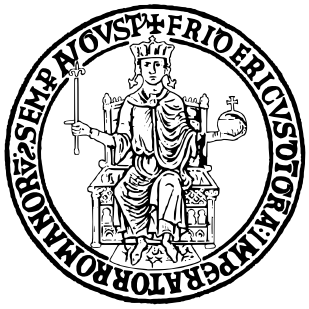

SSM
Scuola Superiore Meridionale

Università degli Studi di Napoli Federico II
Scuola Superiore Meridionale

PhD thesis in

Mathematical and Physical Sciences
for Advanced Materials and Technologies

XXXVI cycle

CERN-THESIS-2024-332
21/01/2025

Observation of the $K^+ \to \pi^+ \nu\bar{\nu}$ decay and measurement of its branching ratio with the NA62 experiment

Supervisor
Prof. Fabio Ambrosino

Candidate
Renato Fiorenza

Academic Years 2020–2024

Abstract

The ultra-rare $K^+ \to \pi^+\nu\bar{\nu}$ decay is a golden mode to test the Standard Model: its branching ratio is predicted very precisely, and it is highly sensitive to indirect effects of physics beyond the Standard Model. The NA62 experiment at CERN has been designed to measure $BR(K^+ \to \pi^+\nu\bar{\nu})$, and provided the most precise measurement of this branching ratio using data collected in 2016–2018.

In this PhD thesis the analysis of data collected by NA62 in 2021–2022 is reported. This dataset was collected after upgrades to the experimental apparatus, and at a higher instantaneous beam intensity with respect to the previous 2016–2018 data taking. Combining NA62 data collected between 2016 and 2022, the first observation of the $K^+ \to \pi^+\nu\bar{\nu}$ decay with a significance exceeding 5σ is reported, and a measurement of $BR(K^+ \to \pi^+\nu\bar{\nu}) = (13.0^{+3.3}_{-3.0}) \times 10^{-11}$ is obtained. To date, this is the smallest branching ratio ever measured with a signal significance above 5σ.

Acknowledgments

This thesis would not have been possible without everyone who has been a part of my life, especially during my PhD studies.

I am deeply grateful to my supervisor Fabio Ambrosino for his invaluable guidance, both scientific and human. I am also extremely thankful to Marco Mirra and all the members of the NA62-Napoli group, including Marco Francesconi, Raffaele Giordano, Paolo Massarotti, Marco Napolitano, Ilaria Rosa and Giulio Saracino, for their extensive support.

Special thanks to Francesco Brizioli and Joel Swallow for being great researchers and teammates.

I thank all the members of the NA62 collaboration for the fruitful interactions and for giving me a great time working on our experiment.

Finally, I thank Michela and my whole family for always being by my side.

Contents

Introduction

The Standard Model of particle physics (SM) is a quantum field theory developed in stages by many scientists worldwide during the course of the second half of the 20th century [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11], and providing a comprehensive description of the elementary particles and their interactions. According to the SM [12], elementary particles and their interactions are described in terms of fermion fields (grouped into *quarks* and *leptons*) and boson fields (*gauge bosons*, plus the *Higgs boson*). The extensive history of experimental achievements [13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26] makes it one of the most successful theories ever. Nonetheless, many fundamental questions about nature are still open: for instance, neutrino oscillations, the observed baryon asymmetry between matter and anti-matter, the existence of dark matter, and the whole theory of gravitation, lack explanation in the SM [27].

Today, the high energy physics community is focused to carry on stringent tests of the SM, in order to search, both directly and indirectly, for evidence of physics beyond the SM. Two complementary approaches are being adopted to this aim. One is to reach increasingly higher energy in the collision of particles (such as the protons at the CERN LHC), in order to directly search for production of new particles not predicted by the SM (*energy frontier*). A different approach consists in collecting as much statistics as possible (*intensity frontier*), in order to search for deviations from the SM predictions in effects that are quantitatively small but accurately known: this allows to probe energy scales well above than what is currently possible with direct searches.

The latter approach is employed by the NA62 experiment at the CERN Super Proton Synchrotron (SPS) [28]. Its goal is to measure the branching ratio (BR) of the ultra rare $K^+ \to \pi^+\nu\bar{\nu}$ decay, predicted very precisely by the SM [29, 30, 31, 32, 33]. This decay is a golden channel for the search of new physics, thanks to the existence of various theories beyond the SM that predict its BR to be significantly different from the SM prediction.

NA62 has collected data during its first run of data taking, between 2016 and 2018, leading to the world's most precise measurement for BR($K^+ \to \pi^+\nu\bar{\nu}$) [34]. During CERN Long Shutdown 2 (between 2018 and 2021), the NA62 apparatus was upgraded to be more resilient against the dominant backgrounds for this decay. Data taking resumed in 2021, at the highest

instantaneous intensity that the NA62 apparatus can sustain. NA62 is currently in its second run of data taking, foreseen to last until CERN Long Shutdown 3.

The subject of this thesis is the analysis of the data collected in 2021 and 2022. As part of the $K^+ \to \pi^+\nu\bar{\nu}$ analysis team, my main work during this PhD course was to implement this analysis, and to provide original contributions to parts of it. A large fraction of my contributions were given in: the K^+–π^+ matching (described in section 3.5 and in appendix A), which controls the main mechanisms for the largest backgrounds of this analysis; the evaluation and correction of the random veto efficiency (section 4.5) which is a crucial parameter to keep under control when dealing with data taken at high intensities; the estimation and validation of the upstream background (section 5.2), the main background of this analysis; the statistical treatment of the results (section 6.3), which is essential to convert the observation of the signal region into an actual BR measurement. Further studies that I performed during my PhD course, less directly related to the $K^+ \to \pi^+\nu\bar{\nu}$ analysis, are reported in appendices B and C.

The preliminary results of this analysis have been presented for the first time in a CERN seminar at the Experimental Physics department in September 2024 [35], and I presented it at the PIC2024 conference [36]. The paper with the final result is in preparation.

The main content of this thesis is organized as follows. Chapter 1 describes the state of the art for the $K^+ \to \pi^+\nu\bar{\nu}$ decay, both from a theoretical and from an experimental point of view. Chapter 2 contains a description of the NA62 experimental apparatus. Chapter 3 illustrates the event selection employed for the analysis. Chapter 4 and 5 are dedicated to the estimations of the signal and the background, respectively. Finally, Chapter 6 contains the results of the analysis.

Chapter 1

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

K mesons, also known as *kaons*, have played an important role in the history of particle physics. Their discovery in cosmic rays in 1947 [37] eventually led to the postulation of the existence of the *strangeness* quantum number and the formulation of the quark model by Gell-Mann [4] and Zweig [5, 6] (independently). One of the first observation of the parity violation, as we now understand it, was in the kaon system: the θ-τ *puzzle*, namely the observation of two apparently identical particles which only differed by intrinsic parity [38], was eventually solved by the discovery of parity violation in the weak interaction [39]. An experiment carried out by Christenson, Cronin, Fitch, and Turlay in 1964 to study neutral kaon decays provided the first clear evidence of CP violation [40]; it won Cronin and Fitch the 1980 Nobel Prize in Physics. More recently, the study of the neutral kaon system led to the first observation of direct CP violation by the experiments NA48 at CERN [41] and KTeV at FNAL [42] in 1999.

Today, experiments on kaons are of crucial importance in the search for physics beyond the Standard Model (BSM). In particular, these experiments mostly aim to explore the intensity frontier: high intensity experiments, typically employing fixed targets, are sensitive to BSM physics on scales up to hundreds of TeV through virtual effects. These scales are to be compared to the ones probed by direct searches for new particles, such as the ones performed at the Large Hadron Collider at CERN, which can only reach few TeV.

Flavour physics studies processes involving transitions of fermions: in the quark sector, quark mixing plays a fundamental role for these processes, and it is parametrized by the Cabibbo-Kobayashi-Maskawa (CKM) matrix. Kaon experiments, which probe the physics of the s quark, can constrain the CKM matrix; in this respect, they are complementary to, and independent from, B meson experiments studying the b quark.

The $K^+ \to \pi^+ \nu\bar{\nu}$ decay, together with its neutral counterpart $K_L \to \pi^0 \nu\bar{\nu}$, is a golden mode to experimentally probe new physics in the flavour sector, as it is extremely sensitive to new physics.

Section 1.1 is a brief introduction to the theoretical background for the discussion of the $K \to \pi\nu\bar{\nu}$ decays. In section 1.2, an overview of the state of the art in the theory of $K \to \pi\nu\bar{\nu}$ decays is given, both in the SM and in BSM physics. The experimental status of the $K \to \pi\nu\bar{\nu}$ decays, with brief historical considerations, is described in section 1.3. Particular emphasis is put on the charged channel throughout this chapter.

1.1 Quark mixing

Quark mixing was first proposed by Cabibbo [43], in the context of the physics of u, d and s quarks, in order to recover the universality of the weak interactions in the quark sector. It was later extended to three quark generations by Kobayashi and Maskawa [44].

1.1.1 The CKM matrix

The fundamental postulate for quark mixing is that the weak eigenstates and the mass eigenstates of the quarks do not coincide; they are related via a unitary matrix. Quark fields are gauge invariant for rotations in the complex plane (namely, they can be redefined up to a complex phase): for this reason, it can be shown that it is sufficient to consider only the mixing of the down-type quarks $(\mathrm{d}, \mathrm{s}, \mathrm{b})$. Their weak eigenstates $(\mathrm{d}', \mathrm{s}', \mathrm{b}')$ can be expressed in terms of the corresponding mass eigenstates by

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

The same gauge invariance exploited above and the unitarity requirement constrain the CKM matrix to have only 4 free parameters: these can be described as three mixing angles and one complex phase which quantifies the CP violation in the quark sector.

The experimental values of the CKM matrix elements exhibit a clear hierarchical structure, which is unexplained in the SM and constitutes one of the open problems within the current theory. This structure is best highlighted with the Wolfenstein parameterization [45]: expanding each element of the matrix as a power series in terms of

$$\lambda = \frac{\left|V_{\mathrm{us}}\right|}{\sqrt{\left|V_{\mathrm{ud}}\right|^2 + \left|V_{\mathrm{us}}\right|^2}} \approx \left|V_{\mathrm{us}}\right| ,$$

the matrix can be expressed as

$$V_{\mathrm{CKM}} = \begin{pmatrix} 1-\frac{\lambda^2}{2} & \lambda & A\lambda^3(\rho-i\eta) \\ -\lambda & 1-\frac{\lambda^2}{2} & A\lambda^2 \\ A\lambda^3(1-\rho-i\eta) & -A\lambda^2 & 1 \end{pmatrix} + O(\lambda^4)\,,$$

where the exact definitions of the (real) Wolfenstein parameters A, ρ and η are given by [27]

$$A\lambda^2 = \lambda \left| \frac{V_{\mathrm{cb}}}{V_{\mathrm{us}}} \right|\,,$$
$$A\lambda^3(\rho - i\eta) = V_{\mathrm{ub}}\,.$$

1.1.2 The unitarity triangle

The unitarity of the CKM matrix imposes

$$\sum_{q=\mathrm{u,c,t}} V_{qj} V_{qk}^* = \delta_{jk} \quad (j,k = \mathrm{d,s,b})\,, \tag{1.1a}$$

$$\sum_{q=\mathrm{d,s,b}} V_{iq} V_{jq}^* = \delta_{ij} \quad (i,j = \mathrm{u,c,t})\,. \tag{1.1b}$$

In particular, the relation

$$\frac{V_{\mathrm{ud}} V_{\mathrm{ub}}^*}{V_{\mathrm{cd}} V_{\mathrm{cb}}^*} + 1 + \frac{V_{\mathrm{td}} V_{\mathrm{tb}}^*}{V_{\mathrm{cd}} V_{\mathrm{cb}}^*} = 0\,,$$

obtained from equation (1.1a) with $j = \mathrm{d}$ and $k = \mathrm{b}$ and dividing by $V_{\mathrm{cd}} V_{\mathrm{cb}}^*$, can be represented in the complex plane as the *unitarity triangle*, shown in figure 1.1. Usually, $\bar{\rho}$ and $\bar{\eta}$ are then defined as

$$\bar{\rho} + i\bar{\eta} = -\frac{V_{\mathrm{ud}} V_{\mathrm{ub}}^*}{V_{\mathrm{cd}} V_{\mathrm{cb}}^*} = (\rho + i\eta)\left(1 - \frac{\lambda^2}{2} + O(\lambda^4)\right)\,,$$

so that they are represented by the apex of the unitarity triangle.

An important goal of many experiments in flavour physics is to constrain the CKM elements and the unitarity of the CKM matrix. The fit to the current experimental results gives [27]

$$\lambda = 0.22501 \pm 0.00068\,, \qquad A = 0.826^{+0.016}_{-0.015}\,,$$
$$\bar{\rho} = 0.1591 \pm 0.0094\,, \qquad \bar{\eta} = 0.3523^{+0.0073}_{-0.0071}\,.$$

Constraints on the $(\bar{\rho}, \bar{\eta})$ plane from various measurements are shown in figure 1.2.

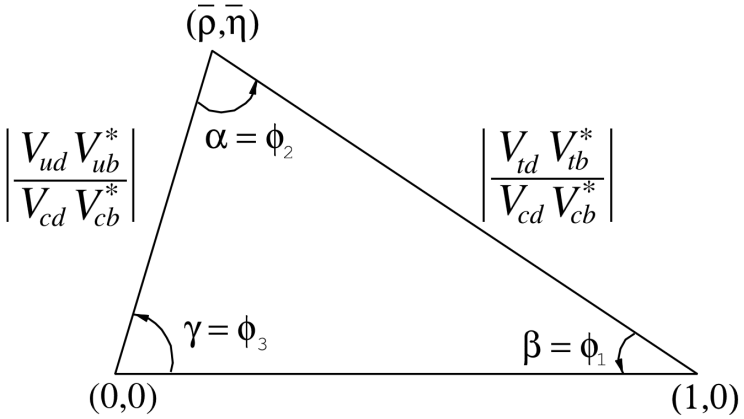

Figure 1.1: The unitarity triangle. [27]

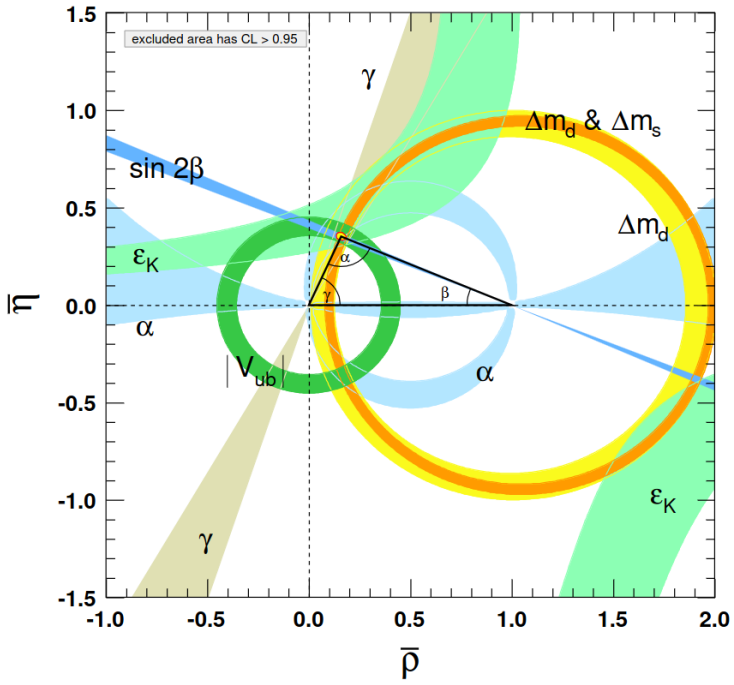

Figure 1.2: Experimental constraints on the $(\bar{\rho}, \bar{\eta})$ plane. [27]

1.1.3 Flavour changing neutral currents

In the lagrangian of the Standard Model, there are no currents involving quarks with same charge but different flavour. In other words, flavour changing neutral currents (FCNCs) are forbidden at tree level. They can occur in loop diagrams, but they are suppressed by the *GIM mechanism*. This mechanism is named after Glashow, Iliopoulos and Maiani, who first proposed it at a time where only three quarks (u, d and s) were known [9]: the existence of a fourth up-type quark was required in order to explain the suppression of the s $\rightarrow$ d FCNC.

Simply put, the s → d transition, at leading order, can only occur via an intermediate transition to an up-type quark (s → u → d, s → c → d, s → t → d): if all the up-type quarks had the same mass, the sum of the three amplitudes would be proportional to $V_{us}V_{ud}^* + V_{cs}V_{cd}^* + V_{ts}V_{td}^*$, which is zero by the unitarity relation (1.1a), that is, the transition would be forbidden. Since the quarks actually have mass differences, the transition is allowed but still heavily suppressed. Of course, a similar argument can be made for all FCNCs.

The suppression of FCNCs in the Standard Model is not due to any first principle: rather, it is an accident due to the structure of weak interactions. Many BSM models predict enhancements of FCNCs, especially at tree level; for this reason, FCNCs are considered powerful probes for BSM physics.

1.2 K → πνν̄ decays: theoretical status

The $K^+ \to \pi^+\nu\bar{\nu}$ and $K_L \to \pi^0\nu\bar{\nu}$ decays both proceed via an s → dνν̄ transition, which features a FCNC. Long-distance contributions to these decays are very small, which makes the theoretical computation very clean: precise calculations of the decay rates can be performed in terms of the SM parameters. At the same time, several BSM models predict deviations in either the charged or the neutral mode of the K → πνν̄ decay, and any experimental deviation from the Standard Model predictions would signal presence of new physics. This is why they are called *golden modes.*

1.2.1 K → πνν̄ in the Standard Model

The box and Z penguin diagrams contributing to the s → dνν̄ transition are shown in figure 1.3. In each diagram, the contribution to the amplitude from the quark $q = u, c, t$ is roughly proportional to $m_q^2 V_{qs}^* V_{qd}$: the largest contribution is therefore due to the t quark, with the c quark giving a small but non-negligible contribution, and the contribution from the u quark being extremely suppressed.

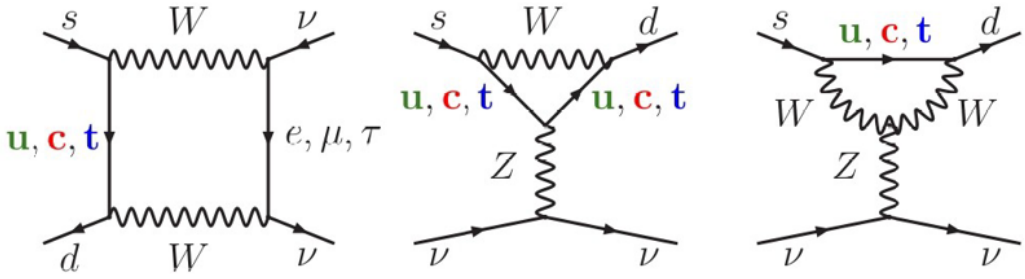

Figure 1.3: Feynman diagrams for s → dνν̄.

The branching ratio of the charged mode can be written as [29, 30]

$$\mathrm{BR}(\mathrm{K}^+ \to \pi^+\nu\bar{\nu}) = \kappa_+(1+\Delta_{\mathrm{EM}})\left[\left(\frac{\mathrm{Im}\left(V_{\mathrm{ts}}^* V_{\mathrm{td}}\right)}{\lambda^5}X_{\mathrm{t}}\right)^2 + \left(\frac{\mathrm{Re}\left(V_{\mathrm{cs}}^* V_{\mathrm{cd}}\right)}{\lambda}(P_{\mathrm{c}}+\delta P_{\mathrm{c,u}}) + \frac{\mathrm{Re}\left(V_{\mathrm{ts}}^* V_{\mathrm{td}}\right)}{\lambda^5}X_{\mathrm{t}}\right)^2\right],$$

(a derivation of this formula can be found in [46]) which exhibits the different ingredients that go into the calculation:

κ_+ is the hadronic matrix element, which can be expressed in terms of very well-measured semileptonic kaon decays (in particular $\mathrm{K}^+ \to \pi^0\mathrm{e}^+\nu$), using the isospin symmetry between π^+ and π^0. Including higher-order chiral corrections, it can be written as [47]

$$\kappa_+ = \left(\frac{\lambda}{0.225}\right)^8 \times (5.173 \pm 0.025) \times 10^{-11}\,.$$

The fact that this value is known so precisely is one of the main reasons for the small final uncertainty on $\mathrm{BR}(\mathrm{K}^+ \to \pi^+\nu\bar{\nu})$;

Δ_{EM} parametrizes the long-distance QED corrections. At next-to-leading order (NLO), it is equal to -0.003 [47];

X_{t} is the loop contribution from the t quark. Because of the large mass of the t quark, effectively it does not depend on the mass of the lepton running in the box diagram. It has been calculated including NLO QCD and EW corrections, and it is [30]

$$X_{\mathrm{t}} = 1.462 \pm 0.017_{\mathrm{QCD}} \pm 0.002_{\mathrm{EW}}\,;$$

P_{c} is the loop contribution from the c quark. Unlike the contribution from the t quark, it depends on the τ mass (the masses of e and μ remaining negligible). It has been calculated at NNLO in QCD and at NLO in electroweak interaction, and can be written as [30]

$$P_{\mathrm{c}} = \left(\frac{0.2255}{\lambda}\right)^4 \times (0.3604 \pm 0.0087)\,;$$

$\delta P_{\mathrm{c,u}}$ parametrizes additional long-distance contributions from c and u quarks, calculated at the one-loop level to be $\delta P_{\mathrm{c,u}} = 0.04 \pm 0.02$ [48].

The branching ratio of the neutral mode can be calculated in a similar way, via the formula [30]

$$\mathrm{BR}(\mathrm{K_L} \to \pi^0\nu\bar{\nu}) = \kappa_L r_{\epsilon_{\mathrm{K}}}\left(\frac{\mathrm{Im}\left(V_{\mathrm{ts}}^* V_{\mathrm{td}}\right)}{\lambda^5}X_t\right)^2,$$

where κ_L is the neutral counterpart to κ_+, and r_{ϵ_K} accounts for the small contribution from indirect CP violation. The cancellation of the contributions from the intermediate c quark state makes the calculation even more precise.

A recent numerical evaluation of these BRs gives [31]

$$\mathrm{BR}(K^+ \to \pi^+\nu\bar{\nu}) = (7.86 \pm 0.61) \times 10^{-11}\,, \tag{1.2a}$$

$$\mathrm{BR}(K_L \to \pi^0\nu\bar{\nu}) = (2.68 \pm 0.30) \times 10^{-11}\,, \tag{1.2b}$$

where more than 70 % of the uncertainty is due to the precision of the knowledge of the SM parameters [30].

In terms of the CKM parameters, the two BRs can be written as [49]

$$\mathrm{BR}(K^+ \to \pi^+\nu\bar{\nu}) \approx (7.92 \pm 0.28) \times 10^{-11} \left(\frac{|V_{cb}|}{41.0 \times 10^{-3}}\right)^{2.8} \left(\frac{\sin\gamma}{\sin 67^\circ}\right)^{1.39}\,,$$

$$\mathrm{BR}(K_L \to \pi^0\nu\bar{\nu}) = (2.61 \pm 0.04) \times 10^{-11} \left(\frac{|V_{cb}|}{41.0 \times 10^{-3}}\right)^{4} \left(\frac{\sin\gamma}{\sin 67^\circ}\right)^{2} \left(\frac{\sin\beta}{\sin 22.2^\circ}\right)^{2}$$

(dependence on λ is weak and not shown). These expressions show clearly the strong dependence of these branching ratios, and of their uncertainties, on $|V_{cb}|$ and on two angles of the unitarity triangle (see figure 1.1). It must be said that the determination of $|V_{cb}|$ is affected by long-standing tensions in experimental results between the inclusive and the exclusive determination. It is therefore possible to recast these equations in terms of different, more precisely measured, flavour physics parameters (one must assume that their measurements are not polluted by contributions from BSM physics): a suitable choice gives

$$\mathrm{BR}(K^+ \to \pi^+\nu\bar{\nu}) = (8.60 \pm 0.42) \times 10^{-11}\,, \tag{1.3a}$$

$$\mathrm{BR}(K_L \to \pi^0\nu\bar{\nu}) = (2.94 \pm 0.15) \times 10^{-11}\,, \tag{1.3b}$$

which are the most precise determinations to date, albeit in slight tension with (1.2) [49, 32]. A very recent paper uses a combination of inclusive and exclusive determinations of $|V_{cb}|$ and obtains $\mathrm{BR}(K^+ \to \pi^+\nu\bar{\nu}) = (8.09 \pm 0.63) \times 10^{-11}$ and $\mathrm{BR}(K_L \to \pi^0\nu\bar{\nu}) = (2.58 \pm 0.30) \times 10^{-11}$ [33].

It is interesting to note that it is possible to recast $\mathrm{BR}(K^+ \to \pi^+\nu\bar{\nu})$ as [27]

$$\mathrm{BR}(K^+ \to \pi^+\nu\bar{\nu}) \approx 1.6 \times 10^{-5} |V_{cb}|^4 \left[(\bar{\rho} - 1.45)^2 + \left(\frac{\bar{\eta}}{\left(1 - \frac{\lambda^2}{2}\right)^2}\right)^2\right]\,,$$

which shows that a measurement of this BR determines an ellipse in the $(\bar{\rho}, \bar{\eta})$ plane. The shift of $\bar{\rho}$ is due to the contribution from the charm quark, which would not enter the unitarity triangle otherwise. The neutral channel, having a purely imaginary amplitude, allows direct extraction of $\bar{\eta}$. Together, a simultaneous measurement of the two BRs constrain the unitarity triangle as shown in figure 1.4.

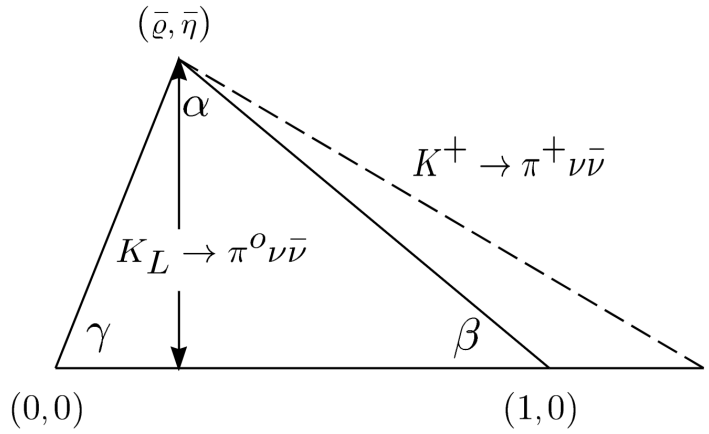

Figure 1.4: Constraints of the unitarity triangle from K → πνν̄ BRs. [50]

1.2.2 K → πνν̄ beyond the Standard Model

It can be shown in a nearly model-independent way (that is, under very mild assumptions, most importantly lepton flavour conservation) that [51, 27]

$$\frac{\mathrm{BR}(\mathrm{K_L} \to \pi^0 \nu\bar{\nu})}{\mathrm{BR}(\mathrm{K^+} \to \pi^+ \nu\bar{\nu})} \lesssim 4.3\,. \tag{1.4}$$

This is because, if lepton flavour is conserved, the neutral mode proceeds mainly via CP violating interference between mixing and decay, and the relative phase between the amplitudes of the mixing and of the decay is subject to trivial trigonometric constraints; (1.4) is obtained by relating between this relative phase and the ratio of BRs through known isospin corrections and the lifetimes of the kaons. (1.4) is known as the *Grossman-Nir (GN) bound*.

$\mathrm{BR}(\mathrm{K_L} \to \pi^0 \nu\bar{\nu})$ and $\mathrm{BR}(\mathrm{K^+} \to \pi^+ \nu\bar{\nu})$ are correlated by the short distance dynamics of $\mathrm{s} \to \mathrm{d}\nu\bar{\nu}$ [52, 53]. Figure 1.5 illustrates the common correlations from different classes of BSM models where a new heavy gauge boson Z' generates FCNC processes at tree level. Depending on the chiral structure of the couplings, these models can be subdivided into three different classes: models with *constrained Minimal Flavour Violation* (MFV), where a CKM-like structure of flavour interactions is present (green region in figure 1.5); models where either left-handed or right-handed currents dominate, in which K → πνν̄ would be strongly correlated to ϵ_K, which gives

the strong constraint shown by the blue region in figure 1.5; models where left-handed and right-handed currents have similar contributions, in which the correlation between ϵ_K and $K \to \pi\nu\bar{\nu}$ is absent, in which in principle the two BRs can assume any value (expanding red region in figure 1.5).

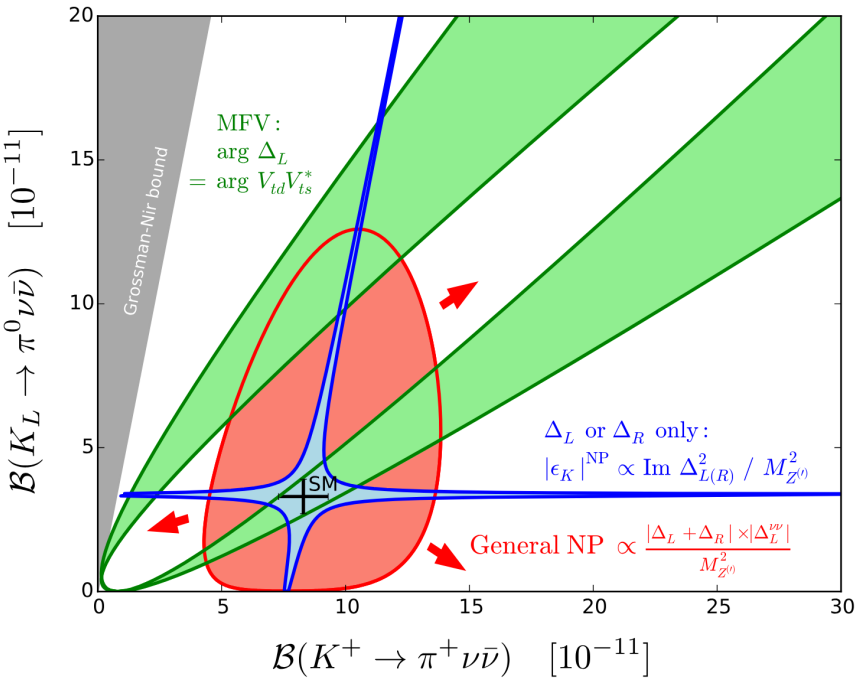

Figure 1.5: Common correlations between $BR(K_L \to \pi^0\nu\bar{\nu})$ and $BR(K^+ \to \pi^+\nu\bar{\nu})$ in BSM scenarios with a Z′-like particle. Δ_L and Δ_R are related to the left-handed and right-handed couplings of Z′. [52]

BSM contributions to $K \to \pi\nu\bar{\nu}$ are also predicted in models including lepton number violation (LNV), such as models predicting Majorana neutrinos [54, 55], and models including leptoquarks [56, 57, 58].

Correlations between B-meson physics (and in particular $B \to K^{(*)}\nu\bar{\nu}$ decays) and $K^+ \to \pi^+\nu\bar{\nu}$ can be found in models where lepton flavour universality (LFU) is violated and new particles are only coupled to the third generations of quarks and leptons [59]. Correlations between these decays are also predicted by models including MFV [60] and LNV [61]. A deviation from the SM prediction of $K \to \pi\nu\bar{\nu}$ can be reinterpreted as LFU violation also for the couplings to the first and second generations of leptons [31].

1.3 Experimental status

Both the charged and the neutral channels of $K \to \pi\nu\bar{\nu}$ are extremely challenging from the experimental point of view. The experimental signature of the charged channel is an incoming kaon and an outgoing pion with missing

energy; the neutral channel is similar, with the additional difficulty of the impossibility to track the particles. The rarity of the processes makes the experiments even harder, as the signal, being rare, must be discriminated also against rare backgrounds. Moreover, high intensity kaon beams are needed in order to collect the necessary statistics, but the high intensity contributes to accidental hits in the detector, which lower the signal acceptance and increase accidental-related background events.

1.3.1 $K^+ \to \pi^+ \nu\bar{\nu}$

The KEK E10 experiment in the late 1970s employed a K^+ stopping target, in order to observe the decay in the kaon center of mass frame. The secondary charged particle was measured by multi-wire proportional chambers and range counters in a hemisphere around the stopping target; analog signals from the range counters were also recorded in order to identify π^+ through the $\pi^+ \to \mu^+ \to e^+$ decay chain. No events were observed, and an upper limit was set to $BR(K^+ \to \pi^+ \nu\bar{\nu}) < 1.4 \times 10^{-7}$ at 90% CL [62].

A similar but improved technique was used by the BNL E787 and E949 experiments between the 1990s and the 2000s. They also used stopped kaons, but they had two stages of electromagnetic separators in the beam which provided a high K^+/π^+ ratio. As in KEK E10, pions were identified via the $\pi^+ \to \mu^+ \to e^+$ decay chain with a waveform analysis. The coverage was increased by using cylindrical range counters and hermetic photon veto counters. A magnetic spectrometer was used in order to measure the pion momentum. E949 was terminated before completion; nonetheless, the BNL experiments, after a blind analysis, observed a total of 7 candidate events (figure 1.6), and measured [63]

$$BR(K^+ \to \pi^+ \nu\bar{\nu}) = (17.3^{+11.5}_{-10.5}) \times 10^{-11} ,$$

which, although with a large uncertainty, does not show deviations from the Standard Model.

The NA62 experiment was proposed in 2005 [28] with the main goal of measuring $BR(K^+ \to \pi^+ \nu\bar{\nu})$ with a precision comparable to the theoretical prediction, by employing a decay-in-flight technique. This technique allows removal of most of the material from the beamline, reducing the rate from accidental pions; muons can be vetoed with an iron shield, and there is no need to wait for the $\pi^+ \to \mu^+ \to e^+$ decay chain, which would cause issues in a high rate environment. The NA62 experimental strategy and detector is described in detail in the next chapter; here the results of the experiment are summarised.

The NA62 experiment collected data in its first run of data taking (Run1) between 2016 and 2018. Results of the analysis of 2016 data [64], 2017 data [65], and 2018 data [34] are summarized in table 1.1. The selected events are described in terms of $m^2_{miss} = (P_K - P_\pi)^2$ and pion momentum,

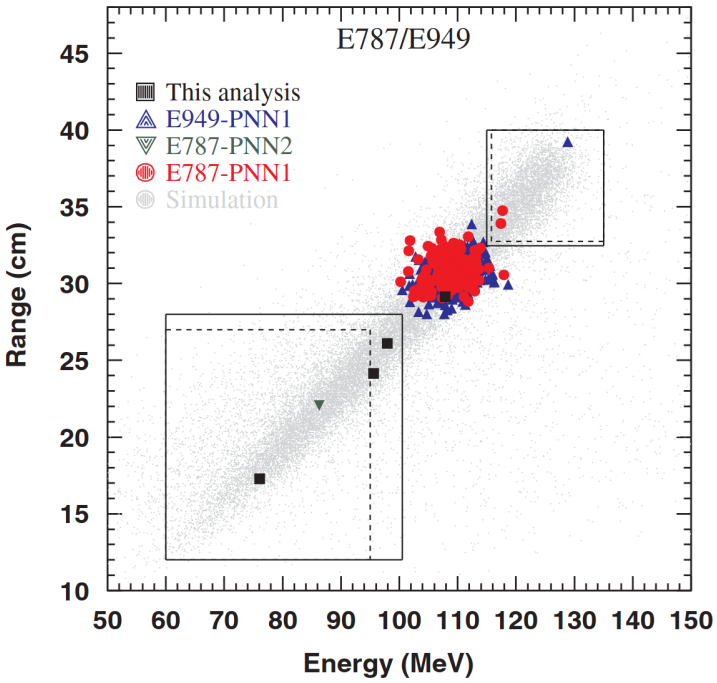

Figure 1.6: Pion range vs kinetic energy of events observed in BNL E787 and E949. Solid (dashed) lines represent signal regions for the E949 (E787) analyses. [63]

where P_{K} and P_{π} are the 4-momenta of the K^+ and of the π^+, respectively. These variables, as discussed in section 2.1, provide kinematic background rejection; their distribution for the selected events are shown in figures 1.7–1.9, where red lines represent signal regions. In total, 20 candidate events were observed, including an estimated background of $7.03^{+1.05}_{-0.82}$, which led to the measurement [34]

$$\mathrm{BR}(\mathrm{K}^+ \to \pi^+ \nu\bar{\nu}) = \left(10.6^{+4.0}_{-3.4}|_{\mathrm{stat}} \pm 0.9_{\mathrm{syst}}\right) \times 10^{-11} \text{ at } 68\% \text{ CL}\,,$$

which is the most precise result so far (figure 1.10).

During CERN Long Shutdown 2 (LS2), between 2018 and 2021, NA62 implemented various hardware improvements. Its second run of data taking (Run2) started in 2021 and is foreseen to last until Long Shutdown 3 [76]. The analysis of the data collected in 2021 and 2022 is the main topic of this thesis. Reports of the data taking in 2021–2024 can be found in [77, 78, 79, 80]. The analysis of the whole Run2 data will allow for a measurement of $\mathrm{BR}(\mathrm{K}^+ \to \pi^+ \nu\bar{\nu})$ with an uncertainty below 20% [81].

1.3.2 $\mathrm{K_L} \to \pi^0 \nu\bar{\nu}$

The neutral channel is sometimes referred to as the *nothing to nothing decay*, as neither the $\mathrm{K_L}$ nor the final state can be directly measured: the π^0 in the final state can only be indirectly detected from its decay products.

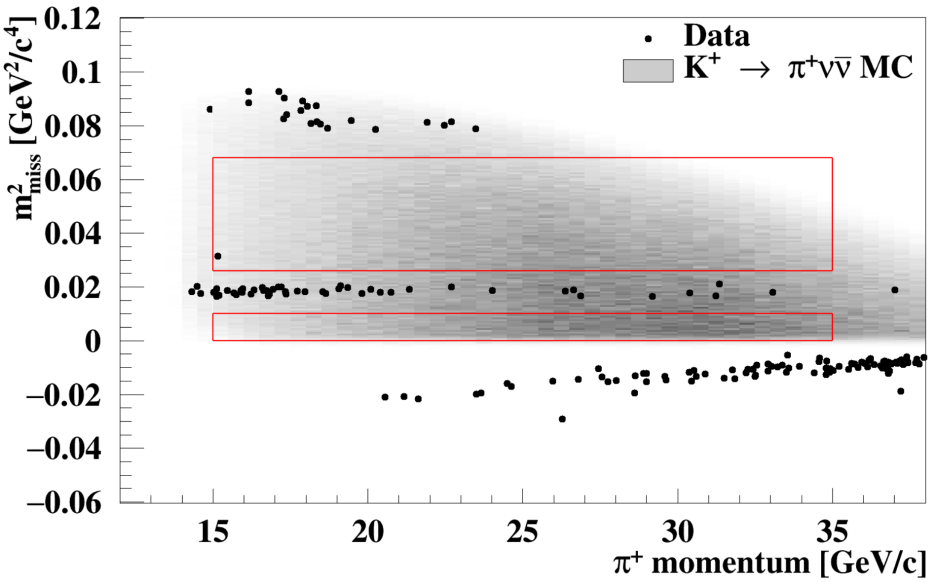

Figure 1.7: Selected events in NA62 2016 data. [64]

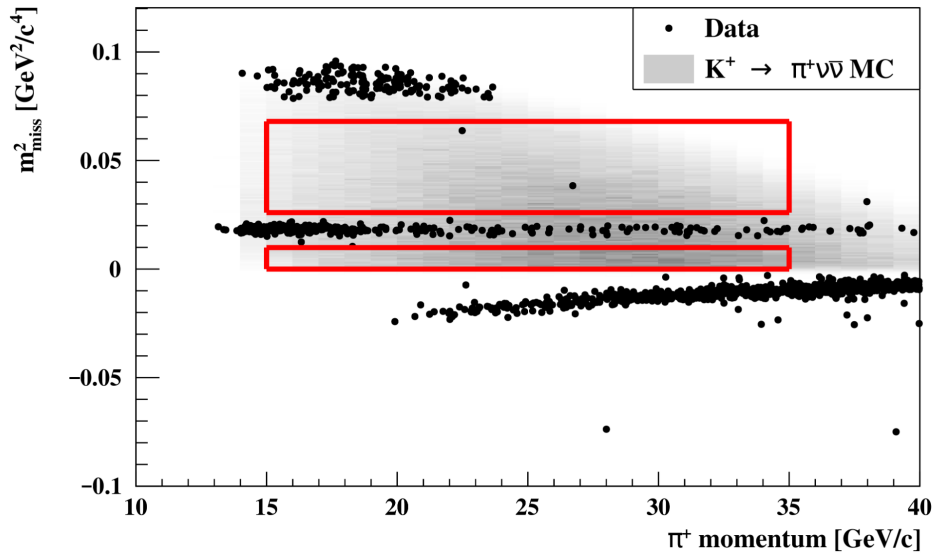

Figure 1.8: Selected events in NA62 2017 data. [65]

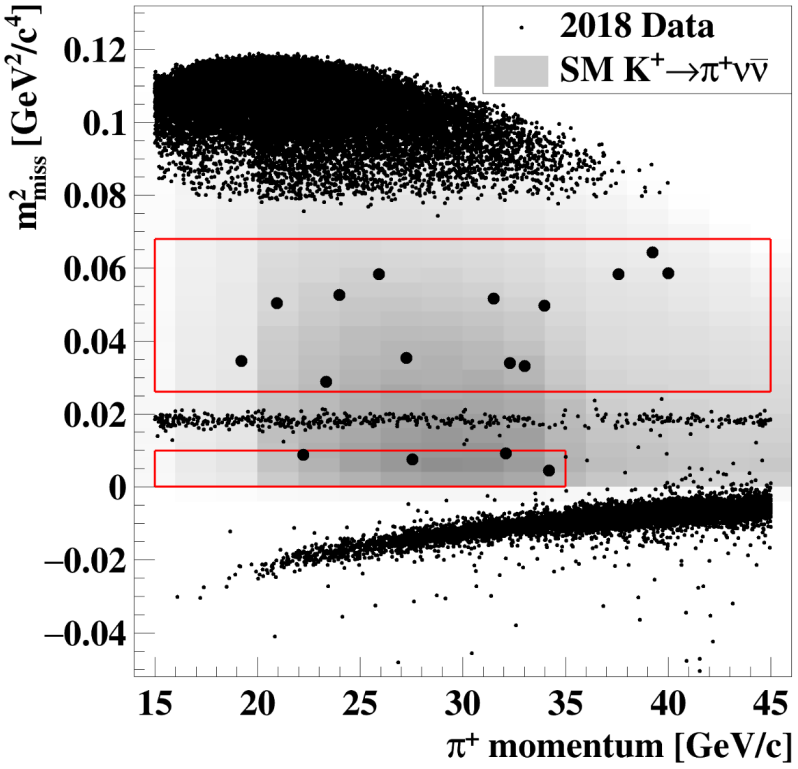

Figure 1.9: Selected events in NA62 2018 data. [34]

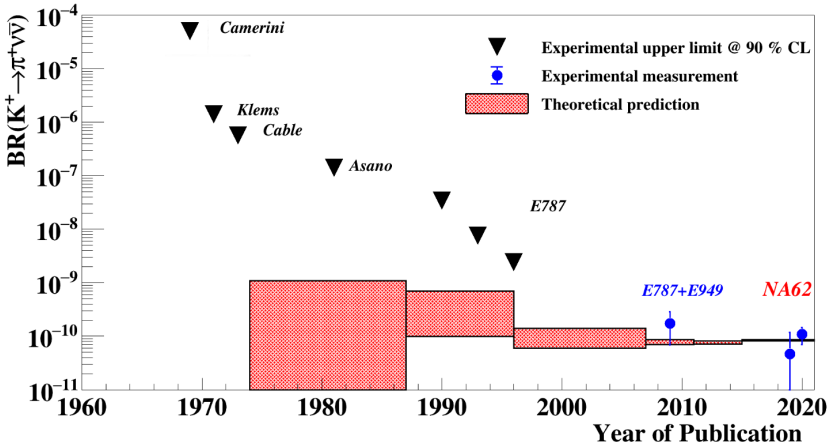

Figure 1.10: Publication history for BR($K^+ \to \pi^+ \nu\bar{\nu}$). [66, 67, 68, 69, 62, 70, 71, 72, 73, 74, 47, 63, 75, 29, 65, 34]

Data	N_b	$N^{\text{exp}}_{\pi\nu\bar{\nu}}$	N
2016 [64]	$0.152^{+0.092}_{-0.033}\vert_{\text{stat}} \pm 0.013_{\text{syst}}$	$0.267 \pm 0.001_{\text{stat}} \pm 0.020_{\text{syst}} \pm 0.032_{\text{ext}}$	1
2017 [65]	$1.46 \pm 0.25_{\text{stat}} \pm 0.21_{\text{syst}}$	$2.16 \pm 0.13_{\text{syst}} \pm 0.26_{\text{ext}}$	2
2018 (S1) [34]	$1.11^{+0.40}_{-0.22}$	$1.56 \pm 0.10_{\text{syst}} \pm 0.19_{\text{ext}}$	2
2018 (S2) [34]	$4.31^{+0.91}_{-0.72}$	$6.02 \pm 0.39_{\text{syst}} \pm 0.72_{\text{ext}}$	15
Run1 [34]	$7.03^{+1.05}_{-0.82}$	$10.01 \pm 0.42_{\text{syst}} \pm 1.19_{\text{ext}}$	20

Table 1.1: Summary of NA62 results. S1 and S2 are two subsets of 2018 data, see [34] for more details. N_b is the number of background events expected, $N^{\text{exp}}_{\pi\nu\bar{\nu}}$ is the number of SM signal events expected, N is the number of events observed. External uncertainties in $N^{\text{exp}}_{\pi\nu\bar{\nu}}$ are due to the SM BR prediction, and do not enter the BR measurement.

