamp-addressed frames will be written into external DRAM memory, using a fixed page allocation. Because an integer number of frames around an L0 trigger timestamp will be read, the frame length should not be too long (25ns being the default). On the other hand, access to SDRAM is optimized with long writes; therefore frames might be grouped for writing, taking care to handle the roll-over when pages are re-used for new data. The number of words stored in each page will be kept in FPGA memory and cleared at roll-over time.

When the SL-FPGA transmits a L0 trigger request to the PP-FPGAs, it will also send timestamp and trigger-type information. The PP-FPGAs will extract the corresponding data frames from DRAM and pack them into FIFOs (together with a word count), from which the SL-FPGA will retrieve them. The

memory controller of the DRAM should arbitrate between write accesses (periodic frame storage) and read accesses (readout following a L0 trigger).

For sub-detectors contributing to the L0 trigger, each PP-FPGA should also transmit to the SL-FPGA trigger primitives information. A separate communication bus between the two types of FPGAs on a board will be used for this kind of data. The SL-FPGA will synchronously merge trigger primitives from the PP-FPGAs. Should a sub-detector require more than one board, these merged primitives must be further merged with similar information transmitted asynchronously from a previous board (e.g. using one of the GbE links), and again transmitted to the next board or to the L0 Trigger Processor. Such inter-board communication requires additional firmware support not required for sub-detectors using a single board (or those not involved in the L0 trigger decision).

The SL-FPGA will handle TTC communication. After receiving a start-of-burst signal from the TTC, it will distribute synchronous reset signals to each PP-FPGA, which in turn will synchronously reset TDC chips and timestamp counters in the TDCC-FPGA. When receiving a L0 trigger, it will retrieve the timestamp and trigger-type information and send L0 trigger requests and the corresponding trigger data to each PP-FPGA. When it receives an end-of-burst signal, it will record the last timestamp count reached.

For the sake of performance checks and the transmission of summary data at the end of bursts, the firmware should permit debugging (enabled only during tests) and monitoring (enabled during normal running, with no performance penalty).

A preliminary version of the firmware, nicknamed “TDCTEST” was developed in 2008-2009 to test the original TELL1 board, and was used in a 2009 RICH test beam. Lacking several required features, it was suitable for use only in simple standalone DAQ tests. It lacked: (a) large latency buffers for data storage; (b) timestamp generation, so that all data in a circular buffer were read out, irrespective of their time; (c) internal triggering, so that an externally generated trigger was required. The production version of the firmware, nicknamed “TDC”, will overcome the limitations of the TDCTEST version and will be suitable for a generic sub-detector's readout and L0 trigger primitives generation. Sub-detector groups will develop dedicated versions, “TDCRICH”, “TDCLAV”, etc., based on the TDC firmware. In most cases, only the PP-FPGA code will require modification.

The board will have other uses in NA62, as well, such as in the LKr L0 trigger system (section 4.1.15.5), and additional TELL1-compatible mezzanine cards will be developed, such as a pair of high-bandwidth asymmetric link cards and a digital data receiver card.

4.1.6 L0 Trigger Processor

At the end of the latency period T(L0P)max, the L0 Trigger Processor (L0TP) will time-match the lists of trigger primitives it has received, checking whether appropriate (programmable) L0 trigger conditions have been satisfied within (flexible and programmable) overlapping time windows. It will simultaneously check multiple trigger conditions, with possibly different time constraints. It will generate down-scaled triggers for control purposes, as well as calibration and monitoring triggers.

Upon determining that conditions are satisfied, the L0 Trigger Processor will issue a L0 trigger. Candidate L0 triggers closer in time than a programmable range will be coalesced into a single trigger, since the time window within which sub-detectors will read data will be significantly larger than the L0 trigger time resolution. The list of remaining L0 trigger candidates, each having both a time tag and a

trigger type word, will be used to generate pulses with a fixed delay from the time of particle crossing, in order to generate a synchronous trigger signal (25ns resolution) to be dispatched by the TTC system (the so-called "L1A" signal in TTC jargon). For each L0 trigger sent in this way the corresponding 6-bit trigger word will follow asynchronously, to be recovered by the TTCrx chip on each system.

In response to every L0 trigger, all sub-detectors will assign each L0-triggered event a sequential event number (uniquely linked to the L0 trigger timestamp, within a spill) and move data from within a certain (possibly trigger-dependent) time window around the trigger time in the form of event frames to permanent buffers (usually in sub-detector L1 PCs), where they can be erased only after a subsequent negative L1 trigger. These actions must be taken even when no data are available within the defined time-window, in which case the event frame will be empty.

Apart from inputs from the participating sub-detectors, the L0 Trigger Processor will also have other inputs, to be used as sources for "forced" synchronous triggers. In response to such pulses the L0TP will arrange, if possible, a special trigger to be sent after an appropriate (constant) time delay. This feature might be used in order to handle, *e.g.,* triggers related to calibration pulses in some sub-detector.

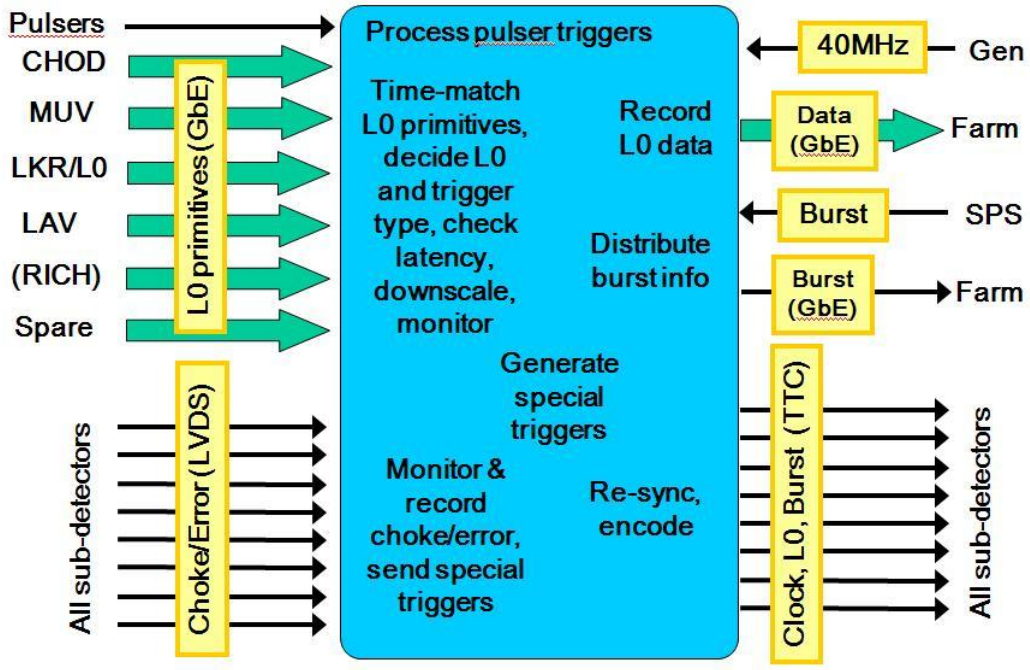

Figure 310: Logical scheme of L0 Trigger Processor.

The L0 Trigger Processor will continuously monitor maskable choke/error lines coming from each sub-system, logging the activity on each for inclusion in the data stream at the end of burst. In response to an unmasked choke signal, the L0 Trigger Processor will dispatch a special **"choke on" trigger** and cease issuing triggers. This special trigger is expected to be serviced as any other by the sub-detectors: the corresponding data frames will provide a means to check that no data were lost prior to the signal. When all (unmasked) choke lines go low, the L0 Trigger Processor will issue a special **"choke off" trigger** to resume normal operations.

In response to an unmasked error signal, the L0 Trigger Processor will dispatch a special **"error on" trigger,** in response to which sub-detectors may be required to send special monitoring data to

diagnose and debug (offline) the situation when the error condition occurred. When all (unmasked) error lines go low again, the L0 Trigger Processor will issue a special **"error off" trigger** before resuming normal operation.

The implementation of the L0 Trigger Processor is not yet defined. Several options are under consideration: a custom FPGA-based design is certainly possible, but the use of a high-performance PC with a real-time operating system, if proven to be feasible, might be simpler, more cost-effective, and easier to program and to maintain. In any case, a custom hardware part will be present, which will take care of re-synchronization of the L0 trigger pulses, and communication to the LTU boards; this will be implemented as a card with a PCI-Express interface, which easily adapts to any L0TP implementation.

4.1.7 CEDAR System

The CEDAR system will not participate in the formation of the L0 trigger, but might contribute to higher trigger levels.

The average kaon flux (neglecting accidentals) will be 50 MHz. Using 240 PMTs, the average singles rate will be about 5 MHz. Around 18 photons per kaon will be detected, while the probability of more than one photon traversing a given photocathode will be of order 1%.

The intrinsic time resolution of the PMTs is 300-400 ps, and a single-kaon time resolution of about 50 ps is required to suppress accidental background, so a minimum of 10 photons should be detected per kaon. Since dead time losses are of the order of 15 ns (11ns from NINO and 5 ns from the HPTDC), the double pulse resolution of the whole system is dominated by the electronics and will be no worse than 15 ns.

Both leading and trailing edge times, and thus the time-over-threshold, of the signal pulse will be measured. These times will be used to determine pulse amplitudes so as to correct for time-slewing induced by amplitude fluctuations and to discriminate against pile-up. Assuming a Gaussian shape for the PMT analog signal, one can show that edge times and time-over-threshold are linearly correlated. It is possible in principle to predict the correlation between time over threshold and amplitude, but for this purpose the contribution of electronic noise to the pulse shape must be negligible.

The readout for the CEDAR will be based on the common TDC/TEL62 system. The processing done in the TEL62 will include counting of the number of PMTs fired per spot and the number of fired spots, to later develop algorithms that use the multiplicity and the pattern of spots to suppress the background. It is also foreseen to record the number of photons in 1ns slots within the readout window, to allow for a precise time coincidence with the trigger signal. At any time, the TEL62 buffers should hold all the data corresponding to a time interval of 1 ms. Given an average of 9 hits per kaon per TEL62, with 32 bits per photon (leading and trailing edge measurements packed into a single word by the TDCs) and a kaon rate of 50 MHz, a maximum input of ≈40 Bytes is expected per kaon per TEL62, which corresponds to ≈200 KB in the buffer of 1ms, well below the size of the buffers available in a TEL62.

The NINO discriminator chip introduces a stretching time of 11ns, while the HPTDC has a dead time which can be set to 5ns. Nevertheless, the major limitation of this electronics is the finite size of the HPTDC hit channel buffer, which can sustain a maximum of 40 MHz over a set of 8 channels, or a maximum of 10 MHz for a single channel, whichever is lower [15]. The 16 HPTDCs in a TDC/TEL62 board are arranged in sets of 8 channels which share some group channel buffers. By using only 2

channels for each set (1/4 of the total number available), the bandwidth will be reduced to a level guaranteeing a small probability of hit loss. Thus, only 32 channels per TDC board (128 channels per TEL62), and the CEDAR readout system will consist of 16 TDC boards and 2 TEL62 boards. The hit loss probability estimated in this configuration is 1%, which corresponds to a detection inefficiency per kaon of a few percent for a minimum of 10 photons distributed over at least 6 spots. The contribution of the readout system to the time resolution is estimated to be ≈70 ps, based on the performances of a prototype system used in RICH tests, a negligible contribution when compared with the intrinsic PMT resolution.

The leading and trailing edge measurements of each pulse can be combined in one word in the TDC, using 7 bits to measure the duration of the signal. Assuming an LSB corresponding to 200 ps, a maximum of 25ns can be measured as the signal length. The readout time window will be (-35 ns, +15 ns), for a total window of 50 ns. The window is not centered on the trigger value, to take into account the 25ns of the maximum signal length quoted above; 2-3 kaons are expected to appear in such time window. With an average of 18 hits per kaon, this corresponds to ≈216 Bytes in the readout window. Assuming a trigger rate of 1MHz, this corresponds to ≈0.2 GB/s of readout rate. Digital signals will be sent from the front-end electronics to the readout electronics through a few metres of high-quality twisted pair cable.

4.1.8 Gigatracker (GTK) System

As shown in the diagram in Figure 311, each GTK sensor is read out by 10 Giga Tracker ASICs (GTK-ASIC) whose output data flows continuously toward the GTK off-detector readout (GTK-RO) cards. These cards provide temporary data storage until the L0 trigger decision, upon which the GTK-RO cards extract trigger matched data from the on-board memory buffers and transmit the data to the on-line PC farm through Gigabit Ethernet switches.

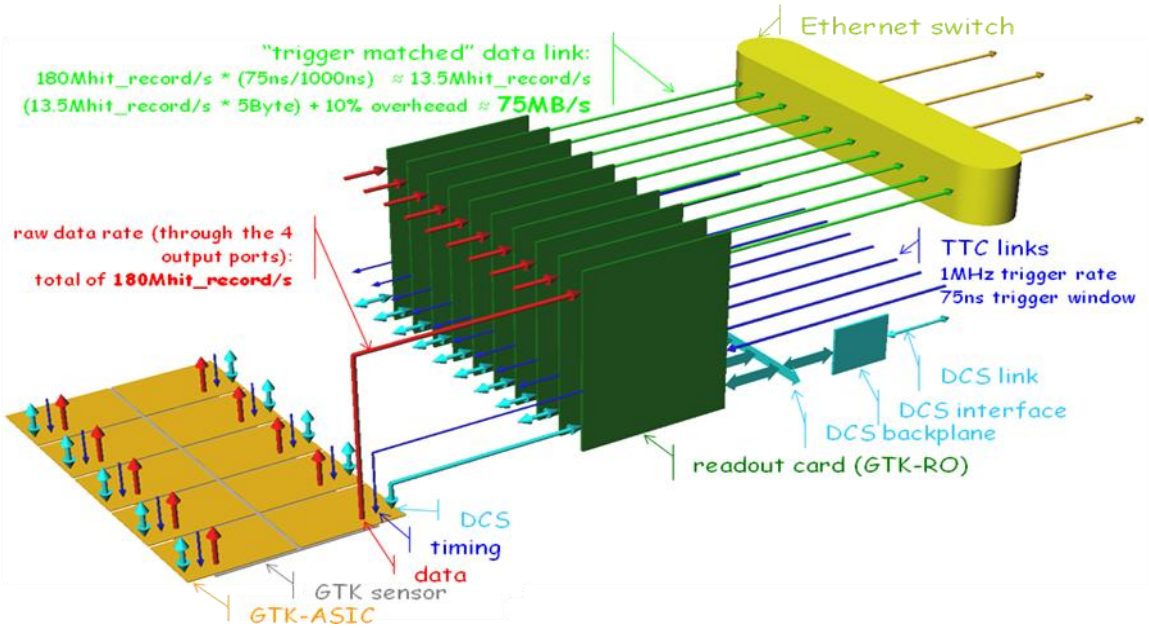

Figure 311 GTK Read-out block diagram.

The most relevant GTK sensor /ASIC parameters for the GTK off detector read-out are summarized below:

- maximum hit rate on centre pixel ≈ 1.5 MHz/mm^2, 140KHz/pixel;
- average hit rate per sensor plane ≈ 750 MHz;
- data word width: 32 bit;
- average centre chip hit rate: 132 MHz;
- average centre chip data rate: 4.3 Gbit/s;
- design data rate (chip rate + contingency): 6 Gbit/s;
- serial links per chip: 2-4;
- readout window: 75ns.

The GTK read-out is based on the following scheme.

The GTK-ASICs continuously send periodic frames to the GTK-RO board. Each frame is made of one header, one hit record and one trailer. The header carries the GTK-ASIC ID and a rolling frame number, which is zeroed at the beginning of the data taking run. The hit record provides the local hit address within the pixel matrix and the time measurement which consists of a fine measurement derived from the TDC in the ASIC and a coarse measurement from the synchronous clock counter in the read-out ASIC. The range of the fine time is one clock period. The roll-over period of the coarse time determines the time interval covered by each output frame. If the EOC read-out architecture is chosen, a second fine time measurement for the trailing ToT signal is sent. The full dynamic TDC range extends to 6.4 µs. The rolling frame number in the header is used to extend the system dynamic range of each hit to more than 10 ms. Each hit record is encoded in 4 bytes for the P-TDC architecture and 5 bytes for the EOC architecture to provide relative leading- and trailing-edge timing information. The trailer carries status information and a CRC-16 checksum.

	N. of bits	Resolution	Range
TAC	6 7	195 ps (56 ps r.m.s) 98 ps (28 ps r.m.s.)	12.5 ns
Coarse counter	11 10	6.25 ns	12.8 µs 6.4 µs
Frame counter	16	6.4 µs	838.9 ms 419.4 ms

Figure 312 Hit format, resolution and range.

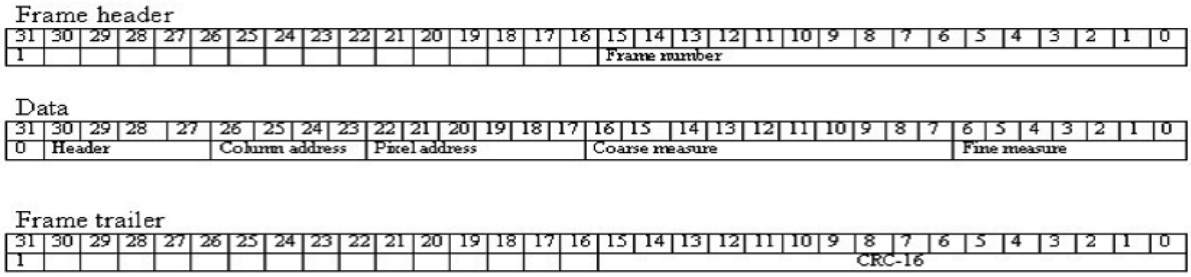

Figure 313 Header, data and trailer format.

Each GTK ASIC features 2 or 4 output ports, depending on final implementation. Depending on the ASIC internal distribution of the data flow and assuming an average hit rate of 4.5 MHz (3.3 MHz + 35 % contingency) per column, each output port transmits up to: $45 * 10^6 * 5 * 8 * (10/8) = 2.25$ Gb/s including the 8b/10b encoding overhead but not including other sources of overhead like, for instance, the frame headers and trailers, of the order of 1%.

Figure 312 and Figure 313 show the hit word and frame format already implemented in the GTK P-TDC demonstrator [19], which easily can be adapted to the EOC architecture.

The requirements of the GTK read-out system follow from the above assumptions and are summarized in Figure 311.

To understand the data rates quoted in Figure 311 one must take into account the assumption that the L0 trigger matching is done on-line, *i.e.* not deferred to the “inter-spill” phase.

A 75MB/s data rate at the output of a GTK-RO card is achievable with a single GbE link if jumbo frames are used and if it is acceptable not to implement the full TCP/IP protocol. Two links could be otherwise foreseen, which would provide also some headroom in case the actual data rates would exceed the expected ones.

The following paragraph summarizes the read-out parameters for all three detector stations and for the configuration where 1 GTK-RO board serves one GTK-ASIC and 2 GbE links are available for one GTK-RO board.

- 750 MB/s of trigger matched data per GTK station
- 2.25 GB/s of trigger matched data for the entire Giga Tracker

Considering that 1 trigger-matched event contains, on average, 180Mhit/s * 1µs * 0.075 = 13.5 hits per GTK-ASIC, and thus 135 hit records per GTK station, one finds

- 405 hits per event for the (3 stations) GTK and thus, assuming 5 Bytes per hit record and a 10% overhead:

- The average size of a GTK event ≈ 2250 Byte

One possible implementation of the GTK-RO board is outlined in the block diagram in Figure 314.

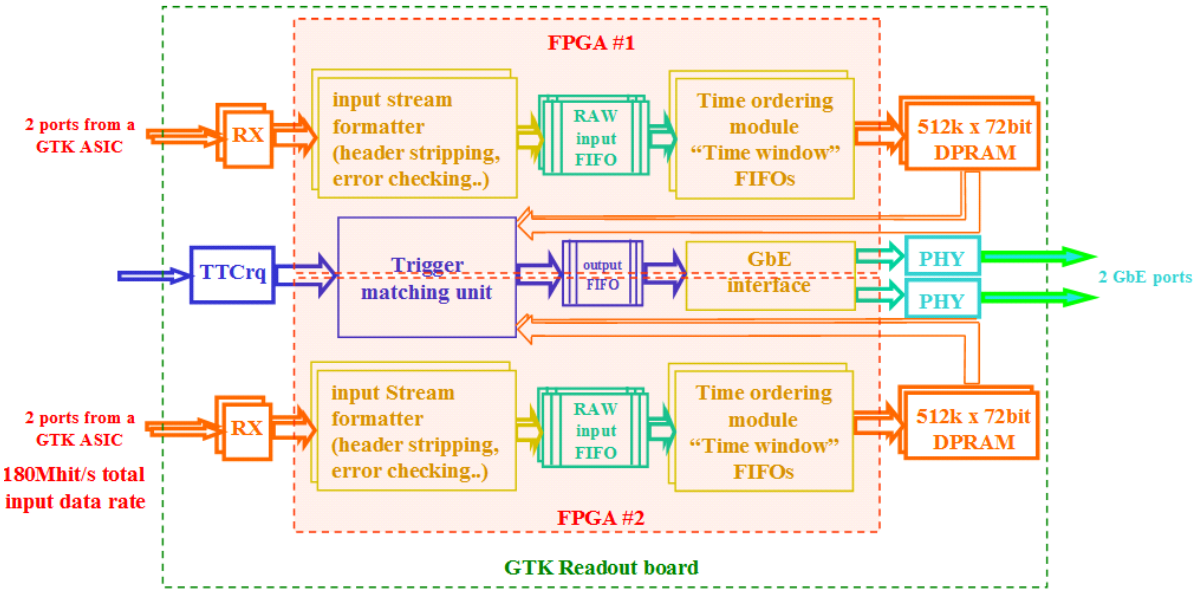

Figure 314 GTK read-out board block diagram

KEY PARTS:
RX: the choice of the deserializer depends on the standard chosen for serial TX on the final GTK ASIC.
TTCrq: TTCrx+QPLL mezzanine daughter card by CERN.
FPGA #1 - 2: Altera Stratix III, 780 pin, 480 I/O, 50k logic elements: EP3SL50F780C3N (≈ $ 525 each).
DPRAM : Cypress synch DPRAM, 167Mhz, 512K x 72bit, 1.8V, CYD36S72V18.
Ethernet PHY #1 - 2: Marvell Alaska GbE PHY.

The key elements of any GTK-RO card architecture are the memory devices used for the temporary storage of the GTK data during the trigger latency time interval. A dual-port RAM (DPRAM) has completely independent read and write timing signals, R/W data buses and R/W address buses; the DPRAM considered in the diagram above has a capacity of 36 Mbits and operates at a maximum clock speed of 250 MHz. As shall be shown later, a GTK-RO card equipped with such devices may cope with a trigger latency of about 6.5 ms.

An alternative memory device such as the QDR-II+ Static RAM, features completely independent R/W timing signals, R/W data buses but a single address bus to be shared among the write and the read accesses. A QDR-II+ SRAM device can operate at frequencies up to 400 MHz, and thus a fast FPGA should be used to access the memory device by multiplexing the write and read concurrent accesses. QDR-II+ SRAM devices offer a higher density/price performance with respect to DPRAMs and become thus particularly effective large memory buffers are necessary to cope with an increased L0 trigger latency time. The following paragraph describes the operations of the individual blocks of the diagram sketched in Figure 314.

The **input stream formatter** controls the packets coming from the GTK-ASIC for CRC errors and checks the frame number against the one calculated on the GTK-RO card. It also strips off the header from the GTK-ASIC and transfers the hit data words into the Raw-Input FIFO appending an End of Frame (EOF)

marker to the data packet for 1 frame. Any error detected in the incoming packets would be coded into a specific field of the “EOF” marker to forward the error information.

The **time ordering module** (see Figure 315) is meant to facilitate the trigger-matching operation. The need for this module arises from the fact that the hit records coming from the GTK ASICs are not necessarily time-ordered. It seems then advisable to pre-order the input data in bins and store the content of each bin in a specific memory page of the buffer memory (be it DPRAM or QDR-II+). A bin is actually a small FIFO (Time-Window FIFO) memory inside the FPGA which is filled as the input data is extracted from the Raw-Input FIFO. Part of the coarse time measurement field of a particular hit record determines the FIFO it will be assigned to. If, for instance, a GTK data frame contains all the hits recorded by a GTK-ASIC in the previous 6.4 µs, and if there are 16 Time-Window FIFOs, each of these would contain, when the Raw-Input FIFO is completely read out, all the hits recorded in a time window of 400 ns. One independent memory buffer is foreseen for each port. At the average input hit data rate for one GTK-ASIC port of ≈ 45 Mhits/s one 6.4 µs frame will contain 288 hits. When one GTK data frame is fully processed, each of the 400 ns Time-Window FIFOs should contain on average 18 hits.

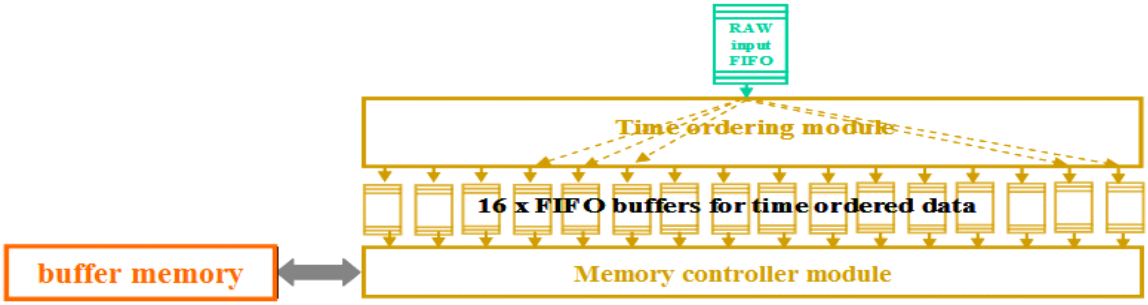

Figure 315 Block diagram of time ordering.

