NEUTRAL PION TRANSITION FORM FACTOR MEASUREMENT AND RUN CONTROL AT THE NA62 EXPERIMENT

by

Nicolas Lurkin

A thesis submitted for the degree of
Doctor of Philosophy

School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham

December 2016

Abstract

The measurement of the π^0 electromagnetic transition form factor (TFF) slope a is performed in the time-like region of momentum transfer using a sample of 1.1×10^6 $\pi^0 \to e^+e^-\gamma$ Dalitz decays collected at the NA62-R_K experiment in 2007. The event selection, the fit procedure and the study of the systematic effects are presented. The final result obtained

$$a = (3.68 \pm 0.51_{\rm stat} \pm 0.25_{\rm syst}) \times 10^{-2} \quad (1)$$

is the most precise to date and represents the first evidence of a non-zero π^0 TFF slope with more than 3σ significance.

The NA62 experiment based at the CERN SPS is currently taking data and aims at measuring the branching fraction of the $K \to \pi\nu\bar{\nu}$ ultra-rare decay with 10 % precision and less than 10 % background. A complex trigger and data acquisition system is in place to record the data collected by the various detectors in use to reach this goal. The Run Control system of the experiment is meant to supervise and control them in a simple transparent way. The choices made to address the requirements for the system and the most important aspects of its implementation are discussed.

DECLARATION OF AUTHOR'S CONTRIBUTION

The theoretical introduction in chapter 1 and the detector description in chapter 2 are included for completeness and to support my own work. With the exception of the software package and the π^0 Dalitz decay MC generator, the neutral pion form factor measurement presented in chapter 3 is the result of my own work. I presented the preliminary results of this work at the HQL 2016 conference[1], published in the proceedings [1]. The final result reported is published in [2].

The NA62 detector, trigger and data acquisition (TDAQ) are presented in chapter 4 to contextualise the work I have done for NA62. The Run Control was presented at the IEEE 2013[2], BEACH 2014[3] and ICHEP 2016[4] conferences and published in the proceedings [3, 4]. It is described in chapter 5 and is entirely the result of my work; I am solely responsible for the design, developement and maintenance of the whole system. In this context I have worked in close collaboration with the trigger and data acquisition working group. Some of the design choices made for the TDAQ are the result of a joint reflection, in particular the interface between the Run Control and the devices involved in the TDAQ process, which I developed myself. I also provide multiple services linked to the data taking to the collaboration: the run conditions database, website and information retrieval tools, shifts database and shifters training, software library for the configuration mechanism, and

[1] XIII International Conference on Heavy Quarks and Leptons, 22–27 May, 2016, Blacksburg, USA

[2] 2013 IEEE Nuclear Science Symposium and Medical Imaging Conference (2013 NSS/MIC), 27 Oct–2 Nov, 2013, Seoul, South Korea

[3] XI International Conference on Hyperons, Charm and Beauty Hadrons, 21–26 Jul, 2014, Birmingham, United Kingdom

[4] 38th International Conference on High Energy Physics, 3–10 Aug, 2016, Chicago, USA

the analysis framework. I have been participating in the data taking campaigns of NA62 since 2012, when I delivered the first operational version of the Run Control.

No other qualifications have been achieved with the work reported in this thesis.

ACKNOWLEDGEMENTS

I would like to express my gratitude to my supervisors Cristina Lazzeroni and Evgueni Goudzovski for the help and support they provided since my arrival in Birmingham, and for their comments and suggestions during the redaction of this thesis. Thanks to Michal Koval, with whom it was a pleasure to work on this analysis, for the many useful discussions and meetings we had. I am also grateful to Fernando Varela Rodriguez for the invaluable help he provided during the initial development of the Run Control at CERN.

I also address a special thanks to my friends and colleagues which made my stay in Birmingham so pleasant and interesting: Angela, Antonino, Francesco, Kostas, Maria-Brigida, Lorenza and Laura. I enjoyed all my time in the NA62 collaboration, surrounded by so many smart and friendly people, spontaneous dinners in the control room, and learning Italian language and culture. Many thanks for the good laughs from the BS group. Thanks to my family and all the friends I met along my PhD years for supporting me even though they could not always understand what I tried to explain to them about my work.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

INTRODUCTION

The Standard Model (SM) of particle physics has been extremely successful in predicting and describing a large number of phenomena over a wide range of energies. It has also been tested to a high degree of precision. However a certain number of questions remain unanswered within this framework and call for new physics. The search for physics beyond the Standard Model (BSM) can be conducted following two different approaches. The LHC experiments at CERN are increasing the energy to produce possible new massive particles. Other experiments are focusing on SM predictions to detect any small deviation from the expected values. To be successful this approach requires excellent precision both on the theoretical and experimental sides.

The latter approach is the one followed by the NA62 experiment at the CERN Super Proton Synchrotron (SPS) [5]. The kaon has been discovered in cosmic rays in 1947 [6] and has been at the heart of many important developments in particle physics over the years. One of the most important has been the discovery of CP violation by the weak interaction: first indirect CP violation in 1964 [7] and later direct CP violation in 1999 [8]. Although already extensively studied in the past and despite the recent shift of interest towards heavier flavours, the kaon system remains an interesting probe of the SM.

Two phases of the experiment are used for this work. The NA62-R_K phase took place in 2007–2008 with a setup similar to the previous experiment NA48/2. The main goal was a test of lepton universality with the measurement of the $R_K = \Gamma(K^\pm \to e^\pm\nu)/\Gamma(K^\pm \to \mu^\pm\nu)$ ratio. The data collected were also used to measure the electromagnetic transition form factor of the π^0 as reported in this thesis.

Following the NA62-R_K data taking campaign, most of the experimental setup has

been dismantled and replaced by new systems, and the beam line has been upgraded. The main goal of the new phase is to measure the branching ratio of the very rare decay $K^+ \to \pi^+\nu\bar{\nu}$. This decay is forbidden at tree level and occurs through Flavour Changing Neutral Current (FCNC) loops. This process is therefore highly suppressed. It is also theoretically very clean: short distance contributions dominate (i.e. heavy particles in the loop) and the hadronic matrix elements can be related to other well measured processes. The SM prediction is $\mathcal{B}(K^+ \to \pi^+\nu\bar{\nu})_{\mathrm{SM}} = (9.11 \pm 0.72) \times 10^{-11}$ [9], where the main error comes from the CKM matrix element. The best measurement for this branching ratio is $\mathcal{B}(K^+ \to \pi^+\nu\bar{\nu})_{\mathrm{exp}} = (1.73^{+1.15}_{-1.05}) \times 10^{-10}$, extracted from the 7 detected events at E787 and E949 experiments at Brookhaven National Laboratory [10]. The achieved precision is not sufficient to conclude on any deviation from the SM.

NA62 has been designed to reach a similar precision to that of the SM by detecting about 50 $K^+ \to \pi^+\nu\bar{\nu}$ SM events per year of data taking with 10% background.

This thesis is organized around the two independent topics carried out in the framework of the NA62 experiment: the first one is the measurement of the π^0 electromagnetic transition form factor using the data collected during the NA62-$\mathrm{R_K}$ phase, and the second one is the development of the run control. The first chapter introduces the motivations and theoretical framework to the physics measurement with the 2007–2008 data and the main objective of NA62 for the 2016–2018 period. The second chapter describes the relevant parts of the experimental setup during the NA62-$\mathrm{R_K}$ phase, including the beam line, detector, trigger configuration and software. The third chapter focuses on the measurement of the π^0 transition form factor performed by the author, including the selection of the data sample, the results obtained and the description of the procedures employed to assess the systematic uncertainties. The fourth chapter describes the NA62 experimental setup for the $K^+ \to \pi^+\nu\bar{\nu}$ measurement, with a particular emphasis on the trigger and data acquisition system, for which most of the work accomplished is dedicated. The last chapter gives a detailed description of the run control of the NA62 experiment, which has been developed by the author of this work.

CHAPTER 1

THEORETICAL FRAMEWORK

1.1 The π^0 Dalitz decay

The neutral pion is the lightest meson in the standard model. Contrary to the charged pions, also discovered in cosmic rays in 1947 [11], the neutral one does not leave a track in the photographic emulsions used for observations. Despite being already inferred from the observation of its decay products from cosmic rays, the first direct detection occurred in 1950 at the University of California's cyclotron [12]. Being amongst some of the effective mediators of the strong nuclear interaction, as proposed by Yukawa in 1935 [13], the pions play an important role in the understanding of the low energy properties of this force. The advent of the quark model in the 1960s [14, 15] allowed identification of the π^0 meson as a composite particle made of a linear combination of the $u\bar{u}$ and $d\bar{d}$ quarks. As for all mesons, the complex internal interactions due to their compositeness are described by the theory of quantum chromodynamics (QCD), for which the fundamental degrees of freedom are the quarks. However due to its non-perturbative nature at low energy, exact computations are difficult, but observations taking place at energies lower than the typical energies of the underlying physics do not see the details of the interaction. Instead the internal interactions can be described as an effective theory, for which the structural properties and the interactions with other particles are encoded in form factors. Because the quarks themselves carry an electric charge, this internal structure also influences the

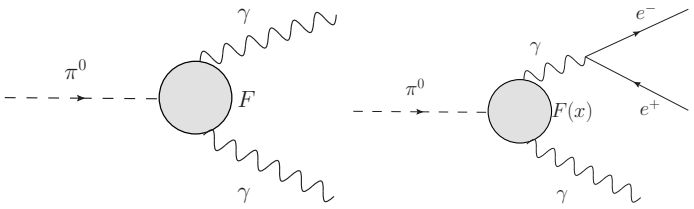

Figure 1.1: Leading order Feynman diagrams for the $\pi^0 \to \gamma\gamma$ process in (a) and for the $\pi^0 \to e^+e^-\gamma$ process in (b).

electromagnetic interactions of the meson. The electromagnetic transition form factor (TFF) describes the electromagnetic structure reflecting the quark distribution in the meson and appearing in the electromagnetic interactions with other particles. The cross-section and differential decay width can be written in the form:

$$\frac{d\sigma}{dq^2} = \frac{d\sigma}{dq^2}\bigg|_{\text{p-l}} \left|\mathcal{F}(q^2)\right|^2, \quad \frac{d\Gamma}{dq^2} = \frac{d\Gamma}{dq^2}\bigg|_{\text{p-l}} \left|\mathcal{F}(q^2)\right|^2 , \tag{1.1}$$

where q^2 is the square of the photon four-momentum transfer and $\mathcal{F}$ is the TFF that can be measured by comparing experimental data with the theoretical quantum electrodynamic (QED) computations for point-like (p-l) particles [16]. If the transferred photon is off-shell ($q^2 \neq 0$), two momentum transfer regions are defined: *time-like* for $q^2 > 0$, and *space-like* for $q^2 < 0$.

The neutral pion has a short lifetime of $(8.52\pm0.18)\times10^{-17}$ s and decays predominantly through the electromagnetic interaction into two photons. The branching ratio for this decay is $\mathcal{B}(\pi^0 \to \gamma\gamma) = (98.823\pm0.034)\,\%$ [17]. The second biggest decay, which was studied by Richard H. Dalitz [18], is the Dalitz decay of the π^0 with $\mathcal{B}(\pi^0 \to e^+e^-\gamma) = (1.174\pm0.035)\,\%$, indicated as π^0_D. These two modes account for most of the decays and all other (observed) channels ($\pi^0 \to e^+e^+e^-e^-$, $\pi^0 \to e^+e^-$) are at least two orders of magnitude smaller. As shown in the Feynman diagram in Figure 1.1, the Dalitz decay comes from the same $\pi^0\gamma\gamma$ vertex as the main channel, but one of the photons is off-shell and further converts into an electron-positron pair.

The following independent variables are commonly used to describe the kinematics of

this decay:

$$x = \left(\frac{M_{ee}}{m_{\pi^0}}\right)^2 = \frac{(p_{e^+} + p_{e^-})^2}{m_{\pi^0}^2}, \quad y = \frac{2p_{\pi^0} \cdot (p_{e^+} - p_{e^-})}{m_{\pi^0}^2(1-x)} \,, \tag{1.2}$$

with the accessible kinematic region defined as

$$r^2 = \left(\frac{2m_e}{m_{\pi^0}}\right)^2 \leq x \leq 1, \quad |y| \leq \beta = \sqrt{1 - \frac{r^2}{x}} \,, \tag{1.3}$$

where m_{π^0} and m_e are respectively the masses of the π^0 and the $e^{\pm}$, M_{ee} is the invariant mass of the electron-positron pair, and the p are the four-momenta of the corresponding particles.

The π^0 electromagnetic transition form factor is an input to some theoretical computations. Two of them are of particular interest as they currently show some tension between theory and experiment. The measured value of the anomalous magnetic moment of the muon $(g-2)_\mu$ is 3σ away from the theoretical expectations [19]. This discrepancy is drawing a lot of attention, and many efforts are devoted to reducing both experimental and theoretical uncertainties. In this context, the second largest theoretical uncertainty comes from the hadronic light-by-light (HLbL) scattering, which features diagrams containing $P\gamma^*\gamma^*$ vertices, where $P = \pi^0, \eta, \eta'$. The π^0 contribution, using the π^0 TFF as an input, accounts for 25 % of the uncertainty coming from the HLbL term [20].

The second tension is the branching fraction of the $\pi^0 \to e^+e^-$ decay, showing a 3.3σ discrepancy between experiment and the predictions. The validity of the model-dependent extrapolation of the π^0 TFF from high energy to low energy (see subsection 1.1.2), used as an input to the theoretical computation of this branching fraction, is questioned [21]. However, there is no alternative model-independent measurement of the π^0 TFF precise enough to confirm the validity of the extrapolation.

The π^0 electromagnetic transition form factor in its general form $\mathcal{F}_{\pi^0\gamma^*\gamma^*}(q_1^2, q_2^2)$, encodes the physics of the $\pi\gamma^*\gamma^*$ vertex for two off-shell photons with four-momenta q_1 and q_2. In the limit where $q_1, q_2 \to 0$ and both photons are real, the vertex becomes equivalent

to the $\pi\gamma\gamma$ one and [22]

$$\Gamma(\pi^0 \to \gamma\gamma) = \left(\frac{m_{\pi^0}}{4\pi}\right)^3 \left(\frac{\alpha}{F_\pi}\right)^2 = \frac{1}{64\pi} m_{\pi^0}^3 \left|\mathcal{F}_{\pi^0\gamma^*\gamma^*}(0,0)\right|^2 \ . \tag{1.4}$$

The following relation gives the normalisation

$$\mathcal{F}_{\pi^0\gamma^*\gamma^*}(0,0) \equiv \mathcal{F}_{\pi^0\gamma\gamma} = \frac{\alpha}{\pi F_\pi} \ , \tag{1.5}$$

where α is the electromagnetic fine-structure constant and F_π is the π^0 decay constant. The semi off-shell form factor is obtained by taking the limit $q_2 \to 0$ and the following definition is obtained: $\mathcal{F}_{\pi^0\gamma\gamma^*}(q^2) \equiv \mathcal{F}_{\pi^0\gamma^*\gamma^*}(q_1^2, 0)$. One can define the normalized TFF $\mathcal{F}(x) = \mathcal{F}_{\pi^0\gamma\gamma^*}(xm_{\pi^0}^2)/\mathcal{F}_{\pi^0\gamma\gamma}$. The leading order (LO) differential decay width can then be written in terms of the $\pi^0 \to \gamma\gamma$ total decay width $\Gamma(\pi^0 \to \gamma\gamma)$ [22]:

$$\frac{d^2\Gamma^{\mathrm{LO}}(\pi_D^0)}{dxdy} = \Gamma(\pi^0 \to \gamma\gamma)\frac{\alpha}{4\pi}\frac{(1-x)^3}{x}\left(1 + y^2 + \frac{r^2}{x}\right)\left|\mathcal{F}(x)\right|^2 \ , \tag{1.6}$$

with x as defined in eq. (1.2). The TFF at low values of q^2 is usually written in the expanded form

$$\mathcal{F} = 1 + ax + \mathcal{O}(x^2) \ , \tag{1.7}$$

where a is the form factor slope parameter.

The shape of the TFF qualitatively follows the distribution given in Figure 1.2. Three distinct regions can be defined. Region I (space-like momentum transfer, $q^2 < 0$) is accessed in the π^0 production cross-section in electron scattering processes (Figure 1.3a), where only momentum is transferred. Region II ($4m_e^2 < q^2 < m_{\pi^0}^2$) is the one of interest for this analysis, specifically accessed in the Dalitz decay of the π^0 (Figure 1.1b). Region III ($q^2 > m_{\pi^0}^2$) is accessible in $e^\pm$ annihilation into π^0 (Figure 1.3b). These last two are said to be time-like, where energy is transferred but no momentum. The experimental techniques detailed in subsection 1.1.2 are specific for each of them, offering complementary ways to measure the form factor over the whole physically accessible range of q^2. In the time-like

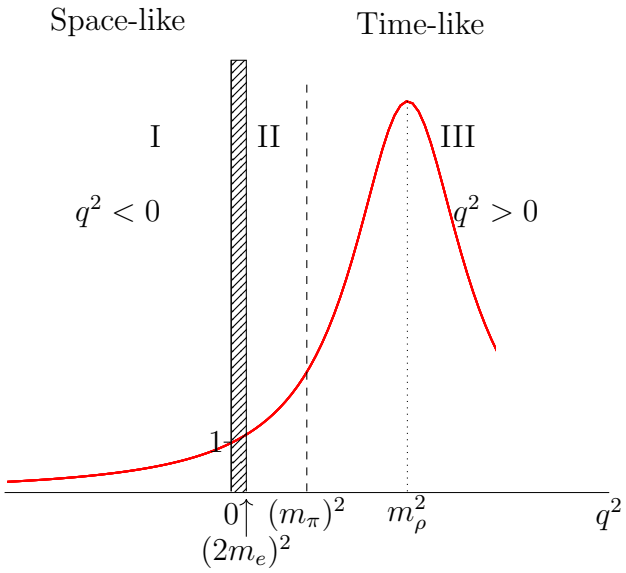

Figure 1.2: General behaviour of the π^0 electromagnetic transition form factor over the accessible physical range of q^2. The TFF is measured for region I and III in e^+e^- experiments, through π^0 creation in scattering and annihilation processes respectively. The region II allows the extraction of the TFF slope from the Dalitz decay of the π^0. The shaded area is kinematically prohibited. The TFF peaks at the neutral vector meson masses due to resonances.

region of momentum transfer, a resonant interaction between the photon and the light vector mesons appears. The value of the form factor is enhanced as q^2 approaches their masses. Building on this effect the vector meson dominance model (VMD) [23] describes the interaction between photons and hadrons in terms of exchange of vector mesons, using them as mediators of the interaction. Using this model M. Gell-Mann and F. Zachariasen [24] estimated that the form factor is dominated by the two lightest resonances: ρ and ω. The estimate for the slope parameter value is: $a \approx m_{\pi^0}^2(m_\rho^{-2} + m_\omega^{-2})/2 \approx 0.03$, a small value that justifies the expansion in eq. (1.7). Figure 1.4 shows the x distribution and the additional slope induced by a form-factor of the size expected from VMD.

1.1.1 Radiative corrections

The next-to-leading order (NLO) radiative corrections to the π_D^0 decay width were originally studied by D. Joseph [25] and the contribution to the total decay width is of the order of 10^{-4}, while they are of the same order as the TFF slope in the differential decay

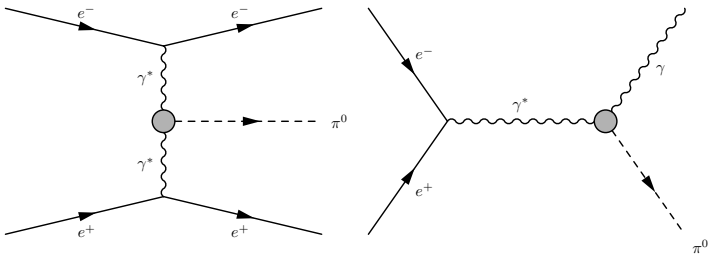

(a) (b)

Figure 1.3: Feynman diagrams for the π^0 production in $e^{\pm}$ scattering (a) and annihilation (b).

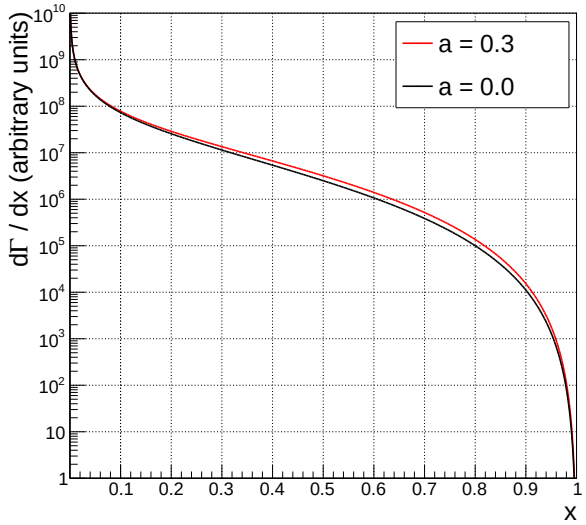

Figure 1.4: Differential decay width $d\Gamma/dx$ of the π^0_D decay for two different TFF slope values. The effect on the x spectrum can be seen for $x > 0.1$.

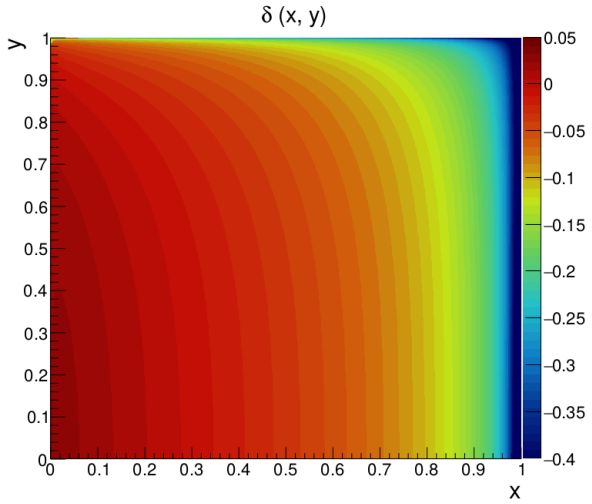

Figure 1.5: Size of the radiative corrections to the π^0_D decay over the x, y plane. The magnitude of the effect goes up to 40 % of the LO decay width and has a stronger variation in the x direction.

width. The radiative corrections in the differential decay width have been theoretically studied first by B.E. Lautrup and J. Smith [26] using the soft-photon approximation, to be soon extended by K.O. Mikaelian and J. Smith [27] to hard photon corrections.

The radiative corrections to the differential decay width are encoded in the function

$$\delta(x,y) = \frac{d^2\Gamma^{\mathrm{NLO}}}{dxdy} \Big/ \frac{d^2\Gamma^{\mathrm{LO}}}{dxdy} \ , \tag{1.8}$$

and introduced as a factor $(1+\delta(x,y))$ to the total differential decay width

$$\frac{d^2\Gamma(\pi^0_D)}{dxdy} = \frac{d^2\Gamma^{\mathrm{LO}}(\pi^0_D)}{dxdy}\,(1+\delta(x,y)) \ . \tag{1.9}$$

The map of values taken by this correction function in the x, y space is shown in Figure 1.5.

The radiative corrections can be classified according to their origin: virtual radiative corrections (Figure 1.6 top), one-photon irreducible contribution (Figure 1.6 middle), and bremsstrahlung (Figure 1.6 bottom):

$$\delta = \delta^{\mathrm{virt}} + \delta^{1\gamma\mathrm{IR}} + \delta^{\mathrm{BS}} \ . \tag{1.10}$$

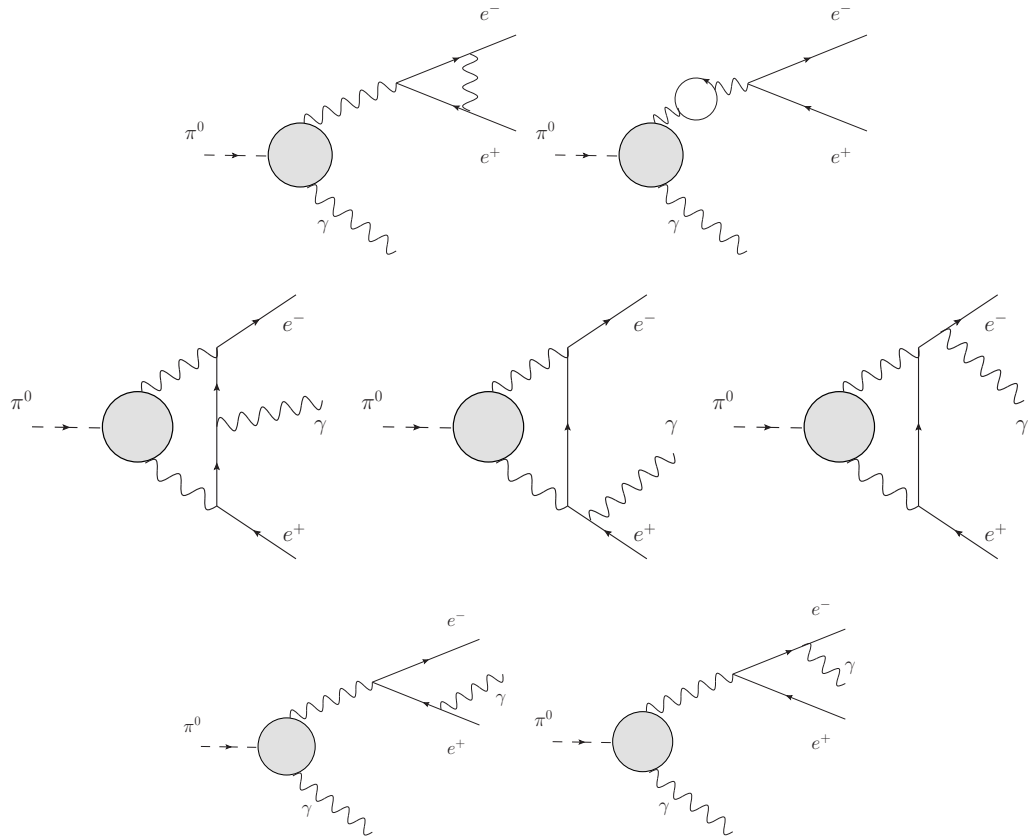

Figure 1.6: Top row: virtual radiative corrections to the π^0_D decay. Middle row: One-photon irreducible contribution to the π^0_D decay. Bottom row: Bremsstrahlung corrections to the π^0_D decay.

The most recent development on this topic, published by T. Husek and others in [22], includes the one-photon irreducible contribution, which was previously neglected. This contribution has an effect of $|\Delta a| \approx 0.5\times10^{-2}$ on the slope of the x spectrum. The remaining bias on the TFF slope introduced by ignoring higher order correction terms in the simulation of the x spectrum is, according to the author of [22], less than 0.01×10^{-2}. The influence of each component on the slope of the x variable can be seen in Figure 1.7.

The δ^{virt} component contains infrared divergences, which are cancelled by the contribution δ^{BS}. The proper treatment of the bremsstrahlung component is one of the crucial points that was required to stabilise the result of this analysis. This component is indeed a four-body decay $\pi^0_D \to e^+e^-\gamma(\gamma)$, described by five independent kinematic variables instead of two. The usual treatment is to integrate over the additional variables to obtain

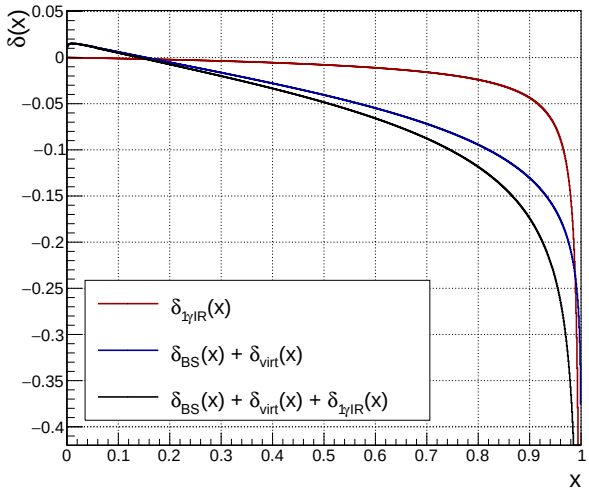

Figure 1.7: Effect of the different components of the radiative corrections on the x spectrum. The black line corresponds to the total radiative correction; the blue line corresponds to its component first computed in [27]; the red line shows the new contribution, introduced in [22].

$\delta^{\mathrm{BS}}(x, y)$. However, it appeared during the analysis that the extra photon is often energetic enough to be detected in the experiment, leading to an instability of the result if it is not included in the simulation. This issue triggered a collaboration with the theory group of the Charles University Prague which resulted in the publication [22].

The most relevant additional kinematic variable is the normalised invariant mass of the two photons

$$x_\gamma = \frac{\left(p_{\gamma 1} + p_{\gamma 2}\right)^2}{m_{\pi^0}^2} \,, \tag{1.11}$$

which is bounded by [22]

$$\frac{\lambda^2}{m_{\pi^0}^2} < x_\gamma \leq 1 + x - \sqrt{4x + y^2\beta^{-2}\left(1-x\right)^2} \,, \tag{1.12}$$

where β is defined in eq. (1.3) and λ is the infrared cut-off used to regularise the divergence. The physical limit corresponds to $\lambda \to 0$. The practical result of the collaboration with the Charles University Prague group is an algorithm providing a correct simulation of the radiative photon using a hybrid 3-body and 4-body generator. The bremsstrahlung contribution is split into two parts according to the x_γ variable, and controlled by the parameter x_γ^{cut}, which is the normalised di-photon invariant mass above which the radiative

photon will be simulated:

$$\delta^{\mathrm{BS}}(x,y) = \int_0^{\min(x_\gamma^{\mathrm{cut}},x_\gamma^{\mathrm{max}})} \delta^{\mathrm{BS}}(x,y,x_\gamma)dx_\gamma + \int_{\min(x_\gamma^{\mathrm{cut}},x_\gamma^{\mathrm{max}})}^{x_\gamma^{\mathrm{max}}} \delta^{\mathrm{BS}}(x,y,x_\gamma)dx_\gamma \ . \tag{1.13}$$

The first term is integrated from $x_\gamma = 0$ to x_γ^{cut} to cancel the infrared divergences. The following ratio

$$R(x_\gamma^{\mathrm{cut}}) = \frac{\int_{x_{\mathrm{min}}}^{x_{\mathrm{max}}} dx \int_{y_{\mathrm{min}}}^{y_{\mathrm{max}}} dy \int_0^{\min(x_\gamma^{\mathrm{cut}},x_\gamma^{\mathrm{max}})} \left[\Gamma^{\mathrm{LO}}(x,y) + \Gamma^{\mathrm{NLO}}(x,y,x_\gamma)\right] dx_\gamma}{\int_{x_{\mathrm{min}}}^{x_{\mathrm{max}}} dx \int_{y_{\mathrm{min}}}^{y_{\mathrm{max}}} dy \int_{\min(x_\gamma^{\mathrm{cut}},x_\gamma^{\mathrm{max}})}^{x_\gamma^{\mathrm{max}}} \Gamma^{\mathrm{LO}}(x,y)\delta^{\mathrm{BS}}(x,y,x_\gamma)dx_\gamma} \ , \tag{1.14}$$

is the total decay width in the lower x_γ region divided by the remaining bremsstrahlung contribution. It represents the fraction of soft radiative photons with respect to harder photons, where the two-photon final state has an invariant mass above the cut-off value x_γ^{cut}. The term

$$\Gamma^{\mathrm{NLO}}(x,y,x_\gamma) = \Gamma^{\mathrm{LO}}(x,y)\left[1 + \delta^{\mathrm{virt}}(x,y) + \delta^{1\gamma\mathrm{IR}}(x,y) + \delta^{\mathrm{BS}}(x,y,x_\gamma)\right] \tag{1.15}$$

contains all the radiative corrections. The ratio of 3-body to 4-body decays was computed numerically by the authors of [22] for the following three values used in the analysis:

$$R(0.05) = 38.366, \quad R(0.01) = 17.397, \quad R(0.005) = 13.700 \ . \tag{1.16}$$

The value $x_\gamma^{\mathrm{cut}} = 0.01$ has been empirically determined to give good results for this analysis. The other two surrounding values generate approximately half (0.05) and twice (0.005) as many soft photons and are used for the validation of the chosen value.

The algorithm implemented in the π_D^0 generator first chooses between the production of a 3-body (one photon) or 4-body (two photons) decay following the probability

$$P(\text{3-body}) = \frac{R}{R+1}, \quad P(\text{4-body}) = \frac{1}{R+1} \ . \tag{1.17}$$

In the 3-body decay case, the acceptance-rejection method [28] is used: the (x,y)

points are randomly picked following the differential decay width similar to eq. (1.9), where the radiative correction $\delta(x, y)$ is replaced by $\delta(x, y, x_\gamma)$. In this modified expression the radiative bremsstrahlung photon contribution $\delta^{\mathrm{BS}}(x, y, x_\gamma)$ is integrated up to x_γ^{cut}. The four-momenta of the e^+, e^- and γ are generated in the π^0 rest frame using the isotropy of the Dalitz decay (π^0 is spin-less) and the fact that the decay products' three-momenta lie in one plane due to momentum conservation. They are then boosted into the laboratory frame.

In the 4-body decay case, the four-momenta of all the final state particles ($e^+e^-\gamma\gamma$) are generated in the π^0 rest frame corresponding to a constant decay amplitude, using the GENBOD algorithm from the CERN software library [29]. Six kinematic variables (of which five are independent)

$$x = \frac{\left(p_{e^+} + p_{e^-}\right)^2}{m_{\pi^0}^2} , \qquad a_i = \frac{p_{e^+} \cdot p_{\gamma_i}}{m_{\pi^0}^2} , \tag{1.18}$$

$$x_\gamma = \frac{\left(p_{\gamma 1} + p_{\gamma 2}\right)^2}{m_{\pi^0}^2} , \qquad b_i = \frac{p_{e^-} \cdot p_{\gamma_i}}{m_{\pi^0}^2} , \tag{1.19}$$

where $i = 1, 2$ is the photon index, are determined and used to compute the squared invariant matrix element for the bremsstrahlung process, which is free of divergence above x_γ^{cut}. The acceptance-rejection method is used to select events according to the matrix element distribution. The momenta are then boosted into the laboratory frame.

When $x_\gamma^{\mathrm{cut}} \to 0$ the number of generated soft photons increases. The simulation will converge for x_γ^{cut} sufficiently small, as the additional photons will become undetectable and the event will thus be indistinguishable from a one-photon event. This is verified and the results are presented in subsection 3.7.2.

1.1.2 Experimental status

Different types of processes and experimental techniques are available to measure the TFF, depending on the momentum transfer region.

For the space-like momentum transfer region ($q^2 < 0$, region I in Figure 1.2) the form factor is obtained by the measurement of the π^0 production cross-section in electron scattering experiments as a function of $Q^2 = -q^2$. The process $e^+e^- \to e^+e^-\gamma^*\gamma^* \to e^+e^-\pi^0$ can be measured in a "non-tagged" or "tagged" way. In the first case the electrons in the final state are not detected, or even requested not to be detected. The virtual photons are nearly on-shell as their squared four-momenta are very low ($q_1^2 \approx q_2^2 \approx 0$). The exact momentum transfer is not known and the Lorentz boost of the π^0 is small.

In the case of the "tagged" events, one or both electrons are detected and the collision happens with a known q_1^2 and/or q_2^2. If one of the electrons is not tagged, the corresponding transferred momentum is small ($q_2^2 \approx 0$), corresponding to the limit where one of the photons is on-shell. However, due to acceptance considerations, the electrons must have a minimum scattering angle to be detected, restricting the use of this technique for higher Q^2 values.

Several experiments have measured the π^0 TFF in this regime and are shown in Figure 1.8a ([30–33]). This technique only allows measurements of $|\mathcal{F}(q^2)|$ far from small $|q^2|$, which is where the *slope* is defined. It is however possible to extrapolate the TFF from high Q^2 values to $Q^2 \to 0$ by choosing a model and making some theoretical assumptions, possibly introducing a systematic bias. The current world average of the value of the π^0 TFF slope, $a = 0.032 \pm 0.004$ [17] is dominated by the CELLO measurement, in which a VMD fit was applied to the space-like data ($0.5\,\mathrm{GeV}^2 < Q^2 < 2.7\,\mathrm{GeV}^2$). The fit was extrapolated to the limit $Q^2 = 0$ resulting in $a = 0.0326 \pm 0.0026_{\mathrm{stat}}$.

In the high part of the time-like momentum transfer region ($q^2 > m_{\pi^0}^2$, region III in Figure 1.2) the measurement of the form factor is performed through the production cross-section in the e^+e^- annihilation process. The cross-section is however small ($10^{-35}\,\mathrm{cm}^2$ to $10^{-34}\,\mathrm{cm}^2$), making the measurement a difficult experimental challenge. The SND and CMD-2 experiments [35, 36] successfully measured the cross-section for high q^2, from which the form factor was measured and is shown in Figure 1.8b.

Finally, in region II in Figure 1.2, $4m_e^2 < q^2 < m_{\pi^0}^2$, the inverse process is used: e^+e^-

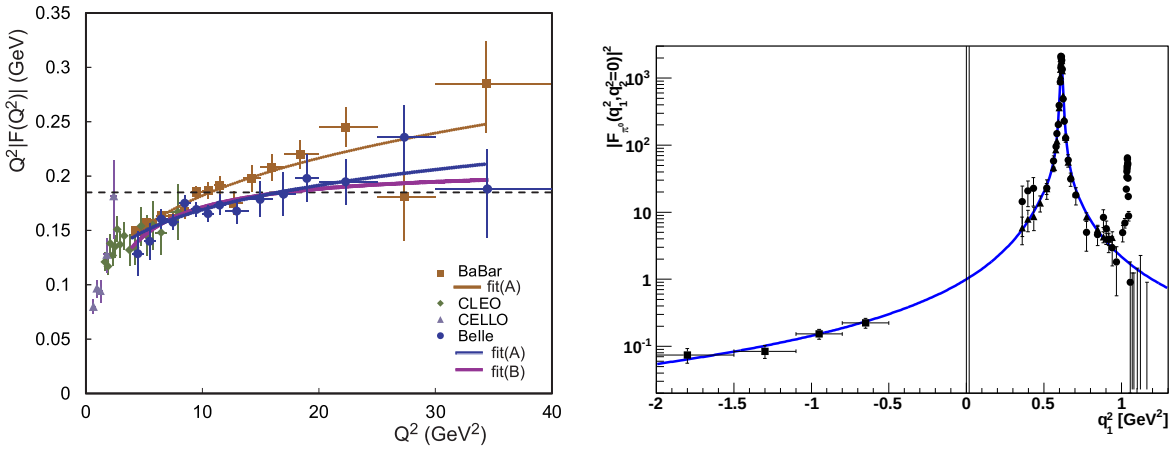

(a) (b)

Figure 1.8: (a) Measurements of the π^0 transition form factor in the space-like momentum-transfer region from electron scattering experiments. The plot taken from [33] shows the product $Q^2\,|\mathcal{F}(Q^2)|$, where $Q^2 = -q^2$. The parametrisation of the fitted curves can be found in [32, 33]. The dashed line shows the asymptotic prediction from perturbative QCD (∼185 GeV). (b) The plot taken from [34] shows the form factor in the time-like momentum-transfer region from the SND [35] and CMD-2 [36] experiments ($q^2 > 0$). The points at $q^2 < 0$ are the data points from the CELLO collaboration [30], as in (a).

pair-production from photon conversion in the Dalitz decay of the π^0: $\pi^0_D \to \gamma e^+ e^-$. This process naturally occurs at small q^2 values, providing an excellent model-independent way of determining the TFF slope. On the other hand, this measurement is limited by the small number of events at high x, which are the ones sensitive to the form factor slope, as seen in Figure 1.4. The first experiments had very limited statistics and did not include the radiative corrections. They were therefore not very successful and could not determine even the sign of the TFF slope. The Geneva-Saclay experiment (S118) performed at CERN in 1977 [37] was the first one including the radiative corrections and leading to a consistent result. The π^0_D was produced in the same decay chain as in the present analysis and the TFF slope measurement result was

$$a = 0.10 \pm 0.03_{\text{stat}} \; . \tag{1.20}$$

Other experiments used the $\pi^- p \to \pi^0 n$ or $\gamma p \to \pi^0 p$ processes as the π^0 production

mechanism and obtained the following results [38–41]:

$$a = -0.110 \pm 0.030_{\text{stat}} \pm 0.080_{\text{syst}} \tag{1.21}$$

$$a = \quad 0.026 \pm 0.024_{\text{stat}} \pm 0.048_{\text{syst}} \tag{1.22}$$

$$a = \quad 0.025 \pm 0.014_{\text{stat}} \pm 0.026_{\text{syst}} \tag{1.23}$$

$$a = \quad 0.030 \pm 0.010_{\text{tot}} \tag{1.24}$$

The average of these five measurements (eq. (1.20) to eq. (1.24)) is

$$a = 0.034 \pm 0.009 \,, \tag{1.25}$$

which is in agreement with the theoretical predictions but does not constitute a clear validation of a positive slope.