The first dedicated experiment to search for this decay was KEK E391a, which employed an electromagnetic calorimeter covering the decay region, and a hermetic photon veto system to suppress $K_L \to \pi^0\pi^0$. With data collected in six months of 2005, and no event observed, they set the upper limit $BR(K_L \to \pi^0\nu\bar{\nu}) < 2.6 \times 10^{-8}$ at 90% CL [82].

The KOTO experiment at J-PARC is the successor of E391a, using a similar technique with improved hardware, and operating since 2013 with the goal of observing this decay. Data collected in 2015, in which no event was observed, led to the upper limit $BR(K_L \to \pi^0\nu\bar{\nu}) < 3.0 \times 10^{-9}$ at 90% CL [83]. In 2021, KOTO published the results of the analysis of 2016–2018 data, where 3 candidate events were observed: after unveiling them, backgrounds from $K^+ \to \pi^0 e^+ \nu$ and $K_L \to \gamma\gamma$ (from scattered K_L) were studied, which led to the background estimation of 1.22 ± 0.26. With this dataset, conservatively considering the observed events as background, they set the upper limit $BR(K_L \to \pi^0\nu\bar{\nu}) < 4.9 \times 10^{-9}$ at 90% CL [84], less stringent than the previous. Efforts, both in hardware and in software development, were then focused in the rejection of these two backgrounds: after the very recently published analysis of 2021 data, no event was observed in this dataset, and the upper limit was set to $BR(K_L \to \pi^0\nu\bar{\nu}) < 2.2 \times 10^{-9}$ at 90% CL, the most stringent so far [85].

1.3.3 Global picture

Figure 1.11 summarizes the current theoretical and experimental status for both the charged and the neutral $K \to \pi\nu\bar{\nu}$ modes. The yellow and green bands show the 68 % CL intervals for the measurements by the BNL experiments and by NA62, respectively, while the orange band shows the 90% CL upper limit from KOTO, which still lies in the zone excluded by the GN

bound. More precision from the experiments is needed, as there is still much room for potential discoveries.

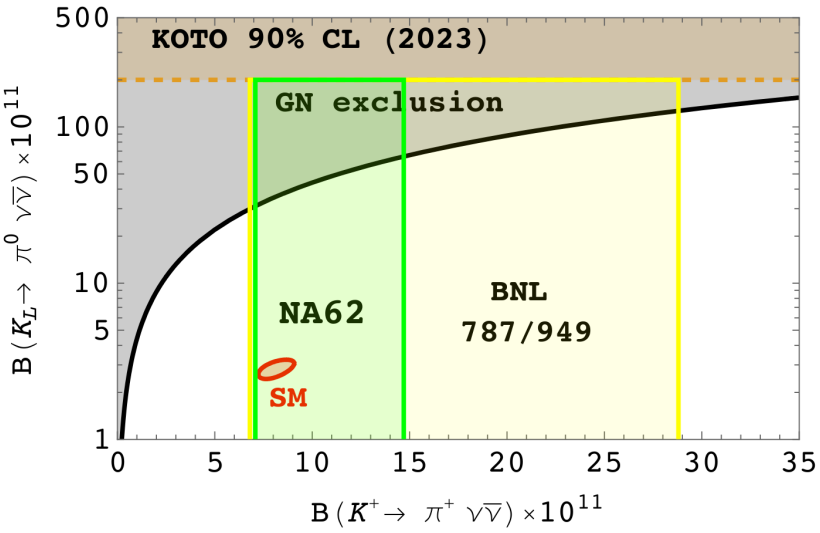

Figure 1.11: Current status for K $\to \pi\nu\bar{\nu}$. [27]

These results can be interpreted as constraints for new physics. Figure 1.12 shows the current upper limits for a two body mode K $\to \pi$X where X is a stable invisible particle. Figure 1.13 shows the current constraints for LFU violation in terms of the BSM contributions to the coefficients δC_L^ℓ ($\ell = \mathrm{e}, \mu, \tau$) of the effective operators responsible for the s $\to$ d$\nu_\ell\bar{\nu}_\ell$ process [31].

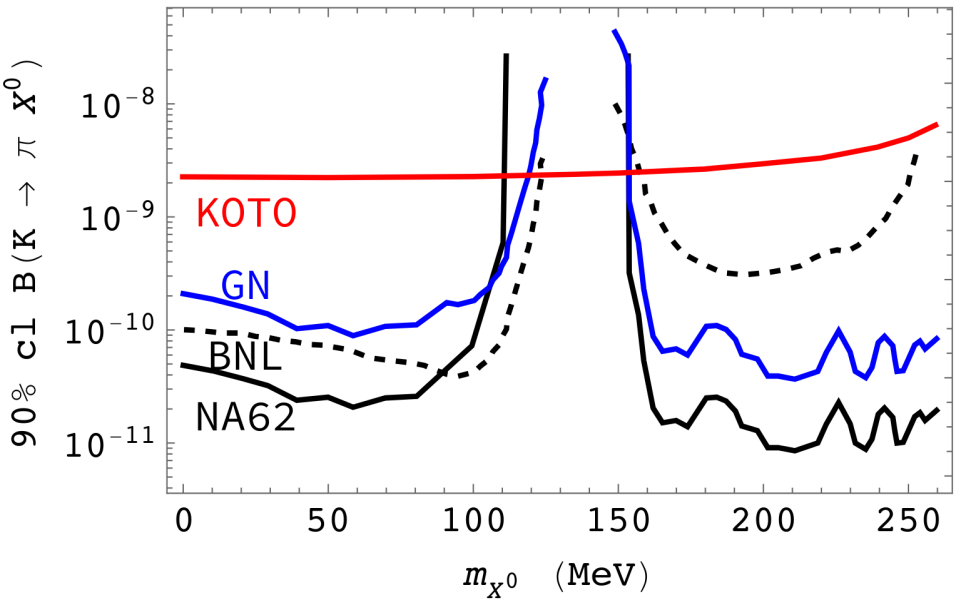

Figure 1.12: Current status for K $\to \pi$X. The solid (dashed) black lines show the exclusion for $K^+ \to \pi^+$X set by NA62 [86] (BNL E787/E949 [63]); the blue line shows the implied limit for $K_L \to \pi^0$X based on the GN bound; the red line shows the exclusion for $K_L \to \pi^0$X set by KOTO [83]. [27]

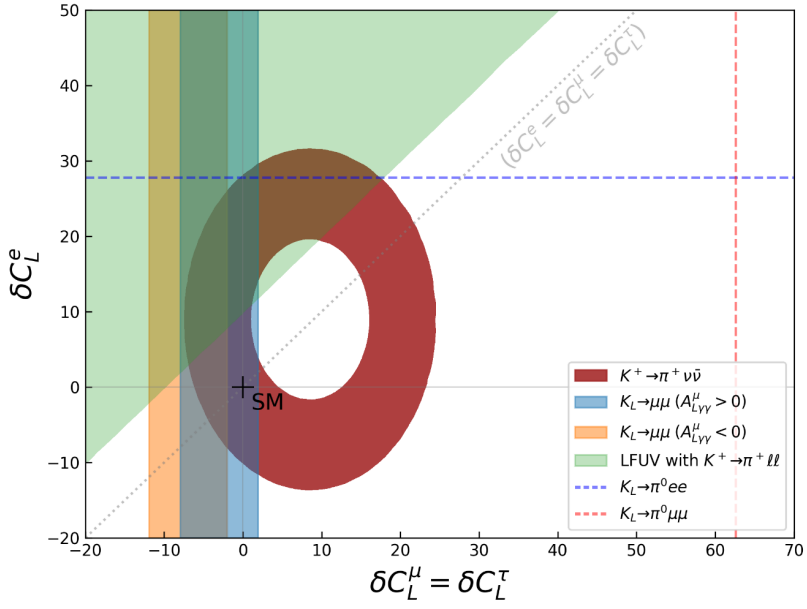

Figure 1.13: Status of LFU violation parameters contraints. [31]

Chapter 2

The NA62 experiment

The NA62 experiment is a fixed-target experiment at the CERN SPS accelerator. It was proposed and designed to measure BR($K^+ \to \pi^+ \nu\bar{\nu}$) with a decay-in-flight technique.

The significant experimental challenges presented by this measurement are the primary reason for the exceptional performance of the NA62 detector. Such a detector can therefore be exploited for the study of a much broader kaon physics programme: NA62 results include precision measurements [87, 88, 89, 90], searches for lepton number and lepton flavour violating decays [91, 92, 93, 94, 95], and searches for production of invisible particles [96, 97, 98, 86, 99]. A *beam dump* configuration for data taking can also be used in order to perform direct searches for new particles [100, 101].

An overview of the experimental strategy employed by NA62 is given in section 2.1. Section 2.2 gives an outline of the NA62 apparatus, and all the following sections are dedicated to the details of the single sub-systems.

2.1 Experimental strategy

The $K^+ \to \pi^+ \nu\bar{\nu}$ decay has a very weak experimental signature: an incoming kaon, an outgoing pion, and missing energy (the neutrino pair escapes undetected). Because of such signature weakness and because of the tiny BR of the signal, many of the kaon decay modes constitute significant sources of background. Table 2.1 summarizes the most important K^+ decays, against which the signal must be discriminated.

The search for $K^+ \to \pi^+ \nu\bar{\nu}$ is a significant experimental challenge because of many different factors. A high intensity beam must be used in order to collect a sufficient number of K^+ decays, i.e. $O(10^{12})$, in a reasonable amount of time. Kaons in the initial state must be correctly identified and matched to the final state, in order to discriminate against events with accidental pions. Charged pions in the final state must be correctly discriminated against muons, in order to suppress backgrounds from K^+ decays to muons, most notably $K^+ \to \mu^+ \nu$. Finally, any particle in the final state other than

Decay mode	Branching ratio [27]
$K^+ \to \mu^+\nu$	$(63.56 \pm 0.11) \times 10^{-2}$
$K^+ \to \pi^+\pi^0$	$(20.67 \pm 0.08) \times 10^{-2}$
$K^+ \to \pi^+\pi^+\pi^-$	$(5.583 \pm 0.024) \times 10^{-2}$
$K^+ \to \pi^+\pi^-e^+\nu$	$(4.247 \pm 0.024) \times 10^{-5}$

Table 2.1: Main K^+ decay backgrounds for the search of $K^+ \to \pi^+\nu\bar{\nu}$.

the π^+ must be efficiently vetoed, in order to suppress backgrounds from K^+ decays to multiple visible particles; most notably, photons from π^0 decays in $K^+ \to \pi^+\pi^0$ and from radiative decays, and additional charged particles such as the ones arising from $K^+ \to \pi^+\pi^+\pi^-$ and $K^+ \to \pi^+\pi^-e^+\nu$.

A key ingredient in the rejection of K^+ decay backgrounds is kinematics. In fact, a powerful variable is

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

where P_K is the 4-momentum of the kaon, and P_π is the 4-momentum of the reconstructed secondary charged particle in the π^+ mass hypothesis. Apart from resolution effects, this quantity is: equal to the invariant mass of the neutrino pair for the signal, therefore positive and broadly distributed; negative for $K^+ \to \mu^+\nu$; equal to $m^2_{\pi^0}$ for $K^+ \to \pi^+\pi^0$; larger than $(2m_\pi)^2$ for $K^+ \to 3\pi$. Figure 2.1 shows that, for a suitable choice of the signal region, the main backgrounds can be suppressed by a factor $O(10^4)$ using kinematics only. The detector must therefore be also able to precisely measure both the magnitude and the direction of the K^+ and of the π^+: kinematic resolution effects, especially from the non-gaussian tails, cause background to enter the signal regions.

2.2 The NA62 detector

A description of the detector can be found in [102, 103]. An overview of the NA62 apparatus, in the configuration used for the Run1 data taking, is given in section 2.2.1, and an outline of the improvements made to the apparatus for the Run2 data taking is given in section 2.2.2.

2.2.1 Run1 configuration

A schematic top view of the apparatus, in the configuration used at the end of Run1, is shown in figure 2.2.

The kaon production target is the origin of the NA62 frame of reference; the beam propagates in the positive Z axis, while the Y axis is vertical, posi-

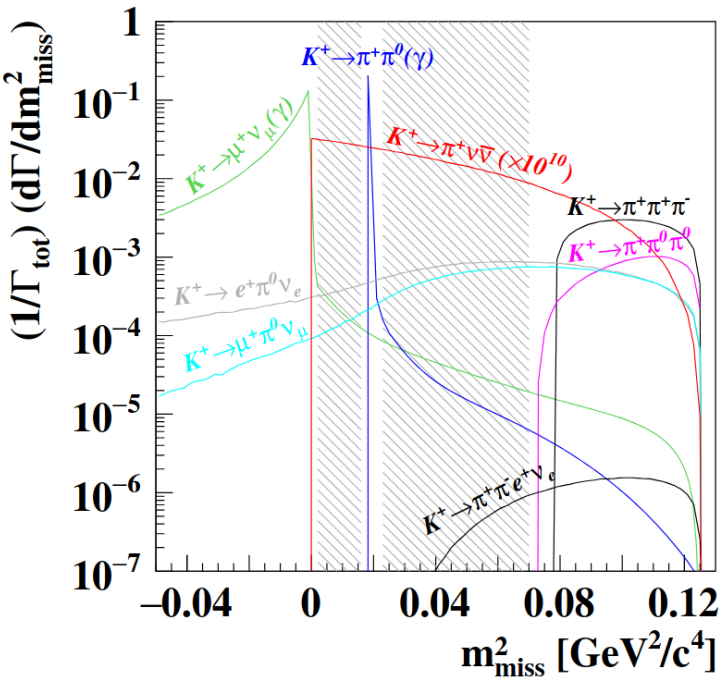

Figure 2.1: Differential decay rate, as function of m^2_{miss}, for $K^+ \to \pi^+\nu\bar{\nu}$ (enhanced by 10^{10} for visibility) and for the main K^+ decays. Hatched areas include the NA62 signal regions.

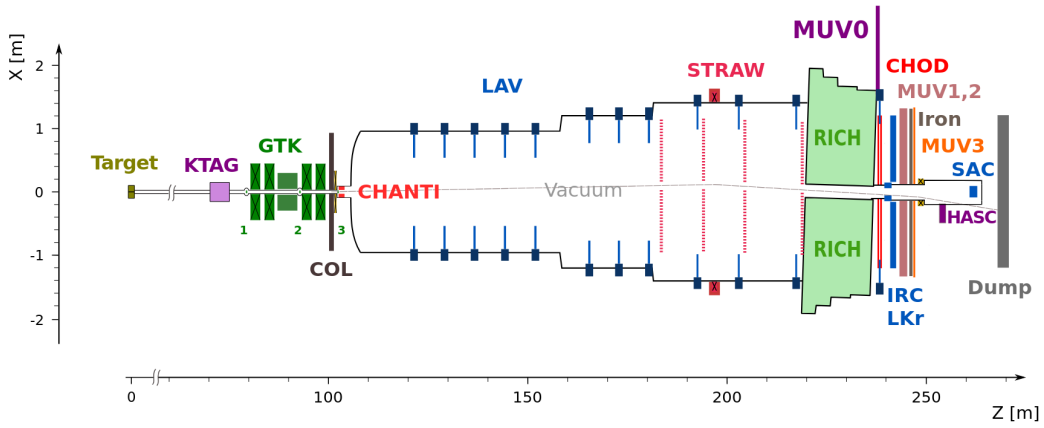

Figure 2.2: Schematic top view of the apparatus at the end of Run1. Boxes with crosses denote dipole magnets. The dashed line shows the trajectory of an undecayed beam particle. [34]

tive upwards; the X axis is defined so that the reference system is orthogonal and right-handed.

A 400 GeV/c proton beam, extracted from the CERN SPS, impinges on a beryllium target, in spills of 4.8 s effective duration (also known as *bursts*). Typical intensities during the Run1 data taking were around 1.7 to 1.9 ×

10^{12} protons per pulse (ppp), corresponding to around 60 % of the nominal intensity. A secondary unseparated hadron beam, with 75 GeV/c momentum, is extracted. It is composed of about 70 % π^+, 23 % p, 6 % K^+, < 1 % μ^+.

The beam then passes through the *kaon tagger* (KTAG), a differential Cherenkov detector for kaon identification, and the *Gigatracker* (GTK), a silicon-based beam spectrometer, composed of three stations (GTK1, GTK2, GTK3). Immediately before GTK3, a final collimator (COL, or TCX) is placed; after it, the *charged anti-coincidence* detector (CHANTI), a plastic scintillator active veto, detects inelastic interactions in GTK3. The typical beam particle rate at GTK3 in 2018 data is around 400 MHz. GTK3 marks the beginning of a 117 m-long vacuum tank: the latter contains the *fiducial decay region*, or *fiducial volume*, defined by 110 m < Z < 170 m, in which 11 % of kaons decay.

Downstream of the fiducial volume, a magnetic spectrometer (STRAW) measures the momentum and direction of charged decay products; a ring-imaging Cherenkov detector (RICH) identifies them; two scintillator hodoscopes (CHOD and NA48-CHOD) provide timing information, and are also used in the trigger. A liquid krypton electromagnetic calorimeter (LKr) complements the RICH for particle identification and serves as main photon veto.

Other than the LKr, photon veto is also provided by the *large-angle veto* system (LAV), comprised of 12 stations of ring-shaped lead-glass calorimeters, and the *small-angle veto* system (SAV), which, in turn, is made of the *intermediate-ring calorimeter* (IRC), placed in front of the LKr, and the *small angle calorimeter* (SAC). The photon veto systems guarantee a hermetic coverage for photons emitted in angles from 0 up to 50 mrad.

The *muon veto system* (MUV) is made of two iron-scintillator hadronic sampling calorimeters (MUV1 and MUV2), and an array of scintillator tiles (MUV3) placed behind 80 cm of iron. Other veto detectors (MUV0, HASC) are installed at optimized locations in order to provide additional coverage for charged particles produced in background kaon decays.

The LKr, MUV1 and MUV2 detectors are read out with flash ADCs; the GTK and the STRAW have specialized TDC boards; all other detectors are read out with TDCs, mounted on custom-made boards (TEL62), which both read out data and provide trigger information. The *level-0 trigger processor* (L0TP) combines logical signals (primitives) from the RICH, CHOD, NA48-CHOD, LKr, LAV and MUV3 detectors into a trigger decision, which is dispatched to all the boards for data readout. An additional *level-1* software trigger (L1) exploits reconstruction algorithms from the KTAG, LAV and STRAW detectors to further reduce the data rate before storing events on disk.

2.2.2 Run2 improvements

The NA62 Run1 result has proven that the decay-in-flight technique and the NA62 apparatus are suitable for the measurement of BR($K^+ \to \pi^+\nu\bar{\nu}$). However, the analysis of the Run1 data employed tight cuts in order to reject the backgrounds. In particular, the largest background, by far, was the *upstream background* (table 2.2), which is due to pions produced upstream of the fiducial volume, either from kaon decays or from scattering against the beamline material. The second largest background was $K^+ \to \pi^+\pi^0$. During LS2, improvements were made in the apparatus to have stronger rejection power against these backgrounds [104].

Background	Estimated number of events
Upstream	$2.76^{+0.90}_{-0.70}$
$K^+ \to \pi^+\pi^0$	0.52 ± 0.05
$K^+ \to \mu^+\nu$	0.45 ± 0.06
$K^+ \to \pi^+\pi^-e^+\nu$	0.41 ± 0.10
$K^+ \to \pi^+\pi^+\pi^-$	0.17 ± 0.08
Total	$4.31^{+0.91}_{-0.72}$

Table 2.2: Backgrounds in 2018 S2, the largest subsample of Run1 data. [34]

A new beamline design [77] was implemented in order to be more resilient to the upstream background (figure 2.3): it is described in the following.

A new GTK station (GTK0) was placed immediately upstream of GTK1: the addition of a station of the beam spectrometer helps in the reconstruction of the upstream part of the events. In particular, it improves the efficiency of reconstruction of the K^+ track and the detection and rejection of pileup tracks.

GTK2 was moved before the scraper, whereas in Run1 it was placed after the scraper. This protects from inelastic interactions of the beam with GTK2, as now the products of such interactions will be absorbed by the scraper; in the previous position, they had a chance of being transmitted towards the fiducial volume.

A new detector (VetoCounter) has been added around the beam pipe, before and after the TCX collimator. This detector is designed to actively veto kaon decays happening upstream of the collimator, by detecting decay products which would otherwise be absorbed in the collimator and remain invisible.

A new hodoscope (ANTI0) has been added at the very entrance of the vacuum tank [105]. This detector, installed mainly to reject the muon halo

in data taken in beam dump configuration, actually complements the Veto-Counter in detecting particles down to about 1 m from the beam line.

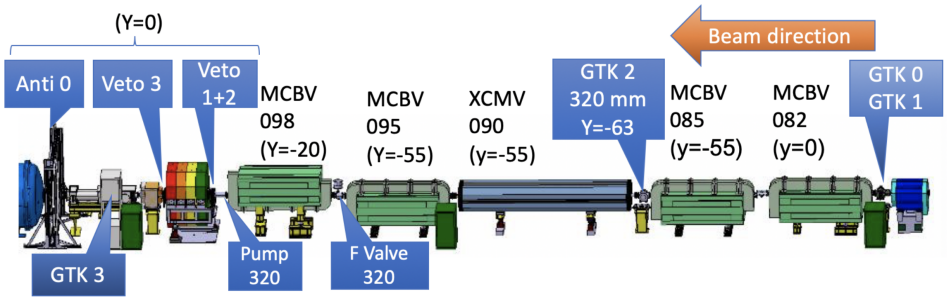

Figure 2.3: Schematic layout of the new achromat and beam line. [77]

In order to reject $K^+ \to \pi^+\pi^0$ more strongly, a new HASC station has been added, which helps in detecting the products of photon interactions with the beam pipe at the RICH.

A sketch of the detector as designed for the Run2 data-taking, including all of the above improvements, is shown in Figure 2.4.

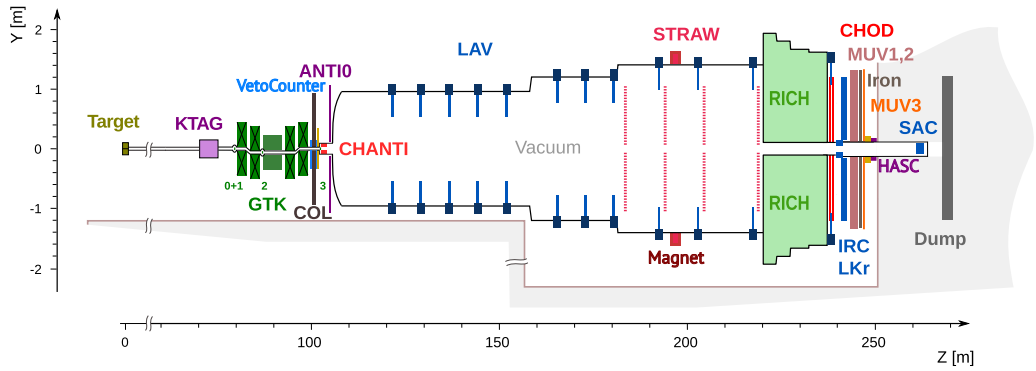

Figure 2.4: Schematic side view of the apparatus during Run2. [36]

Furthermore, after a brief high-intensity test data taking period in 2018, the trigger and data acquisition underwent a small number of fixes and improvements. The system was then deemed suitable to be used at 100 % nominal intensity [106], corresponding to about 3×10^{12} ppp and 600 MHz of beam rate on the GTK (figure 2.5).

2.3 Beam line

The primary proton beam is extracted from the CERN SPS accelerator at 400 GeV/c, and directed towards a 400 mm long, 2 mm diameter beryllium target (T10) via the P42 beam line. A secondary, high-intensity hadron beam

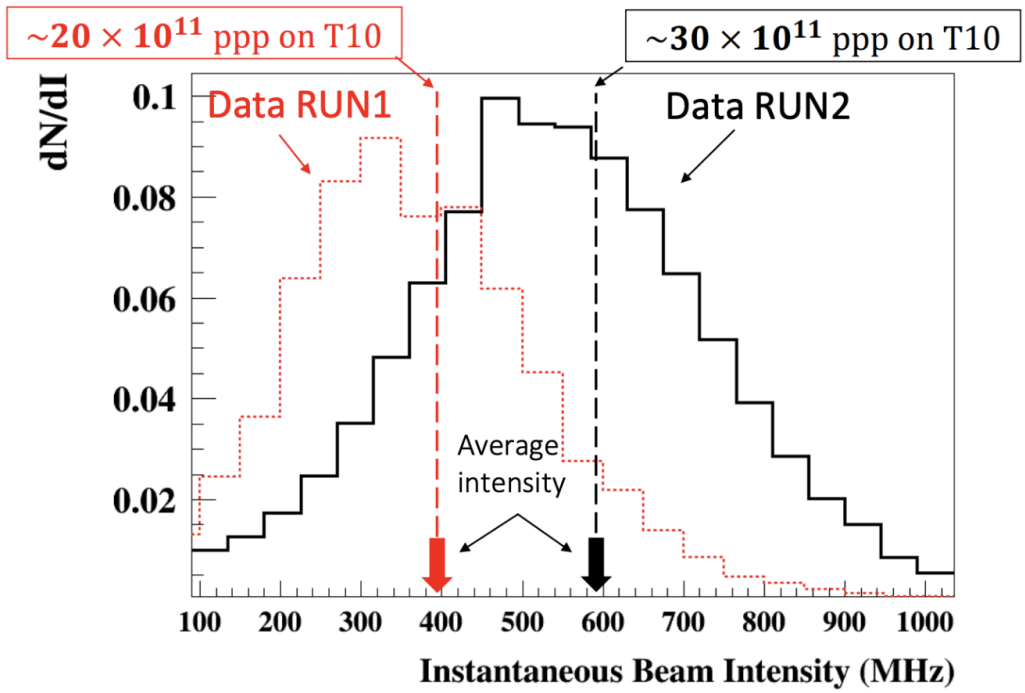

Figure 2.5: Distribution of instantaneous intensity as measured by the GTK in Run1 and Run2. [35]

line (K12), of which a schematic layout is shown in figure 2.6, follows the T10 target at a 0 production angle.

The T10 target is immediately followed by a 950 mm copper collimator with selectable hole size, which absorbs particles produced at extreme angles. A triplet of quadrupole magnets (Q1, Q2, Q3) and an achromat (A1) focus and select beam particles with a central momentum is 75 GeV/c, chosen to maximize the fraction of kaons (about 6 %), and a 1 % rms momentum bite. The achromat consists by two pairs of vertically-deflecting dipole magnets: the first pair displaces the beam by 110 mm, while the second returns the beam on the original axis; between the two, a pair of motorized beam dump units (TAX1, TAX2) absorbs unwanted secondary particles. The positions of the two TAXs can be adjusted in order to change data taking conditions, in particular obtaining a muon beam, used for detector calibration, or to dump the beam completely and perform direct searches for exotic particles.

Between TAX1 and TAX2 there is a radiator, made of tungsten plates of thickness up to 5 mm, which causes electrons to lose energy so that they can be removed from the beam by the second achromat (A2) while minimizing hadronic scattering.

A series of quadrupoles (Q4–Q10) and collimators (C1–C7) perform further focussing and selection of the beam.

The beam line also features a muon sweeper, composed of three dipole magnets (B3) and iron plates inserted between them, with a 40 mm diamater

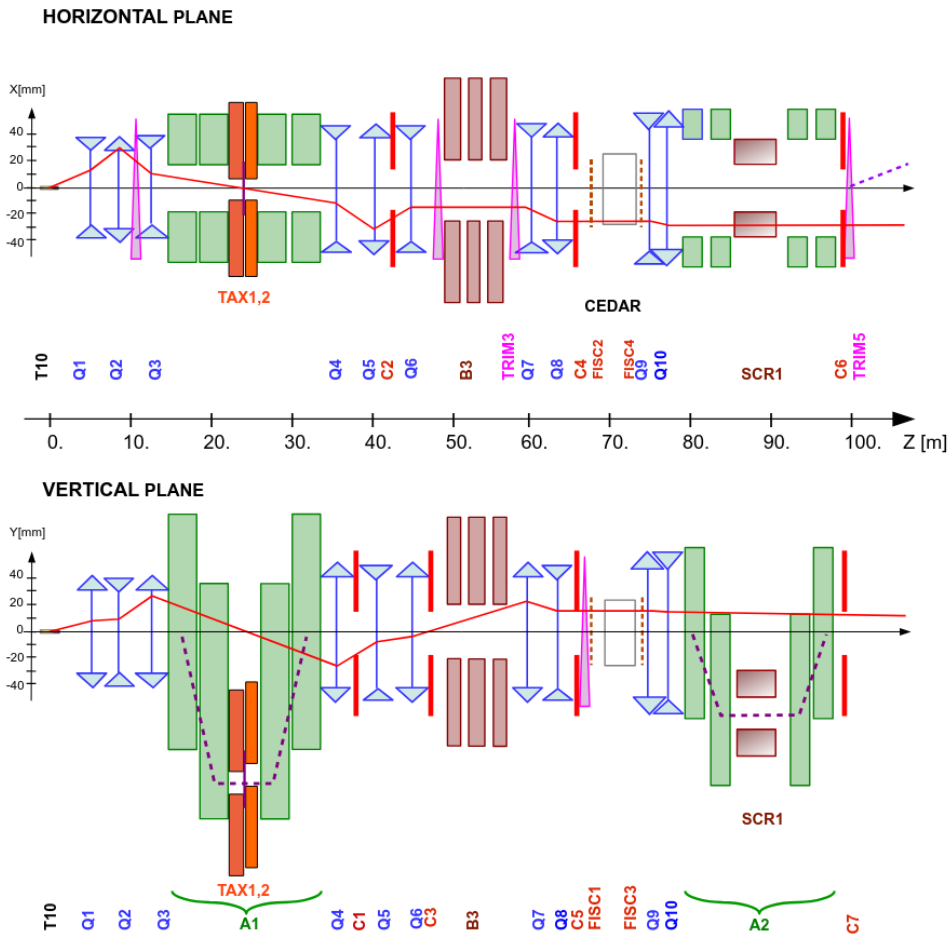

Figure 2.6: Schematic layout of the NA62 beamline as of 2017. The TCX collimator, added in June 2018, is missing. Solid lines indicate the trajectory of a particle leaving the target at 75 GeV/c and at the pictured angle. Dashed lines indicate the trajectory of an initially on-axis particle. [103]

bore: the vertical magnetic field in the iron surrounding the beam removes off-axis muons of both signs.

Inside the achromat A2, the GTK is placed: at GTK2, particles of nominal momentum are displaced vertically by 60 mm; the measurement of the vertical displacement allows for measurement of the momentum of the particles. Between the second and third magnets of the achromat, a toroidal magnetized iron collimator (SCR1) is placed: it defocusses muons so that the following two magnets of the achromat removes them from the beam. The TCX collimator is located immediately upstream of GTK3.

The achromat was optimized before the Run2 data taking. In particular, in Run1 GTK2 was placed immediately after SCR1; during LS2, the positions of GTK2 and SCR1 were exchanged so that the scraper could absorb products of inelastic interactions of the beam with GTK2.

A steering magnet (TRIM5) is placed before the entrance to the fiducial

volume, so that the beam is deflected in the positive X direction by 1.2 mrad: this angle is used to compensate for the spectrometer magnet downstream of the fiducial volume (MNP33) which deflects the beam in the opposite direction by 3.6 mrad, so that the beam is directed back through the central aperture of the LKr, which is on the Z axis. Downstream of MUV3, a bend magnet steers the beam away from the acceptance of SAC and towards the beam dump. Figure 2.7 shows the beam line layout through the decay region.

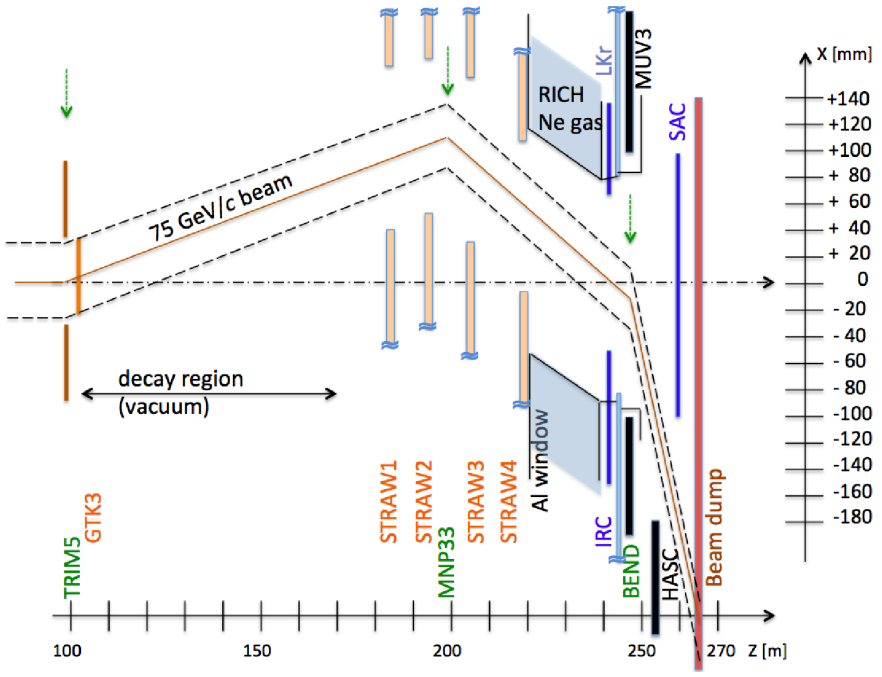

Figure 2.7: Beam line through the decay region, as of Run1. In Run2, the HASC has been duplicated in a symmetrical position with respect to the beam axis. Dashed lines show the 2σ width of the beam profile; vertical arrows indicate the bending center of magnets. [103]

2.4 Kaon tagger (KTAG)

The KTAG [103] is a differential Cherenkov detector, designed to identify kaons in the unseparated hadron beam and measure their time with a resolution of around 70 ps.

The detector is composed of a 5 m long vessel (CEDAR) filled with nitrogen at 1.7 bar at ambient temperature, and a readout system made of 8 sectors, referred to as *octants*, of 48 photomultipliers (PMs) each (Hamamatsu™ R9880 and R7400). The CEDAR was developed in the late 1970s to discriminate protons, kaons and pions in unseparated beams extracted from CERN SPS [107]; its readout has been developed specifically for NA62 in order to meet the high intensity requirements [108]. The excellent time res-

olution is achieved thanks to the very good transit time spread of the PMs used (typically 200 ps) and thanks to NINO chips [109] which allow measurement of leading and trailing edges of the signal and therefore implementation of offline slewing corrections.

Before reaching the readout system, the Cherenkov light passes through an annular diaphragm of fixed radius and tunable aperture, as shown in figure 2.8. The pressure in the vessel is chosen so that the diaphragm lets through only photons emitted by kaons, and blocks photons emitted by pions and protons present in the beam. The aperture of the diaphragm is chosen to maximize the light yield from kaons while minimizing the contamination from light emitted by pions. During data taking, special runs are performed periodically to measure the optimal operating pressure: figure 2.9 shows the results of one such pressure scan. The coincidence of signals in at least 5 octants tags a kaon.

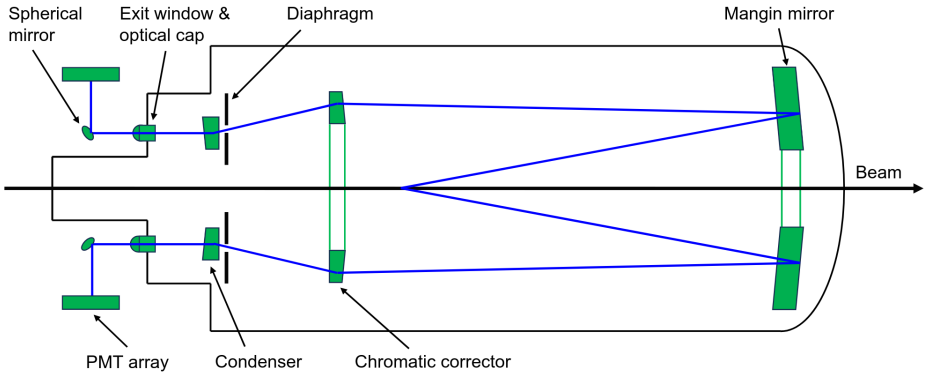

Figure 2.8: Sketch of the KTAG optical system (not to scale), highlighting the optical elements (green areas) and the path of Cherenkov light (blue lines). See [107, 108, 110] for the details of the optics.

The full detector, filled with nitrogen, constitutes a total of $3.4\,\%X_0$ thickness in the path of the beam. Using hydrogen as radiator, this figure reduces to only $0.73\,\%X_0$. In 2023, in order to profit from such a reduction, a new detector (CEDAR-H) was developed and installed in the NA62 beamline [110]. This should improve kaon transmission and reduce particle scattering in the KTAG material.