The contents of each Time-Window FIFO are then transferred to a specific memory page in the buffer memory. The frame number and the ID of the Time-Window FIFO (*i.e.* part of the coarse timing information in the hit record) determine the base address of the target memory page. When the content of a Time-Window FIFO is completely transferred to a memory page an End-Of-Page terminator is also stored. Taking into account that the DPRAM memories are 72 bits wide (36 bit is not enough for the EOC architecture) some of the coarse timing measurement of a hit is already encoded in the address, and 2 hit records could share the same memory buffer location.

Assuming that a memory page depth of 32 locations is reserved to store all the hits in a 400 ns time window (corresponding to a maximum capacity of ≈ 62 hit records, to be compared with the 18 expected in average) then a 512K * 72 bit memory buffer would allow for a trigger latency of 16384 * 0.4 µs ≈ 6.5 ms.

The interface to the memory buffer runs at a clock speed of 200 MHz, so up to 4 clock cycles are needed to fetch and assemble data in “double hit” words. One could then write 300 hit records (expected from each port in a 6.4 µs frame interval) in approximately 3µs (20ns * 150), which is comfortably shorter than the maximum 6.4 µs frame period.

When processing an L0 trigger request, the trigger matching module must evaluate, from the known trigger latency, the memory page. Once the memory page is reached, 18 hits (the average content), can be scanned in ≈ 100ns (10 * 10 ns), assuming that 2 clock cycles are needed to read and check for

the page terminator and assuming also that 2 hits are stored in each memory location. Thus the architecture outlined here seems capable of meeting the 1MHz L0 trigger rate.

The outcome of the trigger matching should be the extraction of an average of 13.5 hits from the whole GTK-ASIC, *i.e.* from the total of the 4 ports. Assuming that the hits are recorded in 5 Bytes, and taking into account an overhead of 10% for the event header, the average size of an event packet should be about 75 Bytes. Multi-event packets, similar to those built by the TEL62 board, should be assembled to optimize the bandwidth of the Gigabit Ethernet connection.

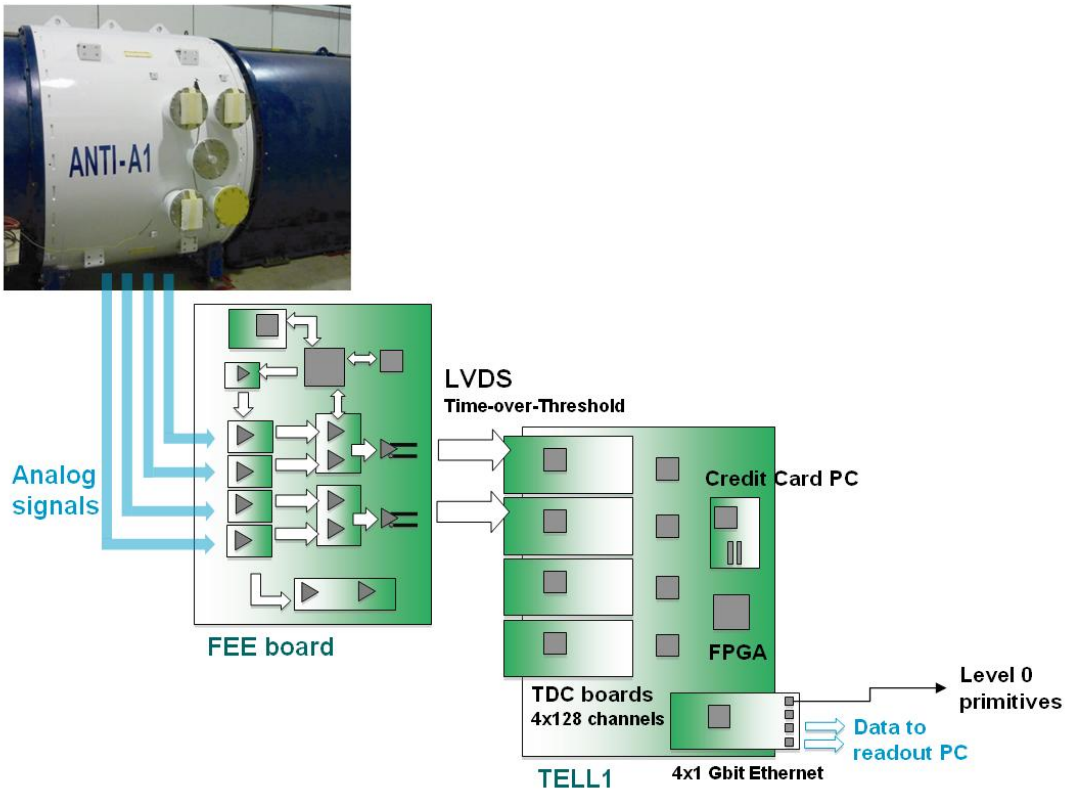

Figure 316: LAV readout scheme

4.1.9 LAV System

The output signals from the PMTs of each of the 12 stations are connected to Front-End Electronics (FEE) cards. These FEE cards discriminate the analogue signals from the PMTs generating digital signals of proper width (equal to the Time-Over-Threshold duration), using the differential LVDS standard.

Given the large dynamic range of the analogue signals, two digital lines are used for each PMT, with different thresholds, to allow a measurement of the signal rise time to correct for time slewing. The corrected time and the signal amplitude are reconstructed from the four (two leading, two trailing) measured times for each PMT pulse. L0 trigger primitives are also evaluated by the system, and sent to the Level 0 Trigger Processor. The general layout is shown in Figure 316.

Each FEE board has two 32-channel LVDS output connectors (a total of 128 wires), corresponding to two TDC board cables (half a TDC board). As mentioned, each PMT corresponds to two TDC signal pairs

corresponding to different thresholds, to allow time slewing correction and to ensure redundancy in case of TDC broken channels.

In Table 57 the number of PMTs per station, stations of each type, and the total number of TDC channels, FEE cards, TDC boards, and TEL62 boards for the whole LAV system are listed.

Table 57: LAV channel counts.

Station type	$N_{PMT/station}$	$N_{stations}$	N_{PMT}	N_{ch}	N_{FEE}	N_{TDCB}	N_{TELL1}
Type 1	160 (5x32)	5	800	1600	25	15	5
Type 2	240 (5x48)	3	720	1440	24	12	3
Type 3	240 (4x60)	3	720	1440	24	12	3
Type 4	256 (4x64)	1	256	512	8	4	1
Total		12	2496	4992	81	43	12

As the Table 57 shows, the LAV has about 2500 PMTs handled by about 5000 TDC channels. To equip the whole system about 90 FEE boards, 50 TDC boards, and 15 TEL62 are needed (including spares).

The dominant component of the rate in the LAV is due to muons, coming both from the beam halo and from K decays. The expected muon rates into the 12 LAV stations due to the beam halo are summarized in Table 58. Assuming that for each muon hit two 32-bit words are recorded from the TDC for each threshold, and that a muon is not firing more than seven blocks in each LAV station, the data rate is also indicated in the Table 58.

Table 58: LAV data rates.

	N_{PMT}	N_{fired}	N_{word}/channel	Rate	N_{bit}	Mbit/s
LAV 1	160	7	4+2	1.77 MHz	32	2378
LAV total	2496	7	4+2	11.2 MHz	32	15053
LAV OR	2496	< 25	4+2	4.13 MHz	32	19842

The number of firing blocks (N_{fired}) has been over-estimated and we introduce a 50% extra hits, therefore the above figures for the data rate should be considered as upper limits. The rate per single station is lower than 1MHz on average, but is ≈2MHz in the station LAV1. This rate can be translated into a single channel rate by assuming azimuthal symmetry for the muon halo, and the muon direction being parallel to the detector axis. In this approximation, for the LAV1, the halo rate is equally shared by 32 blocks of a single ring, resulting in a rate per single channel below 65 kHz.

The expected hit rate for a single channel is easily managed by both the FEE cards and the TDC boards. The total LAV data rate (< 20 Gbit/s) will be divided over the 12 TEL62 boards of the LAV system. Each TEL62 is equipped with four 1Gbit Ethernet interfaces and can therefore tolerate the expected rate.

The data will be sent to a commercial 48 in/48out 1Gbit Ethernet switch connected to the LAV sub-detector PCs.

The raw data coming from the TDC for each hit, namely the leading and trailing edge times for the two different thresholds, shall be converted into a slewing-corrected time (T_0) and charge.

Slewing correction

Define: T_L = leading edge time for the lower threshold, T_H = leading edge time for the higher threshold, T_0 = event time extrapolated to zero amplitude (slewing-corrected), L_{THR} = lower threshold level, H_{THR} = higher threshold level. The correction is obtained with the following formula:

$$T_0 = T_L - L_{THR} \cdot \frac{T_H - T_L}{H_{THR} - L_{THR}}$$

Charge computation

The computation of the charge is obtained according to a polynomial parameterization: after pairing the leading and trailing edge times, the 4th degree polynomial is computed to get the charge.

Both the above operations are performed on each hit and can therefore be executed in parallel in the PP FPGAs (described in Common TDC System)

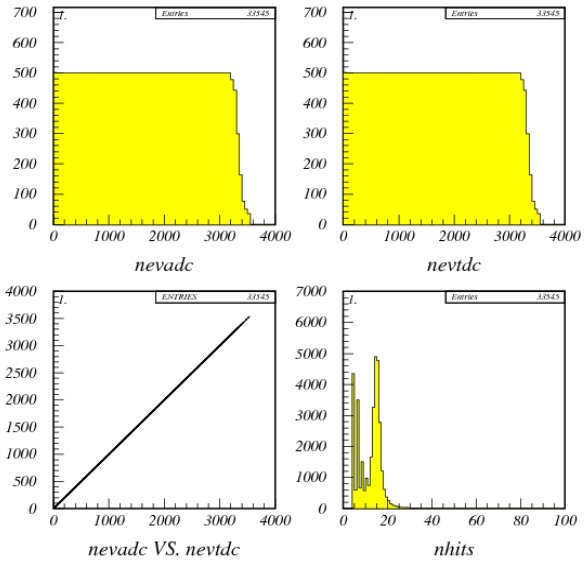

Figure 317: LAV hit multiplicities.

L0 Trigger

The raw information available in the LAV system at L0 trigger evaluation time are the times and charges of single hits in each block of a single station. These are all collected in the TEL62 board serving one station and can be elaborated by the on-board FPGAs to produce for the L0 processor simple primitives containing ring and station summary informations, including the type of particle.

As each station is composed by 4 or 5 ring-shaped layers of crystals, we can combine data to produce simple summary variables containing ring and station information:

Ring primitives (circle of blocks)

- E_bl = reconstructed charge in a single block
- E_ring = charge sum of all the blocks in the ring
- N_ring = number of blocks above threshold in the ring.

Station primitives

- E_tot = total charge of all the blocks in the station
- N_tot = total number of blocks above threshold in the station
- N_cl = total number of clusters, using a proximity algorithm.

Using the above informations and very simple cuts, two LAV triggers primitives can be constructed:

MIP trigger: to identify MIPs (μ or π) and to distinguish them from photons and electrons:

- N_ring ≤ 2 for each of the five or four rings
- 3 ≤ N_tot ≤ 7
- E_bl < 250 MeV for each block over threshold
- E_ring(i) /E_ring(i+1) < 2 for each pair of rings
- N_cl = 1, one cluster (only) in the LAV station

During the LAV-1 test beam in October 2009 at CERN, runs with 2 GeV electrons have been taken. In these runs, a fraction of events are generated by muons from the beam halo. The total TDC hit multiplicity is shown in Figure 317. The separation of electrons, whose hit mean value is ≈15, and muons, with TDC hits ≤ 7, gives a first indication of the capability of distinguishing muons from electrons or photons in a single station by using a simple logic.

High multiplicity trigger: to identify EM-showering particles

- $N_{tot} > 15$ OR $E_{tot} > 20$ GeV
- $E_{ring} > 2.5$ GeV × N_{ring} for al least 2 rings
- $N_{cl} > 2$, more than 2 clusters in the LAV station

The trigger can identify the presence of high energy into one or more LAV stations indicating that the event will contain one or more electrons or photons. It can be used by the L0 central processor to identify π+ π^0 or π^+ π^0 π^0 events having one or more photons into the LAV.

Once the trigger primitives are produced by the L0 FPGA by the TEL62 of the involved station, the board sends this information, together with the event time stamp, to a L0 LAV concentrator using a dedicated Gbit Ethernet interface. In the LAV L0 concentrator the information coming from the 12 stations are combined together (see Figure 318). The aim of the concentrator is to try the association of clusters in different stations coming from the same particle as well as to determine the particle multiplicity in the whole LAV: muons traversing more than one station, showers starting in one station generating an energy leakage in the following ones, or more than one photon/electron/muon firing the LAV stations.

All time and charge information coming from the TEL62 of each station are first aligned in time, to correct for the time of flight, knowing the station position along the blue tube. For all clusters, the times are compared to check the hypothesis that they come from the same decay, and counted. Then, for each cluster the azimuthal position is computed and compared with that of the following station. If clusters match in azimuth, they are merged: the energy is summed, while the time of the particle is computed averaging hit times.

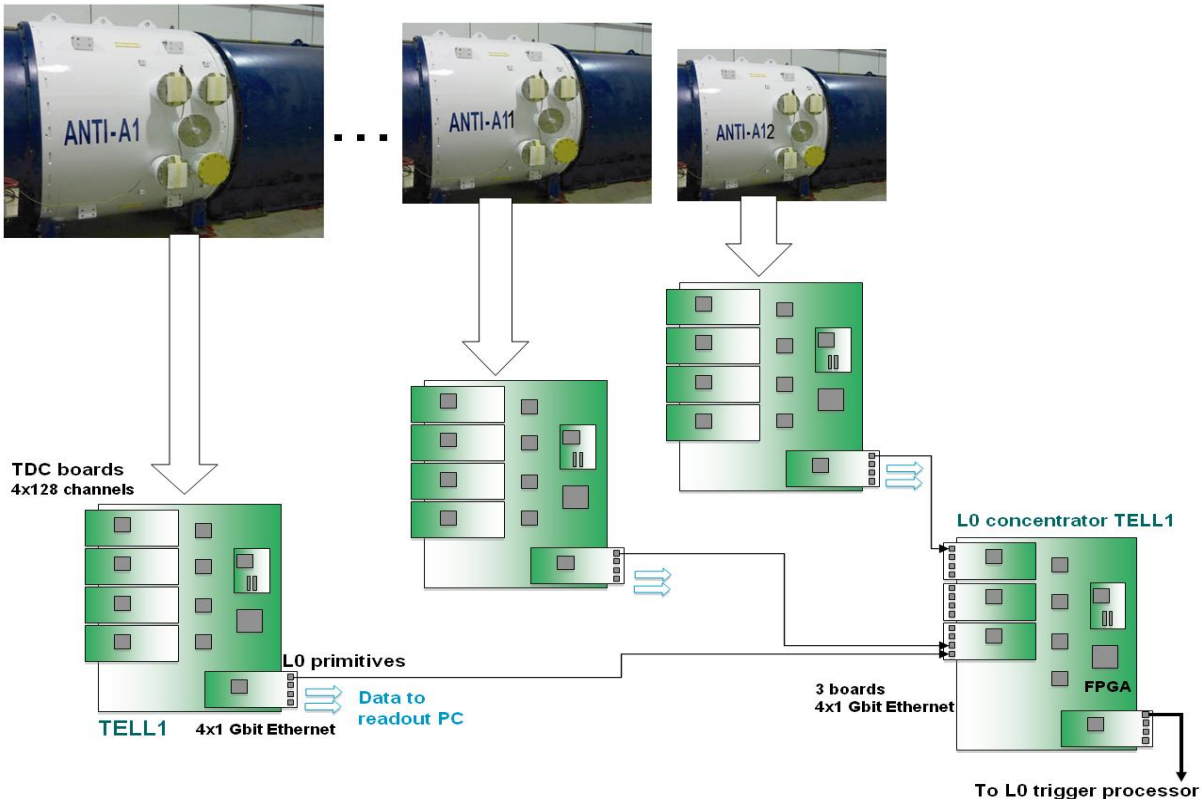

Figure 318: LAV L0 primitive generation scheme.

As results of the computation, a LAV structure is filled for each particle, with the following information:

- Particle time (at the position of a given station, *e.g.* LAV 12)
- Particle total energy deposit
- Total energy deposit in each of the crossed stations
- Azimuthal position (at LAV 12 position)
- Total number of crossed stations
- Total number of hits in each crossed station

Using the above information, better muon identification can be achieved exploiting the number of crossed stations and the energy deposit in each station. In fact, a pion can cross a single LAV station without producing a shower, but the probability of no conversion becomes lower and lower when the number of crossed stations increase.

Once the information are collected, the L0 primitive information encoded in two 32-bit words will be sent to the L0 Trigger Processor using a dedicated GbE link, thus saturating at 16 MHz. Assuming a muon-dominated rate of about 4 MHz (LAV OR), this leaves a safety factor of about 4. The 8 reserved bits can be used to encode the total event charge if needed by L0 central processor. The remaining particle information are stored in a circular memory buffer and sent if required to the L1 central processor.

In the present scheme the concentrator function can be obtained using a TEL62 and 12 Gbit Ethernet links: *e.g.* the 12 Gbit (double-slot) optical receiver mezzanine card developed by the LHCb experiment for the TELL1 (O-Rx card [96]). The computing power required for the data processing can be provided by the SL FPGA on-board the TEL62. The processed data are sent to the L0 central processor by one of the 4 standard TELL1 Gbit Ethernet connections. Compared to a daisy-chain architecture this scheme uses only 1 of the 4 Gbit Ethernet links housed by standard TELL1 and does not need the implementation of inter TELL1 communication protocol.

4.1.10 RICH System

The RICH detector is used in the trigger and offline to enhance the selection of events with a charged pion. The excellent time resolution of the RICH (100 ps) can be exploited in the L0 trigger to determine the reference time of the tracks. At L1 trigger level the RICH will provide the number and the position of the Cerenkov rings, helping to reject events with more than one Cherenkov ring, and at L2 trigger level such information can be combined with the spectrometer information in order to select pion candidates. As described in the corresponding chapter, the RICH detector has two active regions covered with approximately 1000 photomultipliers each, where the Cerenkov light produced by charged particles is focused by a spherical mirror mosaic. After the preamplifier electronics the 2000 signals are sent to 64 boards equipped with 4 NINO chips each (each NINO chip handles 8 input signals), acting as preamplifiers and discriminators. The LVDS output signals, with a time duration proportional to the time over the NINO threshold, are sent to the readout boards by means of 32-pair shielded cables. The cables and the corresponding connectors were chosen to preserve the excellent rise time of the output signal produce by the NINO chips, yielding low jitter time measurement.

Readout

The common NA62 TDC readout system is used for the RICH. Taking into account the requirements for the offline time resolution of the RICH (time resolution for a track better than 100 ps in the momentum range between 15 and 35 GeV/c), the LVDS differential signals, coming from the Front-End electronics (NINO boards), will be digitized by the HPTDCs on the TDCB cards described in section 4.1.5. Since a TDC card can handle 128 channels, two complete TEL62 boards will be used for each active region, for a total of 4 fully-equipped TEL62 boards (2048 channels), allowing for ≈2% of spare channels. The four TEL62 will be housed in the same 9U crate, placed close to the RICH detector.

The default TDC controller FPGA firmware will be used. In order to generate L0 trigger primitives as described below, a pre-processing algorithm can be applied before the transmission of the data in the PP-FPGA, in order to reduce the transit time of the data through the building stages of the trigger.

The TEL62 firmware for the RICH (both for the PP-FPGAs and the SL-FPGA), are modeled on the common default TDC firmware framework, as the main functionality of the TEL62 will be the same for all the detectors exploiting such board. The data will be stored in large memory buffers (more than 2GB for each PP-FPGA) waiting for the trigger. The memory is organized in pages; each page (corresponding to hits within a 25 ns wide time window) is addressed by synchronized timestamp counters in the TEL62 for writing. In case a positive L0 trigger is received, a programmable number of pages around the trigger will be readout from the memory being addressed by the corresponding trigger timestamp value.

During data taking, the monitoring of the board functionalities will be an important activity: in particular the buffer occupancies and other relevant working parameters will be directly evaluated in the FPGAs and transferred to the readout PC's, to spot malfunctioning, both in the readout system and in the frontend electronics. In case of special trigger requests (EOB[53], test triggers, etc.) the FPGA control trigger FSM (Finite State Machine) will provide statistics and reports collected by the internal monitoring blocks.

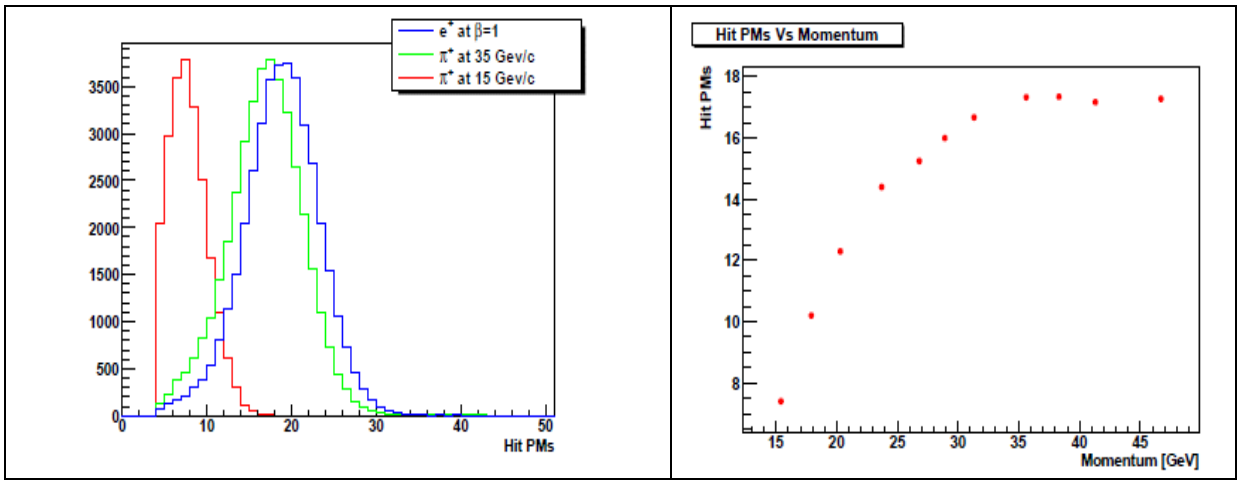

Figure 319: Left: number of firing PMs for pions and electrons.
Right: average number of firing PMTs for pions as a function of pion momentum.

The FPGAs in the RICH TEL62s will be also used to evaluate L0 trigger primitives, in order to contribute to the L0 trigger decision. In this respect both the PP-FPGA and the SL-FPGA contain dedicated blocks specifically designed for the RICH.

The data from the TEL62 will be directly sent to the L1 PCs, through commercial Gigabit Ethernet switches. The time-aligned data from different parts of the RICH will be merged directly in the L1 PCs.

The charged particle rate within the geometrical acceptance of the RICH detector is ≈11 MHz (10 MHz from kaon decays and 1 MHz from muon halo). With 8 bytes of data for each hit (channel number, leading and trailing edges' times), and assuming to have 20 firing phototubes for each particle due to

[53] EOB ≡ End Of Burst

Cerenkov light (the average number measured during the 2007 and 2009 test runs for ultra-relativistic particles being about 18, see Figure 319), a readout bandwidth of ≈1.8 GB/s is needed for the entire RICH. If the phototubes will be adequately mapped onto the TDCs in order to balance their average load, each TDC board will transfer to the corresponding PP-FPGA approximately 110 MB/s of data.

The I/O speed of the TEL62 DDR2 RAM memory is greater than 1.5 GB/s. This bandwidth is more than sufficient for both write and read operations. The buffer space available to store data on the original TELL1 board is 96MB, already enough for 0.9 s of data taking, so that the maximum overall L0 trigger latency for L0 is not an issue. In the new TEL62 the memory will be substituted with a faster and bigger DDR2 SDRAM, allowing bigger safety factors.

Assuming a L0 trigger rate of 1 MHz, the TEL62 output bandwidth will be ≈50 MB/s, which should be sustainable by one Gigabit Ethernet link (max 125 MB/s per channel).

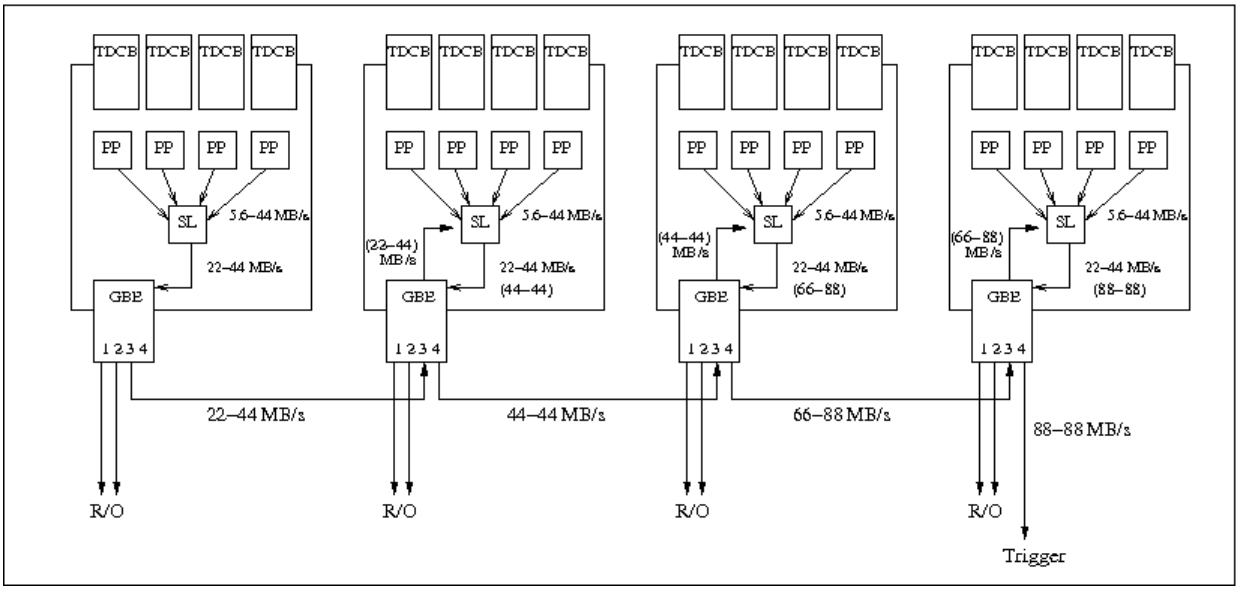

Figure 320: TEL62 daisy-chain connection for L0 primitive information. Bandwidths are indicated for "best" and "worst" cases; total bandwidths in a given stage are indicated in parenthesis.