1.2 The $K^+ \to \pi^+ \nu\bar{\nu}$ decay

The data exploited in this thesis come from the 2007–2008 data-taking period of the NA62 experiment, and the physics studied is that of the pion. However, NA62 is primarily a kaon experiment, and the main aim of the current phase is to measure the branching ratio of the $K^+ \to \pi^+ \nu\bar{\nu}$ decay. The development of the Run Control, which is the subject of the last part of this thesis, as well as the author's participation in the data-taking campaign, were done in the context of the current phase of the experiment. This section introduces the kaon, the theory relevant to the $K \to \pi\nu\bar{\nu}$ decay, and the experimental context.

The kaon is the only strange meson in the standard model. The strangeness quantum number was associated to the kaon after noticing that it is abundantly produced in pairs of strange/anti-strange particles in strong interactions, while its long lifetime suggests the decay reaction is much less probable. This quantum number would be conserved by the

strong interaction while being violated by the weak interaction, and was later understood to be associated to the strange quark composing the kaon.

This phenomenon is now explained by the presence of three families of quarks and the Cabibbo-Kobayashi-Maskawa (CKM) matrix [42] relating the mass eigenstates d, s, b to their weak interaction eigenstates d', s', b', which couple to the mass and weak u, c, t eigenstates:

$$\begin{pmatrix} d \\ s \\ b \end{pmatrix} = V_{\mathrm{CKM}} \begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} . \tag{1.26}$$

The weak universality postulated by Cabibbo [43] requires the unitarity $V_{\mathrm{CKM}} V_{\mathrm{CKM}}^{\dagger} = \mathbb{1}$. After imposing this condition, four free parameters remain: three real parameters and one complex phase responsible for CP violation. The standard parametrisation decomposes the CKM matrix into three rotation matrices representing the mixing between generations:

$$V_{\mathrm{CKM}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{13}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{13}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{13}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{13}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{13}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{13}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{13}} & c_{23}c_{13} \end{pmatrix} , \tag{1.27}$$

where s_{ij} and c_{ij} are the sines and cosines of θ_{ij} $(i, j = 1, 2, 3)$, the mixing angles between the mass and weak eigenstates, and δ_{13} is the complex phase.

The unitarity condition imposes six constraints $\sum_i^{u,c,t} V_{ij} V_{ik}^* = 0$ on the off-diagonal terms of V_{CKM} $(j, k = d, s, b; j \neq k)$. They represent six unitarity triangles in a complex plane, where the area is proportional to the CP violating phase. The most commonly used triangle is

$$V_{ud} V_{ub}^* + V_{cd} V_{cb}^* + V_{td} V_{tb}^* = 0 . \tag{1.28}$$

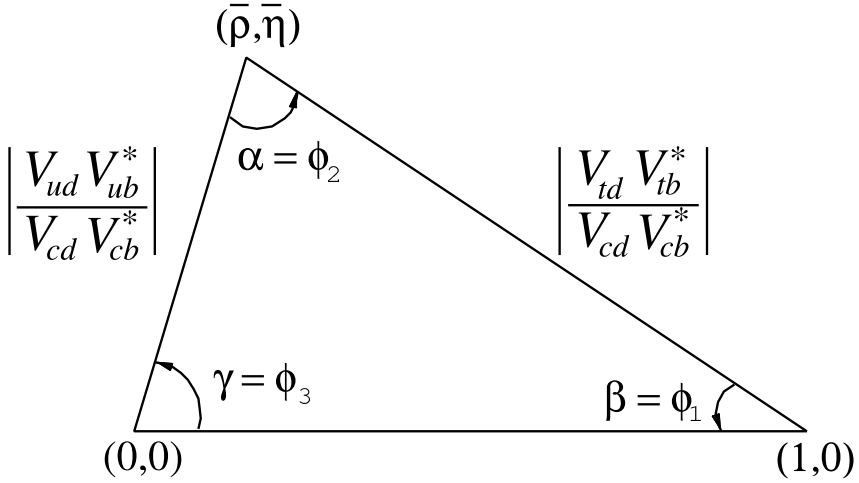

Figure 1.9: The unitary triangle eq. (1.28).

When all sides are normalized to $V_{cd}V_{cb}^*$, their lengths are given by:

$$\left|\frac{V_{ud}V_{ub}^*}{V_{cd}V_{cb}^*}\right| = \bar{\rho} + i\bar{\eta} \ , \tag{1.29}$$

$$\left|\frac{V_{td}V_{tb}^*}{V_{cd}V_{cb}^*}\right| = 1 - \bar{\rho} - i\bar{\eta} \ , \tag{1.30}$$

where $\bar{\rho}$ and $\bar{\eta}$ are the coordinates of one of the vertices of the triangle shown in Figure 1.9, and correspond to the parameters of the Wolfenstein parametrization [44]. The angles of the triangle are:

$$\alpha = \arg\left(-\frac{V_{td}V_{tb}^*}{V_{ud}V_{ub}^*}\right) \ , \quad \beta = \arg\left(-\frac{V_{cd}V_{cb}^*}{V_{td}V_{tb}^*}\right) \ , \quad \gamma = \arg\left(-\frac{V_{ud}V_{ub}^*}{V_{cd}V_{cb}^*}\right) \ . \tag{1.31}$$

The parameters of the CKM matrix are not determined by the theory but are measured by experiments. The current status is [17]:

$$\begin{pmatrix} 0.974\,17 \pm 0.000\,21 & 0.2248 \pm 0.0006 & (4.09 \pm 0.39) \times 10^{-3} \\ 0.220 \pm 0.005 & 0.995 \pm 0.016 & (40.5 \pm 1.5) \times 10^{-3} \\ (8.2 \pm 0.6) \times 10^{-3} & (40.0 \pm 2.7) \times 10^{-3} & 1.009 \pm 0.031 \end{pmatrix} \tag{1.32}$$

In particular the lengths of the sides and the angles of the CKM triangle are measured to test the unitarity of the matrix. The most recent fit to the world constraints is shown in

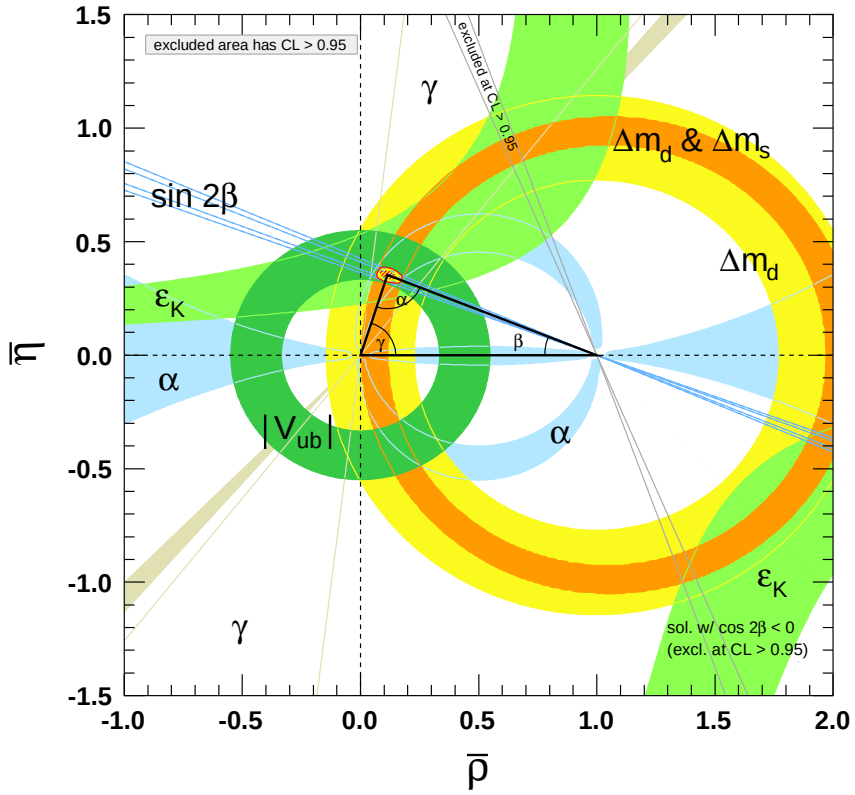

Figure 1.10: CKM triangle constraints in the $(\bar{\eta}, \bar{\rho})$ plane, taken from [17].

Figure 1.10.

The presence of the charm quark in the second quark generation implies the existence of additional terms in the neutral current. This so called GIM mechanism [45], cancels the flavour changing part ($\Delta S = 1$) of the neutral current by the addition of these terms. Because this only allows diagonal transitions in the flavour matrix, the flavour changing neutral current (FCNC) processes are forbidden at tree level and can only occur through loops. They are therefore suppressed and ideal probes of the standard model (SM), where small deviations can be a clear sign of new physics.

Amongst these interactions, the $K \to \pi\nu\bar{\nu}$ decays are considered as golden channels. They are strongly suppressed and theoretically very clean, allowing small effects foreseen by the theory or due to new physics to have the same size as the dominant processes and so be more easily distinguishable. The underlying $s \to d\nu\bar{\nu}$ transition can occur through the box (W boson exchange) and electroweak penguin diagrams (Z boson exchange) shown in Figure 1.11. In the SM the u, c, t quarks contribute to the loops in all diagrams, while the

charged leptons (e, μ, τ) only contribute to the box diagram. New physics particles can also enter these loops, leading to substantial modifications of the SM predictions. The SM branching ratio for the $K^{\pm} \to \pi^{\pm}\nu\bar{\nu}$ decay is calculated theoretically to an excellent precision and includes higher order QCD and electroweak corrections to the top and charm quark contributions. The most recent update gives [9]:

$$\mathcal{B}(K_L \to \pi^0\nu\bar{\nu}) \ = (3.00 \pm 0.30) \times 10^{-11} \ , \tag{1.33}$$

$$\mathcal{B}(K^{\pm} \to \pi^{\pm}\nu\bar{\nu}) = (9.11 \pm 0.72) \times 10^{-11} \ , \tag{1.34}$$

where the total uncertainties are $\lesssim 10\,\%$. The expression of the branching ratio for the charged channel reads:

$$\mathcal{B}(K^{\pm} \to \pi^{\pm}\nu\bar{\nu}) = \kappa_+ \left(1 + \Delta_{\mathrm{EM}}\right) \left[\left(\frac{\mathrm{Im}\lambda_t}{\lambda^5} X(x_t) \right)^2 \right. \tag{1.35}$$

$$\left. + \left(\frac{\mathrm{Re}\lambda_c}{\lambda} P_c(X) + \frac{\mathrm{Re}\lambda_t}{\lambda^5} X(x_t) \right)^2 \right] \ , \tag{1.36}$$

where $x_t = m_t^2/m_W^2, \lambda = |V_{us}|$ and $\lambda_i = V_{is}^* V_{id}$ are the CKM matrix parameters, $X(x_t)$ and $P_c(X)$ are the loop functions for the top and charm quark contributions. The term Δ_{EM} describes the electromagnetic radiative corrections. The hadronic matrix element κ_+ is the theoretically least known input, but it can be extracted from measurement of other well-known semi-leptonic K decay channels. The remaining uncertainties are only due to the CKM matrix parameters $|V_{ub}|, |V_{cb}|$ and γ. The decay rate is dominated by the short-distance contribution from the top loop and a smaller contribution from the charm quark loop. The long distance terms from the charm and up quarks only contribute as a small correction.

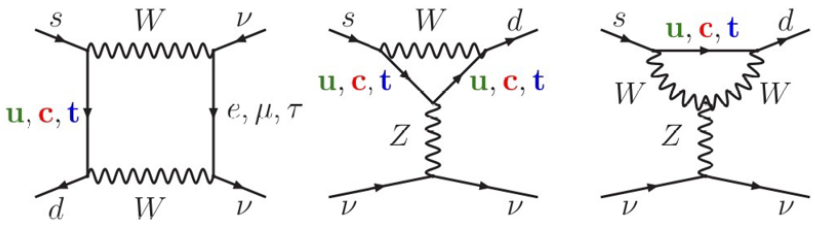

Figure 1.11: Feynman diagrams contributing to the $K^{\pm} \to \pi^{\pm}\nu\bar{\nu}$ decay. The first one is a box diagram involving a W exchange and the last two are electroweak Z penguins.

1.2.1 Experimental status

The existing measurements of these branching ratios are in agreement with the SM expectations:

$$\mathcal{B}(K_L \to \pi^0\nu\bar{\nu}) \; < 2.6\times10^{-8} \quad (90\,\%\ \text{CL})\,, \qquad [46] \tag{1.37}$$

$$\mathcal{B}(K^+ \to \pi^+\nu\bar{\nu}) = \left(17.3^{+11.5}_{-10.5}\right) \times 10^{-11}\,. \qquad [10] \tag{1.38}$$

The limit in the neutral channel was obtained by the E391a experiment at the KEK 12-GeV proton synchrotron. The KOTO experiment [47] at J-PARC aims at reaching the SM sensitivity and providing the first measurement of $\mathcal{B}(K_L \to \pi^0\nu\bar{\nu})$ by 2020. The experiment is served by a high-intensity K_L beam. The neutral kaon beam is generated by impinging a 30 GeV proton beam from the J-PARC Main Ring on a target. It is especially designed to suppress the halo of neutral particles. The signature of the decay consists of two photons coming from the decay of the π^0, with a large transverse momentum balancing the momentum carried away by the neutrinos. The main detector is a *CsI* electromagnetic calorimeter for the detection of the photons. All the remaining subsystems are veto detectors designed to reject charged particles, punch-through (photons not stopped in the calorimeter) and low energy photons coming from background.

The charged channel was investigated by the E787/E949 experiments at the Brookhaven National Laboratory. A 710 MeV/c K^+ beam was slowed down and stopped in a scintillating fibre target. The incoming K^+ was detected by a Cherenkov counter. The charged pion energy, momentum and range were detected, providing a powerful π^+ identification

to reject the background from other charged particles. A Range Stack, consisting of 21 layers of plastic scintillators used to stop the charged particles while measuring their trajectories, was also used to detect the $\pi^+ \to \mu^+ \to e^+$ decay chain. The outermost layer was an electromagnetic calorimeter used to veto any electromagnetic activity and in particular the photons from a π^0 decay, which constitute the main background. Seven candidate $K^+ \to \pi^+\nu\bar{\nu}$ events were identified (shown in Figure 1.12) and used in the branching fraction measurement. The $\sim 60\,\%$ uncertainty does not allow a conclusion on any deviation from the SM predictions.

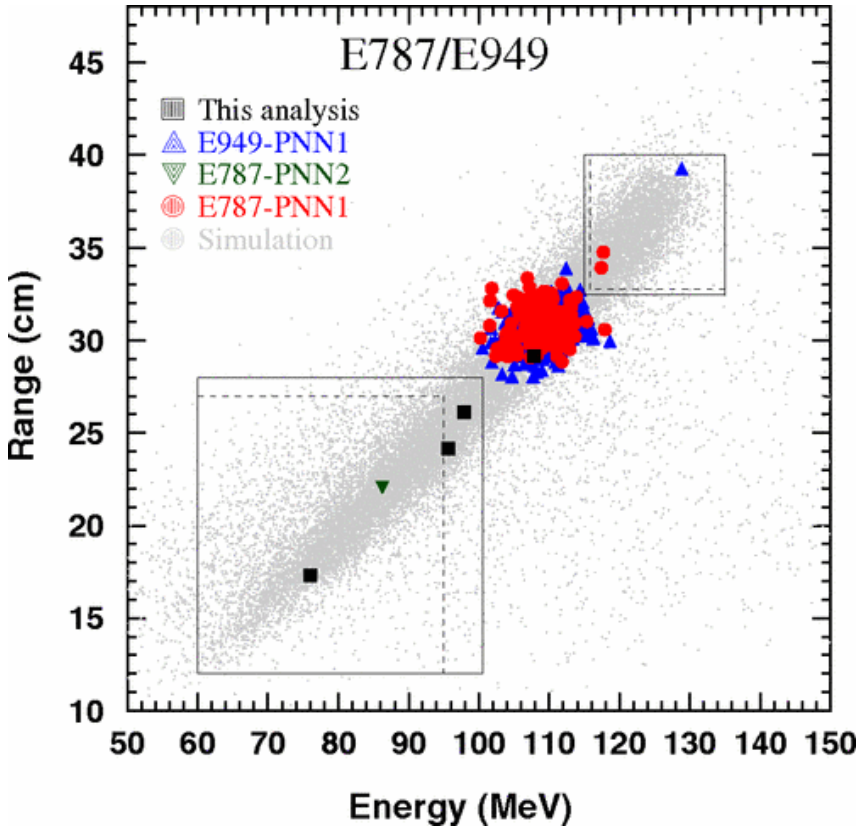

Figure 1.12: Candidate $K^+ \to \pi^+\nu\bar{\nu}$ events detected by the E787 and E949 experiments [48] in their final result. The x-axis is the charged pion kinetic energy, and the y-axis is its penetration depth in the range stack. The dashed squares correspond to the signal region of the E787 experiment while the solid ones correspond to E949. The different markers correspond to events selected in four different analyses (two for E787 and two for E949, below and above the $K^\pm \to \pi^\pm\pi^0$ peak due to two body decay without missing energy). All seven selected events are used for the determination of the branching ratio.

CHAPTER 2

THE NA62-R_K EXPERIMENT AT CERN SPS

The detector used in 2007–08 during the R_K phase of the NA62 experiment was the same as for the NA48/2 experiment [49], depicted in Figure 2.2. On the other hand, the original beam line [50] designed to provide simultaneous K^+/K^- beams was operated at lower intensity, and the trigger logic was changed.

2.1 Beam line

The primary SPS proton beam of 400 GeV/c impinged on a 40 cm long and 2 mm diameter beryllium target. The secondary beam leaving the target at zero angle with respect to the target symmetry axis was directed towards the experiment by the beam line shown in Figure 2.1. A system of four achromat dipole magnets and a momentum-defining slit incorporated into a beam dump selected the final momentum of (74.0±1.4) GeV/c. This 3.2 m thick copper/iron block offered the possibility to block either of the K^+ or K^- beams. A series of quadrupoles and collimators focused the beam horizontally and vertically towards the detector in a charge symmetric way. The beam entered the fiducial decay volume with a composition of approximately 6 % kaons, 62 % pions, 21 % protons, 10 % electrons and 1 % muons. Since the muon halo sweeping system was optimised for the positive beam in 2007, most of the data were recorded with the single K^+ beam in

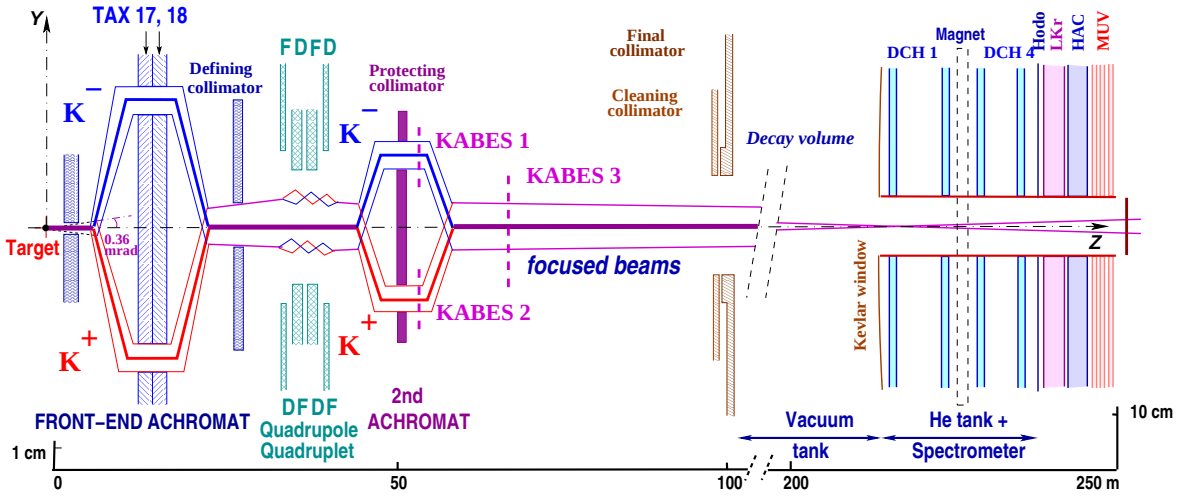

Figure 2.1: The NA48/2 beam line and detector [50].

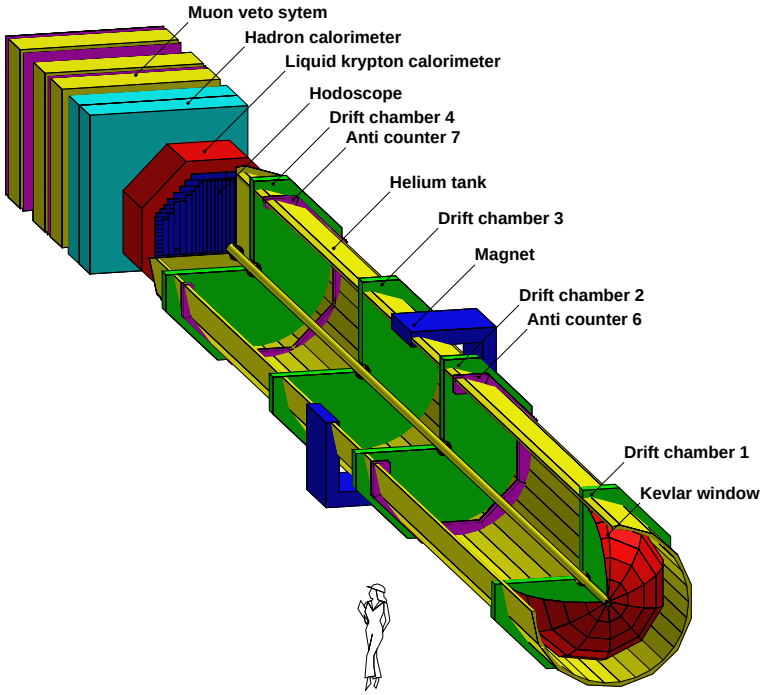

Figure 2.2: Schematic view of the NA48/2 detector [51]. The vacuum decay volume is upstream of the detector. The first element depicted is the helium tank housing the drift chambers (DCHs) and the magnet of the spectrometer, followed by a hodoscope (HOD), the liquid krypton calorimeter (LKr), hadron and muon systems.

order to reduce the halo background. To compensate the spectrometer magnet deflection of ∓ 3.58 mrad, a steering magnet deflected the K^+ and K^- beams at the entrance of the decay fiducial volume at angles of ± 0.23 mrad (± 0.30 mrad in the early stage of data taking) with respect to the detector axis. The polarities of those magnetic fields were regularly reversed simultaneously to reduce the effects caused by an asymmetry in the detector acceptance. With 10^{12} protons on target from the primary beam, the secondary one provided of the order of 2×10^6 particles per spill of 4.8 s length.

The vacuum volume was housed in a 114 m long cylindrical tank whose diameter varied from 1.9 m at the entrance to 2.4 m at the downstream end. The volume was evacuated to a residual pressure of 10^{-4} mbar. The downstream end of the tank was closed off by a thin *Kevlar* window of 0.3 % radiation length (X_0) thickness. It was followed by the spectrometer volume, housing the drift chambers. A central hole was present in all the elements to allow the passage of an aluminium beam pipe of 158 mm outer diameter and 1.1 mm thickness. The leak tightness was ensured by O-rings, decoupling the pipe from the external structure and allowing longitudinal movements.

2.2 Detector

Many detectors were part of the NA62-R_K setup and only the subset that are used for the measurement presented are described. The subsystems were spread longitudinally along the z axis. The main ones were the spectrometer for the measurement of the charged particle momenta, and the Liquid Krypton electromagnetic calorimeter for the measurement of the charged particle and photon energies. In addition a scintillator hodoscope was used in the trigger. They are presented below in the upstream to downstream order.

2.2.1 Magnetic spectrometer

The momenta of charged particles were measured by a magnetic spectrometer composed of four drift chambers (DCH1, 2, 3, 4), housed in a tank filled with helium gas at nearly

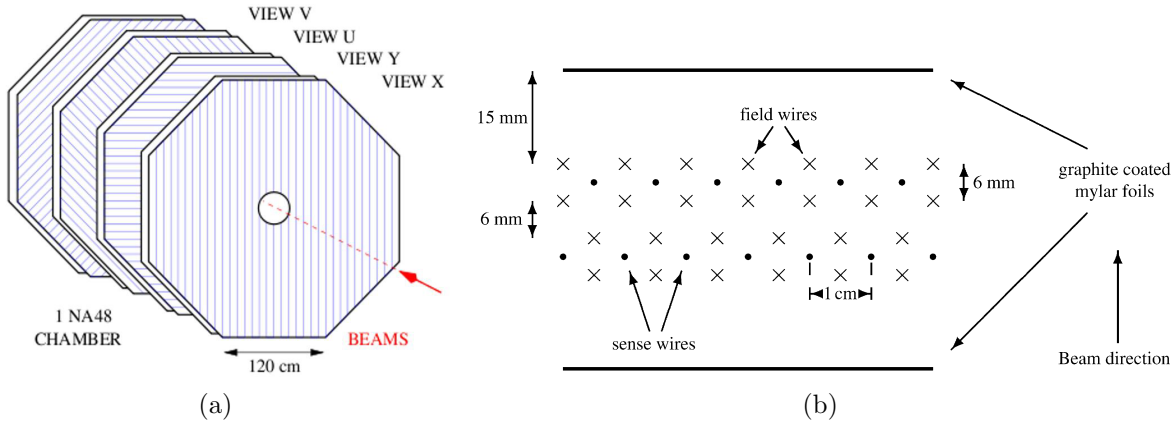

Figure 2.3: (a) Schematic layout of a drift chamber (DCH) made of four views (X,Y,U,V). (b) Geometry of a single drift cell [49].

atmospheric pressure. A dipole magnet placed between the second and third chambers provided a horizontal transverse momentum kick of 265 MeV/c to singly-charged particles. The magnet polarity was regularly inverted to compensate for the different acceptances for the K^+/K^- beams. The chambers shown in Figure 2.3a are of octagonal shape with a radius of 2.9 m and a central hole of 160 mm diameter to accommodate the beam pipe. Each chamber was made of eight planes of sense wires orthogonal to the z axis and oriented at 0° (X), 90° (Y), −45° (U) and 45° (V). Each view of the chamber was composed of two planes staggered by 5 mm (half distance between two wires) as shown in Figure 2.3b, to resolve ambiguities at reconstruction. The spectrometer was operated in a mixture of Argon-Methane (50 %–50 %), allowing a drift time of maximum 100 ns. The resolution on individual hits was better than 100 µm, achieving a momentum resolution of $\sigma_p/p = (0.48 \oplus 0.009 \cdot p)\%$ where the momentum p is in GeV/c. The first term is due to multiple scattering in the helium and *Kevlar* window, and the second is due to the resolution on the hit position. The time resolution on the vertex decay time, using only DCH hits, was about 700 ps.

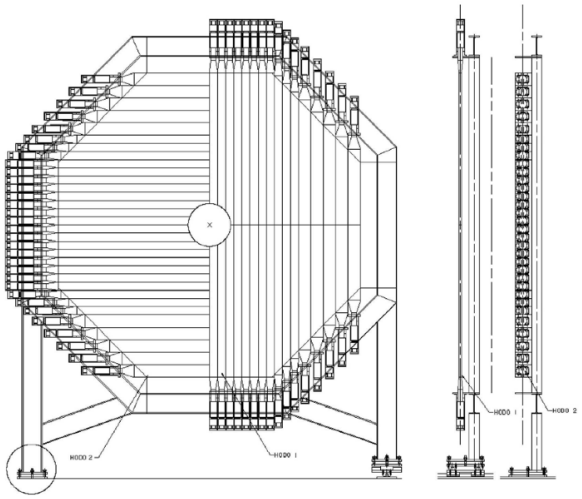

Figure 2.4: Layout of the charged hodoscope (HOD) [49].

2.2.2 Scintillator hodoscope

The charged hodoscope (HOD) is made of two planes of plastic scintillator strips separated by 74 cm and was located just after the spectrometer. It is shown schematically in Figure 2.4. The first plane consists of 64 vertically arranged strips and the second of 64 horizontally arranged strips. The scintillator counters in each plane are grouped in bands of 8×2, whose intersections define 16 mutually exclusive quadrants. The size of the 2 cm thick strips vary from a width of 6.5 cm close to the beam to 9.9 cm further out, and a length of 121 cm to 60 cm. The radius of the central hole is 12.8 cm. The aim of this detector is to provide a fast low-level trigger signal and precise timing for charged particles. The HOD triggers are defined by setting conditions on combinations of hits in the quadrants. The measured time resolution is better than 200 ps.

2.2.3 Liquid Krypton electromagnetic calorimeter

The main detector for measuring the particle energies was the Liquid Krypton electromagnetic calorimeter (LKr), following directly after the HOD. It is an octagonal quasi-homogeneous calorimeter, filled with $\sim 9\,\mathrm{m}^3$ of liquid krypton kept at a constant temperature of 121 K. The liquid krypton has a radiation length of $X_0 = 4.7$ cm and a Molière radius of $R_M = 4.7$ cm. Its properties allow the detector to be compact while

still containing $> 99\,\%$ of electromagnetic showers up to 50 GeV. The active volume has a cross-section of $\sim 5.3\,\mathrm{m}^2$ and a depth of 127 cm (77 X_0). It is divided into 13 248 cells of $\sim 2{\times}2\,\mathrm{cm}^2$ cross-section with no longitudinal segmentation, consisting of copper/beryllium ribbons (Figure 2.5a). About 100 of them were not operating correctly due to faulty electronics and are qualified as "dead". The performances in terms of energy, position and time resolution were:

$$\frac{\sigma_E}{E} = \frac{0.032}{\sqrt{E}} \oplus \frac{0.09}{E} \oplus 0.0042 \;, \tag{2.1}$$

$$\sigma_{X,Y} = \frac{0.42}{\sqrt{E}} \oplus 0.06 \;, \tag{2.2}$$

$$\sigma_t = \frac{2.5}{\sqrt{E}} \;, \tag{2.3}$$

where the units are in GeV, cm and ns. The calorimeter is also used for particle identification, in particular for the trigger in the analysis presented.

The neutral hodoscope (NHOD) consisted of a plane of 256 bundles of scintillating fibres inside the LKr. They were placed vertically at the maximum depth of the shower development for 10 GeV photons ($9X_0$). This detector was mainly used to provide a fast trigger signal, as well as an independent time measurement for the showers. The time resolution for showers above 15 GeV was $\sigma_t = 250\,\mathrm{ps}$.

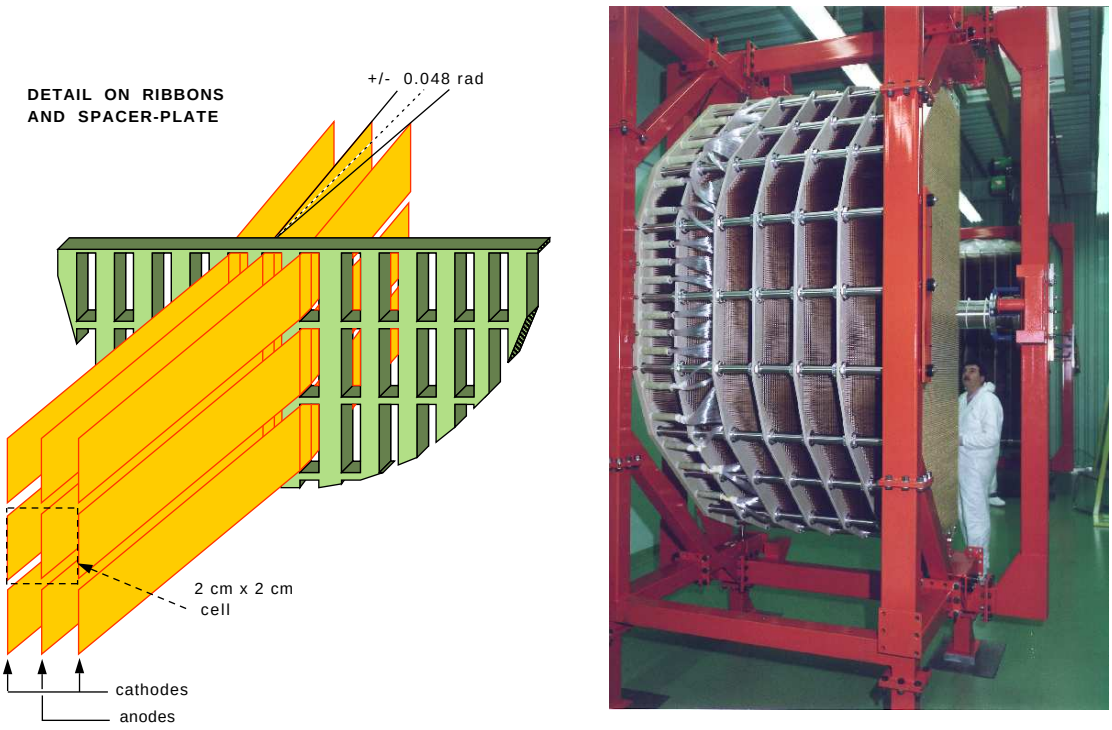

(a) (b)

Figure 2.5: (a) Details of a calorimeter cell structure [49]. (b) Picture of the LKr structure at CERN [52].

CHAPTER 3

MEASUREMENT OF THE π^0 ELECTROMAGNETIC TRANSITION FORM FACTOR

The analysis focuses on a pure sample of $K^{\pm} \to \pi^{\pm}\pi^0; \pi^0 \to e^+e^-\gamma$ ($K_{2\pi D}$) events. A small accidental contamination from other sources of Dalitz decays, $K^{\pm} \to \pi^0_D \mu^{\pm}\nu$ ($K_{\mu 3D}$) and $K^{\pm} \to \pi^0_D e^{\pm}\nu$ (K_{e3D}), is also present and is taken into account in the MC samples. However they do not bias the measurement as the form factor in the π^0_D decay does not distinguish between the different origins of the neutral pion. After the selection criteria have been applied, a fit based on a χ^2 of the data/MC x distribution is performed to extract the TFF slope value. Finally the systematic uncertainties are estimated.

3.1 Data sample and trigger logic

The full data set collected during 4 months of data taking in 2007 is used in this analysis. It was obtained from about 3.5×10^5 SPS spills (with $\sim 10^{12}$ protons on target per spill), corresponding to $\sim 2\times10^{10}$ $K^{\pm}$ decays in the vacuum tank. As a high intensity source of charged kaon decays, the NA62-R_K setup is ideal to study the decays of the π^0. The π^0 is indeed produced in four of the main decay channels of the $K^{\pm}$, whose branching fractions are indicated in Table 3.1. Out of these, $K_{2\pi}$ with its 20% branching ratio is the most obvious one to study. It promises large statistics, while the absence of neutrinos in the

Table 3.1: Main charged kaon decay channels ordered by their branching fractions [17]. The channels marked with a $^{(*)}$ are sources of π^0_D.

Decay		Branching Fraction (%)
$K^\pm \to \mu^\pm\nu$	$(K_{\mu 2})$	63.56±0.11
$K^\pm \to \pi^\pm\pi^0$ $^{(*)}$	$(K_{2\pi})$	20.67±0.08
$K^\pm \to \pi^\pm\pi^+\pi^-$	$(K_{3\pi})$	5.583±0.024
$K^\pm \to \pi^0 e^\pm\nu$ $^{(*)}$	(K_{e3})	5.07±0.04
$K^\pm \to \pi^0\mu^\pm\nu$ $^{(*)}$	$(K_{\mu 3})$	3.352±0.033
$K^\pm \to \pi^\pm\pi^0\pi^0$ $^{(*)}$		1.760±0.023

final state allows strong kinematic constraints to be applied.

Some nomenclature is used to refer to subsamples of data. The burst, corresponding to the data acquired during a single SPS spill, is the smallest grouping of events. They are further grouped into runs, each containing of the order of 1000 bursts and corresponding to a continuous time interval during which the conditions are stable. The runs are organized into 6 periods, between which the data taking conditions can change significantly. The main condition changes are described below and summarized for each period in Table 3.2.

As mentioned in the section 2.1, the beam was used in different modes (K^+/K^-/ Simultaneous) and deflecting angles. A total of 65 % (8 %) of the K^+ (K^-) flux was collected in single-beam mode, while the remaining 27 % was collected with a simultaneous $K^\pm$ beam with a ratio of K^+/K^- fluxes of 2.0.

During the first part of the data taking (55 % of the kaon flux), a lead (Pb) bar was installed between the two HOD planes, approximately 1.2 m in front of the LKr, below the beam pipe. The aim was to study the muon-induced electromagnetic showers in the calorimeter [53] to quantify the misidentification probability of muons as electrons. A total of 11 rows of LKr cells were shadowed by the bar, corresponding to about 10 % of the total number of cells.

The rate of 100 kHz of kaon decays in the vacuum volume during the spill enabled the use of a minimum-bias trigger configuration. The main trigger chain of the experiment was exploited for this analysis and was developed as a high efficiency trigger optimised to select events with at least one electron $e^\pm$ track. The full chain is $Q_1 \times E_{\mathrm{LKr}}(10\,\mathrm{GeV}) \times$

1TRKL(M) $\times$ $L3(K_{e2})$ and is referred to as the K_{e2} ($K^{\pm} \to e^{\pm}\nu$) trigger. The different elements composing this chain were the following:

- Q_1: A coincidence of hits in at least one quadrant in each of the HOD planes.
- $E_{\mathrm{LKr}}(10\,\mathrm{GeV})$: At least 10 GeV of energy deposition in the LKr.
- 1TRLK(M): 1-track condition, related to hits in the DCH. The condition was not applied in period 1. In period 2 to 4 the 1TRKL condition was used, requiring at least one hit in at least two views of the first chamber. In the following period the condition was updated to 1TRKLM, which additionally requires no more than 15 hits in any of the views of DCH1, DCH2 and DCH4.
- $L3(K_{e2})$: At least one track with $p \in (5; 90)\,\mathrm{GeV}/c$ and $E/p > 0.6$, where E is the energy of the associated cluster in LKr and the momentum p is that reconstructed in the spectrometer.

The first three conditions were implemented as a low level hardware trigger, while the last one was a software high level trigger (HLT).

Other downscaled control trigger streams (only 1 event out of N passing the conditions is accepted) were employed and are used to test trigger efficiencies: the $K_{\mu 2}$ ($K^{\pm} \to \mu^{\pm}\nu$) chain corresponding to $Q_1 \times$ 1TRLK(M), the K_{e2} autopass for which the decision of the $L3(K_{e2})$ condition was only flagged, and the NHOD stream triggered by electromagnetic showers in the NHOD detector.

3.2 Software package

The software package (including the simulation, reconstruction, and analysis) has been developed specifically for the NA48 experiment. It consists of a collection of programs intended to simulate the physical event and the detector response, to reconstruct the event and provide all the required information for analysis.