2.5 Gigatracker (GTK)

The GTK [103, 111] is a silicon detector that is designed to measure momentum, time, and direction of incoming beam particles. It needs to be extremely precise in order not to spoil the reconstruction of m^2_{miss} in the event: a precision of 0.2 % in momentum and of 16 µrad in angle is required. The detector takes its name from being able to withstand up to 1 GHz particle rate and

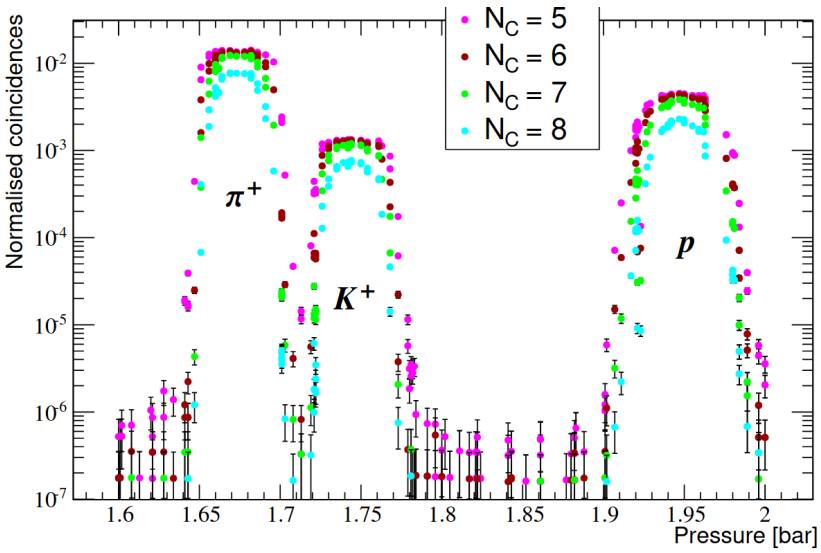

Figure 2.9: Results of a KTAG pressure scan at 1.5 mm diaphragm aperture, as function of the number of octants required to have a coincidence (N_C). [103]

up to 1.5 MHz/mm^2 around the detector center, which requires a hit time resolution better than 200 ps and a track time resolution better than 100 ps.

The detector is composed of three stations in Run1, and four stations in Run2. Each station consists of 18000 pixels of $300 \times 300\,\mu m^2$ or $400 \times 300\,\mu m^2$ area each, corresponding to a total active area of $60.8 \times 27\,mm^2$. Readout is performed by TDCpix chips [112], which are directly attached to the sensor via bump-bonding. The total thickness for each station is around 500 µm, which corresponds to less than 0.5 %X_0. The low thickness is crucial in order to reduce both multiple scattering (which spoils the momentum resolution) and inelastic scattering (which contributes to the upstream background).

The GTK stations are laid out inside an achromat, as shown in figure 2.10, which allows to measure the magnitude of the beam particle momentum by using the vertical displacement at GTK2; the direction of the momentum is measured by the positions at GTK0 (in Run2 only), GTK1 and GTK3. The first and last stations are 22.8 m apart; GTK2 is at 13.2 m from GTK1 in Run2, and at 7.2 m in Run2 (after having been moved upstream of the muon scraper). During LS2, GTK0 was installed 50 mm upstream of GTK1, with the aim of improving tracking performance and pileup rejection.

The detector is traversed by the full beam, and therefore is designed to be as radiation hard as possible. This is achieved thanks to a specifically-designed micro-channel cooling plate, which covers an especially large area compared to regular micro-channel designs, and constitutes the world's first application of such a device in high energy physics. The cooling system keeps the sensor and the front-end electronics at less than 5 °C. The detector mechanics were designed in such a way that sensors can be replaced rapidly if radiation damage becomes unbearable. One of the most irradiated modules

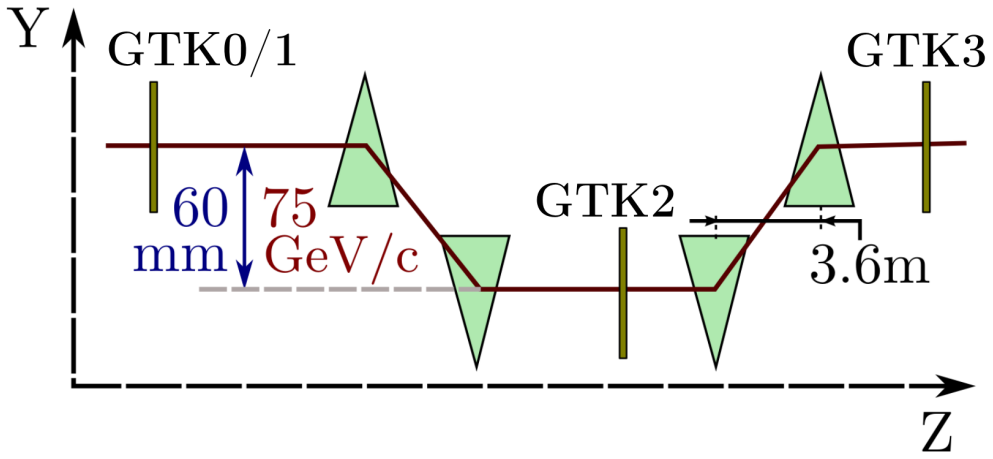

Figure 2.10: Sketch of the GTK layout (not to scale). [103]

reached an average fluence of 0.81×10^{14} $1\text{MeVn}_{\text{eq}}/\text{cm}^2$ by November 2022, but its single hit time resolution was maintained below 160 ps by gradually adjusting the bias voltage [113].

A picture of an assembled GTK station can be seen in figure 2.11.

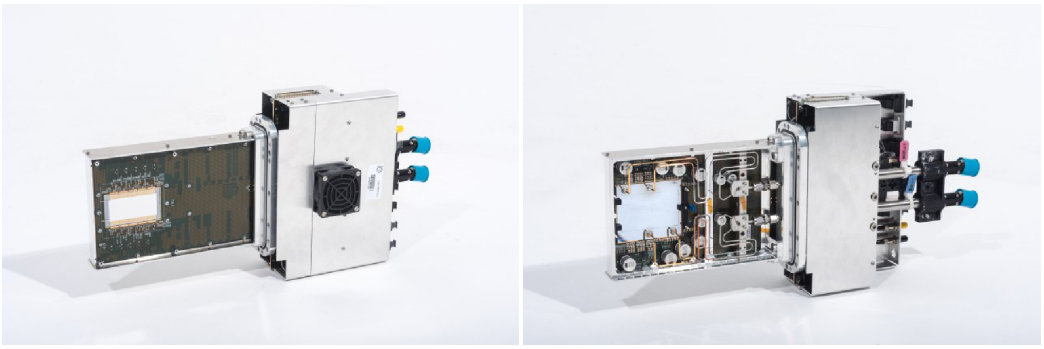

Figure 2.11: Sensor side (left) and cooling side (right) of an assembled GTK station. [103]

2.6 VetoCounter

The VetoCounter [78], developed and installed during LS2, is specifically designed to tackle the upstream background. In particular, it detects secondary particles emitted in kaon decays happening in the GTK achromat, which would otherwise remain undetected and be absorbed in the TCX.

It is composed of three stations: the first two stations (VC1, VC2) are installed before the TCX and separated by a 25 mm lead sheet, while the third station (VC3) is placed after the TCX. Each station is made of 11 scintillator tiles of which 3 tiles are above the beam pipe and 8 below it. Tiles have an active area of $40 \times 120\,\text{mm}^2$, apart from the ones adjacent to

the beam pipe, which are smaller ($20 \times 120\,\mathrm{mm}^2$) in order to avoid excessive rate (figure 2.12) [114]. Tiles are glued to two opposing fishtail light guides, attached to PMs (figure 2.13). The very bottom tile of VC3 is not read out.

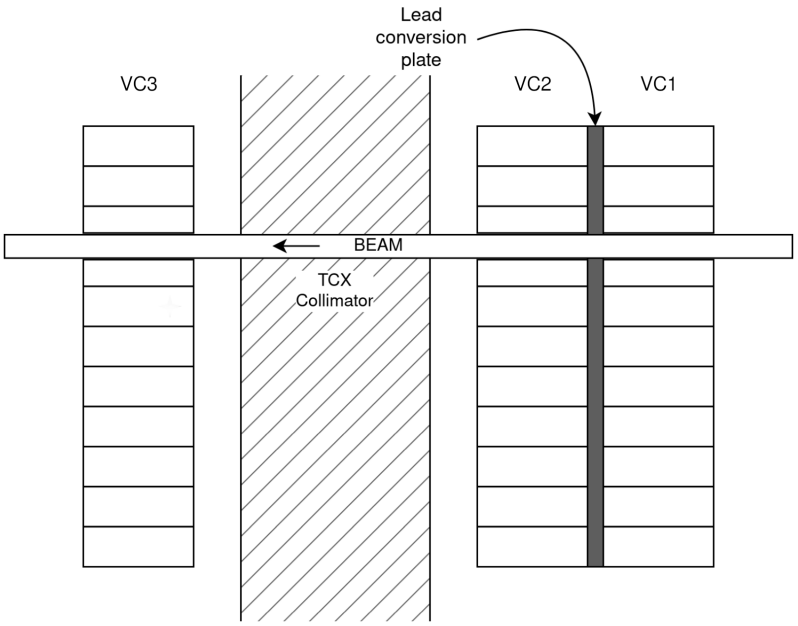

Figure 2.12: Side-view sketch of the VetoCounter layout (not to scale).

Figure 2.13: Bottom part of a VetoCounter station. [104]

The layout of the VetoCounter stations is designed so that it can provide rough particle identification, as shown in figure 2.14: a charged pion (or a positron) gives a signal in VC1 and not in VC3 (a shower arising from the lead converter could be detected in VC2 but all the energy from the pion is absorbed in the collimator); a photon or π^0 converts in the lead and is absorbed in the collimator, giving only a hit in VC2; a halo muon travels through the collimator and gives a hit in VC3. The geometry was studied with help of simulation so that it could give a factor 2 to 3 rejection of the

upstream background from kaon decays [115]. An example event vetoed by the VetoCounter, where a $K^+ \to \pi^+\pi^0$, $\pi^0 \to \gamma\gamma$ decay chain happens in the GTK achromat, an in-time pileup beam particle (which does not leave the beam pipe) is present, and the π^+ reaches the STRAW, is sketched in figure 2.15.

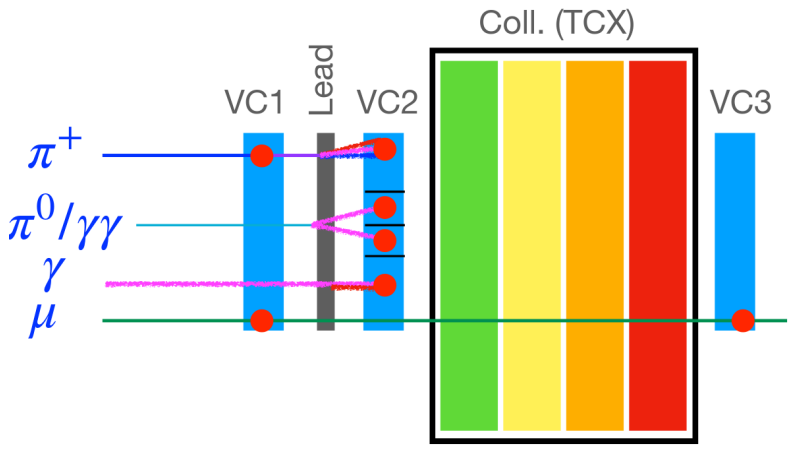

Figure 2.14: Sketch of approximate particle identification capabilities of VetoCounter (not to scale). [35]

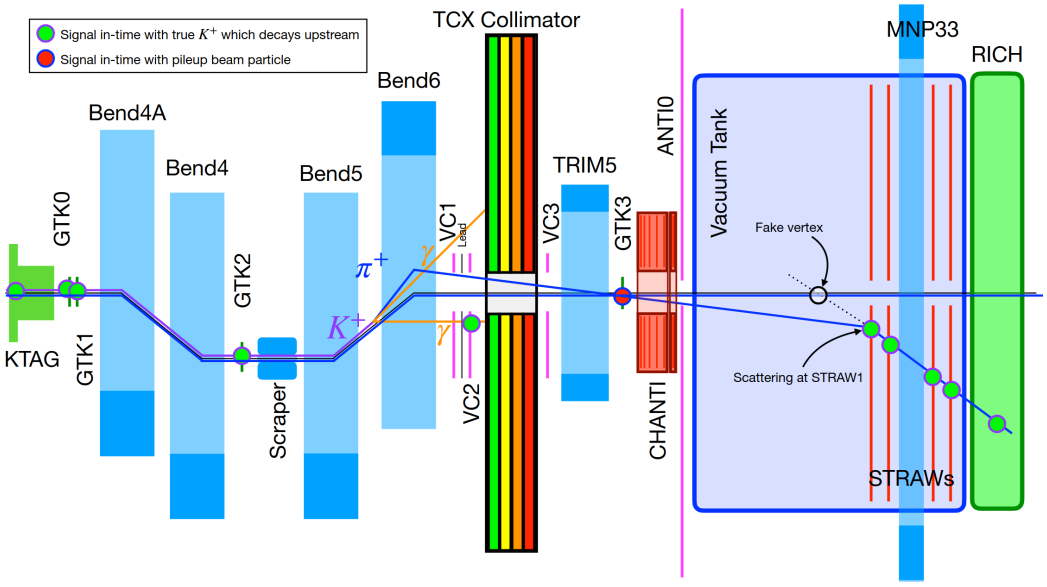

Figure 2.15: Side-view sketch of an upstream background event vetoed by the VetoCounter (not to scale). Blue rectangles represent the relevant magnets. [35]

The VetoCounter is read out by TEL62 boards. A parallel readout based on a prototype FPGA-based TDC system, with higher rate and no dead time,

and a FELIX board [116] was installed and fully commissioned during the 2022 data taking.

2.7 Charged anti-coincidence detector (CHANTI)

The CHANTI [117] provides rejection against background from inelastic scattering of the beam particles with GTK3. A pion emitted in this interaction can enter the fiducial volume with a momentum and direction compatible with a signal event, while other particles emitted at large angles are not detected by the downstream detectors.

The CHANTI is composed of six hodoscope stations of $300 \times 300\,\mathrm{mm}^2$ area, with a $95 \times 65\,\mathrm{mm}^2$ hole to leave room for the beam. Each station is composed of plastic scintillator bars of triangular section, read out with fast wavelength shifting (WLS) Saint-Gobain™ BCF-92 fibers and Hamamatsu™ MPPC S10362-13-050C[1] silicon photomultipliers (SiPM). The triangular cross section is designed so that a particle hitting a plane will almost always traverse two bars (figure 2.16). The scintillator bars are arranged into two planes, one with bars oriented horizontally and one with bars oriented vertically, to form a station (figure 2.17). Figure 2.18 shows the layout of the CHANTI stations, designed in such a way that the detector covers hermetically an angular region between 49 mrad and 1.34 rad for particles emitted at GTK3.

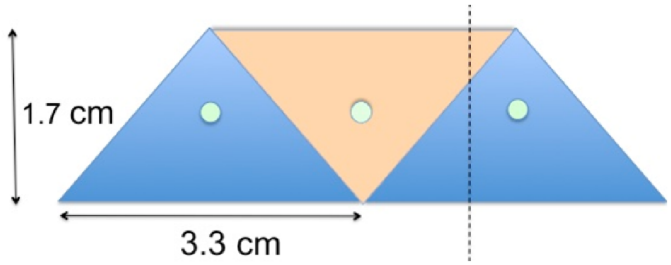

Figure 2.16: Arrangement of the bars to form a plane. The dashed line shows an example trajectory of a particle, which traverses two bars. [103]

The time resolution required in order to keep the probability of accidental veto (caused by pileup) of genuine kaon decays to an acceptable level is of the order of 2 ns. This requirement has been fully satisfied as the time resolution for the coincidence of two planes of a station has been measured to be 830 ps. [117]

In 2024, after the successful commissionining of the new TDC-FELIX based system for the VetoCounter, CHANTI was also equipped with this readout system.

[1] To avoid accumulating radiation damage, these were replaced with newer S13360-1350CS for Run2.

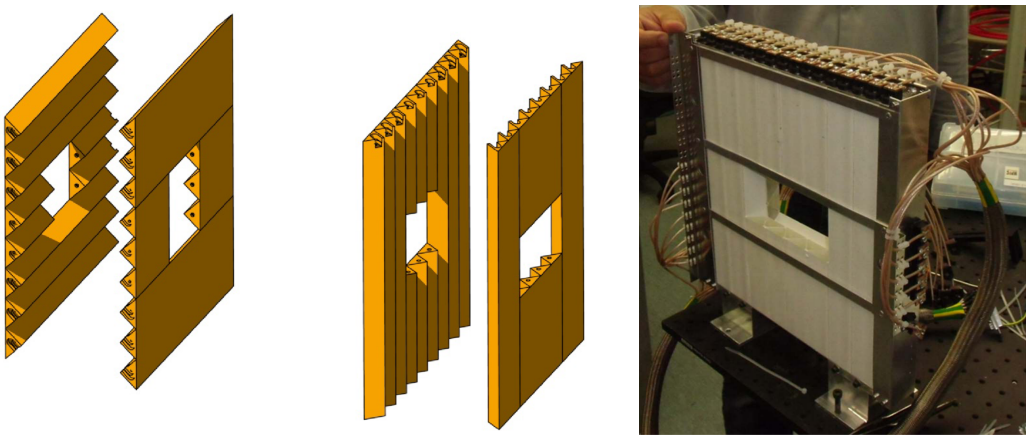

Figure 2.17: A CHANTI station: exploded view (left) [117] and picture of the final assembly (right). [103]

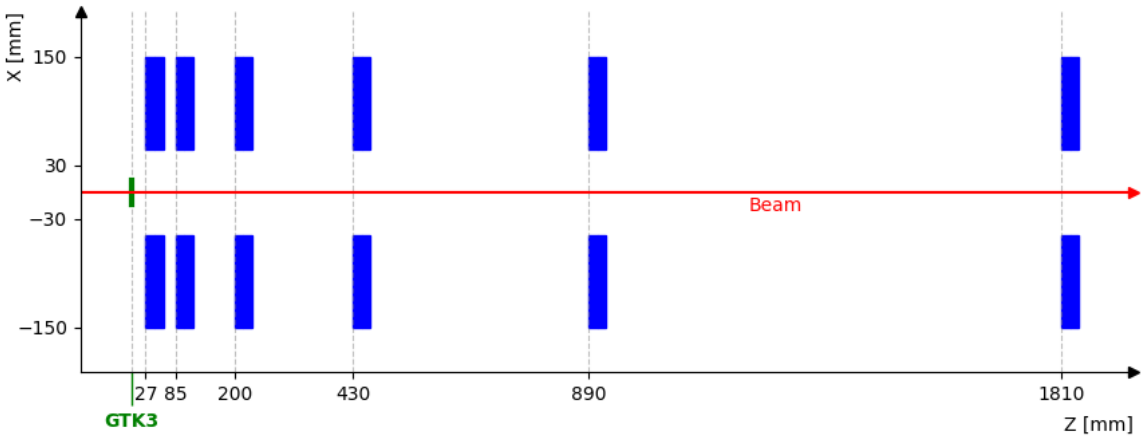

Figure 2.18: Layout of the CHANTI, with respect to GTK3.

2.8 ANTI0

The ANTI0 [105, 78] is a detector installed during LS2, designed to veto events with charged halo particles entering the fiducial decay region. It has crucial importance in data collected in beam dump configuration, and it is an important detector for the rejection of the upstream background.

It is a cell structure hodoscope, covering an area of 1080 mm radius around the beam pipe, similar in diameter to the vacuum tube. It is composed of 280 counters made of $124 \times 124\,\mathrm{mm}^2$ plastic scintillator tiles, each read out by four Hamamatsu™ MPPC S14160-6050HS SiPMs, connected in analog OR logic, into a single electronic channel. Counters are arranged in a checkerboard fashion and in two layers. A schematic view of the ANTI0 detector mainframe can be seen in figure 2.19.

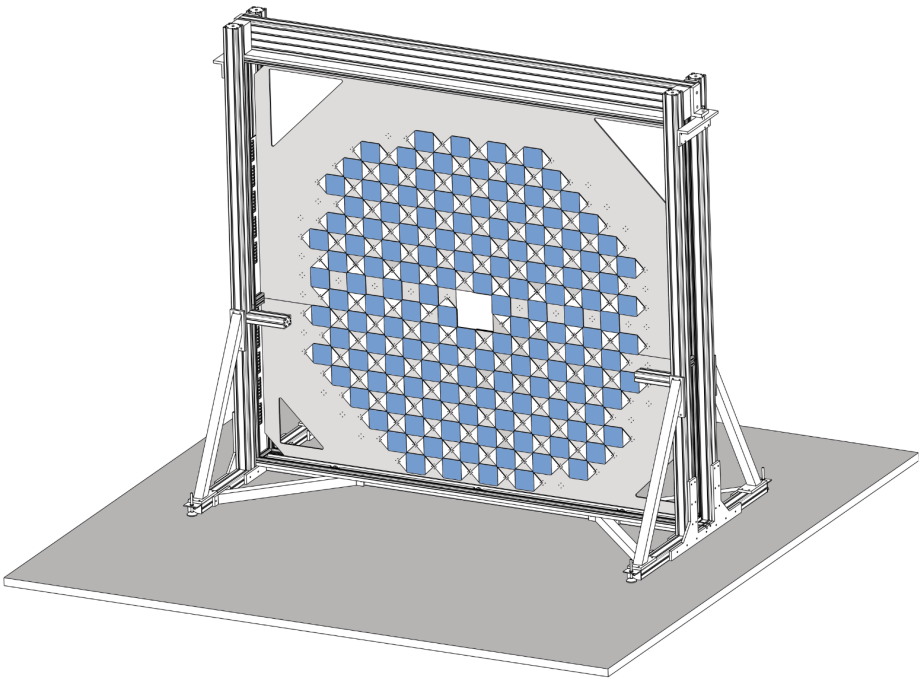

Figure 2.19: Schematic view of the ANTI0 detector mainframe. [105]

2.9 STRAW spectrometer

The STRAW [103] is a magnetic spectrometer based on straw tubes. It measures the trajectories and momenta of charged particles produced in kaon decays happening in the fiducial volume.

The NA62 spectrometer is the first tracking system based on straw tubes operating in vacuum: dedicated tooling was developed to guarantee straight straws positioned to high precision (within $\pm 200\,\mu\text{m}$). It extends over a length of 35 m along the beam line and is composed of four chambers and a large aperture dipole magnet (MNP33) placed between the first and the second pair of chambers. The MNP33 provides an integrated field of 0.9 T m, which gives a horizontal momentum kick of 270 MeV/c; this corresponds to a deflection of -3.6 mrad for a 75 GeV/c particle.

Each straw chamber contains four views: the first two views measure the X and Y coordinate, while the second two measure coordinates along the U and V axes, rotated by $-45°$ and $+45°$ respectively relative to the X view. The use of U and V views guarantees resolution of ambiguities: if, for example, two particles hit a chamber at the same time, the resulting hits in the X and Y views do not uniquely identify two points in the transverse plane. In each view, a few straws are left out so as to leave a hole for the passage of the beam (figure 2.20); due to the slope of the beam line in the horizontal plane, the hole has offsets along the X axis for each chamber (figure 2.7).

Each view has four layers, so that a high detection redundancy is obtained

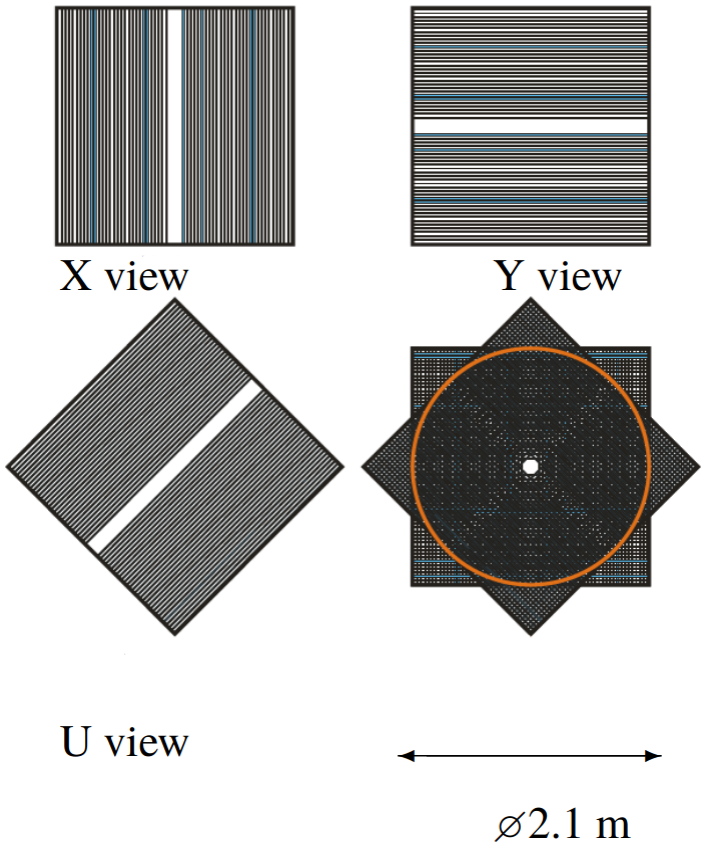

Figure 2.20: STRAW views X, Y, U and resulting active area of a chamber. [103]

and at least two straw crossings per view and per track are guaranteed, as needed to solve the left-right ambiguity. There are 448 straws per view, of 9.82 mm diameter and 2160 mm length, made of 36 µm thick polyethylene terephthalate, coated with 50 nm of copper and 20 nm of gold on the inside; gold-plated tungsten anode wires are 30 µm in diameter. The gas inside is 70 % Ar and 30 % CO_2 at atmospheric pressure. The total amount of material corresponds to 1.8 %X_0.

Both leading edges and trailing edges of the signal are read out. Leading edges are used to measure the drift time, which gives the lateral position of the hit; trailing edges give a measurement of the absolute time of a hit. The time resolution for a track traversing all four chambers is 5.1 ns. Owing to the very large number of channels and the impracticality of a high-resolution TDC system (due to the intrinsically poor time resolution), a dedicated data acquisition system was developed for STRAW, which reduced costs.

The spectrometer is calibrated using $K^+ \to \pi^+\pi^+\pi^-$ events. The total momentum resolution is consistent with

$$\left(\frac{\sigma_p}{p}\right)^2 = (0.30\,\%)^2 + \left(0.005\,\%\frac{p}{\mathrm{GeV}/c}\right)^2 ,$$

and the angular resolution ranges from 60 µrad at 10 GeV/c to 20 µrad at 50 GeV/c, which satisfy the requirements: the kaon mass can be reconstructed from $K^+ \to \pi^+\pi^+\pi^-$ events to a precision better than 1 MeV/c^2.

2.10 Ring-imaging Cherenkov counter (RICH)

The RICH [118, 119, 120] is designed to discriminate muons from pions between momenta of 15 GeV/c and 35 GeV/c.

The radiator vessel is 17.5 m long, and is filled with neon gas at atmospheric pressure and room temperature, which has the perfect index of refraction for the required Cherenkov threshold of charged pions (12.5 GeV/c). At the downstream end, a mosaic of 20 spherical mirrors of 17 m focal length reflects the Cherenkov light into a ring. Each mirror can be aligned by remotely controllable piezoelectric actuators allowing two-axis orientation. Two disks, made of 976 Hamamatsu™ R7400U-03 PMs each, collect the Cherenkov photons: these PMs have a 16 mm wide face and are packed in a hexagonal structure with 18 mm pitch. Figure 2.21 shows a sketch of the RICH.

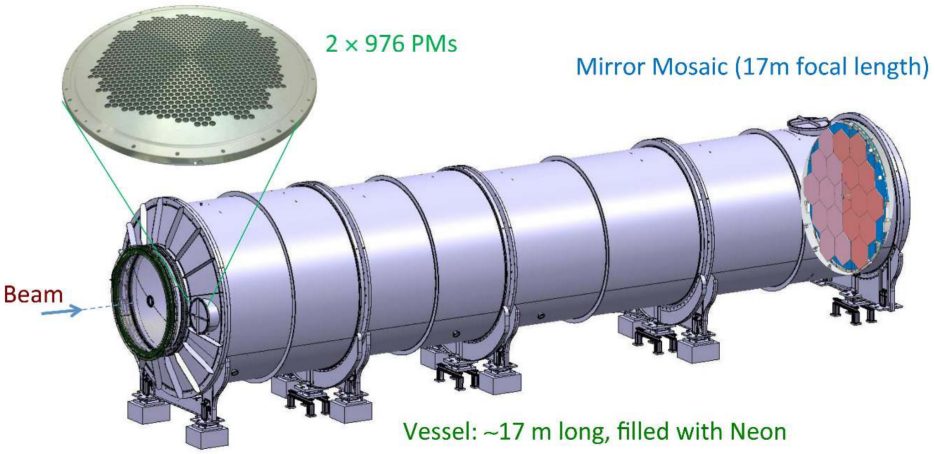

Figure 2.21: The RICH. [103]

The RICH performance is monitored with a clean positron sample, obtained from $K^+ \to \pi^0 e^+ \nu$ decays. The Cherenkov ring radius resolution for positrons has been measured to be 1.47 mm. The single hit time resolution is 0.24 ns, corresponding to a global RICH event time resolution of 70 ps. Due to this excellent timing performance, the RICH is extensively used in the hardware-level trigger chain, providing the reference time. Figure 2.22 shows the distribution of reconstructed ring radii as function of the momentum in single track events; the separation between particles is clearly visible.

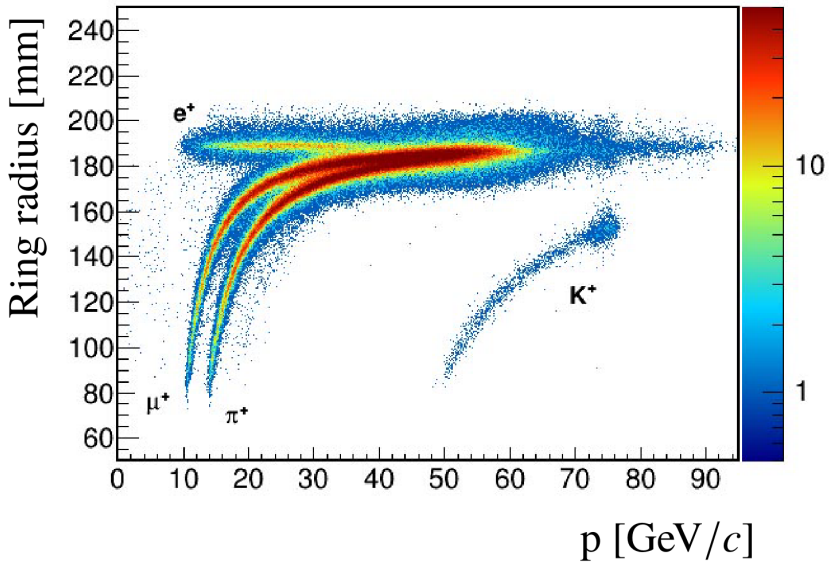

Figure 2.22: RICH ring radius vs momentum of single track events in NA62 2015 data. [103]

2.11 Charged particle hodoscopes (CHOD)

Two scintillator detectors constitute the NA62 hodoscope system: the NA48-CHOD and the CHOD. The former is reused from the NA48 experiment [121], while the latter has been built specifically for NA62 (and is therefore sometimes referred to as NewCHOD). They are 700 mm apart and they are operated independently.

The NA48-CHOD [121] is made of plastic scintillator slabs divided into two views of 64 slabs each, one with slabs oriented vertically and one with slabs oriented horizontally. Views are subdivided into four quadrants of 16 slabs each; each view makes an octagon of 1210 mm apothem. Slabs are 20 mm thick (corresponding to 10 $\%X_0$), whereas their widths are smaller close to the beam pipe (65 mm) and larger in the peripheric region (99 mm), to accomodate for the particle rate. Each slab is read out by a Photonis™ XP2262B PM through a fishtail light guide. A drawing of the NA48-CHOD is shown in figure 2.23.

The NA48-CHOD time resolution approaches 200 ps [103]: this figure is obtained using corrections of time hits for the impact point position on the slabs, which can be calculated through the coincidence between horizontal and vertical views. However, since such corrections are not available at the data acquisition level, the resolution for trigger purposes is around 2 ns. This limits the trigger selection, since, in the high rate environment of NA62, an unacceptable rate of fake coincidences would be produced. Hence, the CHOD was built, with a completely different structure.

The CHOD [103] is made of an array of 152 plastic scintillator tiles, with

30 mm thickness and sizes varying according to the distribution of hit rates. It covers an annulus with inner and outer radii of 140 mm and 1070 mm respectively. Tiles are placed in such a way that they have a 1 mm overlap. The total thickness in the active area is 13 %X_0. The CHOD structure is shown in figure 2.24. Its time resolution, owing to constant fraction discriminators which have been employed since 2017, is better than 500 ps.

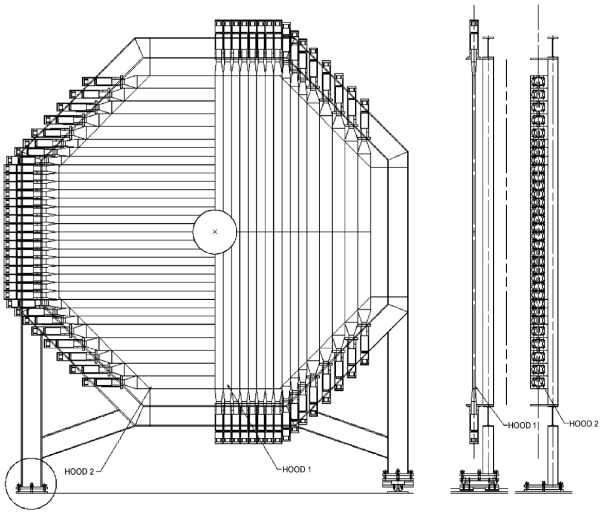

Figure 2.23: The NA48-CHOD. [121]

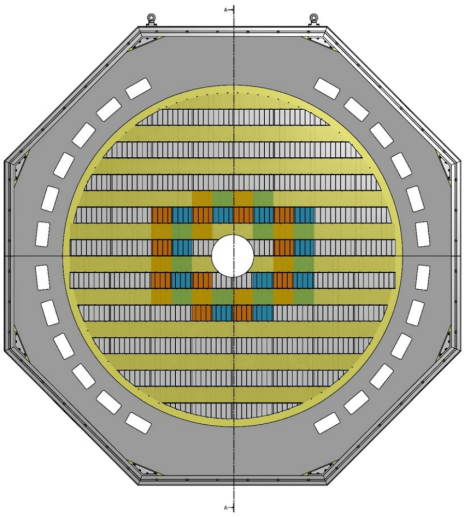

Figure 2.24: The CHOD. [103]

2.12 Liquid krypton calorimeter (LKr)

NA62 reuses the liquid krypton electromagnetic calorimeter that was built for the NA48 experiment [121]. The LKr is used both as main photon veto and as an important aid for particle identification.

The calorimeter is filled with about 9 m^3 of krypton at 120 K, housed inside a octagonal cryostat that extends from the beam pipe (about 8 cm from the Z axis) up to a radius of 128 cm; its depth is 127 cm, for a total thickness of 27X_0. The volume is divided longitudinally by Cu-Be electrodes into cells of $2 \times 2\,cm^2$ cross section. They are placed in a zig-zag shape, in order to minimize the response dependence on the particle impact position (figure 2.25). Each cell works as a ionization chamber by appling a 3 kV bias voltage. The signals are collected by preamplifiers inside the cryostat, and sent out to transceiver boards via vacuum feedthrough connectors on top of the cryostat.

The existing NA48 readout system was phased out in order to satisfy the much higher rate requirements by NA62; a new readout system was built and commissioned. The *calorimeter readout module* (CREAM) employs 14-bit, serial-output 40 MHz flash ADCs, FPGAs for the handling of the data and trigger requests, and large DDR3 memories [122].

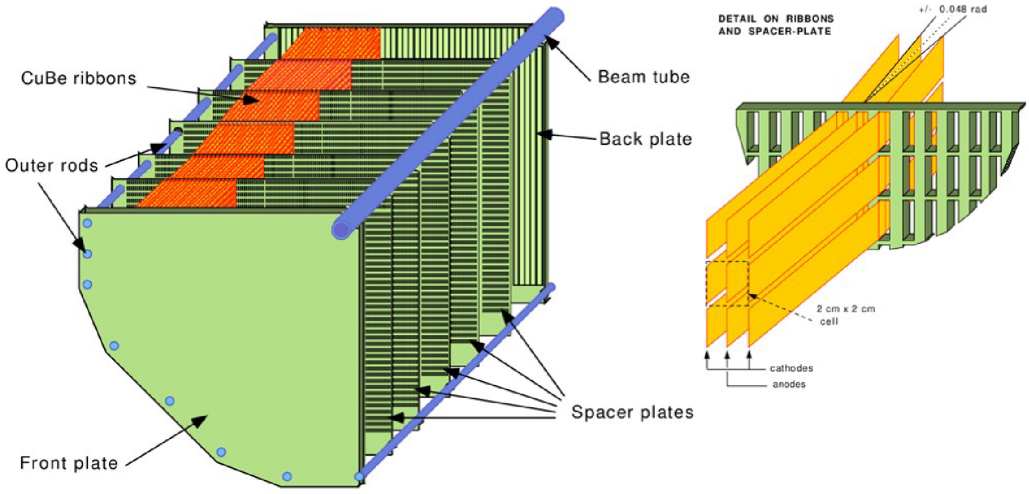

Figure 2.25: Schematic of the calorimeter structure, and detail of a calorimeter cell. [103]

The LKr covers an angular region of polar angles between 1 and 8.5 mrad, with a photon detection inefficiency of $O(10^{-5})$ for photon energies larger than 10 GeV [98]. The energy resolution can be parametrized as function of the cluster energy E as

$$\left(\frac{\sigma_E}{E}\right)^2 = \left(\frac{4.8\,\%}{\sqrt{E/\mathrm{GeV}}}\right)^2 + \left(\frac{11\,\%}{E/\mathrm{GeV}}\right)^2 + (0.9\,\%)^2 \,,$$

while the spatial resolution as

$$\sigma_{x,y}^2 = \left(\frac{4.2\,\mathrm{mm}}{\sqrt{E/\mathrm{GeV}}}\right)^2 + (0.6\,\mathrm{mm})^2 \,.$$

Typical time resolutions are around 0.55 ns for clusters with $E > 3$ GeV.

2.13 Muon veto system (MUV)

The MUV [103] complements the RICH and the LKr for the μ/π discrimination. It is composed of a hadronic calorimeter (MUV1, MUV2), and a fast scintillator plane (MUV3) placed behind a 80 cm thick iron wall, as shown in figure 2.26.

The MUV1 is made of 24 layers of 26.8 mm thick steel plates, interleaved by 23 layers of 9 mm thick, 60 mm wide scintillator strips, alternately aligned horizontally and vertically. Strips in the innermost part are split into two halves, in order to accommodate the higher rate (and the space for the beam pipe). Each strip is wrapped in aluminized mylar foils to minimise light losses, and read out by two WLS fibers, connected to Hamamatsu™ R6095S

PMs to each side (except for the half-length strips). Each PM receives fibers for all the strips that are similarly aligned and have the same transverse position. The detector corresponds to a total of 4.1 interaction lengths, and covers an area of around $2.7 \times 2.6\,\mathrm{m}^2$.

The MUV2 is the refurbished front module of the NA48 hadron calorimeter [121]. It is similar to the MUV1: 24 iron plates of 25 mm thickness are interleaved with plastic scintillator planes made of $1300 \times 119 \times 4.5\,\mathrm{mm}^3$ strips, each spanning half the calorimeter, and alternately aligned in the horizontal and vertical direction. Both the MUV1 and the MUV2 are read out by flash ADCs and CREAM boards.

The MUV3 is placed behind a total of over 14 interaction lengths, provided by the LKr, the MUV1, the MUV2 and the iron wall, and it is used for muon identification. It is built of 148 scintillator tiles, of which the 8 closest to the beam pipe are smaller as required by the high rate, and the remaining 140 have $220 \times 220 \times 50\,\mathrm{mm}^3$ dimensions. Each tile is read out by a pair of two-inch PMs placed behind the tile. The innermost tiles are mostly illuminated by muons coming from decays of beam pions, therefore the time resolution is required to be around 500 ps. This excellent time resolution is also exploited for generation of hardware-level trigger information.

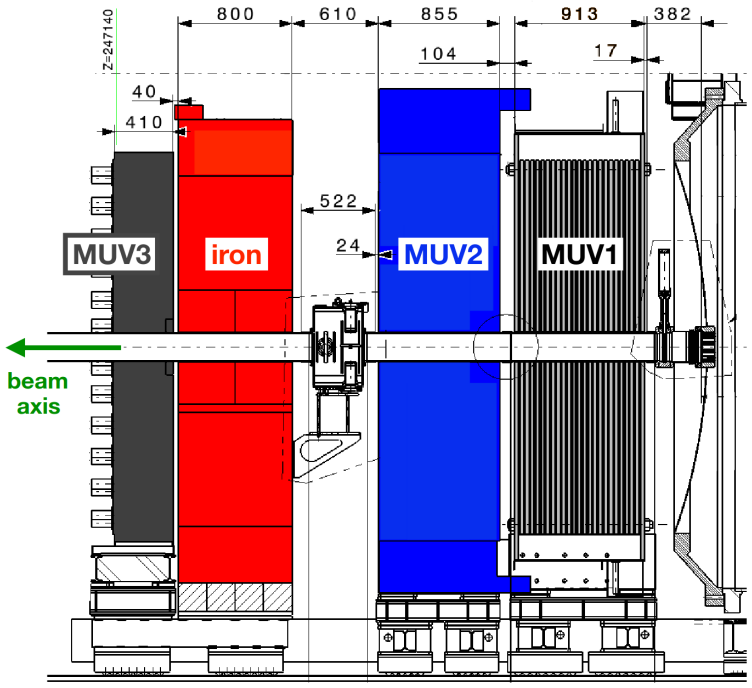

Figure 2.26: The MUV. [123]

2.14 Large-angle veto (LAV)

The 12 LAV stations provide full geometric coverage for photons emitted in the decay volume at angles between 8.5 and 50 mrad. Because of the two-body kinematics, is not possible for both photons emerging out of a $K^+ \to \pi^+\pi^0$ decay in the fiducial volume on the Z axis to be emitted at angles larger than 50 mrad; more than 95 % of photons in the LAV acceptance have energy larger than 200 MeV and must be detected with very low inefficiency.