L0 trigger

The RICH can contribute to the L0 trigger decision by providing a precise time reference for the charged particle. In this case the RICH L0 trigger primitives will be generated directly within the TEL62, by counting the multiplicity of the hits within fine time bins. For each memory page (25 ns wide) stored in the main buffers, the hits are further subdivided in 8 time windows (3.125 ns wide) in a histogram-like fashion. A threshold condition on the maximum number of single bin entries (time coincidences), can be applied in order to define the trigger primitives. A final time histogram for the whole RICH detector has to be constructed in three different steps: the PP-FPGAs (receiving data from 128 channels) define the first partial histograms; four of these are merged together in the SL-FPGA at TEL62 level. Finally the four histograms, produced by each TEL62 have to be merged again to obtain the final

result, from which a decision on the basis of the overall RICH multiplicity can be taken. The last merging can be performed by either connecting the four TEL62 boards in daisy-chain, using for this purpose two of the four GbE links on each board (see Figure 320), or by using an additional (fifth) TEL62 board equipped with Gigabit Ethernet mezzanine receivers (in this case only one GbE link is used for this purpose on each of the four sender boards).

In order to build a flexible trigger algorithm, the time histograms will also allow storing information from hits that were not previously ordered in time. In such a case the expected maximum time disorder, related to the time required for a hit to reach the PP-FPGA from the TDC, defines the minimum time interval for which each histogram must be left "open" for writing, in order to accept further entries.

The minimum reasonable size for each L0 trigger primitive packet is 8 bytes, including timestamp (4 bytes) and the bin contents of the histogram (4 bits x 8 time bins). In the "best case" all the hits from the same event appear on the same PP-FPGA, resulting in a histogram rate of 0.7 MHz per PP-FPGA (= 11 MHz/16), corresponding to 5.6 MB/s from each PP-FPGA to the SL-FPGA. The corresponding SL-FPGA outflow will be 22 MB/s in this case. In the "worst case" the hits from one event are spread over 8 PP-FPGAs maximum (those handling one RICH spot), resulting in a bandwidth of 44 MB/s from the PP-FPGA to the SL-FPGA. In this case the SL-FPGA has to merge the histograms belonging to the same event, and its outflow is 44 MB/s. Using the same logic, in the daisy chain architecture, the last Gigabit Ethernet link would be loaded with 88 MB/s both in the best and worst cases. For the proposed scheme a single GbE link can be used to bring 8 bytes trigger primitives to the L0TP.

To optimize the throughput from the different FPGAs, several geometrical configurations for connecting PMTs to TDCs were considered. Simulation results [99] show that there is no connection configuration in which it is possible to obtain significant reductions in data transfer rate without a significant efficiency loss. In order to maximize the rate of data sent from the TEL62 to the L0 Trigger Processor, L0 trigger primitive words corresponding to several events will be merged in a unique Ethernet frame (including few control words) exploiting jumbo frame support.

To provide L0 primitive time histograms, the SL-FPGA firmware will have to be modified with respect to the default, to include the logic for the histogram merging stage and – in case a daisy-chain connection configuration is chosen – also that to handle the Ethernet links also as receivers for the primitive data from the other boards.

A further use of the RICH in the L0 trigger is given by the possibility to have the analog OR of 8 channels being discriminated by the same NINO chip (precise hardware time alignment of the different channels is mandatory in this case). The set of 250 OR signals could be digitized in two TDC boards (one for each RICH spot) and following the scheme already described, a RICH multiplicity information could be extracted within a single TEL62. The cost for the generation of L0 trigger primitives with this scheme is that a dedicated TEL62 must be employed in addition to the four ones used for readout. Moreover, the signals corresponding to the hardware OR produced by the NINO chips must be recorded, in order to preserve the possibility to have full offline reproducibility of the trigger algorithms.

Finally, additional trigger information from the RICH would also be useful at L0 to generate control triggers or special triggers for other physics studies (*e.g.* triggers on high multiplicity of charged particles).

L1 trigger

The data selected by the L0 trigger will be processed by the L1 PCs in order to apply more sophisticated selection criteria at the single-detector level. Four GbE links with 50 MB/s of data could be handled, in principle, by a single PC equipped with a quad-port Gigabit card (*e.g.* INTEL PRO/1000 GT). An internal PC RAM of 4 GB is enough to store the whole burst, leaving space for buffering the results. The maximum L1 allowed latency (of the order of a few s) would have to be shared between the time used by the L1-RICH PC to compute the trigger primitives, and the time used by the L1 Trigger Processor PC to take the final decision on the event. Assuming a total spill duration of 10 s and a L1 trigger input rate of 1 MHz, the maximum computing time is 1 µs times the number of L1 processes running in parallel (corresponding to the number of L1 CPUs). At L1 the fundamental information to be computed for the RICH is the number of Cerenkov rings in the event (corresponding to the number of charged particles) and the geometrical parameters related to such rings. Preliminary tests show that a ring fit based on a maximum-likelihood approach takes more than 200 µs of computing time per ring on an Intel i7-950 CPU (3.07 GHz, 12 GB RAM), requiring at least 200 CPU cores to be used in order to allow a real-time processing of the data selected by the L0 trigger during the spill (assuming to have events with only one track, which is moreover not the case). Some gain in processing time can be obtained by either improving the fitting procedure or by decreasing the resolution requirements. In any case a multi-core L1 PC cluster must be considered for this purpose. The required size of the PC cluster will be determined after a more complete queue simulation, including other parameters like the bandwidth available from L1 to L2 and the processing capability of the L2 processing farm. Nevertheless, a solution based on the use of GPUs (video processor cores) is under investigation: such an approach dramatically increases the computing power of a single PC.

The following variables will be computed and used to apply selection cuts for the main decay mode and for ancillary triggers:

- N_ring : number of independent reconstructed rings;
- N_hits[i]: number of hits in the i-th ring;
- R_ring[i]: i-th ring radius;
- T_ring[i]: time of the charged particle associated to the i-th ring (average time of the hits participating to the fit);
- Chi2_ring[i]: Quality of i-th ring fit;
- X_ring[i], Y_ring[i]: position coordinates (x,y) of the center of i-th ring.

These quantities could be used to define L1 trigger selection criteria based on track multiplicity, ring quality and ring time, in order to reject events with a large number of reconstructed rings or events with a bad quality of the performed fits. At this stage cuts on the spatial distribution of the hits belonging to the same ring could also be applied, for instance to remove rings that are only partially contained on the RICH acceptance (arcs).

4.1.11 Charged Hodoscope (CHOD) System

The existing charged hodoscope (CHOD) from NA48 will be re-used in NA62. It is placed after the RICH detector to reject photo-nuclear interactions with the RICH material. Depending on the readout architecture the CHOD could also be employed -as backup for the RICH- to define the reference time of all the triggers with at least one track in the final state, (or to measure the RICH trigger efficiency). The CHOD would have the advantage that all channels are handled from a single TEL62 board.

The CHOD readout and trigger architecture could use the common NA62 building blocks. The light generated in plastic scintillator slabs are collected from photomultipliers. Thanks to the similarity of the CHOD signals with respect to those of the RICH, the analog signals will be discriminated in the same front-end described in the RICH chapter, exploiting the 32-channel preamplifier/discriminator boards housing 4 NINO chips each.

Readout

The readout of the digital signals will be based on the common NA62 TDC system described in section 4.1.5. The CHOD has 128 channels, which are arranged in two planes and they could be read out by two TDC boards[54] leaving also some spare channels. For an 11 MHz maximum rate in the CHOD geometrical acceptance, the output rate of each TDC board will be, on average, less than 40 MB/s, which can easily be managed by the corresponding PP-FPGA. The particle hit rate isn't uniform over the hodoscope surface, and some care must be taken in connecting the CHOD channels to TDCs, in order to avoid excessive rates on a single TDC board. In any case, using a single TEL62, the readout bandwidth will be around 100 MB/s, which can be handled by 1 or 2 Gigabit Ethernet links.

Due to the length and the age of the scintillators (slabs of up to 1.21m) the intrinsic time resolution of a single detector plane remains limited to a range between 3 and 5 ns. The offline time resolution using both planes was measured to be 200 ps [15].

L0 trigger

If the CHOD is used as time reference for the L0 trigger the online time resolution should be at the level of 1 ns. In this case the time information from both detector planes has to be used online. Slewing and propagation time corrections should be applied directly in the PP-FPGAs (assuming enough information is available there).

The L0 trigger primitive data are sent to the L0 Trigger Processor in 8-byte frames including detector ID, timestamp and "fine time". To exploit the time resolution of the CHOD (≈1 ns online), the 8 bits reserved to the "fine time" will encode the following information:

- the 5 most significant bits will identify the fine time window (within the 25ns period defined by the timestamp) where the hit(s) occurs with a resolution of 25/32 ns (0.78 ns);
- the 3 less significant bits will contain the hit multiplicity within the fine time window, (*e.g.* 1hit, 2hits, 3 hits, ≥4 hits).

If one or more hits (even if generated by the same decay) happen to occur in adjacent fine time windows, a second frame is generated and sent to the L0 Trigger Processor.

For a rate of 11 MHz in the CHOD acceptance, the L0 primitive data rate from the TEL62 to the L0 Trigger Processor will be 88 MB/s, and such bandwidth could be handled by one of the four available GbE links. In order to reduce the amount of the data to be sent, several frames could be merged into a single packet, allowing including the detector ID only one time (latency issues should be taken into account in this merging, of course).

[54] Ultimately even one TDC board would be sufficient.

Since the online time resolution will be fundamental to trigger on the signal events with good efficiency (this is not trivial to be obtained, given the presence of veto conditions in the L0 trigger algorithm), the time of all counters must be aligned with good precision. Besides the intrinsic offset correction capability of the TDC themselves, this aim can be reached including on the PP-FPGA a look-up table with pre-calculated offsets to be subtracted from the incoming times measured by the TDCs.

The opportunity to measure the L0 trigger time with a single additional TEL62 reading all the counters of the charged hodoscope and the 250 analog OR signals available from the RICH is under investigation. In this case the information from two independent detectors could be merged at L0, contributing to a better online definition of the L0 trigger time.

L1 processing and trigger

Once the data collected by the CHOD will reach the L1 PCs, both the time and the multiplicity of the events can be re-evaluated in software, and given the better resolution on the time measurement ensured by the presence of the complete information, tighter cuts could be applied at this level. For instance, by exploiting the available information on the time-over-threshold, a slewing correction could be applied to the hit time measured by the TDCs, to correct for the bias introduced by the fixed threshold of the discriminator chips.

The size of the PC farm required will be smaller with respect to the RICH case, due to the simpler reconstruction algorithms.

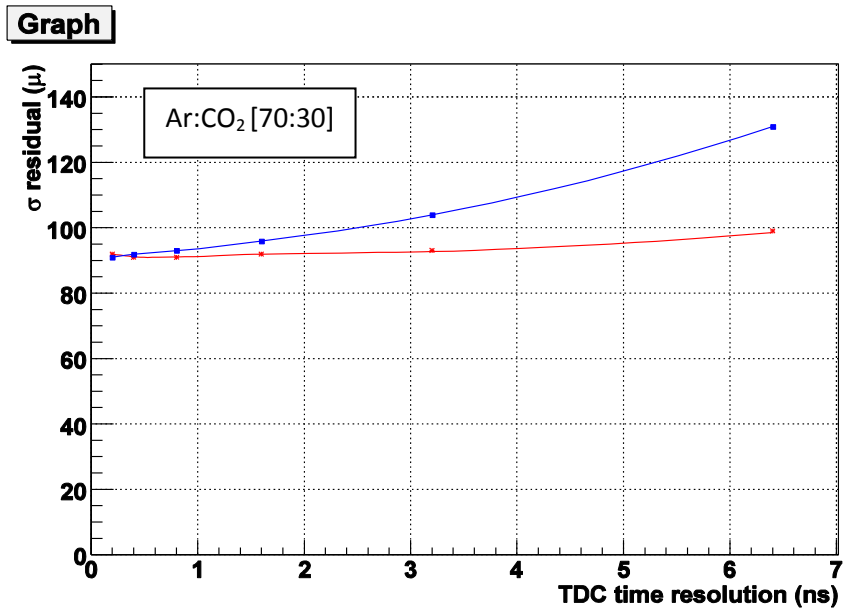

Figure 321: Straw resolution versus TDC time binning
(blue: unknown space-time dependence; red: known space-time dependence).

4.1.12 STRAW System

The Straw readout electronics should provide track data to the NA62 DAQ system in required format, perform online monitoring to check data quality, control the front-end electronics, and possibly trigger on a single charged particle or veto on multiple charged tracks in an event. For data extraction, front-end control and online monitoring a straw detector specific module is used: the Straw Readout Board (SRB). For data collection and event building, handling and selection we will use the readout board [100], which is common to the majority of detectors in NA62.

Time to digital conversion

To understand the required drift time resolution, a dedicated study has been performed with different TDC time bin steps. The targeted position resolution was specified to 130 μm (see Straw Tracker chapter). The result indicates (see Figure 321) that for a straw with known position-time dependence, even a 6 ns time bin of TDC (time to digital converter) would be sufficient, when using an Ar/CO_2 gas mixture.

However, other constraints led to fix the TDC time-binning to a maximum of 3 ns; these include mainly the matching to the GTK, where a time binning of 6ns in the straw detector would require opening a too large window for track fitting. On the lower side, using a very small TDC bin is useless as the space resolution is dominated by multiple scattering of particles (see Straw Tracker chapter). This is the reason why the use of the common TDC system developed for NA62, with much higher time resolution is not needed.

We plan to implement the TDC directly in a FPGA (Field Programmable Gate Array), together with other readout functions. A preliminary study shows that one can achieve 1 ns resolution (1.6 ns bin) with a cost-effective version of FPGA.

Readout

The readout architecture follows the detector partitioning. The smallest readout unit is a module of 16 straws served by 1 front-end board. 30 such boards form a view, and each chamber has 4 views (x,y,u,v). The whole detector has 4 chambers. The data from 15 front-end boards (half a view) is collected by one SRB, which also provides a control for thresholds and test pulses. The chamber (4 views) is thus served by 8 SRBs housed in one VME 9U crate, positioned about 5 meters from the detector. As 2 chambers are within a short distance, it is possible to use a single full VME 9U crate (21 slots) for housing SRBs for two chambers (16 boards). SRBs will receive precise system clock, timing information and control from the common NA62 TTC system, and will time-align the data from the straws by attaching the required timestamps. The data sent from the SRBs to the TEL62 boards will have a fixed format of 40 bits per hit:

Timestamp 30 bit	Fine time 4 bit	Leading/trailing edge 1 bit	Straw ID 4 bit	Control 1 bit

This information will be re-arranged within the TEL62 to provide the final readout data.

For the data extraction and positioning of TDC two approaches are considered (Figure 322).

(1) TDCs are placed on the SRB. For data extraction, control and services for the front-end board 5 meters long halogen-free SCSI cables are used. As these cables are rather heavy and rigid, appropriate attention must be paid to fixation and support, mainly close to the front-end board connection.

(2) TDCs are placed directly on the front-end board. For data extraction, control and services standard halogen-free Ethernet cables are used. Such a solution is mechanically more stable and does not provide a big strain on the PCB connector. Also, such a solution would be advantageous from a noise point of view, as the signal spectra from high-speed serial links lies outside the sensitive bandwidth of front-end electronics. In this case, a SRB collects data over a high-speed serial link running at 320Mbits/s (or more). If needed, for boards close to the beam with high data rate, two of them can be used. There is one dedicated line for clock and one for the control. The control line can use Manchester encoding as the number of bits to transmit is very low; data lines would use 8b/10b encoding to ensure a return-to-zero bit stream.

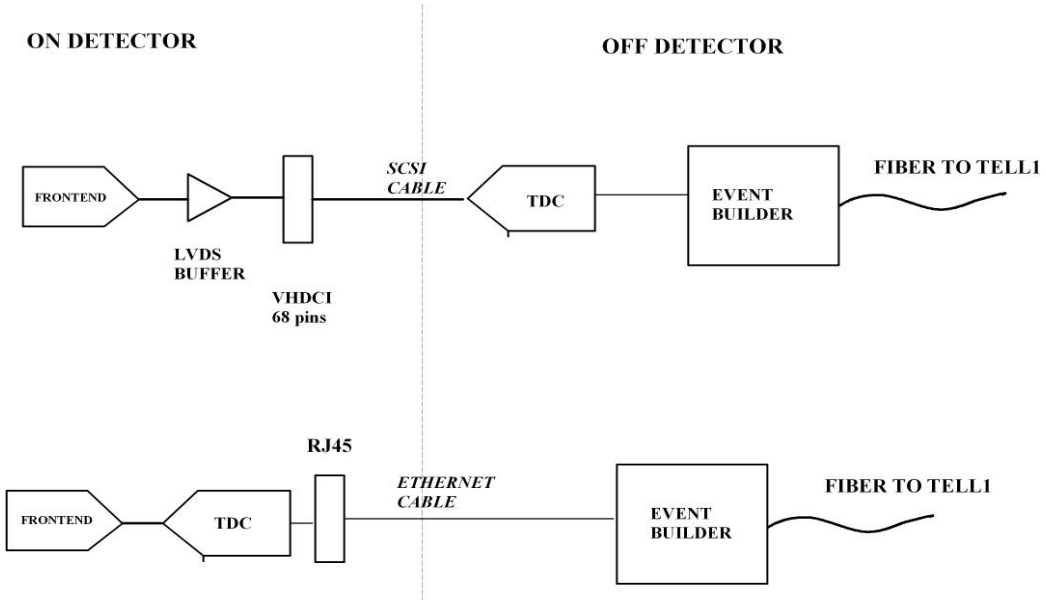

(3)

Figure 322: Two possible architectures for the Straw TDC readout.

All 16 SRBs from one VME crate transmit formatted data to one TEL62 board, so for the whole detector two TEL62 boards are required. For the data transmission from SRB to TEL62 either copper or optical links might be used (see Figure 323).

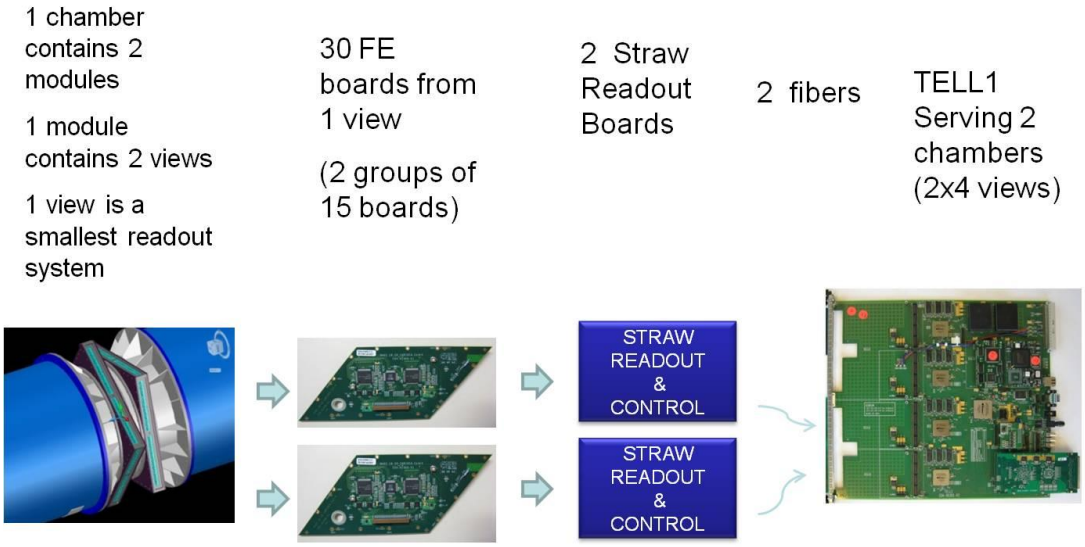

Figure 323: Straw readout scheme for one View (1/2 a Chamber).

Each TEL62 board collects data from 16 SRBs. If optical links are used, the double-width mezzanine board developed for the TEL62 by the LHCb collaboration with 2x12 optical links can be used. If copper links are used, a dedicated TEL62 mezzanine receiver board will have to be developed for this purpose. The firmware on the TEL62 and its control software will be Straw detector specific: its main tasks will be to check the integrity of incoming data, building events from matched timestamps and event management; in case the STRAW detector participates to the L0 trigger formation, the evaluation of trigger primitives would also be handled there.

Upon arrival of the L0 trigger, it sends the requested event with the matched timestamp to the DAQ system. The data sent from the TEL62 board corresponds to hits in straws from two chambers, *i.e.* eight views. When all straws along the particle track generate a signal, the amount of data sent per event is:

$$40\ [\text{bits}] * 4\ [\text{views}] * 2\ [\text{chambers}] * 4\ [\text{straws in view}] * 2\ [\text{lead + trail edge}] =$$
$$= 2688\ \text{bits/event}$$

If neither of the above described approaches works, a fallback solution is identified using the common system of TEL62 boards with mezzanine TDCs. This solution will require a board between front-end and TEL62 for rearranging data lines, providing front-end control, distributing power supply and enabling online and DCS monitoring. The first estimate gives up to four such boards per view, 64 boards for the whole detector; most probably in 9U VME format using VME interface. A TEL62 board with four mezzanine TDCs can handle 512 channels. As one view of a chamber contains 30 front-end boards with 16 channels, one TEL62 board per chamber view is needed, and 16 TEL62 boards for the whole detector.

Expected data rates

The average rate of particles per straw is estimated to be 33 kHz. We expect to have 1 signal from each particle, of which both trailing and leading edges times have to be measured. Thus the average data rate per straw is 40 bits x 2 x 33 kHz = 2.6 Mbit/s, or 42 Mbit/s per front-end board serving 16 straws.

There are few straws close to the beam with much higher particle rates, up to 500 kHz. The maximum data rate will be 40 Mbit/s for such straws. To this data rate one should add extra data caused by noise, spurious hits, etc., which must be kept as low as possible.

The data rate from one TEL62 board to the PCs depends on the L0 trigger rate, *e.g.* for 100 kHz the data rate is 270 Mbit/s. For such a rate, one gigabit Ethernet link would be sufficient, but for the design value of 1MHz L0 trigger rate 3 links are needed. If TEL62 boards could not cope with this data rate, it is possible to add data buffer on SRB boards and transmit data from SRBs to TELL1 only on L0. In that case, building and event management would be pushed back to SRB.

Clock distribution, control and timing of the readout electronics

The Straw detector will use standard NA62 TTC modules, LTU and TTCex, for control and clock distribution. One set of modules is sufficient for the whole detector as there are two VME crates and one TEL62 crate as destinations. TEL62 boards will use the full TTC protocol, while SRB boards need only the clock and “Start Of Burst” signal for synchronization. SRB boards will provide fine tuning of clock delays for timing of the detector. Adjustable delays are needed due to the spread of propagation time through components and cables, and time-of-flight of particles along the beam.

Trigger and vetoing

Both leading and trailing edges from the straws’ signals provide useful information. While from the leading edge time one can obtain the precise crossing position of the particle through the straw, the trailing edges occur at the same time for all straws hit by the same particle, independently from their crossing distance from the wire. The trailing edge time is used as a validation of straws on a track, thus reducing false hits and improving track fitting.

It is also rather straightforward to use the trailing edges times for a fast hardware trigger or veto. For this purpose all the views in the chamber are partitioned into corridors, the size of the partition depending on trigger granularity. Each view has 4 straw planes and the geometry of straws placement guarantees that at least 2 straws are hit by each passing particle. Thus one should expect to have from a minimum of 2 up to 4 straws with signals having the same trailing edge time (see Figure 324). A View Trigger Logic (VTL) opens a time window, which should cover the spread of trailing edges from one view, and if there are at least 2 straws in the window with the same trailing edge it sends the corridor number and timestamp to the Chamber Trigger Logic (CTL). The CTL collects numbers from all views and opens a timing window of the order of 7 ns to compensate for signal propagation time in straws from different views. If there are hits from different views, it sends coordinates and timestamp to the central trigger. Using coordinates and timestamp, a corresponding particle could be searched for in other sub-detectors. Such information could be used for generating: (a) a trigger, when there is only one track; (b) a veto, when there is more than 1 track.