Period	Runs	Beam	Beam Deflection	DCH Trigger	Pb Bar
1	20114 to 20174	K^+,K^-	+0.30 mrad	✗	✓
	20175 to 20203		−0.30 mrad		
2	20209 to 20226	K^+,K^-	+0.30 mrad	1TRKL	✓
	20228 to 20256		−0.30 mrad		
	20268 to 20285		+0.30 mrad		
3	20286 to 20291	K^+,K^-	+0.23 mrad	1TRKL	✓
	20296 to 20303		−0.23 mrad		
	20304 to 20324	K^+	+0.23 mrad		
4	20332 to 20351	K^+	+0.23 mrad	1TRKL	✓
	20352 to 20371		−0.23 mrad		
	20385 to 20386	K^-	−0.23 mrad		
	20387 to 20404	K^+	+0.23 mrad		
5	20410 to 20437	K^+	+0.23 mrad	1TRKLM	✗
	20438 to 20458		−0.23 mrad		
	20459 to 20481		+0.23 mrad		
	20482 to 20485		−0.23 mrad		
6	20487 to 20500	K^-	−0.23 mrad	1TRKLM	✗
	20501 to 20531		+0.23 mrad		
	20525		+0.23 mrad		
	20530 to 20531		+0.23 mrad		

Table 3.2: Summary of conditions for the runs used in this analysis.

3.2.1 Beam simulation

The TURTLE software [54] is used to simulate the optics of the beam line and transport the beam particles to the entrance of the decay volume. The beam was monitored on a run-by-run basis using the $K^{\pm} \to \pi^{\pm}\pi^{+}\pi^{-}$ decay. The average momentum and nominal axis of the beam were extracted from this measurement and are simulated for each run separately. The width of the distribution is however not reproduced exactly, and this needs to be compensated for (see section 3.4).

3.2.2 Decay generator

The generation of the K decay itself is performed by a specific package for a comprehensive list of decay channels relevant to the experiment. In particular it includes the simulation of radiative photons in the $K^{\pm}$ decays [55]. The Dalitz decay of the secondary neutral pion (π^0) produced in the $K^{\pm}$ decay is handled by the specific generator described in

subsection 1.1.1.

3.2.3 Detector simulation

The CMC (Charged Monte Carlo) package implements a detailed description of the detector geometry and material. It is based on GEANT3 libraries [56], which simulate the particle interaction with matter through a realistic list of physical processes. The response of the detector is also simulated by including the inefficiencies of the different subsystems. Since the MC samples are generated on a run-by-run basis, the efficiencies and calibration parameters used can follow accurately the ones measured in the data. The most relevant parameters for this analysis are the DCH chamber misalignment and single wire efficiency. It is worth noting that the simulation of the EM showers in the LKr is sped up by using shower libraries. As a result the energy deposition for hadrons and muons and the shower overlap are not accurately reproduced. The Pb bar (see section 3.1) is also not present in the simulation and should be regarded in the analysis as a blind spot for the runs during which it is present. Any LKr activity behind it must be discarded in both data and MC.

3.2.4 Reconstruction and analysis

Both data and MC events are stored in the COmPACT (or superCOmPACT) data format [57], allowing the reconstruction software to process them indistinguishably. The framework also has a connection to the database providing the data taking conditions and calibration parameters for the event. This is required to compute and apply corrections to improve the data/MC agreement (some are also automatically applied within the framework). A description of those that are applied manually can be found in section 3.4.

After applying basic corrections and a simple filtering based on the recorded trigger word corresponding to the K_{e2} trigger, the event is exported into ROOT files. A multi-step (selection, fitting, plotting) ROOT [58] based code dedicated to this analysis is then used.

Track reconstruction in the spectrometer

The tracks are fully reconstructed using the information from the spectrometer. Contiguous hits from both planes of each view (see subsection 2.2.1) are grouped into clusters. Front segments are built from all possible combinations of clusters in the corresponding views of DCH1 and DCH2. Each of these front segments is a plane representing the track projection in one of the four coordinates (X,Y,U and V). The front tracks are defined by four front segments, one for each view, and their intersections with the chambers are called space points. The front tracks are projected onto DCH4 and are matched with clusters that are within 4 cm. The timing information is then added and the average hit time is used as the track time. The quality of the track is computed as the fraction of hits in time with the track time. As a last step the measured magnetic field map is used in a fitting algorithm to determine the parameters of the track. For each fitted track, the clusters on DCH3 compatible with the extrapolated position of the track are used to refine its momentum.

Energy cluster reconstruction in the liquid krypton calorimeter

The algorithm for the reconstruction of the clusters of energy in the LKr relies on seeds, which are cells with a measured energy $> 100\,\mathrm{MeV}$. Clusters are then identified as seeds that have more energy than the 8 surrounding cells, and with an energy $E_{\mathrm{seed}} > 0.18\,\mathrm{GeV} + 1.8E_{\mathrm{a}}$, where E_{a} is the average energy of the 8 surrounding cells. The energy of each cell with only one cluster within 11 cm and 20 ns (neighbour cluster) is associated to that cluster. If more than one neighbour cluster is found, the cell energy is shared between them according to the expected energy profile. The position of the cluster is defined as the centre of gravity of the energy deposition in a box of 3×3 cells around the seed. Its time and width are computed respectively as the average cell time and as the r.m.s. (in x and y directions separately) of the energy distribution in a box of 5×5 cells around the seed.

3.3 Simulated Monte Carlo Samples

Three main Monte Carlo (MC) samples have been produced for this analysis: 440M events for $K^{\pm} \to \pi^{\pm}\pi^0_D$ ($K_{2\pi D}$), which is the main signal; 120M events for $K^{\pm} \to \pi^0_D \mu^{\pm}\nu$ ($K_{\mu 3D}$) and 120M events for $K^{\pm} \to \pi^0_D e^{\pm}\nu$ (K_{e3D}), which are π^0_D channels with residual events passing the selection. Other smaller samples have also been generated to verify the contamination from the $K^{\pm}$ decays that are the most likely to enter the selection as a background ($300\,\mathrm{M}\ K_{2\pi}; \pi^0 \to \gamma\gamma,\ 90\,\mathrm{M}\ K_{3\pi}$). All π^0_D decays have been generated using the TFF slope value corresponding to the current world average [59]: $a_{\mathrm{sim}} = 0.032$. The π^0_D decay generator described in subsection 1.1.1 was used with the radiative photon cut-off value $x_\gamma^{\mathrm{cut}} = 0.01$. Two other $K_{2\pi D}$ samples were generated with $x_\gamma^{\mathrm{cut}} = 0.005$ and $x_\gamma^{\mathrm{cut}} = 0.05$ to cross-check the convergence of the result when $x_\gamma^{\mathrm{cut}} \to 0$.

The number of simulated events per run is proportional to the number of acquired events during the corresponding run. A small residual difference attributed to varying trigger conditions, MC job failures, and the application of the w_{pk} weights (see subsection 3.4.1) modifies the ratio between data and MC. Systematic effects can result from the imbalance of the data and MC size ratios between runs with different conditions. It is corrected by applying a run-dependent weight $w_r = N_{\mathrm{Data},r}/N_{\mathrm{MC},r}$, where $N_{\mathrm{Data},r}$ and $N_{\mathrm{MC},r}$ are respectively the numbers of selected Data and MC events for the run.

3.4 Corrections for known effects

A number of known detector/reconstruction/simulation effects have already been studied in the past that require corrections to be applied. Some are implemented already in the reconstruction code, while others have to be applied during analysis and are described in this section.

3.4.1 Beam momentum spectrum corrections

The beam momentum spectrum was monitored run by run throughout the data taking with samples of $K^{\pm} \to \pi^{\pm}\pi^{+}\pi^{-}$ decays. The mean value of the momentum was stable around the nominal 74 GeV/c, varying only between (73.6;74.1) GeV/c. The run dependence of this value was reproduced in the MC simulation. However the width of the distribution was kept at a constant fraction of the central value (1.6 GeV/c at $P_K = 74\,\mathrm{GeV}/c$) in MC, while the measured width varied between (1.5;1.8) GeV/c. In order to reduce the possible acceptance effects due to the Data/MC disagreement in the tails of the reconstructed P_K spectrum, as seen in Figure 3.1a, weights are introduced on the MC events to compensate for the discrepancy. The quadratic weights have the following form:

$$w_{pk} = 1 + \lambda_r (p_K - 74\,\mathrm{GeV}/c)^2 \ , \tag{3.1}$$

where p_K is the true kaon momentum and λ_r is a run dependent parameter obtained from a fit of the $K^{\pm} \to \pi^{\pm}\pi^{+}\pi^{-}$ P_K spectrum. Figure 3.2 shows the run-dependent value of λ_r and Figure 3.1b shows the corrected spectrum.

3.4.2 Track momentum corrections

Two residual effects must be corrected for the charged tracks measured by the spectrometer: geometrical misalignment of the chambers, and mis-calibration of the magnetic field integral. The corrected momentum is obtained by applying:

$$p = p_0\,(1+\beta)\,(1+\alpha p_0 q) \ , \tag{3.2}$$

where p_0 is the reconstructed momentum, q is the charge of the track and α, β are the correction coefficients. The α correction corresponds to the chamber misalignment, depends on the charge of the track and is inversely proportional to the track momentum. The coefficient β is a global scale correcting for the mis-calibration of the spectrome-

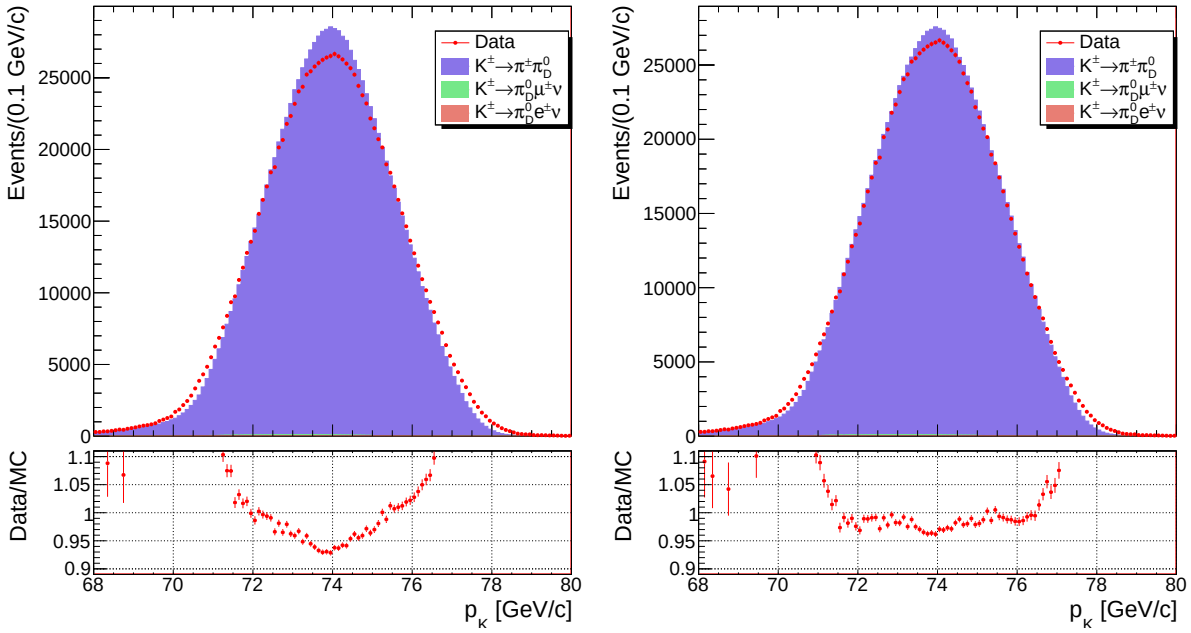

(a) (b)

Figure 3.1: Beam momentum using unweighted MC sample (a); using the w_{pk} weighted MC sample (b).

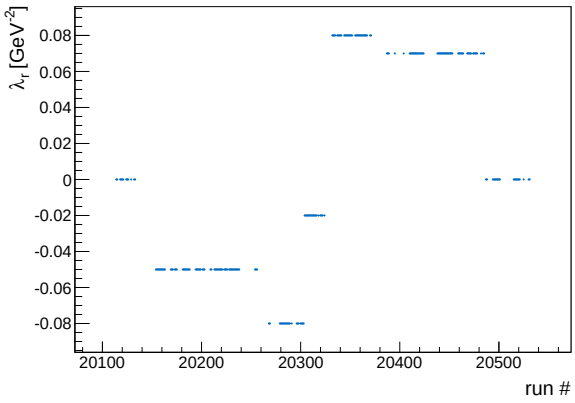

Figure 3.2: Value of the λ_r parameter of the beam momentum spectrum correction w_{pk} as a function of the run number.

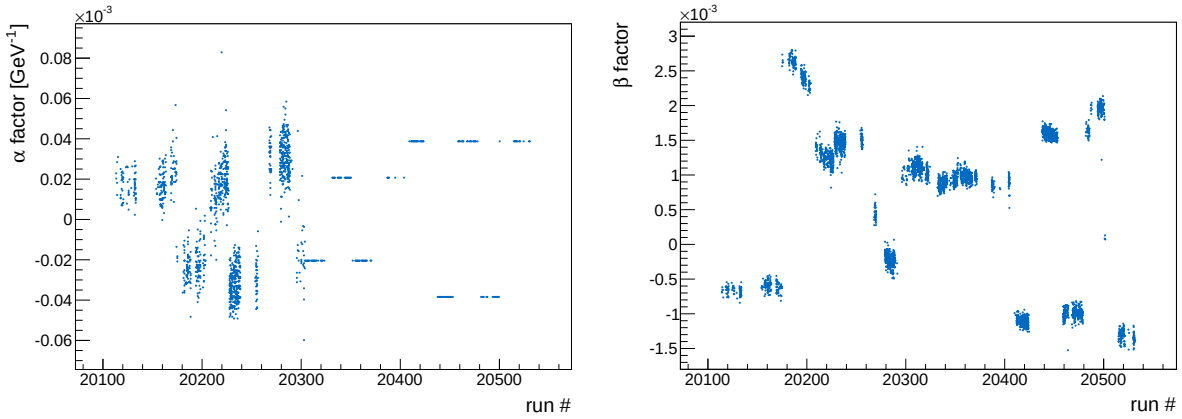

(a) (b)

Figure 3.3: Correction factors α and β applied for spectrometer track corrections as a function of the run number. The fractional interval between two run numbers is divided into as many subintervals as the number of different values the factors take during the run. The β factor, of the order of 10^{-3}, is the dominant contribution to the correction, while the α factor is of the order of $10^{-5}\,\mathrm{GeV}^{-1}$.

ter. The values for these corrections are obtained every ~ 500 bursts from a selection of $K^{\pm} \to \pi^{\pm}\pi^{+}\pi^{-}$ decays. A fit is performed to find the momentum correction parameters, for which the reconstructed 3-pion invariant mass is closest to the nominal $K^{\pm}$ mass. The values of α and β are shown in Figure 3.3. For the first half of the run both K^{-} and K^{+} beams were present, and they are affected differently by the misalignment because of their opposite charges. This allows the individual effects of the α and β factors to be disentangled. On the other hand, α is fixed to an optimal value for the later (single beam) runs and only β is fitted.

In addition, a stray magnetic field of $\sim 50\,\mu\mathrm{T}$ was present in the decay region due to the Earth's magnetic field and this was mapped with a Hall probe. The track momentum direction at the vertex is corrected by propagating the track backwards from the DCH position to the vertex z coordinate, taking into account the measured Earth's field.

3.4.3 LKr non-linearity energy correction

For each event, only the LKr cells above a certain threshold were read out (zero-suppression mechanism). As a result a non-linearity in the energy response of the calorimeter affected clusters of energy deposit below 11 GeV. For such clusters the energy must be modified as:

$$E = E_0 f(E_0) \ , \tag{3.3}$$

where E_0 is the raw cluster energy and $f(E_0)$ is a third order polynomial. The typical correction was of the order of 0.1 %, about 10 times smaller than the energy resolution. This mechanism is not implemented in the MC, for which this correction should not be applied.

3.4.4 LKr cell-by-cell energy re-calibration

The energy response of each individual LKr cell has been studied from a sample of $K^{\pm} \to \pi^0 e^{\pm} \nu$ (K_{e3}) decays [53]. In order to improve the global resolution and uniformity of the calorimeter, a run-dependent energy scaling is applied to the cluster energy depending on the cell hit. This correction, applied only for data, again has a typical size of 0.1 %.

3.4.5 Cluster position correction

The LKr was built in such a way that the axes of the ionization cells focus on a projectivity point located on the z axis at a distance $D = 109.98$ m in front of the calorimeter front plane. The trajectories of particles originating nearby this point are close to parallel to these axes and the measured transverse positions (x_0, y_0) do not depend on the depth at which the shower develops. When originating further from this point, a correction depending on the shower depth must be applied to the transverse position of the cluster. The shower depth in cm for electron and photon clusters is given by $d_{\rm cl} = 16.5\,{\rm cm} + 4.3\,{\rm cm} \cdot \ln(E/1\,{\rm GeV})$, where the reconstructed energy E is given in GeV. The resulting

projectivity correction to the coordinates is:

$$(x_p, y_p) = (x_0, y_0) \cdot \left(1 + \frac{d_{\rm cl}}{D}\right) \ , \tag{3.4}$$

$$z_p = z_{\rm LKr} + d_{\rm cl} \ , \tag{3.5}$$

where x_0, y_0 are the reconstructed coordinates of the cluster and $z_{\rm LKr}$ is the z position of the LKr front plane.

An additional transformation of the coordinates is required to correct a residual misalignment between the LKr and the spectrometer coordinate systems. The transformation contains a constant shift and a relative rotation angle, which are measured using a sample of electrons. The position of the track measured in the DCHs and extrapolated onto the LKr is compared to the position of the associated cluster of energy deposition. The complete corrections for data, including the projectivity, the shift and the rotation angle of $\phi \approx 0.87{\times}10^{-3}$ radians is:

$$x = \left(x_0 + 0.136\,{\rm cm} + 0.87{\times}10^{-3} y_0\right) \cdot \left(1 + \frac{d_{\rm cl}}{D}\right) \ , \tag{3.6}$$

$$y = \left(y_0 + 0.300\,{\rm cm} - 0.87{\times}10^{-3} x_0\right) \cdot \left(1 + \frac{d_{\rm cl}}{D}\right) \ , \tag{3.7}$$

$$z = z_{\rm LKr} + d_{\rm cl} \ . \tag{3.8}$$

The geometry implemented in the Monte Carlo is not identical and does not include the rotation. The corrections applied to MC are:

$$x = (x_0 - 0.013\,{\rm cm}) \cdot \left(1 + \frac{d_{\rm cl}}{D}\right) \ , \tag{3.9}$$

$$y = y_0 \cdot \left(1 + \frac{d_{\rm cl}}{D}\right) \ , \tag{3.10}$$

$$z = z_{\rm LKr} + d_{\rm cl} \ . \tag{3.11}$$

The size of the modification of the transverse cluster positions is typically about 0.5 cm.

3.5 $K^{\pm} \to \pi^{\pm}\pi^0_D$ event selection

The final state that must be selected is composed of three charged tracks coming from the same vertex, two of them having the same charge and one of them having a charge opposite to the incoming kaon, and one cluster of energy deposition in the LKr calorimeter not associated with any charged track. Assuming the cluster of energy is produced by a photon coming from the same vertex as the tracks, the kinematics of the event must be compatible with the signal chain. To ensure that the data and simulated events receive similar treatment and can be compared reliably, the cuts applied in the trigger are replicated.

The following conditions are applied on the events which are not extracted from a burst tagged as *bad*, and which are collected by the K_{e2} trigger chain described in section 3.1.

3.5.1 Vertex selection

The vertex is reconstructed using a Kalman filter algorithm [60] applied on the track segments reconstructed in the upstream part of the spectrometer. The algorithm propagates the tracks backwards into the decay region, taking into account the multiple scattering in the helium and the *Kevlar* window of the spectrometer and the stray magnetic fields due to the Earth's and parasitic vacuum tank fields. The parameters describing the tracks are iteratively adjusted until the convergence of the fit. The resolution in the longitudinal position of the vertex is of the order of $\sim 1\,\mathrm{m}$. The events are required to have exactly one 3-track vertex located in the fiducial decay region $z_{\mathrm{vtx}} \in (-17; 90)\,\mathrm{m}$. The lower limit is the entrance of the decay volume, which corresponds to the last collimator of the beam located at $z = -18\,\mathrm{m}$, within the vertex resolution. The upper limit discards events decaying too close to the spectrometer, which would not be reconstructed properly, but all events satisfying the selection condition in fact have $z_{\mathrm{vtx}} < 80\,\mathrm{m}$, as seen in Figure 3.4. The vertex must be of good quality with $\chi^2 < 25$, where the χ^2 is the sum of the reduced χ^2 of the three tracks forming the vertex. The electric charge of the vertex, defined as

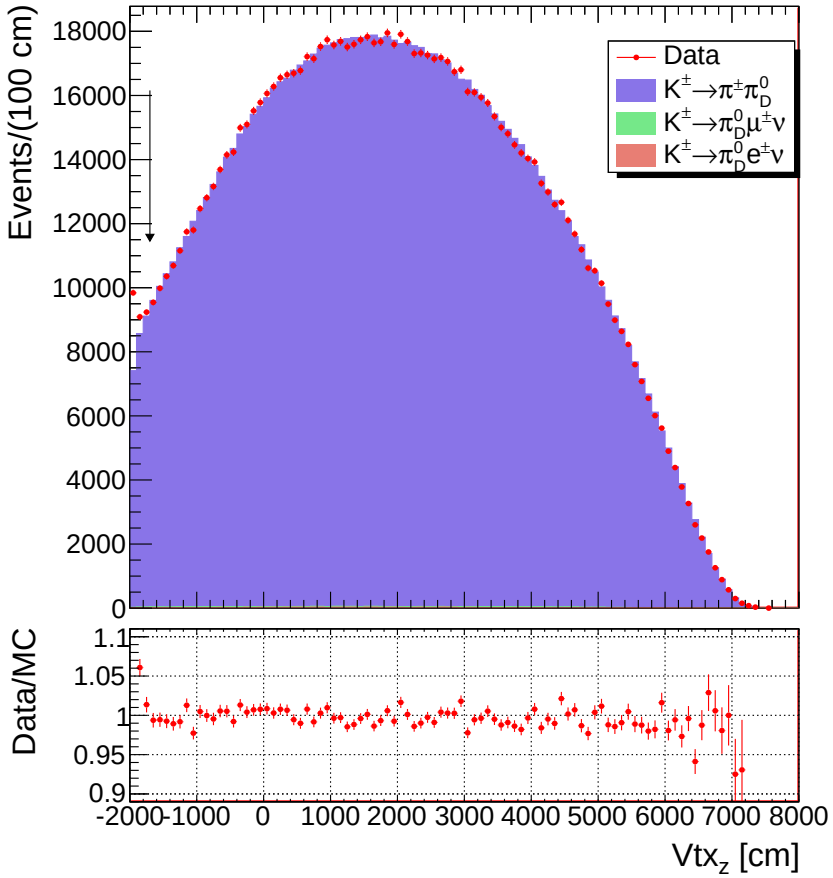

Figure 3.4: z coordinate of the vertex for data, MC and the data/MC ratio. The arrow indicates the standard value for the cut applied in the $K_{2\pi D}$ selection. The upper cut is outside of the z range visible in this plot (9000 cm); no event reaching this part of the selection is seen above 8000 cm.

the sum of charges of all tracks composing the vertex, must be $|q_{\text{vtx}}| = 1$. In runs for which the charge of the incoming kaon was known (i.e. single beam runs), the sign of q_{vtx} must agree with the beam charge. The vertex time t_{vtx}, used in the following conditions, is defined as the average time of the tracks forming the vertex.

3.5.2 Charged track selection

The event is searched for "extra" tracks that satisfy the following conditions:

- The track does not belong to the vertex;

- The track is within 20 ns of t_{vtx} (Data only);
- The reconstructed momentum is below 74 GeV/c;
- The closest distance of approach (CDA) between the track and the z-axis is smaller than 10 cm and in the fiducial decay region $z_{\text{CDA}} \in (-20; 90)$ m.

These tracks are typically particles coming from an accidental kaon decay such as $K_{\mu 2}$. Events containing such tracks are discarded as they can affect the identification of the event through additional energy clusters or hits shared in the track reconstruction.

The remaining tracks belonging to the vertex must satisfy the following properties:

- The fraction of in-time hits with respect to the average DCH hit time > 0.7.
- The Data/MC ratio on the distribution of track momenta shows that the simulation of tracks below 2 GeV/c does not describe the data correctly (Figure 3.5). The reconstructed momenta should therefore be in the range $p \in (2; 74)$ GeV/c, with the upper limit corresponding to the nominal beam momentum.
- All tracks are extrapolated to the drift chamber z planes, and the impact points (x, y) on DCH1, DCH2 and DCH4 should be within the acceptance: 12 cm $< \sqrt{x^2 + y^2} <$ 110 cm.
- The distance between any pair of tracks in the DCH1 plane should be at least 2 cm in order to suppress photons converting into e^+e^- pairs in the material in, or in front of, the spectrometer.
- All track times are required to be within 25 ns of the trigger time, and they should not be separated by more than 15 ns from each other. The track time is given by $t_i = t_{\text{R},i} - t_{\text{DCH}}$, where $t_{\text{R},i}$ is the reconstructed track time given by the DCH and t_{DCH} is the offset of the spectrometer with respect to the HOD trigger time. All time related cuts are applied on data only as the simulation does not provide the information on the timing.

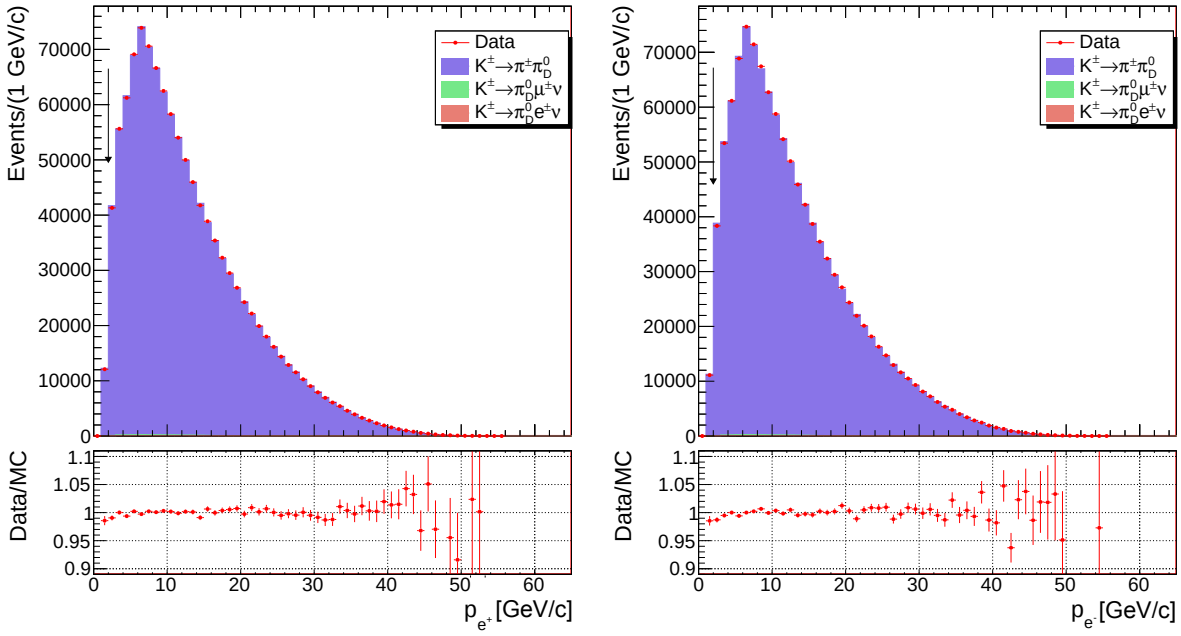

(a) (b)

Figure 3.5: Electron (a) and positron (b) track momentum for data, MC and data/MC ratio. A 2 % deficit of data is observed below 2 GeV/c. The arrows indicate the standard value for the cut applied in the $K_{2\pi D}$ selection.

3.5.3 Photon selection

Only one photon cluster in the LKr calorimeter is allowed in the event. This is considered to be the only cluster in the LKr satisfying the following conditions:

- Not being geometrically associated to a track. The track impact points on the LKr front plane are extrapolated taking into account the bending by the spectrometer magnet. The distance between the cluster and each of the two same-sign tracks should exceed 20 cm, and the distance to the third, opposite-sign track should be greater than 10 cm. This difference of treatment is due to the larger transverse size of the pion cluster, which necessarily has the same sign as the vertex. These cuts are not applied if the track impact point is located inside a radius of 10 cm around the centre of the calorimeter, corresponding to the beam hole. This results in a relative gain of acceptance of $\sim 0.01\,\%$.

- Not being geometrically associated to an undeflected track. The impact points on the LKr are extrapolated without considering the spectrometer magnetic field and the distance from the cluster should be above 20 cm. This condition discards clusters due to internal or external bremsstrahlung occurring before the magnet.

- Several selection criteria are applied due to the known differences between data and MC. The cluster energy should be above 2 GeV as the LKr non-linearity and the simulation of the bremsstrahlung photons are in agreement with the data only above this limit (see subsection 3.7.2). The Pb bar (see section 3.1), present during periods 1–4, was not included in the simulation and any cluster present behind it, with $y \in (-33.575; -11.850)$ cm, is ignored.

- The cluster in data events should be in-time with the event, within 10 ns of the vertex time.

Once a suitable cluster of energy deposit is identified, the photon candidate is reconstructed assuming its origin is located at the 3-track vertex position. The following

criteria are further applied:

- The cluster position in the LKr front plane must lie within the acceptance of the LKr, at least 8 cm away from areas affected by a missing or malfunctioning digitiser, and more than 2 cm away from any dead cell.
- The energy is required to be above 3 GeV.
- The radius of the extrapolated photon position on DCH1 must be larger than 11 cm to exclude events where the photon traverses the DCH inner flanges.

3.5.4 Kinematic particle identification

In order to pursue the analysis, the three tracks need to be identified. The usual particle identification (PID) method in the NA62-R_K experiment relies on the ratio E/p. The track momentum p is measured in the spectrometer. The energy E is measured from the cluster of energy deposit in the LKr calorimeter, which is geometrically associated to the track by matching its extrapolated position on the LKr front plane with the position of the cluster. Electrons produce an electromagnetic shower, which is mostly absorbed in the LKr, leading to $E/p \approx 1$. Pions are minimum ionising particles (MIPs) that release little energy in the calorimeter, unless initiating a hadronic shower through the strong interaction. The E/p of $\pi^\pm$ peaks at low values of the ratio but covers the whole available range, leading to a non-negligible probability of mis-identifying the $\pi^\pm$ as $e^\pm$ (about 2 %) and vice-versa (about 8 %). This identification technique also requires all charged particles to be within the acceptance of the LKr, reducing the overall acceptance for the events.

An alternative kinematic particle identification technique has been developed and relies on the specific properties of the $K_{2\pi D}$ decay chain. The final state contains $\pi^\pm e^+ e^- \gamma$, and only one of the tracks has a charge opposite to the beam, which is necessarily the $e^\mp$ candidate, with an energy of $\sqrt{|\vec{p}|^2 + m_e^2}$. Two hypotheses are possible for the $\pi^\pm$ and $e^\pm$ assignments of the two same-sign tracks. They are tested by reconstructing the kinematics of the event under both assumptions. An hypothesis is considered valid if the

reconstructed variables are in agreement with the expected values for a $K_{2\pi D}$ event:

$$M_{\pi^0}^{\text{low}} < M_{ee\gamma} < M_{\pi^0}^{\text{high}} , \tag{3.12}$$

$$M_K^{\text{low}} < M_{\pi ee\gamma} < M_K^{\text{high}} , \tag{3.13}$$

$$x, |y| < 1 , \tag{3.14}$$

where the invariant mass $M_{ee\gamma} = \sqrt{(p_{e^+} + p_{e^-} + p_\gamma)^2}$ corresponds to the reconstructed π^0 mass, the invariant mass $M_{\pi ee\gamma} = \sqrt{(p_{\pi^\pm} + p_{e^+} + p_{e^-} + p_\gamma)^2}$ corresponds to the reconstructed $K^\pm$ mass, and x and y are the Dalitz kinematic variables. Events with a single valid hypothesis are selected.

This kinematic identification has been tested using the $K_{2\pi D}$ and the $K_{\mu 3D}$ MC samples by associating the MC particles to the reconstructed tracks using their momentum direction. The MC particle and the track are positively associated if the angle between the true and reconstructed direction is $\theta < 2{\times}10^{-4}$ rad. The event is considered fully identified when both e^+ and e^- MC particles are uniquely assigned to tracks, the $\pi^\pm$ being assigned to the remaining track. About 91 % of the MC $K_{2\pi D}$ events reaching the PID method are fully identified. The missing 9 % are mainly due to too small an opening angle between the e^+ and e^- momenta.

The distibutions of the reconstructed π^0 and $K^\pm$ masses for both assignment hypotheses are shown in Figure 3.6. Despite a small overlap, the bulks of the distributions for the correct and wrong mass assignments are well separated. The oblique and vertical bands in Figure 3.6a are due to resolution effects on the $e^\pm$, the photon, and the $\pi^\pm$. The horizontal band is due to events where an emitted hard radiative photon is selected instead of the photon emitted from the Dalitz decay.

The mass cuts were varied in the ranges $M_K^{\text{low}} \in (0.395; 0.485)\,\text{GeV}/c^2$, $M_K^{\text{high}} \in (0.505; 0.595)\,\text{GeV}/c^2$ and $M_{\pi^0}^{\text{low}} \in (0.035; 0.13)\,\text{GeV}/c^2$, $M_{\pi^0}^{\text{high}} \in (0.14; 0.235)\,\text{GeV}/c^2$. Multiple combinations of the $K^\pm$ and the π^0 mass cuts were tested on the MC sample. Their effect is assessed by looking at the fraction of events correctly identified, incorrectly

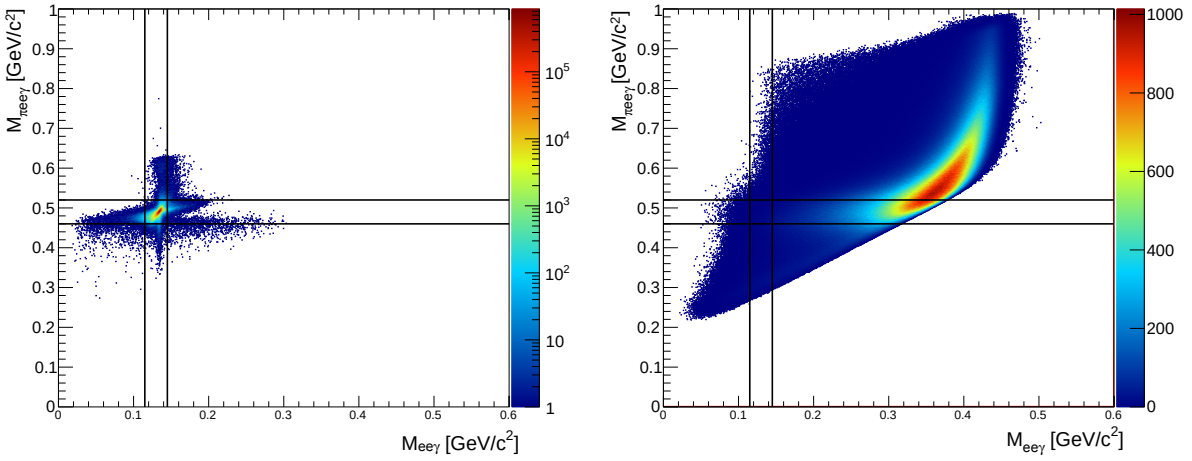

(a) (b)

Figure 3.6: Phase spaces for the reconstructed π^0 and $K^{\pm}$ masses in the PID process. (a) for the correct particle mass assignment, and (b) for the wrong one. The vertical and horizontal black lines show the applied cuts on M_{π^0} and M_K respectively. The identified events are those located inside the box in (a) but outside in (b); all other events are rejected.

identified, or rejected because none or both of the hypotheses were valid. The fraction of identified events varied between 93.95 % and 99.74 %, the events with no good hypothesis between 0.05 % and 6.02 %, and the events with both valid hypotheses between 0.02 % and 0.36 %. Out of the identified events, the probability of mis-identification (i.e. wrong mass assignment) varied between 0.01 % and 0.02 %.

The final selected cuts are a trade-off between identification efficiency and misidentification probability:

$$0.115\,\mathrm{GeV}/c^2 < M_{ee\gamma} < 0.145\,\mathrm{GeV}/c^2\,, \tag{3.15}$$

$$0.460\,\mathrm{GeV}/c^2 < M_{\pi ee\gamma} < 0.520\,\mathrm{GeV}/c^2\,, \tag{3.16}$$

represented in Figure 3.6 by the solid black lines. For these values, the fraction with correct identification is 99.6 % and that of mis-identified events is 0.02 %. None of the hypotheses is valid for 0.3 % of the events and both in 0.1 % of the cases. A total of 3.83×10^6 data events reach the PID stage, out of which 91.5 % are identified with a single

valid hypothesis, 8.4 % are rejected with no valid hypothesis, and 0.1 % are rejected due to both hypotheses being valid. The number of events with no valid hypothesis is consistent with the number of $K_{3\pi}$ decays reaching this stage of the selection, as estimated from a MC $K_{3\pi}$ sample. The acceptance for $K_{2\pi D}$ events is increased of a factor of more than two with respect to the E/p particle identification, and the systematic uncertainties due to particle misidentification and the bias on the x distribution shape are decreased. This is due in small part to the detection efficiency of the LKr calorimeter ($\sim 1\,\%$) and for the bigger part to the release of the requirement for the three tracks to be in the LKr acceptance.

Since all particles are detected in this decay chain, the event is requested to be compatible with no missing momentum (including transverse momentum with respect to the beam axis), within the resolution:

$$70\,\mathrm{GeV}/c < \left|\vec{P}_{\mathrm{tot}}\right| < 78\,\mathrm{GeV}/c\,, \tag{3.17}$$

$$\left|\vec{P}_t\right|^2 < 5\times10^{-4}\,(\mathrm{GeV}/c)^2\,, \tag{3.18}$$

where $\vec{P}_{\mathrm{tot}} = \vec{p}_\gamma + \sum_{i=1}^{3}\vec{p}_i$ is the total momentum corresponding to the reconstructed $K^\pm$ and $\vec{P}_t$ is its transverse component with respect to the time-dependent beam axis. The kaon beam momentum direction was independently measured from $K_{3\pi}$ decay samples every $\mathcal{O}(500)$ bursts, separately for K^+ and K^-. The variations, due to changes in the beam configuration and magnet polarities, are of the order of 10^{-4} radians.

3.5.5 Trigger conditions

The last step, in order to obtain comparable data and MC samples, is to apply the trigger conditions described in section 3.1 on both samples. To eliminate edge effects caused by different calibrations and resolutions between online and offline analyses, tighter variants of the criteria are applied to both data and MC samples.

The energy deposition in the LKr is dominated by the $e^\pm$ and the photon, and the

offline trigger condition corresponding to $E_{\mathrm{LKr}}(10\,\mathrm{GeV})$ conservatively requires that they deposit together more than 14 GeV of energy in the LKr. The difference between the online and offline energy cuts takes into account the resolution which is $\sigma_E \approx 1\,\mathrm{GeV}$ ($\sigma_E = 0.14\,\mathrm{GeV}$) online (offline) for 10 GeV clusters. The LKr acceptance variable for the electron and positron $A_{\mathrm{LKr}}(e^{\pm})$ is defined as 1 if the track impact point on LKr is within the geometrical acceptance, not behind the Pb bar, more than 8 cm away from any missing or malfunctioning digitiser, and further than 2 cm from a dead cell; else it is equal to 0. This definition allows the offline condition to be written in the following concise form:

$$E_{\mathrm{LKr}}^{\mathrm{offline}} : \left(E_{\gamma} + A_{\mathrm{LKr}}(e^{+}) \cdot E_{e^{+}} + A_{\mathrm{LKr}}(e^{-}) \cdot E_{e^{-}}\right) > 14\,\mathrm{GeV}\ . \tag{3.19}$$

The photon acceptance $A_{\mathrm{LKr}}(\gamma)$ is not explicitly required as the conditions are already fulfilled by the definition of the photon candidate (see subsection 3.5.3).