LAV stations are ring-shaped, and are built of lead-glass blocks which were originally used for the OPAL electromagnetic calorimeter barrel [124] and refurbished for the use in NA62. The blocks are made of about 75 % PbO by weight, and emit Cherenkov light when an electromagnetic shower is produced in them. Each of these blocks is around 37 cm long, and its front and rear faces are $10 \times 10\,\mathrm{mm}^2$ and $11 \times 11\,\mathrm{mm}^2$, respectively. Light is collected by a 4 cm long cylindrical light guide and read out by a 76 mm diameter Hamamatsu™ R2238 PM. Blocks are arranged in rings by aligning them radially, and multiple rings, staggered in azimuth, make a LAV station (figure 2.27). Different LAV stations have different radii in order to satisfy the hermetic geometry requirements. The minimum effective depth for a particle incident on any station is $21X_0$; most particles are intercepted by four or more blocks, which add up to $27X_0$. The time resolution for photons of 1 GeV energy is about 1 ns, while the inefficiency of detecting a photon of energy larger than few hundreds of MeV is at the permille level [98].

2.15 Small-angle veto (SAV)

The SAV, composed of the intermediate-ring calorimeter (IRC) and the small-angle calorimeter (SAC), extend the photon detection coverage down to 0 mrad (figure 2.28). This is crucial for $\pi^0 \to \gamma\gamma$ decays where the two photons are extremely unbalanced and one of the two has an energy too low (and therefore angle too large) to be inside the LAV acceptance, as the other one will have a large boost and can stay inside the beam pipe [103].

The IRC is a cylindrical shashlik calorimeter, and it surrounds the beam pipe just in front of the LKr. It is made of two longitudinal modules, of 89 mm and 154 mm depth respectively, separated by 40 mm. Both modules contain ring-shaped layers of 1.5 mm lead absorber and 1.5 mm scintillator plate, for a total thickness of $19X_0$. The scintillator rings are divided into four sectors which are optically isolated; Saint Gobain™ BCF-92 WLS fibers traverse both modules, and each sector is read out by a Hamamatsu™ R6427 PM.

The SAC is also a shashlik calorimeter, made of 70 layers of alternating $205 \times 205 \times 1.5\,\mathrm{mm}^3$ lead and plastic scintillator planes, for a total depth of $19X_0$. The scintillation light is collected by Kuraray™ Y-11(250) MSJ WLS fibers, which are bent into a U shape and threaded through two holes

Figure 2.27: LAV1. [103]

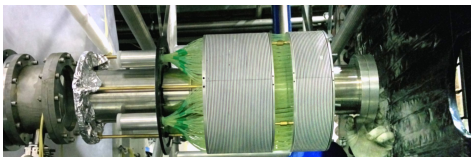

(a) The IRC. [103]

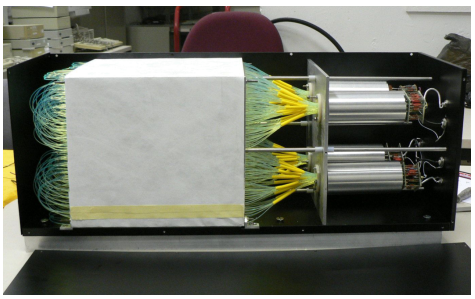

(b) The SAC. [103]

Figure 2.28: The SAV.

each. Fiber ends are subdivided into four bundles, each read out by a 25 mm Hamamatsu™ R6427S PM. There is no actual optical separation of the four readout channels, so the SAC is effectively a single-channel detector. The SAC is installed in the beam vacuum, just upstream of the beam dump. As discussed in section 2.3 and shown in figure 2.7, a dipole magnet ensures that the beam is directed out of the SAC acceptance, so that only neutral particles travelling along the Z axis can impact the SAC. The SAC is aligned at a 23 mrad angle with respect to the Z axis, in order to prevent photons incident on the Z axis from traversing the detector along fibers without encountering lead or scintillator material.

The IRC and the SAC are readout both with CREAM boards and TEL62 boards for extra redundancy.

2.16 MUV0 and HASC

The MUV0 is a scintillator hodoscope, designed to detect π^- from $K^+ \to \pi^+\pi^+\pi^-$ decays which have such a small momentum ($\lesssim 10\,\mathrm{GeV}/c$) that the spectrometer dipole magnet deflects them out of the RICH acceptance. It is mounted on the downstream flange on the RICH and covers the area defined by $1545\,\mathrm{mm} < X < 2945\,\mathrm{mm}$ and $|Y| < 700\,\mathrm{mm}$; it is composed of 48 plastic scintillator tiles, grouped in 9 super-tiles, each read out by a Hamamatsu™ R7400 PM [103].

The *hadronic sampling calorimeter* (HASC) is made from 9 identical longitudinal modules of 60 layers of 16 mm lead and 4 mm scintillator, with $100 \times 100\,\mathrm{mm}^2$ transverse area [103]. It is used to detect π^+ from $K^+ \to$

$\pi^+\pi^+\pi^-$ decays with high momenta ($\gtrsim 50\,\mathrm{GeV}/c$) which travel along the beam pipe, escaping the STRAW acceptance, and are deflected by the bend magnet in front of the SAC. It is also very useful to detect positrons emitted along the Z axis in photon conversions in the downstream end of the RICH, either on the mirrors or on the beam pipe.

In some rare cases, the e^+e^- pair from such a photon conversion is very unbalanced, so that the e^+ has too high energy to be bent into the HASC acceptance. Therefore, during LS2, HASC2, a perfect copy of the HASC, was made and installed in a symmetrical position with respect to the beam pipe. The HASC2 placement is designed to detect the e^- of such an unbalanced e^+e^- pair, as shown in figure 2.29.

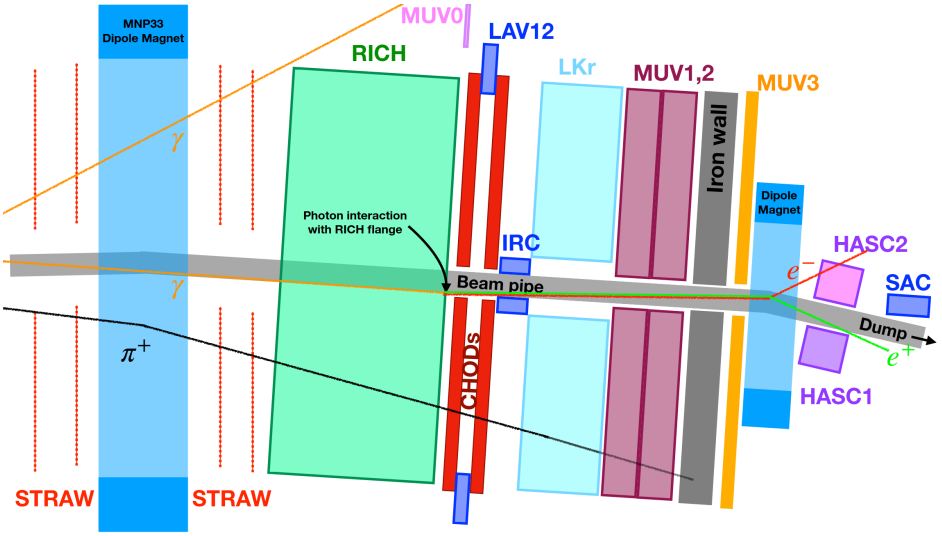

Figure 2.29: Top-view sketch of a $K^+ \to \pi^+\pi^0$ event where HASC2 plays a crucial role in the veto (not to scale). [35]

2.17 Trigger and data acquisition system

The trigger and data acquisition system of NA62 is rather unique in that it provides scalability, large integration, and excellent performances.

A general-purpose carrier board, called TEL62 [125], hosting high-precision TDCs [126], was developed to provide a common readout for most sub-detectors, scalability, efficient online selection and lossless high rate readout. TEL62 boards are capable of both producing trigger primitives and to handle data coming from the daughter TDC boards, sending to the online PC farm a programmable number of 25 ns wide time slots around the trigger time at a request of a level-0 trigger.

The hardware level-0 (L0) trigger must withstand the 10 MHz rate of beam K^+ and π^+ decays in the fiducial volume and reduce it to 1 MHz [127]. The system has been designed to be able to evaluate a trigger decision based

on the full set of digitized data which are transmitted from the detectors to L0, without any dedicated sub-system trigger. This provides extreme flexibility, no design limitations on the availability of detector information, reduction of hardware, and full control and offline reproducibility of the trigger algorithms. The L0 trigger decision is taken by the L0 trigger processor (L0TP), implemented on a commercial FPGA development board [128]. The L0TP can apply a *downscaling* to each trigger line: for a downscaling of D, only one every D events satisfying the trigger conditions will be actually passed to the next trigger stage.

A high-level trigger, or level-1 (L1) trigger, is implemented in software running in the online PC farm. The L1 trigger has access to the complete data stream of all sub-detectors, and can reconstruct higher-level objects like tracks or clusters, allowing a data rate reduction to 100 kHz. A small fraction (around 2 %) of events are marked as *autopass* and are stored regardless of the L1 trigger decision.

Figure 2.30 shows the trigger rates per spill as function of the beam intensity, based on data collected in 2018 [129], while table 2.3 summarizes all the relevant rates in the NA62 experiment.

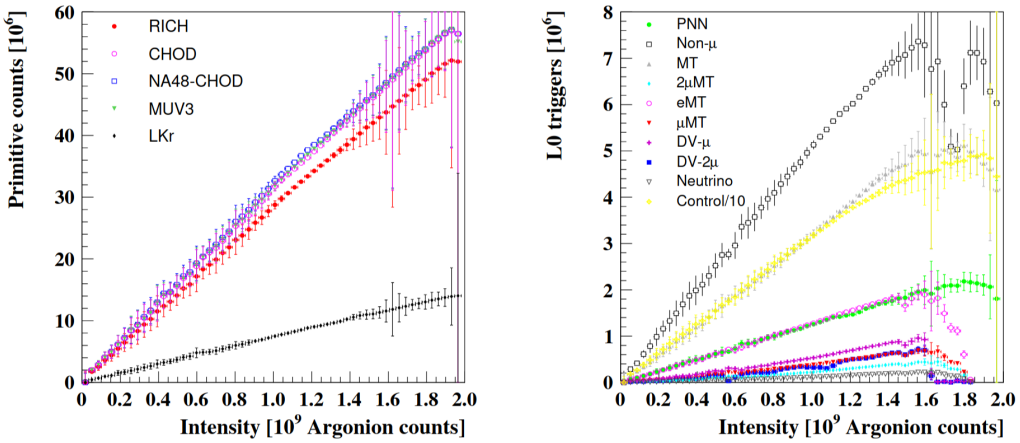

Figure 2.30: Number of primitives for each relevant detector (left) and L0 triggers before downscaling for each trigger line (right) per spill as function of the beam intensity, as measured by an Argon ionization chamber (Argonion) placed at the end of the experimental hall. See [129] for the complete definition of the trigger lines.

Nominal beam particle rate	600 MHz
Beam K^+ rate	36 MHz
Beam K^+ decay rate in the fiducial volume	5 MHz
Beam π^+ decay rate in the fiducial volume	5 MHz
L0 trigger rate	1 MHz
L1 trigger rate	100 kHz
PNN trigger rate	50 kHz

Table 2.3: Summary of relevant rates in the NA62 experiment. The PNN trigger is the trigger line dedicated to the collection of the signal events, see section 3.3.

Chapter 3

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

The signal selection has generally the same structure as the one already used for the analysis of Run1 data [34], but several changes were made, both to improve on existing methods and to adapt to the larger intensity of the dataset as well as the addition of the new detectors.

A $K^+ \to \pi^+ \nu\bar{\nu}$ decay is selected by building event information in steps, and discarding the event when it does not fit the signal description. The first step is the reconstruction and selection of a downstream track with positive charge, and the association of other downstream sub-detector signals. The K^+ is then reconstructed by associating a KTAG signal and a GTK track, using both timing and space information. In this way, a candidate vertex between the two tracks is formed. Using the measured momentum of the two tracks, m^2_{miss} is calculated, which, as discussed in section 2.1, is the key kinematic variable used for background rejection. Vetoes are applied in the upstream part of the detector, including anti-coincidence with VetoCounter and ANTI0, and rejection of spurious hits in the GTK. Finally, particle identification is enforced on the downstream track, and photon veto and multiplicity vetoes are applied.

For the final BR extraction, the number of observed events must be normalized to an effective number of kaon decays. This number is estimated by counting the events from a normalization channel, which has been chosen to be $K^+ \to \pi^+\pi^0$ followed by $\pi^0 \to \gamma\gamma$ (denoted $K^+ \to \pi^+\pi^0_{\gamma\gamma}$ in the following), because it shares many properties with the signal, resulting in similar selections; this leads to systematic uncertainty cancellations. For the signal selection, only a set of photon and multiplicity vetoes is applied on top of the normalization selection. The effect of these vetoes in the signal selection is evaluated separately.

The analyzed dataset is described in section 3.1, and a brief description of the simulated samples used for this analysis is given in section 3.2. Sections 3.3–3.10 go through the steps of the signal selection, starting from the very first selection applied, namely the trigger lines used. Finally, section 3.11 describes the selection of the normalization events.

3.1 Dataset

Data used for this analysis were collected in a total of about 300 days during 2021 and 2022. A total of about 2.5×10^{18} protons on T10 (POT) were delivered, and about 550×10^{3} spills were recorded. For comparison, during Run1 the total POT delivered were 2.2×10^{18} [81].

Data collected in 2021 presented a peculiar feature: in the first second of each spill spikes of very high instantaneous intensity (reaching 6000 MHz, as measured by the GTK) were observed; after the first second, the intensity was stable around the nominal average of 600 MHz (figure 3.1). The spikes were due to an issue in the beam extraction from the SPS accelerator, which was fixed during the yearly winter shutdown between 2021 and 2022.

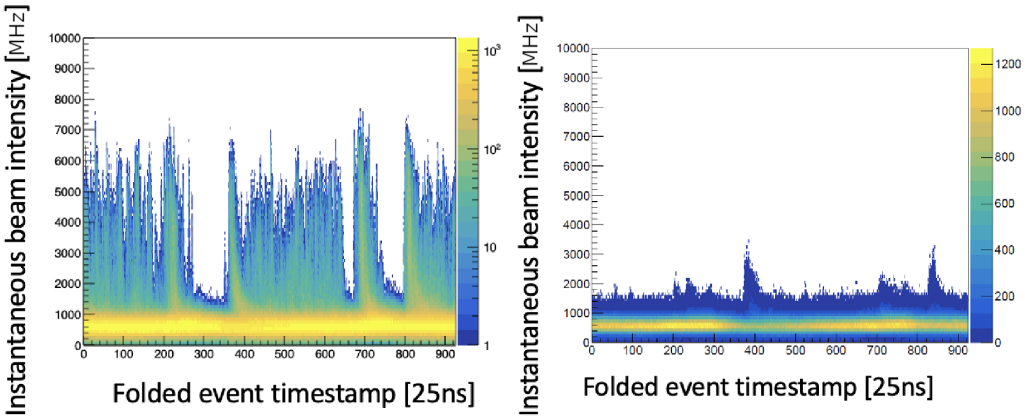

Figure 3.1: Instantaneous intensity as function of the event timestamp in the spill modulo 25 ns, for all events in the spill (left) and only for events after the first second of spill (right). [78]

The intensity spikes caused the data acquisition system to suffer, as the apparatus could not sustain the large instantaneous rate of events. The large instantaneous intensity was also a problem for the data quality, resulting in large inefficiencies and increased backgrounds in this part of the spill. Therefore, it was decided to discard the early part of each spill recorded in 2021 from the analysis. This corresponds to approximately 2 % of the combined 2021–2022 recorded dataset.[1]

Not all recorded spills are considered for the analysis. Data quality is monitored during data taking and after the data reprocessing: spills in which sub-detector performances were not optimal are discarded for the final analysis. The fraction of discarded spills amounts to around 25 % for the 2021–2022 dataset. Some of these can be reused for different analyses which do not have

[1]The corresponding POT loss is difficult to quantify, since for a large fraction of this time the data acquisition system had large inefficiencies, which had the effect of not recording events on disk. Therefore, the number of recorded events lost does not scale proportionally to the POT loss.

data quality requirements as stringent as the $K^+ \to \pi^+\nu\bar{\nu}$ analysis.

A summary of the dataset used for the present analysis is reported in table 3.1.

	2021	2022
Days of data taking	85	215
Protons delivered on T10	0.7×10^{18}	1.8×10^{18}
Number of spills recorded	145×10^3	403×10^3
Number of spills analyzed	103×10^3	311×10^3

Table 3.1: Summary of the 2021 and 2022 datasets.

3.2 Simulated samples

The analysis is designed to be as data-driven, and independent of simulations, as possible, because it is very difficult to rely on models and simulations for extremely rare events. Nonetheless, some simulated samples are used for estimation of quantities for which no data sample can be cleanly isolated.

The relevant kaon decays are simulated with Monte Carlo (MC) techniques. Beam optics without particle interactions are simulated with the `TURTLE` package [130]. Kaon decays are simulated in the kaon rest frame using the SM predictions for the relevant matrix elements and form factors, then boosted to the laboratory frame. The detector geometry and response is simulated using the `Geant4` toolkit [131]. Some aspects of the simulation, such as signal formation in sub-detectors or detector inefficiencies, rely on parametrizations from data. Simulated samples are then reconstructed and analysed in the same way as data.

Events are simulated in an effective zero-intensity environment, where no event pileup is present. In ancillary studies which need the parametrization with respect to the intensity, multiple events are superimposed in time, according to intensity templates measured from data.

3.3 Trigger lines

Four trigger lines were used in Run2 to collect data relevant for the $K^+ \to \pi^+\nu\bar{\nu}$ analysis.

The PNN trigger is the trigger line dedicated to the collection of $K^+ \to \pi^+\nu\bar{\nu}$ events. The trigger conditions applied at L0 are:

RICH: at least three hits in $RICH^2$, which define the trigger reference time;

CHOD: at least one signal in CHOD within 5 ns of the reference time;

UTMC: no more than 5 coincident signals in CHOD within 5 ns of the reference time;

$\overline{\text{Qx}}$: no coincident signals in diagonally-opposite CHOD quadrants;

$\overline{\text{MUV3}}$: no signal in MUV3 within 4 ns of the reference time;[3]

$\overline{\text{LKr}}$: no more than 40 GeV of energy deposited in LKr, and veto of events that have at least 2 clusters with a total of more than 5 GeV energy.

Conditions applied at L1 are:

KTAG: a kaon identified in KTAG within 5 ns of the reference time;

$\overline{\text{LAV}}$: at most 2 coincident signals in LAV[4] within 6 ns of the L0 trigger time;

STRAW: at least one positively-charged track forming a vertex with the nominal beam axis upstream of the first STRAW chamber, with momentum smaller than 65 GeV/c;

STRAW-PNN[5]: STRAW tracks must be in acceptance of the MUV3 detector, and no two tracks must form a vertex.

The NORM trigger is the control trigger line dedicated to events with at least one pion track. With reference to the conditions as defined above, it requires RICH, CHOD and $\overline{\text{MUV3}}$ at L0, and KTAG and STRAW at L1. It is downscaled by a factor 400.[6] This trigger line is mainly used to collect normalization $K^+ \to \pi^+\pi^0_{\gamma\gamma}$ events.

The MB trigger is the control trigger line with minimum bias, requiring only at least two hits in the RICH and the CHOD condition at L0, and no conditions at L1. It is downscaled by a factor 600. The MB trigger is mainly used to collect control $K^+ \to \mu^+\nu$ samples.

The PNN, NORM and MB trigger lines are designed to be progressively less stringent, so that efficiency ratios between them benefit from systematic uncertainty cancellations.

Finally, the CTRL trigger is implemented to be independent of all other trigger lines and only requires a signal in the NA48-CHOD. It is downscaled by a factor 2000.

[2]For part of the 2021 dataset, the RICH condition only required two hits.

[3]Apart from a portion of the 2021 dataset, signals from a specific tile adjacent to the beam pipe are ignored, as this tile is mostly illuminated by muons from beam pion decays, and mostly generates accidental veto of events.

[4]LAV stations 1 and 12 are ignored to avoid vetoing events with inelastic interactions in GTK and RICH, respectively.

[5]Not included in a portion of the 2021 dataset.

[6]In small portions of the 2021 dataset, this downscaling is 100 or 300.

3.4 Downstream particle

A candidate downstream STRAW track must fulfill geometric and quality criteria. Quality criteria are based on hit requirements (all four chambers must have been hit, and no hit must be shared with a different candidate track) and on the quality of the fit (evaluated via the residuals of the fit). Geometric criteria impose that the extrapolation of the track must remain within the geometric acceptance of the downstream detectors (NA48-CHOD, CHOD, LKr, MUV1, MUV2, MUV3) and stay outside of the geometric acceptance of the LAV, IRC and SAC detectors.[7]

A NA48-CHOD coincidence is required between the horizontal and vertical plane. The best hit pair is chosen using a discriminant based on $T_V - T_H$ (where T_V and T_H are the times of the hits in the vertical and horizontal plane, respectively; both are corrected for time walk and light propagation) and on $\frac{1}{2}(T_V + T_H) - T_{\text{trig}}$ (where T_{trig} is the trigger reference time). A CHOD coincidence is also required, and the best hit is chosen based on both time and space coincidence with the STRAW track.

The association of the track to the RICH detector is performed by requiring at least 4 hits that are in time among them (within 2 ns) and in an annulus around the expected ring center, as extrapolated by the STRAW track. The times of such selected hits are averaged to calculate the RICH time of the downstream particle.

The downstream particle time is defined by the mean of the STRAW time, the CHOD time, and the RICH time. The mean is weighted by the approximate resolutions of the detectors (10 ns, 0.5 ns, 0.2 ns respectively), and is therefore effectively dominated by the RICH time.

3.5 Kaon selection and GTK track matching

A time coincidence with KTAG is required. In particular, the KTAG coincidence which is closest in time with the downstream particle is required to consist of hits in at least 5 sectors. The average time of the KTAG hits defines the kaon time.

Once the downstream track is built, and the kaon time is defined, a GTK track must be associated, in order to measure the kaon momentum and calculate m^2_{miss}. This is a crucial part of the selection: at a 600 MHz particle rate, it is fairly likely that more than one beam particle is compatible in time with the downstream particle and the kaon time, and errors in this association can easily lead to background events. Errors can be of two different types:

wrong association: this happens whenever the downstream track is coming from a K^+ decay in the fiducial volume, but an unrelated GTK track is

[7] In geometric acceptance requirements, tolerances are applied for extrapolation errors and for scattering in the material.

matched to it. This leads to a mismeasurement of the kaon momentum, and therefore of m^2_{miss}, enhancing the probability of main kaon decays entering the signal region;

accidental association: this happens if the downstream track is uncorrelated with the GTK tracks in the event (and therefore no parent kaon is actually reconstructed in the GTK), but a GTK track is associated anyway. If such events are not identified and discarded, they enter the signal region as what is called the *upstream background*.

3.5.1 *Bayesian* K^+–π^+ matching algorithm

The algorithm to match the K^+ track to the downstream particle and to the kaon time was completely revised for the present analysis. Further details about the studies that led to the present formulation of the algorithm are given in appendix A.

Not all GTK tracks are considered as potential candidate K^+ tracks. In particular, for each track, a χ^2 discriminant compares the measured 3-momentum to the average beam 3-momentum, and a threshold discards tracks that are not compatible with the average beam. The beam parameters (mean and standard deviation of momentum and angles) are measured in independent $K^+ \to \pi^+\pi^+\pi^-$ samples via the total momentum of the three pions as measured by the STRAW. In addition, a condition on the quality of the GTK track fit is required: another χ^2 discriminant is calculated by evaluating the residuals of the linear fit of the X and Y positions of the hits and of the hit times. This condition removes tracks that are deemed fake. The number of tracks that satisfy these requirements in a loose time window (corresponding to a half-width of 3 ns around the average of the RICH and KTAG times), referred to as *good* GTK tracks in the following, is denoted by N_{GTK}, and its distribution is shown in figure 3.2.

A prior probability of any given good GTK track to be the correct K^+ track is assigned as

$$\mathrm{P}(\mathrm{K}^+) = \frac{\overline{N_{\mathrm{K}^+\text{reco}}}(N_{\text{GTK}})}{N_{\text{GTK}}}\,, \tag{3.1}$$

where $\overline{N_{\mathrm{K}^+\text{reco}}}$ is the expected number of parent kaons among the good GTK tracks, and is evaluated as a function of N_{GTK}; the division by N_{GTK} reflects the lack of bias assigned to any specific GTK track, and ensures that the sum of $\mathrm{P}(\mathrm{K}^+)$ over all GTK tracks equals $\overline{N_{\mathrm{K}^+\text{reco}}}$. Since, by definition, $\overline{N_{\mathrm{K}^+\text{reco}}}$ cannot exceed 1, it can be reinterpreted as the probability, $P_{\mathrm{K}^+\text{reco}}$, of the K^+ being reconstructed in the GTK as a good track. This quantity can be measured in a $K^+ \to \pi^+\pi^+\pi^-$ control sample: its dependency on N_{GTK} is shown in figure 3.3.

The good GTK tracks are then evaluated for space and time matching to the event, based on two variables. The spatial matching to the STRAW

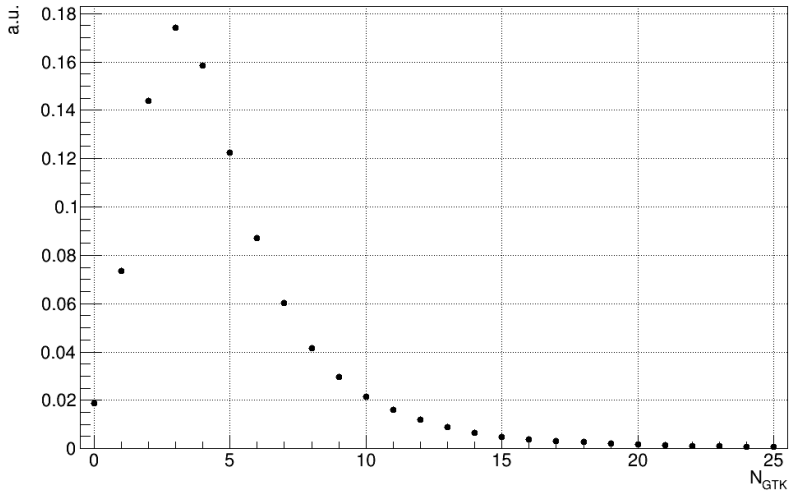

Figure 3.2: Distribution of the number of good GTK tracks.

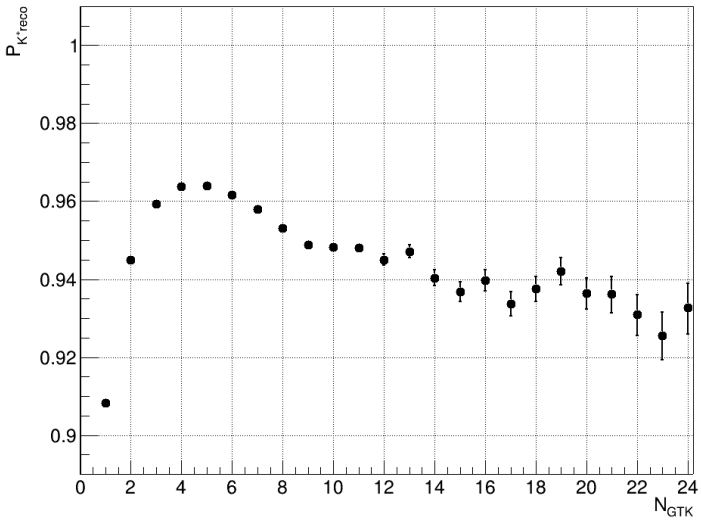

Figure 3.3: Probability of reconstructing the K^+ in the GTK as a good track, as a function of the number of good GTK tracks.

track is quantified by the *closest distance of approach* (CDA), evaluated by extrapolating both the STRAW track and the GTK track to the fiducial volume. The time matching of the GTK track to the event is quantified by

$$\begin{aligned} \Delta T_{\text{match}} &= \frac{\Delta T_{\text{RICH,GTK}} + \Delta T_{\text{KTAG,GTK}}}{\sqrt{2}} \\ &= \sqrt{2}\left(\frac{T_{\text{RICH}} + T_{\text{KTAG}}}{2} - T_{\text{GTK}}\right), \end{aligned} \tag{3.2}$$

where T_{RICH} and T_{KTAG} are the RICH and KTAG times, calculated as discussed above, T_{GTK} is the time of the GTK track, and $\Delta T_{x,\mathrm{GTK}} = T_x - T_{\mathrm{GTK}}$ ($x = \mathrm{RICH}, \mathrm{KTAG}$). Probability density functions (PDFs) of these two variables, $f_{\mathrm{CDA}}(\mathrm{CDA})$ and $f_{\Delta T}(\Delta T_{\mathrm{match}})$, both in the correct K^+ track hypothesis and in the accidental pileup track hypothesis, are fitted from the $\mathrm{K}^+ \to \pi^+\pi^+\pi^-$ control sample, and shown in figure 3.4.

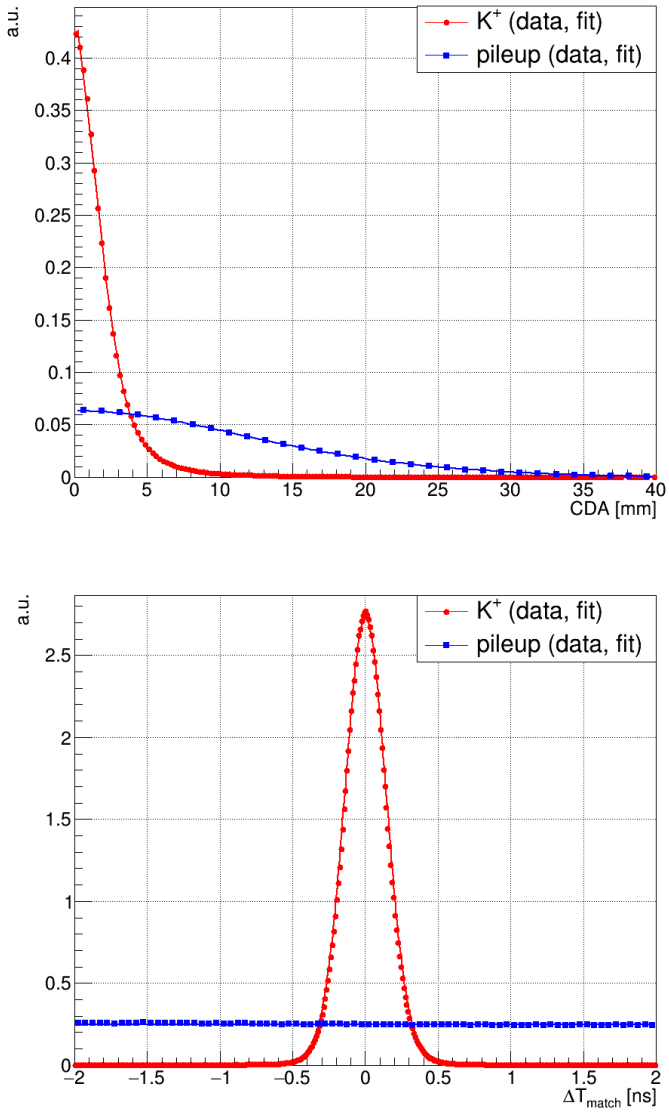

Figure 3.4: PDFs for CDA (top) and for $\Delta T_{\mathrm{match}}$ (bottom). Red is for kaon hypothesis, blue is for pileup track hypothesis.

A log likelihood-ratio test statistic ρ, for each GTK track, between the hypotheses of correct K^+ track and accidental pileup track, quantifies the matching quality of the track. The two likelihoods are calculated via Bayes'

theorem:

$$\begin{aligned}\rho &= \rho_0 + \log \frac{\mathrm{P}(\text{pileup}|\mathrm{CDA}, \Delta T_{\text{match}})}{\mathrm{P}(\mathrm{K}^+|\mathrm{CDA}, \Delta T_{\text{match}})} \\ &= \rho_0 + \log \frac{f_{\mathrm{CDA}}(\mathrm{CDA}|\text{pileup}) f_{\Delta T}(\Delta T_{\text{match}}|\text{pileup})\,\mathrm{P}(\text{pileup})}{f_{\mathrm{CDA}}(\mathrm{CDA}|\mathrm{K}^+) f_{\Delta T}(\Delta T_{\text{match}}|\mathrm{K}^+)\,\mathrm{P}(\mathrm{K}^+)}\,,\end{aligned} \tag{3.3}$$

where $\mathrm{P}(\text{pileup}) = 1 - \mathrm{P}(\mathrm{K}^+)$ and ρ_0 is a translation factor, which makes all possible values of ρ positive.

The track with the smallest value of ρ is chosen to be the GTK match. It should be noted that the factor $\mathrm{P}(\text{pileup})/\,\mathrm{P}(\mathrm{K}^+)$ does not depend on the GTK track and therefore does not enter the choice of the GTK match, but it does change the absolute value of ρ, which can then be used as a proxy for the quality of the matching.

If $N_{\mathrm{GTK}} > 1$, a further log likelihood-ratio test statistic $\Delta\rho$ is computed to compare the two best tracks. The two hypotheses being compared are that the best match is the correct K^+ track (and the second best match is an accidental pileup track), and that the second best match is the correct K^+ track (and the best match is an accidental pileup track). Using indices 1 and 2 for the best and the second best matches, respectively, the test statistic is

$$\begin{aligned}\Delta\rho &= \rho_2 - \rho_1 \\ &= \log \frac{\mathrm{P}(\mathrm{K}^+|\mathrm{CDA}_1, \Delta T_{\text{match},1})\,\mathrm{P}(\text{pileup}|\mathrm{CDA}_2, \Delta T_{\text{match},2})}{\mathrm{P}(\mathrm{K}^+|\mathrm{CDA}_2, \Delta T_{\text{match},2})\,\mathrm{P}(\text{pileup}|\mathrm{CDA}_1, \Delta T_{\text{match},1})}\,.\end{aligned} \tag{3.4}$$

High values of $\Delta\rho$ correspond to larger separation between the best and the second best matches.

3.5.2 K^+–π^+ matching selection criteria

Once a GTK track has been associated, a set of selection criteria is applied to reject events wrong associations and accidental associations.

A first criterion is $\mathrm{CDA} < 4\,\mathrm{mm}$, which rejects accidental associations.

Timing is also considered for the association; in particular, the KTAG time, the RICH time and the GTK time are all required to be within 0.5 ns from each other. Furthermore, the GTK time and the NA48-CHOD time cannot be more than 1.1 ns apart.

The GTK track matching is also evaluated via ρ_1 and $\Delta\rho$. For events with $N_{\mathrm{GTK}} = 1$, $\Delta\rho$ cannot be defined, and the criterion $\rho_1 < 2.825$ is applied in order to reject accidental associations, which generally have large ρ_1 (namely, a large likelihood of the only track being a pileup track). For events with $N_{\mathrm{GTK}} > 1$, a cut is applied in the $(\rho_1, \Delta\rho)$ plane, consisting of criteria on two linear combinations: one is $\rho_1 - \frac{\Delta\rho - 1.05}{530} < 6.30$, that mainly rejects the accidental associations (which generally have large ρ_1), and the other is $\Delta\rho - \frac{\rho_1 - 2.45}{4.05} < 0.1$, that mainly rejects the wrong associations (which generally have small $\Delta\rho$). This criterion is graphically represented in figure 3.5.

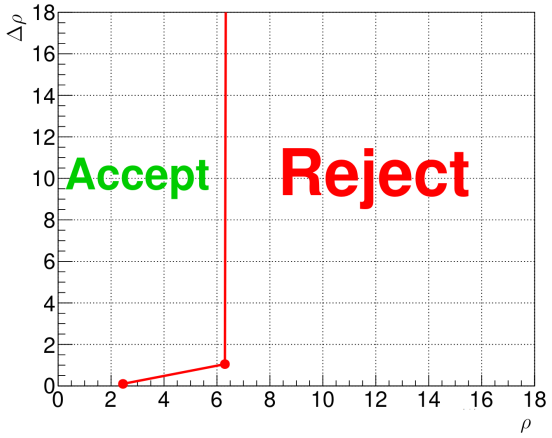

Figure 3.5: Graphical representation of the matching quality criterion for $N_{\mathrm{GTK}} > 1$ events.

The average efficiency of K^+ identification is 80 %, while the wrong association probability is 1.5 % and the accidental association probability is 5.9 %. Figure 3.6 shows the performance of the K^+–π^+ association, evaluated through a $\mathrm{K}^+ \to \pi^+\pi^+\pi^-$ control sample, as function of the instantaneous intensity, measured by the GTK. The saturation with intensity of the wrong and accidental association probabilities (without a significant drop of efficiency) can be attributed to the dependence on N_{GTK} through the prior probability (3.1), which effectively biases the algorithm to discard events with too many tracks, where errors are more likely to happen. This constitutes one of the main achievements of this K^+–π^+ matching algorithm.

3.6 Fiducial volume

With a GTK track tagged as a kaon, the K^+–π^+ vertex can be calculated as the midpoint between the extrapolations of the two tracks at their closest approach. Criteria on the coordinates of this vertex, denoted X_{vtx}, Y_{vtx}, Z_{vtx}, are applied as follows.

The vertex must be inside the fiducial volume: $110\,\mathrm{m} < Z_{\mathrm{vtx}} < 170\,\mathrm{m}$. For STRAW track momenta $p_{\pi^+} < 20\,\mathrm{GeV}/c$, the criterion $Z_{\mathrm{vtx}} < 165\,\mathrm{m}$ is also applied, in order to reject $\mathrm{K}^+ \to \pi^+\pi^0$ decays where the π^0 has large momentum and the two photons coming out of its decay are very unbalanced: in this case, the lower energy photon could go undetected through LAV stations, while the higher energy photon can go into the SAV, which has slightly higher photon detection inefficiency than the LKr.

Furthermore, a cut on the track distance from the beam axis at the first STRAW chamber, denoted R_{STRAW1}, is applied. This cut depends both on p_{π^+} and on Z_{vtx}, and is designed to remove most of the $\mathrm{K}^+ \to \pi^+\pi^+\pi^-$

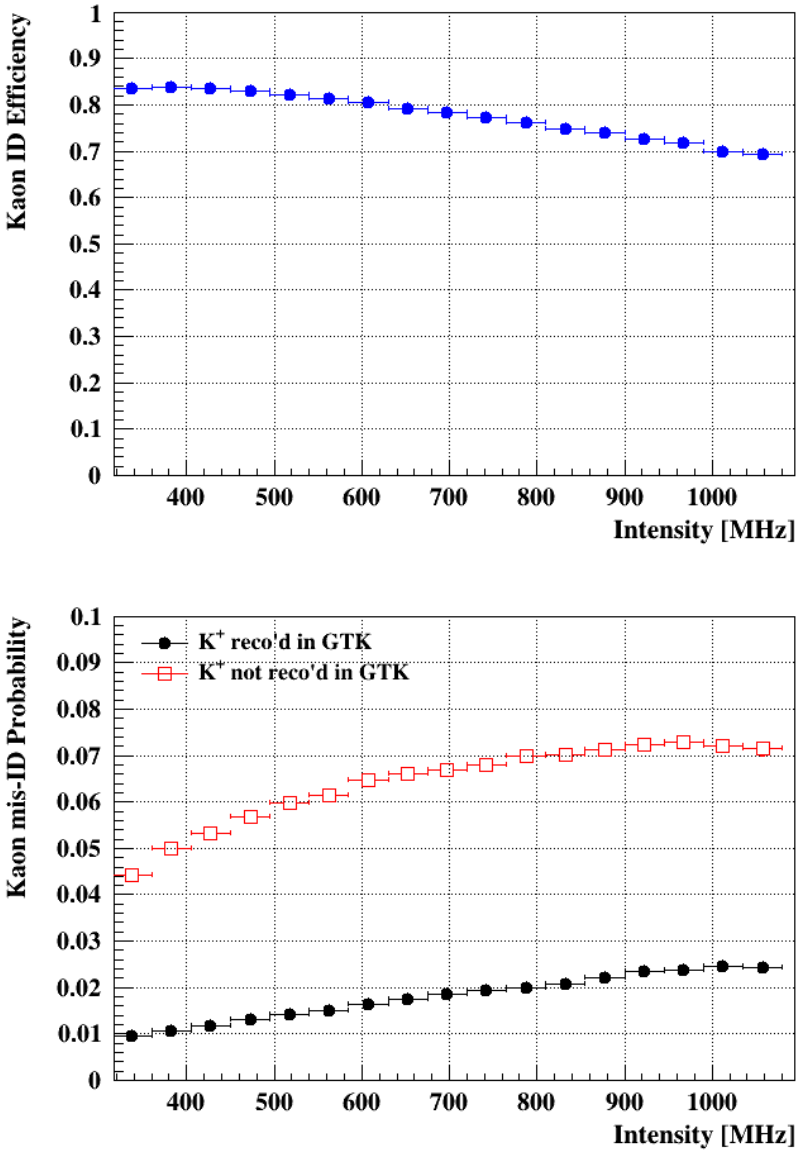

Figure 3.6: Efficiency (top) and error probabilities (bottom) for the K^+–π^+ association. In the bottom figure, red is for accidental associations, black is for wrong associations.

decays. These are indeed more likely to enter the signal region if they happen at the beginning in the fiducial volume and have a π^+ hitting the first STRAW chamber at small distances from the beam axis, while the other pions are emitted at large angles and escape the geometric acceptance of the detector. The dependence of this cut from p_{π^+} is one of the additions to the present analysis with respect to the Run1 analysis, aimed to increase the signal acceptance without losing background rejection power.