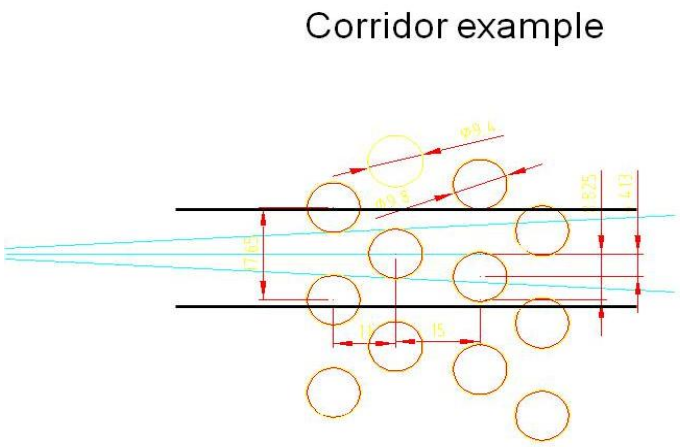

Figure 324: Smallest possible track corridor in 1 view.

4.1.13 MUV System

In total (with MUV3 option A), 560 channels have to be read-out (Table 59) and it is foreseen to use the standard NA62 read-out approaches. In this respect, for MUV1 and MUV2 stations two data acquisition choices exist:

1. **The "LAV" readout**: after suitable discrimination, the PMT signals are routed to TDC boards, housed on the TEL62 boards. The signal amplitude is determined by measuring the time-over-threshold in the TDCs. To cover the dynamic range of the analogue signals, two digital lines are used, with different thresholds. The corrected time and signal amplitude are reconstructed from four measured times (two leading and two trailing) per PMT. With 128 channels per TDC board and a maximum of 4 TDC boards per TEL62 module, one or two TEL62 boards could be used to read MUV1 and MUV2.

2. **The "LKr" readout**: the CREAM readout of the LKr calorimeter is based on flash ADCs with 40 MHz sampling, 14 bit resolution, large data buffering and optional zero suppression. To be able to interface the CREAM read-out directly the MUV1+2 signals will have to be shaped to a comparable timing than those of the LKr, *i.e.* 20 ns rise-time and 2.7 µs fall-time with maximum amplitude of 1 V. The signals are than re-shaped, so that the FADC receives a signal with 40ns rise-time and 70ns FWHM. One additional VME crate with 6 (3) CREAM modules would be adequate to acquire the data from MUV1 (MUV2)

The MUV3 system is used as fast muon veto, mostly for the L0 trigger, and only timing information is needed. The signals are discriminated with existing Constant Fraction Discriminators[55] compensating jitter from amplitude fluctuations. The time measurements are done with the standard TEL62/HPTDC readout boards. One board with 512 channels can receive the data from the full MUV3 station.

[55] From the NA48 AKL detectors.

Table 59 Number of read-out channels of the MUV detector.

Module	Number of Channels
MUV1	176
MUV2	88
MUV3 Design A (B)	296 (252)
Total	**560 (516)**

4.1.14 Charged ANTI (CHANTI) System

The conceptual scheme of the CHANTI readout is shown in Figure 325.

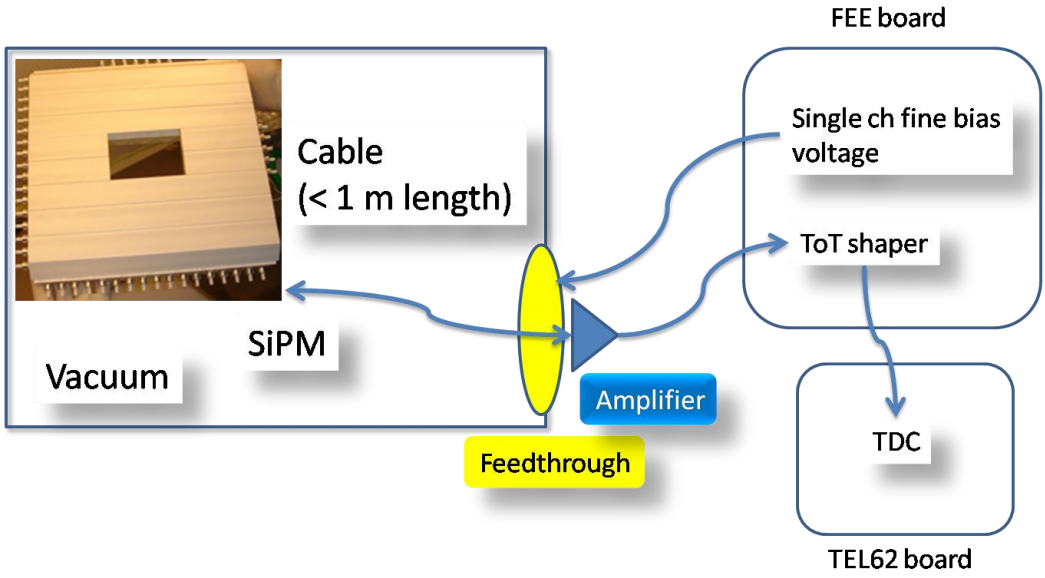

Figure 325 Schematic view of the CHANTI readout.

Each CHANTI station is composed by 46 scintillator bars, for a total number of 276 channels. After amplification, signals are processed by FEE boards similar to the one used for the LAV system.

The 9 FEE boards (32 channels each), will have to provide, for each channel:

- a way to control the V_{bias} with O(10 mV) accuracy
- a fast, DC coupled, conversion to a Time Over Threshold-LVDS signal output
- a temperature and/or a dark current monitor for slow control adjustment of the V_{bias}

Thresholds and V_{bias} settings will be controlled using the CANOpen standard. Since the dynamical range of the signals is expected to be small (the detector is essentially sensitive to MIPS) a single threshold will be used.

The LVDS output will be directly sent to a standard TEL62 system equipped with TDC boards for both leading and trailing edge measurement. One TEL62 board equipped with three TDC boards will be enough for the entire detector, allowing also for a good number of spare channels.

L1 trigger

The CHANTI is not providing primitives for the L0 trigger.

An evaluation of the maximum L0 data rate can be done as follows. The highest multiplicity in the CHANTI is expected from the beam halo events crossing all six stations. If fully efficient on these events, the system will give at most 4 x 6 = 24 hits or 192 bytes per event. These events occur at 1 MHz rate. Inelastic interactions in the GTK will be detected at approximately the same rate, with comparable hit multiplicity. One can safely estimate a maximum of 200 bytes at 2 MHz. Assuming a 1 MHz L0 trigger rate and a readout time window of O(100 ns) one expects O(200 kHz) of such events in coincidence with a trigger. This would generate a data rate of about 40 MB/s well below the TEL62 specifications.

After zero suppression done at TEL62 level, the L1 farm checks data integrity and performs further data reduction. Raw times will be corrected for individual channel calibration constants and corrected for time propagation along fibres by using x-y coincidence algorithms.

Moreover, time slewing corrections could be implemented by exploiting the ToT-amplitude correlation. This will allow a tighter time window cut, which will be optimized according to the maximum allowed L1 rate. If necessary beam halo events could be recognized at this level by means of appropriate multiplicity and collinearity algorithms.

One multi-core PC farm should be envisaged for L1 processing. If compared to the RICH system, given the less demanding L1 algorithm complexity and the 1 MHz expected L0 rate, the total number of CPU cores will be considerably less.

4.1.15 LKr (LKr) Readout system

The 13’248 channels of the LKr electromagnetic calorimeter are continuously digitized with 40 MHz flash ADCs. The calorimeter is therefore the largest single data producer in NA62, generating 800 GB/s of raw data.

The initial current is derived from the charge measured by a preamplifier mounted inside the cryostat at liquid Kr temperature and connected to the anode electrode by a blocking capacitor. The integration time constant chosen for the charge preamplifier is 150 ns. The signal from the preamplifier is transmitted to a combined receiver and differential line driver mounted outside the calorimeter close to the signal feed-through connectors. The receiver amplifies the preamplifier signal and performs a pole-zero cancellation. Those two elements were built for the NA48 experiment and will remain untouched for NA62. The dynamic range (50 GeV of deposited energy) corresponds to an amplitude of ±1 V (on a 100 ohm termination) at the input of the digitizer electronics. The required signal to noise ratio is 15000 to 1. The LKr readout chain scheme is sketched in Figure 326.

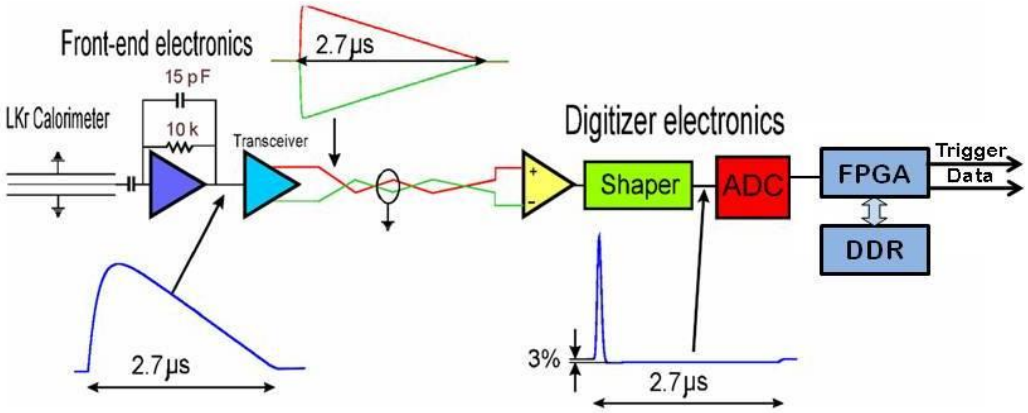

Figure 326 LKr readout scheme.

The Calorimeter REAdout Module (CREAM) is the new back-end part of the NA62 LKr data acquisition chain. CREAMs provide 40 MHz sampling of 13248 calorimeter channels, data buffering during the SPS spill, optional zero suppression, and programmable digital trigger sums for the LKr L0 trigger system. The digitization has 14-bit resolution, and the event readout rate is up to 1 MHz.

The Timing, Trigger, and Control (TTC) system [87] developed for the LHC experiments is used by all NA62 sub-detectors. The LKr-specific TTC and CREAM modules are single-width 6U VME64x units. One TTC module and 16 CREAM modules are housed in one VME64x crate and controlled by a VME-PCI bridge master, which is used for configuration, monitoring and test purposes. The entire LKr back-end system is composed of 28 readout crates.

LKr clock distribution

High-speed applications using ultra-fast data converters often require an extremely clean clock signal to make sure that an external clock source does not contribute with undesired noise to the overall dynamic performance of the system.

The experiment TTC system provides distribution of 40.087 MHz clock, Level 0 trigger information, broadcasts and individually-addresses control signals. The system is used down to the CREAM crates. One LKr-TTC module interfaces TTC commands to all the boards in one crate through the VME backplane. A CERN standard TTCrq [93] mezzanine board is used to decode and convert optical inputs. This board incorporates a TTCrx ASIC chip [91] [92] [100] and a QPLL [92], which provide three programmable output clocks (40, 60/80 and 120/160 MHz) with a time jitter below 50ps. The measurements done by the ATLAS LArg group show that the TTCrx clock signal noise (related to trigger and B-channel information sent after every trigger) can be reduced below 20 ps by turning off all B-channel activity on the chip providing the ADC sampling clock. In addition, the TTCrx has an option to deskew the clock signals in step of 104 ps, which might be useful for the data alignment and calibration. For a good performance is it planned to use two TTCrx chips, one to provide a clean clock by switching off all the B-channel activity, and the other to derive the trigger information

The VME backplane P0 connector is used to deliver all TTC signals from the LKr-TTC module to all CREAMs in the crate via LVDS links.

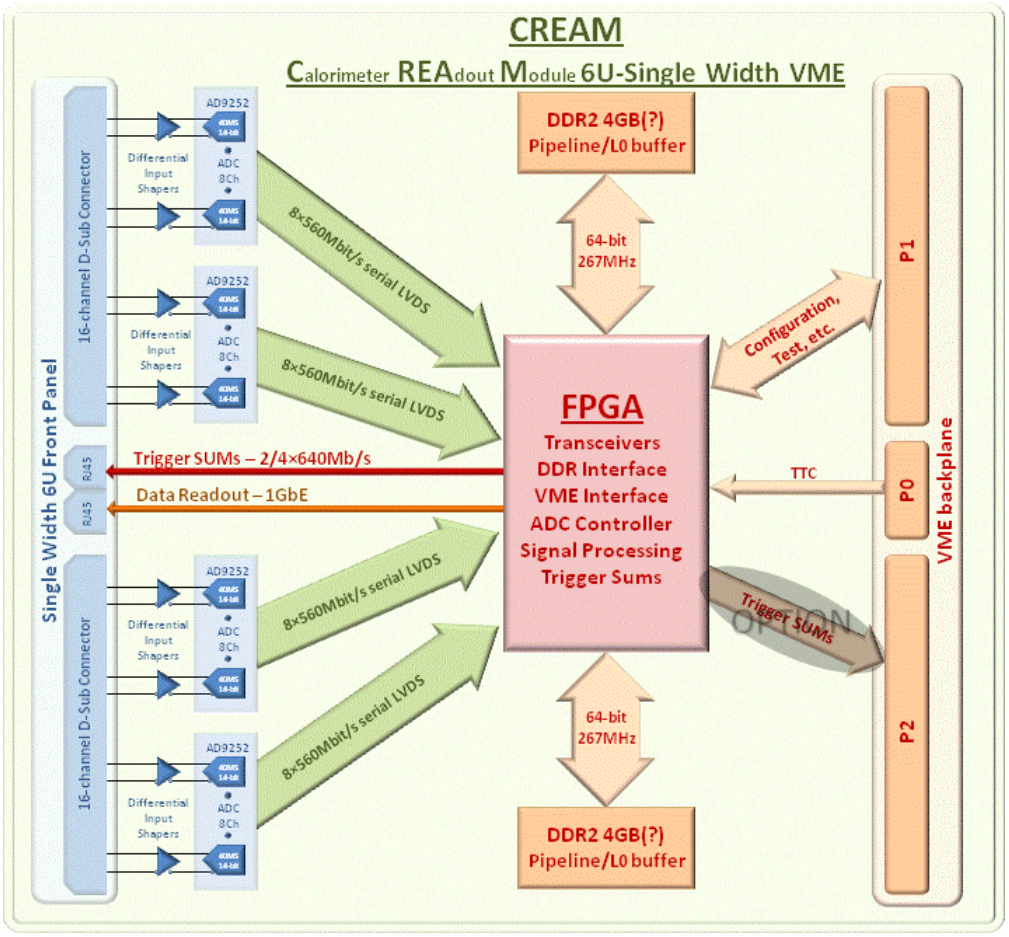

Figure 327: CREAM block diagram

4.1.15.1 The Calorimeter REAdout Module (CREAM)

The CREAM is a 1-slot wide 6U VME64x module. It houses 32 calorimeter channels (4 octal ADC chips).

Each ADC readout channel has a dedicated serial LVDS link to a FPGA serving 32 channels. The FPGA acquires ADC raw data at 40 MHz, performs primary treatment-formatting, and interfaces data to a DDR2 SODIMM memory module for temporary storage. The block diagram of the module is shown in Figure 327.

The VME bus (P1) is used for configuration, monitoring and test purposes. Experiment triggers and clocks are supplied via the dedicated VME backplane P0 connector. The trigger sums are sent to the trigger system via a front panel RJ45 connector. The module should be able to operate in standalone mode using on-board generated clocks and triggers, and the event memories should be writable and readable via standard VME facilities.

The CREAM data and trigger processing requirements will evolve during the experiment running, so it is expected to have upgrades of the FPGA firmware.

Analog input circuit

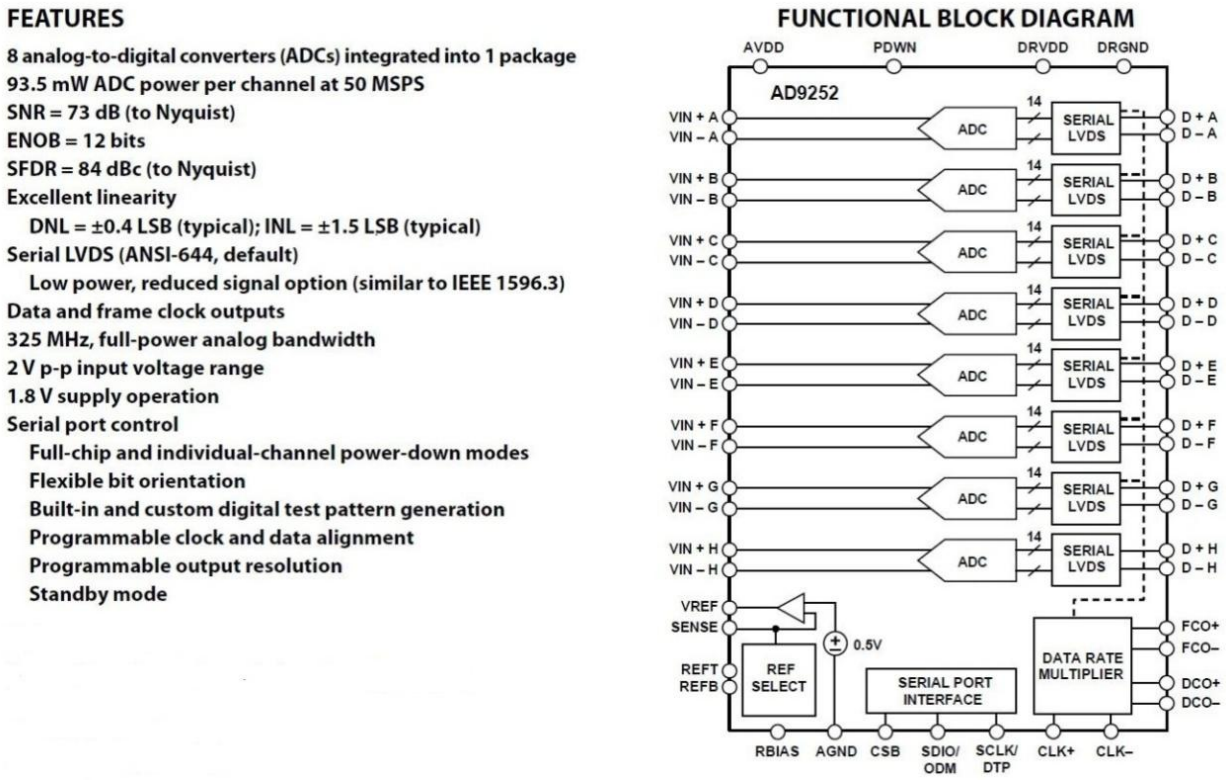

FEATURES

8 analog-to-digital converters (ADCs) integrated into 1 package
93.5 mW ADC power per channel at 50 MSPS
SNR = 73 dB (to Nyquist)
ENOB = 12 bits
SFDR = 84 dBc (to Nyquist)
Excellent linearity
DNL = ±0.4 LSB (typical); INL = ±1.5 LSB (typical)
Serial LVDS (ANSI-644, default)
Low power, reduced signal option (similar to IEEE 1596.3)
Data and frame clock outputs
325 MHz, full-power analog bandwidth
2 V p-p input voltage range
1.8 V supply operation
Serial port control
Full-chip and individual-channel power-down modes
Flexible bit orientation
Built-in and custom digital test pattern generation
Programmable clock and data alignment
Programmable output resolution
Standby mode

Figure 328 AD9252 ADC diagram and features.

The signal at the input of the CREAM module has a 20 ns rise-time, a 2.7 µs fall-time and a ±1 V maximum amplitude. Each of the 32 channels consists of a differential line receiver and a pulse shaper. The shaped ADC input is a differential semi-Gaussian signal with a 40 ns rise time and a 70 ns FWHM.

ADC

The AD9252 ADC [101] from Analog Devices satisfies the experiment requirements. It is an octal, 14-bit, 50 MSPS ADC with an on-chip sample-and-hold circuit and one serial data link per channel. The circuit block diagram and main features are shown in Figure 328.

The ADC requires a single 1.8 V power supply and LVPECL-/ CMOS-/LVDS-compatible sample rate clock for full performance operation. No external reference or driver components are required for many applications.

The ADC contains several features designed to maximize flexibility and minimize system cost, such as programmable clock and data alignment and programmable digital test pattern generation. The available digital test patterns include built-in deterministic and pseudorandom patterns, along with custom user-defined test patterns entered via the Serial Port Interface (SPI).

FPGA controllers

Most of the CREAM functionality is implemented within two FPGAs. The first one handles the VME-bus protocol, the module configuration and provides internal communication with the other circuits on board. The “main” FPGA defines the mode of operation and configures the ADCs via SPI. It handles the ADC serial output links, reads, formats and locally stores the data in a circular buffer with a depth corresponding to at least 1ms (L0 trigger latency).

The minimum 8 data samples from each of the 14k calorimeter cells are read out from the CREAMs on every **Level 1** trigger (and not L0 as most other sub-detectors), at the rate of up to 100 kHz. Each sample consists of 14 bits, for a total of about 200 Gbit/s of data for the entire calorimeter. This data volume is demanding in term of the bandwidth of the subsequent parts of the DAQ system, as well as in term of storage capacity, therefore a zero suppression scheme is foreseen.

In parallel to the readout, the CREAM performs the summation of the selected channel samples, and sends the resulting “Super Cell” data to the LKr L0 trigger system.

Upon receipt of the L0 trigger signal, the relevant control part extracts the corresponding 8 (at least) data samples from the “pipeline” or “circular” part of the buffer and stores them in a readout buffer large enough to store up to 1.6×10^7 64-byte data samples (corresponding to a 1MHz trigger rate during twice the 1 s latency time of the L1 trigger). Two DDR2 interfaces provide the control for two SODIMM memory modules of 2 GB each, where both linear and circular buffers are implemented. With the above figures, the size of the readout buffer is 1 GB.

A Gigabit Ethernet link is managed by the FPGA for the CREAM data readout at the L1 trigger reception, and to receive L1 triggers.

L0 and L1 data storage DDR2 memory

The DDR2 memory should be virtually split into two parts. The first one serves as a pipeline (or “circular buffer”), and insures continuous data storage at 40 MHz with programmable latency (up to 1ms). The second part of the memory is used as a L0 trigger buffer. The L0 is effectively a random signal with about 1 MHz average rate, occurring a fixed time (latency) after the event took place. If a L0 command is received by the module, the corresponding data is extracted from the pipeline. A programmable number of adjacent time slices (stored in the pipeline at an offset corresponding to the L0 time) are copied from the pipeline into the second part of the memory (L0 buffer) and then read out from the module after a positive L1 trigger is received on the GbE link.

Trigger sums

During the data acquisition, digitised signals from the selected channel are summed up to build a Trigger Sum (Super Cell) to be sent to the LKr L0 trigger system. The selection of the channels contributing to a particular Super Cell as well as the number of Cells is programmable: in order to match with the foreseen granularity of the LKr L0 trigger system (section 4.1.15.5), Super-Cells composed of either 4x4 or 2x8 cells have been considered. The readout of the sums data is performed via commercial serializers and standard Ethernet cable (cat.5) at 640 Mbit/s per link. The maximum number of links is 4, allowing up to 4 Super Cells from each CREAM module.

4.1.15.2 Data handling inside the FPGA

Zero suppression

The data copying mechanism described above will also allow to read interesting events without zero suppression, in order to have, at a later stage of the analysis, all the original data available. This scenario will allow for a non-zero suppressed readout while keeping a simpler network structure and a lower transfer rate. Reading only L1-selected data reduces the data rate by a factor at least 10, since the L1 rate is expected to be O(100 kHz).

An additional data rate reduction is possible by applying a simple zero suppression algorithm to individual channels. Since for each event a large fraction of channels will only contain pedestal counts, one can discard channels where the difference between the maximum and the minimum value of the samples is below a predefined value (programmable and possibly different for each channel). This will add at least a factor 8 of data reduction, and it will reduce and simplify the configuration of the LKr event building farm.

Pre-processing at this stage could flag each sample, using the two free bits to pack a sample into a 16 bit data word, e.g. to mark whether it is consistent with a pedestal, or if it is the maximum sample, to speed up later processing. This could help in speeding up further processing, allowing to quickly search for interesting channels in the early phases of the processing, before any refined reconstruction. Header and packet formatting info should be added as well.

Possible additional computations

A more sophisticated use of the readout using the processing power of the FPGA could be investigated. Repetitive calculations on the data samples are can be performed in the FPGAs rather than in the PC farm, thereby improving the throughput. A digital filter in the chain could perform computations on many samples and produce a result for each time slice. As an example, this type of filter could make raw computations of energy and time on each 25 ns sample, or it could also average the 8 samples and possibly flag the corresponding data as “pedestal”, using the 2 extra bits available in the DDR2 memory for each channel.

4.1.15.3 Network connections

Network architecture

We are working in the hypothesis that there is no need to transfer information at the single cell layer to the DAQ at the L0 rate. In the approach described above, selected L0 fragments are stored in an intermediate buffer, waiting for L1 to be transferred to the event building PCs. This hypothesis assumes also that the L1 rate is around 100 kHz. With this rate, no zero suppression and 10 s readout time, the aggregated throughput from each CREAM is 400 Mbit/s. Multiplexing the links from 16 CREAMs in local switches to one 10 GbE link does not limit the transfer rate: 28 x 10GbE links are used from the experimental hall to the NA62 control room. There a large switch (O(128) ports) houses the

28 links, a farm of multi-core PCs with 10GbE cards and a suitable number of 10GbE links for the transfer of complete events to the final NA62 event building farm.

The large switch in the NA62 control room is a part of the larger NA62 network infrastructure switch, both for cost reasons and for the capability of transferring LKr built events to the final event building farm using the switch backplane. The same switch could host the LKr trigger readout machines and additional service connections to the LKr readout control PC.

This implementation is dimensioned to allow a full-rate non-zero suppressed data transfer. The baseline mode of operation is however to apply a mild zero suppression to the events, except for a fraction of those (random events, calibrations, downscaled control triggers). In this scenario, one can gain on the number of event builder PCs. Special non-zero suppressed runs could still be performed at a reduced rate.