The equivalent of the $L3(K_{e2})$ trigger condition for the electron or positron can be written:

$$L3^{\mathrm{offline}}(e^{\pm}) : A_{\mathrm{LKr}}(e^{\pm}) \cdot (p_{e^{\pm}} > 5.5\,\mathrm{GeV}/c) \cdot \left(\frac{E_{e^{\pm}}}{p_e^{\pm}} > 0.8\right)\ . \tag{3.20}$$

The different values with respect to the online trigger (see section 3.1) take into account the difference of resolution on the cluster energy and track momenta ($\sigma_p \approx 0.02\,\mathrm{GeV}/c$ for both online and offline). The event is selected if either $L3^{\mathrm{offline}}(e^{+})$ or $L3^{\mathrm{offline}}(e^{-})$ is satisfied.

3.5.6 Signal region and DCH acceptance cut

The $x = (M_{ee}/m_{\pi^0})^2$ variable is highly correlated to the opening angle between the $e^{\pm}$ tracks. The requirement of a minimum distance of 2 cm between the track impact points in DCH1 steeply reduces the acceptance for small opening angles, and therefore for events at low x. Figure 3.7 shows that the $K_{2\pi D}$ acceptance is not well reproduced in the simulation for $x \leq 0.01$, with a data deficit of about 1 %. Since events in this region are not sensitive

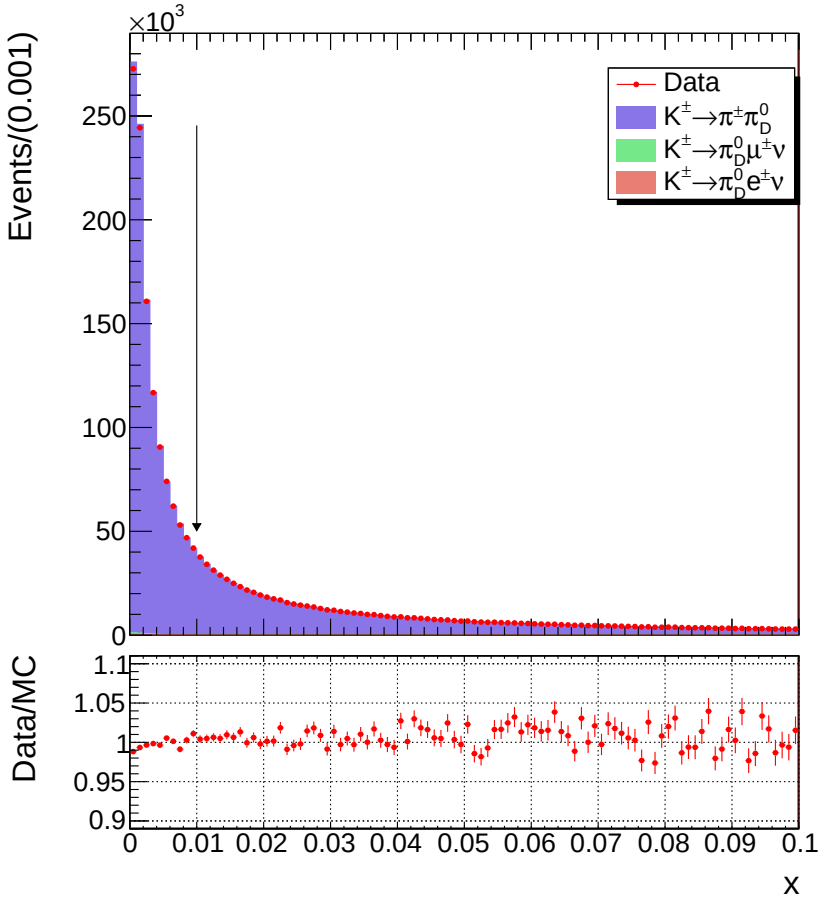

Figure 3.7: Zoom on the $x < 0.1$ region of the x variable for data, MC and their ratio. The deficit of data events can be seen for $x < 0.01$. The arrow indicates the lower x cut value in the $K_{2\pi D}$ selection.

to the form factor, and to avoid any acceptance related effect, the signal region is defined as $x > 0.01$, equivalent to $M_{ee} > 13.5\,\mathrm{MeV}/c^2$.

A disagreement between data and simulated events can be seen for events with small opening angles between the $e^{\pm}$ tracks in the central region of the DCH. This is particularly visible in the distributions of the distance $d_{ee,DCH1}$ between the $e^{\pm}$ impact points in the first drift chamber (Figure 3.8a), and their radius with respect to the centre of the chamber $R_{e^+,DCH1}, R_{e^-,DCH1}$(Figure 3.8b). The deviation of the Data/MC ratio from a constant is seen for radii below 20 cm, and for inter-track distances smaller than 10 cm. Furthermore these variables are highly correlated to the $e^{\pm}$ track opening angles and the x variable, leading to instabilities of the fit result.

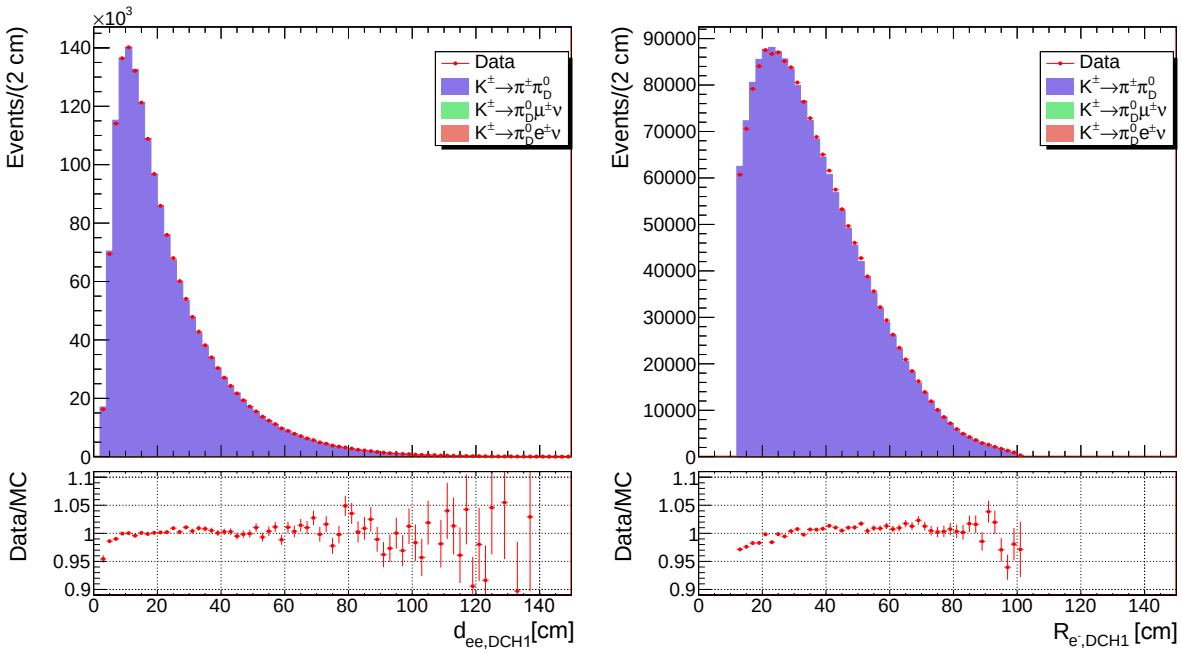

(a) (b)

Figure 3.8: Distributions of (a) $d_{ee,DCH1}$ and (b) $R_{e^-,DCH1}$ defined in subsection 3.5.6. The $x > 0.01$ cut is applied, but not the DCH acceptance cut. The deficit of data events in the lowest part of the distributions ($d_{ee,DCH1} < 10\,\text{cm}$ and $R_{e^-,DCH1} < 20\,\text{cm}$) can be seen.

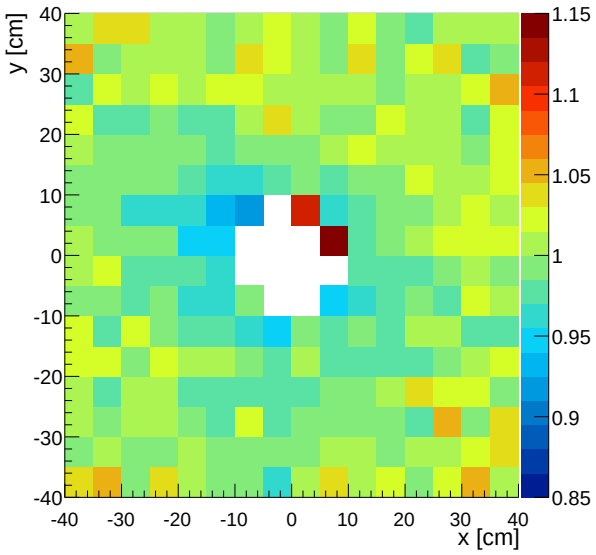

Figure 3.9: Data/MC ratio of the e^+ impact point in the central region of DCH1. The region $-20\,\text{cm} < x < 20\,\text{cm}$; $-20\,\text{cm} < y < 20\,\text{cm}$ shows a deficit of data events. The two bins with a high ratio (red) are an artefact due to the low statistics in these bins.

The hypothesis of an inefficient area in one of the DCH is tested by defining a lattice on the first drift chamber. Sub-samples of events are created based on the cell impacted by the e^+ and e^- tracks and the $d_{ee,DCH1}$ distribution is plotted for each sub-sample. The observations show that the inefficient area is located in the central region around the beam hole. It can be seen in Figure 3.9, showing the data/MC ratio of the e^+ impact point in the central region of DCH1. Figure 3.10 shows the difference between an x slice in the beam hole region (left) with a deficit of data events up to 5 %, and an x slice in the outer region (right) which is compatible with a flat distribution. The extended size of the area where the effect can be seen and its absence in the $\pi^\pm$ distributions dismisses the hypothesis of an inefficiency due to one or a group of wires.

The resulting additional cut is introduced:

$$\left|x_{\text{DCH1}}(e^\pm)\right| > 20\,\text{cm}\ , \tag{3.21}$$

$$\left|y_{\text{DCH1}}(e^\pm)\right| > 20\,\text{cm}\ . \tag{3.22}$$

It reduces the acceptance by 30 % but the data/MC agreement is restored, as can be seen

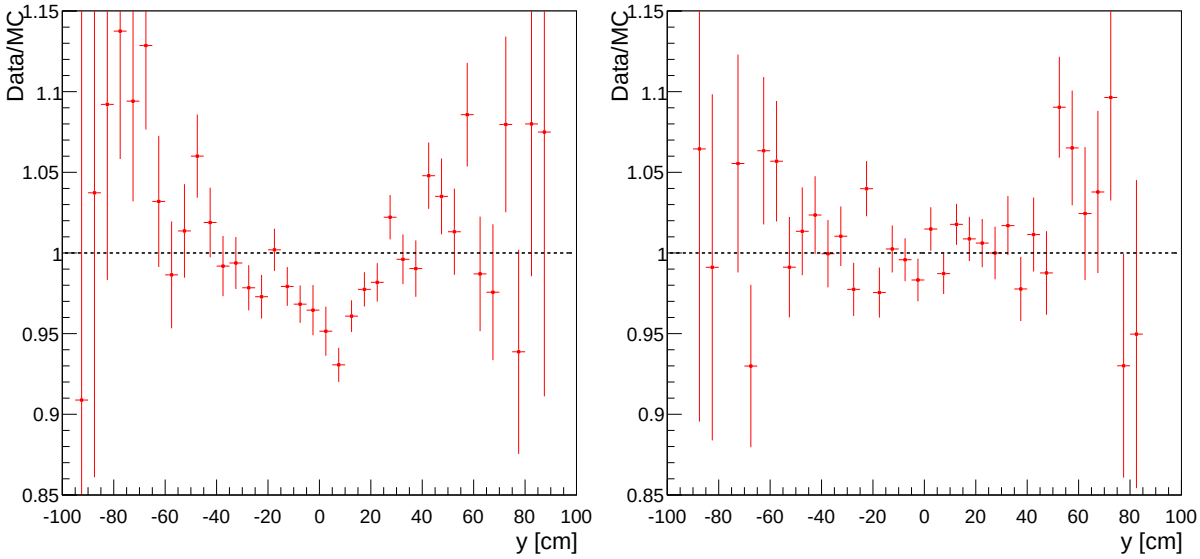

(a) (b)

Figure 3.10: Slices in x of Figure 3.9. (a) corresponds to the slice $-15\,\text{cm} < x < -10\,\text{cm}$, close to the beam pipe and shows a clear deficit of data events. (b) corresponds to the slice $20\,\text{cm} < x < 26\,\text{cm}$, further away from the beam pipe, where the data and MC distributions agree better.

in Figure 3.11, and the fit result is no longer sensitive to small variations of the cuts.

The final selected data sample amounts to 1.1×10^6 π^0_D events. The overall acceptances of the above $K_{2\pi D}$ selection, evaluated with MC simulations, are 1.90 % for the $K_{2\pi D}$ decays, 0.02 % for the $K_{\mu 3D}$ decays and 0.01 % for the K_{e3D} decays. The $K_{2\pi D}$ acceptances for periods with and without the Pb bar installed are 2.23 % and 1.64 %, respectively. The reconstructed distribution of the x variable and the acceptances for the simulated samples are shown in Figure 3.12. The e^+e^- mass resolution, determined from the $K_{2\pi D}$ MC sample, can be approximated by $\sigma_{ee} = 0.9\,\% \cdot M_{ee}$, which translates into the resolution on the x variable as $\sigma_x = 1.8\,\% \cdot x$ (see subsection 3.7.3). The reconstructed $e^+e^-\gamma$ and $\pi^\pm\pi^0$ invariant mass spectra are shown in Figure 3.13; the mass resolutions are $1.5\,\text{MeV}/c^2$ and $3.7\,\text{MeV}/c^2$, respectively.

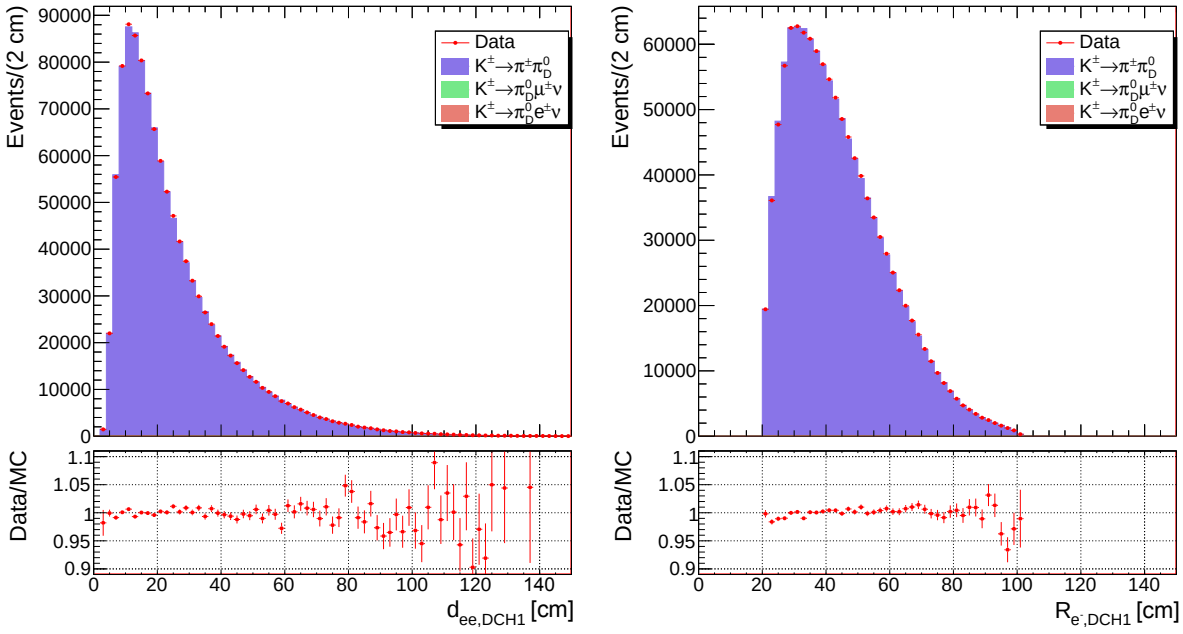

(a) (b)

Figure 3.11: Distributions of (a) $d_{ee,DCH1}$ and (b) $R_{e^-,DCH1}$ defined in subsection 3.5.6. The DCH acceptance cut is applied and the data/MC agreement is good.

3.6 Fit procedure

The TFF enters the decay width as a factor of $|\mathcal{F}(x)|^2 = |1+ax|^2 = 1+2ax+a^2x^2$. The principle of the measurement is to triplicate each MC event into three samples, with weights corresponding respectively to the constant, linear, and quadratic terms of the TFF. They can then be combined together to give the total contribution, assuming a certain value of the TFF slope. These three samples all contain the three simulated decay chains ($K_{2\pi D}, K_{\mu 3D}, K_{e3D}$) taking into account the respective branching fractions. The reconstructed MC x distribution obtained is compared to the data and the fit result corresponds to the value of a with the best data/MC agreement.

The comparison is performed using a binned χ^2 estimator, which does not work for lowly populated or empty bins. To avoid this issue at large x, where only 1 % of the events have $x > 0.5$ (see Figure 3.12), an equipopulous binning is chosen with the bin limits determined from the data sample. Various numbers of bins have been tested, as

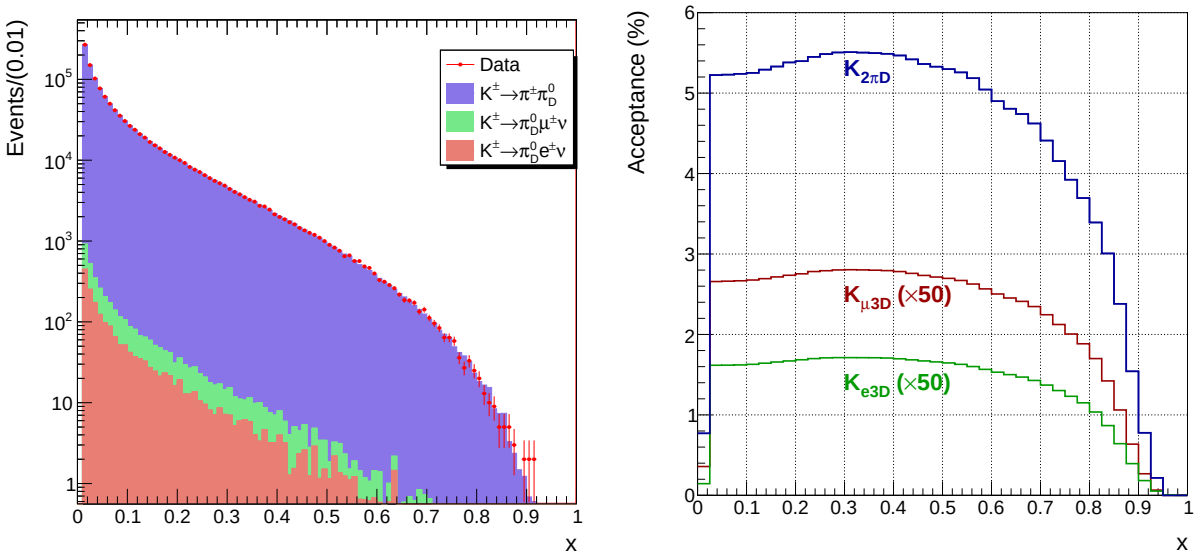

(a) (b)

Figure 3.12: (a) Spectra of the x variable for the three $K^{\pm}$ decay modes considered ($K_{2\pi D}, K_{\mu 3D}, K_{e3D}$), after the $K_{2\pi D}$ selection. (b) Acceptances of the $K_{2\pi D}$ selection for the three decay channels as a function of x. The acceptance is maximal for the main channel, while the two other π^0_D channels are suppressed (scaled by 50 for visibility). The acceptances for all other kaon decay channels are suppressed to at least 10^{-7} (see subsection 3.7.7). The drop in the first bin is due to the signal region definition ($x > 0.01$).

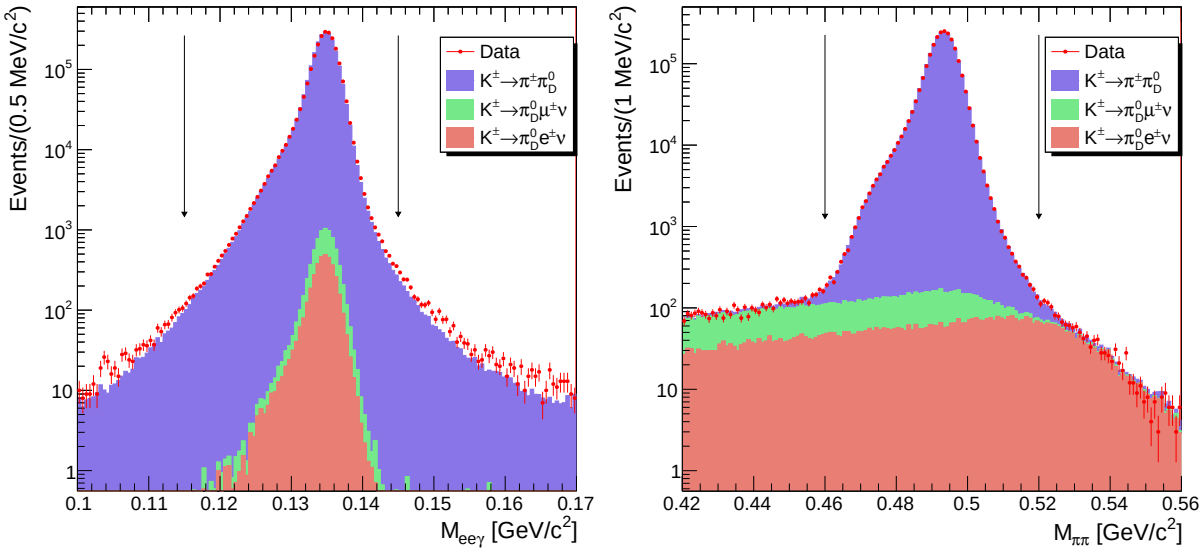

(a) (b)

Figure 3.13: Reconstructed (a) $e^+e^-\gamma$ and (b) $\pi^\pm\pi^0$ invariant mass distributions for data and simulated components. The radiative shoulders in the reconstructed masses are well reproduced in the MC thanks to the simulation of the radiative photon. The arrows correspond to the default cut values for the kinematic PID procedure applied in the $K_{2\pi D}$ selection.

reported in subsection 3.7.9. A default of 50 bins with widths ranging from $\Delta x = 0.006$ to $\Delta x = 0.59$ was chosen, giving a stable statistical uncertainty. The χ^2 estimator [61, 62], evaluating the hypothesis of homogeneity of the data and MC (i.e. the probability that both data and MC histograms describe the same distribution), is a modified version of the standard one [63], which takes into account the normalisation and weights of the MC histograms. Its expression is:

$$\chi^2(a) = \sum_{i=1}^{r} \frac{(n_i - N\hat{p}_i)^2}{N\hat{p}_i} + \sum_{i=1}^{r} \frac{(w_i - W\hat{p}_i)^2}{s_i^2} \,, \tag{3.23}$$

where r is the number of bins, n_i is the number of data events in the ith bin and N is the total number of events in the data sample. The first term is identical to the standard χ^2 for unweighted histograms, where the number of events in a bin follows a Poisson distribution. The second term in eq. (3.23) corresponds to the χ^2 estimator for a weighted histogram, where the sum of weights in the ith bin, w_i, is expected to follow a gaussian distribution

of mean Wp_i and variance σ_i^2. The sum of all the MC event weights is $W = \sum_{i=1}^{r} w_i$ and p_i is the probability for an event to be found in the ith bin ($\sum_{i=1}^{r} p_i = 1$). The estimator of the variance is $s_i^2 = \sum_{i=1}^{r} w_i^2$, the sum of squares of weights of events in the ith bin, which properly takes into account the MC statistical uncertainties. The maximum likelihood estimator of p_i under the hypothesis of homogeneity is

$$\hat{p}_i = \frac{Ww_i - Ns_i^2 + \sqrt{(Ww_i - Ns_i^2)^2 + 4W^2 s_i^2 n_i}}{2W^2} \,. \tag{3.24}$$

The standard χ^2 test of two unweighted histograms is equivalent when one of them is considered as a weighted histogram with all weights equal to 1.

The weighting technique used allows many values of a to be tested without generating specific MC samples. The constant value $a_{\text{sim}} = 0.032$ corresponding to the PDG world average [59] is used at the generator level. Each MC event should receive the weight:

$$w(a) = \frac{(1 + a x_{\text{MC}})^2}{(1 + a_{\text{sim}} x_{\text{MC}})^2} \,, \tag{3.25}$$

where x_{MC} is the true MC generated value of the x variable. The histograms must be recreated for every tested value of a. However, building the three aforementioned constant (h_c), linear (h_l) and quadratic (h_q) histograms is much more efficient computation-wise. The events are filled in the histograms with the following weights, respectively:

$$w_c = \frac{1}{(1 + a_{\text{sim}} x_{\text{MC}})^2}, \quad w_l = \frac{x_{\text{MC}}}{(1 + a_{\text{sim}} x_{\text{MC}})^2}, \quad w_q = \frac{x_{\text{MC}}^2}{(1 + a_{\text{sim}} x_{\text{MC}})^2} \,. \tag{3.26}$$

The total distribution for a certain value of a is obtained by summing them as $h(a) = h_c + 2ah_l + a^2 h_q$.

The value of the slope is found by a minimisation of the χ^2 with respect to a using the MINUIT MIGRAD algorithm [64]. The statistical uncertainty δa corresponds to the 1σ interval: $\chi^2(a \pm \delta a) = \chi^2(a) + 1$.

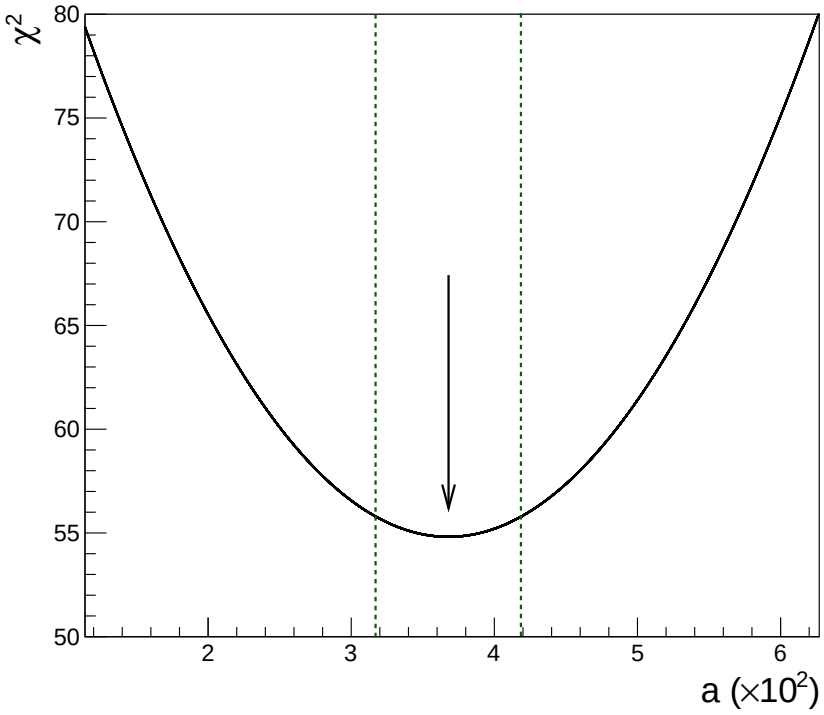

Figure 3.14: χ^2 as a function of the TFF slope value. The number of degress of freedom (ndf) is 49. The arrow indicates the fit central value (minimum χ^2), and the dashed lines represents the 1σ statistical uncertainty band.

The measured value of the π^0 transition form factor slope using the NA62 2007 data sample is

$$a = (3.68 \pm 0.51) \times 10^{-2} \,, \tag{3.27}$$

where the total statistical uncertainty is quoted. The uncertainty is dominated by the size of the data sample ($\delta a_{\rm data} = 0.48 \times 10^{-2}$). The contribution of the MC sample size is smaller ($\delta a_{\rm MC} = 0.18 \times 10^{-2}$). The χ^2/ndf corresponding to the data/MC(a) comparison is 54.8/49, which has a p-value of 26.4 %. The $\chi^2(a)$ function is shown in Figure 3.14.

Figure 3.15 shows an illustration of the fit result, where the ratio data/MC($a = 0$) is plotted. For clarity only 20 equipopulous bins are used, with the markers placed at the bin barycentres. The red solid line is the $|\mathcal{F}(x)|^2$ function with a slope equal to the fit central value and the dashed lines indicate the 1σ band.

The result is not sensitive to the inclusion of the quadratic term in the expansion $|\mathcal{F}(x)|^2 = 1 + 2ax + bx^2$: the central value for a changes by $\Delta a = +0.03 \times 10^{-2}$, while

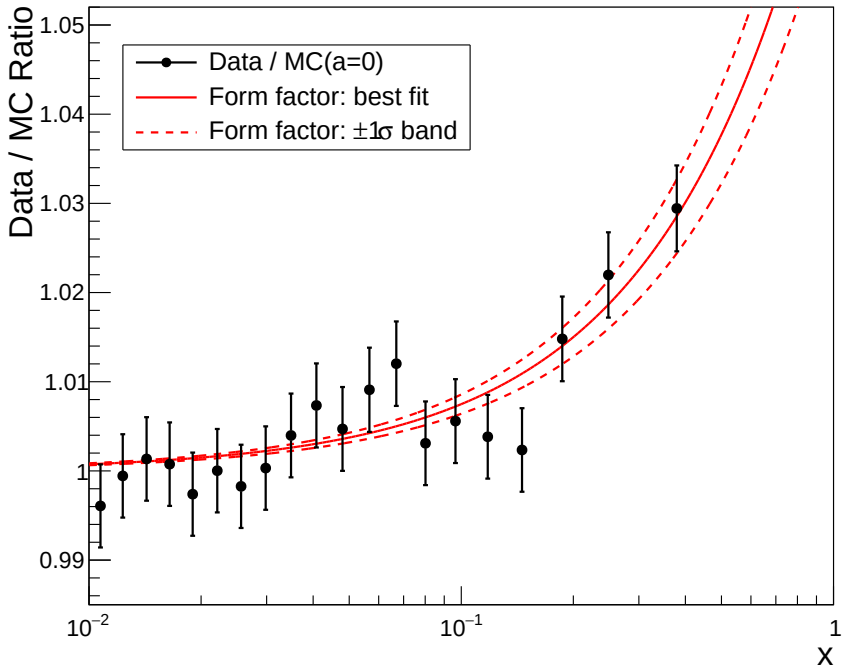

Figure 3.15: Illustration of the fit to the TFF. The effect of a positive TFF slope ($a > 0$) is clearly seen in the ratio of the reconstructed data and MC distributions, with the MC sample weighted to obtain $a = 0$. Data and MC events are divided into 20 equipopulous bins; the horizontal positions of the markers correspond to the bin barycentres. The solid line represents $|\mathcal{F}(x)|^2$ with the measured slope value. The dashed lines indicate the $\pm 1\sigma$ band. Only the statistical uncertainties are included.

$b = 0.00 \pm 0.19$.

3.7 Systematic effects

3.7.1 Trigger efficiency

The trigger used in this analysis is the main trigger line of the experiment designed to select events with electron ($e^{\pm}$) tracks, as described in section 3.1. The π^0_D events contain two electron tracks and the trigger efficiency is therefore expected to be high. This is verified by testing the relative efficiencies between the main trigger stream and the control and secondary streams. Any event selected by a trigger line also contains the decisions of all other streams, allowing the fraction of events for which a trigger condition is inefficient to be measured. In all cases tested, corresponding to each individual trigger condition of the trigger stream, no inefficient event was found and the upper limits on the trigger inefficiencies are given. They are computed from the total numbers of events in the control trigger samples, assuming a binomial distribution.

Q_1 trigger

The Q_1 inefficiency was measured in the R_K analysis [53] for a sample of $K_{\mu 2}$ events triggered by the NHOD control trigger. The inefficiency was found to be $(1.4 \pm 0.1)\,\%$. The same technique is used here, applying the $K_{2\pi D}$ selection to the events selected by the NHOD trigger stream. A total of 5280 events are selected and all of them are tagged with a positive decision from the Q_1 trigger condition. The upper limit on the inefficiency is $0.06\,\%$ at $90\,\%$ CL. Contrary to the $K_{\mu 2}$ decay with only a single charged track, the $K_{2\pi D}$ decay features three charged tracks in the final state, greatly increasing the probability to detect at least one of them.

To further test this condition, a simulation of the efficiency is used on the $K_{2\pi D}$ MC events. A model of the HOD is used, where a 0.02 cm gap is present between the quadrants.

The tracks are extrapolated to the HOD position and a detection efficiency is assigned. If the hit position is inside a gap, the detection efficiency is set to 0; otherwise it is set to 99.76 %. The gap width and the efficiency were tuned by a data/MC comparison in other decay channels. A track is considered detected/undetected according to a random distribution following the assigned efficiency. The event is discarded if all tracks are undetected. The fraction of rejected events is 2.4×10^{-4} and the change of the TFF slope fit result is $|\Delta a| < 0.01\times10^{-2}$.

1TRKL(M) trigger

Similarly, this condition was studied in the R_K analysis and the inefficiency is reported to be smaller than the Q_1 inefficiency. The lower limit on the number of hits was meant to detect single-track events and is always satisfied for π^0_D events containing three tracks. However, the upper limit of 15 hits in a single view of DCH, active during periods 5 and 6, could be an issue for events in which a non-detected energetic photon initiates a shower in the spectrometer or the beam pipe.

The same sample of events as above, passing the NHOD trigger chain, is used. All the 2910 selected candidates in periods 5 and 6 were also tagged as accepted by the 1TRKLM condition. The upper limit on the inefficiency due to the hit multiplicity cut-off is 0.10 % at 90 % CL. When processing the sub-samples for periods 1 to 4 and 5 to 6 separately, a comparison of several DCH related distributions does not show any significant difference and the fit results are statistically compatible, as seen in Table 3.7.

Another verification is performed by removing from the $K_{2\pi D}$ MC sample the events with an energetic radiative photon that could generate such a shower. The MC events where the less energetic of the two generated photons has at least 0.5 GeV of energy and an extrapolated radius in the first drift chamber of $R_{\mathrm{DCH1}} < 11$ cm are removed. In such a configuration the photon has the possibility to interact with the beam pipe or the DCH flanges and initiate a shower. The fraction of events removed is 0.10 % (matching the above upper limit for inefficiency) with no impact on the measured TFF slope.

E_{LKr} trigger

The $E_{\mathrm{LKr}}(10\,\mathrm{GeV})$ condition is not present in the downscaled $K_{\mu 2}$ trigger chain ($Q_1 \times$ 1TRKL(M)). The $K_{2\pi D}$ selection is applied to this stream and 12 108 events are selected. All of them are also positively tagged by the online E_{LKr} condition and the upper limit on the efficiency is 0.03 % at 90 % CL.

A potential effect on the measurement is investigated by removing from the data sample the 0.03 % of events with the lowest LKr energy deposit. These events have the highest probability of being inefficient. The effect on the TFF slope is due to the correlation between the LKr energy deposit and the reconstructed x variable and amounts to $|\Delta a| = 0.06\times10^{-2}$. This shift is taken as a systematic uncertainty.

$L3(K_{e2})$ trigger

A control trigger stream called $L3(K_{e2})$-autopass, which disregards the decision of the $L3(K_{e2})$ algorithm, was acquired. From the 11 095 events of this autopass sample passing the $K_{2\pi D}$ selection, all are correctly tagged as passing the online $L3$ trigger. The upper limit on its inefficiency is 0.03 % at 90 % CL.

In a similar way as discussed above, the 0.03 % of events most likely to fail that condition are removed from the data sample. Since this algorithm consists of two conditions ($p_{\mathrm{min}} > 5\,\mathrm{GeV}/c$ and $E/p > 0.6$) the test is done separately for each one. In one case the 0.03 % of the events with the smallest $p' = \max\{p_{e^+}, p_{e^-}\}$ are discarded, corresponding to those most likely to fail the first condition. In the other one, 0.03 % of the events with the smallest $(E/p)' = \max\{(E/p)_{e^+}, (E/p)_{e^-}\}$ are rejected, corresponding to those most likely to fail the second condition. For both samples the change in the fit result $|\Delta a| < 0.005\times10^{-2}$ is negligible.

3.7.2 π^0 Dalitz decay generator

As explained in subsection 1.1.1, the π^0 Dalitz decay generator implemented for this analysis [22] contains a cut-off parameter x_γ^{cut} controlling the generation of the extra radiative photon in the final state. A lower value for this parameter means more low energy radiative photons are simulated (4-body decay), while for a higher value only hard radiative photons are simulated, integrating the contribution of soft photons into the initial Dalitz event distribution (3-body decay). The physics is in principle correctly reproduced for $x_\gamma^{\text{cut}} \to 0$ but practical difficulties arise as the matrix element squared diverges as $x_\gamma^{\text{cut}} \to 0$. The computational time required for the acceptance-rejection algorithm increases and a trade-off between the amount of time required to produce the MC sample and the accuracy of the simulated process must be chosen. However the TFF fit result should converge as the softer photons are generated and the 4-body decay events become experimentally indistinguishable from the 3-body decay events.

The convergent behaviour is tested by using the true MC generated observables of three $K_{2\pi D}$ MC samples generated with $x_\gamma^{\text{cut}} = 0.05, 0.01, 0.005$. Only the events with two photons in the final state (one from π_D^0 decay and a radiative one) are of interest for this purpose. Since the two photons are indistinguishable, they are sorted according to their energies in the laboratory frame. The less energetic is labelled as the radiative photon, while the most energetic is called the Dalitz decay photon.

The $K_{2\pi D}$ selection described in section 3.5 includes a cut at 2 GeV on the reconstructed energy of the clusters in the LKr calorimeter. The energy distribution of the radiative photons is shown in Figure 3.16. The two samples with $x_\gamma^{\text{cut}} = 0.01, 0.005$ agree well between each other for radiative photon energies above 2 GeV. On the other hand the third sample with $x_\gamma^{\text{cut}} = 0.05$ only agrees above 3 GeV, rendering the fit unstable.