Finally, in order to minimize the non-gaussian tails of the $m^2_{\rm miss}$ resolution,

cuts are applied in both the $(X_{\text{vtx}}, Z_{\text{vtx}})$ and the $(Y_{\text{vtx}}, Z_{\text{vtx}})$ planes, which remove the far tails of the distributions of the positions of the vertices.

3.7 Upstream background rejection

A set of veto conditions are applied to reject upstream background events.

A signal in the CHANTI is defined as the coincidence between a hit in the horizontal and a hit in the vertical views of a CHANTI station. No signal in the CHANTI must be present within 3 ns of either the track time, or the KTAG time, or the GTK time.

No hits must be present in the ANTI0 which are consistent with the position of the STRAW track when back-extrapolated to the ANTI0 plane, and which are in time with the track time within 3 ns.

No signals must be present in the VetoCounter within 2 ns of the track time, except if a single hit in a single PM is recorded (in which case the signal is tagged as noise), or exactly three hits, one per station and all with the same Y coordinate, are present (in which case the signal is tagged as a muon and not considered dangerous for background).

No good GTK tracks which are within 0.6 ns of the KTAG time should form a vertex with the downstream track in the region $100\,\text{m} < Z < 105\,\text{m}$, upstream of the fiducial volume. The back-extrapolation of the downstream track at the GTK3 plane must also not be within the GTK3 sensor. Furthermore, no extra hits in the GTK, which are not associated with a track, must be present within 0.4 ns of the KTAG time or the RICH time. GTK hits with high time-over-threshold (> 23 ns) are also considered as extra activity to be vetoed.

A specific category of events, where the position of the downstream track at the first STRAW chamber is in the region defined by $75\,\text{mm} < R < 200\,\text{mm}$ (R is the distance from the Z axis) and $X > 150\,\text{mm}$, and its back-extrapolation at the TRIM5 plane is in the region defined by $400\,\text{mm} < R < 500\,\text{mm}$ and $X < -300\,\text{mm}$, is thought to be due to beam pions entering through the beam pipe and scattering in the first STRAW chamber. These events are also vetoed.

Finally, a boosted decision tree (BDT) classifier is used to reject upstream background events [132]. This classifier uses spatial information about the π^+ track and the vertex: the (X, Y) position of the track extrapolation at the final collimator, the (X, Y) position of the track at the first STRAW chamber, the vertex coordinates and the slope of the STRAW track. The BDT classifier is trained on signal Monte Carlo samples and on a background data sample, which is obtained by enlarging the fiducial volume cut to $105\,\text{m} < Z_{\text{vtx}} < 170\,\text{m}$, and inverting both the CDA and the timing conditions on the K^+–π^+ matching quality described in section 3.5.2. This BDT-based criterion is the most powerful against upstream background events.

3.8 Particle identification

The identification of the charged particle is performed based on quantities measured independently by the RICH, the calorimeters (LKr, MUV1, MUV2) and the MUV3. The particle must be identified as a pion in order to reject muons and positrons.

3.8.1 RICH

The RICH is leveraged for particle identification with two different approaches, giving rise to two distinctly applied criteria.

In a first approach, the hits that are compatible with the STRAW track (as defined in section 3.4) are fitted in order to extract the ring radius, and the mass of the particle m_{RICH} is calculated based on the STRAW momentum. It is required that $m_{\mathrm{RICH}} < 200\,\mathrm{MeV}/c^2$, and that it is larger than a momentum-dependent threshold ranging from $112\,\mathrm{MeV}/c^2$ to $125\,\mathrm{MeV}/c^2$.

In a different approach, expected rings are calculated based on the expected ring center and on the STRAW momentum assuming several mass hypotheses (positron, muon, pion, kaon), and for each mass hypothesis a likelihood (L_{e}, L_{μ}, L_{π}, L_{K}) is computed based on the observed hits (the calculation closely follows the algorithm used in the WA89 experiment [133]). A likelihood L_0 is also computed for the hypothesis that the hits are only due to noise. A condition is applied to the ratio $\max\{L_{\mathrm{e}}, L_{\mu}, L_{\mathrm{K}}, L_0\}/L_{\pi^+}$, which must be less than a momentum-dependent value (between 9 % and 16 % for $p_{\pi^+} < 35\,\mathrm{GeV}/c$, and 82 % for $p_{\pi^+} > 35\,\mathrm{GeV}/c$).

Both conditions are looser at high momenta, where the RICH performance degrade, and stricter conditions would cost in terms of signal efficiency; background rejection at higher momenta is recovered by exploiting different variables, especially kinematics.

3.8.2 Calorimeters

An energy deposit (*cluster*) is required to be associated in the LKr, and it must be compatible with the STRAW track both in time (within 20 ns) and in space (the extrapolation of the track at the LKr plane must lie within 100 mm of the cluster). Hits in MUV1 and MUV2 that are in time (within 10 ns of the LKr cluster time) and close to the STRAW track extrapolation (within 300 mm) are also collected. With information from the three calorimeters, the total calorimetric energy is built.

The ratio between the calorimetric energy and the STRAW track momentum must be less than $1.2\,c$, and also the ratio between the energy deposited in the LKr and the STRAW track momentum must be less than $0.8\,c$. Muons are rejected by a condition based on a discriminant which detects energy deposits that are compatible with a minimum ionizing particle (MIP) in at least one of the three calorimeters. Additional conditions reject events where extra

energy is deposited in MUV1 or MUV2 far from the hits associated with the track (more than 7 slabs away). Further specific criteria reject penetrating particles depositing energy in MUV2 but not in MUV1 (by checking the fraction of energy deposited in MUV1 and MUV2), and muon decays before LKr (using the shape of the LKr cluster).

Finally, a BDT classifier is used to quantify the probability of the particle being a pion based on all the calorimetric variables discussed above [123]. A momentum-dependent threshold on this probability is applied, ranging between 85 % at 15 GeV/c and 95 % at 45 GeV/c.

3.8.3 MUV3

A veto on MUV3 signals is applied. Specifically, events are rejected if they have a hit in MUV3 within 5 ns of the track time.

A specific MUV3 tile, adjacent to the beam pipe, is mostly illuminated by muons from beam pion decays, therefore it is ignored for the veto condition, except if the STRAW track extrapolates near it.

3.8.4 Performance

High-purity control samples of pions and muons are selected in order to evaluate the particle identification performance. The pion sample is obtained from a $K^+ \to \pi^+\pi^0$ sample collected with the MB trigger, by tagging the π^0 both kinematically (via m^2_{miss}) and actively in the LKr (by requiring two extra clusters); the muon sample is obtained from $K^+ \to \mu^+\nu$ decays, also collected with the MB trigger, tagged kinematically.

Figure 3.7 shows the measured performance for the particle identification, as function of the track momentum, separately for the RICH and for the combined LKr and MUV systems (discontinuities are due to the momentum-dependent criteria). For the whole 15 to 45 GeV/c range, an efficiency for pion identification of (73.00 ± 0.01) %, and a probability of muon misidentification of $(1.3 \pm 0.2) \times 10^{-7}$, are measured.

3.9 Photon and multiplicity rejection

Photon and multiplicity rejection are a crucial part of the selection, rejecting kaon decays with extra visible particles in the final state other than the π^+, such as $K^+ \to \pi^+\pi^0$ or $K^+ \to \pi^+\pi^-e^+\nu$ decays.

3.9.1 Photon veto

Photon vetoes exploit information from the LAV, LKr, IRC, and SAC detectors.

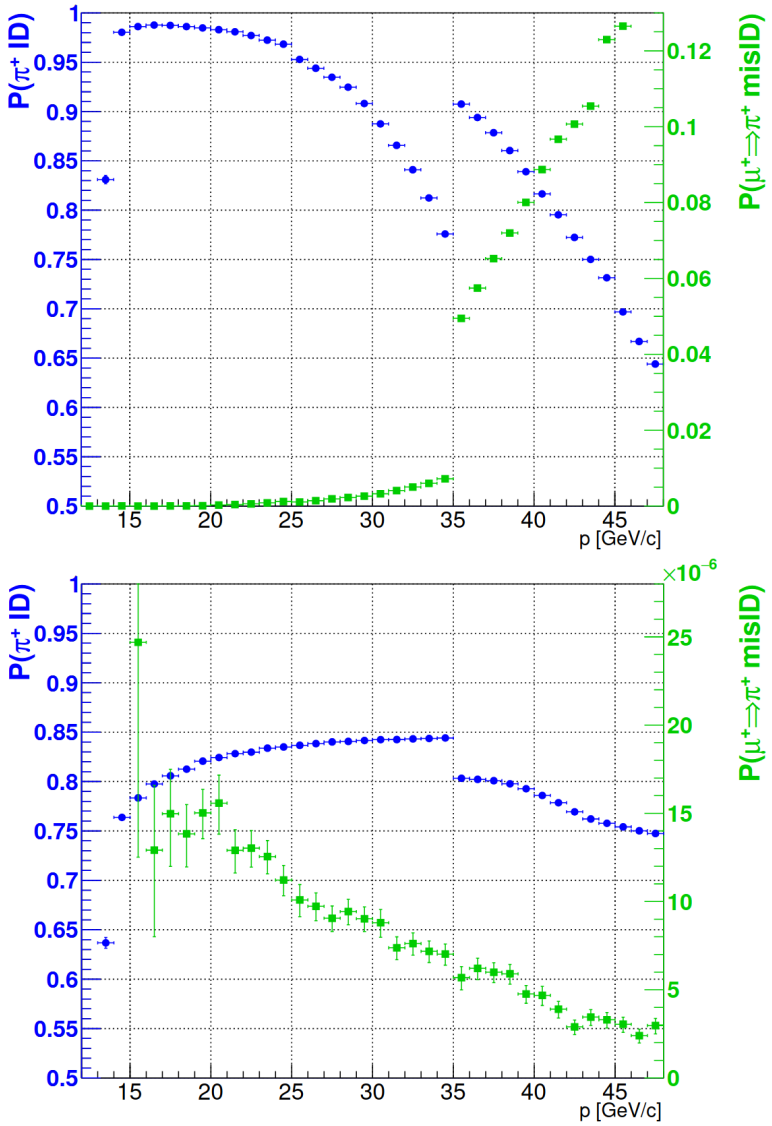

Figure 3.7: Particle identification performance: pion identification efficiencies (blue) and muon misidentification probabilities (green) for the RICH (top) and for the combination of the LKr and MUV detectors (bottom).

No signals must be present in the LAV within 3 ns of the track time. Only LAV stations downstream of the reconstructed decay vertex are considered for the veto.

An LKr cluster is identified as a candidate photon if it is farther than 100 mm from the π^+ track impact point on the LKr front plane, and is in time with the track time, within a time window that depends on the cluster energy, ranging between 5 and 12 ns. This time window was greatly reduced with respect to the Run1 analysis (in which it could reach up to 50 ns) as a result of a complete refurbishment of the LKr cluster reconstruction algorithm. The new algorithm is much more robust to clusters overlapping in space, whereas

the previous version could artificially shift the cluster times by tens of ns.

Both the TEL62 and the CREAM readouts are considered for the analysis of signals of the IRC and the SAC. Signals read by the TEL62 system must not be within 7 ns of the track time; specific conditions on the time-over-threshold are applied, in order to be resilient against pileup. The CREAM system must not read any signal corresponding to an energy larger than 1 GeV within 5 ns of the track time.

3.9.2 Multiplicity veto

The multiplicity veto conditions reject extra activity compatible with charged particles being present downstream of the fiducial volume. This is not only important for the rejection of kaon decays with more than one charged particle in the final state, but also for additional rejection of $K^+ \to \pi^+\pi^0$ decays (and decays involving photons) where photons convert in the downstream material (for example, near the RICH mirrors).

For all other STRAW tracks in the event, the closest distance of approach (CDA) to the candidate π^+ track is evaluated, and their track time is also calculated based on possible associations with other detectors. If any track is found with small CDA and with time close to the candidate π^+ time, the event is rejected. The thresholds of CDA and time difference for event rejection vary, depending on the timing information available from the association to the other detectors. Furthermore, a dedicated algorithm searches for partial tracks compatible with particles coming from kaon decays and not leaving signals in all four STRAW chambers.

Conditions on the multiplicity of signals not associated with the candidate π^+ in the CHOD are applied; extra coincident signals in at least two of the CHOD, NA48-CHOD, and LKr detectors, are also vetoed. Signals in MUV0, HASC or HASC2 are also considered as activity from extra particles and veto the event.

3.9.3 Performance

The performance of the π^0 rejection is evaluated using a $K^+ \to \pi^+\pi^0$ sample where the π^0 is tagged kinematically. Since the rejection is so powerful, in order not to be limited by statistics, the number of $K^+ \to \pi^+\pi^0$ events passing the photon and multiplicity vetoes is evaluated on data collected in the PNN trigger, while the number of total $K^+ \to \pi^+\pi^0$ events is, in fact, the same normalization used for the final $BR(K^+ \to \pi^+\nu\bar{\nu})$ extraction.

Overall, the π^0 rejection inefficiency is measured to be $(1.72 \pm 0.07) \times 10^{-8}$. The π^0 rejection inefficiency is plotted in figure 3.8 as function of π^+ momentum. The trend is explained by considering that in asymmetric $K^+ \to \pi^+\pi^0$ decays, where the π^0 has either large or small momentum, photons are most likely to be emitted with large or small angles: this leads them to point

to LAV or SAV, which have the largest photon detection inefficiency. Mid-range momentum π^0, instead, decay more frequently into photons pointing to the LKr, which is more powerful in vetoing them.

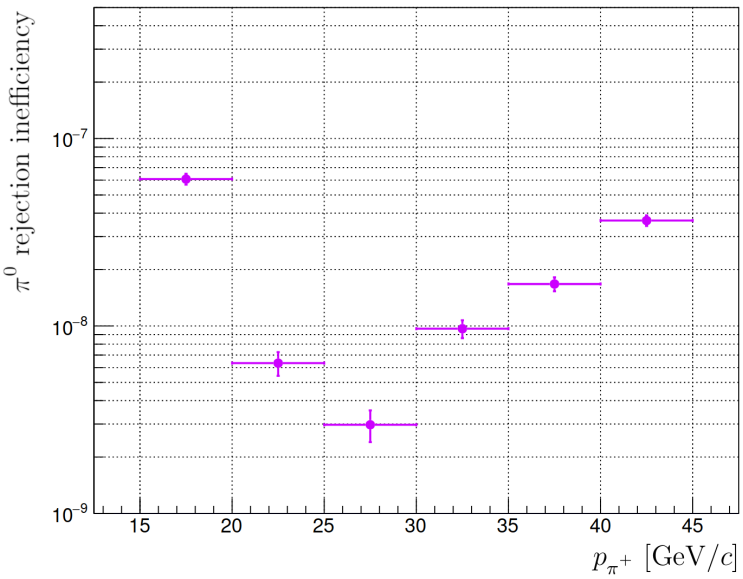

Figure 3.8: π^0 rejection inefficiency in $K^+ \to \pi^+\pi^0$ events as function of π^+ momentum. [35]

3.10 Kinematic selection

The last step in the selection of the signal is the selection in the *kinematic plane* defined by the variables p_{π^+} and m^2_{miss}.

In order to define regions in this plane, it is important to take into account the kinematic limit for the main kaon decays, as discussed in section 2.1, and the experimental resolution on m^2_{miss}, which is about $10^{-3}\,\text{GeV}^2/c^4$ for the $K^+ \to \pi^+\pi^0$ and $K^+ \to 3\pi$ decays. It is also useful to define the true value of the m^2_{miss} in $K^+ \to \mu^+\nu$ decays, which is function of the track momentum p_{π^+}:[8]

$$m^2_{\text{miss},\mu\nu} = \left(m^2_{\pi^+} - m^2_{\mu^+}\right)\left(1 - \frac{p_{K^+}}{p_{\pi^+}}\right),$$

where $p_{K^+} = 75\,\text{GeV}/c$. Furthermore, the experimental resolution on m^2_{miss} for $K^+ \to \mu^+\nu$ decays is about $\sigma_{\mu\nu} = 1.2 \times 10^{-3}\,\text{GeV}^2/c^4$.

Other than p_{π^+} and m^2_{miss}, two more variables, independent of the GTK and of the STRAW momenta, respectively, are used as measurements of the

[8] This formula is actually valid up to terms of $O\left(\left(m^4_{\pi^+} - m^4_{\mu^+}\right) p_{K^+}/p^3_{\pi^+}\right)$.

squared missing mass, in order to further control resolution effects on m^2_{miss} due to the spectrometers. These variables are:

$$m^2_{\text{miss,beam}} = (P_{\text{beam}} - P_{\pi^+})^2 ,$$
$$m^2_{\text{miss,RICH}} = (P_{\text{K}^+} - P_{\text{RICH}})^2 ,$$

where P_{beam} is the measured average 4-momentum of the beam, and P_{RICH} is the 4-momentum of the candidate π^+ where the magnitude of the 3-momentum is measured via the RICH ring radius (assuming the π^+ mass) rather than the STRAW. They have a worse resolution than m^2_{miss}, because either one of the spectrometers is not used.

With the notation discussed above, the conditions defining the two signal regions R1 and R2 are reported in table 3.2, while table 3.3 defines the background regions, used to tag kinematically the main kaon decays, and the control regions (CRs), used at the end of the analysis to validate the background estimations. A particular control region, CR3D, specifically used to cross-check the definition of the signal region in terms of $m^2_{\text{miss,beam}}$ and $m^2_{\text{miss,RICH}}$, is defined by considering events that satisfy both the p_{π^+} and m^2_{miss} criteria of R1 or R2, but do not satisfy either of the criteria on $m^2_{\text{miss,beam}}$ or $m^2_{\text{miss,RICH}}$. Therefore, it overlaps the signal regions in the kinematic plane (it is still kept independent from them by the different conditions on $m^2_{\text{miss,beam}}$ and $m^2_{\text{miss,RICH}}$). A graphical description of all the regions of the kinematic plane is shown in figure 3.9.

	R1			R2
m^2_{miss} (GeV$^2/c^4$)	$[0.000, 0.010]$			$[0.026, 0.068]$
p_{π^+} (GeV/c)	$[15, 20]$	$[20, 25]$	$[25, 35]$	$[15, 45]$
$m^2_{\text{miss,beam}}$ (GeV$^2/c^4$)	$[-0.005, 0.0135]$		$[0.000, 0.0135]$	$[0.024, 0.068]$
$m^2_{\text{miss,RICH}}$ (GeV$^2/c^4$)	$[0.000, 0.010]$	$[0.000, 0.020]$	$[-0.005, 0.020]$	$[0.020, 0.070]$

Table 3.2: Definition of the signal regions.

The events from the 2021–2022 dataset passing all the selection criteria in the background regions are shown in figure 3.10. Discontinuities around momentum values that are multiples of 5 GeV/c are a byproduct of the optimization of the selection in bins of π^+ momentum.

3.11 The normalization selection

As mentioned at the beginning of this chapter, the $\text{K}^+ \to \pi^+\pi^0_{\gamma\gamma}$ decay channel is used for normalization. The event selection for the normalization sam-

	p_{π^+} (GeV/c)	$m^2_{\rm miss}$ (GeV$^2/c^4$)
$\pi^+\pi^0$	$[15, 45]$	$[0.015, 0.021]$
$\mu^+\nu$	$[15, 45]$	$\left[m^2_{{\rm miss},\mu\nu} - 3\sigma_{\mu\nu}, m^2_{{\rm miss},\mu\nu} + 3\sigma_{\mu\nu}\right]$
3π	$[15, 45]$	> 0.072
CR1	$[15, 45]$	$[0.010, 0.015]$
CR2	$[15, 45]$	$[0.021, 0.026]$
CRmu	$[15, 45]$	$\left[m^2_{{\rm miss},\mu\nu} + 3\sigma_{\mu\nu}, \min\{m^2_{{\rm miss},\mu\nu} + 6\sigma_{\mu\nu}, 0\}\right]$
CRmu2	$[15, 45]$	$\left[\min\{m^2_{{\rm miss},\mu\nu} + 6\sigma_{\mu\nu}, 0\}, 0\right]$
CRmu3	$[35, 45]$	$[0, 0.010]$
CR3pi	$[15, 45]$	$[0.068, 0.072]$
CR3D	$[15, 45]$	same as signal regions; see text

Table 3.3: Definition of background regions and control regions.

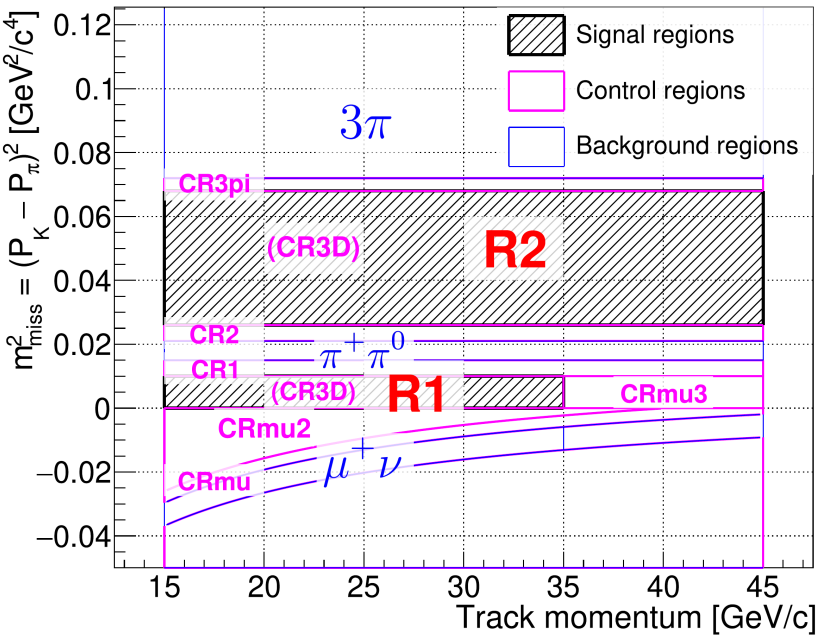

Figure 3.9: Definition of kinematic regions. CR3D overlaps R1 and R2 in the kinematic plane, but it has inverted conditions on $m^2_{\rm miss,beam}$ and $m^2_{\rm miss,RICH}$.

ple follows almost the same flow as the signal selection, which implies extensive cancellation of systematics. The only differences with respect to the signal selection are the following:

- the normalization selection is applied on data collected with the NORM trigger (rather than the PNN trigger);

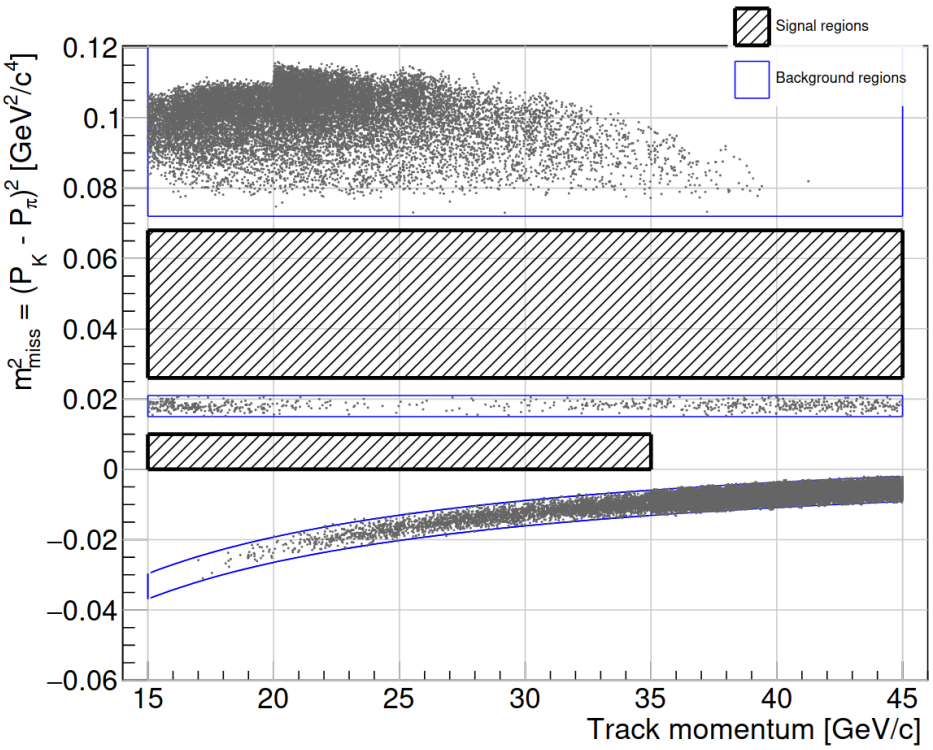

Figure 3.10: Events passing the $K^+ \to \pi^+\nu\bar{\nu}$ selection criteria in the background regions. [35]

- no photon or multiplicity veto criteria (as defined in section 3.9) are applied, except for a condition on the number of STRAW tracks forming a vertex, which suppresses the Dalitz decay $\pi^0 \to e^+e^-\gamma$;
- the kinematic region is defined by $0.010\,\mathrm{GeV}^2/c^4 < m^2_{\mathrm{miss}} < 0.026\,\mathrm{GeV}^2/c^4$, namely it is the union of CR1, the $\pi^+\pi^0$ background region, and CR2.

Chapter 4

Signal expectation

The number of expected signal events is calculated based on an estimation of the integrated kaon flux, and the efficiency of selecting signal events. The use of a normalization channel reduces the calculation of efficiencies, in fact, to a calculation of efficiency ratios, where some terms cancel, reducing the final uncertainty.

Section 4.1 explains the strategy for the calculation. Sections 4.2–4.5 are dedicated to the evaluation of the single factors that enter the calculation; results are given in section 4.6. Further considerations on the analysis performance are given in section 4.7.

4.1 Single event sensitivity

The number of normalization events $N_{\pi\pi}$, multiplied by the NORM trigger downscaling D_{NORM},[1] can, in principle, be used in order to estimate the number of kaon decays collected N_{K} via the BR of the $\mathrm{K}^+ \to \pi^+\pi^0_{\gamma\gamma}$ decay, denoted $\mathrm{BR}_{\pi\pi}$, and the total efficiency of the normalization selection $\varepsilon_{\pi\pi}$:

$$N_{\mathrm{K}} = \frac{N_{\pi\pi} D_{\mathsf{NORM}}}{\mathrm{BR}_{\pi\pi}\varepsilon_{\pi\pi}} \,. \tag{4.1}$$

The number of expected signal events, given an assumed signal branching ratio $\mathrm{BR}_{\pi\nu\bar{\nu}}$, is then calculated as

$$N_{\pi\nu\bar{\nu}} = N_{\mathrm{K}}\mathrm{BR}_{\pi\nu\bar{\nu}}\varepsilon_{\pi\nu\bar{\nu}} = N_{\pi\pi} D_{\mathsf{NORM}} \frac{\mathrm{BR}_{\pi\nu\bar{\nu}}}{\mathrm{BR}_{\pi\pi}} \frac{\varepsilon_{\pi\nu\bar{\nu}}}{\varepsilon_{\pi\pi}} \,,$$

where $\varepsilon_{\pi\nu\bar{\nu}}$ is the global efficiency of the signal selection. (The downscaling of the PNN trigger is equal to 1, and not made explicit in the equation.)

[1] Since the downscaling briefly changed during the data taking (see also footnote 6 on page 56), the product $N_{\pi\pi} D_{\mathsf{NORM}}$ should actually be understood as a weighted count of normalization events, where each event is weighted by the downscaling applied in its data taking period. An effective number of observed normalization events can be defined as $N_{\pi\pi,\mathrm{eff}} = N_{\pi\pi} D_{\mathsf{NORM}}/400$.

The *single event sensitivity* (SES) is defined as the proportionality constant between the branching ratio of the signal and the number of signal events expected $N_{\pi\nu\bar{\nu}}$:

$$\mathrm{SES} = \frac{\mathrm{BR}_{\pi\nu\bar{\nu}}}{N_{\pi\nu\bar{\nu}}} = \frac{1}{N_{\mathrm{K}}\varepsilon_{\pi\nu\bar{\nu}}} = \frac{\mathrm{BR}_{\pi\pi}}{N_{\pi\pi}D_{\mathsf{NORM}}}\frac{\varepsilon_{\pi\pi}}{\varepsilon_{\pi\nu\bar{\nu}}}\,.$$

The SES depends only on the experimental conditions and not on the signal branching ratio: in fact, it is used to extract the measured branching ratio from an observed number of signal events.

Each of the two efficiencies $\varepsilon_{\pi\pi}$ and $\varepsilon_{\pi\nu\bar{\nu}}$ is, in principle, the product of three factors: the trigger efficiency of the trigger line used to collect the relevant sample, the acceptance of the selection, and the efficiency factor accounting for accidental activity vetoing the event (*random veto*). However, for the calculation of the SES and the signal expectation, only the ratio $\varepsilon_{\pi\nu\bar{\nu}}/\varepsilon_{\pi\pi}$ is needed. Owing to the design of the trigger and to the choice of the normalization channel, many effects cancel out in the ratio, such as: random veto effects in the upstream background rejection, reconstruction inefficiencies in the downstream detectors, trigger conditions common to the NORM and PNN triggers. This has the effect of reducing the systematic uncertainty on the SES.

In practice, the efficiency ratio is expressed as

$$\frac{\varepsilon_{\pi\nu\bar{\nu}}}{\varepsilon_{\pi\pi}} = \frac{A_{\pi\nu\bar{\nu}}}{A_{\pi\pi}}\varepsilon_{\mathrm{trig}}\varepsilon_{\mathrm{RV}}\,,$$

where $A_{\pi\nu\bar{\nu}}$ and $A_{\pi\pi}$ are the acceptances of the signal and of the normalization selections respectively, $\varepsilon_{\mathrm{trig}}$ is the ratio between the trigger efficiencies of the NORM and of the PNN trigger, and $1-\varepsilon_{\mathrm{RV}}$ is the probability of a signal event being rejected by the photon or the multiplicity veto criteria because of accidental activity. Indeed, the photon and multiplicity veto criteria are the only ones that are applied only in the signal selection but not for the normalization, and do not cancel in the ratio. $\varepsilon_{\mathrm{RV}}$ is referred to as *random veto efficiency*.

The SES is therefore calculated as

$$\mathrm{SES} = \frac{\mathrm{BR}_{\pi\nu\bar{\nu}}}{N_{\pi\nu\bar{\nu}}} = \frac{\mathrm{BR}_{\pi\pi}}{N_{\pi\pi}D_{\mathsf{NORM}}}\frac{A_{\pi\pi}}{A_{\pi\nu\bar{\nu}}}\frac{1}{\varepsilon_{\mathrm{trig}}\varepsilon_{\mathrm{RV}}}\,, \tag{4.2}$$

and the number of expected SM signal events is

$$N^{\mathrm{SM}}_{\pi\nu\bar{\nu}} = N_{\mathrm{K}}\mathrm{BR}^{\mathrm{SM}}_{\pi\nu\bar{\nu}}\varepsilon_{\pi\nu\bar{\nu}} = N_{\pi\pi}D_{\mathsf{NORM}}\frac{\mathrm{BR}^{\mathrm{SM}}_{\pi\nu\bar{\nu}}}{\mathrm{BR}_{\pi\pi}}\frac{A_{\pi\nu\bar{\nu}}}{A_{\pi\pi}}\varepsilon_{\mathrm{trig}}\varepsilon_{\mathrm{RV}}\,. \tag{4.3}$$

To reflect the optimization of the selection, the analysis is performed in ($5\,\mathrm{GeV}/c$ wide) bins of π^+ momentum. It is therefore important to evaluate equations (4.2) and (4.3) in each bin of momentum, as the values in the single bins are the ones that ultimately enter the final BR measurement.

4.2 Number of normalization events

Counting the number of events passing the normalization selection gives a proportion of the number of collected kaon decays, as per equation (4.1).

The use of the $K^+ \to \pi^+\pi^0_{\gamma\gamma}$ decay as normalization channel has many advantages. First, its selection is similar to the signal: the only crucial difference is the application of photon and multiplicity vetoes, and, as a consequence, this is the only effect that does not cancel in the ratio of selection efficiencies. Second, its branching ratio is precisely known (relative uncertainty $O(10^{-3})$), and this reflects in a negligible contribution from $\mathrm{BR}_{\pi\pi}$ to the uncertainty on the SES. Third, its branching ratio is sufficiently large that the number of normalization events collected is abundant, therefore the relative statistical uncertainty on $N_{\pi\pi}$ is also negligible ($O(10^{-4})$). Finally, this decay can be selected with high purity, so that the normalization sample can be considered almost background-free: a data-MC comparison shows that it has a contamination of about 0.2 % from $K^+ \to \pi^+\pi^0$ decays followed by the Dalitz decay $\pi^0 \to e^+e^-\gamma$. This amount is therefore subtracted from $N_{\pi\pi}$, and a corresponding identical systematic uncertainty is assigned.

The result of the normalization selection is shown in figure 4.1 as function of π^+ momentum. In total,

$$N_{\pi\pi} D_{\mathsf{NORM}} = (7.815 \pm 0.016) \times 10^{10} \,.$$

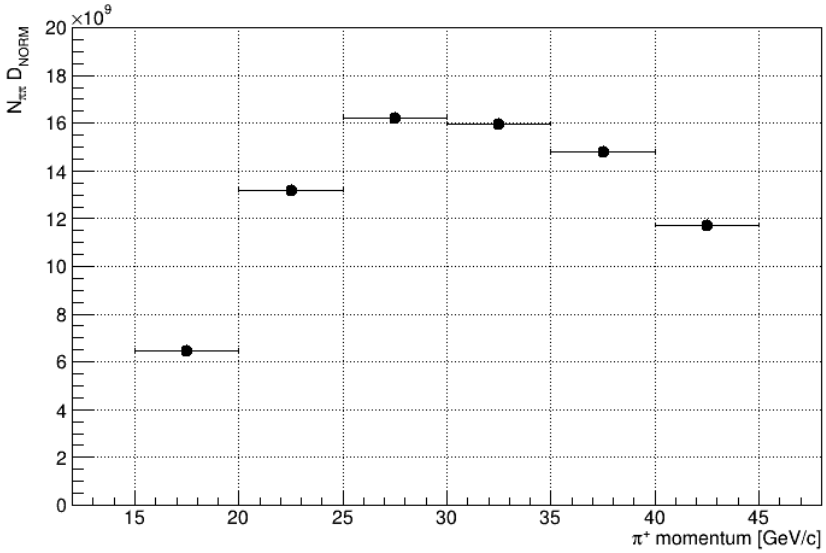

Figure 4.1: $N_{\pi\pi} D_{\mathsf{NORM}}$ as function of π^+ momentum.

4.3 Signal and normalization acceptances

The acceptance of the signal and of the normalization selections are evaluated on simulations. MC samples of $O(10^7)$ events are generated for both the decays $K^+ \to \pi^+\nu\bar{\nu}$ and $K^+ \to \pi^+\pi^0_{\gamma\gamma}$ happening in the fiducial volume: the fraction of events passing the full selections is evaluated. No intensity effects are simulated: only the decaying kaon is generated in each event. This prevents double-counting of the efficiency loss due to random veto, which is evaluated separately (in a data-driven way) via $\varepsilon_{\mathrm{RV}}$.

Figure 4.2 shows the distribution of the selected MC events, namely the acceptances, for both signal and normalization selections, in bins of π^+ momentum. The statistical error, due to the finite MC sample size, is negligible.

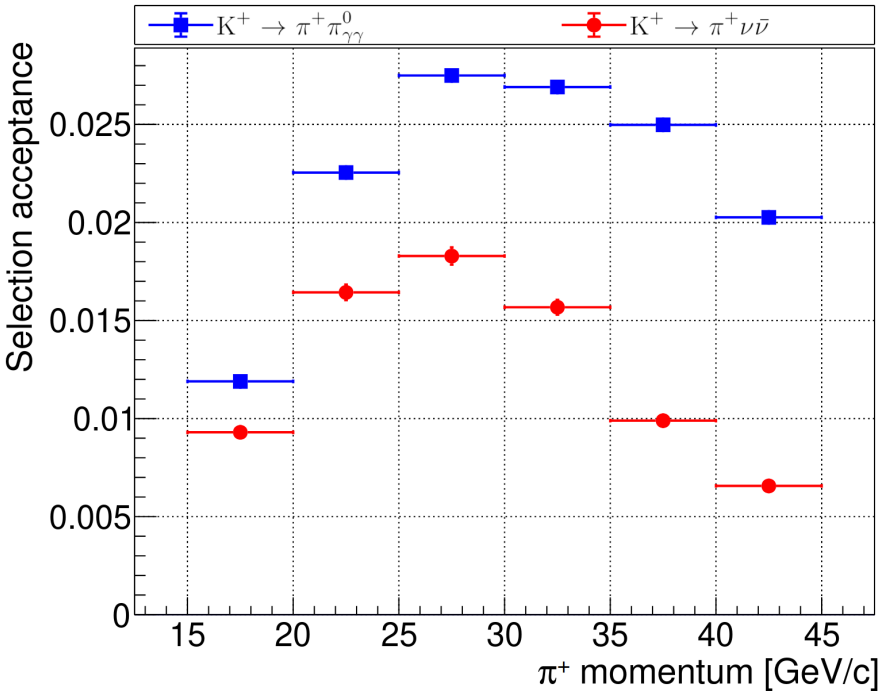

Figure 4.2: Acceptance of the normalization (blue) and signal (red) selections as function of π^+ momentum. Systematic uncertainties are not shown. [35]

Systematic uncertainties can be assigned to both acceptances for mismodelling in the simulation. In particular, data-MC comparisons of control $K^+ \to \pi^+\pi^+\pi^-$ samples show a relative difference of 5 % in the K^+–π^+ matching efficiency; similarly, a relative discrepancy of 1.5 % in the π^+ particle identification efficiency between data and MC has been observed. However, these systematic uncertainties are not propagated to the SES uncertainty, because they are fully correlated in the normalization and in the signal selection, and therefore cancel in the ratio.

An effect that does not cancel in the ratio is the mismodelling of the interactions of the π^+ with the detector material that cause a signal event

to be rejected by the photon or multiplicity criteria. For this reason, a relative systematic uncertainty of 2.8 %, evaluated comparing a data sample of $K^+ \to \pi^+\pi^0$ events with two photons in the LAV with a similar simulated sample, is assigned to $A_{\pi\nu\bar{\nu}}$ and not to $A_{\pi\pi}$, and therefore propagated to the SES.

Table 4.1 reports the summary of the acceptances and their uncertainties.

	Normalization	Signal
A, acceptance	0.134	0.0762
δA, statistical	$< 10^{-4}$	$< 10^{-4}$
δA, K^+–π^+ matching	0.007	0.0038
δA, π^+ identification	0.002	0.0011
δA, veto from π^+ interactions	—	0.0022
δA, total	0.007	0.0045
δA, propagated to SES	—	0.0022

Table 4.1: Summary of the acceptances and their uncertainties.

4.4 Trigger efficiencies

The trigger efficiencies for the NORM and the PNN trigger lines can be expressed as the product of the efficiencies of the single trigger requirements, as their application is performed independently:

$$\varepsilon_{\mathsf{PNN}} = \left(\varepsilon_{\mathsf{RICH}}\,\varepsilon_{\mathsf{CHOD}}\,\varepsilon_{\overline{\mathsf{MUV3}}}\,\varepsilon_{\mathsf{UTMC}}\,\varepsilon_{\overline{\mathsf{Qx}}}\,\varepsilon_{\overline{\mathsf{LKr}}}\right)\left(\varepsilon_{\mathsf{KTAG}}\,\varepsilon_{\overline{\mathsf{LAV}}}\,\varepsilon_{\mathsf{STRAW}}\,\varepsilon_{\mathsf{STRAW\text{-}PNN}}\right),$$
$$\varepsilon_{\mathsf{NORM}} = \left(\varepsilon_{\mathsf{RICH}}\,\varepsilon_{\mathsf{CHOD}}\,\varepsilon_{\overline{\mathsf{MUV3}}}\right)\left(\varepsilon_{\mathsf{KTAG}}\,\varepsilon_{\mathsf{STRAW}}\right),$$

where brackets separate L0 from L1 conditions. Each condition raises a hardware bit whose status is recorded in the data, therefore the efficiency of each condition can be measured in a reference sample where the trigger requirement is expected to be passed, and evaluating the fraction of events where the corresponding bit is raised.

The RICH, CHOD and $\overline{\mathsf{MUV3}}$ requirements are studied with $K^+ \to \pi^+\pi^0$ events collected with the CTRL trigger, and selected similarly to the normalization events. The total efficiency from these requirements is around 97 %; however, since they are applied equally to signal and normalization, these efficiencies cancel in the ratio, and do not affect the SES.

The UTMC and $\overline{\mathsf{Qx}}$ requirements are studied with the normalization sample, by additionally imposing the multiplicity veto conditions described in section 3.9.2. They are found to have a combined efficiency above 99 %.

The largest contribution to the L0 trigger inefficiency of the PNN trigger is from the $\overline{\text{LKr}}$ condition. This requirement is studied with a reference sample of $K^+ \to \pi^+\pi^0_{\gamma\gamma}$ events, collected with the NORM trigger, where both photons are detected in the LAV. In this sample the trigger efficiency is studied as a function of the energy deposition in the LKr, E_{LKr}. In order to convert such dependence into a dependence on the π^+ momentum p_{π^+}, and to account for the different p_{π^+} bias of the reference sample with respect to the signal (the reference sample is biased towards low momentum photons and therefore towards high p_{π^+}), the E_{LKr}/p_{π^+} distribution as function of p_{π^+} is exploited. This distribution is measured from a sample of pions selected from $K^+ \to \pi^+\pi^+\pi^-$ decays. The result of this evaluation is $\varepsilon_{\overline{\text{LKr}}} = (90.1 \pm 1.5)$ %, where the uncertainty is purely systematic, assigned to account for the description of the momentum dependence by considering the momentum dependence of $\varepsilon_{\overline{\text{LKr}}}$ itself.