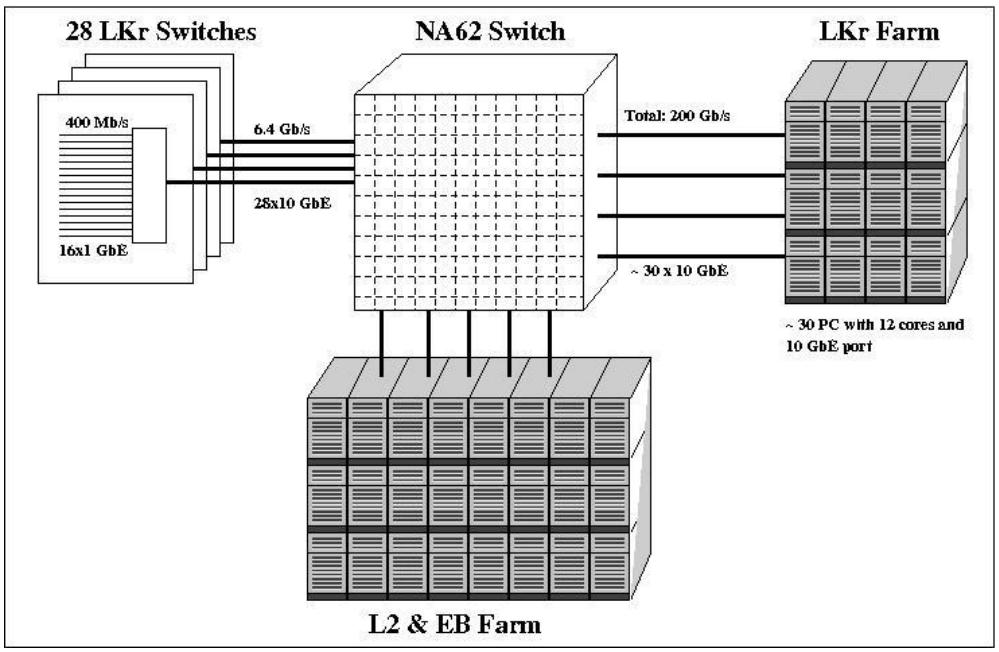

Figure 329: Layout of network connections and LKr PC farms.

PC farm

The main function of the LKr PC farm (see Figure 329) is that of collecting the event fragments from the 432 CREAM boards and building one single LKr event out of them. This implies that for each L1 readout trigger (or a block of them), the LKr control PC will also distribute the address of the destination PC (as distributed from the L1 Trigger Processor). As additional tasks, the LKr PC farm could perform operations like “halo-expansion” of clusters (*i.e.* the flagging of cells which must be readout, independently from their sample content, because they are close to a region with cells above some threshold) and zero suppression (if not already done in the FPGA).

Assuming a figure of 4 ms to build a LKr event, about 30 PCs (12 cores each) are required to do all the LKr event building at an input rate of 100 KHz. Increasing this number will give more CPU power for the additional tasks described above or for pre-computation for the L2 trigger.

4.1.15.4 Temporary Readout System

This system will be used until the end of 2012 for beam surveys and technical runs if needed.

The NA48 LKr readout system, based on CPD modules [71], was consolidated in 2008 to replace the optical links, the Data Concentrator and the VME RIO readout for maintenance reasons. Data from CPDs are read and stored in DDR2 memories in the Smart Link Modules (SLM) [102] and then read via Ethernet to a series of PCs. All activity on the LKr before the installation of the new readout described in this document relies on this CPD-SLM readout. The system was shown to be capable of reading non-zero suppressed data up to a trigger rate of 10 kHz. For operation of the LKr before the installation of the new NA62 system, the consolidation of the power supplies is needed, as well as the development of a trigger interface to connect the new NA62 trigger system with the old RIO-based LKr trigger system.

Existing hardware

The heart of the new system is the SLM. Each SLM receives data from the CPDs via a parallel 20 bit LVDS path implemented on multiple standard RJ45 cables. The readout protocol from the CPDs is handled by the SLM FPGA, and each event is stored in a 1 GB DDR2 memory. The interface with the outside word is via an Ethernet connection. The readout protocol is based on basic MAC packet transfer, and special protocol commands are implemented for the interaction with a PC. The typical sequence of network operations is the following:

- initialize the SLM at the start of spill;
- read the number of collected events at the end of spill;
- loop on the number of events, reading them one after the other and process them.

An internal timestamp is also implemented, with a counter incremented by a 1 MHz clock: this was intended to match fragments from different SLMs, but differences between quartz crystal frequencies in each SLM prevent the use of this feature.

It should be remarked that the NA48 LKr readout system requires the NA48 clock distribution system to be running, which should therefore be maintained for the interim period, and guaranteed to be phase-synchronous with the TTC one (*e.g.* by generating both from the same master source).

NA62 to NA48 trigger interface

In the SLM system the trigger distribution is essentially the one of NA48 [103]. The trigger requests and information are received from the NA48 Trigger Supervisor, which will be not maintained for NA62. The 64-bit information arrives to the RIO-TIC processor through one differential PECL pair, and is managed by the TAXI chip (AMD Am7969). The NA62 L0 Trigger Processor (replacing the NA48 Trigger Supervisor) communicates decisions and information through the TTC system. To allow interoperability, the trigger receiver stage of the old LKr readout system needs to be redesigned, in order to feed the CPD system with the TTC information. The easier way to do this is to use a TEL62 board, which is equipped with a TTC receiver (TTCrx), and build a new daughter-card for it housing a TAXI chip, implementing a protocol conversion (see Figure 330). The PP-FPGA and SL-FPGA in the TEL62 are used to re-encode the information, adding or modifying the data format to match the protocol specifications. In addition this new daughter-card contains other inputs, to manage XOFF signals (back-pressure control signals) from the read-out boards and to receive calibration triggers.

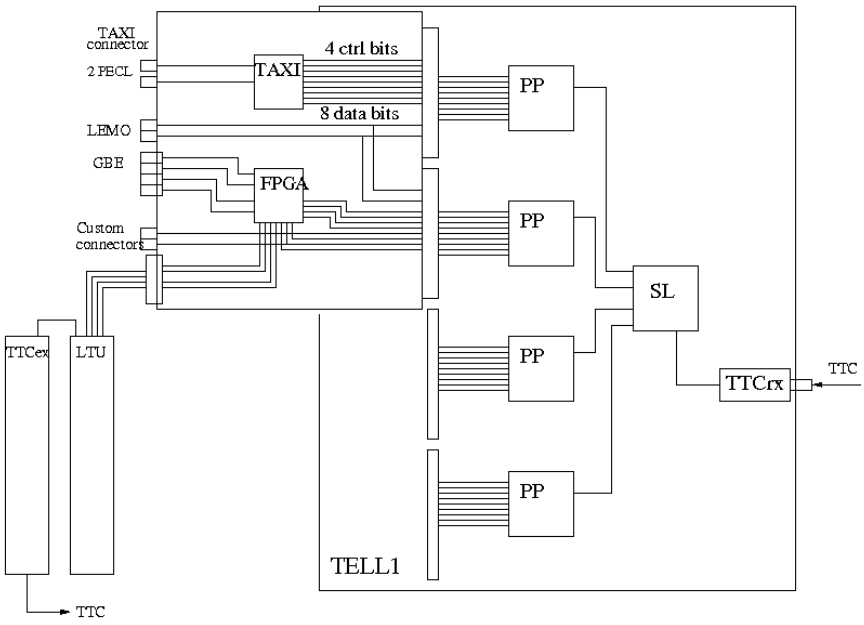

Figure 330: Scheme of the trigger interface for the SLMs.

Furthermore, in order to provide a small-scale standalone trigger supervisor replacement system, the new daughter-card hosts 4 Gigabit Ethernet cables to possibly collect trigger primitives from other detectors, and other inputs for direct connection of reference or trigger counters. The TEL62 FPGAs are used to implement the logic decisions. This functionality will be very useful to test the new trigger system and to provide a working system for test beams in the early phase of the data taking or during technical runs. The output from the card is in a format suitable for the LTU module, to provide the TTC trigger distribution to any NA62 sub-systems.

Farming

The current LKr readout PC farm is composed of 12 Super Micro SC808T-980 PCs and of 2 (with a possible connection of two more) Elonex power servers. Each PC is equipped with one 4-port GbE interface card. The existing software is able to read at end of burst all the data from the SLMs, to re-format it, to optionally apply a zero suppression algorithm, to format the output data in a buffer and to write it on disk.

With this software the system was tested using two complete CPD racks, and it was shown that with the Super Micro machines (faster and with more memory) it is possible to read all events at a 10 kHz trigger rate. In addition, running on those machines multiple copies of the same program, it is possible to reduce the time needed to process a burst by a factor 2. With other PCs (slower and with less memory), the same results cannot be obtained: possible solutions are either their replacement or the installation of two additional PCs of the same kind, wiring their Ethernet cards to only 2 SLMs.

For the operation of a system in a possible NA62 run, several tools need to be prepared:

- a control mechanism to start centrally all the readout programs on all PCs;
- a data-merging mechanism to build complete events from the various fragments;
- a safe mechanism to match LKr events with event fragments from other detectors;

- a monitoring system to supervise and control the operations.

CPD power supplies

The NA48 CPD power supplies were produced in the mid-80s with special low-noise switching supplies for the PS195 experiment and refurbished for the NA48 application, which required higher output currents. The actual power requirements are: +5.2V 350A, -5.2V 90A, -2V 8A, +15V 15A, -15V 1A. The power supplies have reached their expected life cycle and, in order to have an operational readout while the new CREAM system is being built, at least a fraction of them should be replaced with new ones, using the same mechanics and newer power modules. Requests have been sent and prototypes are going to be tested to verify the functionality and the coherent noise level.

4.1.15.5 LKr L0 Trigger System

Introduction and overview

The Level 0 LKr electromagnetic calorimeter trigger identifies electromagnetic clusters in the calorimeter and prepares a time-ordered list of reconstructed clusters together with the arrival time, position, and energy measurements of each cluster. As such, the system also provides a coarse-grained readout of the LKr that can be used in L1/L2 software trigger levels and off-line as a cross-check for the standard readout (see Figure 331).

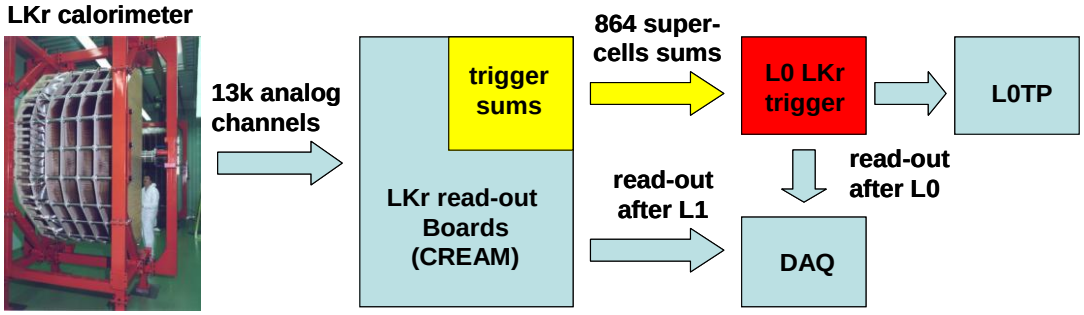

Figure 331: Scheme of the LKr L0 trigger system.

The trigger processor continuously receives from the LKr readout modules signals corresponding to tiles of 16 calorimeter cells (super-cells). Electromagnetic cluster search in the electromagnetic calorimeter is executed in two steps with two one-dimensional (1D) algorithms. From the trigger point of view the calorimeter is divided in slices parallel to the horizontal axis (assuming for the moment super-cells of 2x8 cells, 2 vertical, 8 horizontal). In the first step pulse peaks in space and time are searched independently in each slice with a 1D algorithm, along such axis. In the second step different peaks which are close in time and space are merged and assigned to the same electromagnetic cluster.

The LKr L0 trigger processor is a three-layer parallel system, composed of Front-End (FE) and Concentrator boards, both based on the TEL62 cards (Figure 332).

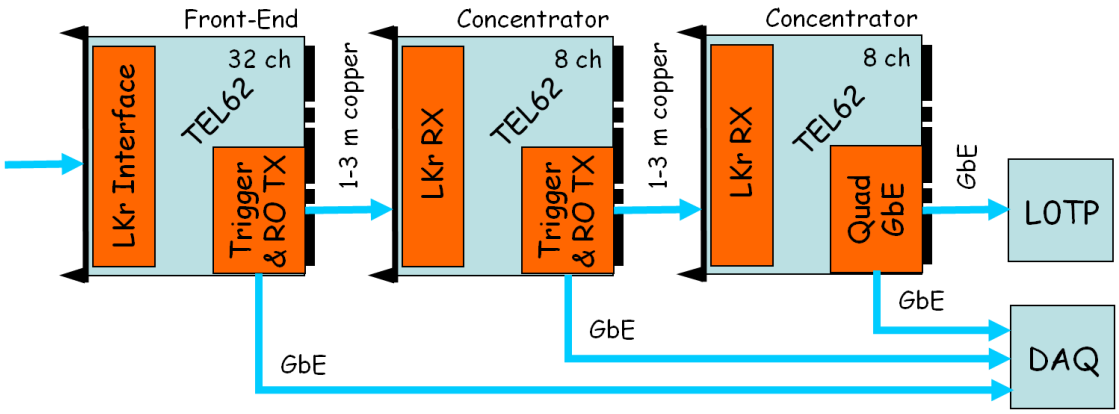

Figure 332: Preliminary implementation of LKr L0 trigger system.

Each FE board receives 32 tiles (Super Cells) from the LKr readout modules and performs peak search in space and computes time, position and energy for each detected peak[56].

The concentrator board receives trigger data from up to 8 FE boards and combines peaks detected by different front-end boards into a single cluster.

The implementation of the trigger processor -assuming input tiles (Super Cells) of 2x8 calorimeter cells (2 cells along the horizontal axis and 8 cells along the vertical axis)- will be discribed here. The extension to the other configurations (8x2 or 4x4) is straightforward.

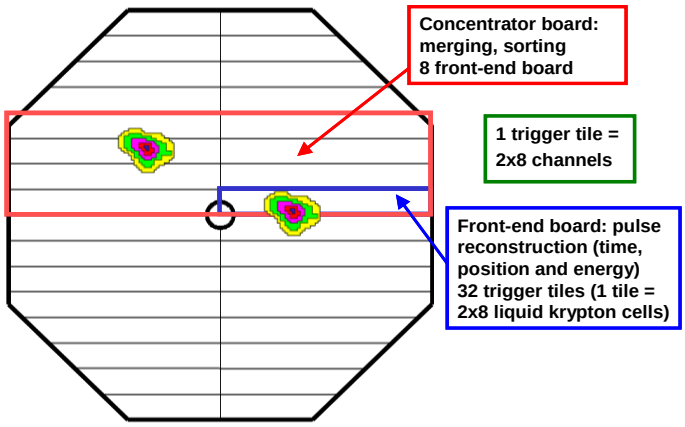

Figure 333: Trigger processor break-down

In total, the system will be composed of 36 TEL62 boards, up to 175 mezzanine cards and 192 high-performance FPGAs. A summary of the main parameters of the L0 LKr trigger is given in Table 60.

[56] In a previous concept the LKr readout system was assumed to provide analog sums, and ADC mezzanine boards were housed in the front-end boards.

Table 60: LKr L0 trigger parameters.

Input channels (tiles)	864
Trigger output channels	1
Readout output channels	28 (raw data) + 7 (reconstructed clusters)
Electronic modules (TEL62)	28 front-end+ 7 concentrator + 1 final concentrator
Latency	< 100 μs

Front-End boards

The Front-End boards continuously receive 864 trigger sums from the readout system, each one corresponding to 16 calorimeter cells. 28 FE boards are foreseen for the whole LKr L0 trigger system.

Each FE board receives 32 tiles trigger sums from a readout module, performs the peak search algorithm and transmits reconstructed peaks to the Concentrator boards.

Raw data received by the readout modules are also stored in L0 latency memories, to be readout after a positive L0 trigger is received.

The peak search algorithm is executed in parallel on all the tiles in the following steps (Figure 334):

- Peak search in space. A peak in space is defined by the following condition:
  $E_{i-1}[n] < E_i[n]$ AND $E_i[n] > E_{i+1}[n]$
  where E is the ADC count, i is the the tile number and n is the sample number.

- Peak search in time. A peak in time is defined by:
  $E_i[n-2] < E_i[n-1] < E_i[n]$ AND $E_i[n] > E_i[n+1]$

- Threshold check:
  $E_i[n] > E_{th}$

- Parabolic interpolation in time around sample maximum, using samples n-1, n and n+1 to get an estimate of peak height (Emax).

- Linear interpolation in energy between samples n-2 and n-1 to get the fine time corresponding to a programmable fraction between 0 and 1 of Emax.

Information on reconstructed peaks (maximum energy Emax, coarse time and fine time) are transferred from the Front-End board to the Concentrator boards on a low-latency dedicated trigger link.

A preliminary version of the above peak reconstruction algorithm was simulated and implemented on an ALTERA Stratix I FPGA (the device used as PP-FPGA on the original TELL1 boards, see section 4.1.5). With 12-bit pulse height resolution a non optimized version of the algorithm can process one peak (5 samples) at a rate in excess of 80 MHz, corresponding to 62.5 ns to process one peak. The maximum acceptable peak rate in a single PP-FPGA is thus 16 MHz for this model of FPGA[57].

[57] The new TEL62 board will have more powerful devices.

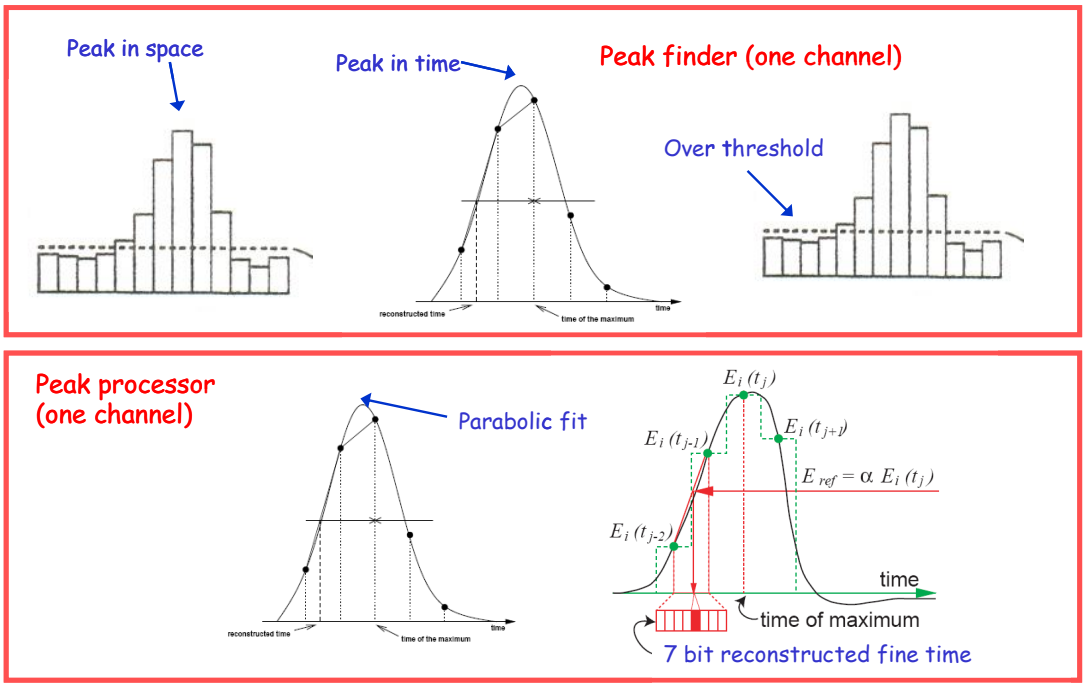

Figure 334 Peak reconstruction algorithm of LKr L0 trigger.

Simulations were performed using the following pulse shape

$$A\ [\ 1 + \sin(2\pi\ t/T - 3\pi/2)\ ]$$

with T = 175 ns to check the algorithm, obtaining satisfactory theoretical performances. A more sophisticated algorithm will be implemented for the real data-taking.

Two custom mezzanines are foreseen for the FE TEL62 boards. The readout module interface mezzanine will continuously receive tile signals from the LKr readout modules. The transmitter mezzanine will transmit high-priority trigger data to the Concentrator boards on a custom trigger link, and low-priority readout data to PCs on a standard Gigabit Ethernet copper cable.

Concentrator boards

The second-stage (Concentrator) boards receive trigger data form the FE boards, possibly combine peaks detected by different FE boards into a single cluster, and prepare time-ordered trigger primitives for the L0 Trigger Processor.

Reconstructed clusters are also stored in L0 latency memories and readout after a positive L0 trigger signal is received via TTC. Eight concentrator boards are foreseen, the last one dedicated to the L0 Trigger Processor interface.

Each concentrator board receives data from up to 8 FE boards, covering a region of 4 tiles along the vertical axis and 64 tiles along the horizontal axis. Each Concentrator board will:

- perform the peak reconstruction algorithm for clusters at the boundary between two FE boards around the central vertical axis of the calorimeter (a cluster in this region will be split between two neighbouring FE boards along the horizontal axis);
- merge information from different FE boards along the vertical axis, by associating reconstructed pulses along such axis to the same cluster. The region covered by a Concentrator board is divided in an inner region and an outer region: only clusters with a maximum along the vertical axis in the inner region are managed by the Concentrator board, to avoid double-counting of the same cluster (see Figure 335).

One custom mezzanine on each Concentrator board is foreseen to receive high-priority trigger data from the FE boars.

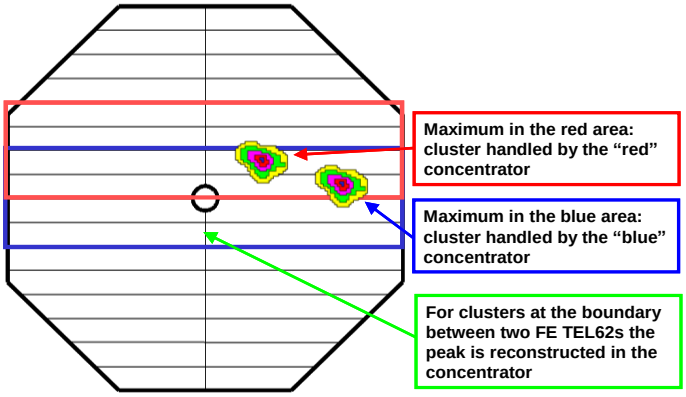

Figure 335: Concentrator board action.

Connectivity and crates layout

Connectivity between LKr readout modules and FE boards is implemented with low-latency point-to-point links in such a way that a single FE board will receive a slice of 32 contiguous trigger sum tiles. 8 FE boards are connected to a Concentrator board, with some overlap between neighbouring Concentrators to guarantee that each cluster will be fully contained in at least one Concentrator board.

The LKr L0 trigger system will be hosted in one VME crate and two or three TEL62 crates. Two 21-slot TEL62 crates (like those built for the LHCb experiment) will host 14 FE boards (half of the calorimeter) each. The remaining eight Concentrator boards will be either divided between the two above crates or hosted in a third crate, depending on power consumption and cabling requirements.

Readout modules interface

At least two tiles will be transmitted from each CREAM readout module, with 16 bits per tile transmitted every 25 ns to fully exploit the resolution of the CREAM on-board 14 bit ADCs (see section 4.1.15). Data can be transmitted from the CREAM to the FE boards using either optical fibres or copper cables: the decision will define the required mezzanine receiver card on the FE boards. A widely used solution has been identified, based on a Texas Instruments TLK2501 serializer chip and an optical transceiver. In this case 16 data bits will be transmitted together with two control bits. Assuming a 100 MHz transmission clock a bandwidth of 1.6 Gbps can be obtained, exceeding the minimum required bandwidth needed to transmit 16 bits at 80 MHz. A digital transmission option using copper cables is still under study.

L0 Trigger Processor interface

Trigger primitives are produced by the Concentrator boards and sent to the L0 Trigger Processor. Each Concentrator board receives trigger data (portions of clusters in the calorimeter) from eight FE boards, but only clusters with a centre in the four inner FE boards are reconstructed by each Concentrator.

The reconstructed cluster rate for Concentrator boards covering the central region of the calorimeter can be (over)estimated as

Reconstructed cluster rate =
= Instantaneous hit rate x multiplicative factor for central regions x (4 / 28) =
= 30 MHz x 3 x (4 / 28) = 13 MHz

corresponding to

64 x 13 Mbps = 0.8 Gbps

The average trigger primitive rate per Concentrator board can be estimated as

Reconstructed cluster rate per Concentrator = Instantaneous hit rate x (4 / 28) =
= 30 MHz x (4 / 28) = 4 MHz

corresponding to

64 x 4 Mbps = 256 Mbps.

Trigger primitives correspond to cluster multiplicities in a given time slot comparable to time resolution of the trigger processor. Time matching between different clusters (e.g. two simultaneous photons originating from the same particle) is done in the LKr L0 trigger system. A final Concentrator board collects and handles the trigger output of the seven Concentrators boards. A single link connects the last Concentrator board to the L0 Trigger Processor.