The second cut applied to the photon clusters, which is strongly influenced by x_γ^{cut}, is the limit on the distance between the photon cluster and the impact point of the undeflected $e^\pm$ tracks on the LKr front plane. The true distributions for photons with $E > 2$ GeV are shown in Figure 3.17. Again the two samples with smaller x_γ^{cut} values

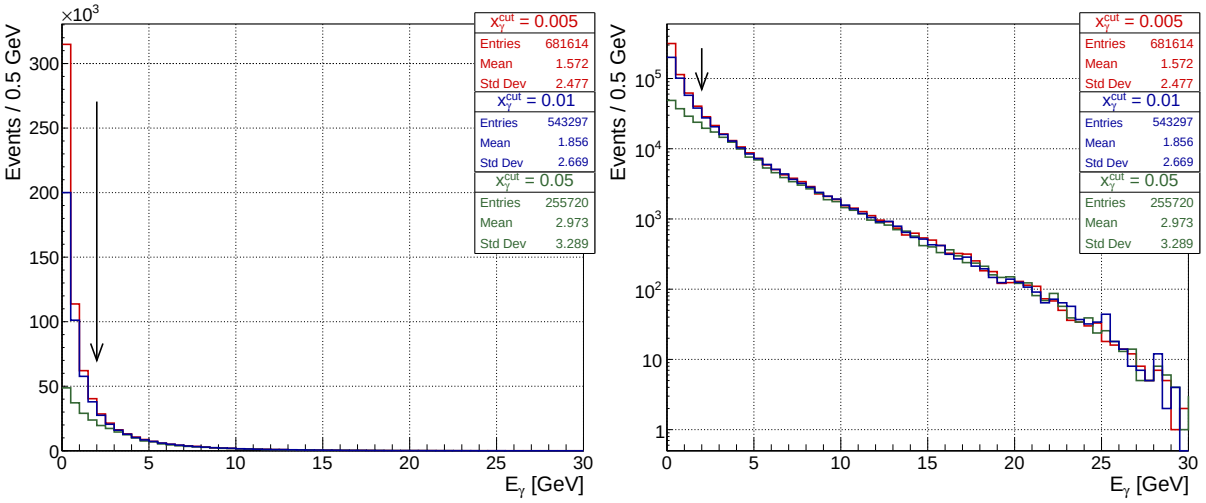

(a) (b)

Figure 3.16: True MC energy of the radiative photon for the three samples generated with different x_γ^{cut} values in (a) linear and (b) logarithmic scale. The arrow represents the cut applied on the photon energy in the selection (2 GeV). The samples with $x_\gamma^{\text{cut}} = 0.005, 0.01$ (red and blue) are in agreement at that energy.

agree well above 15 cm, while the convergence of the third sample is not satisfactory.

In this context, the sample generated with $x_\gamma^{\text{cut}} = 0.05$ is deemed not to be suitable for this analysis. On the contrary, using the generator with $x_\gamma^{\text{cut}} = 0.01$ produces a stable result. The convergence is confirmed by the other sample with $x_\gamma^{\text{cut}} = 0.005$, which is equivalently good in the context of this analysis. The fit result with this sample yields

$$a = (3.46 \pm 0.54) \times 10^{-2} \, . \tag{3.28}$$

The MC contribution to the uncertainty is $\delta a_{\text{MC}} = 0.25 \times 10^{-2}$. As the statistical uncertainty due to data is fully correlated between the two fits, only the MC statistical uncertainties are uncorrelated. The estimate of the uncorrelated statistical error between this result and the default one ($\delta a_{\text{MC}} = 0.18 \times 10^{-2}$) is therefore their sum in quadrature: 0.31×10^{-2}. The shift between the two TFF slopes obtained is $\Delta a = -0.22 \times 10^{-2}$, and therefore the two results are statistically compatible with each other and it is concluded that no systematic effect is observed due to the cut-off parameter x_γ^{cut}.

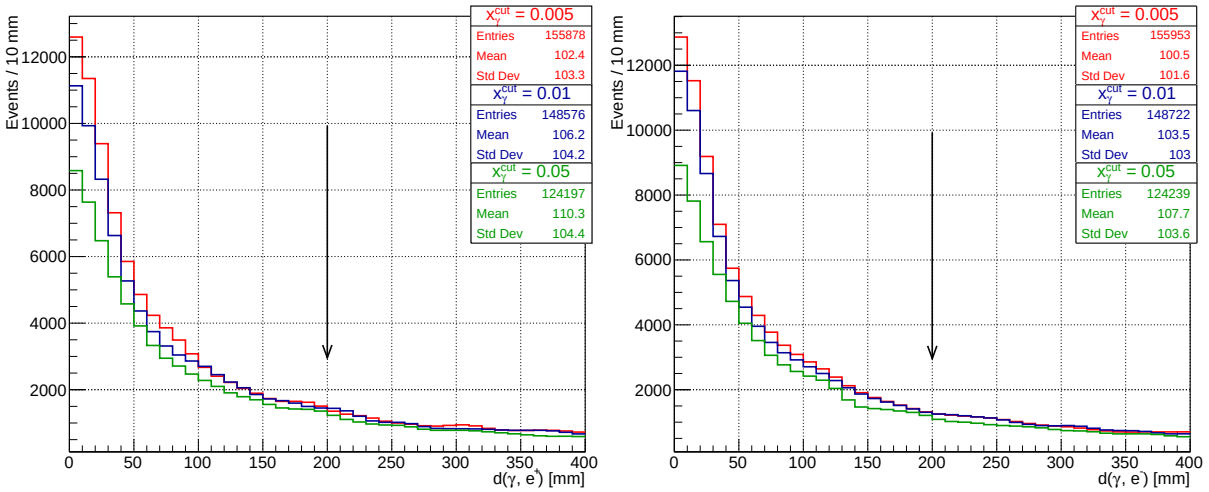

(a) (b)

Figure 3.17: Distance between the radiative photon cluster and the undeflected (a) e^+ (b) e^- track impact point on the LKr front plane. The arrows represent the cut applied on this distance in the selection. The samples with $x_\gamma^{\text{cut}} = 0.005, 0.01$ (red and blue) are in good agreement between each other.

3.7.3 Spectrometer calibration and resolution

The (α, β) factors correcting for the magnetic field integral and spectrometer chamber misalignments, applied to the tracks measured in the spectrometer, are explained in detail in section 3.4. This approach is however an approximation relying on a model of the spectrometer, and more parameters could be necessary to fully describe the calibration. The parameters are also measured statistically for intervals of time that might not discern faster changes. It can potentially affect the measurement of the TFF slope due to the correlation between the $e^\pm$ momenta and the x variable. A 1 % effect on the momentum calibration (data only or MC only) generates an additional slope of $|\Delta a| = 7.8{\times}10^{-2}$ in the x distribution. The sensitivity of the fit to a residual miscalibration is assessed by switching on/off the corrections for data and MC samples independently, as shown in Table 3.3.

The shift of the central value is dominated by the β correction. Th effect of the α contribution alone is found to be negligible, $\Delta a < 0.01{\times}10^{-2}$, which is expected since the factor α is two orders of magnitude smaller than the factor β. Switching off the

Table 3.3: Summary of the spectrometer correction effects on the fit result. The (α, β) corrections are applied independently to data and MC in all possible combinations. The correct treatment is to turn them On/Off simultaneously for data and MC (first and last row).

(α, β) correction	Fit result ($\times 10^2$)	$\Delta a(\times 10^2)$
Data:On, MC:On	3.68	0.00
Data:On, MC:Off	3.81	0.13
Data:Off, MC:On	3.40	-0.28
Data:Off, MC:Off	3.52	-0.16

corrections for data only or MC only is not the proper treatment, and therefore the shift from the last row of the table, where the (α, β) corrections are turned off for both data and MC, is considered as the systematic uncertainty estimate: $\delta a = 0.16 \times 10^{-2}$.

The spectrometer momentum scale correction extracted from $K_{3\pi}$ samples ensures that the reconstructed mass peak is not shifted. This also means that the probable reason for the slight shift of the reconstructed π^0 and K mass peaks observed in the $K_{2\pi D}$ sample (Figure 3.18) is a mis-calibration of the photon energy measured in the LKr.

The effect of the spectrometer resolution on the resolution of the x variable is studied from the MC $K_{2\pi D}$ sample. From the distribution of reconstructed and true x values shown in Figure 3.19a, the resolution as a function of x is parametrised as $\sigma(x) = 0.018x$ (see Figure 3.19b). The detector response simulation implemented in CMC (Section 3.2) should reproduce most of the resolution effects present in the data sample. However a residual mismatch between the simulation and the data might be present. Since there is no closed kinematics decay channel providing an abundance of electron-positron pairs in the final state, the mismatch is estimated from the comparison of the reconstructed invariant mass of a $K^{\pm} \to \pi^{\pm}\pi^{+}\pi^{-}$ sample. A simple event selection is applied to a data sample from the $K_{\mu 2}$ trigger chain and a $K_{3\pi}$ MC sample.

The reconstructed $M_{3\pi}$ invariant mass and the run-dependent resolution on $x_{3\pi} = M_{3\pi}^2/M_K^2$ are shown in Figure 3.20. The sample is virtually free of background since $K_{3\pi}$ is the dominant decay with three charged tracks in the final state. The resolution is generally overestimated by the MC and the maximal difference of the ratio $\sigma(x_{3\pi})_{\rm data}/\sigma(x_{3\pi})_{\rm MC}$ from

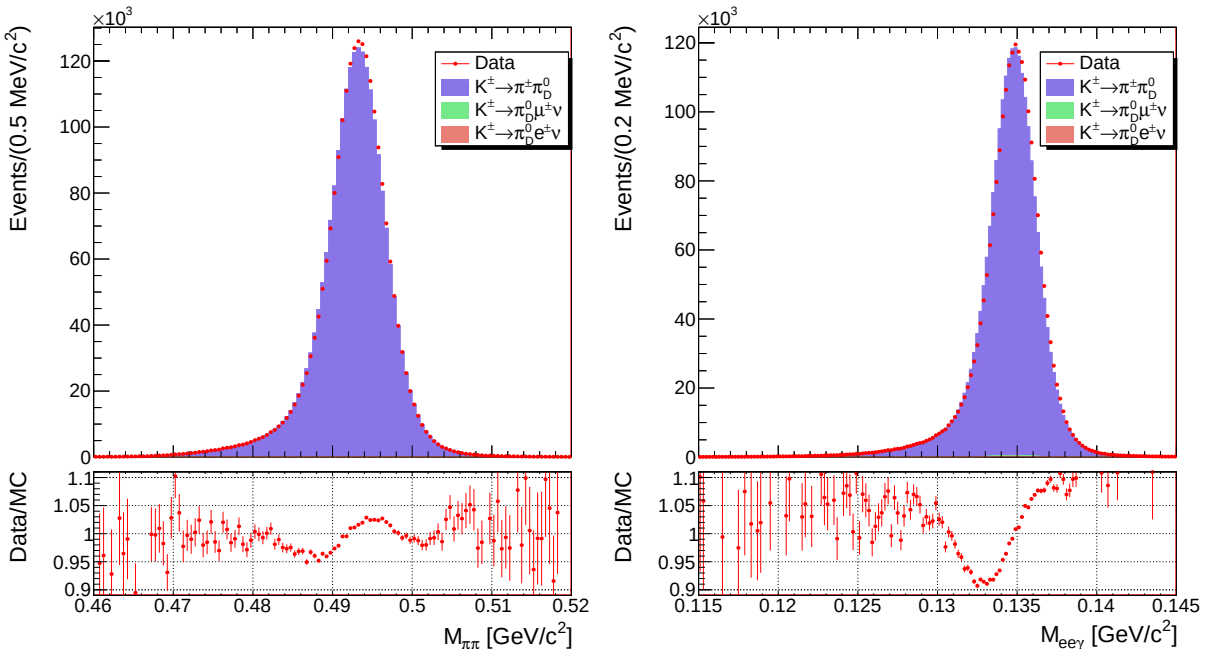

(a) (b)

Figure 3.18: (a) Reconstructed kaon mass, (b) reconstructed π^0 mass. A small shift observed between data and MC is attributed to a mis-calibration of the photon energy.

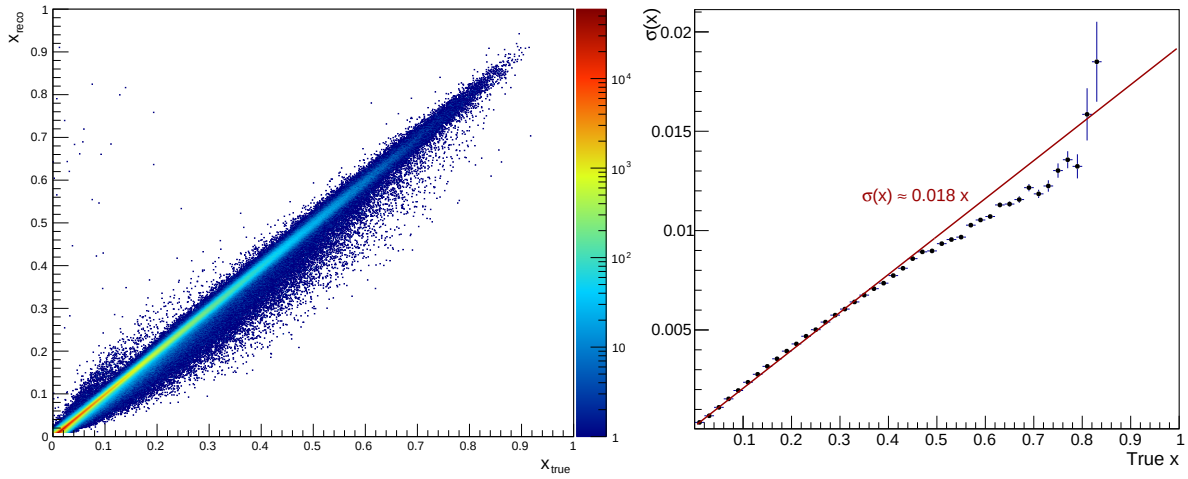

(a) (b)

Figure 3.19: (a) Distribution of the $K_{2\pi D}$ MC events in the $(x_{\text{true}}, x_{\text{reco}})$ plane. (b) Resolution on the x variable estimated from the same sample.

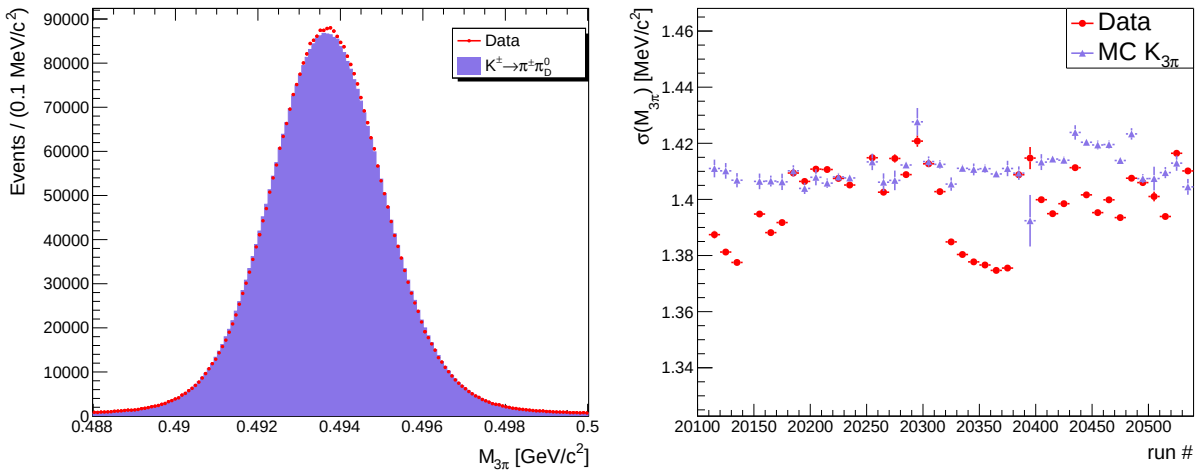

(a) (b)

Figure 3.20: (a) Distribution of the $M_{3\pi}$ invariant mass for data and MC $K_{3\pi}$ samples. (b) Resolution on $M_{3\pi}$ for data and MC as a function of the run number.

unity is ~ 0.02.

The π^0 TFF slope sensitivity to the resolution mismatch effect can be estimated by an introduction of a MC resolution scaling factor K. The reconstructed x value of each selected MC event is replaced in the fit procedure by the following combination of true (x_{true}) and reconstructed (x_{reco}) values:

$$x = x_{\text{true}} + K(x_{\text{reco}} - x_{\text{true}}) \, . \tag{3.29}$$

Setting $K = 1$ corresponds to the reconstructed x distribution. Figure 3.21 shows the sensitivity of the fit result to the MC scaling factor. It is assumed that the maximum relative resolution discrepancy between data and MC is the same for both Dalitz x and $x_{3\pi}$ variables: $\sigma(x)_{\text{Data}}/\sigma(x)_{\text{MC}} = 0.98$. The difference between the default fit result and the one corresponding to $K = 0.98$ ($a = 3.63{\times}10^{-2}$) is considered as the systematic uncertainty: $\delta a = 0.05{\times}10^{-2}$.

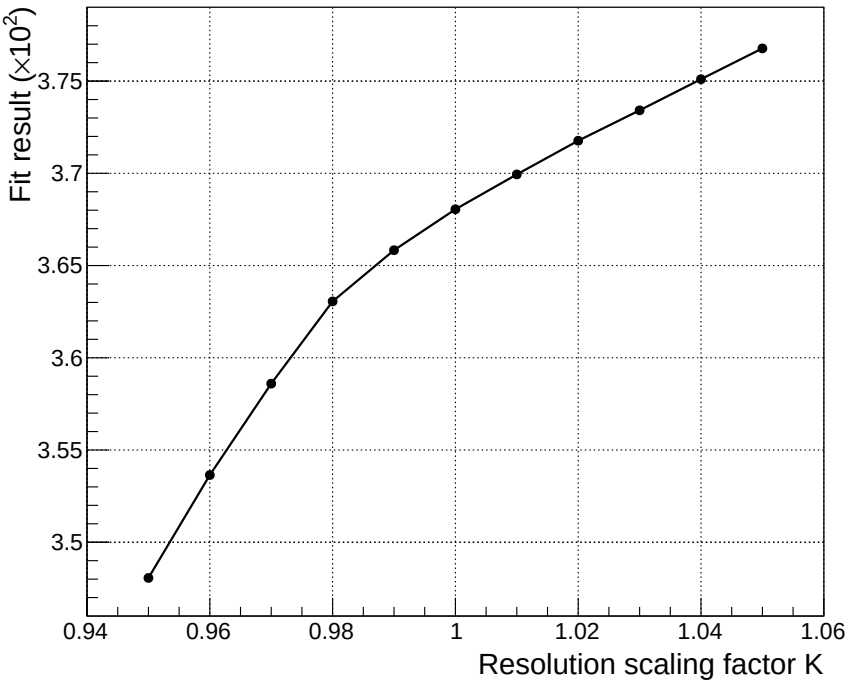

Figure 3.21: Effect of the scaling factor K on the fit result.

Table 3.4: Summary of LKr correction effects on the fit result.

Correction	Fit result ($\times 10^2$)	$\Delta a(\times 10^2)$
Standard	3.68	0.00
Non-linearity correction: Off	3.61	-0.07
Non-linearity correction: Alternative	3.71	0.03
Cell by cell re-calibration: Off	3.68	0.00
Global MC E_{LKr} scaling factor	3.70	0.02
Cluster position projectivity correction: Off	3.68	0.00

3.7.4 LKr calorimeter calibration

A mentioned in subsection 3.7.3, a mis-calibration of the cluster energy is the most probable cause for the small shift of the reconstructed M_K and M_{π^0} mass distributions (see Figure 3.18). It can only enter the final result indirectly through a modification of the acceptance, and no strong effect on the TFF slope is expected.

A similar method to that described in subsection 3.7.3 is used, by switching on/off the individual corrections. The effect of the non-linearity function (see subsection 3.4.3) is also tested by using an alternative parametrisation.

The largest effect in Table 3.4 is observed when the non-linearity correction is removed,

changing the acceptance for low energy photons. A MC global scaling factor of 1.001 is sufficient to correct the shift observed in the M_K and M_{π^0} distributions. This scaling factor is of the same order as the run-dependent cell by cell recalibration (0.999 with RMS width of 0.004).

The shifts of the fitted central value due to the alternative non-linearity correction function and the global MC energy scale are used as the estimate of the systematic uncertainty due to the LKr energy calibration: $\delta a = 0.04 \times 10^{-2}$.

3.7.5 Accidentals

Events with accidental activity can sometimes display a signature similar to a π^0_D event and reach the final data sample. These could be for example $K_{3\pi}$ events with an accidental in-time photon. A possible effect on the fit result is investigated by releasing the cuts related to the suppression of accidental activity.

Cluster timing condition

To estimate the number of accidental clusters entering the event time-window, the cluster timing condition is reversed. Instead of being in a 10 ns wide window around the trigger, the cluster is requested to be between 15 ns and 35 ns away from the trigger time. In this way only 11 events with accidental calorimetric activity are selected in this window, which is twice as large as the usual one, accounting for 5.5 effective events in the standard selection. No impact on the fit result is observed by their addition to the standard data sample.

Extra tracks condition

An additional 1896 data events are selected by allowing extra tracks in the event, not related to the main vertex. Their x distribution is shown in Figure 3.22a. The normalised number of MC events with at least one extra track is 103. The addition of both data and

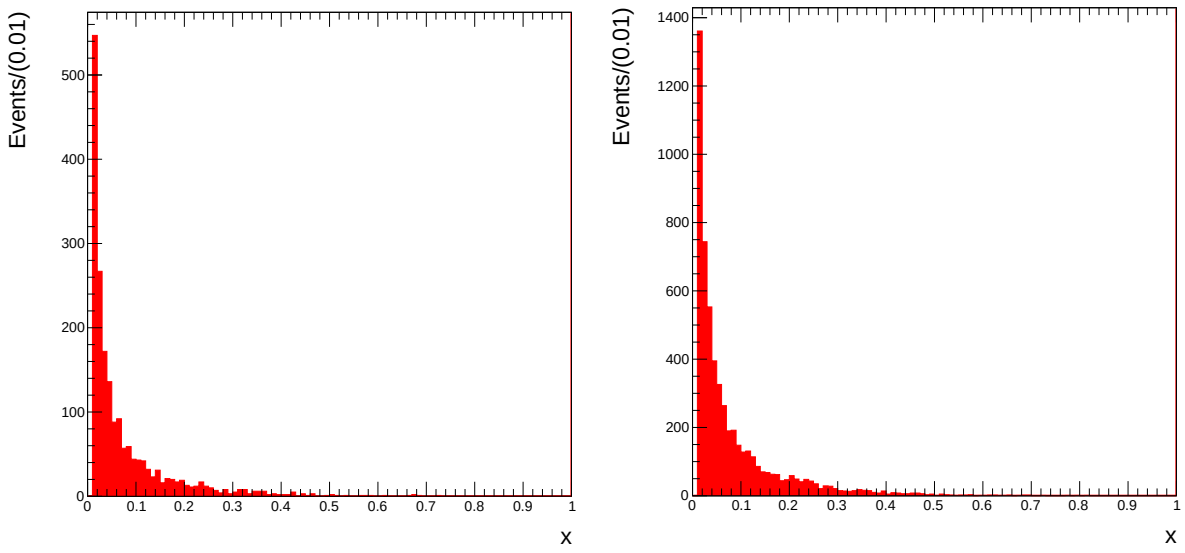

(a) (b)

Figure 3.22: (a) x distribution of the 1896 additional data events selected when releasing the extra track condition. (b) x distribution of the 5600 additional data events selected in the modified selection allowing multiple vertices.

MC events to the standard sample shifts the TFF slope fit result by $\delta a = -0.06 \times 10^{-2}$.

Multiple vertices

The standard selection requires a single reconstructed three-track vertex. Ghost tracks are non-physical tracks built from DCH hits produced by different charged particles. Additional three-track vertices containing such fake tracks are sometimes reconstructed. A modification of the selection allowing multiple vertices is considered: the whole process is executed on each vertex separately and the remaining vertex with the lowest χ^2 is selected in the final stage. The number of data events with at least two good vertices is 5600 and their x distribution is shown in Figure 3.22b. The number of selected MC events is 1105. The fit result on the samples containing these events changes by $\Delta a = -0.09 \times 10^{-2}$ with respect to the standard result.

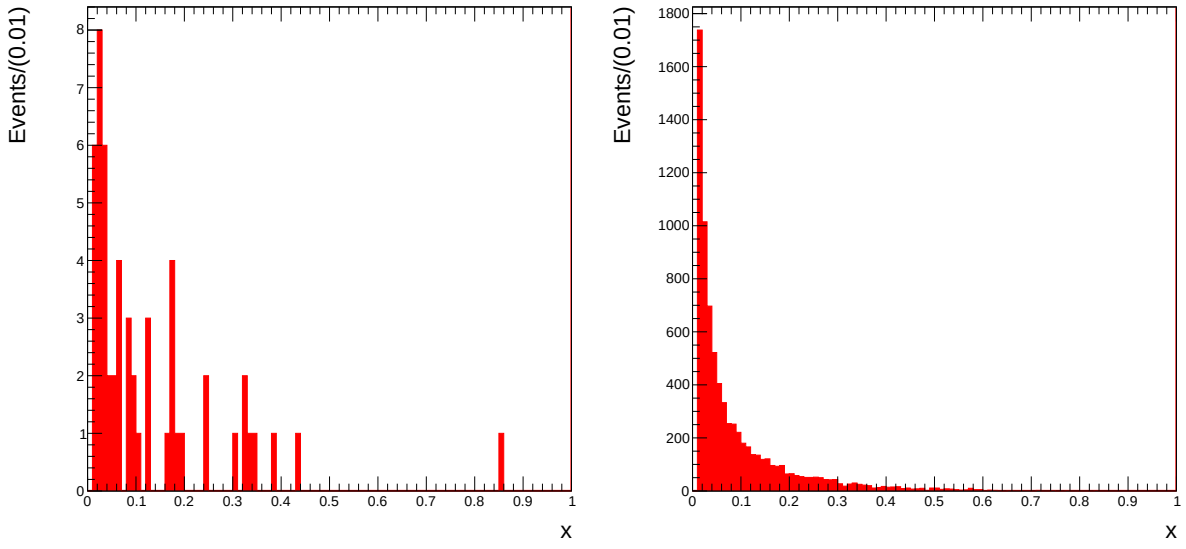

(a) (b)

Figure 3.23: (a) x distribution of data events with $q_{\text{vtx}} = -q_{\text{beam}}$. (b) x distribution of data events with out-of-time tracks.

Vertex charge

The standard selection requires the vertex charge to be $|q_{\text{vtx}}| = 1$, and the sign to be equal to the beam kaon charge sign during single-beam runs. Releasing the latter condition for the data sample adds 54 events to the sample (Figure 3.23a), changing the TFF slope fit result by $\Delta a = 0.01 \times 10^{-2}$.

Track timing

The spectrometer tracks are also subject to effects due to accidental activity. This is checked by releasing the constraint for the reconstructed tracks to be within 25 ns of the trigger time. The maximal time differences of 15 ns between tracks and 10 ns between the cluster and the vertex time are kept. An additional 5870 π^0_D events, which were out-of-time with respect to the trigger, are selected (Figure 3.23b and Figure 3.24). The effect of their addition into the fit is a shift of $\Delta a = 0.10 \times 10^{-2}$.

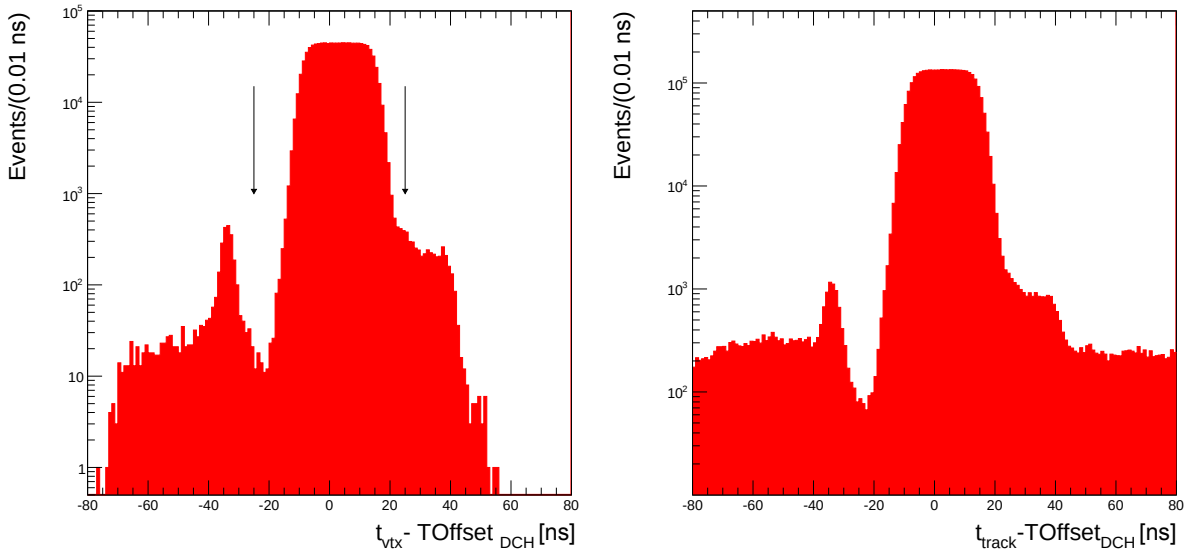

(a) (b)

Figure 3.24: Absolute time of reconstructed vertex (a) and track (b). The arrows represent the default timing cut condition (± 25 ns), which is applied to individual track times. The cut value is given by the 25 ns clock period used in the experiment. The peak at -36 ns, which is present in the data sample for all three-track events, can be attributed to the dead time of the HOD L0 trigger.

The systematic uncertainty assigned for the accidentals is taken as the sum in quadrature of all the fit result shifts summarized in Table 3.5: $\delta a = 0.15 \times 10^{-2}$

3.7.6 Particle identification

The approach taken in this analysis for the particle identification relies only on the kinematics of the event and is independent of the E/p distributions of individual $e^{\pm}$ and $\pi^{\pm}$ tracks. It is more robust, stable and increases the acceptance with respect to the tradi-

Table 3.5: Summary of fit result differences for the variations of the selection regarding accidental activity.

Accidental check	Fit result ($\times 10^2$)	$\Delta a (\times 10^2)$
Cluster timing	3.68	0.00
Extra track	3.62	-0.06
Multiple vertices	3.59	-0.09
Vertex charge	3.69	0.01
Track timing	3.78	0.10

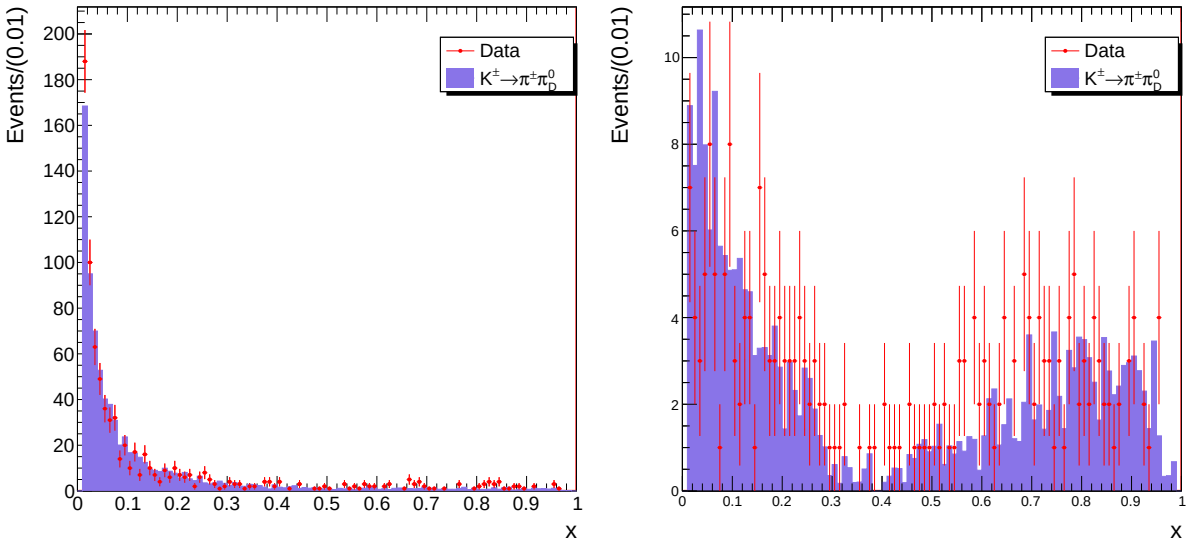

(a) (b)

Figure 3.25: Reconstructed x distribution of the additional events passing the modified selection allowing two valid PID hypothesis (a) and for the LNV selection (b).

tional approach by a factor of more than 2. It can however introduce a mis-identification in two different ways. In the first case the tracks of genuine π^0_D events are swapped ($\pi^{\pm} \leftrightarrow e^{\pm}$), because both kinematic configurations are similar and the correct one is slightly outside one of the bounds. In the other case three-track non-π^0_D events are compatible with the π^0_D kinematics when the wrong particle content is assigned (e.g. $K_{3\pi}$ with two $\pi^{\pm}$ mis-identified as $e^{\pm}$). Two independent modifications of the standard selection are made to select mis-identified events and test their effect on the fit result.

The first modification allows events where both PID hypotheses are valid, which constitutes less than 0.1 % of the sample. A total of 593 data events and 575 normalised $K_{2\pi D}$ MC events with both valid hypotheses are selected. Their x distributions are shown in Figure 3.25a. To continue the selection and perform the fit, one of the hypotheses is chosen randomly, assuring that approximately half of the additional events have swapped $e^{\pm} \leftrightarrow \pi^{\pm}$. The variation of the TFF slope is $\Delta a = 0.01 \times 10^{-2}$.

The second modification builds a selection for the lepton number violating (LNV) decay $K^{\pm} \to e^{\pm}e^{\pm}\pi^{\mp}\gamma$. The track with charge opposite to q_{vtx} is assigned the charged

pion mass, while the two same-sign tracks are assigned the electron mass. Any event passing this selection can be considered as mis-identified. A total of 188 data events are selected, for 145.45 normalised $K_{2\pi D}$ MC events. Genuine π^0_D events therefore pass the LNV selection, leaving about 42±18 data events as possible background. As seen in Figure 3.25b, the x distribution of events passing the LNV selection is different from the correctly reconstructed π^0_D events, but it is similar between data and MC, mostly cancelling the effect in the fit. The shift of the fit resulting from the addition of the selected events to the default sample is $\Delta a = 0.06\times10^{-2}$.

The total systematic uncertainty assigned due to the particle identification is the sum in quadrature of these two deviations: $\delta a = 0.06\times10^{-2}$.

3.7.7 Neglected π^0 Dalitz decay sources

Along with the main channel producing π^0_D ($K_{2\pi D}$), two other channels were simulated ($K_{\mu 3}$ and K_{e3}). The effects of the latter two channels on the final fit result are limited due to their small branching fractions and small acceptances with respect to the main channel. Neutral pions are, however, produced in other kaon decays, e.g. $K^\pm \to \pi^\pm\pi^0\pi^0$. Because they all have smaller branching ratios, their contribution is expected to be even more suppressed than that of K_{e3} and $K_{\mu 3}$. The total branching fraction of the neglected π^0_D sources is 1.8 %. The requirements on the additional clusters and the kinematic constraints strongly reduce their acceptance. An upper limit is set by removing the $K_{\ell 3}$ ($\ell = \mu, e$) samples from the MC distributions used in the fit. The results are summarized in Table 3.6. The estimate of the systematic uncertainty due to neglected π^0_D sources is conservatively taken to be equal to the removal of the K_{e3} sample: $\delta a = 0.01\times10^{-2}$.

As mentioned in section 3.3, two additional samples, $K_{2\pi}$ followed by $\pi^0 \to \gamma\gamma$ and $K_{3\pi}$, were used to test for a possible contamination. The first one can possibly be selected as a $K_{2\pi D}$ event if one of the photons converts into an e^+e^- pair before the DCH and is compatible with a 3-track vertex topology. No such event is observed, confirming the efficiency of the DCH track distance and LKr cluster distance cuts in the selection. The

Table 3.6: Contribution of the various MC Dalitz samples to the fit result.

MC samples included	Fit result ($\times 10^2$)	$\Delta a(\times 10^2)$
$K_{2\pi}, K_{\mu 3}, K_{e3}$	3.68	0.00
$K_{2\pi}, K_{\mu 3}$	3.67	-0.01
$K_{2\pi}, K_{e3}$	3.66	-0.02
$K_{2\pi}$	3.65	-0.03

$K_{3\pi}$ decay can be mistaken as a $K_{2\pi D}$ event if it is compatible with the kinematic PID when two of the $\pi^{\pm}$ tracks are assigned an electron mass. Again, no $K_{3\pi}$ event is observed passing the selection.

3.7.8 Beam momentum spectrum simulation

As explained in section 3.4, the beam momentum receives a first order correction by introducing p_K-dependent weights. A remaining data/MC discrepancy in the tails of the beam momentum can be seen in Figure 3.1b. It can be improved by applying weights w_{pk2} corresponding to a second order correction to the MC events, as seen in Figure 3.26. These weights were computed during the R_K measurement [53] and are applied simultaneously to the run-dependent weights (as described in section 3.3), taking into account the new effective number of MC events. The TFF slope result changes by $\Delta a = 0.03 \times 10^{-2}$, which is considered as the systematic uncertainty.

3.7.9 Stability checks

In order to see a possible systematic bias due to different data-taking conditions, the sample is divided into groups of sub-samples. The groups are divided according to a specific condition:

- The sample is split between the data-taking periods 1–4 and 5–6. During the first periods, the beam deflection changed from ± 0.30 mrad to ± 0.23 mrad, the Pb bar was installed between the HOD and the LKr, and the DCH trigger was 1TRKL (lower hit multiplicity bound only). The last two periods were taken with a con-

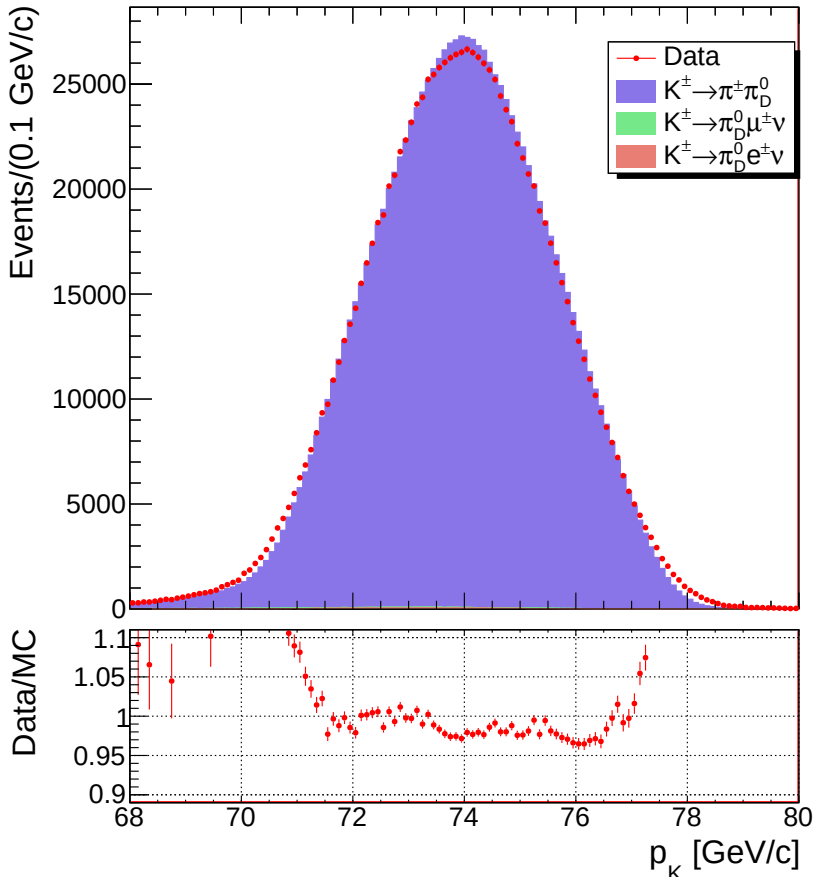

Figure 3.26: Beam momentum after applying the second order weights w_{pk2}.

stant beam deflection of ± 0.23 mrad, the Pb bar was removed and the DCH trigger included the upper hit multiplicity cut (1TRKLM).

- The sample is split between data taken with positive or negative spectrometer magnet polarity.

- The splitting is performed between events with positive or negative vertex charge. The events are not balanced in this group, because of the choice to run mostly with positively charged kaons, driven by the sweeping system optimized for positive muons.

- The second kinematic variable y of the $\pi^0 \to e^+e^-\gamma$ decay is not used in this analysis since it does not enter the form factor. The event acceptance is however sensitive to it and the samples are divided into two roughly equal sub-samples ($y < 0.4$ and $y \geq 0.4$).