The L1 trigger conditions are studied with a sample of $K^+ \to \mu^+\nu$ decays, collected with the MB trigger (which has no L1 conditions), and selected similarly to the signal but exploiting particle identification and kinematics to tag the muon (the sample is the same as the one used for the ε_{RV} measurement, see section 4.5). The KTAG and STRAW conditions, which are common in the NORM and PNN trigger and cancel in the ratio, have a combined inefficiency of less than 1 %. The STRAW-PNN condition is also found to have an inefficiency smaller than 1 %. The largest contribution to the L1 trigger inefficiency comes from the $\overline{\text{LAV}}$ condition, which is found to have about 97 % efficiency, with negligible statistical uncertainty.

Finally, the ratio $\varepsilon_{\text{trig}} = \varepsilon_{\text{PNN}}/\varepsilon_{\text{NORM}}$, which enters the SES, is calculated as the product of the efficiencies of the conditions that are not common between the two trigger lines:

$$\varepsilon_{\text{trig}} = \varepsilon_{\text{UTMC}}\, \varepsilon_{\overline{\text{Qx}}}\, \varepsilon_{\overline{\text{LKr}}}\, \varepsilon_{\overline{\text{LAV}}}\, \varepsilon_{\text{STRAW-PNN}} \,.$$

The trigger efficiencies are reported in table 4.2, and $\varepsilon_{\text{trig}}$ is plotted in figure 4.3 as function of π^+ momentum.

	Efficiency
$\varepsilon_{\text{NORM}}$	96.7 %
ε_{PNN}	(83.4 ± 1.4) %
$\varepsilon_{\text{trig}}$	(85.9 ± 1.4) %

Table 4.2: Trigger efficiencies. Statistical uncertainties are all negligible.

A control sample of $K^+ \to \mu^+\nu$ decays, collected with the PNN trigger, is obtained by applying the signal selection apart from the RICH particle identification criteria and selecting the $\mu^+\nu$ kinematic region. Applying the same

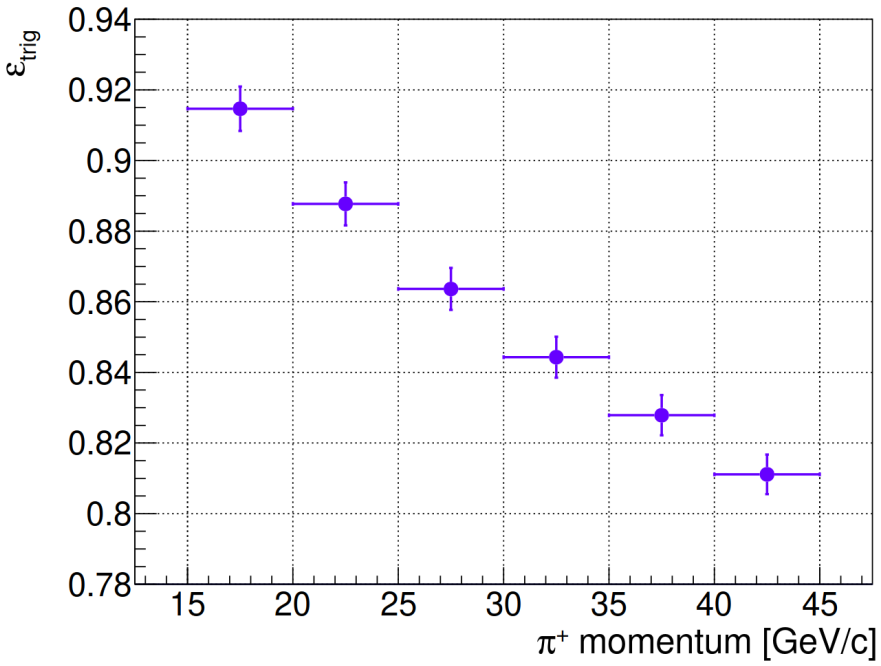

Figure 4.3: Trigger efficiency ratio as function of π^+ momentum. [35]

selection to data collected with the MB trigger, and rescaling by the downscaling of the MB trigger and the ratio of the trigger efficiencies $\varepsilon_{\mathsf{PNN}}/\varepsilon_{\mathsf{MB}}$ (where $\varepsilon_{\mathsf{MB}} = \varepsilon_{\mathsf{RICH}}\,\varepsilon_{\mathsf{CHOD}}$), an expected number of events in the control sample is calculated. An excellent agreement is found, as shown in figure 4.4. The agreement is especially significant because the ratio $\varepsilon_{\mathsf{PNN}}/\varepsilon_{\mathsf{MB}}$ contains more trigger components than $\varepsilon_{\text{trig}}$. For this reason, no further systematic uncertainties are applied to $\varepsilon_{\text{trig}}$.

4.5 Random veto efficiency

The random veto efficiency quantifies the effect of accidental vetoes due to the instantaneous intensity. It is intrinsically independent from the π^+ momentum, as it is by definition due to activity that is uncorrelated with the π^+ itself. However, it does depend on the beam intensity. As discussed in section 4.1, ε_{RV} is defined as the random veto efficiency due only to the photon veto and multipicity criteria described in section 3.9, which are the only ones that are not common between the signal and the normalization selections, and therefore can give efficiency loss effects that enter the SES (random veto effects due to all the other criteria cancel in the ratio).

The random veto efficiency is evaluated using a control sample of $K^+ \to \mu^+\nu$ decays. This decay is chosen because it can be selected in a similar way to the signal, as it shares the key feature of having a single charged track. The sample is selected from the MB trigger line, by applying the full signal selection with modified particle identification and kinematic criteria:

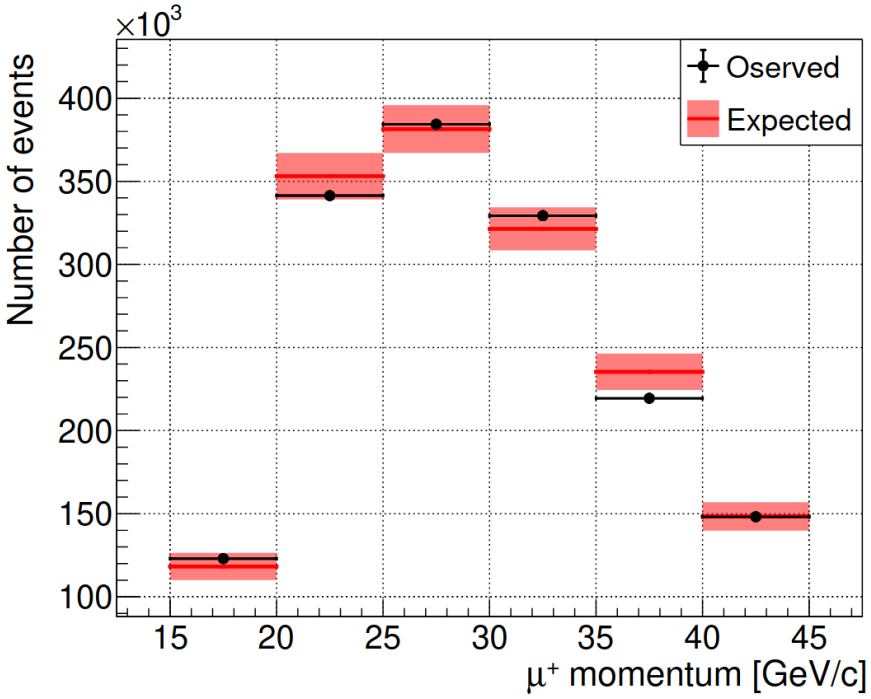

Figure 4.4: Validation of the trigger efficiency ratio measurement.

a MUV3 signal is required to be in time with the track within 5 ns; the calorimetric BDT classifier must report a probability of the particle being a muon larger than 99 %; the event must lie inside the $\mu^+\nu$ kinematic region. The fraction of these events passing the photon veto and multiplicity criteria is a measurement of the random veto efficiency.

Some effects due to the muon itself, rather than to accidental activity, can veto the event. One such effect is the muon producing delta rays in the RICH mirrors, leading to activity in the LAV and in the CHOD; a more important effect, especially for Run2 data, is that the muon can go into the acceptance of the HASC2, vetoing the event without any additional activity. Since the latter case would be a significant systematic effect in the $\varepsilon_{\mathrm{RV}}$ measurement, events for which the muon track points to the HASC2 are removed from the sample.

In order to quantify the residual effect of activity due to the muon itself, the same procedure is also applied to a MC sample of $\mathrm{K}^+ \to \mu^+\nu$ decays at zero effective intensity, obtaining a value denoted $\varepsilon_{\mathrm{RV}}^{\mathrm{MC}} = (98.8 \pm 0.1)$ %. Since this MC sample contains no pileup event, $1 - \varepsilon_{\mathrm{RV}}^{\mathrm{MC}}$ is used as a measure of the contribution from the muon to the random veto inefficiency. Denoting the value obtained from data with $\varepsilon_{\mathrm{RV}}^{\mathrm{data}}$, the final random veto efficiency is calculated by

$$\varepsilon_{\mathrm{RV}} = \varepsilon_{\mathrm{RV}}^{\mathrm{data}} / \varepsilon_{\mathrm{RV}}^{\mathrm{MC}} \,, \tag{4.4}$$

in the assumption that the veto from the muon and the veto from accidental activity are independent.

As discussed, $\varepsilon_{\mathrm{RV}}$ should not depend on the track momentum, as it is fundamentally independent from the track; however, $\varepsilon_{\mathrm{RV}}^{\mathrm{data}}$ and $\varepsilon_{\mathrm{RV}}^{\mathrm{MC}}$ can, and do, depend on the track momentum, precisely because they include the effects due to the muon itself. After performing the ratio in equation (4.4), the cancellation of the momentum dependence is found not to be complete, and a relative systematic uncertainty of 1 % is assigned to account for this.

The final result is

$$\varepsilon_{\mathrm{RV}} = (63.2 \pm 0.6)\ \% .$$

The contributions of the single photon vetoes, as well as the total random veto efficiency, are plotted in figure 4.5 as a function of instantaneous intensity.

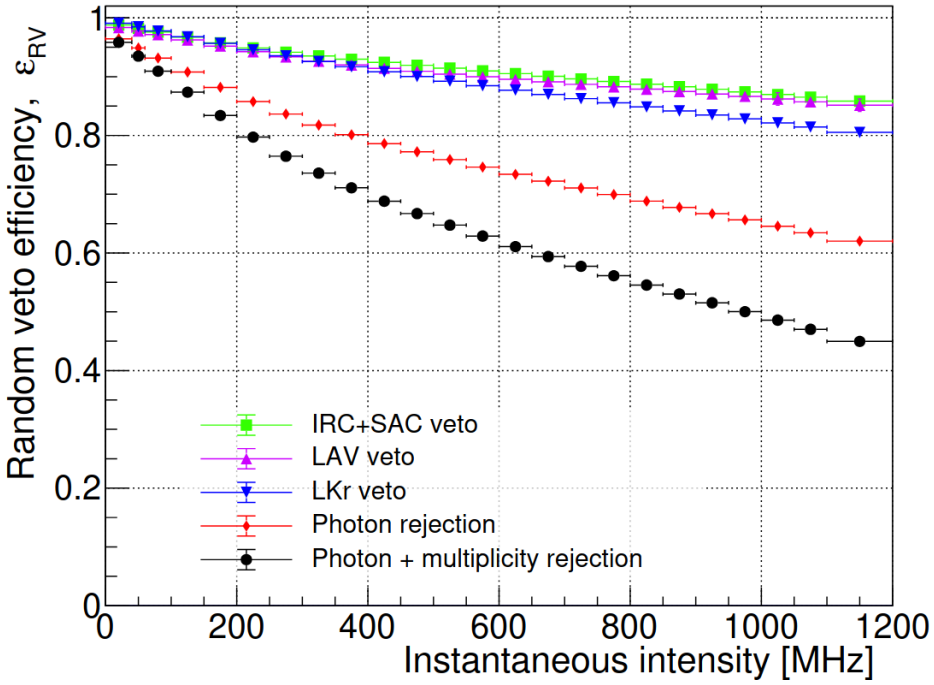

Figure 4.5: Random veto efficiency as a function of instantaneous intensity. [35]

4.6 Results

Despite equation (4.1) giving an expression for the total number of kaon decays in the fiducial volume, it is impractical to evaluate, as it contains random veto effects from the normalization selection. An *effective* number of kaon decays can be more easily calculated as

$$N_{\mathrm{K,eff}} = \frac{N_{\pi\pi} D_{\mathrm{NORM}}}{\mathrm{BR}_{\pi\pi} A_{\pi\pi}} = (2.86 \pm 0.15) \times 10^{12} ,$$

keeping in mind that it does not account for either the trigger efficiency nor the random veto effects from the normalization selection, and that the largest

contribution to the uncertainty is from the systematic uncertainty on $A_{\pi\pi}$, which does not enter the SES.

The SES and its inputs, evaluated over the whole range of π^+ momentum, are reported in table 4.3. It is important to stress that all the relevant uncertainties are systematic, with the statistical uncertainties being negligible. The dominant contribution to the uncertainty comes from the mismodelling effects in $A_{\pi\nu\bar{\nu}}$.

$N_{\pi\pi} D_{\mathsf{NORM}}$	Normalization events	$(7.815 \pm 0.016) \times 10^{10}$
$\mathrm{BR}_{\pi\pi}$	$\mathrm{BR}(\mathrm{K}^+ \to \pi^+\pi^0) \times \mathrm{BR}(\pi^0 \to \gamma\gamma)$ [27]	(20.43 ± 0.09) %
$A_{\pi\nu\bar{\nu}}/A_{\pi\pi}$	Acceptance ratio	(56.8 ± 1.6) %
$\varepsilon_{\mathrm{trig}}$	Trigger efficiency ratio	(85.9 ± 1.4) %
$\varepsilon_{\mathrm{RV}}$	Random veto efficiency	(63.2 ± 0.6) %
SES	Single event sensitivity	$(0.848 \pm 0.029) \times 10^{-11}$

Table 4.3: SES calculation.

In order to compute the number of expected SM signal events according to equation (4.3), $\mathrm{BR}^{\mathrm{SM}}_{\pi\nu\bar{\nu}} = 8.4 \times 10^{-11}$ [29] is used[2], despite not being the most up to date prediction, in order to facilitate comparisons with the Run1 analysis [34]. The total number is

$$N^{\mathrm{SM}}_{\pi\nu\bar{\nu}} = 9.91 \pm 0.34\,;$$

the computation in bins of π^+ momentum is reported in table 4.4.

p_{π^+} bin [GeV/c]	15–20	20–25	25–30	30–35	35–40	40–45
$N_{\pi\pi} D_{\mathsf{NORM}} \times 10^{-9}$	6.45(2)	13.13(3)	16.17(3)	15.92(3)	14.78(3)	11.70(2)
$A_{\pi\pi} \times 10^2$	1.190(2)	2.255(2)	2.750(3)	2.691(3)	2.498(2)	2.026(2)
$A_{\pi\nu\bar{\nu}} \times 10^2$	0.93(2)	1.64(4)	1.83(4)	1.57(4)	0.99(2)	0.66(2)
$\varepsilon_{\mathrm{trig}} \times 10^2$	91.5(6)	88.8(6)	86.4(6)	84.4(6)	82.8(6)	81.1(6)
$\varepsilon_{\mathrm{RV}} \times 10^2$			63.2 ± 0.6			
$N^{\mathrm{SM}}_{\pi\nu\bar{\nu}}$	1.20(4)	2.21(7)	2.41(7)	2.03(6)	1.26(4)	0.80(3)

Table 4.4: Calculation of the number of expected SM signal events. Only uncancelled uncertainties are quoted for the acceptances.

[2]Since this number does not enter the final BR extraction, it is assumed without uncertainty for simplicity.

4.7 Comparison with Run1 and beam intensity considerations

Table 4.5 shows a comparison of the relative uncertainty on the SES inputs between the Run1 analysis (in particular, in the 2018 S2 subsample, which is the most sensitive sample in Run1, accounting for 60 % of the total signal yield) and the present analysis. The relative uncertainty on the number of normalization events is negligible. The largest contribution to the uncertainty used to be from the trigger efficiency, which has been greatly improved by using a dedicated trigger mask for the normalization, allowing for cancellation of the most critical contributions. Systematic uncertainties on $\varepsilon_{\mathrm{RV}}$ and on the acceptance ratio have also been decreased by using improved evaluation strategies. Overall, the relative uncertainty on the SES has improved by a factor 1.8.

	2018 S2	2021–2022
$\frac{\delta(A_{\pi\nu\bar{\nu}}/A_{\pi\pi})}{A_{\pi\nu\bar{\nu}}/A_{\pi\pi}}$	3.5 %	2.8 %
$\frac{\delta\varepsilon_{\mathrm{trig}}}{\varepsilon_{\mathrm{trig}}}$	5.6 %	1.6 %
$\frac{\delta\varepsilon_{\mathrm{RV}}}{\varepsilon_{\mathrm{RV}}}$	1.5 %	0.9 %
$\frac{\delta\mathrm{SES}}{\mathrm{SES}}$	6.4 % (6.3 % for Run1 combined)	3.5 %

Table 4.5: Comparison between SES uncertainties of 2018 S2 [34] and the present analysis.

A comparison of the SES inputs between Run1 and the present analysis is reported in table 4.6. The acceptances of the signal and normalization selections have been improved by a relative amount of 15 % and 20 %, respectively. The trigger efficiencies are compatible if compared with their uncertainties; however, some reduction in 2021–2022 is expected due to intensity effects, such as random veto from the $\overline{\mathsf{MUV3}}$ and the $\overline{\mathsf{LAV}}$ conditions. The random veto efficiency has been kept comparable to the Run1 value by retuning veto conditions, despite the significantly larger intensity which otherwise would have had a larger impact. In the end, the total signal yield for the present analysis is expected to be very similar to the signal yield of the Run1 data analysis.

The analysis of the signal yield (of the present analysis) as function of beam intensity results in the plot shown in figure 4.6, where the slope near

	Run1		2021–2022
$A_{\pi\nu\bar{\nu}}$	(2018 S2)	6.37 %	7.62 %
$A_{\pi\pi}$		11.8 %	13.4 %
$\varepsilon_{\text{trig}}$		89 %	85.9 %
ε_{RV}		66 %	63.2 %
$N_{\text{K,eff}}$	(Run1 combined)	2.7×10^{12}	2.86×10^{12}
$N_{\pi\nu\bar{\nu}}^{\text{SM}}$		10.01	9.91

Table 4.6: Comparison of the SES of the Run1 analysis [34] and the present analysis. Uncertainties are not reported (see table 4.5).

0 intensity depends mainly on the acceptance of the selection, and the saturation is due to the a combination of data acquisition effects and offline selection, which both worsen at higher intensity, and can be characterized as contributing to a paralyzable dead time. This analysis showed that the optimal intensity for the NA62 apparatus is about 450 MHz, or 75 % of the nominal: since August 2023, this has been adopted as the new reference value for the intensity.

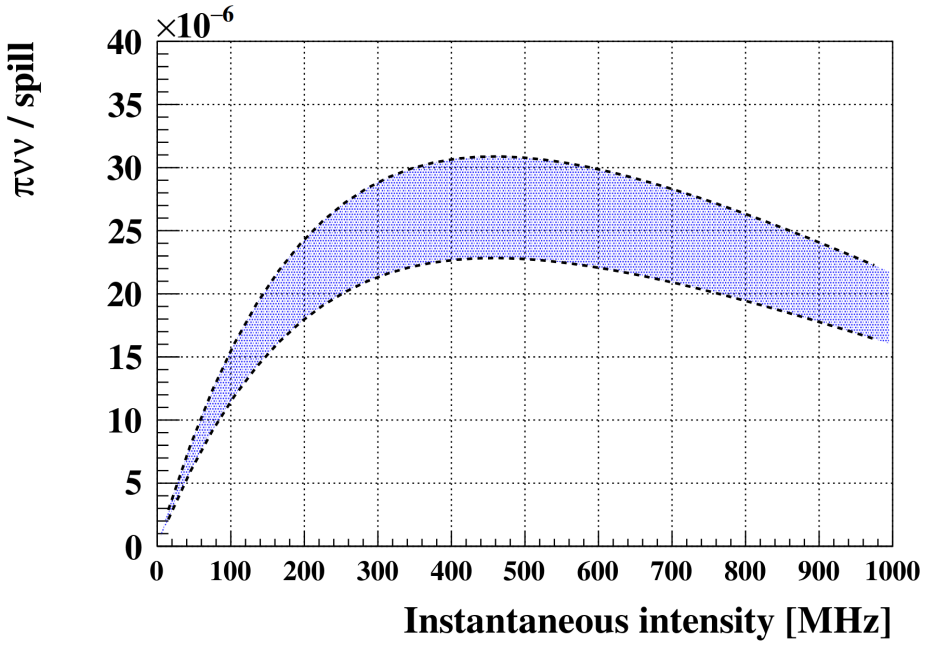

Figure 4.6: SM signal yield per spill as a function of beam intensity. [80]

Chapter 5

Background estimation

The estimation of background contamination in the signal regions is the counterpart to the evaluation of the SES for the extraction of the BR from the number of observed events. In a simplified discussion, if the background were perfectly known, it could be simply subtracted from the observation in order to calculate the number of signal events, which could then be converted into a measured BR via the SES. In practice, the Poissonian fluctuation and the uncertainty of the estimation cannot be eliminated. It is therefore crucial to develop and validate procedures for background estimation which are as precise as possible. This is accomplished by using a data-driven approach whenever possible: using simulated samples as little as possible reduces systematic uncertainties on the mismodelling of the backgrounds. It is in fact generally difficult to quantify the mismodelling of the mechanisms that can make a background event mimic a signal as rare as $O(10^{-10})$.

Background contamination generally arises in two different ways. The conceptually simpler class of backgrounds stems from the misreconstruction of the final state of a K^+ decay in the fiducial volume, and is described in section 5.1. A different background arises from misreconstruction of the initial state, where the detected π^+ originates upstream of the fiducial volume: this is the upstream background, which differs from all the other backgrounds in both the mechanism and in the estimation strategy. The upstream background is discussed in section 5.2. A summary of the backgrounds is presented in section 5.3.

5.1 Kaon decays in the fiducial volume

Each kaon decay can contaminate the signal region with its own mechanism; however, two categories for these backgrounds can be identified.

A first category consists of those decays that are kinematically not allowed in the signal region. In this case, m^2_{miss} must be misreconstructed in order for the events to contaminate the signal. This can happen via the non-gaussian tails of the m^2_{miss} resolution, which are due to effects such as elastic

interactions in GTK or STRAW, incorrect matching between K^+ and π^+ candidate tracks, position mismeasurements in the spectrometers, etc.

For this category of decays, the estimation strategy is based on the assumption that the kinematic rejection is independent from the rest of the selection: this is generally true as m^2_{miss} depends mostly on the 3-momenta as measured by the spectrometers, while photon vetoes, multiplicity vetoes and particle identification use information from independent detectors. In this assumption, a normalization factor (which takes into account the misreconstruction of the final state) can be obtained by counting events in a kinematic region where the background decay is dominant; an appropriate rescaling factor, also called the *kinematic tail fraction*, can be estimated from suitable control samples to account for the probability of misreconstructing m^2_{miss}.

The second category of background events from kaon decays in the fiducial volume consists of kaon decays that are kinematically allowed in the signal region. Owing to the design of the experiment and of the signal regions, these decays have smaller BR, or smaller probability of having their final state misreconstructed. Generally, these backgrounds can be estimated with a strategy which is analogous to the signal estimation described in section 4.1, by relying on simulated samples to give a measurement of the acceptance of the selection for these decays.

Background kaon decays, together with their respective mechanisms for entering the signal region, are listed in table 5.1 and discussed in the following.

Background	BR	Mechanism
$K^+ \to \pi^+\pi^0(\gamma)$	20.7 % [27]	m^2_{miss} misreco, γ veto inefficiency
$K^+ \to \mu^+\nu(\gamma)$	64.2 % [27]	m^2_{miss} misreco, μ^+ misidentification
$K^+ \to \pi^+\pi^+\pi^-$	5.6 % [27]	m^2_{miss} misreco, multiplicity veto inefficiency
$K^+ \to \pi^+\pi^-e^+\nu$	4.2×10^{-5} [27]	Multiplicity veto inefficiency
$K^+ \to \pi^0\ell^+\nu$	8.4 % [27]	ℓ^+ misidentification, γ veto inefficiency
$K^+ \to \pi^+\gamma\gamma$	9.6×10^{-7} [90]	γ veto inefficiency

Table 5.1: Summary of background sources from K^+ decays in the fiducial volume.

5.1.1 $K^+ \to \pi^+\pi^0(\gamma)$

$K^+ \to \pi^+\pi^0$ decays can enter the signal region by m^2_{miss} misreconstruction together with an inefficiency of the photon veto system which lets the photons from the π^0 decay go undetected. As discussed in section 5.1, the two

effects are assumed to be independent. Therefore, the number of expected background events in any region R of the kinematic plane is calculated as

$$N_{\pi^+\pi^0}(R) = N(\pi^+\pi^0) f_{\pi^+\pi^0}(R)\,, \tag{5.1}$$

where $N(\pi^+\pi^0)$ is the number of observed events in the $\pi^+\pi^0$ region after the signal selection (apart from the kinematic requirements), $f_{\pi^+\pi^0}(R)$ is the $K^+ \to \pi^+\pi^0$ kinematic tail fraction for region R, and the equation is understood to be evaluated in bins of track momentum.

The kinematic tail fraction is estimated using a control sample that is independent from the signal. Specifically, the $K^+ \to \pi^+\pi^0$ control sample is selected from the NORM trigger by tagging the two photons from $\pi^0 \to \gamma\gamma$ in the LKr: this constitutes an effectively background-free sample of $K^+ \to \pi^+\pi^0$ decays without any constraint on the kinematic variables. $f_{\pi^+\pi^0}(R)$ is evaluated by taking the ratio of the number of events in this sample that fall in region R by the number of events in the same sample that fall in the $\pi^+\pi^0$ region. Figure 5.1 shows the distribution of m^2_{miss} in the control sample: the kinematic tail fraction for the signal region is obtained by taking the ratio of the total events in R1 and R2 (red in the figure) to the events in the $\pi^+\pi^0$ region (yellow in the figure). A systematic uncertainty, estimated via data-MC comparisons and of the order of few percent depending on the momentum bin, is applied to account for a possible bias of the m^2_{miss} distribution in the control sample.

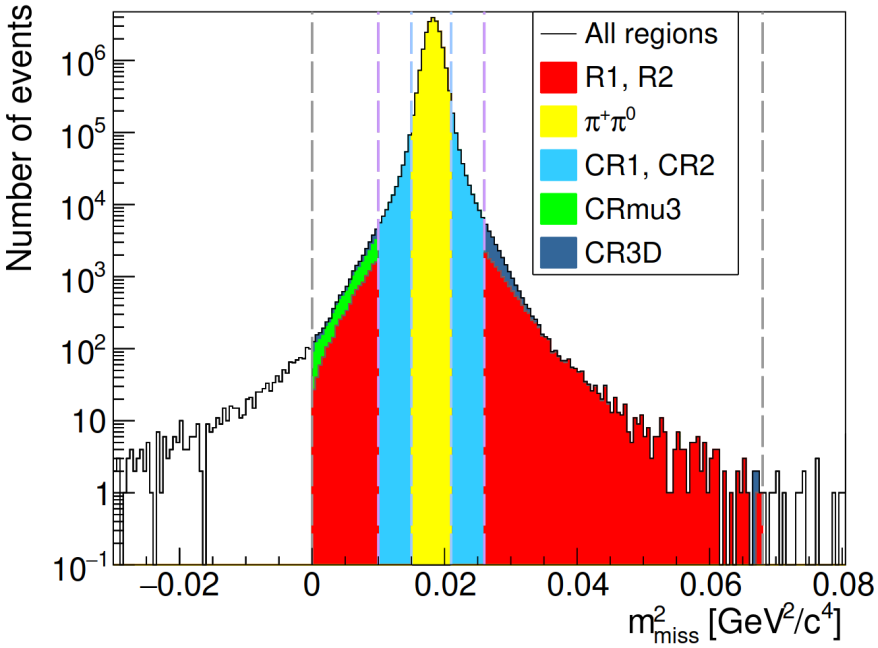

Figure 5.1: Distribution of m^2_{miss} in the $K^+ \to \pi^+\pi^0$ control sample, with breakdown into kinematic regions. [35]

After the signal selection, $N(\pi^+\pi^0) = 722$ is found in the whole $\pi^+\pi^0$ region. On average, $f_{\pi^+\pi^0}$ in the signal region is equal to $(1.20 \pm 0.01) \times 10^{-3}$; its uncertainty is dominated by the small statistics of $K^+ \to \pi^+\pi^0$ control events in the signal region. The total number of expected $K^+ \to \pi^+\pi^0$ background events, when summing the result of equation (5.1) over the full momentum range, is estimated to be 0.76 ± 0.04, where the uncertainty is mostly from the Poissonian fluctuation of $N(\pi^+\pi^0)$.

Since, in the $K^+ \to \pi^+\pi^0$ control sample, the π^0 is required to be fully reconstructed in the LKr by the presence of exactly two electromagnetic clusters (LAV and SAV photon vetoes are applied), radiative $K^+ \to \pi^+\pi^0\gamma$ decays where the radiative photon is detected in LAV, LKr, IRC or SAC are not included in the background estimation[1]. The kinematic rejection of these decays is degraded by the presence of an additional photon, but the photon veto rejects them more strongly than $K^+ \to \pi^+\pi^0$. It has been shown, by using simulated samples and single photon detection inefficiencies evaluated via a tag-and-probe method [98], that the rejection of $K^+ \to \pi^+\pi^0\gamma$ is about 30 times stronger than $K^+ \to \pi^+\pi^0$ [65]. This factor can be applied to a kinematic tail fraction calculated from simulated $K^+ \to \pi^+\pi^0\gamma$ events (to which a 100 % systematic uncertainty is assigned), to obtain a number of expected $K^+ \to \pi^+\pi^0\gamma$ background events of 0.07 ± 0.01 in the signal region.

5.1.2 $K^+ \to \mu^+\nu$

$K^+ \to \mu^+\nu$ can contaminate the signal region if m^2_{miss} is misreconstructed, and the μ^+ is misidentified as a π^+. The two effects are assumed independent, and the number of background $K^+ \to \mu^+\nu$ events in a region R is given by

$$N_{\mu^+\nu}(R) = N(\mu^+\nu) f_{\mu^+\nu}(R)\,,$$

where, analogously to equation (5.1), $N(\mu^+\nu)$ is the number of observed events in the $\mu^+\nu$ region after the signal selection (apart from the kinematic requirements), $f_{\mu^+\nu}(R)$ is the $K^+ \to \mu^+\nu$ kinematic tail fraction for region R, and the equation is evaluated in bins of track momentum.

The $K^+ \to \mu^+\nu$ control sample from which the kinematic tail fractions are evaluated is obtained from the MB trigger by using particle identification to tag the muon. In particular, the downstream particle must be identified as a muon by the LKr and MUV systems. The RICH-based particle identification depends on m_{RICH}, which in turn depends on the STRAW track momentum (see also section 3.8.1): therefore, in order to keep the kinematics as similar as possible to the signal, the criteria on RICH signals are left unchanged with respect to the signal selection. In other words, the RICH is required to tag the downstream track as a pion in the $K^+ \to \mu^+\nu$ control sample, and

[1] In the following, $K^+ \to \pi^+\pi^0\gamma$ actually denotes only this specific class of radiative decays.

this prevents biasing the m^2_{miss} spectrum of this sample with respect to the $K^+ \to \pi^+\nu\bar{\nu}$ selection.

The m^2_{miss} distribution of the $K^+ \to \mu^+\nu$ control sample is shown in figure 5.2. Data-MC comparisons show a 35 % disagreement on the kinematic tail fraction in the CRmu region, which is assigned as relative systematic uncertainty to $f_{\mu^+\nu}$. Events where the muon decays upstream of the RICH are not included in the control sample because of the calorimetric particle identification, therefore they are studied with MC samples: their contribution to this background is found to be negligible ($< 10^{-2}$).

A total of $N(\mu^+\nu) = 31372$ events are found in the $\mu^+\nu$ region after the signal selection, with a very strong bias towards high momentum due to the RICH particle identification, which is stronger in rejecting muons at low momenta. The average kinematic tail fraction in the signal region is $(1.6 \pm 0.6) \times 10^{-5}$, and the total number of expected background events is 0.87 ± 0.19.

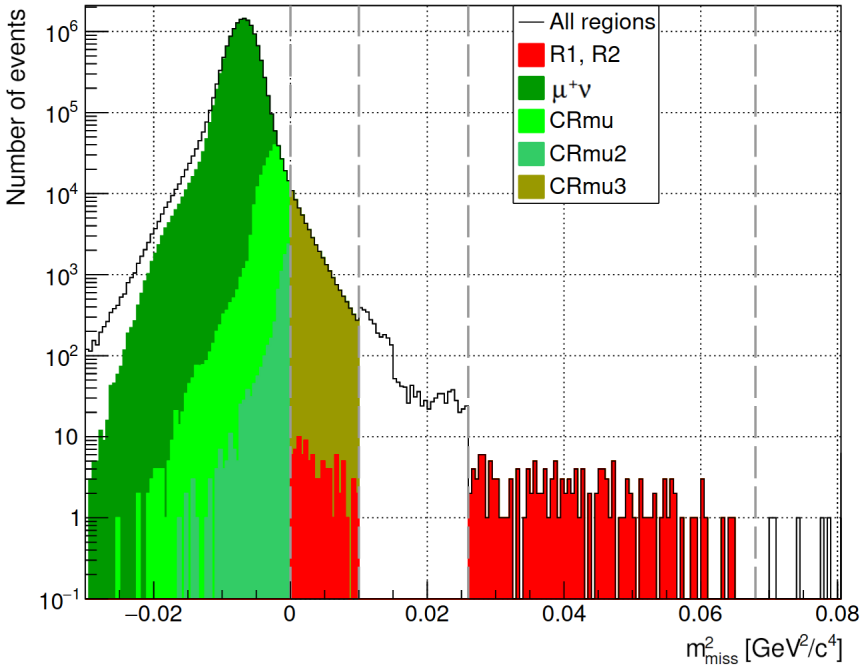

Figure 5.2: Distribution of m^2_{miss} in the $K^+ \to \mu^+\nu$ control sample, with breakdown into kinematic regions. [35]

5.1.3 $K^+ \to \mu^+\nu\gamma$

Unlike the $K^+ \to \pi^+\pi^0$ control sample described in section 5.1.1, the $K^+ \to \mu^+\nu$ control sample described in section 5.1.2 does not have modifications of the photon veto requirements with respect to the standard signal selection. Therefore, it includes radiative decays, which consequently enter the $K^+ \to \mu^+\nu$ background estimation.

However, a specific class of $K^+ \to \mu^+\nu\gamma$ events eludes the above estimation. These are the events, denoted $K^+ \to (\mu^+\gamma)\nu$ in the following, where the muon and the photon hit the LKr close to each other (at a distance $\lesssim 20\,\text{mm}$), forming a single cluster which has an enhanced probability of being misidentified as a pion. A sketch of such an event is shown in figure 5.3. Figure 5.4 shows the event display of a $K^+ \to (\mu^+\gamma)\nu$ event in the LKr, where it is rather clear that a MIP-like deposit near the STRAW track extrapolation is close to a broad photon-like cluster.

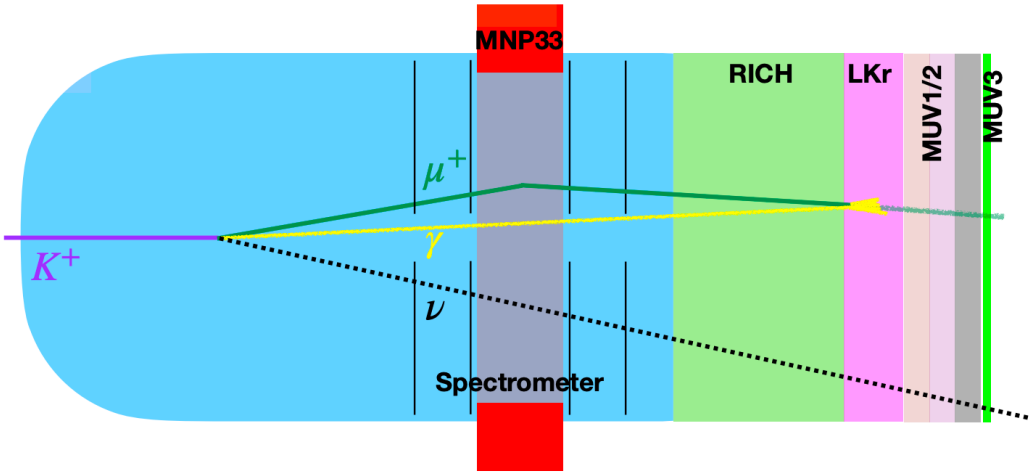

Figure 5.3: Sketch of a $K^+ \to (\mu^+\gamma)\nu$ event (not to scale). [35]

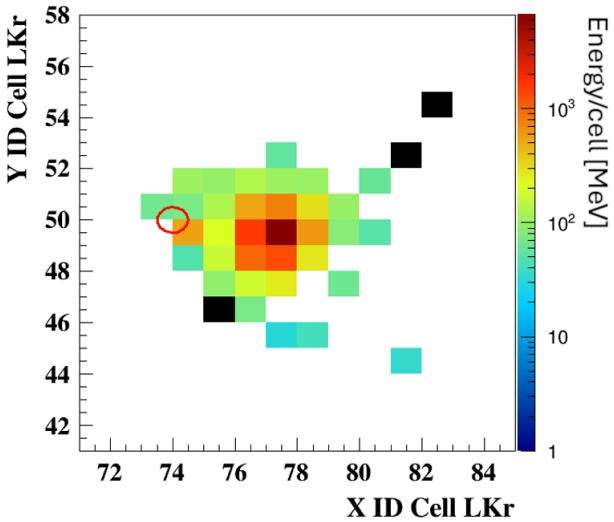

Figure 5.4: Event display of a $K^+ \to (\mu^+\gamma)\nu$ event in the LKr. The red circle shows the STRAW track extrapolation; black denotes dead cells. [35]

Due to the misidentification of the LKr cluster, these events do not enter the evaluation of $f_{\mu^+\nu}$, since, in the $K^+ \to \mu^+\nu$ control sample, the calorimetric muon identification is required. The enhancement of the misidentification

probability is especially significant if the muon has sufficiently high momentum ($\gtrsim 35\,\mathrm{GeV}/c$) that the RICH-based particle identification is degraded (see figure 3.7), and the photon has sufficiently high energy ($\gtrsim 5\,\mathrm{GeV}$) that the energy deposit in LKr is compatible with one produced by a pion.

Selection of $\mathrm{K}^+ \to (\mu^+\gamma)\nu$ decays is found to be significantly more likely in 2021–2022 data than in Run1 data due to a degradation in particle identification performance at high intensities, particularly due to the higher occupancy in MUV1 and MUV2. Indeed, before applying an additional $\mathrm{K}^+ \to (\mu^+\gamma)\nu$ veto, a statistically significant excess of events, with respect to Run1 data, was found in region R2 at $p_{\pi^+} > 35\,\mathrm{GeV}/c$. This background was then identified and studied by using independent data control samples and simulations.

In the assumption that the LKr cluster is due to the overlap of a muon and a photon, the photon 4-momentum can be reconstructed: indeed, the direction of flight is given by the position of the LKr cluster and the K^+–μ^+ vertex, and its energy is the reconstructed energy deposited in the LKr (the nominal MIP energy deposit of 0.6 GeV is subtracted from it). A suitable squared missing mass variable is then calculated as

$$m^2_{\mathrm{miss},\mu\nu\gamma} = (P_\mathrm{K} - P_\mu - P_\gamma)^2\,,$$

where P_K is the 4-momentum of the kaon as measured by the GTK, P_μ is the 4-momentum of the muon as measured by the STRAW in the μ^+ mass hypothesis, and P_γ is the 4-momentum of the photon, calculated as described above.

From MB data, a $\mathrm{K}^+ \to (\mu^+\gamma)\nu$ control sample is selected by applying the full signal selection, but requiring a signal in MUV3 and lifting calorimetric BDT requirements. The distribution in the $(m^2_\mathrm{miss}, m^2_{\mathrm{miss},\mu\nu\gamma})$ plane of this sample is shown in figure 5.5, where the $\mathrm{K}^+ \to (\mu^+\gamma)\nu$ population is clearly visible as events accumulating around $m^2_{\mathrm{miss},\mu\nu\gamma} = 0$, especially in R2. The resolution on $m^2_{\mathrm{miss},\mu\nu\gamma}$ is found to be around $2.3 \times 10^{-3}\,\mathrm{GeV}^2/\mathrm{c}^4$.

A $\mathrm{K}^+ \to (\mu^+\gamma)\nu$ veto condition is therefore defined by rejecting events with $\left|m^2_{\mathrm{miss},\mu\nu\gamma}\right| < 0.01\,\mathrm{GeV}^2/c^4$, reconstructed LKr energy larger than 5 GeV, and $\max\{L_\mathrm{e}, L_\mu, L_\mathrm{K}, L_0\}/L_{\pi^+} > 5\,\%$ (L's are RICH-based likelihoods, see section 3.8.1 for the definition of the symbols). This veto condition is found to suppress this background by a factor 20 while decreasing the signal acceptance only by a relative amount of 0.4 %; it is therefore applied to the signal and normalization selections on top of all the conditions described in chapter 3.