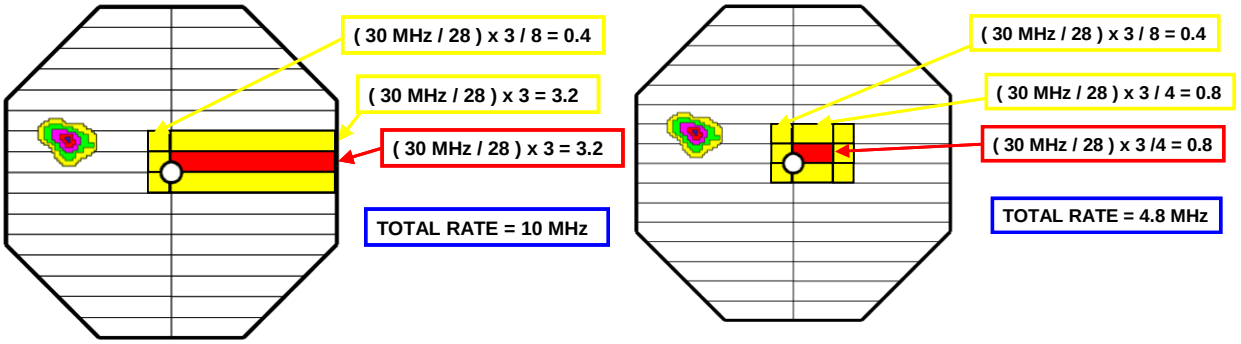

Figure 336 LKr L0 system hit rates

Hit rates, dataflow and latency

The hit input rate is one of the most important constraints for the design of the LKr L0 trigger system, determining the minimum required computing power for each FPGA and the minimum required transmission link bandwidth for trigger and readout data. Hit rates were estimated under the following (conservative) assumptions: instantaneous design hit rate 30 MHz (1); all rates in the central region 3 times the average hit rate in the calorimeter, all particles (including muons) generating a shower of 256

calorimeter cells. Under these assumptions, and taking into account showers originating in the neighbouring regions, the rate for a FE board and for a PP-FPGA in a FE board are shown in Figure 336.

These rates must be matched to the PP-FPGA peak processing power and the various bus bandwidths inside the TEL62 board.

A new TEL62 output card was designed, housing two low-latency high-speed links for the trigger data, and two Gigabit Ethernet links for the readout data, in order to allow large (4.8 Gbps) data rates from the FE boards to the Concentrator boards.

The dedicated trigger link is based on a Camera Link cable assembly and a National Instruments 48-bit channel link chipset running at 100 MHz in one direction, and a 7-bit channel link chipset running at 66 MHz in the opposite direction. A high-quality industrial-grade halogen-free Camera Link cable is available, but relatively expensive. A custom cable solution with halogen-free individually-shielded twin-axial pairs assembled in the lab is being investigated. The bandwidth from a FE board to the Concentrator board is 48 bits at 100 MHz, while the bandwidth form a Concentrator to a FE board is 7/14 bits at 66 MHz.

Given the asynchronous architecture of the LKr L0 trigger system, latency is not constant and can thus only be estimated. Total latency can be divided in the following components: peak reconstruction in the FE PP-FPGA, data transmission from the FE PP-FPGA to the FE output connector, data transmission from the FE board to the Concentrator board over the dedicated trigger link, cluster reconstruction in the Concentrator board, data transmission from the Concentrator board to the L0 Trigger Processor. Each one of these tasks will contribute a few clock cycles to the total latency, giving a total latency of the order of few μs. Of course the total latency can increase in case of hit input rate fluctuations giving pile-up at the input of the processing elements and of the transmission links. A measurement of the capability of the system to absorb hit input rate fluctuations is thus given by the margin between the expected average processing and transmission capabilities at each stage of the processor and the maximum ones.

Readout

After reception of a L0 trigger, the computed data are readout from the L0 LKr trigger to be eventually logged on tape together with sub-detector data. Two different kinds of data will be readout from the L0 LKr trigger:

(a) raw data from each FE board and
(b) reconstructed clusters (with reconstructed time, position, energy and shape) from each Concentrators board.

The readout bandwidth from each FE TEL62 can be estimated as:

Readout BW (FE) = L0 trigger rate x tiles x samples x 16 bit =

= 1 MHz x 32 x 5 x 16 = 2.56 Gbps

which can be easily matched to the maximum 2 Gbps maximum acceptable rate available from each FE board using some compression algorithm (to be developed).

The readout bandwidth from each concentrator TEL62 can be estimated as:

Readout BW (Concentrator) = L0 trigger rate x clusters x 256 bit =

= 1 MHz x 1 x 256 = 256 Mbps

with very conservative assumptions of one reconstructed cluster per board per event, and a maximum of 256 bits to encode time, position and energy for a cluster in the calorimeter. This data can be used by the L1/L2 software trigger as a seed for more elaborate trigger algorithms.

4.1.16 SAC/IRC System

The design of the readout electronics is determined by the rate the detectors have to survive. The estimated rate for photons and muons in the SAC is about 1 MHz. With a signal length of 60 ns, the probability to have a second hit overlapping the first hit in the detector is about 6%. This is because the SAC, with its long attenuation length, can be considered as a single-channel detector. Considering the IRC with a rate of approximately 5 MHz the overlapping probability is about 25%. This means that the ability to distinguish two consecutive pulses is of great importance for the electronics design. A few possibilities can be envisioned: set a high threshold (due to the fact that most of the rate is caused by muons, one can set a threshold above the minimum ionizing particle (MIP) with a resulting loss of efficiency); construct a segmented detector (a reflective material, for example Tyvek, could be put between the individual segments in order to prevent crosstalk); or use a fast ADC readout (waveform digitizer) which will allow the observation of the different pulses by looking at the shape of the signal. The last solution has been chosen for both the SAC and IRC readouts.

The readout requirements for the SAC and the IRC can be met by means of high performance 1 GHz FADC operating as a mezzanine board on TEL62 general readout board designed for the LHCb experiment as shown schematically in Figure 337. One can effectively build such a system at a reasonable cost due to the availability of commercial 8 bit FADC chips operating at 1GHz.

At least two manufacturers have recently developed dual 8-bit 1 GHz FADC converters using CMOS technology[58]. These circuits can be operated as either as a dual channel 1 GHZ or as one single 2 GHz FADC (2 GHz operation is possible using the leading and trailing edge of the 1 GHz clock. Every second input would not be used. The change from 1 to 2 GHz is made via the serial control interface). All features of the FADC are controlled with a serial interface. This eliminates many external components necessary for gain and offset tuning. For example it is possible to program the pedestals, gain and delay of each channel via this interface. The power consumption of both chips is in the order of 1.7 W and they are available in 20 x 20 mm LQFP packages. As presently the price of the circuit [90] is about a factor two lower and it has therefore been chosen for this application.

Figure 339 shows the schematic of four of the 16 channels of the GHz FADC mezzanine board. Only 3 ¼ integrated circuits are needed to implement 4 FADC channels. The mezzanine board will be equipped with 4 ADC and 2 FPGA circuits on each side of the PCB. The analogue inputs to FADC must be

[58] *AT84D001B from www.atmel.com and FADC08D1000 from www.national.com*

differential and the manufacturer recommends connecting the inputs with a 1:1 transformer to FADC inputs. The transformer can cause a shift in the pedestal of about 3 LSB levels at the signal rate of 3 MHz but this can be compensated with the FADC internal offset circuit and pre-samples of the pedestal. The transformer has a bandwidth of 0.4 MHz to 500 MHz. The input impedance is 100 ohms suitable for high quality twisted pair cables like CAT7[59]. The clock is distributed individually to each FADC with a low-jitter <1 ps rms clock fan-out circuit and must be AC coupled. The data outputs of each FADC are de-multiplexed to half to clock frequency (500 MHz) and are transmitted with LVDS levels to the frontend FPGA. One FPGA circuit can handle the outputs from 4 FADC channels. The task of this frontend FPGA is to reduce the rate of the FADC (8 Gbytes/s) to one suitable for transmitting to the TEL62. There are 160 data lines available in the 200 pin connector of the mezzanine board between the frontend FPGA and the PP-FPGA of the TEL62. A rate of up to 1.6 Gbytes/s can be transmitted to the TEL62 PP-FPGA per connector. The front-end FPGA has therefore to reduce amount of raw FADC data with a factor 5 by zero-suppression.

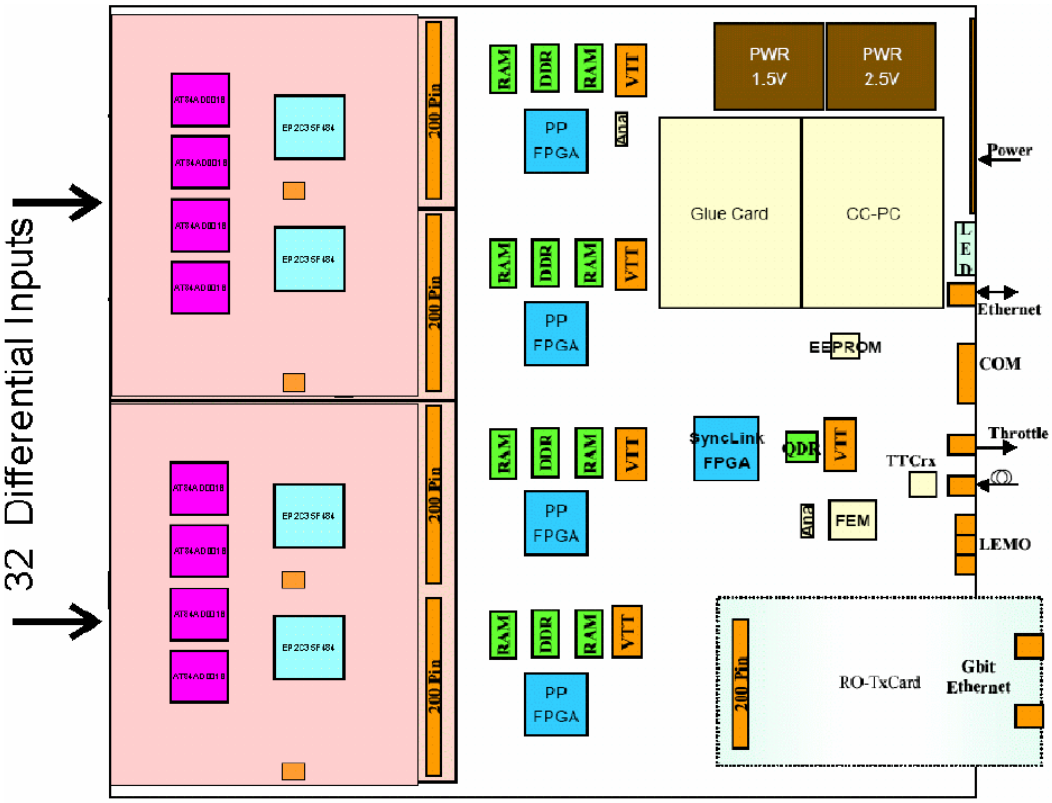

Figure 337 TEL62 board equipped with two FADC mezzanine cards (16 channels each)

Front-End electronics

The amplitude of the signal from the PMT varies between 10 mV for minimal ionizing particle to approximately 3 V for a 75 GeV photon (a dynamic range of about 300). A frontend electronics is necessary in order to provide the necessary differential input to the FADC chip (see Figure 338). It will

[59] *For example, UNINET 7702 4P; www.daetwyler.net or www.disdata.ch Art. Nr. 686350*

be placed on the PCB mounted on the PM itself. The requirements are the low power consumption since it will operate in vacuum (for the SAC) and the cooling might not be very effective. The signal will be transmitted to the FADC input with twisted pair cables.

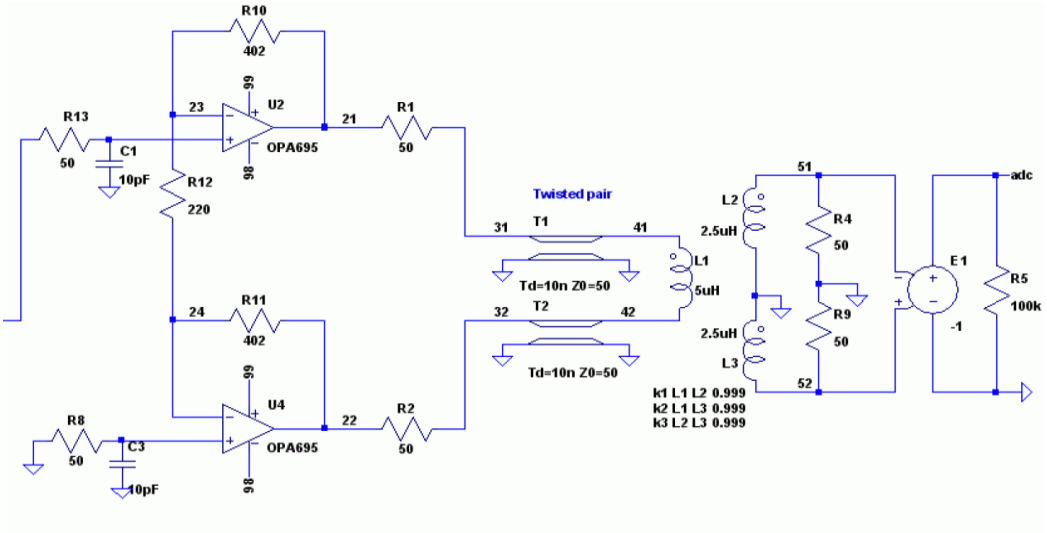

Figure 338 Frontend electronics and amplifier from the PMT to the ADC board.

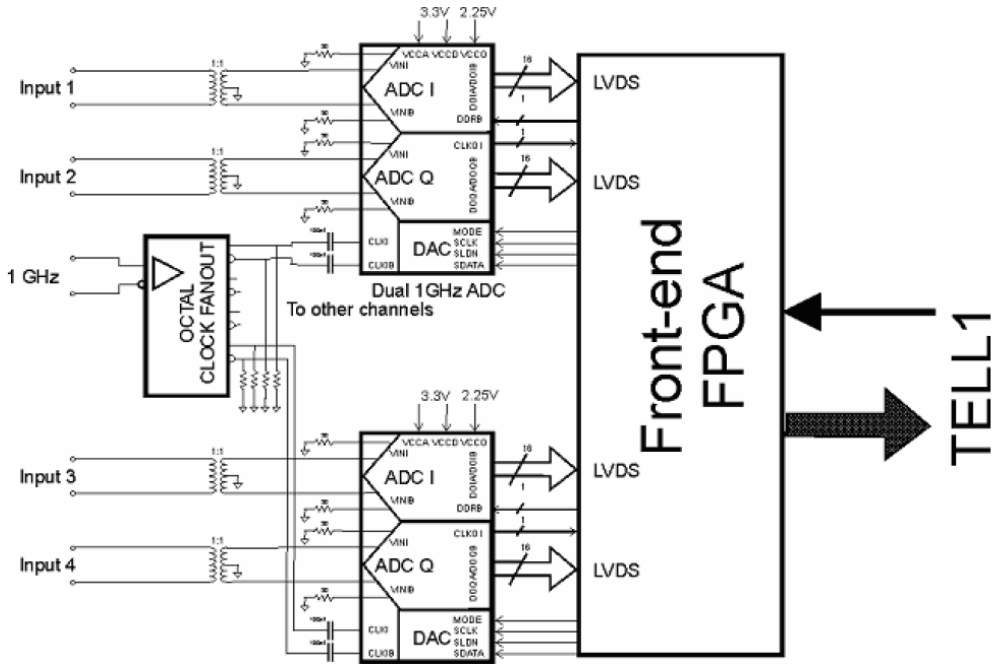

Figure 339 Schematics of the ADC board with four channels.

Readout electronics

The SAC and the IRC will use a TEL62 based readout. The clock and the trigger will be distributed to the TEL62 using TTC. The data will be transmitted from the TEL62 to a readout PC and afterwards, using Gigabit Ethernet, to the central processing system for event construction.

4.1.17 GPU Improvements

A R&D program was started in Pisa to evaluate the possibility of improving the performance and cost-effectiveness of the TDAQ above the baseline solution described here, by massively exploiting GPUs (video card processors) in a hard real-time environment [104].

GPUs have the advantage over standard CPUs of a massive parallelism and huge computing power for parallelizable tasks which do not require complex control, and the advantage with respect to FPGAs of a much simpler programmability and scalability, being commodity devices. The present generation of devices can provide up to 1 Teraflop computing power with a memory bandwidth of 100 GB/s. Thanks to this performance the processing time heavily depends on the speed of data transfer towards the links to bring the data to the GPU.

While the use of GPUs in scientific computing is now quite established, its use in a data-acquisition system has not been tried, but preliminary studies within NA62 are encouraging in this respect and will continue.

A full discussion of this program goes beyond the scope of this document, and here we only mention some points where the use of GPUs could provide some advantage for NA62.

- GPUs could be hosted in L1/L2 PCs to provide a significant boost in the computing power for non-time-critical computations; this is straightforward.
- GPUs could be used to implement the L0 Trigger Processor, most likely requiring such device to be implemented in a PC, which is at present required to handle the GPU.
- GPUs could be used to replace the hardware systems evaluating L0 trigger primitives; this is the most challenging task, as the transfer and handling of all the sub-detector data is required, and a hard real-time response of both the GPU and the controlling CPU are mandatory. Some algorithms for computing L0 trigger primitives on GPUs have been developed and timed, showing that with due care to the hardware architecture involved the data processing capability and latency are not an issue.

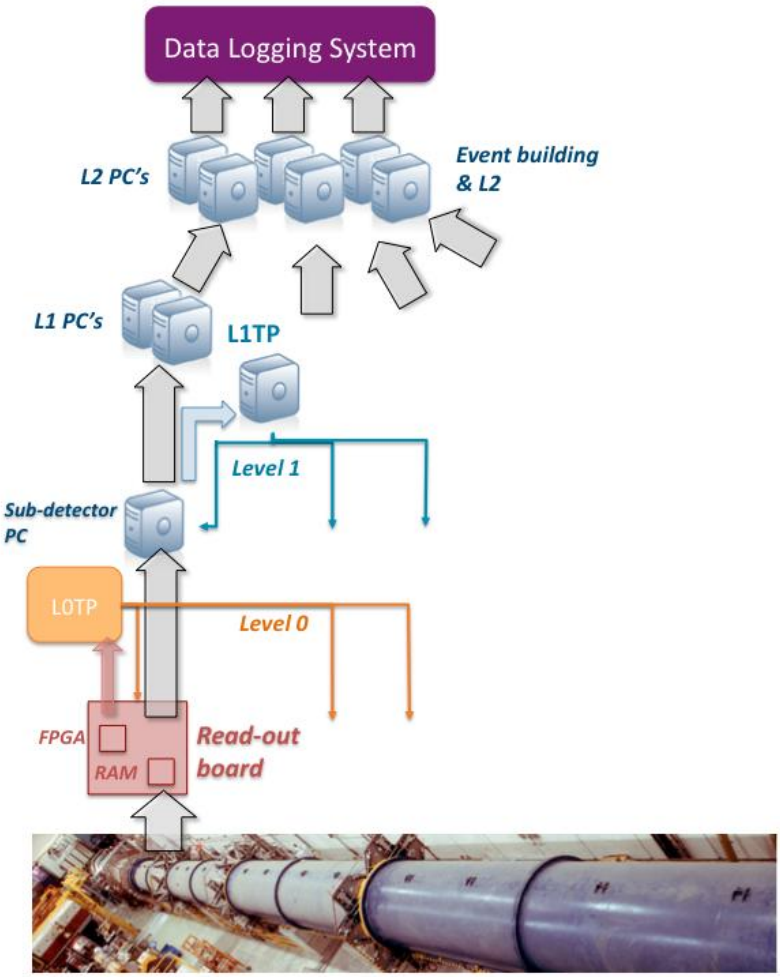

- *Figure 340 Logical scheme of the data-flow, from the detector to the permanent storage.*

4.2 Online Computing System

Data processing model

The data processing model is strongly connected with the Trigger & DAQ architecture, described in the previous sections. The general scheme of the data-flow, depicted in Figure 340 is summarized in the following.

The flow of data, from the sub-detectors to the permanent data storage, can be schematically summarized in the following steps

- The **readout boards** – for most sub-detectors the TEL62 – get data from the front-end electronics of the detectors, store the information in buffer memories (RAM), and (if the sub-detector is contributing to the L0 trigger) generate Level 0 Trigger primitives, as shown in Figure 341 the GTK is indeed not partecipating to the L0 trigger and has its own readout system. The LKr has a different data-flow, due to the large amount of data produced by the 40 MHz flash ADCs. The calorimeter readout modules (CREAM) are capable of holding events without (or with a very loose) zero-suppression in large RAM memories, that are downloaded to the online farm only on a positive Level-1 trigger decision. In order to partecipate to the

Level-0 trigger, reduced granularity information, coming from a separate readout is used, based on sums of super-cells continuously digitized and pipelined in order to perform cluster-finding and photon-counting algorithms and send information to the Level 0 Trigger Processor (L0TP).

- The **Level 0 Trigger Processor** (**L0TP**) elaborates the L0 primitives from different sub-detectors, and produces the L0 trigger decision.
- The L0 positive decision is distributed to all the readout boards, and when this is received, the data are transferred to dedicated PCs (**sub-detector PCs**).

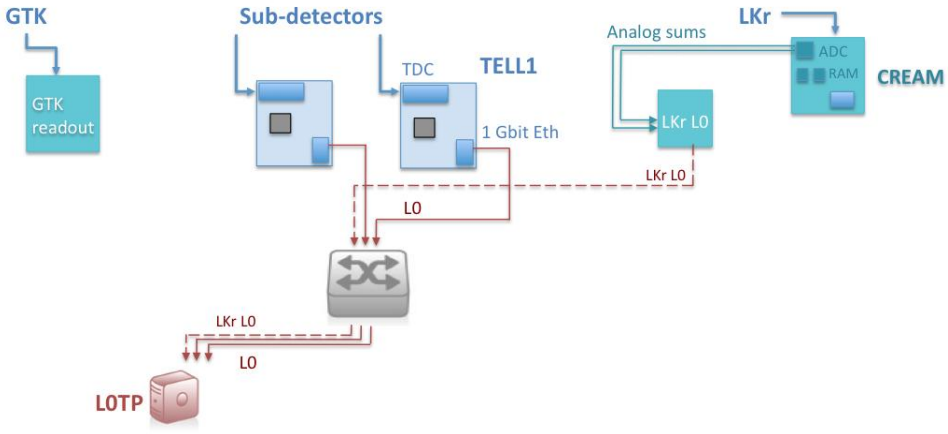

Figure 341 Logical scheme of the data-flow from the detectors to the readout boards and of L0 primitives to the L0 trigger processor.

This is shown in Figure 342 where the flow of data packets is shown, from different readout boards of a given sub-detector until they reach a PC in the L1 farm, that will assemble all the fragments of a given event for the sub-detector in a “sub-event”.

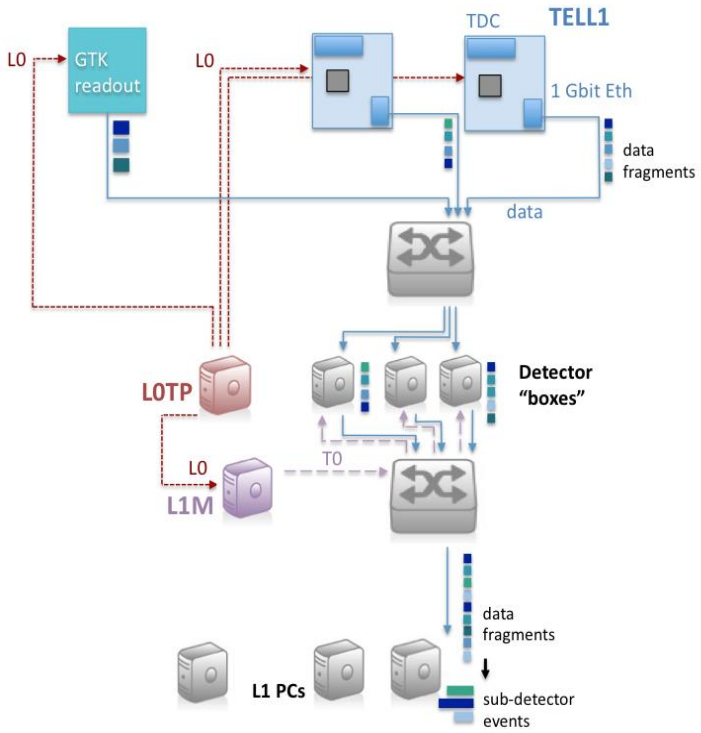

Figure 342 Logical scheme of the data-flow from the readout boards to the L1 farm, upon arrival of positive L0 decision.

The "detector boxes" shown in such figure can be either ordinary PCs or intelligent network boards dispatching the data fragments of a given event all to a single L1 PC, according to an algorithm (that could even be as simple as a round-robin table) produced by a Level-1 manager (L1M) machine. This task can even be performed by the readout boards (*e.g.* the TEL62 credit-card-PC) or by the main switch (placed in the electronics barrack and connecting all the readout boards of all sub-detectors) using IP packet inspection and redefinition of the final address of the data fragments.

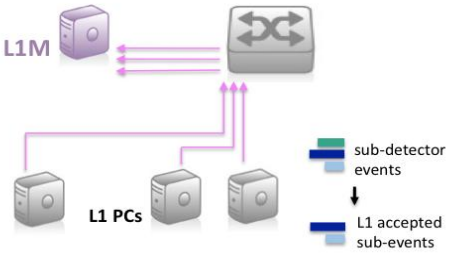

Figure 343 Logical scheme of assembly of sub-events in the L1 farm and building of L1 decision in the L1 Manager (L1 Trigger Processor).