- The geometric acceptance is tested by dividing events according to the reconstructed z position of the vertex.

The sub-samples within a group are complementary and statistically independent. The TFF slope fit is performed on each of them and the results are summarized in Table 3.7. The slope values agree between each other within the statistical uncertainties.

To conclude the checks for systematic uncertainties, the cuts introduced in the selection process are varied around their nominal values. For each of the variations the TFF slope fit is performed to verify that no instability can be observed. Figures 3.28 to 3.42 show the fit results and the numbers of selected events for each tested value of each cut. The x-axis indicates the value of the cut for the quantity mentioned in the caption. The fit result is presented with the full statistical uncertainty (red line) and the estimated uncorrelated statistical uncertainty (blue line). The dotted green line indicates the nominal value. No instability is observed, despite small variations of the TFF slope compatible with the uncorrelated statistical uncertainty estimates.

Table 3.7: Fit results for different splittings of the total sample.

Sub-sample	$a \times 10^2$	$\delta a \times 10^2$	# K_{e2} candidates
Full data sample	3.68	0.51	1 114 243
Periods 1–4	3.53	0.73	537 222
Periods 5–6	3.83	0.72	577 021
Spectrometer magnet polarity $B+$	3.61	0.70	600 179
Spectrometer magnet polarity $B-$	3.76	0.76	514 064
Vertex charge $q_{\text{vtx}} = +1$	3.87	0.56	922 857
Vertex charge $q_{\text{vtx}} = -1$	2.84	1.23	191 386
Dalitz variable $y < 0.4$	3.82	0.81	504 085
Dalitz variable $y \geq 0.4$	3.64	0.67	610 158
Vertex position $z_{\text{vtx}} \in (-17; 10)$ m	4.12	0.85	386 265
Vertex position $z_{\text{vtx}} \in (10; 30)$ m	2.86	0.91	349 707
Vertex position $z_{\text{vtx}} \in (30; 90)$ m	4.04	0.90	378 271

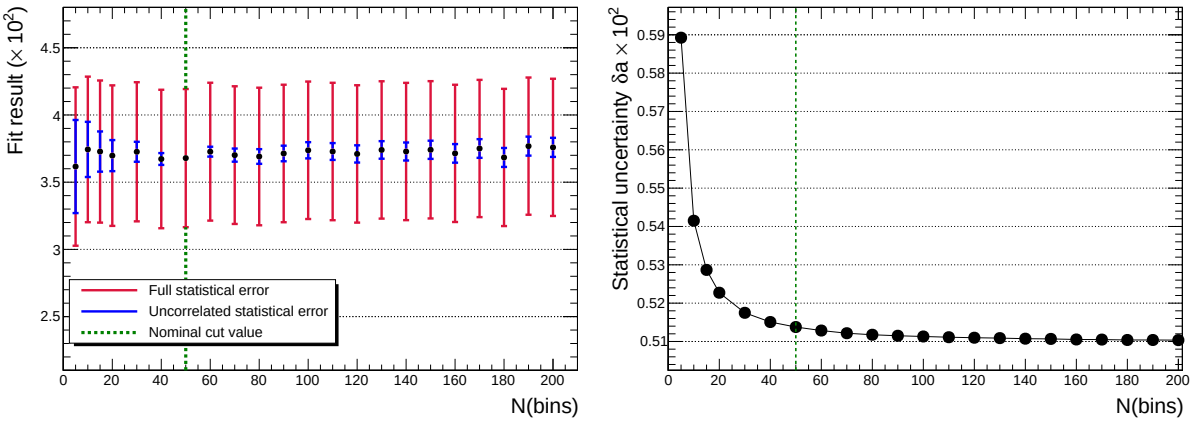

(a) (b)

Figure 3.27: Fit result and statistical uncertainty as a function of the number of bins used in the procedure. The dashed vertical line denotes the default value for this analysis.

The stability of the fit result with respect to the number of equipopulous bins used in the procedure has been assessed. Figure 3.27 shows the fitted central value and the statistical uncertainty as a function of the number of bins. The central value is stable in all configurations and the statistical uncertainty decreases when increasing the number of bins. It stabilises at $\delta a = 0.51 \times 10^{-2}$ from $N_{\text{bins}} \sim 40$ onwards, justifying the choice of 50 bins.

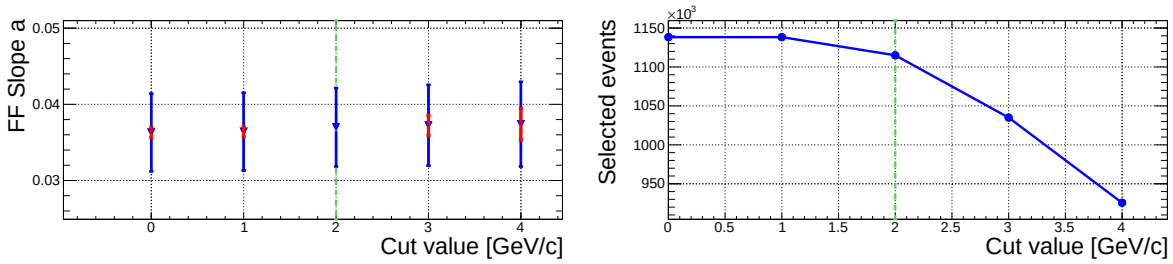

(a) (b)

Figure 3.28: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the minimum $e^{\pm}$ momentum. This cut is applied simultaneously on all tracks in the selection, but the cut value is modified only for tracks identified as $e^{\pm}$ for the purpose of this scan.

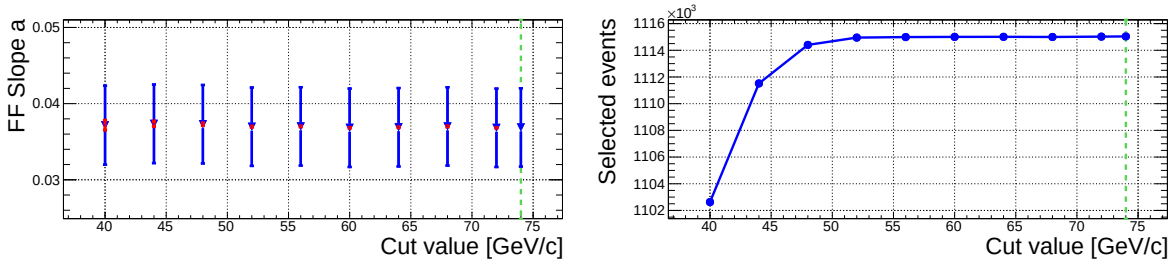

(a) (b)

Figure 3.29: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the maximum $e^{\pm}$ momentum. This cut is applied simultaneously on all tracks in the selection, but the cut value is modified only for tracks identified as $e^{\pm}$ for the purpose of this scan.

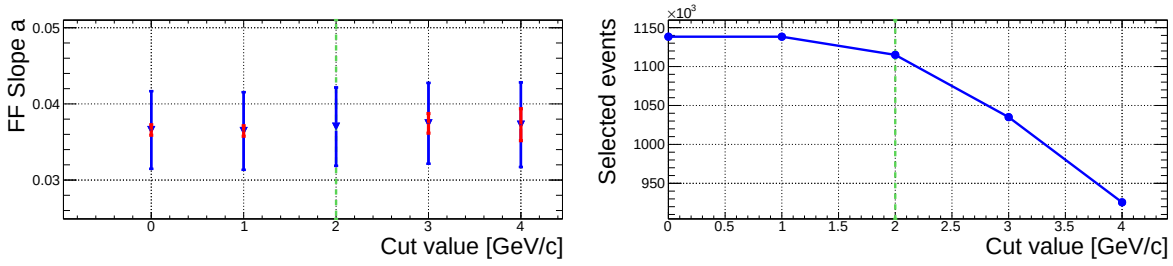

(a) (b)

Figure 3.30: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the minimum $\pi^{\pm}$ momentum. This cut is applied simultaneously on all tracks in the selection, but the cut value is modified only for tracks identified as $\pi^{\pm}$ for the purpose of this scan.

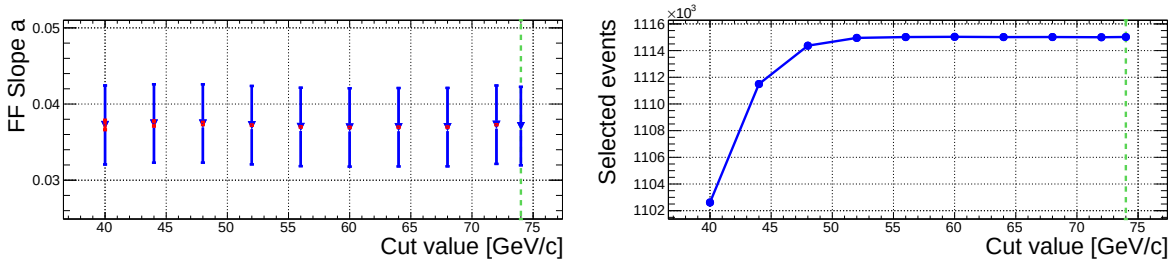

(a) (b)

Figure 3.31: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the maximum $\pi^{\pm}$ momentum. This cut is applied simultaneously on all tracks in the selection, but the cut value is modified only for tracks identified as $\pi^{\pm}$ for the purpose of this scan.

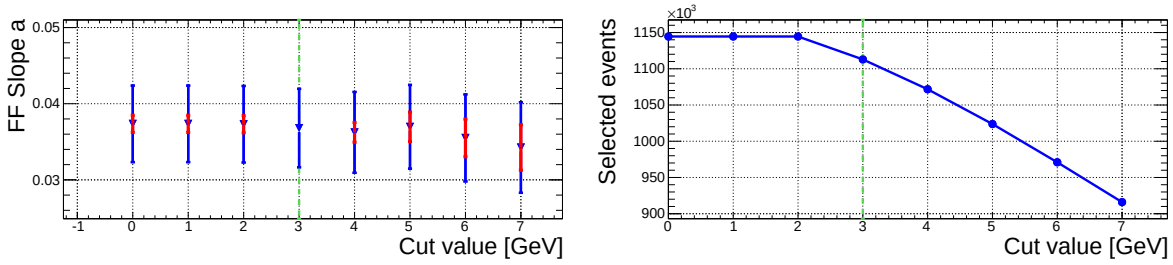

(a) (b)

Figure 3.32: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the minimum γ energy.

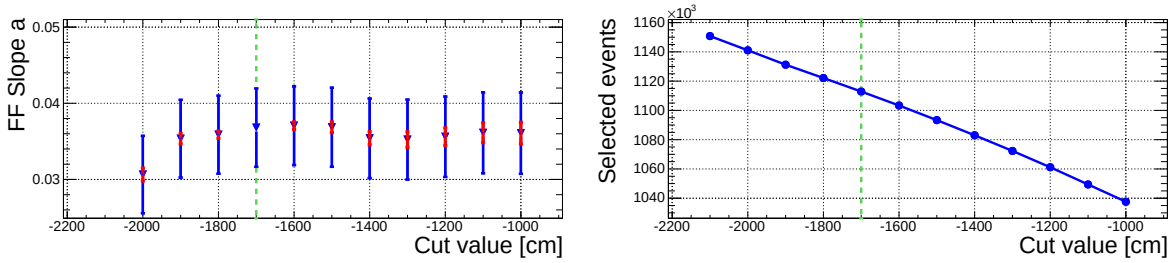

(a) (b)

Figure 3.33: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the minimum vertex z position.

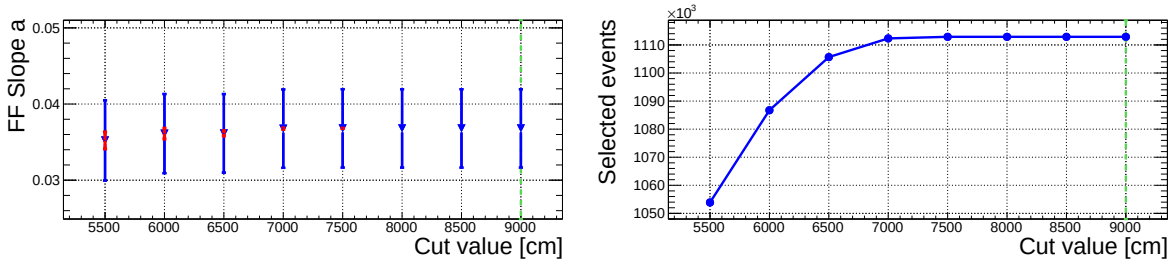

(a) (b)

Figure 3.34: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the maximum vertex z position. The standard cut value of $z_{\mathrm{vtx}} = 9000\,\mathrm{cm}$ is very loose and effectively cutting away no event beyond $z_{\mathrm{vtx}} = 8000\,\mathrm{cm}$.

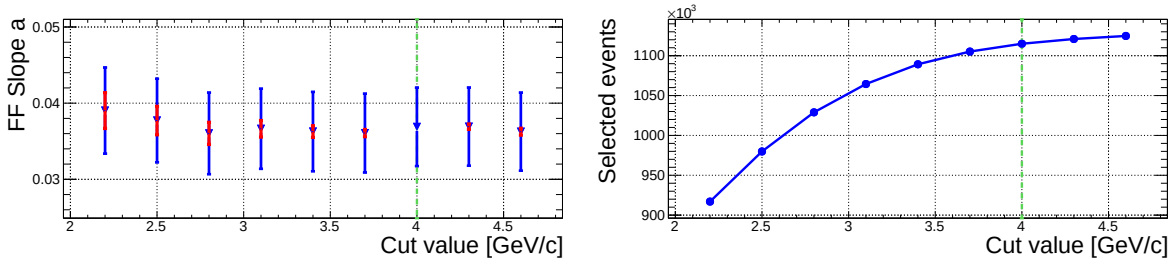

(a) (b)

Figure 3.35: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the maximum allowed deviation of the total reconstructed momentum from the nominal $74\,\mathrm{GeV}/c$ beam momentum.

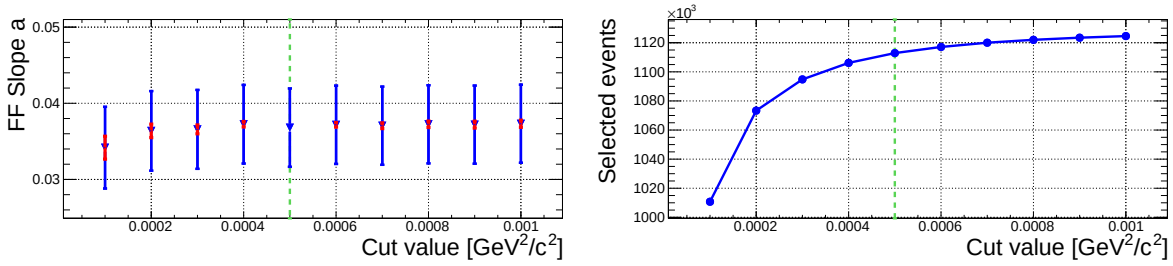

(a) (b)

Figure 3.36: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the maximum allowed square transverse momentum.

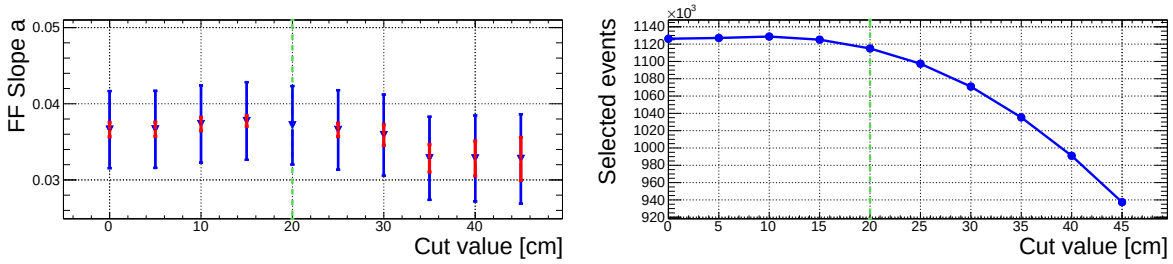

(a) (b)

Figure 3.37: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the minimum distance between photon cluster and track impact point on LKr. The drop of acceptance for the cut at 0 cm (effectively removing the condition) is due to a larger fraction of events with more than one good reconstructed cluster.

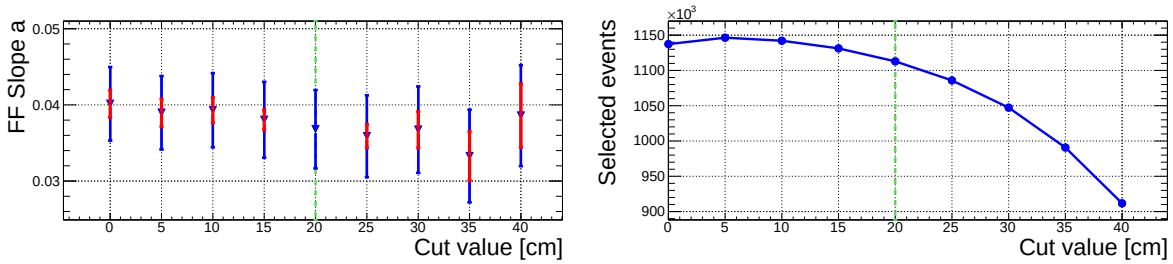

(a) (b)

Figure 3.38: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the minimum distance between photon cluster and undeflected track impact point on LKr. The drop of acceptance for the cut at 0 cm (effectively removing the condition) is due to a larger fraction of events with more than one good reconstructed cluster.

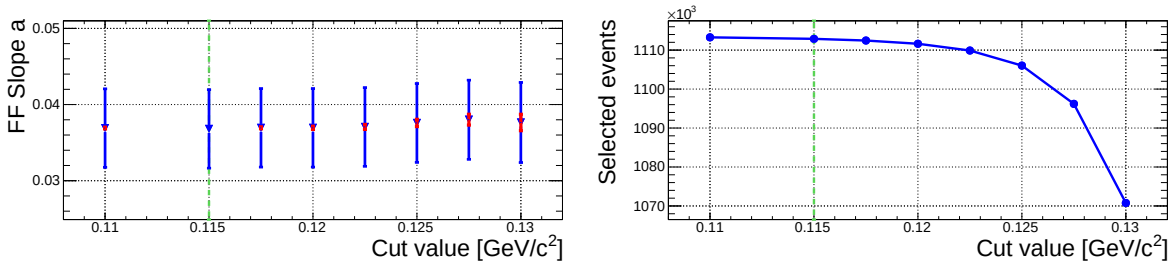

(a) (b)

Figure 3.39: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the minimum reconstructed $ee\gamma$ invariant mass.

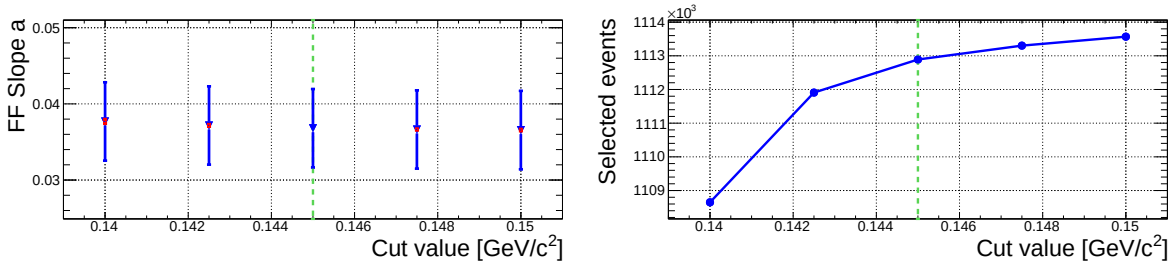

(a) (b)

Figure 3.40: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the maximum reconstructed $ee\gamma$ invariant mass.

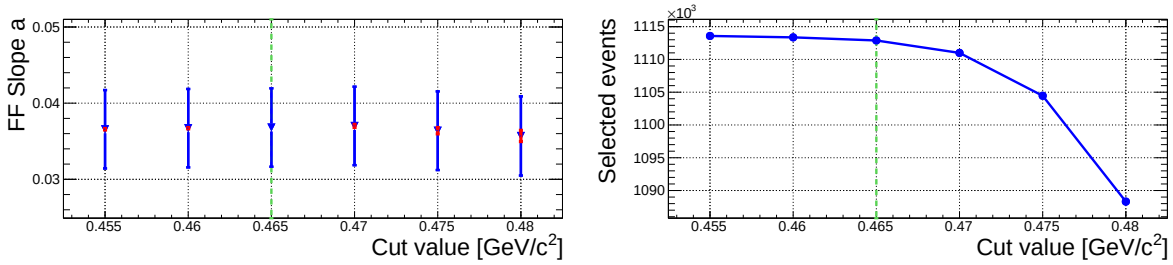

(a) (b)

Figure 3.41: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the minimum reconstructed $\pi^+\pi^0$ invariant mass.

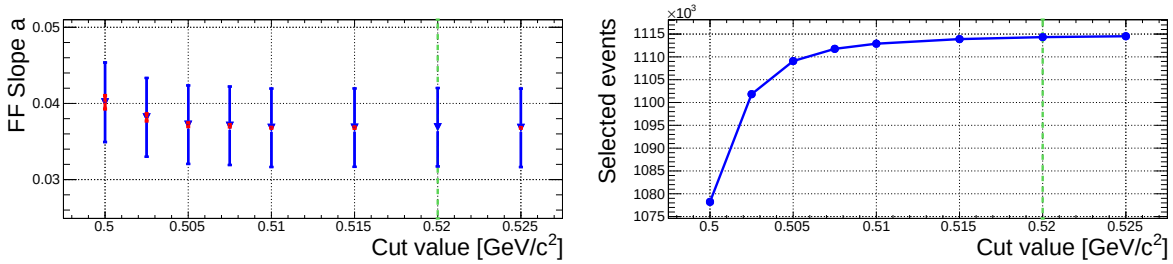

(a) (b)

Figure 3.42: (a) Stability of the fit result and (b) number of selected events as a function of the cut value on the maximum reconstructed $\pi^+\pi^0$ invariant mass.

3.8 Error budget and the final result

The statistical and systematic uncertainties are reported in Table 3.8. The dominant contribution to the total uncertainty is statistical and due to the limited data sample size. The main systematic uncertainties are due to spectrometer calibration and background from accidental activity.

The final result

$$a = (3.68 \pm 0.51_{\text{stat}} \pm 0.25_{\text{syst}}) \times 10^{-2} = (3.68 \pm 0.57) \times 10^{-2} \tag{3.30}$$

is compatible with the theoretical predictions, summarised in section 1.1, being $\sim 1\sigma$ above the VMD expectation. The 15 % relative uncertainty represents an improvement of a factor 2 with respect to the previous best measurement [41]. The Figure 3.43 shows a comparison of this result with respect to previous measurements in the time-like region of momentum transfer.

The uncertainty achieved on the TFF slope is about twice as large as that of the extrapolation from CELLO data (see subsection 1.1.2), which is currently used in the hadronic light-by-light term computation for $(g-2)_\mu$, and in the computation of $\mathcal{B}(\pi^0 \to e^+e^-)$. Therefore, the better precision will not significantly improve the related uncertainties, but rather confirms the validity of the treatment of the CELLO data. In both cases, the shift of the central value can also contribute to a shift in the theoretical value towards the experimental ones, slightly reducing the tension between them.

Source	$\delta a \times 10^2$
Total statistical	0.51
Statistical - Data	0.48
Statistical - MC	0.18
Total systematics	0.25
Spectrometer momentum scale	0.16
Accidental background	0.15
Particle mis-ID	0.06
Calorimeter trigger efficiency	0.06
Spectrometer resolution	0.05
LKr non-linearity and energy scale	0.04
Beam momentum spectrum simulation	0.03
Neglected π^0_D sources	0.01

Table 3.8: Summary of the uncertainties.

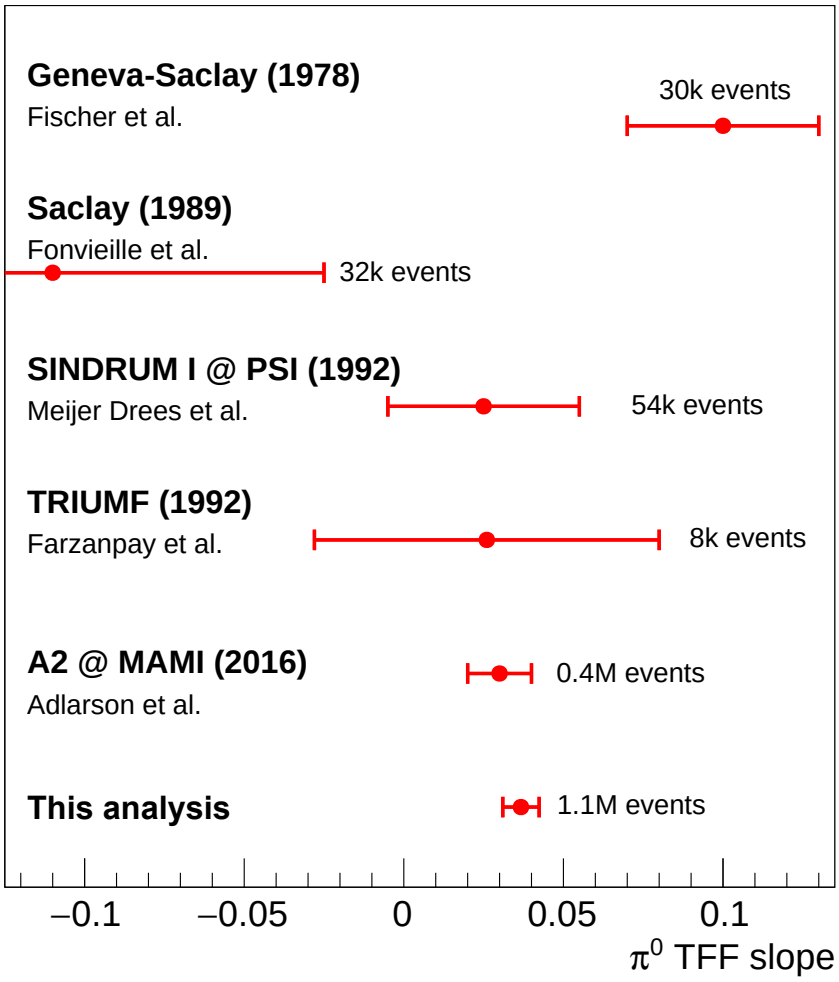

Figure 3.43: Results on the π^0 TFF slope from π^0_D decay spectrum measurements.

CHAPTER 4

THE NA62 EXPERIMENT AT CERN SPS

4.1 The experimental principle

Collection of ~ 100 $K^+ \to \pi^+\nu\bar{\nu}$ events is an experimental challenge. A branching ratio of $O(10^{-10})$ and 10 % acceptance means that $O(10^{13})$ kaon decays must occur in the fiducial volume. The desired precision of 10 % on the measurement requires the background to stay below 10 %, translating into a background rejection factor greater than 10^{12}. The signature of a signal event consists of one incoming charged track, one outgoing charged track, and nothing else. Most of the reduction of the background rate is therefore achieved with the help of the veto detectors installed as part of the experimental setup. They have been demonstrated to reach extremely high efficiencies. Any additional signal in one of these veto detectors indicates a different decay channel, featuring particles such as photons or muons.

The last factor 10 of rejection of the background events is achieved using the kinematics of the event, sketched in Figure 4.1. The only measurable observables are the momenta of the K^+ (P_K) and π^+ (P_π), and the angle $\theta_{K\pi}$ between them. They are combined into a single convenient observable, the squared missing mass, under the assumption that the

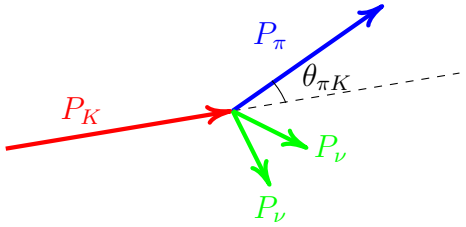

Figure 4.1: Kinematics of the $K^+ \to \pi^+\nu\bar{\nu}$ decay [5].

charged track in the final state is a pion:

$$m^2_{\text{miss}} = (E_K - E_\pi)^2 - (P_K - P_\pi)^2 = \tag{4.1}$$

$$m_K^2 + m_\pi^2 - 2E_K E_\pi - 2|P_K||P_\pi|\cos(\theta_{K\pi}) \ , \tag{4.2}$$

with $E_K = \sqrt{m_K^2 + P_K^2}$ and $E_\pi = \sqrt{m_\pi^2 + P_\pi^2}$.

Table 3.1 shows the main decay channels of the charged kaon and their branching fractions. The squared missing mass of the kinematically constrained ones, corresponding to more than 80 % of the total K^+ branching fraction, is shown in Figure 4.2a. The dominant one is the $K^+ \to \pi^+\pi^0$ two-body decay which, neglecting detector resolution effects, displays a sharp peak at $m^2_{\text{miss}} = m^2_{\pi^0}$. The $K^+ \to \pi^+\nu\bar{\nu}$ signal region is divided in two; region I below and region II above this peak. The $K^+ \to \mu^+\nu_\mu$ channel, albeit a two-body decay, does not feature a peak at $m^2_{\text{miss}} = 0$ because the muon is wrongly assigned a pion mass in the evaluation of m^2_{miss}. However, the upper kinematic limit is 0 and, neglecting detector resolution effects, it does not overlap with the signal. The $K^+ \to \pi^+\pi^+\pi^-$ channel is a three-body decay with a continuous distribution sharing part of the allowed range of the signal channel, so region II is defined below the $K_{3\pi}$ threshold. Figure 4.2b also shows the other non-kinematically constrained decays, which cannot be fully suppressed with the m^2_{miss} variable.

The NA62 experiment uses in-flight decay of high momentum (75 GeV/c) kaons. It is, however, difficult at this momentum to separate kaons from other particles in the beam (pions and protons). The KTAG Cherenkov tagging detector, upstream of the decay volume, is designed to be sensitive only to kaons. The integrated rate of particles seen

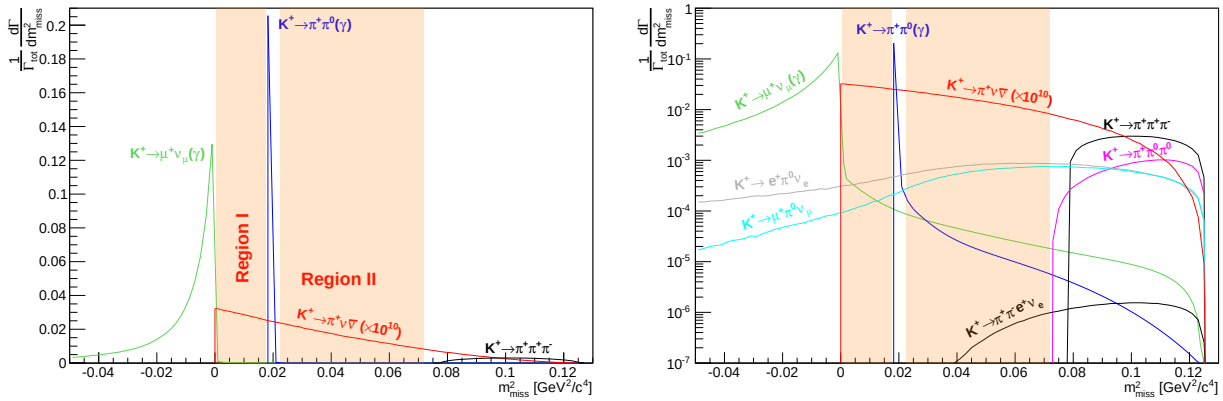

(a) (b)

Figure 4.2: The m^2_{miss} distributions for the signal and background channels [5]. The shaded orange areas are the signal regions, which are free from the dominant background sources ($K_{2\pi}, K_{\mu 2}, K_{3\pi}$). The background channels are normalised to their branching ratio, and the signal is multiplied by a factor 10^{10}. (a) Kinematically constrained background sources. (b) All background decay channels, adding the non-kinematically constrained ones with respect to (a).

by the upstream detectors goes up to 750 MHz, which is a technical challenge that led to the development of new technologies for the beam tracker. A schematic layout of the 270 m long experimental setup is shown in Figure 4.3, in which the particle identification detectors are visible (KTAG, RICH) together with the trackers (Gigatracker, STRAW) and the vetos (CHANTI, LAV, LKr, MUV, IRC/SAC, CHOD). These subsystems are described in the following sections.

4.2 Beam line

The 400 GeV/c primary beam from the SPS accelerator at CERN impinges on a beryllium target (400 mm long and 2 mm in diameter) to create secondary particles. A copper collimator with variable aperture is installed after the target. It absorbs some of the unwanted particles at angles larger than 6 mrad. The subsequent optics of the 102 m long K12 beam line selects the momentum and guides the remaining particles towards the experiment. The first element is a set of focusing quadrupole magnets, which collect

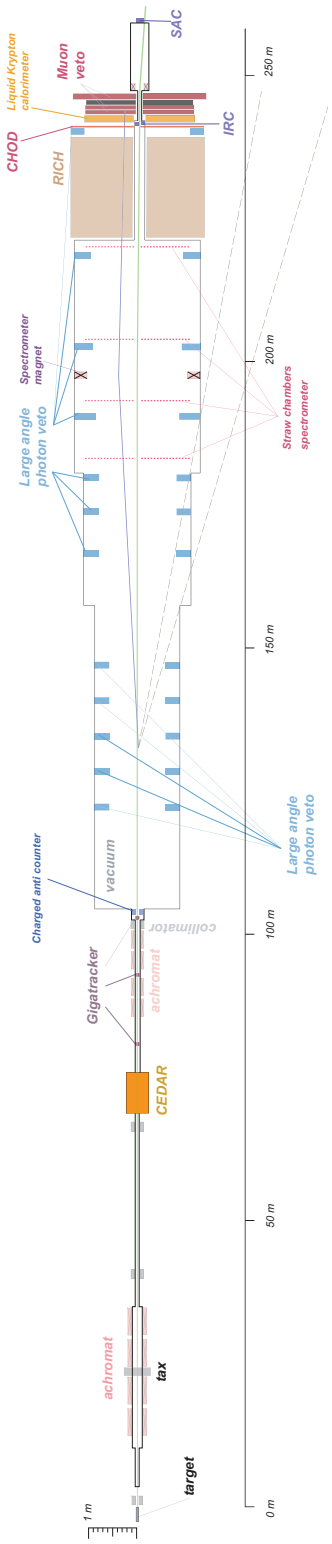

Figure 4.3: Schematic view of the NA62 experiment, showing the main sub-detectors [5].

particles with a large angular acceptance. They are followed by an achromat composed of four dipole magnets. Coupled to two collimators (TAX1 and TAX2) placed in the central part of the achromat, it can select only the 75 GeV/c component of the beam. The aperture of the collimators defines the narrow momentum band of $\Delta p/p = 1.0\,\%$ around the central value of 75 GeV/c. In between the two TAXes a radiator can be introduced in the beam to degrade the momentum of the e^+ component, which can be absorbed later. Further quadrupole magnets refocus the beam and make it parallel in both horizontal and vertical planes to pass through the CEDAR detector. A final achromat is associated with the Gigatracker detector for the measurement of the kaon momentum. The last piece of optics before entering the decay volume is the TRIM5 magnet, which is used to give a final 90 MeV horizontal kick to the beam and deflect it from the detector axis, such that the spectrometer magnet directs the beam through the central hole in the LKr calorimeter.

The advantages of the high momentum in-flight decay technique used by NA62 are multiple. The number of kaons produced per primary proton increases with the square of the kaon momentum; therefore the required kaon intensity is reached with fewer primary protons. Because the detector efficiencies are maximum beyond a few GeV of energy deposit, the detection efficiency for the kaon decay products also improves when increasing the available energy. The choice of a 75 GeV/c kaon beam optimises the number of kaon decays inside the fiducial region per proton.

The main disadvantage of this technique is the difficulty to separate the kaons from the other components of the beam. The different components are roughly 6 % K^+, 20 % protons, 70 % π^+ and the rest are μ^+ and e^+.

The beam received by NA62 is not continuous but has a specific time structure. It alternates periods with and without particles. The periods with particles are called "bursts" and are of a typical length of 3 s, with 10 s–20 s inter-burst periods. The intensity of particles varies, rising after the start of a burst, reaching a maximum during the burst and then falling towards the end of the burst. This structure comes from the SPS accelerator "slow extraction" [65]. Signals are sent to the experiment to indicate the start of burst

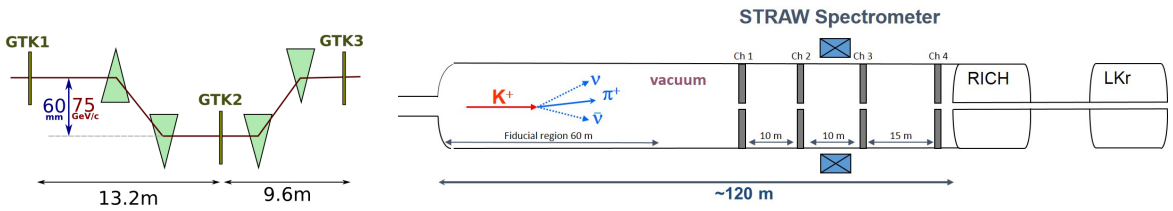

(a) (b)

Figure 4.4: (a) Schematics of the GTK layout with the three stations GTK1, GTK2, GTK3. The green triangles represent the dipole magnets. (b) Schematic of the vacuum volume showing the four straw spectrometer chambers around the magnet. It is followed by the RICH and LKr detectors, traversed by the beam pipe directly connected to the vacuum tank [66].

(SOB) and the end of burst (EOB).

4.3 Detector

4.3.1 Tracking system

The Gigatracker (GTK) is a silicon pixel detector that provides the measurement of the beam kaon momenta. It is constituted of three stations spaced by approximately 10 m, located around four dipole magnets forming an achromat system, as seen in Figure 4.4a. Each station is 200 µm thick to reduce as much as possible the amount of matter (0.22 % of X_0). The system is able to measure the kaon momentum with a resolution $\sigma_p/p = 0.2\,\%$, an angular resolution $\sigma_\theta = 16$ µrad and a time resolution better than 200 ps. This detector is placed in the beam and must be able to sustain the very high and spatially non-uniform rate of particles up to 750 MHz, with a peak rate of 1.3 MHz/mm^2 in the central area, as well as the high radiation environment. At nominal beam intensity the detectors can be exposed to a fluence corresponding to 4×10^{14} MeV n_{eq} cm^{-2} in one year of data taking.

The purpose of the Straw spectrometer in Figure 4.4b is to accurately measure the momentum of the charged secondary particles originating from the decay region. It is formed of four straw chambers of four views each, for a total of 7168 straw tubes. A

large aperture dipole magnet inserted between the second and third chambers generates a vertical magnetic field of 0.38 T, providing a transverse momentum kick of 270 MeV/c to charged particles. The measurements in each individual view are combined to obtain the 2D coordinates of the track impact point. A 12 cm diameter hole in the chambers allows the beam to pass through without leaving hits. The precision reached on the momentum is $\sigma_p/p = 0.32\,\% \oplus 0.009\,\%p/(\mathrm{GeV}/c)$ and on the angle $\sigma_\theta \leq 60\,\mu\mathrm{rad}$. As for the Gigatracker, this detector has a very low material budget (0.5 % X_0 per station) to minimize interactions.

4.3.2 Particle identification

As it is not possible to separate kaons from other particles in the beam (mainly pions and protons), the CEDAR/KTAG detector shown in Figure 4.5 is used. It is a differential Cherenkov detector placed in the beam line upstream of the Gigatracker, which allows individual kaons to be tagged. The CEDAR part is a tank filled with nitrogen or hydrogen in which the beam particles generate a Cherenkov light cone. It is essential for the CEDAR operation that the beam is parallel to the optical axis, and with a divergence smaller than 100 µrad. A diaphragm with a very thin slit allows the light cone generated by a kaon to be transferred to the KTAG, which is the sensitive part of the detector. The light is spread and directed by the optics towards 8 light boxes lined with 48 photomultipliers each to minimise the individual rate detected by each channel. The light originating from protons or pions is blocked and does not reach the KTAG. The detector has so far been operated with nitrogen, but is designed to work also with hydrogen, which minimizes Coulomb scattering, reducing the detector activity induced by beam interactions. The total thickness is 39 (5) radiation lengths for the CEDAR filled with nitrogen (hydrogen). The achieved time resolution $\sigma_t = 70$ ps enables the signals from downstream detectors to be matched to the tagged kaon.