The $\mathrm{K}^+ \to (\mu^+\gamma)\nu$ control sample is used to estimate the number of $\mathrm{K}^+ \to (\mu^+\gamma)\nu$ background events: the number of events in the control sample satisfying $\left|m^2_{\mathrm{miss},\mu\nu\gamma}\right| < 0.01\,\mathrm{GeV}^2/c^4$ is rescaled by the downscaling of the MB trigger, the ratio of the PNN and MB trigger efficiencies, and the probability of misidentifying the LKr cluster as a pion. The latter probability is calculated in a similar $\mathrm{K}^+ \to (\mu^+\gamma)\nu$ control sample, where the RICH

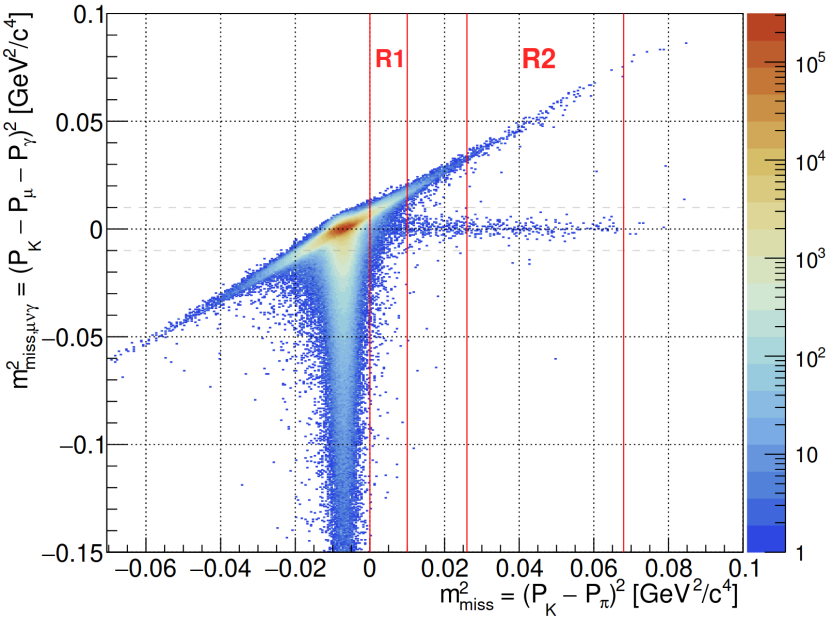

Figure 5.5: Distribution in the $(m^2_{\text{miss}}, m^2_{\text{miss},\mu\nu\gamma})$ plane of the $K^+ \to (\mu^+\gamma)\nu$ control sample. m^2_{miss} intervals containing the signal regions are indicated by the red lines. [35]

muon rejection is relaxed, as the fraction of events satisfying the calorimetric pion identification. As a result, the number of expected $K^+ \to (\mu^+\gamma)\nu$ background events in the signal region is 0.82 ± 0.43.

To validate this prediction, the signal selection without $K^+ \to (\mu^+\gamma)\nu$ veto nor calorimetric BDT probability criteria is applied, obtaining a set of validation samples in the sidebands of the calorimetric BDT pion probability, shown in figure 5.6.

5.1.4 $K^+ \to \pi^+\pi^+\pi^-$

A $K^+ \to \pi^+\pi^+\pi^-$ decay in the fiducial volume requires both m^2_{miss} misreconstruction and STRAW and multiplicity veto inefficiencies in order to enter the signal region. Unlike $K^+ \to \pi^+\pi^0$ and $K^+ \to \mu^+\nu$ decays, this is a three-body decay, meaning that its m^2_{miss} distribution is not peaked. Therefore, in order to apply an approach based on kinematic tail fractions, the 3π region is subdivided into ($0.02\,\text{GeV}^2/c^4$ wide) bins of m^2_{miss}, and the number of $K^+ \to \pi^+\pi^+\pi^-$ background events in a region R is calculated as

$$N_{3\pi}(R) = \sum_{m^2_{\text{miss}} \in 3\pi} N(m^2_{\text{miss}}) f_{3\pi}(R, m^2_{\text{miss}}) \,, \tag{5.2}$$

where the sum is over m^2_{miss} bins in the 3π region, $N(m^2_{\text{miss}})$ is the number of selected events in the m^2_{miss} bin, and $f_{3\pi}(R, m^2_{\text{miss}})$ is the kinematic tail

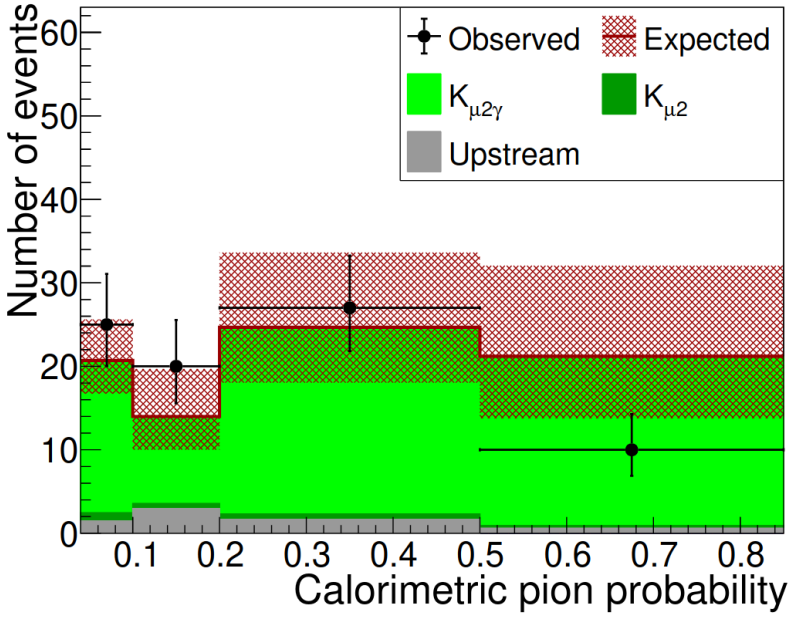

Figure 5.6: Validation of the $K^+ \to (\mu^+\gamma)\nu$ background estimate. In the legend, $K_{\mu 2}$ and $K_{\mu 2\gamma}$ denote $K^+ \to \mu^+\nu$ and $K^+ \to (\mu^+\gamma)\nu$, respectively. [35]

fraction of region R for the m^2_{miss} bin (namely, the expected ratio between the number of events in region R and the number of events in the m^2_{miss} bin).

The evaluation of $f_{3\pi}(R, m^2_{\text{miss}})$ relies on a sample of MC generated $K^+ \to \pi^+\pi^+\pi^-$ events. This sample is used in order to build the joint distribution of the true value of m^2_{miss} (calculated from the MC truth), denoted m^2_{miss}(True), and the reconstructed value of m^2_{miss} (which is the one entering equation (5.2)), also denoted m^2_{miss}(Reco), shown in figure 5.7. From this distribution it is possible to extract the scaling between the number of events where m^2_{miss}(True) lies in region R and the number of events in each m^2_{miss}(Reco) bin. An additional convolution is performed to convert the number of events where $m^2_{\text{miss}}(\text{True}) \in R$ in the number of events where $m^2_{\text{miss}}(\text{Reco}) \in R$. The results for the two signal regions are shown in figure 5.8.

Overall, the number of background $K^+ \to \pi^+\pi^+\pi^-$ events in the signal region is estimated to be 0.11 ± 0.03.

5.1.5 $K^+ \to \pi^+\pi^- e^+\nu$

The $K^+ \to \pi^+\pi^- e^+\nu$ decay is a 4-body decay, whose m^2_{miss} distribution is spread over R2 (see figure 2.1) and no kinematic rejection is possible based on m^2_{miss}. The number of expected background events from this decay can be calculated via a formula which is analogous to the signal estimation (see

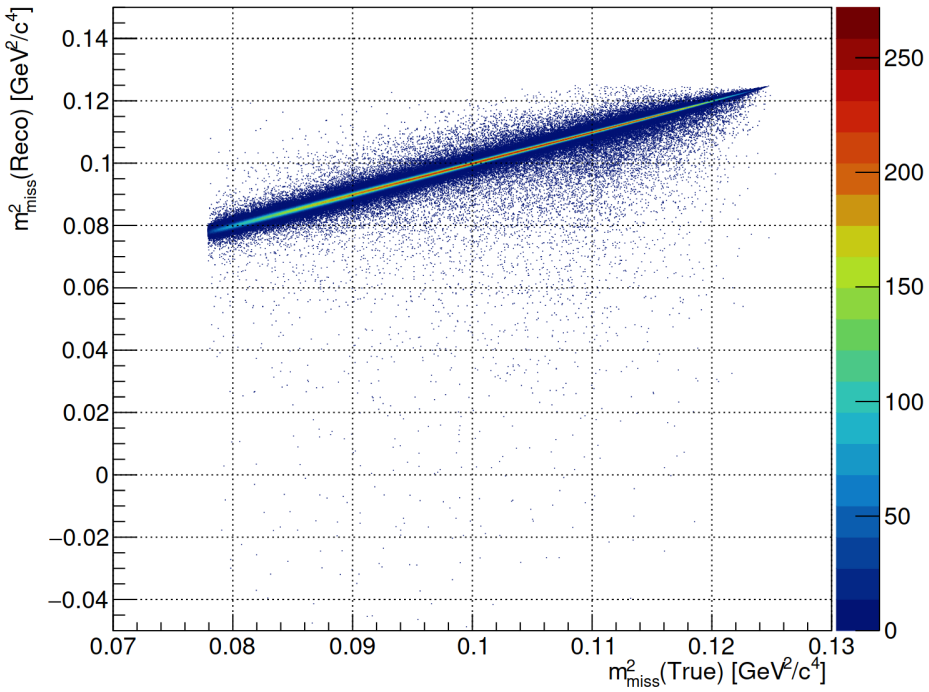

Figure 5.7: Distribution of reconstructed and true m^2_{miss} for simulated $K^+ \to \pi^+\pi^+\pi^-$ events.

equation (4.3)):

$$N_{\pi\pi e\nu} = \frac{N_{\pi\pi} D_{\text{NORM}}}{\text{BR}_{\pi\pi} A_{\pi\pi}} \text{BR}_{\pi\pi e\nu} A_{\pi\pi e\nu} \varepsilon_{\text{trig}} \varepsilon_{\text{RV}} \,,$$

where $\text{BR}_{\pi\pi e\nu} = (4.247 \pm 0.024) \times 10^{-5}$ [27] and $A_{\pi\pi e\nu}$ are the branching ratio and the acceptance of $K^+ \to \pi^+\pi^- e^+\nu$, respectively (other symbols are defined in section 4.1). The acceptance is calculated to be $A_{\pi\pi e\nu} = (1.3 \pm 0.3) \times 10^{-8}$ using a MC sample of 2×10^9 events; figure 5.9 shows the distribution in the kinematic plane of the 76 simulated events that satisfy all selection criteria except the kinematic requirements. The uncertainty on $A_{\pi\pi e\nu}$ is due to the finite MC sample size. The total number of $K^+ \to \pi^+\pi^- e^+\nu$ background events is $0.89^{+0.33}_{-0.27}$.

5.1.6 Other kaon decays

The $K^+ \to \pi^+\gamma\gamma$ decay is relatively rare ($\text{BR}(K^+ \to \pi^+\gamma\gamma) = (9.61 \pm 0.17) \times 10^{-7}$ [90]), but its distribution in the kinematic plane can overlap R2, and the photon veto power is different to that of $K^+ \to \pi^+\pi^0$ decays because of the different kinematics. Out of $O(10^7)$ simulated events, none pass the full signal selection: assuming independence of the photon veto from the rest of the signal selection, a factorization method is applied to calculate the acceptance, and a strategy similar to $K^+ \to \pi^+\pi^- e^+\nu$ is

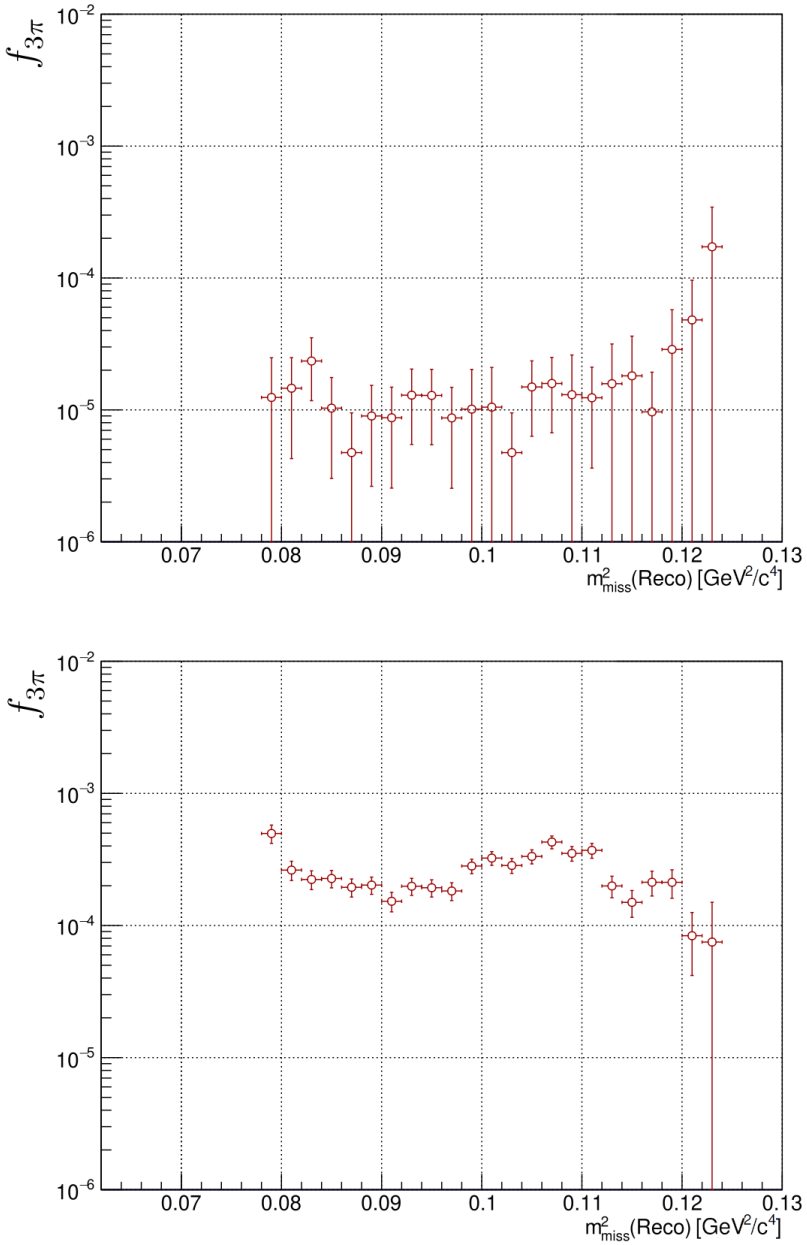

Figure 5.8: Kinematic tail fraction for $K^+ \to \pi^+\pi^+\pi^-$ in signal regions R1 (top) and R2 (bottom), as a function of the m^2_{miss} bin in the 3π region.

applied to evaluate the number of $K^+ \to \pi^+\gamma\gamma$ background events. This is estimated to be 0.01 ± 0.01, where the uncertainty is systematic.

The semileptonic $K^+ \to \pi^0\ell^+\nu$ decays have relatively large BRs ($\text{BR}(K^+ \to \pi^0 e^+\nu) = (5.07 \pm 0.04)$ % and $\text{BR}(K^+ \to \pi^0\mu^+\nu) = (3.352 \pm 0.033)$ % [27]), and they do not have strong constraints in the kinematic plane, but they are extremely suppressed by the combination of the photon veto and the particle identification criteria. Out of $O(10^7)$

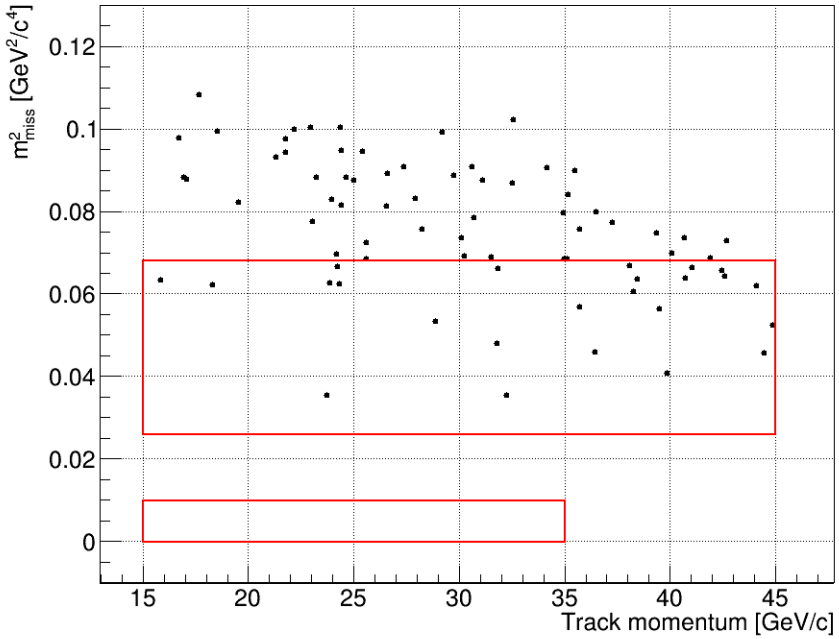

Figure 5.9: Distribution of selected simulated $K^+ \to \pi^+\pi^-e^+\nu$ events in the kinematic plane. Red boxes show the signal regions.

simulated events, none pass the full signal selection: an estimation is obtained by assuming independence among the kinematic criteria, the photon veto, the RICH-based particle identification, and the calorimetric particle identification. The contamination of semileptonic decays in the signal region is calculated to be $< 10^{-3}$ events, and therefore negligible.

5.2 Upstream background

The upstream background is composed of events involving a decay or interaction of a beam particle upstream of the fiducial volume. These events mimic a signal event if a π^+ is produced and reaches the downstream detectors, with no other additional particles associated to the π^+ production being detected, and a candidate K^+ is reconstructed in the KTAG and GTK, and matched (both in time and space) to the downstream π^+.

For an *upstream* event to pass the signal selection, multiple inefficiencies or misreconstructions must occur across the detector:

- the STRAW track, when extrapolated to the GTK3 plane, must not lie within the GTK3 active area (see section 3.7): if the π^+ enters the fiducial volume from the beam pipe, an elastic scattering or a position misreconstruction at the first STRAW chamber (STRAW1) must occur;

- the K^+–π^+ matching criteria (defined in section 3.5.2) must be satisfied, namely the CDA must be less than 4 mm (*good* CDA), the timing between KTAG, GTK and RICH must be consistent, and the conditions on the Bayesian matching parameters ρ and $\Delta\rho$ must be satisfied: if a scattering or misreconstruction happened at STRAW1, at least the CDA (and possibly the timing) must be good only by accident;
- the vertex must be reconstructed inside the fiducial volume: such vertex is fake by definition, in the assumption that the π^+ originated upstream;
- no additional activity must be detected for any veto condition, either upstream (see section 3.7) or downstream (section 3.9): some inefficiency (possibly geometric in nature) must occur, as the π^+ production is expected to be associated to other particles.

5.2.1 Mechanisms

Many different mechanisms can be imagined for the production of the π^+ and the candidate K^+ track reconstruction. Two main categories can be identified: events where the matched GTK track is uncorrelated with the π^+, and events where it is correlated with the π^+. Previous studies indicate that the dominant mechanisms are of the first kind [65, 34].

An example class of events entering the first category is constituted by events where the π^+ is produced in the decay of a beam K^+ upstream of GTK3, or in an inelastic interaction of a beam particle (either a K^+ giving coincident signals in the KTAG, or, less likely, an additional pileup beam particle). The π^+ is then detected downstream and associated to a pileup (possibly fake) GTK track. Here, the relevant kaon decays are $K^+ \to \pi^+\pi^0$ and $K^+ \to \pi^+\pi^+\pi^-$, as the misidentification of the downstream particle is suppressed by many orders of magnitude, while the photon and multiplicity vetoes are less powerful in detecting (neutral or charged) pions emitted upstream of the fiducial volume. The VetoCounter was designed to aid in the rejection of this class of events by detecting the additional particles. A different π^+ production mechanism is the photonuclear interaction of a halo muon with the material upstream of the fiducial volume: these events should be detected and rejected by the ANTI0.

The second category of mechanisms includes events where a beam particle (likely a K^+ giving coincident signals in KTAG) is correctly reconstructed up to GTK3, and then interacts with it. The interaction can produce the π^+ promptly, or a long-lived particle decaying in the fiducial volume. In this case, no other particles must be detected by the CHANTI for the event to enter the signal region.

5.2.2 Estimation strategy

The estimation procedure for the upstream background follows similar principles compared to the one used for the Run1 analysis [65, 34], but it was revised for the analysis of 2021–2022 data, both to make it completely data-driven (and therefore more robust) and to adapt it to the Bayesian K^+–π^+ matching.

All discussed mechanisms require upstream background events to pass the K^+–π^+ matching conditions in a fundamentally different way from physical K^+ decays in the fiducial volume: in particular, the spatial association (and possibly the timing) between the downstream and the upstream track must be accidental. Therefore, the number of upstream background events is expected to scale proportionally to the number of events that do not pass the CDA requirement, where the proportionality constant must account for the fraction of events that pass accidentally the K^+–π^+ matching conditions (including the one on the CDA).

Concretely, the *upstream reference sample* (URS), independent from the signal, is defined by selecting events that pass the full signal selection, but with lifted K^+–π^+ matching requirements, and inverted CDA requirement, namely CDA > 4 mm (*bad* CDA). The URS is assumed to be representative of the joint[2] $(\Delta T_{\text{match}}, N_{\text{GTK}})$ distribution of upstream background events, and to account for the probability of an upstream background event to pass the full selection apart from the K^+–π^+ matching conditions.

It is assumed that the CDA is independent from the ΔT_{match} and N_{GTK} variables as defined in section 3.5.1, namely, that spatial and timing informations are measured independently; then, the number of upstream background events can be calculated as

$$N_{\text{ups}} = \sum_{\Delta T_{\text{match}}, N_{\text{GTK}}} N_{\text{URS}}(\Delta T_{\text{match}}, N_{\text{GTK}}) f_{\text{CDA}} P_{\text{match}}(\Delta T_{\text{match}}, N_{\text{GTK}}) \,, \tag{5.3}$$

where the sum is over 2D $(\Delta T_{\text{match}}, N_{\text{GTK}})$ bins, and:

$N_{\text{URS}}(\Delta T_{\text{match}}, N_{\text{GTK}})$ is the number of events in the URS lying in the relevant bin;

f_{CDA} is the ratio between the number of upstream background events with good CDA and the number of those with bad CDA;

$P_{\text{match}}(\Delta T_{\text{match}}, N_{\text{GTK}})$ is the probability for an upstream background event lying in the relevant bin with good CDA to pass the rest of the K^+–π^+ matching requirements (namely the ones on ρ and $\Delta\rho$, described in section 3.5.2).

Sections 5.2.3–5.2.5 are dedicated to the detailed description of these quantities.

[2] Here and in the following, ΔT_{match} is understood to be, in fact, its absolute value: it is assumed that there is no dependence on its sign.

It must be noted that no kinematic variables enter equation (5.3). In fact the equation can be evaluated in any region of the kinematic plane (meaning that the URS can be defined in any kinematic region), in the assumption that ΔT_{match} and N_{GTK} are independent on the kinematics. Results in the following are given for the whole signal region, but, since the analysis is performed in bins of π^+ momentum, the final measurement will depend from the values obtained in each bin of π^+ momentum rather than from the integral number.

5.2.3 Upstream reference sample

After applying the selection as described above, the URS for the signal region is found to contain a total of 51 events. Its distribution in the $(\Delta T_{\text{match}}, N_{\text{GTK}})$ plane is plotted in figure 5.10. Grid lines in the figure show the 2D bins used for the evaluation of equation (5.3).

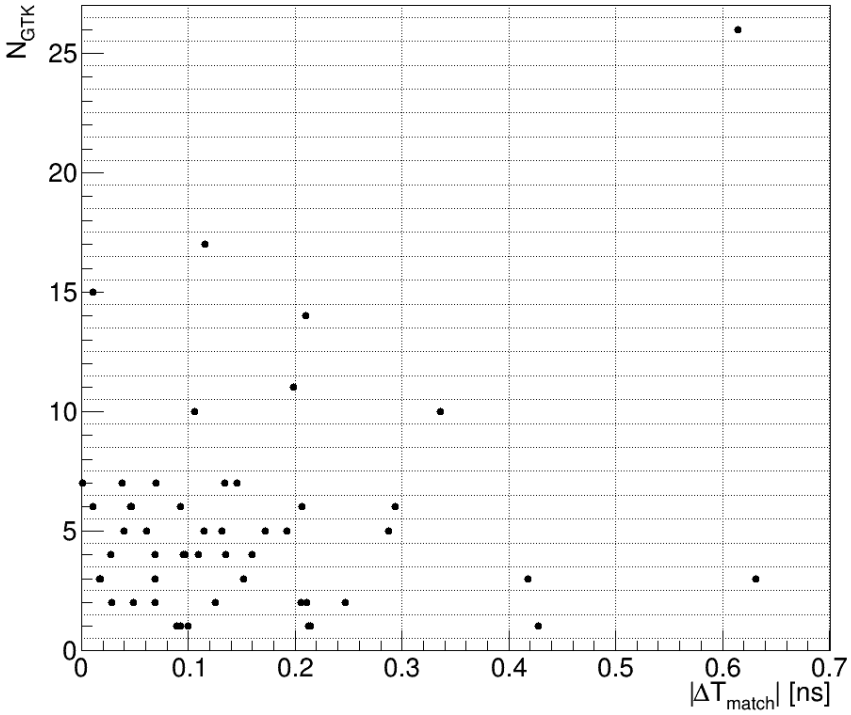

Figure 5.10: $(\Delta T_{\text{match}}, N_{\text{GTK}})$ distribution of the URS.

5.2.4 f_{CDA}

The f_{CDA} factor is defined as $\text{P}(\text{CDA} < 4\,\text{mm})/\,\text{P}(\text{CDA} > 4\,\text{mm})$, where the probabilities are understood to be conditioned to the upstream background

hypothesis. By the assumptions as discussed, it is independent by ΔT_{match} and N_{GTK}.

f_{CDA} is estimated in a data-driven way: specifically, it is extracted from the CDA distribution of the URS sample itself, shown in figure 5.11. Effectively, f_{CDA} must be understood as the ratio between the number of upstream background events in the first (4 mm wide) CDA bin and the total number of events in the URS.

In order to estimate it, the assumption is made that the CDA distribution in upstream background events is reasonably flat up to 12 mm. This is in agreement with the hypothesis of the accidental spatial matching and the PDF for pileup GTK tracks as shown in figure 3.4. f_{CDA} is therefore calculated as

$$f_{\text{CDA}} = \frac{\frac{1}{2} N_{\text{URS}}(4 < \text{CDA/mm} < 12)}{N_{\text{URS}}},$$

namely half of the number of events in the URS that fall in the second and third bins of CDA, divided by the total number of events in the URS. A systematic uncertainty is assigned as

$$(\delta f_{\text{CDA}})_{\text{syst}} = \frac{\frac{1}{2} \left| N_{\text{URS}}(4 < \text{CDA/mm} < 8) - N_{\text{URS}}(8 < \text{CDA/mm} < 12) \right|}{N_{\text{URS}}},$$

to account for the possible deviation of the second and third bins from flatness.

For the whole signal region, $f_{\text{CDA}} = 0.20 \pm 0.03$.

5.2.5 P_{match}

The product $N_{\text{URS}} f_{\text{CDA}}$ is an estimate of the number of upstream background events with good CDA. However, for an event to enter the signal region, it has to also pass the K^+–π^+ matching requirements on ρ and $\Delta\rho$: P_{match} is the probability accounting for this.

P_{match} depends on ΔT_{match} and N_{GTK} as these variables enter the definition of ρ and $\Delta\rho$, while it is assumed not to depend on CDA, since the CDA distribution of the upstream background is assumed flat for CDA < 4 mm, as discussed above. $P_{\text{match}}(\Delta T_{\text{match}}, N_{\text{GTK}})$ is measured in each bin from the normalization sample, as the fraction of events that pass the K^+–π^+ matching criteria, reweighted for the CDA distribution in order to remove the bias given by the fact that the normalization events are genuine kaon decays and have a CDA distribution that falls steeply between 0 and 4 mm (see figure 3.4). Figure 5.12 shows the resulting $P_{\text{match}}(\Delta T_{\text{match}}, N_{\text{GTK}})$. Statistical uncertainties are negligible.

In principle, given an assumed distribution of CDA, $P_{\text{match}}(\Delta T_{\text{match}}, N_{\text{GTK}})$ can be calculated mathematically via toy Monte Carlo simulations. This has been performed as a cross-check of the results in figure 5.12, and no significant deviation was found.

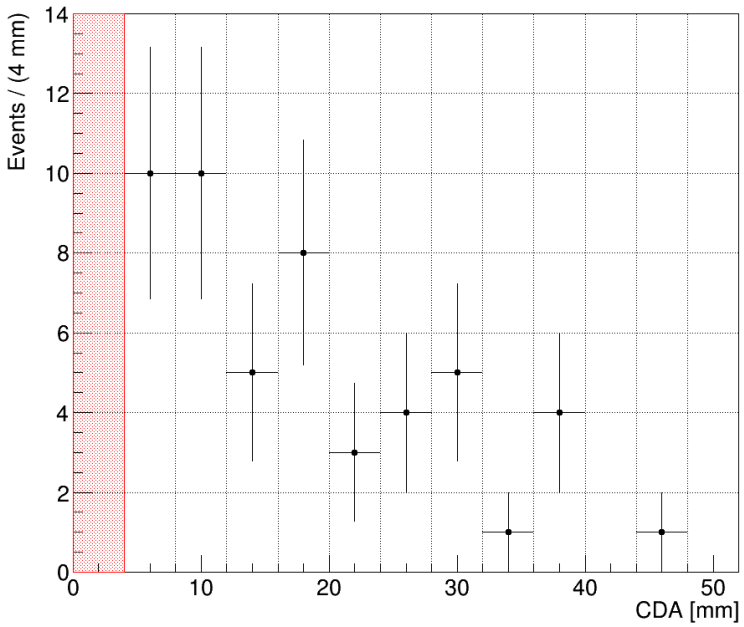

Figure 5.11: CDA distribution of the URS. The red box contains the signal region.

5.2.6 Validation

A set of 10 upstream background validation samples is defined, which are all upstream-background dominated, and independent from each other and from the signal sample. Table 5.2 describes each of the samples: in particular, samples 1 and 7 are enriched of beam interactions; samples 3 and 5 are enriched of upstream decays; even-numbered samples are defined by selecting a kinematic region which does not include well-reconstructed physical K^+ decays in the fiducial volume.

Table 5.3 reports the results of the validation. Expectations and observations are found to be in excellent agreement across all the samples. A further cross-check is performed by splitting each validation sample in π^+ momentum bins. The very good agreement shown in figure 5.13 is further proof that the estimation procedure is robust and is reliable for the analysis as performed in π^+ momentum bins.

5.2.7 Results

The final evaluation of equation (5.3) gives a total of $7.4^{+2.1}_{-1.8}$ expected upstream background events in the signal region. The asymmetric uncertainty accounts for the small statistics in the URS, which is the largest source of uncertainty. The next largest contribution to the uncertainty is from f_{CDA}, which is at the level of 15 to 20 %, depending on the π^+ momentum bin.

Sample	Description
1	Inversion of GTK-based rejection of inelastic interactions, VetoCounter and CHANTI vetoes lifted
3	Inversion of GTK veto on extra hits (VetoCounter and CHANTI vetoes applied)
5	Inversion of VetoCounter veto, GTK veto on extra hits lifted (CHANTI veto applied)
7	Inversion of CHANTI veto, GTK veto on extra hits lifted (VetoCounter veto applied)
9	Inversion of BDT-based upstream background rejection
$n+1$	Same as sample n, but with signal region replaced by $m^2_{\text{miss}} < -0.005\,\text{GeV}^2/c^4$

Table 5.2: Definition of upstream background validation samples.

Sample	Expected	Observed
1	14.2 ± 5.3	13
2	12.3 ± 2.2	12
3	5.5 ± 1.5	6
4	13.4 ± 5.1	11
5	10.4 ± 2.0	10
6	69.4 ± 4.7	76
7	25.6 ± 3.7	27
8	16.5 ± 6.4	13
9	366 ± 18	404
10	5.9 ± 3.7	3

Table 5.3: Results of upstream background validation.

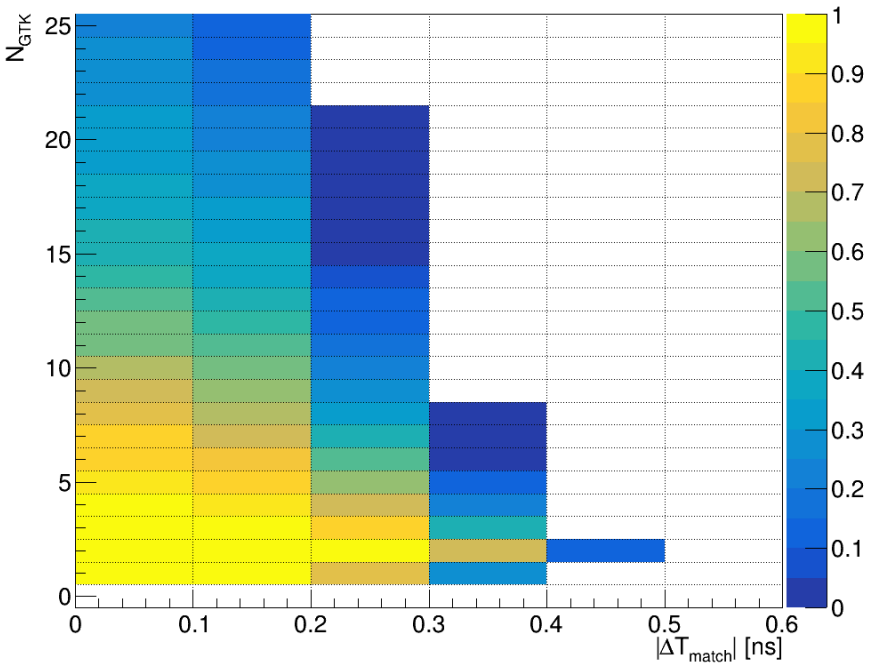

Figure 5.12: Probability of an upstream background event with good CDA to pass the K^+–π^+ matching requirements.

While the method for estimating f_{CDA} may seem crude, its strength lies in the use of a simple model with an uncertainty that is easy to calculate and, therefore, reliable, as proven by the large set of validation samples. In previous analyses, a fit to the CDA distribution was used to estimate f_{CDA}: potential issues such as small statistics or fluctuations of the far tail of the distribution causing instabilities of the fit, or the difficulty of defining a reliable systematic uncertainty, are addressed by the present method.

An interesting observation is that validation sample 5 contains 9 events which were not rejected by the veto on extra GTK signals: these events would have entered the signal region if not vetoed by VetoCounter. Comparing this number to the total upstream background expectation gives a crude *a posteriori* confirmation that the VetoCounter rejects about half of the upstream background events, as was expected by simulations [115].

5.3 Background summary

The summary of all the background expectations in the signal region is reported in tables 5.4 and 5.5.

Compared to the analysis of 2018 S2 data [34], the upstream background is still dominant and constitutes a comparable fraction (about 65 %) of the

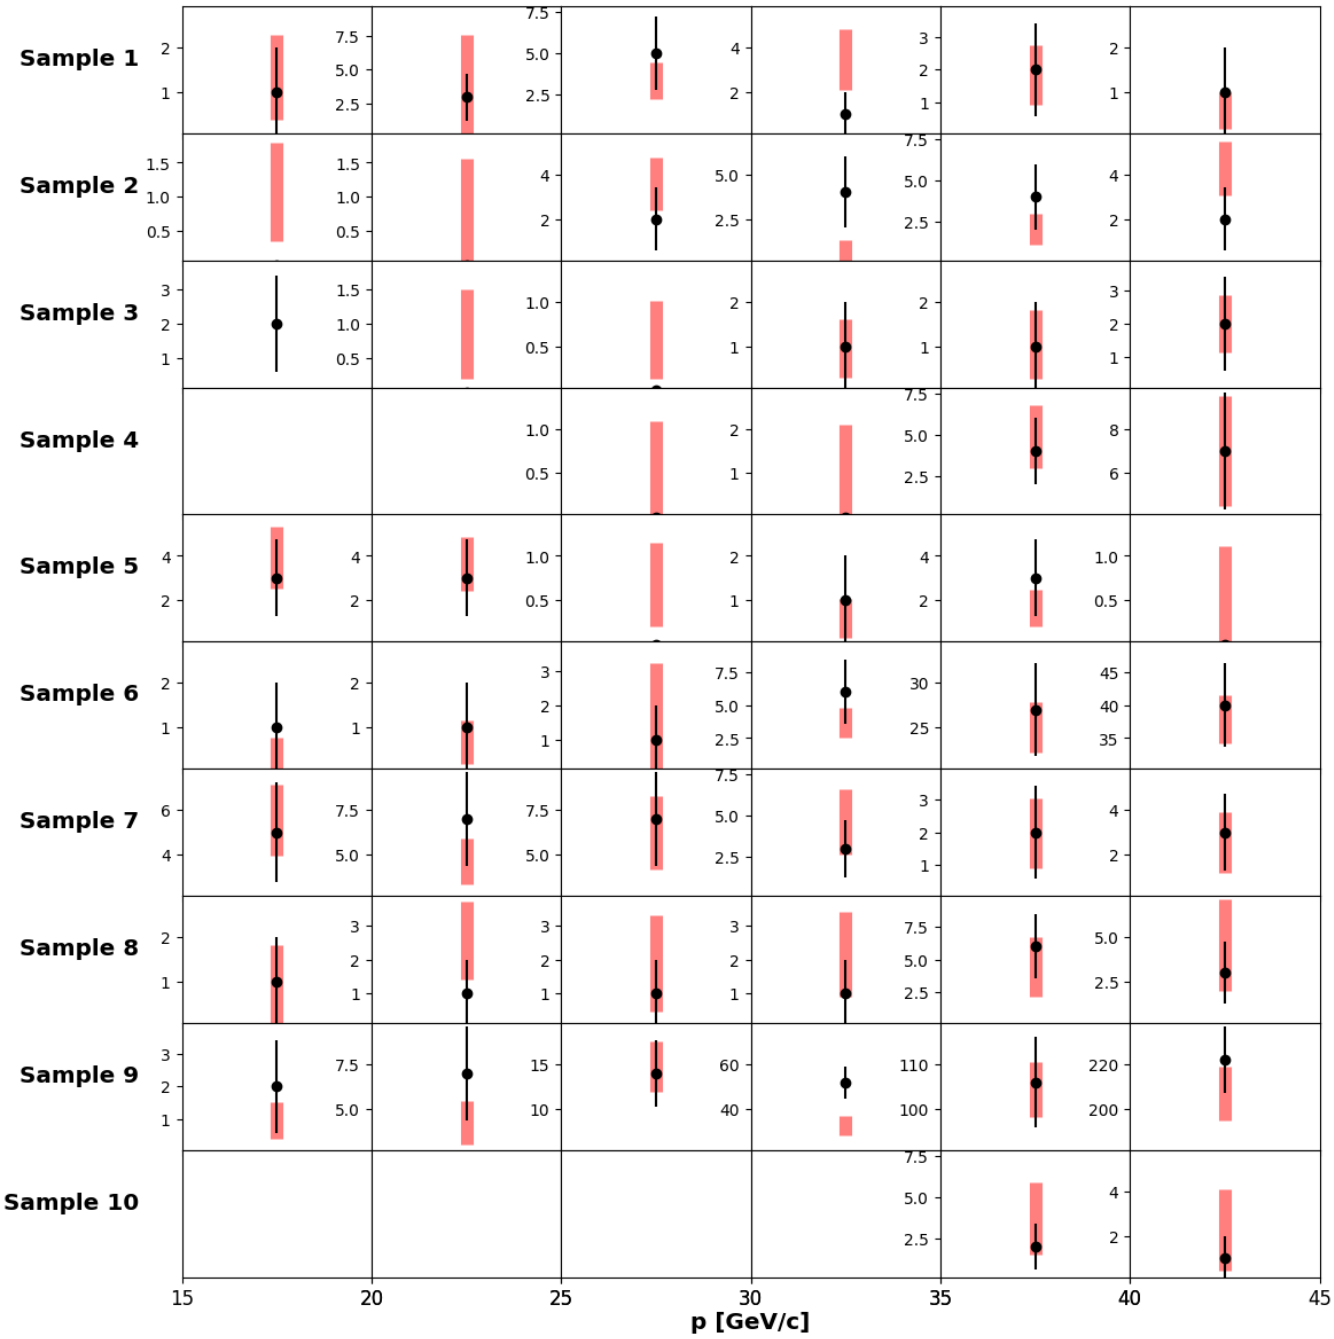

Figure 5.13: Validation of upstream background, in bins of π^+ momentum. Red boxes indicate expectations with their uncertainties, dots indicate observations.

total background. The composition of kaon decay backgrounds has slightly changed, with the $K^+ \to \mu^+\nu(\gamma)$ accounting for a total of 15 % of the background (compared to the 10 % of the last part of 2018 data), and $K^+ \to \pi^+\pi^0(\gamma)$ accounting for 8 % (to be compared to 12 %). Overall, $\sqrt{S+B}/S \approx 0.5$ has been maintained.

$K^+ \to \pi^+\pi^0$	0.76 ± 0.04
$K^+ \to \pi^+\pi^0\gamma$	0.07 ± 0.01
$K^+ \to \mu^+\nu(\gamma)^*$	0.87 ± 0.19
$K^+ \to (\mu^+\gamma)\nu$	0.82 ± 0.43
$K^+ \to \pi^+\pi^+\pi^-$	0.11 ± 0.03
$K^+ \to \pi^+\pi^-e^+\nu$	$0.89^{+0.34}_{-0.28}$
$K^+ \to \pi^0\ell^+\nu$	< 0.001
$K^+ \to \pi^+\gamma\gamma$	0.01 ± 0.01
Upstream	$7.4^{+2.1}_{-1.8}$
Total	$11.0^{+2.1}_{-1.9}$

Table 5.4: Summary of background expectations, integrated over the whole signal region.