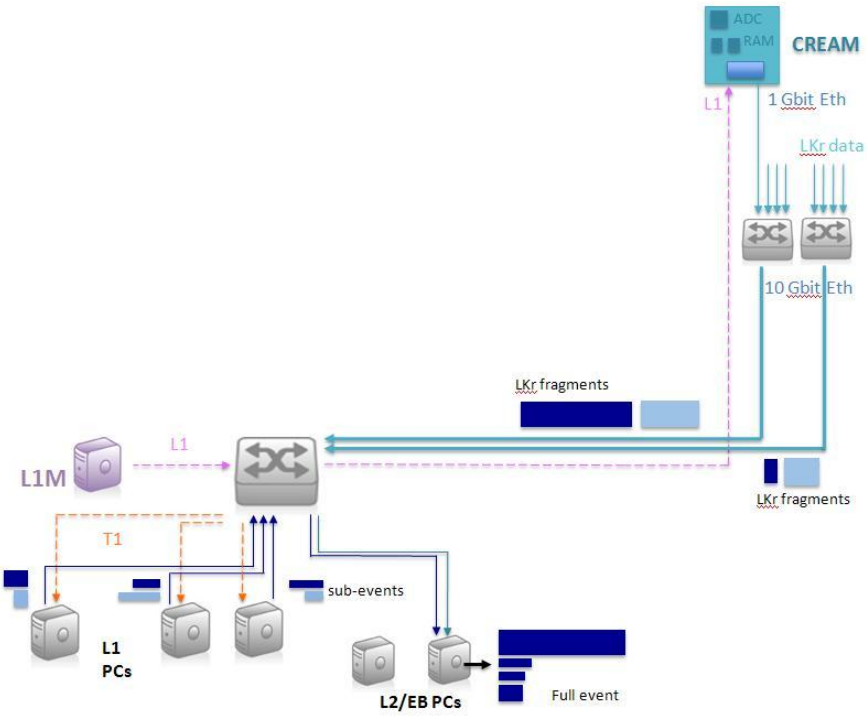

Figure 344 Logical scheme of the L1 distribution, data-flow from the L1 to the L2 machines and full readout of LKr data from the CREAM modules memories to the L2 farm.

- The **Level 1** software trigger is based on primitives produced by the sub-detector PCs, each handling the complete event data from a single sub-detector; the **L1M** trigger manager (or L1 Trigger Processor) will then merge those primitives in order to assert a Level 1 decision, as shown in Figure 343.
- The L1 trigger decision is distributed to the sub-detector PCs and (assuming a positive decision) sub-event data (L1 accepted sub-events) is transferred to the **Event Builder PCs.** These processes, running on part of the online-farm machines (labelled **L2/EB PCs** in Figure 344, assemble full events by merging together all the information coming from the different sub-detectors. Data from the CREAM (calorimeter readout modules) are also downloaded on a positive L1 decision, from the buffers onto the appropriate L2 machine. Once events are fully built, a **Level 2** trigger can be evaluated and issued.

- At this point, L2-accepted events can be transferred to the Data Logging System (DLS), possibly running remotely (e.g. in the CERN main data-center or even off-site), in order to be registered on a permanent storage (e.g. tapes), as schematically shown in Figure 345.

Putting together all the logical steps of the trigger and event-building, the general picture of the data-flow is the one shown in Figure 346.

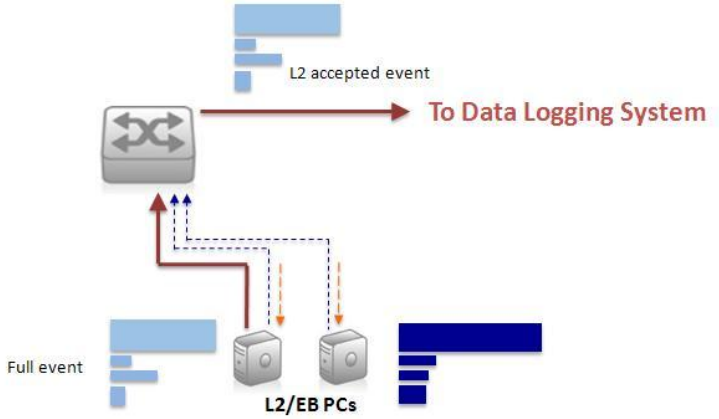

Figure 345 Logical scheme of the data-flow of fully reconstructed and L2-selected events from the L2 farm to the Data Logging System.

Event data model

The Event Data Model defines the data formats for the experiment. We can identify at least two basic data formats:

- **RAW**: the binary data readout by the different sub-detectors and assembled by the event-building processes in the on-line farm. It contains all the information produced by the sub-detectors electronics and the trigger modules, possibly pre-processed by dedicated processors (*e.g.* the Liquid Krypton calorimeter system farm) and filtered by the Level 2 trigger. The size of the event is estimated on the basis of the average number of active channels in each sub-detector and on the number of words needed to store the data of each channel, that we generically call **element**: an element is a single hit in the spectrometer chambers, a fired cell in the Liquid Krypton calorimeter, a fired pixel in one station of the GTK, etc. In general, the minimal information for one element is given by a time measurement word.
- **RECO**: the output of the reconstruction of the elements: tracks from the tracker hits, clusters from the calorimeters cells, etc.

Other data formats will be needed by the experiment, in particular we can identify at least the following additional formats:

- **THIN**: summary data extracted from the fully reconstructed events, easier to handle by analysis programs, complemented with meta-data coming from the conditions data-base and calibration tasks; and possibly even more skimmed events, that we can call **Super-THIN**, that can be used for specific analysis purposes.

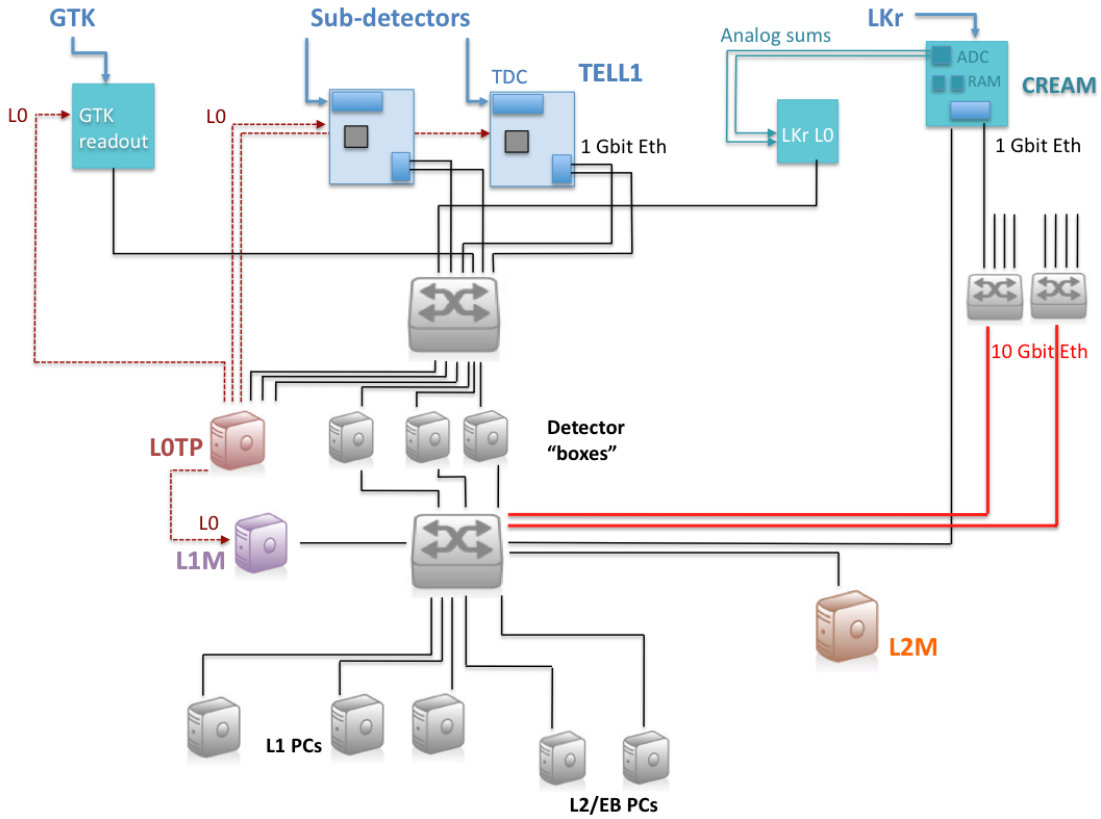

Figure 346 Logical scheme of the data-flow, trigger distribution and online farm.

- **CALB**: data from single sub-detectors or group of sub-detectors, possibly produced by earlier stages of the trigger (*e.g.* event-fragments from the output of the Level 1 trigger), needed for monitoring and calibration purposes. A typical example are events flagged as containing muon candidates by the Level 1 trigger of the photon or muon veto detectors, and partially reconstructed in order to perform monitoring with straight muons.

In addition to data coming from the readout of the sub-detectors, meta-data should also be considered. The three main types of meta-data produced by the experiment will be:

- **COND**: conditions of the detector, stored in the appropriate data-base, connected to the Detector Control System (DCS), such as high- and low-voltage settings of the sub-detector elements, temperature and pressure readouts, status of the readout boards and crates, etc.
- **CONF**: configurations of the run, including active detectors, trigger configurations, beam conditions, etc.
- **COST**: calibration parameters computed by calibration tasks running on (partially or fully) reconstructed data.

Event reconstruction model

The event data model is strongly connected with the model for event reconstruction: they are schematically shown in Figure 347.

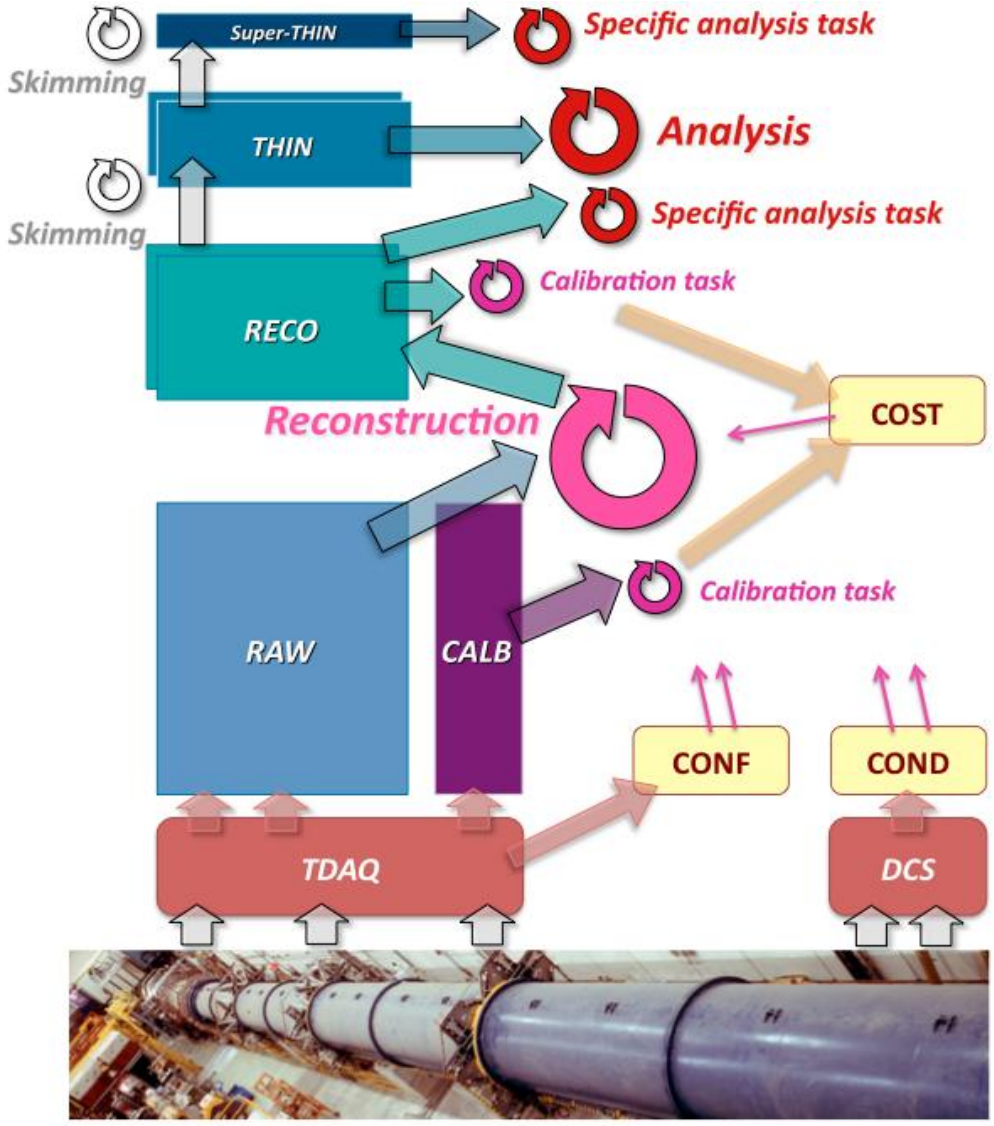

Figure 347 Event data model and event reconstruction model.

The Event Reconstruction Model indeed defines the architecture, both from the logical and implementation points of view, with which data coming from the sub-detectors are first assembled (event-building, EB), reconstructed and possibly filtered (for L2 trigger purposes or for the generation of data streams), and finally permanently stored by the data-logging system (DLS). The model is strictly interconnected with the trigger scheme of the experiment, since some of elaboration and filtering of the partial or full events can be done at the earlier stages of the trigger and EB.

The data coming from the readout of the sub-detector are in general stored in buffer memories until a positive Level 0 trigger decision is taken by the Level 0 trigger processor or the L0 trigger latency has elapsed; in the former case the readout boards send data packets to the sub-detector/L1 PCs, where a first, rough reconstruction of the sub-detector is performed. The maximum rate at the output of the L0 is defined to be **1 MHz**.

The implementation of the Level 1 will be fully in software, and the machines receiving the data and running the L1 reconstruction and decision will constitute the Level 1 farm. The L1 farm will be connected to the readout boards through Gigabit Ethernet links. In order to assemble the data fragments coming from different readout boards of a sub-detector in a given PC, the data packets from the readout boards belonging to a specific event should be routed to the same machine. This can be done assigning the destination IP address of the packets on the fly on a event-by-event basis (*e.g.* using a round-robin mechanism and a look-up-table for the sub-detector PCs to be used by the readout boards). An alternative scheme can be that of concentrating the links from the readout and to the L1 processors into a dedicated switch, handling the packet traffic. At the output of the L1 stage, we consider a maximum rate of **100 kHz**.

The implementation of the L1 farm will be done using rack-mounted PCs, placed in the experiment computing room, upstairs in the Building 918. Network switches concentrating Gigabit Ethernet links from the sub-detector readout systems in the experimental hall will be connected through 10 Gigabit Ethernet fibres to a main concentrator switch upstairs in the computing room. An additional, similar system of network concentrator switches connected by high-speed fibres (from downstairs in ECN3 to upstairs in the computing room) will be dedicated to the readout of the Liquid Krypton calorimeter, which is handled differently from most other sub-detectors because of its large raw data volume. The network layout is described in the following.

On the sub-events selected by the L1 trigger further operations should be performed:

- The sub-events should be correctly re-assembled in complete events by an **event-building** process (EB); this will include the assembling of the (possibly non zero suppressed) LKr data;
- A **fast reconstruction** (FR) software, not necessarily using the most detailed information from the detector (*e.g.*, possibly using heavily zero-suppressed data from the Liquid Krypton calorimeter), nor calibrations, and based on simplified reconstruction algorithms, should run on the assembled events;
- On a positive L1 decision, the complete Liquid Krypton calorimeter data should be downloaded from the readout buffers and appended to the complete event data structure;
- A **L2 trigger** decision could be taken for further rejection of background events;
- At this stage event selection algorithms can be used on the reconstructed events, in order to have one or more **data streams** for storage. A minimal set of streams should include:
    - Golden events, fulfilling a first, loose selection on the main signal;
    - Muon events, selecting decay or halo muons with a (variable) pre-scaling factor.

The EB, FR, L2 and streaming processes, should run on a sub-set of the online-farm.

Additional machines will be hosted in the experimental hall, in order to handle the DCS of the different sub-detectors, and will also be connected by Ethernet links to a network switch downstairs, linked with 10 Gigabit/second (Gbps) links to the main network switch in the computing room.

All the machines in the L1 trigger farm will be on an experiment private network, while dedicated gateway machines will be connected to the CERN General Purpose Network (GPN).

Table 61: Estimated event sizes and data rates at the input of the event building/L2 farm. Additional data from L0 trigger systems and Trigger Processors is not included.

Detector	Event size (B)	L1 trigger rate	Through-put at L2 input (MB/s)
CEDAR	216	100 kHz	21
GTK	2250		215
CHANTI	192[60]		18
LAV	160[61]		15
STRAW	768[62]		73
RICH	160[63]		15
CHOD			
MUV	768[64]		73
IRC & SAC	576[65]		55
Total	≈5k		≈500
LKr	222 k[66]		

Event size and rates

The maximum L0 trigger rate is fixed to be 1 MHz, and we assume 100 kHz at the output of the Level 1 trigger. On the other side of the data processing chain, the Data Logging System, the maximum speed of event logging to tape can be fixed at **100 MB/s** (per tape drive, we assume one fully dedicated tape drive for NA62). The maximum event rate of the DLS, in an operational model in which raw data is assembled by the event-building processes and then directly transferred to tape (with suitably dimensioned disk buffers in order to profit from the low duty cycle of the SPS), is fixed then by the event size:

DLS logging rate (events/s) = tape logging bandwidth (B/s) / event size (B).

The maximum Level 2 output rate is then fixed by the duty-cycle of the slow extraction of the SPS:

---

6 stations * 2 views * 2 strips/view * 2 words * 32bit
20 fired crystals * 2 words * 32bit
4 stations * 4 views * 4 straws * 1.5 average multiplicity * 2 words * 32bit
20 fired PMTs * 2 words * 32bit
48 scintillator planes * 2 tiles/plane * 2 words * 32bit
32 IRC sectors + 4 SAC PMTs * 8 samples *16bit
[66] 13248 cells*8 samples*16bit

L2 rate (events/s) = DLS logging rate (events/s) / SPS duty-cycle

Assuming a SPS cycle made of a 9.6 flat-top spill each 40 seconds, a duty-cycle of **0.25** can be considered.

In order to estimate the size for the different formats of data, we start from the **RAW** and make educated guesses for **RECO**, **THIN** and the other data formats described above.

The data throughput at the input of the event building/L2 farm can be computed from the sub-event data format and the expected rates (100 kHz L1 trigger rate) for each of the sub-systems. This is summarized in Table 61.

Networking

NA62 plans to use switched Ethernet technology to transfer data from the sub-detector electronics to the online processing farm. Several network links will be needed in the detector area: links for connecting readout boards to sub-detector/L1 PCs, to transmit data, and to the L0 Trigger Processor, to transmit L0 Trigger primitives; links for connecting L1 PCs to the L1 Trigger Processor PC, among themselves and to the L2 farm fabric; and links for slow control of the readout boards themselves. To have a uniform system we intend to use Gigabit Ethernet (GbE) everywhere, even for tasks that require less bandwidth (DCS for example). One possible exception to this rule will be the embedded PCs on TEL62 boards for which 100Mbps links are used.

An overall sketch of the TDAQ system and of the network interconnections can be seen in **Error! eference source not found.**. To facilitate the maintenance of the computers, we plan to keep as many as possible of them upstairs, *i.e.* in a dedicated computing room in building 918 next to the NA62 main control room (approximately half of the surface formerly dedicated to the NA60 control room, as shown in Figure 348 rather than downstairs, *i.e.* in the ECN3 experimental area.

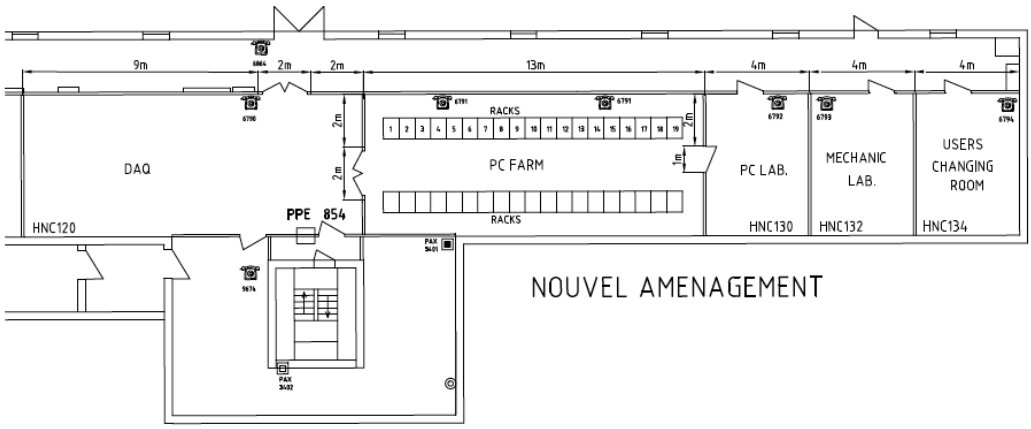

Figure 348 Proposed layout of building 918. The NA62 upstairs computing room will be hosted in a part (approximately one half) of the area formerly used as NA60 control room.

This requires a large amount of fibres connecting the experimental area to the computer room in building 918. Moreover, due to the position of the sub-detectors all along the decay region, the readout boards are not concentrated in the same area, but rather spread along the more than 100 m long underground experimental zone. For this reason we plan to split the network in two branches: an

experimental hall (ECN3) networking structure and a ground level (building 918) one, with optical fibres connecting the two. In the experimental hall we plan to have a central switch in the electronics barrack sitting on top of the Liquid Krypton calorimeter. Starting from there we will have connections to all sub-detectors along the NA62 detector setup.

The following table summarizes the expected number of network links required up to the level of L1 PCs.

We are planning to have several fibres per station/sub-detector running 1 Gbps Ethernet instead of just one (or more) fiber running 10 Gbps Ethernet. The 10 Gbps approach would require to place a concentrator switch at each station/sub-detector (TEL62 network cards are 1 Gbps) and we feel the combined cost would be greater compared to the 1 Gbps solution.

The central switch will act as a sort of aggregator/transceiver, allowing us to reduce the number of fibres connecting the experimental hall to the control room using 10 Gbps ones. The estimated number of 10 Gbps fibers needed is 60.

Table 62: Expected number of network links up to the L1 PC farm. Links internal to sub-system hardware not included.

Detector	**GbE links** to L1 PCs (**data**)	**GbE links** to other detector boards (**trigger primitives**)	**10 GbE links** to farm **(data)**	**GbE links** to L0TP **(trigger primitives)**	**DCS FE links** (100 Mbps)
CEDAR	8	0	0	0	1
GTK	60	0	0	0	
CHANTI	1	0	0	0	1
LAV	48	11	0	1	12
STRAW	64	0	0	0	2
RICH	32	3	0	1	4
CHOD	4	0	0	1	1
MUV	2	0	0	1	1
IRC/SAC		0	0	0	1
LKr			28		
LKr L0 trigger	1	0		1	35
L0TP	1	0		0	1

The ground level (building 918) network infrastructure will consist of a main switch placed inside the computer room. This will be connected to the 10 Gbps fibers going to the experimental hall and to the roughly 300 computers forming the online system (L1 PCs, L2 and event reconstruction farm). The switch will also be connected to a 10 Gbps optical fibre going to the main IT facility, used to transfer the data to be logged onto tape.

Level 1 farm

At this level, the data are merged for each sub-detector. In the general case sub-detectors will be readout following a L0 trigger by separate, independent devices, so that a first building of the event at the sub-detector level is needed. The L1 farm is responsible for receiving data from the sub-detectors front-end electronics, for generating trigger primitives for the Level 1 Trigger Processor (L1TP), and for storing the data pending L1TP response. The data selected by the L1TP will then be passed on to the Level 2 farm, where the event building for the entire detector will be performed and a more complex selection will take place. The only exception to this scheme is given by the LKr sub-detector where, due to the very large event data size involved, a dedicated farm is used.

The L1 farm needs to provide enough aggregate network bandwidth and CPU resources to handle the foreseen data rate, and of course we plan to take full advantage of modern multi-core CPUs and fast 10Gb network links. Still, depending on sub-detector, in can be difficult to accommodate the data from one entire burst in just one computer, generating the trigger primitives at the same time. For this reason we plan to use several computers for each sub-detector, subdividing the load between them with a plain round robin algorithm. To reach this target we count on the ability of the front-end electronics to send different events to different computers. This should be easily accomplished by the TEL62 readout boards used by the majority of the sub-detectors.

Table 63: Estimation for the number of sub-detector/L1 PC and links.

Sub-detector	Number of readout boards	Number of readout links	Number of sub-detector PCs
CEDAR	2	2	1
GTK	tbd	30(60)	15(30)
CHANTI	1	tbd	tbd
LAV	12	24	12
STRAW	2	8	tbd
RICH	4	8	200 cores
CHOD	1	4	1
LKr	432	28 (10GbE)	30 (12-core)
LKr L0 trigger	35	7	tbd
MUV	1	2	1
IRC & SAC	1	tbd	tbd

In case this solution is not viable, we plan to investigate the possible use of network processors to be placed in between the front-end electronics and the L1 farm to modify the ethernet packet at wire speed.

Both solutions guarantee an easily scalable system, where more computers may be added depending on the needs, the main limitation being the central switch.

The number of sub-detector/L1 PCs and links is roughly estimated Table 63.

4.3 Online Control System

The online control system will be a loosely connected set of hardware and software with the main task of controlling the initialization and monitoring the progress of a data-taking run ("run control"), as a thin layer interfacing the sub-system specific electronics on one side and the common infrastructure - such as the SPS timing signals and the PC farms - on the other. The online system should normally not interfere with the sub-detector electronics during a data-taking run, but it may be allowed some action during the inter-spill period. The online system should also interface to the Detector Control System (*a.k.a.* "slow control"), and to the data monitoring system, and present a common user interface for all these items. It is expected that the DIM software [94] will be heavily used for the development of most parts of the online control system.