The second particle identification detector is the RICH, shown in Figure 4.6, an 18 m long tank filled with neon, designed to separate charged pions, muons and electrons in

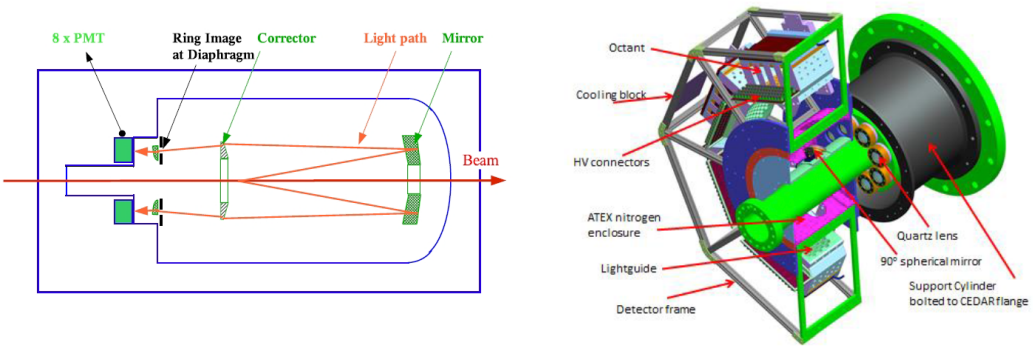

(a) (b)

Figure 4.5: (a) Sketch of the CEDAR vessel filled with the radiator gas (nitrogen or hydrogen) and optical components [66]. The Cherenkov light cone generated by the beam particles is reflected on the mirrors and passes through the elements for focussing and correcting the chromatic errors. The kaon light passes the diaphragm and enters the KTAG in (b) through the quartz lenses [67]. The light is reflected radially and spread onto eight light boxes equipped with 48 photomultiplier tubes each.

the momentum range $15\,\mathrm{GeV}/c < p < 35\,\mathrm{GeV}/c$. Secondary charged particles passing through the neon with a momentum above the threshold, $p_{\mathrm{min}} = m/\sqrt{n^2-1} \approx 90m$, where n is the refractive index of the medium and m the mass of the particle, produce a cone of Cherenkov light. The light is reflected on the 17 m focal length mirrors at the end of the tank towards the photomultipliers. Each charged particle fires a ring of photomultipliers from which the opening angle of the cone can be reconstructed. The relationship $cos\theta_c = \frac{1}{n\beta}$ links this angle to the particle velocity β. Used in combination with the momentum P reconstructed in the Straw spectrometer, the identity $P = m\beta/\sqrt{1-\beta^2}$ allows the particle to be identified by its mass.

4.3.3 Photon, muon and accidental interactions veto system

Since the signal channel is characterized by missing energy and the background decay rates are several orders of magnitude larger, it is extremely important to add veto detectors to reject events with additional particles. The first set of subsystems is designed to provide hermetic coverage for photons. It is composed of different detectors that cover angles from 50 mrad down to 0 mrad. It is especially important for the $K^+ \to \pi^+\pi^0$ decay channel

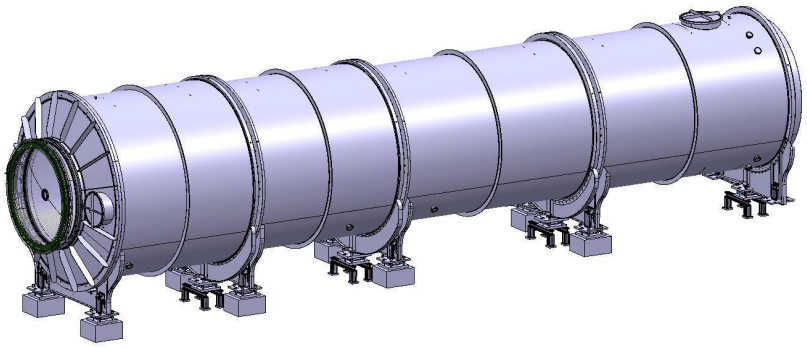

Figure 4.6: Drawing of the RICH vessel [68]. The mosaic of mirrors at the end of the vessel reflects the Cherenkov light onto two detection planes on each side of the beam pipe. One of them is visible on the entrance cap on the side of the junction with the vacuum tube.

with a branching ratio of 20 %, where the probability of missing two photons should be reduced to at least 10^{-8}.

The 12 Large Angle Veto (LAV) stations are placed all along the experimental setup, between 120 m and 240 m, covering angles between 8.5 mrad $< \theta <$ 50 mrad. The first 11 stations are part of the vacuum decay tube and are leak-free to maintain the vacuum conditions. The basic building blocks are lead glass crystals with photomultipliers, mounted in rings with 4 or 5 rings per station. The rings are staggered to cover holes left between crystals in the previous ring. A schematic view of a LAV station is shown in Figure 4.7a.

The Liquid Krypton Calorimeter (LKr), already described in subsection 2.2.3, is a key element for photon detection as it intercepts most of the secondary photons. It covers angles 1 mrad $< \theta <$ 8.5 mrad and is required to have detection inefficiencies lower than 10^{-5}. The LKr calorimeter has been used since the original NA48 experiment as a veto and to measure photon energy, but the readout has been replaced to cope with the larger rates.

The remaining angles below 1 mrad are covered by the IRC (Inner Ring Calorimeter) detector located around the beam pipe within the inner LKr radius, and the SAC (small angle calorimeter) located at the end of the experiment on the beam axis. The SAC is illuminated only by the neutral particles travelling in the beam pipe (from decay or beam particles), because the beam is deviated by a magnet before reaching it. These are

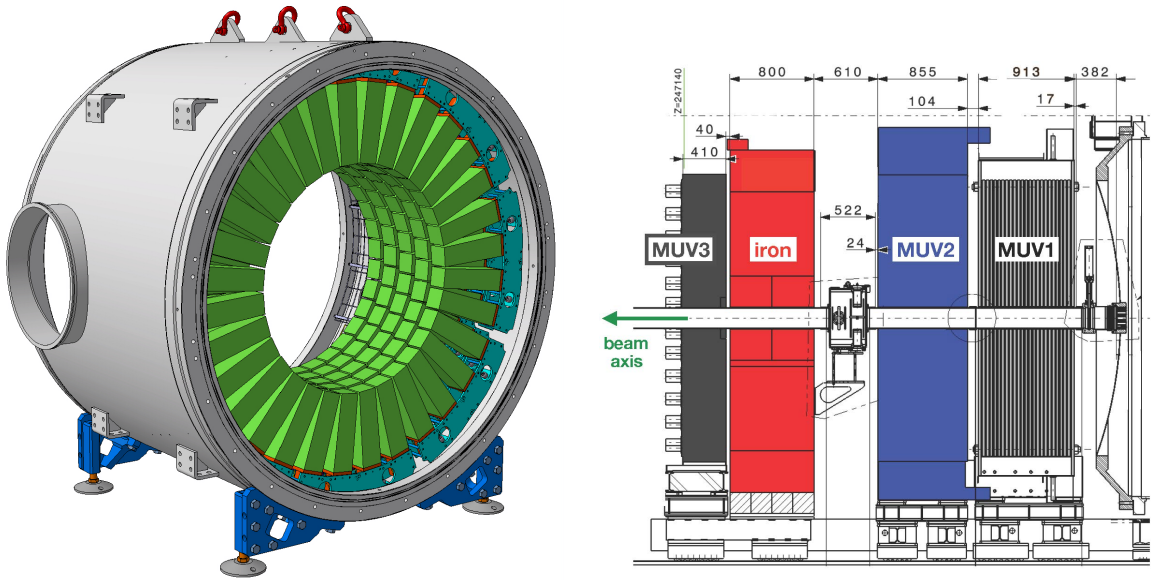

(a) (b)

Figure 4.7: (a) Drawing of a LAV station with the lead glass crystals in green and their readout in light blue [66]. (b) Drawing of the three MUV stations: the hadronic calorimeters MUV1 and MUV2, and the muon veto MUV3 after the iron wall [66].

detectors with few readout channels, but they still need to provide detection inefficiencies lower than 10^{-4}. They are sampling calorimeters made of alternating layers of lead and scintillator.

The purpose of the CHarged ANTI (CHANTI) detector is to veto events for which inelastic interactions occur in the last station of the GTK. These interactions can produce charged particles that can be detected by the spectrometer, faking a signal event. The veto is performed by detecting particles, at high angle compared to the beam at the exit of the GTK3 station, with scintillator planes.

The three MUV detectors shown in Figure 4.7b are required to achieve the necessary further 10^5 muon rejection factor compared to the pions. The MUV1 and MUV2 are hadronic calorimeters located downstream of the LKr calorimeter. They are made of sandwiches of iron and scintillator strips with alternating orientation between layers. The third station is placed after an 80 cm thick iron wall and serves as a fast muon veto. The MUV3 is made of tiles with direct optical connection to the photomultipliers, providing a time resolution of 250 ps.

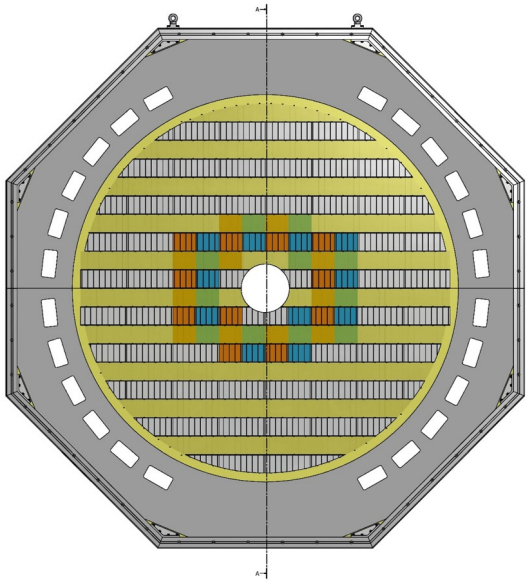

Figure 4.8: Schematic view of the NewCHOD [66]. The short tiles are placed in the central part around the beam pipe (orange and blue), surrounded by the large ones (twice the width of the short ones). The rows are located alternately on the opposite sides of a carbon fibre plate, overlapping vertically by 1 mm.

Finally the CHarged HODoscopes (CHOD and NewCHOD) are used in the trigger and as vetos for nuclear interactions in the RICH elements. The CHOD has already been described (referred to as HOD in subsection 2.2.2) and has not been modified with respect to the previous setup. Only the front-end readout has been replaced to cope with the higher expected rates. The NewCHOD (Figure 4.8), currently operating in parallel with the CHOD, is made of a plane of scintillator tiles individually read out by two silicon photomultiplier channels used in coincidence, increasing the granularity of detection and optimising the rates for each channel. A total of 152 tiles of 30 mm thickness, 108 mm height and a width of 134 mm (around the beam pipe) or 268 mm (on the outer part) are arranged into four quadrants.

Furthermore all these veto detectors provide an accurate timing, with better than 1 ns resolution, as they need to be able to pinpoint single events over the length of the detector despite the high rate of beam background.

4.4 Trigger and data acquisition

The Trigger and Data Acquisition (TDAQ) system is an essential part of the experiment. It refers to the equipment involved in the chain of processes handling the data, from the collection of the analog signal coming from the detectors to the writing of complete events on disk.

The data taking process starts with the collection of the analog signal from the detector channels. The signal is processed by the acquisition board to transform it into digital signal. The TEL62 board [69] is derived from the LHCb TELL1 board [70] and was designed to be the common standardized readout board for most of the subsystems. The detector channels are clustered into a few TEL62 boards that can host up to 4 TDCB [71] daughter cards able to receive a maximum of 128 channels each. A simple pattern recognition algorithm runs on the data to classify the event into one of the predefined interesting detection patterns. When a positive match is found, indicating the detection of a particle, a trigger primitive is generated. The primitive consists of a 25 ns time granularity timestamp accompanied by a fine time of 1/256 of the clock period (roughly 100 ps). The data are written in windows of 25 ns into an internal buffer memory.

The rate of events seen in the NA62 detector at nominal intensity is above 10 MHz. The computing infrastructure cannot cope with such a high rate and it must be reduced. The purpose of the TDAQ is to discriminate quickly between potentially interesting events and uninteresting ones, rejecting the latter to reduce the rate to a sustainable level. The trigger is a real-time system and the decision must be taken within a limited time. This is done in three steps, increasing the complexity with each level:

- A hardware level 0 (L0) trigger using only a limited amount of information from a few subdetectors.
- A software level 1 (L1) trigger using information from each subdetector independently.
- A software level 2 (L2) trigger running more complex algorithms correlating infor-

mation from multiple subsystems.

The L0 is a hardware trigger implemented on a Field-Programmable Gate Array (FPGA) and intended to reduce the rate from 10 MHz to 1 MHz. The Level 0 Trigger Processor (L0TP) is the board dedicated to this task. The time allocated for generating and dispatching the trigger is 1 ms. The decision process is based on the time-matching of the primitives received from the acquisition boards. When one of the interesting combinations of primitives is detected, a positive decision is taken. The L0TP generates a trigger containing the timestamp of the event and broadcasts it to all acquisition boards. Upon its reception, the boards extract the data from the internal memory. To be certain to have all the information corresponding to this event, a defined number of 25 ns time windows around the trigger timestamp are extracted and sent to the next level of trigger. The detectors included in the L0 are the LAV, spectrometer, RICH, (New)CHOD, LKr and MUV3.

The Local Trigger Unit (LTU) is installed in the cavern, one for each subsystem. It is responsible for the distribution of the main 40 MHz clock that defines the time, and the SPS clock synchronizing with the start- and end-of-burst. The clock signals are distributed through optical fibre. When a subsystem is included in the global data taking of the experiment it also relays the L0 triggers from the L0TP to the acquisition boards. When a subsystem is in local mode, it can emulate triggers.

The level 1 (L1) and subsequent trigger levels are software triggers implemented in the PC farm. The PC receiving the detector data sent after a positive L0 trigger is picked from a round-robin over all the available PCs. This round-robin is reset and synchronized at the start of burst. To reduce the rate to 100 kHz, the L1 trigger processes every event in a more sophisticated way than the L0 and must be fully executed within 1 s. All the detectors are still considered independently and no cross-detector analysis is performed. In the case of a positive decision, a data request is broadcasted to the LKr and GTK readout. After receiving the L1 detector data, the L2 software takes a decision based on the correlation of data from multiple sub-detectors. The L2 should reduce the rate to

around 10 kHz–15 kHz. This final stage of processing can extend into the inter-burst period but must be finished before the new burst starts. Once a positive trigger is acknowledged, the event is sent to the merger PCs, which assemble all the events of the burst and write them on disk.

There are two special cases in this standard trigger process. The LKr detector with its 13248 channels generates a sizable rate of raw data (800 GB/s), which cannot be processed entirely and used at L0. The digitized data coming from the readout board of the LKr are cascade-processed through three layers of TEL62 boards. The final layer generates the trigger primitive that is sent to the L0TP. During the processing of the L0 and L1 triggers, the data are stored in the front-end electronics. They are sent to the PC farm only after a positive L1 trigger decision, thus reducing the data transfer rate to about 100 kHz and making the full LKr information only available starting from L2. The second mechanism used to reduce the LKr data bandwidth makes use of the fact that most channels are empty in an event. This "zero-suppression" mechanism transmits data only for channels where a sizeable signal is recorded. The GTK has a triggerless architecture: its channels are continuously read out and the data are sent to the off-detector readout board. This detector does not generate primitives and does not participate in the L0. Similarly to the LKr, the data are stored until an L1 trigger is received.

Different trigger chains can be implemented at every trigger level. At L0, the main $K^+ \to \pi^+\nu\bar{\nu}$ channel is triggered when a single track is found in (New)CHOD, no more than 1 cluster is found in LKr and no activity is detected in MUV3. In addition, control triggers are required to collect other samples. These samples will help in understanding the detectors and the data collected, checking the agreement of Monte Carlo simulation and understanding the efficiencies of the trigger chains by comparing well-known normalisation channels. These trigger chains correspond to different combinations of conditions in data, excluding event topologies corresponding to background. Several trigger chains are run in parallel, being limited by the maximum trigger rate and bandwidth. In order not to exceed the limit, the chains can be downscaled, rejecting a defined fraction of the events

that would otherwise pass the trigger.

CHAPTER 5

THE NA62 RUN CONTROL SYSTEM

The previous chapter shows that the TDAQ of NA62 is a complex system involving a wide range of equipment, developed by a number of teams. Most of the shift-takers, who have to run the experiment, do not know the details of its operation. It is therefore necessary to organize this system into a coherent and transparent structure that can easily be controlled as a whole.

The NA62 Run Control is responsible for the configuration, control and monitoring of all equipment and processes linked to the TDAQ. It provides a hierarchical and homogeneous control of all the subsystems at different levels of granularity, from the experiment as a whole down to a single board. It manipulates abstract objects representing different parts of the detector. This abstraction enables the standardisation of the operation of the otherwise very different equipment. The abstract object communicates with the real hardware through the network and a standard interface. Each device is then responsible for decoding the standard instructions according to its own specificities. Similarly the reported internal state of the device is mapped onto the abstract object maintained by the Run Control.

This chapter details how the Run Control internally represents and communicates with the TDAQ system. The architecture and technologies used to achieve these goals are also described.

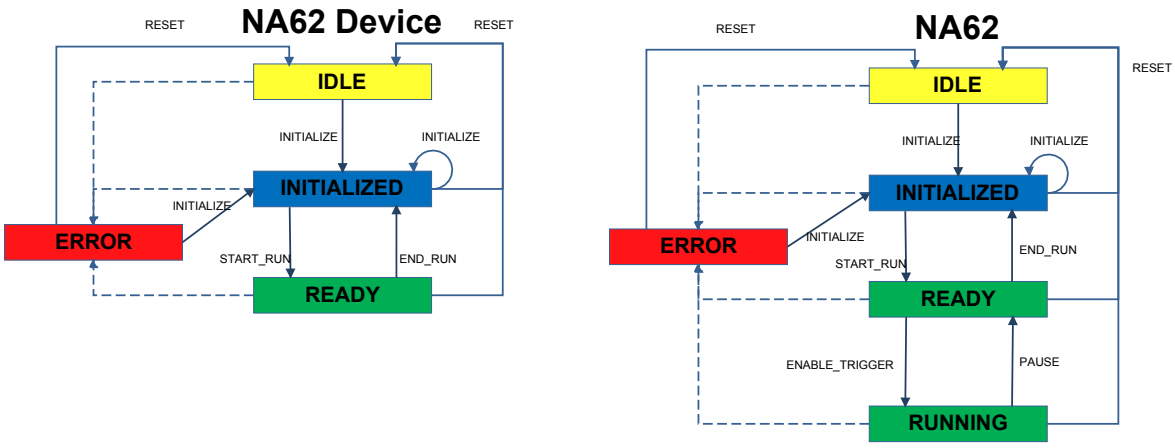

(a) (b)

Figure 5.1: Main Finite State Machine diagrams of the NA62 Run Control. (a) Standard FSM diagram for the devices. It is followed by the TEL62, LKr, Straw, IRC/SAC and GTK control softwares. (b) FSM diagram of the root node representing the global state of the NA62 Trigger and Data Acquisition (TDAQ).

5.1 Modelling of the TDAQ in the Run Control

The abstract objects manipulated by the Run Control are Finite State Machines (FSMs). An FSM is defined by a set of states and transitions. The machine at any point has to be in one and only one of the defined states, referred to as the current state. A state is defined by a set of variables. A transition to another state can be triggered by external events such as commands, or by a change in the result of the evaluation of the variables through logical conditions. This concept is well suited for electronic devices as the electronic components can only be in a limited number of logical states, and these can easily be mapped onto the abstract FSM.

Each controlled device (board, PC, software) is modelled as an FSM. The states are computed according to the evaluation of logical expressions depending on a set of parameters provided by the device. The transitions between states can be spontaneous, following a change of one of these values, or it can be triggered by one of the defined commands. Figure 5.1 shows the standard FSM used in the Run Control. The accessible states are:

- **IDLE**: Initial state after starting or resetting the FSM and the devices.

- **INITIALIZED**: When all the devices have been configured and the TDAQ is ready to start data-taking.
- **READY**: All the devices are in their data-taking state but waiting for triggers. Triggers are not yet generated by the trigger processor.
- **RUNNING**: The trigger processor is out of the paused state and generates triggers.
- **ERROR**: This state can be reached from any other state whenever a problem occurs.

Each state allows a list of abstract commands:

- **LOAD_RECIPE**: Apply a specific configuration from the configuration database.
- **INITIALIZE**: Request initialisation of all the devices.
- **START_RUN**: Request all the devices to start the run and move into a state where they are able to acquire data.
- **ENABLE_TRIGGER**: Request the trigger processor to start generating triggers at the next burst.
- **PAUSE**: Request the trigger processor to stop generating triggers at the next burst.
- **END_RUN**: Request all devices to stop taking data and end the current run.
- **RESET**: Immediately move to an idle/initialized state, stopping the current run if it was ongoing.

Due to their specificities some devices follow a different FSM (or a modified version of the standard one). The L0TP alone defines the difference between the **READY** and **RUNNING** states and therefore follows the same FSM as the global NA62 TDAQ in Figure 5.1b.

The PC farm nodes are composed of two elements: one part is hardware, the other is software. The hardware monitoring is done using the Farm Monitoring and Control

(FMC) component (see subsection 5.3.1). It monitors the power state, memory, cpu[1], file-system, network and processes of the node. These monitored values are assembled in an FSM to assess whether the node is in **ON**, **OFF** or **ERROR** state. A second component monitors the state of the NA62-specific farm software. This piece of software has been developed to process the data coming after the L0 trigger and to generate the L1 and L2 triggers. The FSM for this software part is also a simple **ON**, **OFF** or **ERROR** state. The two FSM states are assembled to decide if the node is ready to accept data (both parts are in the **ON** state) or not.

The RunManager is a special device node that has been designed to manage the run information and ensure the coherence of the configuration. It is a software device that runs within the Run Control. Its state machine contains only the **INITIALIZED**, **READY** and **ERROR** states. Its role is to load configurations from the database and to request information from the operator when starting and ending a run. More details on how the RunManager functions are given later.

Describing and manipulating every device as an FSM homogenises their control. The hierarchical supervision is achieved by linking each FSM in a tree-like structure. A partial outline of this structure is shown in Figure 5.2. The tree is formed of hierarchically-linked nodes. The external nodes (with no child) are called leaf nodes. The internal nodes (with at least one child) are called branch nodes. The root node is the top-most node, having no parent.

The device FSMs described before are all leaf nodes (also called device nodes). The internal nodes (also called logical nodes) group the devices into logical entities representing subsystems of the experiment. These nodes are modelled as the FSM in Figure 5.1b. Their states are computed according to the logical evaluation of the states of their child FSMs (e.g. "If all children are in **READY** state, move to **READY**"). Thus their role is to summarize the state of all elements belonging to this logical grouping. The root of the tree is the logical node that represents the global state of the TDAQ.

[1]central processing unit

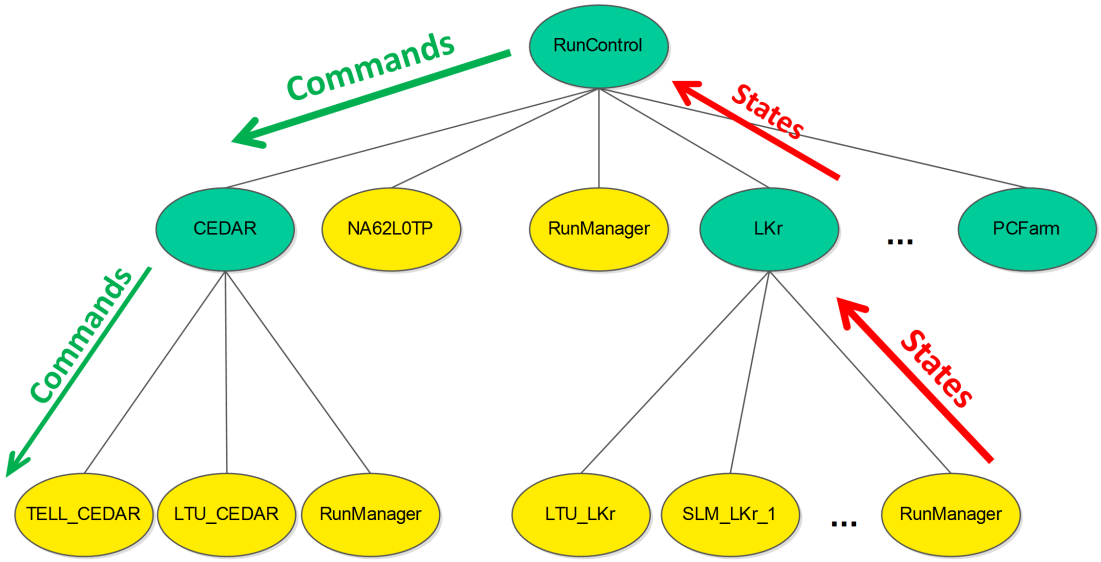

Figure 5.2: Tree-like hierarchy of units. Green units are logical units and yellow are device units. Commands are propagated downwards and states are propagated upwards and summarized in the logical units.

The tree transfers information between parent and child nodes. State parameters are received and updated from the device nodes and the FSM state is updated accordingly. The state is then propagated upward from child to parent, every logical node updating its own state until the root node is reached. In the opposite direction, a command can be issued at any node (device or logical). The command is propagated downwards from parent to children. A command is only accepted by a child if it is allowed to do so, which is determined by the state it is in. The command either continues its way down the tree in the case of a logical node, or the actions defined for the command are executed in the case of a device node. These actions depend on the type of device unit and are described later. This process of state/command transmission is sketched in Figure 5.2.

The logical node and all its attached device children form what is called a domain. Domains can always be detached from the main tree and become an independent subtree. When a subtree is detached, the logical node that has been cut becomes the root node of the subtree. States are no longer propagated upwards to the main tree and commands issued in the main tree are no longer transferred to the subtree. This mechanism permits

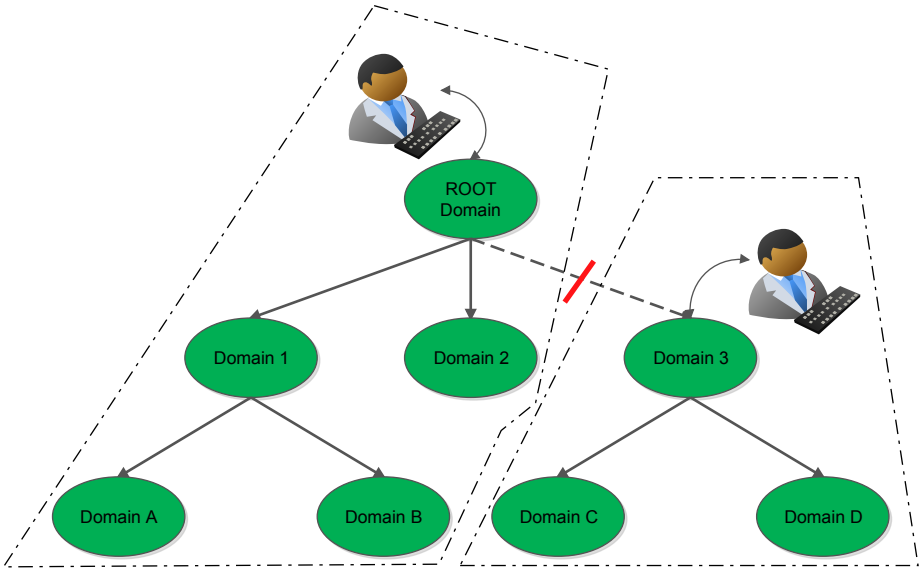

Figure 5.3: A domain and all its subdomains can be taken out of the main tree. The connections between the main tree and the subtree are cut and the states and commands are not propagated any more. A different user can take control of the new independent tree.

multiple users to concurrently control parts of the experiment without interfering with each other. This is illustrated in Figure 5.3.

In this abstraction the RunManager nodes have an important role to play. Every domain contains a RunManager (i.e. every logical node has a RunManager node as a child). When a command reaches a RunManager, inputs are requested from the user. However because the user does not want to input the information for each domain, and to ensure coherence between domains, a single manager acts as a master. The master is the RunManager belonging to the current root node. The commands available to the RunManager are the following:

- **LOAD_RECIPE**: The master RunManager requests the user to select one of the predefined configurations. The corresponding configuration is loaded from the configuration database and cached, to be later applied individually by the devices.

- **INITIALIZE**: The slave RunManager checks the state of the master. If the master is in **INITIALIZED** state, it waits for a notification that the initialisation

is finished and then copies the information from the master. If the master is in **READY** state the run has already been started and the configuration is not allowed to be changed. It is therefore assumed that the master already contains the correct configuration and it is immediately copied.

- **START_RUN/END_RUN**: The master RunManager requests the user to enter information about the run. The slave waits for the notification from the master and then copies the information.

Every standard device can receive one of the commands defined by the FSM (Figure 5.1a). When necessary, the device requests the current configuration from the local RunManager. The call returns directly when the RunManager is in **READY** state as the configuration is assumed to be correct. Otherwise the call is blocked until the RunManager receives a configuration type, either from the user or by copying the master RunManager. The configuration is loaded from the database during the **LOAD_RECIPE** command. Some device-specific processing is executed and the command is transferred to the device through the network, where it is decrypted and understood by the local control software, which executes the relevant actions. This is depicted in Figure 5.4.

5.2 Standard interface

Because of the diversity of equipment to be configured it would be difficult to know the details of the possible configurations for each of them individually and to program them in the Run Control. Instead a flexible XML configuration file scheme has been developed and a procedure has been devised to transmit these files and make sure they are correctly loaded.

The configuration files can be edited and each version is stored in a database. During configuration the file is transmitted to the device according to the procedure described in this section. After reception, the file is reverse-processed and loaded into the device memory by the local control software. Immediately after a start or an end of run, the

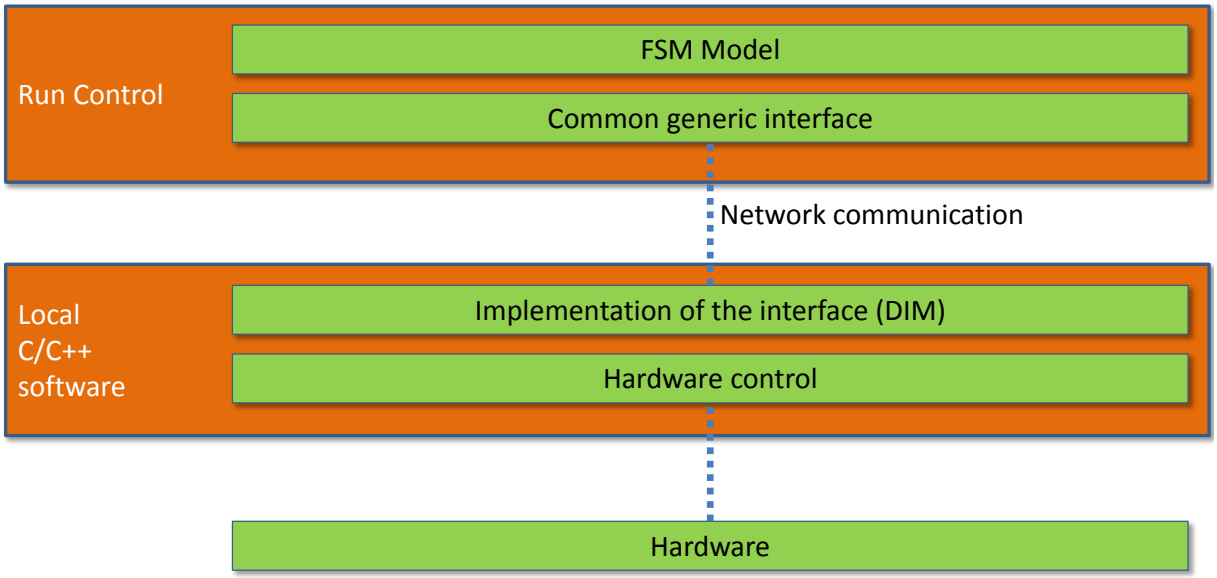

Figure 5.4: The commands to a device are formed by the Run Control and sent through DIM [72] to the local control software that decodes it and executes the corresponding sequence of actions.

actual configuration is requested from the device to verify that the configuration has been properly loaded, and for bookkeeping purposes.

The configuration procedure relies on a series of standard commands and services that need to be implemented on the interface by each device's local control software. It works on a client-server model where the same instance can be both server and client. In a client-server model, a server can publish a list of services or values, to which any number of clients can subscribe. On the other hand, a client can subscribe to any number of servers to receive or request information. The server part implements command and service ports. The clients can push commands/requests to the device via the command ports (input). The minimum set of commands to be implemented to interface with the Run Control is the following:

CMD.1 *Command*: To transmit the commands described in subsection 5.2.1.

CMD.2 *FileContent*: To transmit the configuration files according to the procedure described in subsection 5.2.2.

CMD.3 *RequestConfig*: To request the device to report back its current configuration to the

Run Control.

The service ports provide information to the clients connected to the device (output). With these ports, the Run Control is able to determine the exact state of the hardware, transmit information to the user and log it in the database. The minimum set of services to be implemented is the following:

SVC.1 *State*: Report the FSM state. For the devices following the standard FSM there should be a one-to-one correspondence with the FSM diagram in Figure 5.1a. The standard convention is shown in Table 5.1. The use of any other value for the ERROR state allows different types of errors to be defined, identified by the error code.

SVC.2 *Info*: Transmit output to the user. This service should be the equivalent of the terminal output when the control software runs in command line mode. It is only used to inform the user currently working on this specific equipment and is not recorded.

SVC.3 *Logging*: Report the information to be logged by the mechanism as described in subsection 5.2.4. Contrary to the previous service, this one should only provide important and summarized information that is stored in the offline database to keep track, for future records, of important problems that happened during the run. This service provides an identifier for the detector generating the message, a severity code and the message itself.

SVC.4 *Config*: Report back to the Run Control the current values of the configuration parameters. See subsection 5.2.2 for details.

The Run Control itself provides several services that can be used by any devices:

- *InRun*: Indicates whether a run is currently ongoing.
- *RunNumber*: Run number of the current run. Updated after the **START_RUN** command.

Table 5.1: Mapping between value of service **SVC.1** and Run Control device FSM state.

FSM State	Value
IDLE	0
INITIALIZED	1
READY	2
ERROR	Other (-99 reserved)

- *BurstNumber*: Number of bursts in the current run. Updated at each new burst.
- *PCFarmOptions*: List of options to be used by the PC farm software.
- *EnabledDetectors*: List of enabled detectors for the current run and the number of boards producing data for each detector.
- *L1EnabledDetectors*: List of enabled detectors to be read out at L1 for the current run and the number of boards producing data for each of them.
- *CREAMCrates*: List of active CREAM boards (calorimeter readout) able to send data to the PC farm.
- *EnabledPCNode*: List of enabled PC farm worker nodes that should be used to receive data from DAQ boards.
- *EnabledMerger*: List of enabled PC mergers that should be used to merge bursts.
- *RunningPCNode*: List of active PC farm worker nodes able to receive data from DAQ boards.
- *RunningMerger*: List of active PC mergers able to merge events.
- *L1TriggersNumber*: Number of L1 triggers during the previous burst.
- *L2TriggersNumber*: Number of L2 triggers during the previous burst.
- *BEAMINFO*: Block of data containing information about the beam. This block is inserted in the data written onto disk.

5.2.1 Run Control commands

The Run Control does not directly control the devices. Instead it manipulates the FSM objects representing them, then the FSM objects generate and transmit commands to the device. As the Run Control is not aware of the internal operation of any device, the commands are very generic and the device is expected to understand them and execute the appropriate sequence of actions specific to itself. After the execution of the associated action the devices should report back to the Run Control, notifying the success or failure of the action by updating/refreshing its state variables on service **SVC.1**.

The commands are sent to the command port **CMD.1** and can include any number of parameters. The minimum set of commands to be understood and implemented is the following:

- initialize [parameters]
- startrun [parameters]
- endrun [parameters]
- resetstate [parameters]

After receiving the *initialize* or *startrun* command, if a configuration file is needed, the device should be ready to receive a single file according to the procedure described in subsection 5.2.2. For the commands followed by a configuration file, the state should be updated/refreshed only after the configuration file has been applied. If no configuration file is expected, the command can be executed immediately and anything received on the *FileContent* command port **CMD.2** can be discarded.

5.2.2 Configuration

Every single device in the experiment requires a different configuration, which is likely to change over time. The configuration mechanism of the Run Control takes advantage of the existing recipe mechanism of the JCOP [73] framework (see subsection 5.3.1). In order

to hide the internal details of the devices from the Run Control, the actual configuration parameters/values are contained in a configuration file. The content of the file is written into the JCOP database, which contains a list of recipes describing the configuration modes (e.g. calibration run, muon run, physics run, ...). Each recipe contains the set of configuration files necessary to all the devices in the experiment. When loading a recipe, the file is transmitted to the device, which is again responsible for decoding it and applying the values.

The file transfer mechanism is:

1. After sending one of the *initialize* or *startrun* commands described in subsection 5.2.1, the Run Control extracts the content of the corresponding configuration file from the local cache (previously loaded from the database) and writes it to a common network disk space available to all devices. The path to the file is sent in the command slot **CMD.2**.

2. Upon receipt of the file path, the device reads it and starts processing it (the details are left to the specific device implementation).

3. When the processing of the configuration file is over and the values have been applied, the device updates its state parameters **SVC.1**, indicating the end of the configuration procedure.

At any point in the procedure, if an error occurs or an anomaly is detected by the device, the system moves to an **ERROR** state that interrupts the procedure.

4. To maintain the consistency of the conditions database, the device can be requested to report back the current configuration by sending a path to the command port **CMD.3**.

5. The device has to generate a file containing all the current parameter values in the same format as the configuration file. The file is transmitted to the Run Control by writing it to the common network disk space at the path provided by the Run

Control and it notifies the Run Control that the operation has been completed by sending back the path to the file in the service port **SVC.4**. The Run Control takes care of storing it in the offline database. Two possibilities are given to the subsystem:

- When applying the configuration files, keep track of the real value of each parameter (the ones that have been applied) and report this list of values, assuming that these values were effectively correctly loaded into the hardware.
- Request the hardware to report the actual value of each parameter and report this list of true values.

5.2.3 XML configuration

To help in generating, maintaining and applying the XML configuration files within the C/C++ control software, the XMLPreProcessor tool was created. The aim is to map the elements of the XML file onto an internal data structure in the configuration software. The configuration is rarely complete: as an example the TEL62 boards are generic and not all parameters are relevant to all subsystems; some parameters need to be configured only during initialization, others only at the start of run, and others can be freely moved from one stage to another. Nevertheless, the internal data structure is common and should contain all the available parameters. This led to a specific requirement of the tool, which was to give the possibility to apply partial configuration files, ruling out the simple option of serialization of the structure into the XML file. Since several completely different subsystems are using it, the tool must also be able to act transparently on any configuration structure.