The configuration, initialization, monitoring and control of sub-detector specific hardware is of course responsibility of each sub-detector group. However, when a sub-detector is included in the active list of sub-detectors for running, each sub-system (thus including the Trigger Processors, etc.) should be able to respond to a common set of commands issued by the NA62 TDAQ system via network. Each sub-detector will identify a *single* Linux machine which will handle the communication tasks with the common NA62 TDAQ system. The remote commands to which each sub-detector system should respond must include the following:

- **Enable/disable global NA62 control**: when enabled the sub-system can only be controlled from the NA62 TDAQ system, for normal running; when disabled the sub-system is under local control (for tests and debugging) and is completely ignored for running and data collection purposes.
- **Perform cold start initialization**: this task forces all basic (possibly time consuming) sub-detector configuration actions to be started, and must be used after a power-up or a reset. No parameter is passed from the TDAQ to the sub-system and therefore no conditional configurations can be performed at this time.
- **Start a new run**: three parameters are passed, the first one being the run number, the second being a common (equal for all sub-system) TDAQ configuration string, and the third being a sub-system specific configuration string. After successfully completing this task, the sub-detector is in running mode from the point of view of the TDAQ system, and it becomes ready to receive triggers as soon as a "start of burst" trigger is received through the TTC system.
- **End a run**: after completing this task, the sub-detector system replies with a set of summary data which the TDAQ system will take care to have stored together with the data.
- **Query status**: this command can be dispatched at any time, and the system should reply with the information on its current state, which should include at least the number of the last run started and the number of the last burst started.
- **Reset**: this command can be dispatched at any time, and forces the system to go back to an uninitialized state (its main use being to get back to a known situation in case some serious error occurred). Dispatching this command during a run can have unpredictable effects on the data being collected.

Most likely, during the development phase or for testing purposes, sub-detectors will also want to be able to run also with different setups (*e.g.* in the lab, or at the experiment using a separate "private" clock/trigger distribution system), and therefore they might be interested in running parts of the NA62

online software on different machines. The online group will provide support for the common NA62 online software to run on “NA62 standard configuration” hardware.

Bibliography

1. BUCHALLA, G.; BURAS, A. J. The rare decays K --> pi nu nubar, B --> X nu nubar and B --> l+l-: An update. **Nucl. Phys. B548 (1999)**, 1999.

2. MISIAK, M.; BROD, J. QCD corrections to FCNC decays mediated by Z-penguins and W-boxes. **Phys. Lett. B451**, 1999.

3. BURAS, A. J. et al. Charm quark contribution to K+ --> pi+ nu nubar at the next-to-next leading order. **JHEP 11 (2006) 002**, 2006.

4. BROD, J.; GORBHAN, M. Electroweak Corrections to the Charm Quark Contribution to K+ --> pi+ nu nubar. **Phys. Rev. D78 (2008) 034006**, 2008.

5. ISIDORI, G.; MESCIA, F.; SMITH, C. Light-quark loops in K--> pi nu nubar. **Nucl. Phys. B718**, 2005.

6. BURAS, A. J. et al. Patterns of Flavour Violation in the Presence of a Fourth Generation of Quarks and Leptons. **TUM-HEP-750-10; e-Print: arXiv:1002.2126**, 2010.

7. BLANKE, M. et al. FCNC Processes in the Littlest Higgs Model with T-parity: a 2009 Look. **Acta Phys.Polon.B41:657,2010 e-Print: arXiv:0906.5454 [hep-ph]**, 2009.

8. BLANKE, M. et al. Rare K and B Decays in a Warped Extra Dimension with Custodial Protection. **JHEP 0903:108, e-Print: arXiv:0812.3803 [hep-ph]**, 2009.

9. ALTMANNSHOFER, W.; BURAS, A. J.; GORI, S.; PARADISI, P.; STRAUB, D. M. ET AL. Anatomy and Phenomenology of FCNC and CPV Effects in SUSY Theories. **Nucl. Phys. B830**, 2010.

10. E949 COLLABORATION; ARTAMONOV, A.V. ET AL. New Measurement of the K+ --> pi+ nu nubar branching ratio. **Phys. Rev. Lett. 101 191802**, 2008.

11. NA62 COLLABORATION AMBROSINO F. ET AL. Proposal to Measure K -> pi nu nu rare decay at the CERN SPS. **CERN-SPSC-2005-013**, 2005.

12. NA62 COLLABORATION AMBROSINO F. ET AL. NA62 Status Report. **CERN-SPSC-2007-035**.

13. M. FIORINI ET AL. **An ultra fast silicon pixel detector for the NA62 experiment:** The Gigatracker. KAON09. Tsukuba: PoS. 2009.

14. LENTI, M. The NA62 RICH detector. **Nucl.Phys.Proc.Suppl. 197**, 2009. 117-120.

15. NA48 COLLABORATION; ANVAR, S. ET AL. The Beam and Detector for the NA48 neutral kaon CP violation experiment at CERN. **Nucl. Instrum. Methods A 574**, 2007. 433-471.

16. ATHERTON H. W. , ET AL. **Precise Measurements of Particle Production by 400 GeV/c Protons on Beryllium Targets**. Geneva. 1980. (CERN Yellow Report: CERN 80-07).

17. THE NA48/2 COLLABORATION; BATLEY J. R. ET AL. Search for direct CP violating charge asymmetries in K± -> π± π+ π- and K± -> π± π0 π0 decays. **Eur. Phys. J. C 52**, 2007. 875-891.

18. BRIANTI, G.; DOBLE, N. **The SPS North Area High Intensity Facility: NAHIF**. CERN. Geneva. 1977. (SPS/EA 77-2; SPSC/T-18).

19. BONNET, Y. **Layout TCC2/TT81--->85**. CERN. Geneva. 2006. (CERN Drawing Directory (CDD) : SPSXLTC20001).

20. DOBLE, N. Beam Documentation for NA48/NA62: Geometry-BEATCH, Optics-TRANSPORT, Beam Simulations-TURTLE and HALO. **CERN**, 2010. Disponivel em: <http://doble.web.cern.ch/doble/>.

21. Y., B. **Layout TCC8/ECN3**. CERN. Geneva. 2009. (CERN Drawing Directory (CDD); SPSXLECN0002).

22. BROWN K.L., ET AL. **TRANSPORT, A Computer Program for designing Charged Particle Beam Transport Systems**. Geneva. 1980. (CERN Yellow Report: CERN 80-04).

23. BOVET C. ET AL. **The CEDAR counters for Particle Identification in the SPS Secondary Beams**. CERN. Geneva. 1982. (CERN Yellow Report : CERN 82-13).

24. BROWN, K. L.; ISELIN, C. **DECAY TURTLE, A computer Program for simulating Charged Particle Beam Transport Systems**. Geneva. 1974. (CERN Yellow Report: CERN 74-2).

25. ISELIN, C. **HALO, A Computer Program to calculate Muon Halo**. Geneva. 1974. (CERN Yellow Report: CERN 74-17).

26. GIGATRACKER WORKING GROUP. The GIGATRACKER: Addendum to the NA62 proposal. **Internal Note NA62-07-08**, Nov. 2007.

27. FIORINI, M. et al. Test of Silicon Sensors for a High Rate Pixel Detector for the NA62 experiment. **Internal Note NA62-08-01** , Geneva, 2008.

28. ANELLI ET AL. Optimum Segmentation and Thickness of Silicon Pixel Detectors for Signal to Noise Ratio and Timing Resolution. **Nuclear Science Symposium Conference Record IEEE, Vol 2**, 2006. 671-680.

29. FACCIO, F.; CERVELLI, G.; MARCHIORO, A. Single Event Effects in 130 nm, CMOS technologies. **LNL annual report 2005, ISSN 1828-8545, 70** , 2005. ISSN LNL Annual Report, ISSN 1828-8545, 70.

30. A. KLUGE ET AL. **An ultra fast 100ps, 100micron 3D-pixel imager**. Proceedings of the 2008 SPIE conference. San Jose, USA: SPIE. 2009.

31. P. JARRON ET AL. **Pixel Read-Out Architectures for the NA62 GigaTracker**. TWEPP08. Naxos

(Greece): Heliotopos Conferences. 2008.

32. CHRISTIANSEN, J.; MOTA, M. A High-Resolution Time Interpolator Basedon a Delay Locked Loop. **IEEE Journal of Solid-State Circuits**, v. 34, n. 10, p. 1360-1366, Oct. 1999.

33. MESTER, C. **A multi-channel 24.4 ps bin size Time-to-Digital Converter for HEP applications**. TWEPP08. Naxos, Greece: Heliotopos Conferences. 2008.

34. HO, R. Efficient on-chip global interconnects. **Symp. VLSI Circuits**, 2003.

35. BASHIRULLAH, R. A hybrid current/voltage mode on-chip signaling scheme with adaptative bandwidth capability. **IEEE Trans. VLSI**, 2004.

36. BARTH, P. Silicon fusion bonding for fabrication of sensors, actuators and microstructures. **Sens. Act. A,** v. 23, n. 1-3, p. 919-926, 1990.

37. J. THOME ET AL. CMOSAIC: 3D Stacked Architectures with Interlayer Cooling. Disponivel em: <http://www.nano-tera.ch/projects/67.php>.

38. Y. TEMIZ ET AL.. **Two-Phase Cooling of 3D Chip Stacks**. Proceedings. of the EPFL-CMI MicroNanoFabrication Annual Review Meeting. Lausanne: EPFL. 2010.

39. A. RADADIA ET AL.. The fabrication of all-silicon micro gas chromatography columns using gold diffusion eutectic bonding. **J. Micromech. Microeng.**, v. 20, n. 015002, 2010.

40. A. MAPELLI ET AL. Low material budget microfabricated cooling devices for particle detectors and front-end electronics. **to be published in Nucl. Phys. B,** 2010.

41. WALLIS, G.; POMERANTZ, D. Field Assisted Glass-Metal Sealing. **J. Appl. Phys**, v. 40, n. 10, p. 3946-3949, 1969.

42. S. ILIE ET AL. **Chemical and radiolytical characterization of perfluorocarbon fluids used as coolant for LHC experiments**. Geneva. (CERN - EDMS doc N°842110).

43. BRUUS, H. **Theoretical Microfluidics**. USA: Oxford University Press, 2008.

44. WEBB, R. L. **Effect of Manifold Design on Flow Distribution in Parallel Micro-Channels**. International Electronic Packaging Technical Conference and Exhibition. Maui, Hawai: ASME. 2003.

45. D0 COLLABORATION. The upgraded D0 Detector. **Nucl. Instrum. Methods A565,413**, 2006.

46. MINERVA COLLABORATION. The Physics Case and Technology of the MINERvA Experiment. **MINERvA Document 218-v4**, 2006.

47. BEZNOSKO, D. ET AL. Nuclear Science Symposium Conference Record - IEEE, 790-793, vol.2, 2004.

48. SAINT-GOBAIN. **http:** //www.detectors.saint-gobain.com.

49. F. RONCHETTI (ON BEHALF OF THE ALICE COLLABORATION). **Journal of Physics: Conference Series 160 012012**, 2009.

50. BUZHAN, P. ET AL. **Nucl. Instrum. Methods A504, 48**, 2003.

51. SADYGOV, Z. ET AL. **Nucl. Instrum. Methods A504, 301**, 2003.

52. GOLOVIN, V.; SAVELIEV, V. **Nucl. Instrum. Methods A518, 560**, 2004.

53. ANGELONE, M. ET AL. Silicon Photo-Multiplier radiation hardness tests with a beam controlled neutron source. **Submitted to Nucl. Inst. Meth. A**. ISSN arXiv: 1002.3480.

54. NASA. **http:** //outgassing.nasa.gov/.

55. A. ANTONELLI ET AL. Study of the out-gassing rate of the NA62 Large-Angle Photon Veto system. **Internal Note NA62-09-01** , May 2009.

56. JENNINGER, B.; VANDONI, G. A Vacuum system for NA62. **CERN Internal Note TE technical note (in preparation)**, Geneva, 2010.

57. BERGSMA, F. et al. Comparison of field calculations and measurements of a spectrometer magnet, v. A 361 (1995) 466-471, n. Nuclear Instruments and Methods in Physics Research A, 1995. ISSN SSDI 0168-9002(95)00142-5.

58. OPAL COLLABORATION, K. AHMET ET AL.. **Nucl. Instrum. Methods A, 305 (1991) 275**, 305, 1991. 275.

59. F. AMBROSINO ET AL. IEEE Nuclear Science Symposium Conf. Record. Honolulu: IEEE. 2007. p. N05-6. arXiv:0711.3398.

60. PALLADINO, V. Proc. XIII Conf. On Calorimetry in High Energy Physics (CALOR 2008). Pavia: Institut of Physics Publishing. 2008.

61. ANTONELLI, A. et al. Study of the outgassing rate of the NA62 Large Angle Photon Veto System. **Internal Note NA62-09-01**, May 2009.

62. FANTI, V. ET AL. **Nucl. Instrum. Methods A,** v. 574, p. 433, 2007.

63. G. ATOYAN ET AL. **Nucl. Instrum. Meth. A320**, 1992. 144-154.

64. G.S. ATOIAN ET AL. **Nucl. Instrum. Meth. A531**, 2004. 467.

65. B. HALLGREN ET AL. **Nucl. Instrum. Meth. A 419**, 1998. 680-685.

66. MOLDED scintillators technical data. **IHEP Protvino**, 2010. Disponivel em:

<http://www.ihep.su/scint/mold/product-e.htm>.

67. BREKHOVSKIH, V. ET AL. **LHCb Calorimeter Technical Note 2000-40**, Geneva, 2000.

68. BREKHOVSKIH, V. ET AL. **LHCb Calorimeter Technical Note 2000-39**, 2000.

69. ALLSPACH D. ET AL. **CKM Vacuum Veto System Detector Outgassing Studies**. Fermilab. Chicago. March 2003. (Technical Memorandum CKM-81).

70. G.S. ATOIAN ET AL. **Nucl. Instrum. Meth. A 584**, 2008. 291-303.

71. B. HALLGREN ET AL. For the CPD (Calorimeter Pipeline Digitizer module) see: The NA48 LKr calorimeter digitizer electronics chain - Wire Chamber Conference 1998, Vienna. **CERN preprint EP/98-48**, Geneva, 1998.

72. MOVCHAN, S. **Straw Tracker Prototype for the precise measurement of very rare decay K->pi + nn**. 8th. International Conference on POsition Sensitive Detectors. Glasgow: PSD8. 2008. p. 307-309.

73. BLUM, W.; RIEGLER, W.; ROLANDI, L. **Particle Detection with Drift Chambers**. 2nd. ed. Berlin: Springer, 2008.

74. CATIANCCIO, A. PIPES UNDER INTERNAL PRESSURE AND BENDING. **CERN PH-EP-Tech-Note-2009-004**, Geneva.

75. BONIVENTO, W. et al. Development of the CARIOCA front-end chip for the LHCb muon detector, v. 491, n. 1-2, P233, 2002.

76. MORAES, D. ET AL. The CARIOCA Front End Chip for the LHCb muon chambers. **Internal Note LHCb-MUON 2003-009**.

77. ROMANIOUK, A. **Toshiba wire validation**. [S.l.]: https://edms.cern.ch/document/113748/1.

78. DOROVICH D. ET AL. Device for measuring of wire tension in drift tubes. **JINR, 2001, P13-2001-201, 11 p**.

79. WASEM, A. **Two Gas Single Pass Mixer**. CERN. Geneva.

80. BURCKHART, H. ELMB. **CERN**. Disponivel em: <http://elmb.web.cern.ch/ELMB/ELMBhome.html>.

81. DANIELSSON, H. Procedures and test results from the assembly of the 64-straw prototype. **CERN**. Disponivel em: <https://edms.cern.ch/document/1085111/1>.

82. T. AKESSON ET AL. Aging studies for the ATLAS Transition Radiation Tracker (TRT). **Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment**, v. 515, n. 1-2, p. 166-179, December 2003.

83. SHEMANSKI, D. E. CO2 Extinction Coefficient 1700-3000A. **The Journal of Chemical Physics, V56, N4**, 1972.

84. ANGHINOLFI, F. ET AL. NINO: an ultrafast low-power front-end amplifier/discriminator ASIC designed for the multigap resistive plate chambers. **Nucl. Instrum. Methods A533**, 2004. 183-187.

85. ANZIVINO, G. ET AL. Construction and test of a RICH prototype. **Nucl. Instrum. Methods A593**, 2008. 314-318.

86. ANGELUCCI, B. ET AL. Pion-Muon separation with a RICH prototype for the NA62 experiment. **Nucl. Instrum Methods A621**, 2010. 205-211.

87. HTTP://TTC.WEB.CERN.CH/TTC/INTRO.HTML. **TTC:** Timing, Trigger and Control Systems for LHC Detectors.

88. SOZZI, M. A concept for the NA62 Trigger and Data Acquisition. **Internal Note NA62-07-03**, April 2007.

89. KRIVDA, M. Information on the ALICE trigger module. Disponivel em: <http://epweb2.ph.bham.ac.uk/user/krivda/alice/>.

90. TAYLOR, B. G. **TTC laser transmitter (TTCex, TTCtx, TTCmx)**. [S.l.]. (User Manual (old TTCex version, new version in preparation, August 2010)).

91. CHRISTIANSEN, J. ET AL. Receiver ASIC for Timing, Trigger and Control distribution in LHC experiments. **IEEE Trans. Nucl. Sci., 43**, 1996. 1773-1777.

92. MOREIRA, P. QPLL manual; http://proj-qpll.web.cern.ch/proj-qpll/images/qpllManual.pdf. **CERN-EP/MIC**, 2005.

93. TTCRQ mezzanine board. **CERN-EP/MIC; http://proj-qpll.web.cern.ch/proj-qpll/ttcrq.htm**.

94. C. GASPAR ET AL. **DIM: Distributed Information Management System**. CERN. Geneva. (http://dim.web.cern.ch/dim/).

95. HAEFELI, G. ET AL. TELL1, specification fo a common readout board for LHCb. **LHCb note 2003-007 - http://lphe.epfl.ch/tell1/**, 2005.

96. WIEDNER, D. Optical 12 input Receiver Card IF14-1 for the LHCb TELL1 Board. **CERN / EDMS**, 2006. Disponivel em: <https://edms.cern.ch/document/758517/1>.

97. COLLAZUOL, G. ET AL. **Proc. 11th EUROMICRO Conference on Digital System Design Architectures, Methods and Tools**. Parma 3-5 Sept.2008. 2008 (DSD08). (405, IEEE).

98. CHRISTIANSEN, J. HPTDC, High Performance Time to Digital Converter version 2.2. **CERN-EP/MI; March 2004; http://tdc.web.cern.ch/tdc/hptdc/hptdc.htm**, 2004.

99. ANGELUCCI, B. **Master thesis, University of Pisa**. Pisa. 2010.

100. CHRISTIANSEN J. ET AL. (RD12 COLLABORATION). TTCrx reference manual. **CERN EP/MIC note; http://ttc.web.cern.ch/TTC/TTCrx_manual3.9.pdf**, 2004.

101. ANALOG Devices AD9252 data sheet. **Analog Devices**, 2010. ISSN http://www.analog.com/static/imported-files/data_sheets/AD9252.pdf.

102. HALLGREN, B. et al. **The NA62 Liquid Krypton Calorimeter Data Acquisition Upgrade**. Dresden: IEEE. 2008.

103. Z. GUZIK ET AL. RISC mezzanines for controlling data acquisition in the NA48 experiment at CERN. **Nucl. Instrum. Methods A452**, 2000. 289.

104. G. LAMANNA ET AL. **GPUs for fast triggering and pattern matching at the CERN experiment NA62**. Vienna. (Vienna Conference on Instrumentation 2010 submitted to Nucl. Instr. Meth. Phy. Res. A)).

105. P.S. COOPER ET AL. **Charged Kaons at the Main Injector**. FNAL. Chicago. 2001.

NA62 Acronyms and Abbreviations

ADC	Analog to Digital Converter
AM	Absorber Module
APD	Avalanche PhotoDiode
ASIC	Application Specifc Integrated Circuit
BCRST	Bunch Counter ReSeT
BEATCH	Progran to provide coordinates of all beam elements as input for alignment
BEND	BENDing magnet or dipole
BIF	Barrier Improvement Factor
CALB	Sub-detector data format for calibration
CCPC	Credit-Card PC: commercial processor on TELL1/TEL62 boards
CEDAR	Cerenkov Differential counter with Achromatic Ring Focus: differential Cerenkov detector developed at CERN
CHANTI	Charged ANTI
CHOD	Charged HODoscope
CKM	Cabibbo–Kobayashi–Maskawa matrix
CM	Circulation Module
COLL	COLLimator
COND	Condition data of the detector (extracted from DCS)
CONF	configurations of the run, including active detectors, trigger configurations, beam conditions, etc.
COST	calibration parameters computed by calibration tasks running on (partially or fully) reconstructed data.
COTS	Commercial-Of-The-Shelf
CPD	Calorimeter Pipeline Digitizer module
CREAM	Calorimeter Readout Module
CTL	Chamber Trigger Logic for straws L0 trigger system
DAQ	Data Acquisition System
DCS	Detector Control System
DDR2	Double Data Rate SDRAM (memory chips)
DIM	Distributed Information Management system
DLL	Delay Locked Loop
DLS	Data Logging System
DM	Distribution Module
DPRAM	Dual Ported RAM
DRAM	Dynamic RAM
EB	Event-Building
ECN3	Experimental Cavern housing the NA62 experiment
ECRST	Event Counter ReSeT
EDX	Energy-Dispersive X-ray spectroscopy
EE	End of Ejection
EOB	End Of Burst
EOC	End Of Column option for GTK chip architecture

EoC	End-of-Column part in the P-TDC chip architecture of the GTK
EOF	End Of Frame
FADC	Flash Analog to Digital Converter
FE	Front-End
FEE	Front-End Electronics
FEM	Finite Element Model
FIFO	First In First Out buffer
FISC	Filament Scanner, a beam profile detector inside the beam vacuum system
FNAL-NICADD	Fermi Nat Lab - Photo injector Lab (18Mev electron linac)
FPGA	Field Programmable Gate Array
FR	Fast Reconstruction
GbE	Gigabit Ethernet
GIM	Glashow–Iliopoulos–Maiani mechanism which suppresses flavour-changing neutral currents)
GOL	Gigabit Optical Link transmitter
GPN	General Purpose Network
GPU	Graphic Processing Unit
GTK	GigaTracker
HALO	A beam simulation program to calculate muon HALO rates
HPTDC	High Performance Time to Digital Converter
HV	High Voltage
IRC	Intermediate Ring Calorimeter
JTAG	Joint Test Action Group protocol
L0	Level 0 Trigger
L0TP	Level 0 Trigger Processor
L1	Level 1 trigger
L1TP	Level 1 Trigger Processor
L2	Level 2 trigger
LAV	Large Angle Veto
LED	Light-Emitting Diode
LG	Lead Glass
LGTS	Lead Glass Test Station
LKr	Liquid Kripton calorimeter
LTU	Local Trigger Unit
LV	Low-Voltage
LVDS	Low-Voltage Differential Signalling
MBPL-TP	Dipole Bending Magnet with Tapered Pole
MEPs	Multi-Event Packets
MIP	Minimum ionising particle
MNP33	NA62 Experimental Magnet
MUV	Muon Veto System
NAHIF	North Area High Intensity Facility
NIM	Nuclear Instrumentation Module
NINO	Fast front-end preamplifier-discriminator chip developed by ALICE
NNLO	Next-to-next -to-leading order
PCB	Printed Circuit Board

PDE	Photon Detection Eff.
Pe	Photo-electron
PECL	Positive Emitter-Coupled Logic
PEI	PolyEtherImide
PET	PolyEthylene Terephthalate
PMT or PM	Photomultiplier Tube
POPOP	1,4-bis(5-phenyloxazol-2-yl) benzene organic scintillator
PP-FPGA	Pre-Processing FPGA in the TELL1/TEL62 boards
PPO	2,5-Diphenyloxazole organic scintillator
PTP	Para-TerPhenyl
PVSS	Object-oriented process visualization and control system by ETM (a commercial SCADA system)
QCD	Quanten Chromodynamik
QDR-II	Quad Data Rate II memory
QPLL	Quartz-crystal based Phase-Lock Loop
QUAD	QUADrupole
RECO	Data format for fully reconstructed events
RICH	Ring Imaging Cherenkov
SAC	Small Angle Calorimeter
SAV	Small Angle Veto
SDRAM	Synchronous Dynamic Random Access Memory
SEM	Scanning Electron Microscope
SiPM	Silicon PhotoMultiplier
SOB	Start of Burst
SL-FPGA	Sync-Link FPGA in the TELL1/TEL62 boards
SLM	Smart Link Modules
SM	Standard Model of particle physics
SPI	Serial Port Interface
SPR	Single Photoelectron Response
SRAM	Static Random Access Memory
SRB	Straw Readout Board
TCC8	Target Chamber Cavern upstream of ECN3
TDAQ	Trigger and Data Acquisition system
TDC	Time to Digital Converter
TDCB	Time to Digital Converter Board
TDCC-FPGA	TDC Controller FPGA in the TDC boards
TEL62	Trigger and Data Acquisition board developed for NA62, based on TELL1 design
TELL1	Trigger ELectronics for L1 trigger: readout board developed by LHCb
THIN	Data format for summary data from fully reconstructed events
TRIM	Steering Magnets for the Beam
TTC	Timing, Trigger and Control
TTCex	TTC encoder VME board developed by CERN
TTCrq	TTC receiver mezzanine card developed by CERN
TTCrx	TTC receiver ASIC developed by CERN
TURTLE	Trace Unlimited Rays Through Lumped Elements, a beam tracking and simulation program

VME	Electronic bus and rack standard
VTL	View Trigger Logic for straws L0 trigger system
WE	Warning of Ejection
WLS	Wave-Length Shifting
WWE	Warning of Warning of Ejection