Since a C/C++ program cannot dynamically query itself to understand which variables are defined, a two-part program was written. The first is an XML parser, based on libxml2 [74], able to parse and navigate the XML structure. It executes both the read operation from the XML file into the data structure, and the write operation from the

Figure 5.5: Illustration of the XMLPreProcessor program. The data structure representing all the possible configuration variables is defined in *struct.h*. The pre-processor reads and understands it, then generates the code in *structProxy.cc* and *structProxy.h*, using the routines provided by the XML parser. During run-time, complete (*struct.xml*) or partial (*partial_struct.xml*) XML files can be loaded into the structure through the generated code. This generated code is also used to report the board configuration by creating a complete XML file from the values found in the data structure.

data structure into the XML file. It provides a set of routines in both C and C++ to to set/extract values based on the full hierarchical name of the parameter (e.g. "exampleStruct.my_substruct.my_integer"). It is also possible to retrieve the list of parameters modified by the configuration file. The second element of this program is a pre-processor, which parses and understands the configuration data structure and generates the necessary code, making use of the routines provided by the XML parser. This generated code maps the content of the XML structure onto the data structure and vice-versa, ignoring parameters that are not common to both sides. Figure 5.5 illustrates the result of the XMLPreProcessor tool.

5.2.4 Logging

Logging of events and errors is an important part of the data-taking process. During later analysis, it allows users to be aware of any problems that happened and to discard some data that was affected by a problem. It also allows users to get concrete and detailed feedback on the devices generating the logs for future work and improvements.

The Run Control records all the relevant information it generates and collects it into two databases: the online database, storing information related to the run and the instantaneous state of the data taking and subsystems; and the offline database, which contains the subset of the online values that are relevant to future analysis and the configuration for each run.

In addition to run value logging, there is a centralized device logging. This logging is displayed in the control room and recorded in the database. Every device implements the service port **SVC.3** providing the source identification, a severity code and the text message itself. Four severity levels are allowed:

- 0 for **Info**: an information that the user might want to know but no action is needed.
- 1 for **Warning**: an event has occurred and the user might want to take action even if it is not mandatory.
- 2 for **Error**: an event has occurred and the user should take action.
- 3 for **Fatal**: a fatal error has occurred and the run should be aborted.

5.3 Technologies

5.3.1 Development framework and tools

The Run Control is built using the WinCC Open Architecture[1] (WinCC OA) SCADA[2] software. WinCC OA provides the basic tools to build the Run Control: graphical user interfaces, scripting language, distribution of the system over multiple machines. On top of this, the JCOP [73] and UNICOS [75] frameworks developed at CERN provide an extensive range of components developed for the needs of the LHC experiments. The Run Control therefore relies on well-proven technologies centrally developed and maintained by CERN.

Many of these components are currently in use in the Run Control; the look and feel of the application is similar to that of the LHC experiments. The main window of the graphical user interface shown in Figure 5.6 is a typical example. It displays: the detector FSM tree described previously on the left; the menus giving access to the framework options are on top; while the right side is dedicated to panels specific to the currently selected tree node. This interface allows the whole detector to be controlled, and a run to be started, in only a few mouse clicks, even for those who do not know the details of the TDAQ system. A colour code is used throughout the whole user interface to quickly identify which system is not performing well by intuitively following the red nodes in the FSM hierarchy.

Other windows are tailored for NA62, such as the "BigScreen" window (Figure 5.7) displayed in the control room. It constitutes the main reporting and diagnostic tool for operators to understand the status of the run. The central picture displays a schematic view of the detector and the state of all subsystems with coloured icons. The "Trigger Flow" box on top shows a summary of the numbers of events received at the different trigger levels. The numbers within columns should be similar, while the decision of the

[1] Previously known as PVSS

[2] Supervisory Control And Data Acquisition (SCADA) are industrial control systems designed to control and monitor large scale infrastructure

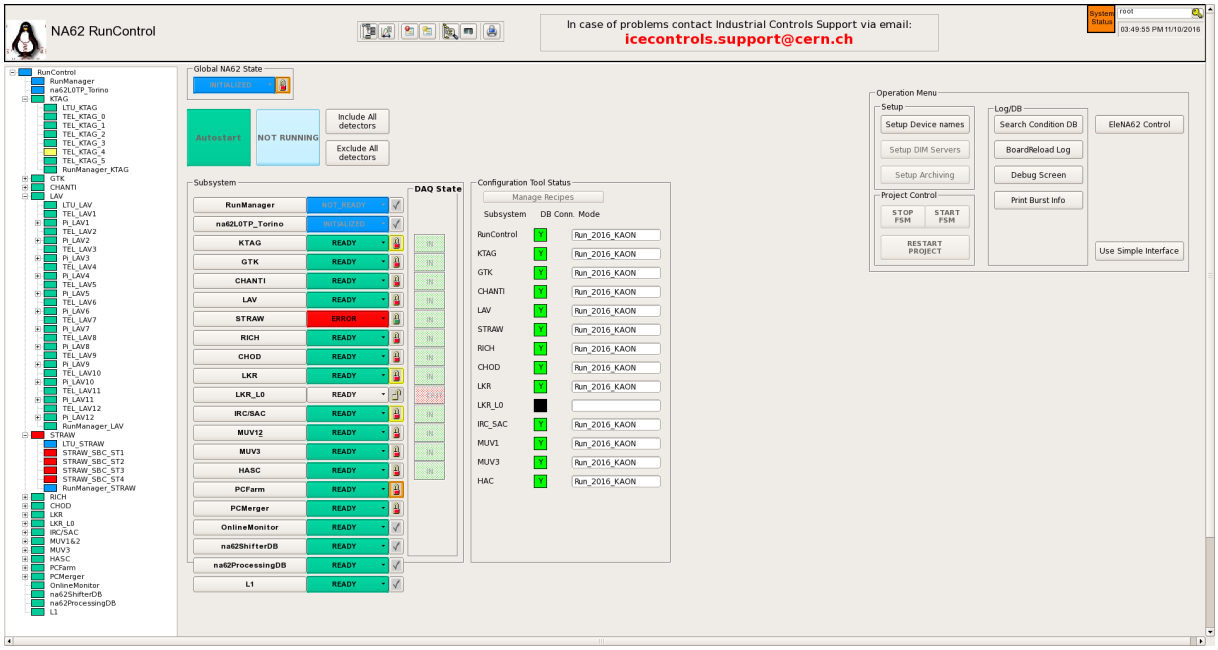

Figure 5.6: Main window of the Run Control graphical user interface. On the left is the detector tree, and on top the menus giving access to the framework options. The right side is dedicated to panels specific to the currently selected tree node.

trigger algorithms reduces the number of events from one column to the next one. The numbers are automatically compared at the end of each burst and the "flow" is coloured in red for stages where events are missing. The "PCFarm" block shows the number of data packets received by the PC farm from each detector. Once a problem is detected in the trigger flow, it allows quick identification of which detector is responsible for the missing events. The detectors displayed in yellow in the table send data after a positive L1. The "Beam infos", "Exp. Scalers" and "Primitives Count" boxes display information about the raw numbers of particles seen in the detector. The first two boxes show the total number of particles in the beam at different points along the experiment. The "primitives count" corresponds to the number of primitive triggers sent by the L0 detectors, and the number of L0 triggers generated for each trigger stream. The merger box shows the burst being processed by the merger PCs and the number of events written onto disk. The disk usage is also closely monitored to avoid losing data due to too small a disk space being available. Finally, in the top right corner, the "Run Info" block displays the metadata information related to the run itself.

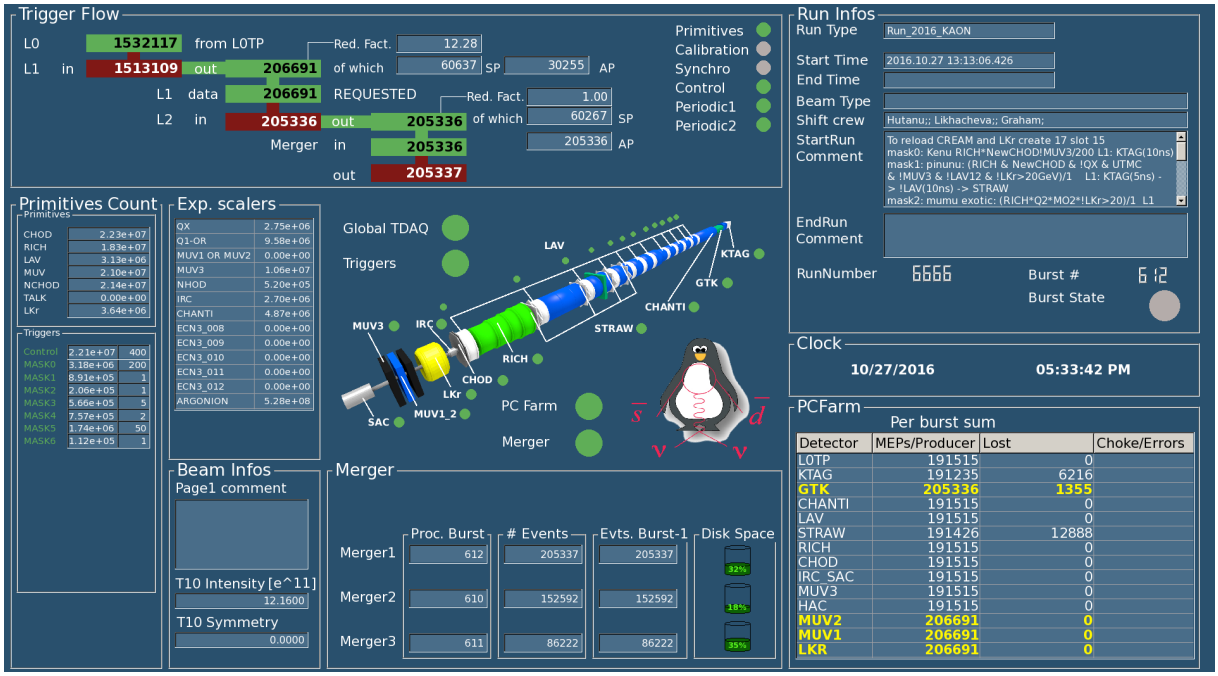

Figure 5.7: "Big Screen" of the Run Control displayed on a screen on the wall of the control room. It is the main tool reporting the status of the run, and allows quick detection of a problem in the data taking process.

All communications between the Run Control and any device are mediated through DIM [72] (Distribution Information Management), providing a platform-independent and network-transparent inter-process communication layer. C/C++ software encapsulating all the knowledge about the device is written by system experts and implements the DIM layer, complying with the interface detailed in the previous section.

The FSM toolkit allows the different FSMs of the Run Control to be built by defining the different states, the transitions rules and the actions associated with each command. It is also the tool that permits these FSMs to be arranged in a hierarchical multiple level tree. This tool then manages each FSM and the interactions between them with SMI++ [76] (state management interface) processes and an underlying DIM communication. The FSM toolkit can be used to automatize recovery procedures and its partitioning capabilities enable control of every part of the detector in a concurrent way for standalone operation.

The Farm Monitoring and Control (FMC) component is divided into two parts: a set of processes running on every worker node of the PC farm and a set of panels and WinCC OA processes running in the Run Control. Their interaction allows the PC farm to be

fully monitored and controlled from a single software application. The FMC component includes: CPU, memory, network and file-system monitoring; a process manager from which processes can be automatically (re)started and monitored for failures; and a system logger. The power management interface of nodes is also included in this scheme so that the Run Control is able to operate the power and report any hardware failure. This component is designed to be deployed on a large computer farm. For this purpose it relies on a master-slave model. A few PC nodes are defined as masters and each is allocated a list of slave nodes, which run a small program that can receive commands from the master. The master runs a process manager that monitors all its slave nodes. Its configuration controls which processes should be running on the slaves and when the monitoring software detects that one of them is not running, a command is sent to restart it. The FMC component also allows the configuration of each node to be modified.

Finally, keeping records in the database is simplified by the built-in ability to communicate with Oracle databases. One database is used to store and version all the different configurations of the experiment (recipes) and the frameworks help with storing, retrieving and applying these configurations. It also helps with storing important values in the conditions database, in such a way that it is always possible to know exactly what the data taking conditions were during future analysis.

5.3.2 Architecture

WinCC OA works as a series of managers, where each has a dedicated purpose: maintaining the database, displaying the graphical interface, centralising and dispatching information between managers, communication with third-party components, etc. The managers communicate with the TCP/IP protocol, making it easy to distribute the system over different computers on a network. In this way the managers can also be made redundant: when a host running one of the managers dies, the manager can easily be restarted on a different computer and so resume its operation transparently.

The Run Control is spread across different machines in the experiment [3, 4], as shown

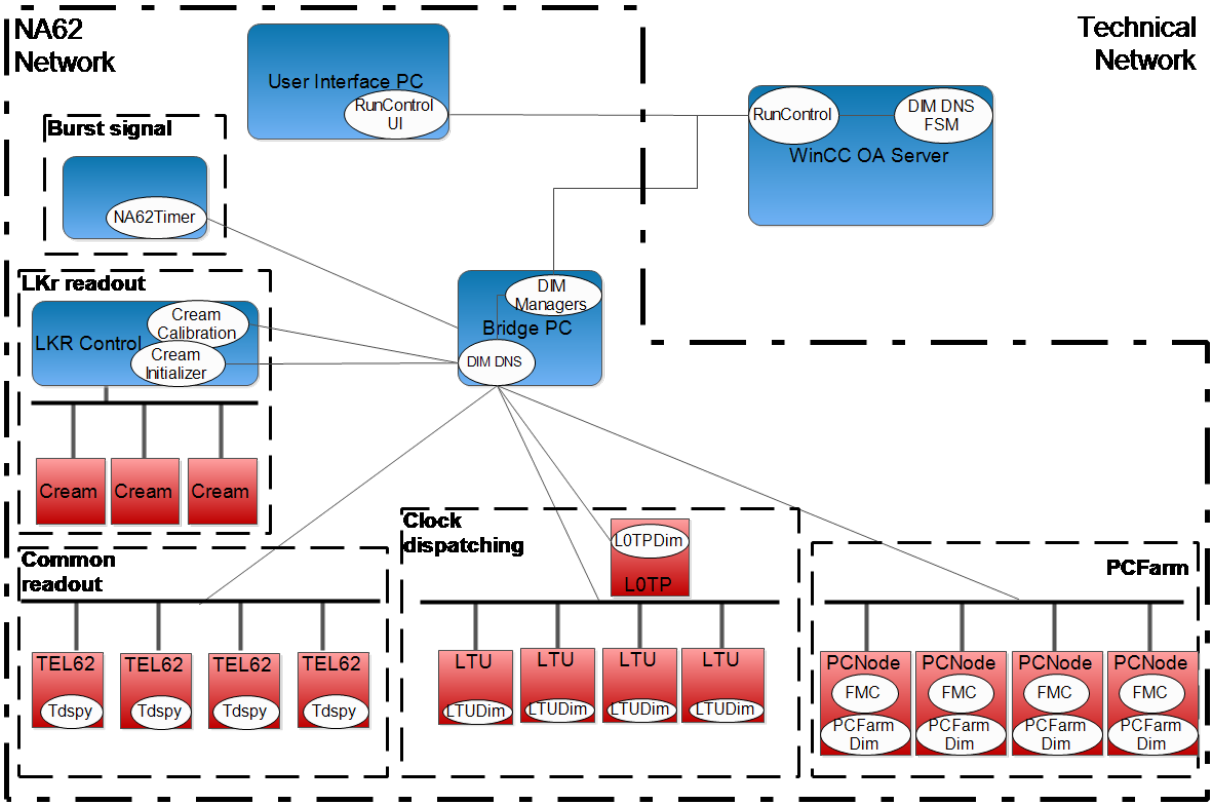

Figure 5.8: Schematic of the infrastructure of the DAQ and Run Control of the NA62 experiment. Red elements are for manipulating data. Blue elements are part of the control system of the TDAQ.

in Figure 5.8. The core of the system is located on a dedicated WinCC OA data server on the technical network. The rest of the equipment is on a different network, requiring a machine bridging both. This bridge machine hosts the DIM managers of the Run Control, which are responsible for loosely binding all the devices in the experiment and the Run Control. In case of a network problem, or if this node is cut off from one or the other network, the DAQ of the experiment continues to run correctly by itself. Once communication is re-established, the Run Control resumes control. This bridge also receives external information such as the start and end of burst signals, or information about the current run from the accelerator operators. The user interface itself runs on a different computer in the control room and is remotely connected to the main system.

5.4 Automatization of the Run Control: EleNA62

During 2013–2015 (runs, dry runs[1] and test runs), the basic functionalities of the Run Control (basic configuration, starting/ending a run, archiving of conditions, collection and display of information for diagnostics) were implemented and have been demonstrated to work correctly. The system has also been continuously updated with new subsystems and a number of shift-takers and experts have used the Run Control successfully, accumulating invaluable experience.

It has nevertheless been clear from early on that the system is relatively complex. Some problems can have small but important effects and be difficult to spot, locate and fix. The attention of the operators, whether in the daytime or nighttime, can also be diverted. Issues with the data taking could therefore easily stay unnoticed for a long time.

The EleNA62 (Electronic Eye of NA62) system, currently in an early stage of development, aims to assist the operators in their task, replacing human intervention when possible or trying to catch the attention of the operator by highlighting detected issues. It is built as a core component, providing tools and basic functionalities, and modules that can be enabled/disabled at any time. The modules provide the monitoring and control for specific elements of the TDAQ. Three of them currently exist for the PC farm, the beam line magnets and the decay volume vacuum.

5.4.1 Core component

On startup the core of EleNA62 loads all the existing modules and their configuration. If the configuration is modified while running, this is detected and the configuration is reloaded. The core of EleNA62 provides the main event loop of the system. Every module provides the frequency at which it expects to run its main loop and the core takes care of calling it when required.

[1] runs without beam

More importantly, the core provides a framework for logging and reporting events on three levels. The log records a timestamped message every time one of the monitoring functions of a module is triggered, or an automatic action is taken. They are kept for two days and are useful to experts for debugging and deciphering later possible problems not understood by the operator.

The notification mechanism is used to report events to the operator. A dedicated panel displays all the active notifications, accompanied by a message, a timestamp and an acknowledgement state. They are displayed visually to attract the attention of the operator, who has the possibility to "acknowledge" it, meaning actions will be taken to resolve the issue, or to solve it, meaning actions have been taken and the problem has been resolved.

The last reporting mechanism is audio. A playlist has been defined, so that the audio file corresponding to events in the Run Control is played in the control room. These events can be an operation of the Run Control (start of run, end of run, loss of connectivity, etc) or in the case of a critical error requiring immediate intervention (loss of vacuum, loss of a magnet in the beam).

5.4.2 PC farm module

The PC farm module has an important role to play in the data acquisition by ensuring the lowest possible downtime of the acquisition software. It constantly monitors the state of all PC nodes during a run. When one of them crashes and does not recover after the end of a burst, it is restarted. A count of the number of restarts is kept for each node to detect patterns that would indicate a faulty node or a systematic corruption of the data. If a change in the number of data sources is detected, all the nodes are automatically restarted, as it needs to be taken into account.

The FMC component, described in subsection 5.3.1, is used to monitor the hardware of all farm nodes (including the mergers). When a hardware failure is detected, a fixed number of attempts are made to power cycle the node at 10 min intervals. If not successful,

a notification of failure is sent to the operator for actions to be taken.

Finally the disk usage of the three merger PCs is monitored. If the disk is full, the data gathered during the burst cannot be written and are discarded, resulting in data loss. A notification is sent to the operator when the disk usage reaches 90 %. At 98 % a second notification is sent, the merger is automatically excluded from the data acquisition and the entire farm is restarted.

5.4.3 Magnets module

The magnets of the experiment and the beam line are controlled and monitored by a dedicated application provided by the CERN beam department. However, following a series of incidents where a faulty magnet or manual override did not trigger the protection mechanism which is supposed to shut down the beam, a redundant monitoring system was implemented as an EleNA62 module. The expected and actual electrical currents in the magnets are provided by the beam department through DIP (a read-only version of DIM) and are used by this module to assess the state of all magnets. If an inconsistency is detected, an alarm is rung in the control room and a notification is sent to the operator, requesting the beam to be stopped immediately.

5.4.4 Vacuum module

The vacuum pressure is monitored by a dedicated application provided by the CERN vacuum group. Despite not being critical to the safety of the experiment, a failure of the vaccum system is still harmful to the quality of the data. Following an incident where such a failure went unnoticed for more than two days, a redundant module was implemented in EleNA62. The expected values of the pressure in the evacuated volumes are encoded in the system and the actual readings from the sensors are acquired through DIP. If a safety threshold around the expected values is exceeded, a notification is sent on the Run Control and the alarm is rung in the control room.

5.4.5 Future developments

The aim of all these developments is to minimise the time during which good data cannot be acquired. An important part of the knowledge assimilated during the operation of the Run Control is related to recurrent failures of the acquisition boards. Some of these failures are linked to radiation or other systematic issues that cannot be eliminated, but for which a simple recovery procedure is known. Similarly many operating procedures have been devised and tested for normal operation. As a support for the operators it is foreseen to implement them directly in EleNA62, which will automatically execute those that do not require human intervention (e.g. reset of a stuck board). For the others, a visual guidance of the steps requiring input will be presented to them (for example in Figure 5.9).

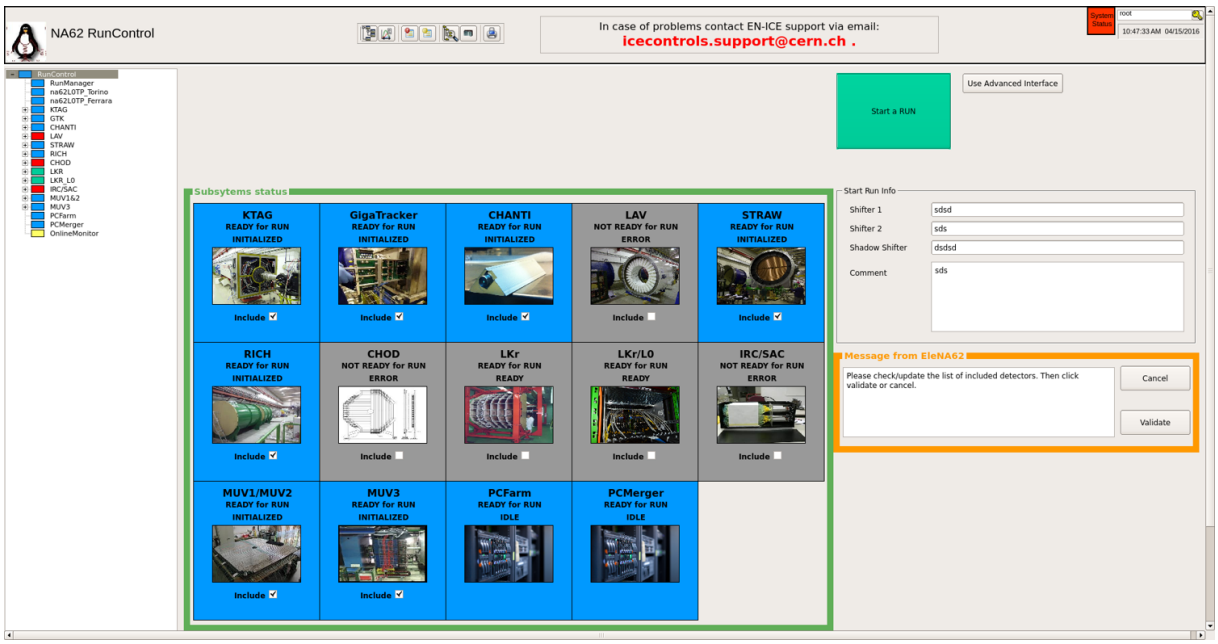

(a)

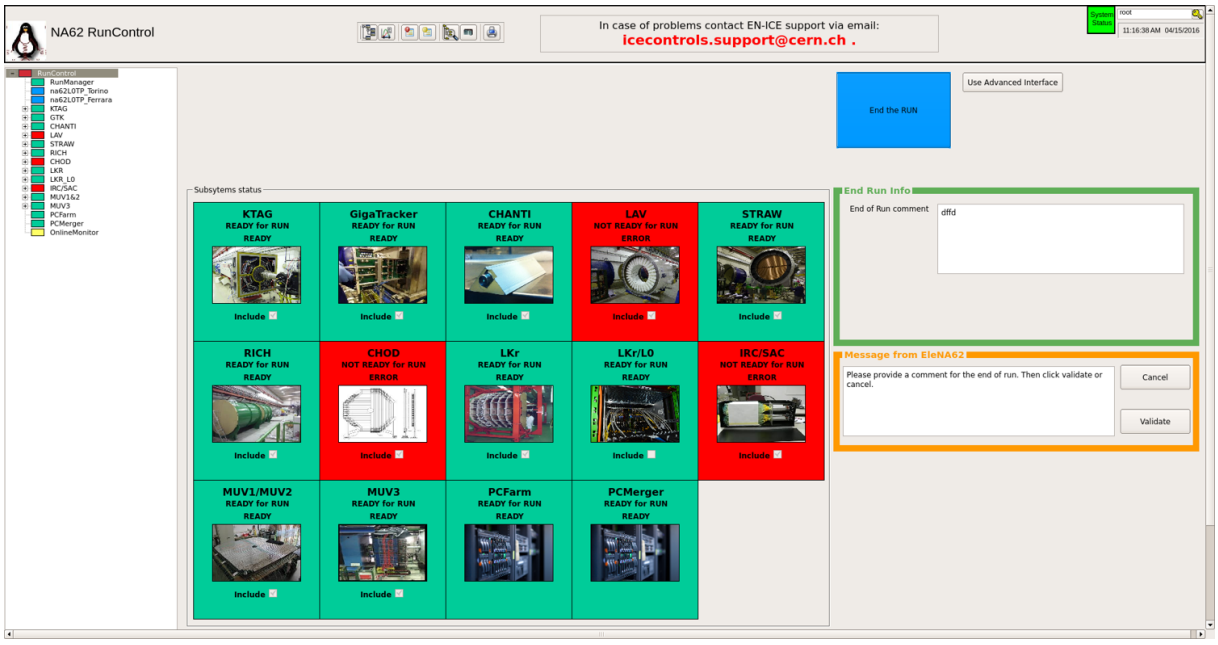

(b)

Figure 5.9: Example of the new interface, where a single button is used to automatise the start and end of run procedure. Little input is requested from the operator, and they receive textual instructions and coloured visual guidance through coloured boxes indicating where the input is needed. Starting a run in (a), and ending a run in (b).

CONCLUSIONS

About 1.1 million fully reconstructed $\pi^0 \to e^+e^-\gamma$ events from the data sample acquired in 2007 by the NA62-R_K experiment are studied. The slope of the π^0 electromagnetic transition form factor is measured to be

$$a = (3.68 \pm 0.51_{\text{stat}} \pm 0.25_{\text{syst}}) \times 10^{-2} = (3.68 \pm 0.57) \times 10^{-2} \, . \tag{5.1}$$

The uncertainty is dominated by the limited size of the data sample. This result is the most precise measurement of the form factor slope in the time-like region of momentum to date, with a factor two improvement in precision with respect to the previous measurement [41]. A positive slope is observed with a significance of more than 6σ, and is compatible with the theoretical predictions. The central value is compatible within 1σ with the vector meson dominance expectations.

The NA62 experiment started physics data taking in 2015 with the aim of measuring the branching ratio of the ultra-rare decay $K^+ \to \pi^+\nu\bar{\nu}$. Thanks to a hierarchy of finite state machines, the Run Control binds the complex Trigger and Data Acquisition system of the experiment in a coherent and uniform way, allowing its control without knowledge of all its components. Standard technologies, supported and developed at CERN for the LHC experiments, are used, which greatly accelerated the initial development of the system and ensures its reliability and future maintenance. The Run Control is tailored to the specific needs of NA62, with important efforts in particular for the configuration of the equipment. It has already achieved a high level of realiability and efficiency and is constantly improving, following requests and suggestions from experts and operators.

LIST OF REFERENCES

[1] N. Lurkin. Recent results from NA48/2 and NA62 experiments at CERN. *PoS*, HQL2016, 2016.

[2] C. Lazzeroni et al. Measurement of the π^0 electromagnetic transition form factor slope. *Physics Letters B*, 768:38 – 45, 2017.

[3] N. Lurkin. The NA62 Run Control. In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2013 IEEE*, pages 1–4, Oct 2013.

[4] N. Lurkin. The NA62 run control. *JPCS*, 556(1):012074, 2014.

[5] The NA62 collaboration. *NA62 Technical Design Document*, 2010.

[6] G.D. Rochester and C.C. Butler. Evidence for the Existence of New Unstable Elementary Particles. *Nature*, 160:855–857, Dec 1947.

[7] J.H. Christenson, J.W. Cronin, V.L. Fitch, and R. Turlay. Evidence for the 2π Decay of the K_2^0 Meson. *Phys. Rev. Lett.*, 13:138–140, Jul 1964.

[8] A. Lai et al. A precise measurement of the direct CP violation parameter $\mathrm{Re}(\varepsilon'/\varepsilon)$. *Eur. Phys. J. C*, 22(2):231–254, 2001.

[9] A.J. Buras, D. Buttazzo, J. Girrbach-Noe, and R. Knegjens. $K^+ \to \pi^+\nu\bar{\nu}$ and $K_L \to \pi^0\nu\bar{\nu}$ in the Standard Model: status and perspectives. *J. High Energy Phys.*, 2015(11):1–34, 2015.

[10] A.V. Artamonov et al. Study of the decay $K^+ \to \pi^+\nu\bar{\nu}$ in the momentum region $140 < P_\pi < 199$ MeV/c. *Phys. Rev. D*, 79:092004, May 2009.

[11] C.M.G. Lattes, G.P.S. Occhialini, and C.F. Powell. Observations on the Tracks of Slow Mesons in Photographic Emulsions. *Nature*, 160:453–456, 1947.

[12] R. Bjorklund, W.E. Crandall, B.J. Moyer, and H.F. York. High Energy Photons from Proton-Nucleon Collisions. *Phys. Rev.*, 77:213–218, January 1950.

[13] H. Yukawa. On the interaction of elementary particles. *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series*, 17:48–57, 1935.

[14] G-M. Murray. A Schematic Model of Baryons and Mesons. *Phys. Lett.*, 8:214–215, 1964.

[15] G. Zweig. An SU_3 model for strong interaction symmetry and its breaking. (CERN-TH-412):80 p, Feb 1964.

[16] N.M. Kroll and W. Wada. Internal Pair Production Associated with the Emission of High-Energy Gamma Rays. *Phys. Rev.*, 98:1355–1359, Jun 1955.

[17] K.A. Olive et al. Review of Particle Physics. *Chin. Phys. C*, 40(10):100001, 2016.

[18] R.H. Dalitz. On an Alternative Decay Process for the Neutral π-Meson. *Proceedings of the Physical Society. Section A*, 64(7):667, 1951.

[19] F. Jegerlehner and A. Nyffeler. The muon g-2. *Physics Reports*, 477(13):1 – 110, 2009.

[20] A. Nyffeler. Precision of a data-driven estimate of hadronic light-by-light scattering in the muon $g-2$: Pseudoscalar-pole contribution. *Phys. Rev. D*, 94:053006, Sep 2016.

[21] A.E. Dorokhov and M.A. Ivanov. Rare decay $\pi^0 \to e^+e^-$: Theory confronts KTeV data. *Phys. Rev. D*, 75:114007, Jun 2007.

[22] T. Husek, K. Kampf, and J. Novotný. Radiative corrections to the Dalitz decay $\pi^0 \to e^+e^-\gamma$ revisited. *Phys. Rev. D*, 92(5):054027, 2015.

[23] J.J. Sakurai. Theory of strong interactions. *Annals of Physics*, 11(1):1 – 48, 1960.

[24] G-M. Murray and Z. Fredrik. Form Factors and Vector Mesons. *Phys. Rev.*, 124:953–964, Nov 1961.

[25] D.W. Joseph. Electron pair creation in π+p capture reactions from rest. *Il Nuovo Cimento*, 16(6):997–1013, 1960.

[26] B.E. Lautrup and J. Smith. Radiative corrections to decays with a Dalitz pair. *Phys. Rev. D*, 3:1122–1135, 1971.

[27] K.O. Mikaelian and J. Smith. Radiative corrections to the decay $\pi^0 \to \gamma e^+ e^-$. *Phys. Rev. D*, 5:1763–1773, 1972.

[28] *Random Number Generation and Monte Carlo Methods.* Springer-Verlag New York, 2003.

[29] F. James. Monte-Carlo phase space. Report CERN-68-15, CERN, 1968.

[30] H.J. Behrend et al. A measurement of the π^0, η and η' electromagnetic form-factors. *Z. Phys. C*, 49:401–410, 1991.

[31] J. Gronberg et al. Measurements of the meson - photon transition form-factors of light pseudoscalar mesons at large momentum transfer. *Phys. Rev. D*, 57:33–54, 1998.

[32] B. Aubert et al. Measurement of the $\gamma\gamma^* \to \pi^0$ transition form factor. *Phys. Rev. D*, 80:052002, 2009.

[33] S. Uehara et al. Measurement of $\gamma\gamma^* \to \pi^0$ transition form factor at Belle. *Phys. Rev. D*, 86:092007, 2012.

[34] *MesonNet Workshop on Meson Transition Form Factors*, 2012.

[35] M.N. Achasov et al. Experimental study of the $e^+e^- \to \pi^0\gamma$ process in the energy region $\sqrt{s} = 0.60 - 0.97$ GeV. *Phys. Lett. B*, 559(34):171 – 178, 2003.

[36] R.R. Akhmetshin et al. Study of the processes $e^+e^- \to \eta\gamma$, $\pi^0\gamma \to 3\gamma$, in the c.m. energy range $600 - 1380$ MeV at CMD-2. *Phys. Lett. B.*, 605(12):26 – 36, 2005.

[37] J. Fischer et al. Measurement of the Neutral Pion Electromagnetic Form-Factor Slope. *Phys. Lett. B*, 73:359, 1978.

[38] H. Fonvieille et al. Dalitz Decay: $\pi^0 \to \gamma e^+e^-$ and the π^0 Electromagnetic Transition Form-factor. *Phys. Lett. B*, 233:65–68, 1989.

[39] F. Farzanpay et al. Measurement of the slope of the π^0 electromagnetic form-factor. *Phys. Lett. B*, 278:413–418, 1992.

[40] R. Meijer Drees et al. Measurement of the π^0 electromagnetic transition form-factor. *Phys. Rev. D*, 45:1439–1447, 1992.

[41] P. Adlarson et al. Measurement of the $\pi^0 \to e^+e^-\gamma$ Dalitz decay at MAMI. *arXiv:1611.04739 [hep-ex]*, 2016.

[42] M. Kobayashi and T. Maskawa. CP Violation in the Renormalizable Theory of Weak Interaction. *Prog. Theor. Phys.*, 49:652–657, 1973.

[43] N. Cabibbo. Unitary Symmetry and Leptonic Decays. *Phys. Rev. Lett.*, 10:531–533, Jun 1963.

[44] L. Wolfenstein. Parametrization of the Kobayashi-Maskawa Matrix. *Phys. Rev. Lett.*, 51:1945–1947, Nov 1983.

[45] S.L. Glashow, J. Iliopoulos, and L. Maiani. Weak Interactions with Lepton-Hadron Symmetry. *Phys. Rev. D*, 2:1285–1292, Oct 1970.

[46] J.K. Ahn et al. Experimental study of the decay $K_L^0 \to \pi^0\nu\overline{\nu}$. *Phys. Rev. D*, 81:072004, Apr 2010.

[47] T. Yamanaka. The J-PARC KOTO experiment. *Progr. Theor. Exp. Phys.*, 2012(1), 2012.

[48] A.V. Artamonov et al. New Measurement of the $K^+ \to \pi^+ \nu\overline{\nu}$ Branching Ratio. *Phys. Rev. Lett.*, 101:191802, Nov 2008.

[49] V. Fanti et al. The beam and detector for the NA48 neutral kaon CP violation experiment at CERN. *Nucl. Instrum. Meth. Phys. A*, 574(3):433 – 471, 2007.

[50] J.R. Batley et al. Search for direct CP violating charge asymmetries in $K^\pm \to \pi^\pm\pi^+\pi^-$ and $K^\pm \to \pi^\pm\pi^0\pi^0$ decays. *Eur. Phys. J. C*, 52(4):875–891, 2007.

[51] NA48/2 collaboration website. http://na48.web.cern.ch/NA48/NA48-2/Overview/NA48_2Description.html. Accessed: 2016-11-30.

[52] NA48/2 collaboration website. http://na48.web.cern.ch/NA48/Welcome/images/calorimeter/calorimeter.html. Accessed: 2016-11-30.

[53] C. Lazzeroni et al. Precision Measurement of the Ratio of the Charged Kaon Leptonic Decay Rates. *Phys. Lett. B*, 719:326–336, 2013.

[54] K.L. Brown and F.C. Iselin. Decay TURTLE (Trace Unlimited Rays Through Lumped Elements): A Computer Program for Simulating Charged Particle Beam Transport Systems, Including Decay Calculations. 1974.

[55] C. Gatti. Monte Carlo simulation for radiative kaon decays. *Eur. Phys. J. C*, 45(2):417–420, 2006.

[56] R. Brun, F. Bruyant, M. Maire, A.C. McPherson, and P. Zanarini. GEANT3. 1987.

[57] C. Biino et al. *COmPACT 7.3 Users Guide*, 2007.

[58] R. Brun and F. Rademakers. ROOT An object oriented data analysis framework. *Nucl. Instrum. Meth. Phys. A*, 389(1):81 – 86, 1997.

[59] K.A. Olive et al. Review of Particle Physics. *Chin. Phys. C*, 38:090001, 2014.

[60] R. Frühwirth. Application of Kalman filtering to track and vertex fitting. *Nucl. Instrum. Meth. Phys. A*, 262(2):444 – 450, 1987.

[61] N.D. Gagunashvili. Comparison of weighted and unweighted histograms. *PoS(ACAT)054*, 2006.

[62] N.D. Gagunashvili. Chi-square tests for comparing weighted histograms. *Nucl. Instrum. Meth. Phys.*, 614(2):287 – 296, 2010.

[63] K. Pearson. On the Theory of Contingency and Its Relation to Association and Normal Correlation. Drapers Co. Memoirs, Biometric Series No. 1, London, 1904.

[64] F. James and M. Roos. Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations. *Comput. Phys. Commun.*, 10:343–367, 1975.

[65] Verena Kain, Karel Cornelis, and Ewald Effinger. New Spill Control for the Slow Extraction in the Multi-Cycling SPS. In *Proceedings, 7th International Particle Accelerator Conference (IPAC 2016): Busan, Korea, May 8-13, 2016*, page TUPMR051, 2016.

[66] The NA62 collaboration. The NA62 Detector and Beam at the CERN SPS Accelerator. To be published.

[67] E. Goudzovski et al. Development of the kaon tagging system for the NA62 experiment at CERN. *Nucl. Instrum. Meth. A*, 801:86–94, 2015.

[68] D. Aisa et al. The RICH detector of the NA62 experiment at CERN. *Nucl. Instrum. Meth. A*, 766:1 – 4, 2014.

[69] B. Angelucci, E. Pedreschi, M. Sozzi, and F. Spinella. TEL62: an integrated trigger and data acquisition board. *J. Instrum.*, 7(02):C02046, 2012.

[70] G. Haefeli et al. The LHCb DAQ interface board TELL1. *Nucl. Instrum. Meth. Phys. A*, 560(2):494 – 502, 2006.

[71] E. Pedreschi et al. A High-Resolution TDC-Based Board for a Fully Digital Trigger and Data Acquisition System in the NA62 Experiment at CERN. *IEEE Trans. Nucl. Sci.*, 62(3):1050–1055, June 2015.

[72] C. Gaspar, M. Dönszelmann, and Ph. Charpentier. DIM, a portable, light weight package for information publishing, data transfer and inter-process communication. *Comput. Phys. Commun.*, 140(12):102 – 109, 2001. CHEP2000.

[73] M. Gonzalez-Berges. The Joint COntrols project framework. *eConf*, C0303241:THGT006, 2003.

[74] D. Veillard. http://www.xmlsoft.org/index.html.

[75] H. Milcent et al. UNICOS: An open framework. In *ICALEPCS 2009, Kobe, Japan*, 2009.

[76] B. Franek and C. Gaspar. SMI++ object oriented framework for designing and implementing distributed control systems. In *Nuclear Science Symposium Conference Record, 2004 IEEE*, volume 3, pages 1831–1835 Vol. 3, Oct 2004.