*Includes radiative decays but not $K^+ \to (\mu^+\gamma)\nu$.

p_{π^+} bin [GeV/c]	15–20	20–25	25–30	30–35	35–40	40–45
$K^+ \to \pi^+\pi^0(\gamma)$	0.326(35)	0.071(13)	0.043(10)	0.123(19)	0.108(18)	0.163(23)
$K^+ \to \mu^+\nu(\gamma)$	0.056(51)	0.029(21)	0.070(34)	0.56(19)	0.35(11)	0.64(24)
$K^+ \to \pi^+\pi^+\pi^-$	0.041(12)	0.052(15)	0.015(5)	0.003(1)	< 0.001	< 0.001
$K^+ \to \pi^+\pi^-e^+\nu$	$0.11^{+0.13}_{-0.07}$	$0.21^{+0.17}_{-0.12}$	$0.07^{+0.11}_{-0.06}$	$0.16^{+0.15}_{-0.11}$	$0.25^{+0.18}_{-0.13}$	$0.21^{+0.17}_{-0.12}$
Upstream	$0.56^{+0.43}_{-0.31}$	$1.84^{+0.89}_{-0.69}$	$1.09^{+0.67}_{-0.50}$	$1.38^{+0.77}_{-0.58}$	$0.77^{+0.56}_{-0.41}$	$1.59^{+0.82}_{-0.63}$
Total	$1.12^{+0.46}_{-0.34}$	$2.23^{+0.90}_{-0.71}$	$1.32^{+0.68}_{-0.51}$	$2.26^{+0.81}_{-0.64}$	$1.51^{+0.59}_{-0.46}$	$2.63^{+0.86}_{-0.69}$

Table 5.5: Summary of background expectations in each π^+ momentum bin.

Chapter 6

Measurement of $\mathbf{BR(K^+ \to \pi^+ \nu\bar{\nu})}$

Once the estimation of signal and of each background component in the signal region is completed and validated, one last step of validation is to compare the expectation and the observation inside the control regions. Finally, the number of events in the signal regions is counted; by using statistical inference, the measured value of $\mathrm{BR}(\mathrm{K}^+ \to \pi^+ \nu\bar{\nu})$ is then computed.

Results from the control and signal region observations are reported in sections 6.1 and 6.2, respectively. The BR calculation is described in section 6.3, and the preliminary results are reported in section 6.4.

6.1 Control regions

Control regions were defined in table 3.3 and shown in figure 3.9.

Expectations of background and signal in each of them is performed by applying the full analysis, as described in the previous chapters, but replacing the signal region with the control region at hand. The expected composition of the control regions is summarized in table 6.1. Each control region is mostly populated by a specific main kaon decay, which allows to cross-check each background estimation; CR3D is mostly populated by events with a mismeasurement of the momentum of a track, and therefore is populated by most backgrounds (and by the signal). The total expectations are reported in table 6.2.

Events selected in the control regions are shown in the kinematic plane in figure 6.1. The comparison with the observation is reported in figure 6.2 as function of the π^+ momentum, and in table 6.3 as integral in the full momentum range. A good agreement is found, validating the background and the signal expectations.

CR1 CR2	$K^+ \to \pi^+\pi^0(\gamma)$, upstream, $K^+ \to \pi^+\nu\bar{\nu}$
CRmu CRmu2	$K^+ \to \mu^+\nu(\gamma)$, upstream
CRmu3	$K^+ \to \mu^+\nu(\gamma)$, $K^+ \to \pi^+\pi^0(\gamma)$, upstream, $K^+ \to \pi^+\nu\bar{\nu}$
CR3pi	$K^+ \to \pi^+\pi^+\pi^-$, $K^+ \to \pi^+\pi^-e^+\nu$, upstream, $K^+ \to \pi^+\nu\bar{\nu}$
CR3D	$K^+ \to \mu^+\nu(\gamma)$, $K^+ \to \pi^+\pi^0(\gamma)$, $K^+ \to \pi^+\pi^-e^+\nu$, upstream, $K^+ \to \pi^+\nu\bar{\nu}$

Table 6.1: Expected composition of the control regions.

p_{π^+} bin [GeV/c]	15–20	20–25	25–30	30–35	35–40	40–45	Signal
CR1	3.2 ± 0.7	1.3 ± 1.2	0.5 ± 0.3	1.4 ± 0.8	2.8 ± 1.4	4.5 ± 1.4	1.6
CR2	5.4 ± 0.7	1.2 ± 0.3	0.6 ± 0.1	2.4 ± 0.8	4.0 ± 1.3	5.5 ± 1.2	1.4
CRmu	0.9 ± 0.4	7.8 ± 2.8	22 ± 7	46 ± 15	310 ± 100	288 ± 95	< 0.01
CRmu2	0.3 ± 0.3	1.6 ± 0.7	3.4 ± 1.2	7.2 ± 2.2	18.6 ± 5.7	—	0.1
CRmu3	—	—	—	—	45 ± 12	96 ± 27	1.1
CR3pi				0.4 ± 0.2			0.5
CR3D	0.7 ± 0.6	0.7 ± 0.7	0.7 ± 0.6	2.6 ± 1.0	0.5 ± 0.4	1.1 ± 0.9	1.0

Table 6.2: Summary of control region expectations of background (per π^+ momentum bin) and SM signal (integrated, with negligible uncertainty compared to the background). In CR3pi, each single bin is expected to have $\lesssim 0.1$ background events, and therefore only a total is reported.

	Expectation	Observation
CR1	15.8 ± 3.9	20
CR2	20.9 ± 3.0	17
CRmu	675 ± 140	515
CRmu2	31.4 ± 6.6	32
CRmu3	141 ± 31	115
CR3pi	0.9 ± 0.2	1
CR3D	7.5 ± 2.7	9

Table 6.3: Summary of control region results.

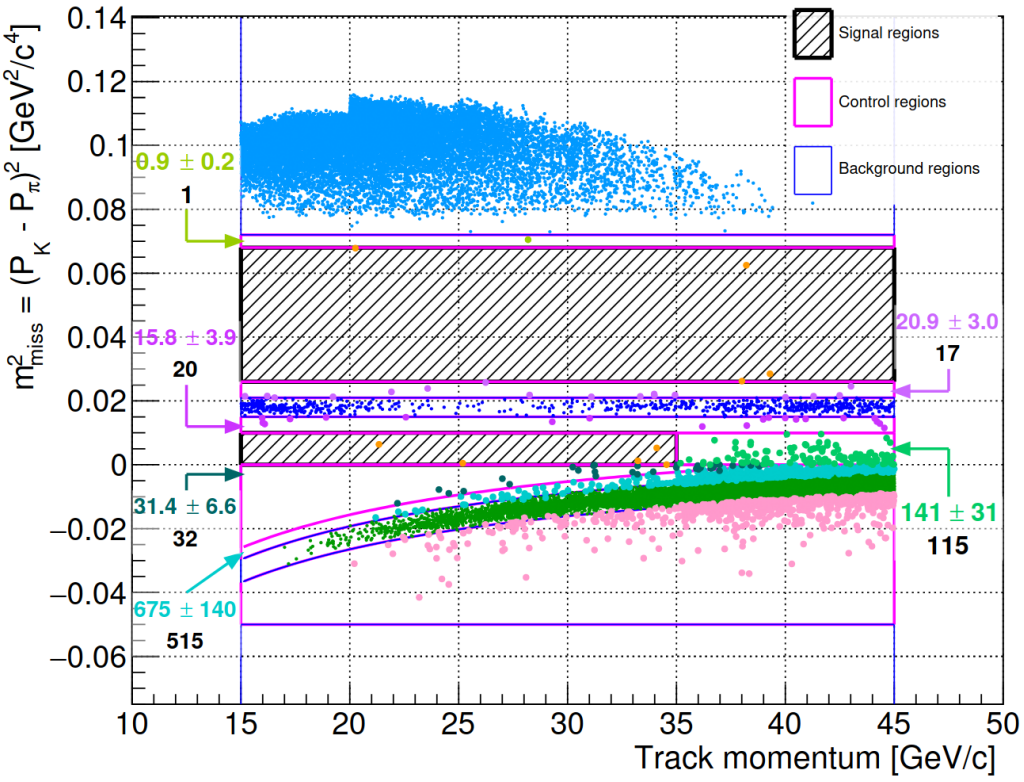

Figure 6.1: Distribution of the events selected in the background and control regions in the kinematic plane. Orange points, overlapping the signal regions, belong to the CR3D region. [35]

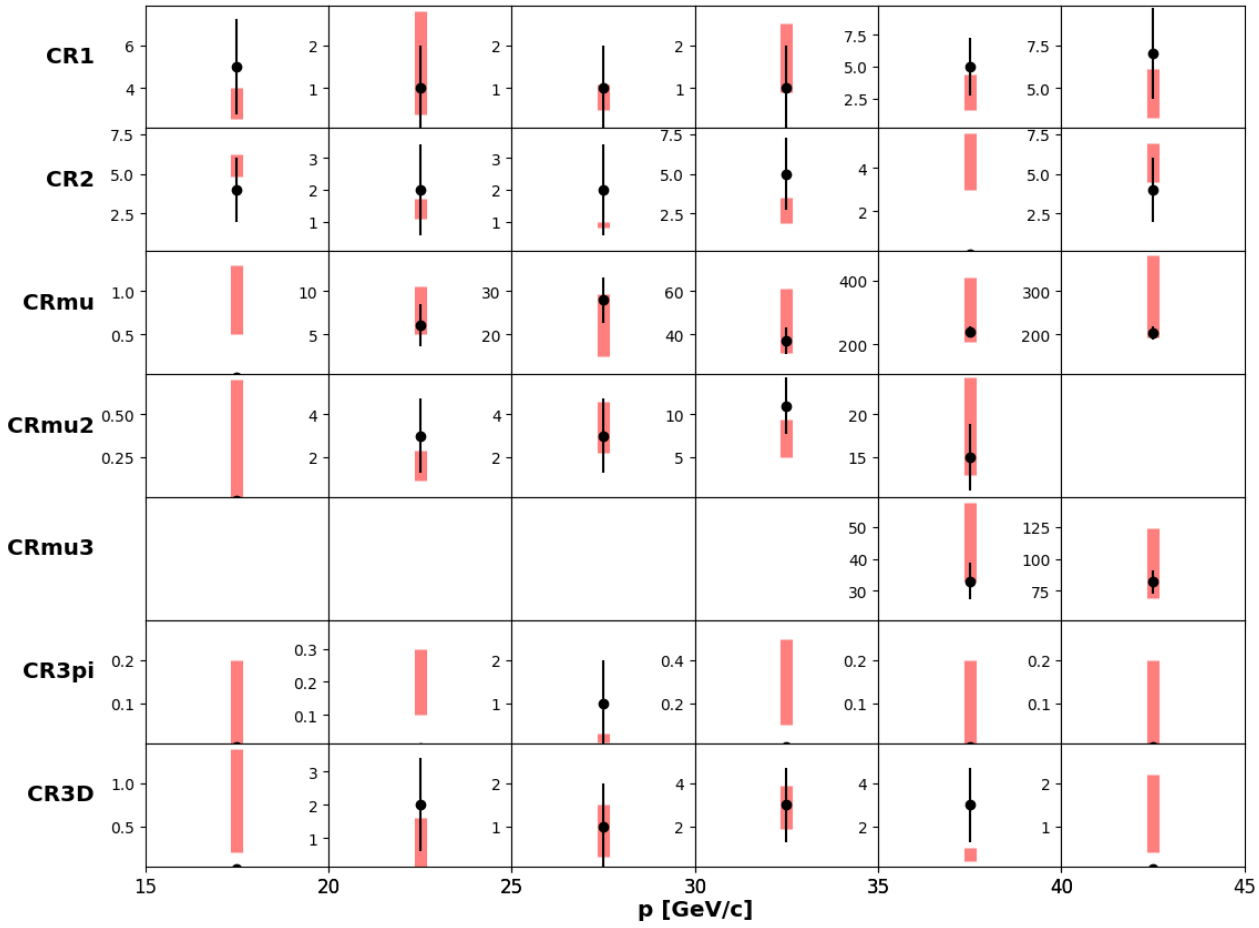

Figure 6.2: Control region results in bins of π^+ momentum.

6.2 Signal region

After the full selection, $N = 31$ events are found in the present dataset. The final distribution in the kinematic plane is shown in figure 6.3.

For comparison with the observation, the total expected SM signal (as calculated in section 4.6) is $N^{\mathrm{SM}}_{\pi\nu\bar{\nu}} = 9.91 \pm 0.34$, while the total estimated background (as calculated in 5.3) is $N_b = 11.0^{+2.1}_{-1.9}$. A detailed comparison in bins of momentum of the expected signal and background (as detailed in tables 4.4 and 5.5 respectively), is reported in table 6.4. A slight excess of 50 % more events than expected is found.

6.3 Statistical interpretation

In order to extract a measurement of the BR and its uncertainty from the observations and expectations of the single momentum bins (*categories*), a statistical treatment is needed.

The notation used in the following is:

n_i: number of observed events in the i-th category;

$s_i \pm \sigma_{s,i}$: number of expected signal events, with its uncertainty, in the i-th category;

$b_i \pm \sigma_{b,i}$: number of estimated background events, with its uncertainty, in the i-th category;

$\theta = \mathrm{BR}_{\pi\nu\bar{\nu}}/\mathrm{BR}^{\mathrm{SM}}_{\pi\nu\bar{\nu}}$: *signal strength*, the parameter to infer statistically.

In each category, the expected number of observed events is given by $\mu_i(\theta) = b_i + \theta s_i$, and its uncertainty is given by[1] $\sigma_i(\theta)^2 = \sigma^2_{b,i} + \theta^2\sigma^2_{s,i}$.

The distribution of the observed events $N_i(\theta)$ is modelled by

$$\begin{aligned} \Lambda_i(\theta) &\sim \mathcal{N}(\mu_i(\theta), \sigma_i(\theta))\,, \\ N_i(\theta)|\Lambda_i(\theta) &\sim \mathrm{Poisson}(\Lambda_i(\theta))\,, \end{aligned} \tag{6.1}$$

namely the distribution of $N_i(\theta)$ is obtained by drawing an expected number of events $\Lambda_i(\theta)$ from a Gaussian[2] with mean $\mu(\theta)$ and standard deviation $\sigma(\theta)$, and then drawing from a Poissonian with mean $\Lambda_i(\theta)$. The likelihood of θ can therefore be expressed as

$$L(\theta; \{n_i\}) = \prod_i \int_0^{+\infty} d\lambda f_{\mathcal{N}}(\lambda; \mu_i(\theta), \sigma_i(\theta))\, \mathrm{P}_{\mathrm{Poisson}}(n_i; \lambda)\,. \tag{6.2}$$

[1] A symmetrization procedure is applied in order to deal with the asymmetric uncertainties.

[2] The Gaussian is, in fact, truncated, in order to make the probability of $\Lambda_i(\theta) < 0$ vanish.

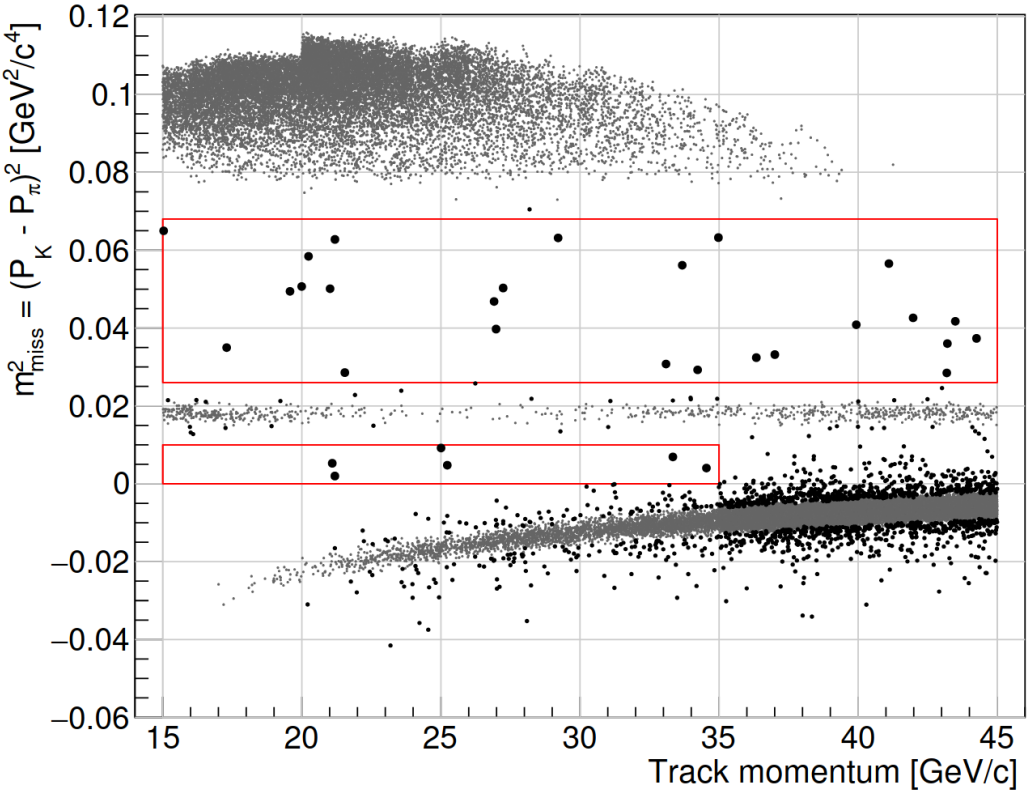

Figure 6.3: Distribution of events after the full signal selection. [35]

p_{π^+} bin [GeV/c]	15–20	20–25	25–30	30–35	35–40	40–45
$N^{\mathrm{SM}}_{\pi\nu\bar{\nu}}$	1.20(4)	2.21(7)	2.41(7)	2.03(6)	1.26(4)	0.80(3)
N_b	$1.12^{+0.46}_{-0.34}$	$2.23^{+0.90}_{-0.71}$	$1.32^{+0.68}_{-0.51}$	$2.26^{+0.81}_{-0.64}$	$1.51^{+0.59}_{-0.46}$	$2.63^{+0.86}_{-0.69}$
N	4	6	6	6	3	6

Table 6.4: Expectations and observations in the signal region.

In practice, the test statistic

$$q(\theta) = -2 \ln L(\theta; \{n_i\})$$

is calculated, for computational efficiency.

A p-value for the observation of the signal can be calculated by generating pseudo-experiments according to the assumed statistical model (6.1) in the background-only hypothesis (namely, with $\theta = 0$), and defining p as the fraction of pseudo-experiments where the resulting likelihood is smaller than the observed likelihood.

The fitted value of the signal strength $\hat{\theta}$ is defined as the one minimizing $q(\theta)$.

By Wilks' theorem [134],

$$\Delta q(\theta) = -2 \ln \frac{L(\theta; \{n_i\})}{L(\hat{\theta}; \{n_i\})}$$

is assumed to follow a χ^2 distribution, so that the 68 % (resp. 95 %, 99.7 %) confidence interval for θ can be calculated by solving the equation $\Delta q(\theta) = 1$ (resp. 4, 9). The procedure to find confidence intervals is validated by generating pseudo-experiments according to the assumed statistical model (6.1) (with $\theta = \hat{\theta}$) and verifying that the confidence intervals obtained from the quantiles of the resulting distribution of the maximum likelihood estimator are consistent. Further validation is performed by checking the confidence interval coverage for pseudo-experiments where θ is set to arbitrary values.

The total uncertainty on the BR is given by the 68 % confidence interval as discussed above. In order to extract a statistical-only uncertainty, the procedure is repeated in the assumption that the signal and the background are known exactly, namely reducing equation (6.2) to the product of Poisson probabilities.

6.4 Preliminary results

Following the procedures described in the previous section, results are calculated for the 2021–2022 NA62 dataset.

The p-value in the background-only hypothesis is found to be 1.4×10^{-4}, leading to a significance of 3.6σ. The test statistic q is shown in figure 6.4 as function of the BR, and the resulting measurement of the branching ratio is

$$\begin{aligned} \mathrm{BR}(\mathrm{K}^+ \to \pi^+ \nu\bar{\nu})_{2021\text{–}2022} &= (16.2^{+5.1}_{-4.5}) \times 10^{-11} \\ &= \left(16.2^{+4.9}_{-4.3}|_{\mathrm{stat}}\ {}^{+1.4}_{-1.4}|_{\mathrm{syst}}\right) \times 10^{-11} . \end{aligned} \tag{6.3}$$

The post-fit comparison between expectation and observation is shown in figure 6.5.

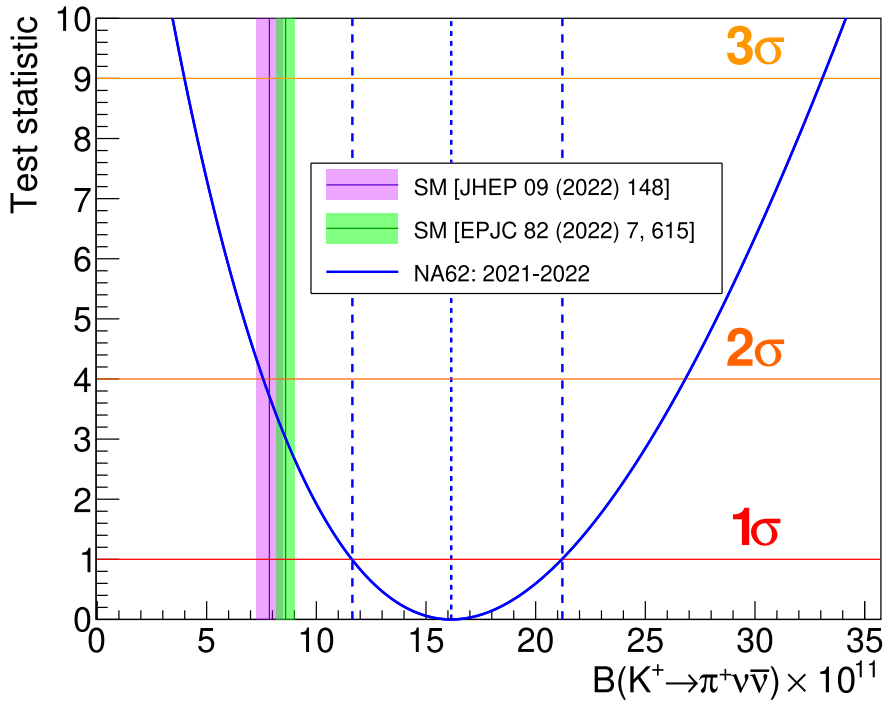

Figure 6.4: Test statistic Δq as function of the BR, for 2021–2022 data.

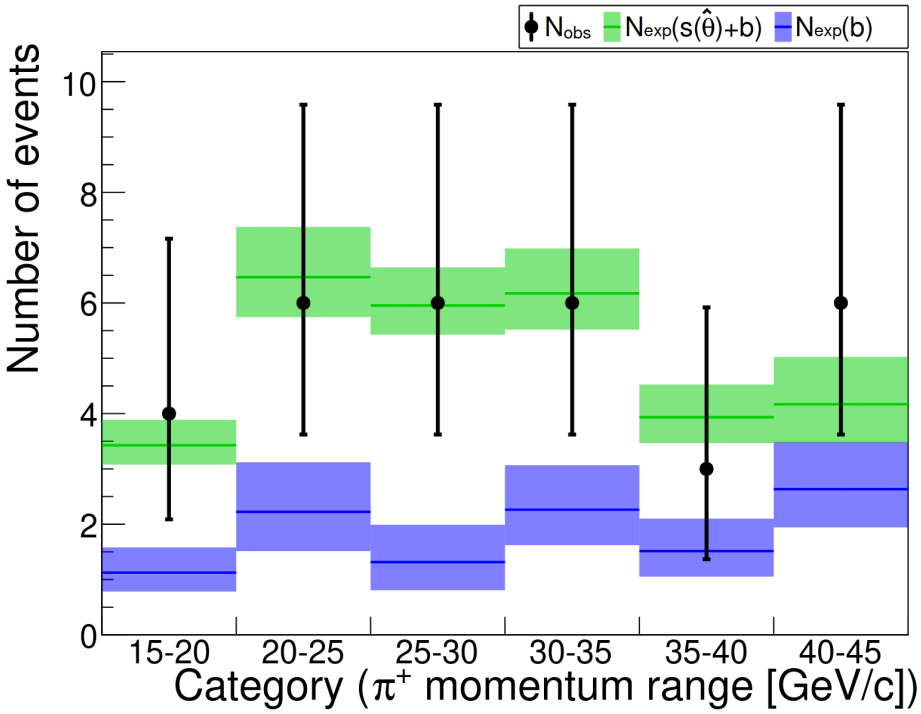

Figure 6.5: Comparison between post-fit expectation and observation in 2021–2022 data. [35]

The 2021–2022 categories can be combined with the existing analysis categories from the results of the 2016 [64], 2017 [65] and 2018 [34] data analyses. Table 6.5 shows the expected and observed events for all the categories. In the full NA62 2016–2022 dataset, a total of 51 events have been observed, with a background expectation of 18^{+3}_{-2}.

Category	Sample	p_{π^+} range (GeV/c)	$N^{\mathrm{SM}}_{\pi\nu\bar{\nu}}$	N_b	N
1	2016	[15, 35]	0.267 ± 0.020	$0.152^{+0.093}_{-0.035}$	1
2	2017	[15, 35]	2.16 ± 0.13	1.46 ± 0.33	2
3	2018 S1	[15, 45]	1.56 ± 0.10	$1.11^{+0.40}_{-0.22}$	2
4	2018 S2	[15, 20]	0.56 ± 0.04	$1.14^{+0.78}_{-0.30}$	1
5		[20, 25]	1.43 ± 0.09	$1.02^{+0.67}_{-0.28}$	4
6		[25, 30]	1.53 ± 0.10	$0.41^{+0.32}_{-0.10}$	2
7		[30, 35]	1.32 ± 0.09	$1.09^{+0.52}_{-0.30}$	6
8		[35, 40]	0.69 ± 0.04	$0.29^{+0.31}_{-0.10}$	1
9		[40, 45]	0.48 ± 0.03	$0.35^{+0.41}_{-0.12}$	1
10	2021–2022	[15, 20]	1.20 ± 0.04	$1.12^{+0.46}_{-0.34}$	4
11		[20, 25]	2.21 ± 0.07	$2.23^{+0.90}_{-0.71}$	6
12		[25, 30]	2.41 ± 0.07	$1.32^{+0.68}_{-0.51}$	6
13		[30, 35]	2.03 ± 0.06	$2.26^{+0.81}_{-0.64}$	6
14		[35, 40]	1.26 ± 0.04	$1.51^{+0.59}_{-0.46}$	3
15		[40, 45]	0.80 ± 0.03	$2.63^{+0.86}_{-0.69}$	6

Table 6.5: NA62 analysis categories.

Using all the NA62 categories as input for the statistical inference, a p-value for the background-only hypothesis of 2×10^{-7} is obtained. The background-only hypothesis is therefore rejected with a significance above 5σ, and therefore this marks the first *observation* of the $K^+ \to \pi^+\nu\bar{\nu}$ decay.

The test statistic q for the full NA62 dataset is shown in figure 6.6, and the branching ratio result is

$$\begin{aligned} \mathrm{BR}(K^+ \to \pi^+\nu\bar{\nu})_{2016\text{–}2022} &= (13.0^{+3.3}_{-3.0}) \times 10^{-11} \\ &= \left(13.0^{+3.0}_{-2.7}|_{\mathrm{stat}}\ ^{+1.3}_{-1.3}|_{\mathrm{syst}}\right) \times 10^{-11}\,. \end{aligned} \tag{6.4}$$

This includes an additional systematic uncertainty of 0.5×10^{-11}, accounting for the combination of the different datasets.

The results of the $\mathrm{BR}(K^+ \to \pi^+\nu\bar{\nu})$ measurements are summarized in figure 6.8, showing the comparison among the results from BNL [63], NA62

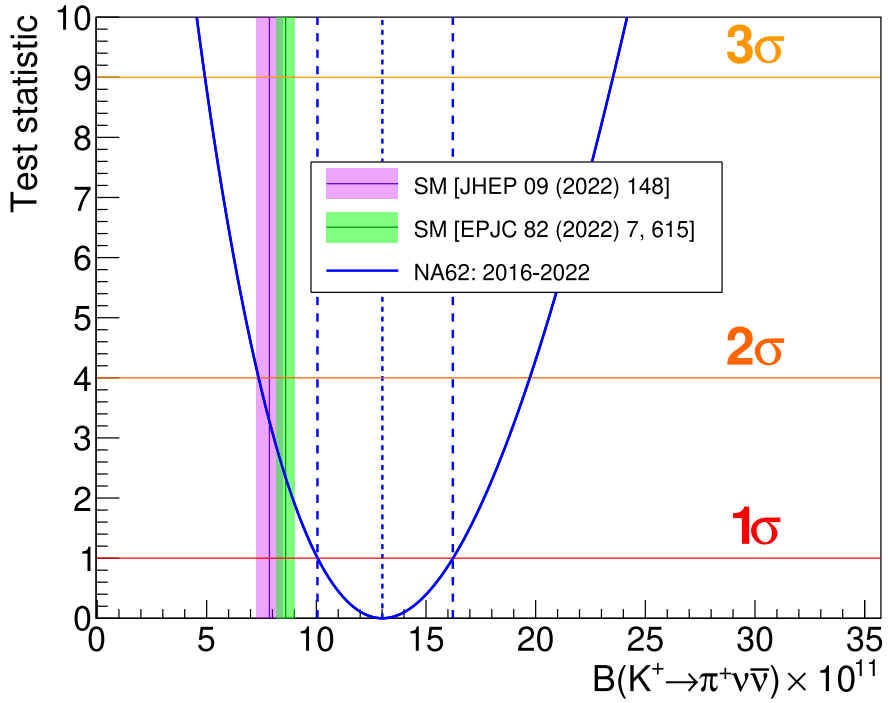

Figure 6.6: Test statistic Δq as function of the BR, for 2016–2022 data.

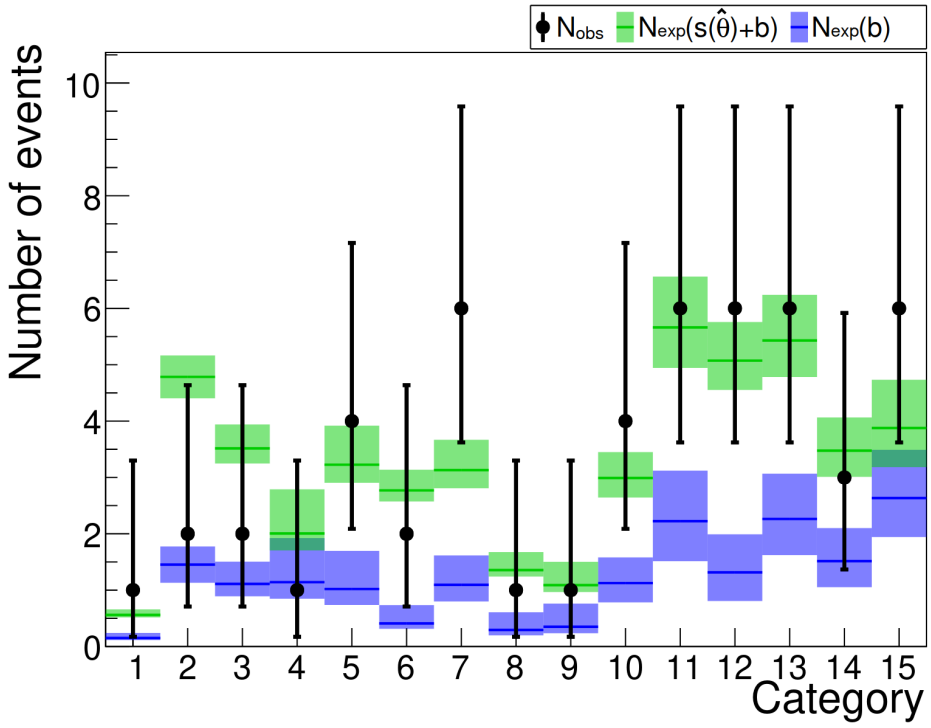

Figure 6.7: Comparison between post-fit expectation and observation in 2016–2022 data. [35]

Run1 [34], NA62 Run2 (6.3), NA62 combined (6.4), and theoretical predictions (1.2a), (1.3a) [31, 32]. The updated landscape of $K \to \pi\nu\bar{\nu}$ decays is shown in figure 6.9.

Very recently, it has been shown that a minimally broken $U(2)_q$ flavour symmetry would generate a small excess $K^+ \to \pi^+\nu\bar{\nu}$ which is compatible with this result and with other experimental results in $B \to D^{(*)}\tau\nu$ and $B \to K^{(*)}\nu\bar{\nu}$ decays [33].

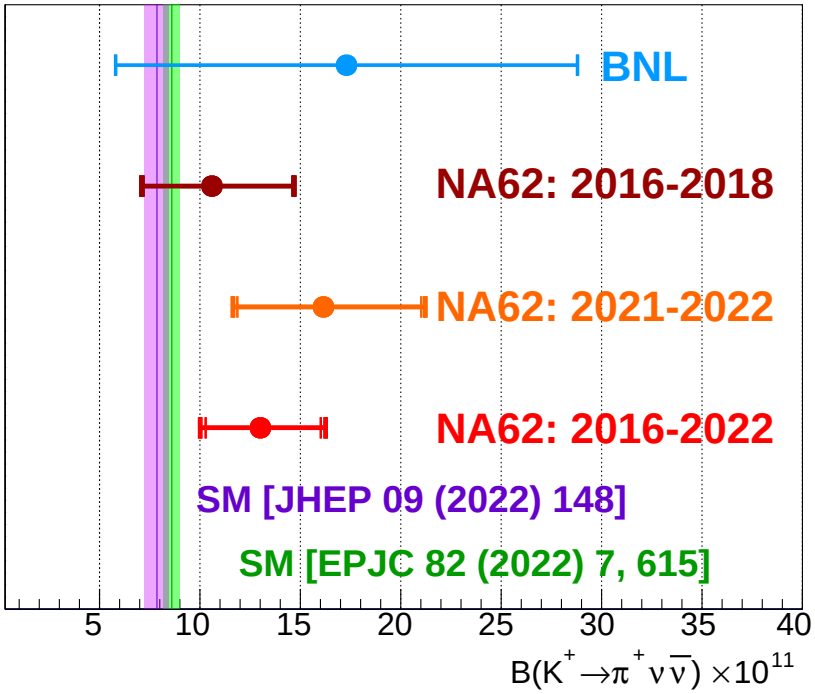

Figure 6.8: Summary of measurements of $BR(K^+ \to \pi^+\nu\bar{\nu})$. Thinner and thicker vertical bars represent statistical and total uncertainties, respectively.

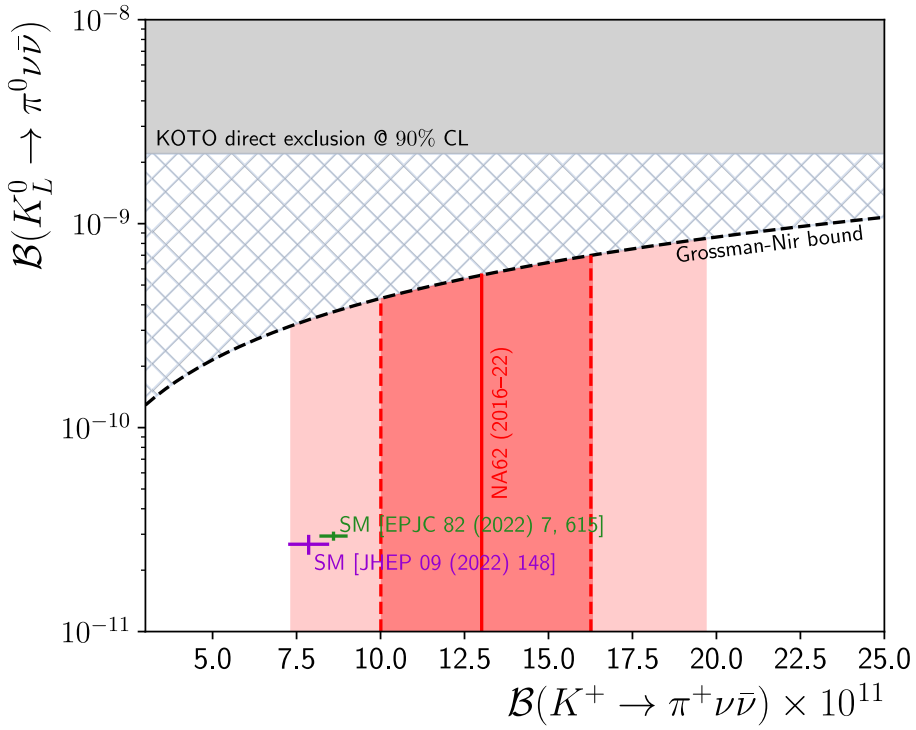

Figure 6.9: Updated experimental status of K $\to \pi\nu\bar{\nu}$. Darker and lighter bands show the 68 % and 95 % confidence level intervals, respectively. [35]

Conclusion

After having obtained the world's most precise measurement ever of $\mathrm{BR}(\mathrm{K}^+ \to \pi^+\nu\bar{\nu})$ with data taken during its first physics run, the NA62 experimental apparatus has been upgraded during CERN Long Shutdown 2 in order to be more resilient against the dominant backgrounds for this analysis. NA62 has then collected data in 2021 and 2022 at its highest instantaneous beam intensity ever.

With the inclusion of this dataset, owing to several improvements in the analysis, the signal yield was doubled with respect to the previous NA62 result, while keeping the background under control: the new detectors proved to be crucial in this regard. The analysis of this dataset also allowed to determine the optimal beam intensity for the operation of the NA62 experiment.

The result of this analysis, and the combination with previous data, is the observation, with a significance above 5σ, of the $\mathrm{K}^+ \to \pi^+\nu\bar{\nu}$ decay: to date, it is the decay with the smallest branching ratio ever observed at this level of significance. Compared to the previous NA62 result, the precision in the BR measurement improved from 40 % to 25 %. NA62 measurements are compatible with each other, and with the BNL measurement. The new measured BR is compatible with the SM prediction within 1.7σ, while being approximately 50 % larger.

NA62 is expected to collect data until CERN Long Shutdown 3, and aims to reach a precision in the measurement of $\mathrm{BR}(\mathrm{K}^+ \to \pi^+\nu\bar{\nu})$ better than 20 %. With HIKE [135, 136] not having been approved, the final NA62 result will likely be the last result on ultra-rare K^+ decays for several years: it will be crucial to exploit the full NA62 dataset in order to clarify agreement or tension with the SM. The ongoing investigation of the $\mathrm{K_L} \to \pi^0\nu\bar{\nu}$ decay by the KOTO experiment and the proposed KOTO-II project [137] will be also essential to probe the correlation between charged and neutral decay modes.

Appendix A

Study of K^+–π^+ matching

This appendix is dedicated to give further details on how the parameters and the selection criteria of the K^+–π^+ matching, described in section 3.5, were defined.

A.1 The $K^+ \to \pi^+\pi^+\pi^-$ sample

A control sample of $K^+ \to \pi^+\pi^+\pi^-$ events can be selected from MB data. Events are selected where exactly 3 tracks satisfy the same conditions as the π^+ from the signal (namely, the ones described in section 3.4). The three tracks must have a total charge of +1, not be associated (both in space and time) with a MUV3 signal, and they must be within 3 ns from each other. A KTAG signal must also be present within 3 ns of each track, and the reconstructed vertex of the three tracks must be inside the fiducial volume $105\,\text{m} < Z < 170\,\text{m}$. LAV and SAV vetoes are applied as in the signal selection (see section 3.9.1). A final requirement on the reconstructed invariant mass of the three pions is also applied. The total number of selected $K^+ \to \pi^+\pi^+\pi^-$ events is denoted $N_{3\pi}$ in the following.

This sample provides an excellent way to test the association between STRAW and GTK tracks, and, in general, to test the GTK reconstruction. Indeed, in each event the total 3-momentum of the three pions is measured by the STRAW and gives a measurement of the kaon momentum independently of the GTK. Specifically, for each GTK track, a χ^2 discriminant compares its momentum as measured by the GTK with the total 3π momentum, and the track with smallest discriminant is tagged as the *true* K^+ track.

A.2 Parameters

The probability of the true K^+ track being a good track (as defined in section 3.5.1), as a function of the number of good tracks N_{GTK}, can thus be evaluated, and the result is shown in figure 3.3.

Then, in each event a random π^+ is chosen, and the distribution of the CDA and of ΔT_{match} (see equation (3.2)) of the true K^+ track can be evaluated and fitted in order to extract the PDFs for these variables in the K^+ hypothesis (shown in figure 3.4, red). A sample of tracks that are uncorrelated with the downstream activity is obtained by artificially shifting all GTK tracks in time, namely by considering out-of-time GTK tracks with readjusted time. This sample of pileup tracks is used to evaluate the distribution of the CDA and of ΔT_{match} of the pileup tracks, which then are fitted to obtain the PDFs in the pileup hypothesis (shown in figure 3.4, blue).

A.3 Definition of performance metrics

Once the parameters have been calculated, the matching quality parameter ρ (see equation (3.3)) can be calculated for each track. The efficiency of K^+ identification ε_K is measured as the fraction of events where the GTK match (namely, the track with the smallest ρ) is the true K^+ track. The wrong association probability η_w is measured as the fraction of events where the GTK match is not the true K^+ track. It must be noted that, even without any selection criteria on the matching quality, these two do not add up to 1: this is because the GTK match is constrained to be a good track, and therefore $1 - \varepsilon_K - \eta_w$ is the fraction of events where no good GTK track is reconstructed. On the other hand, the accidental association probability η_a is calculated on the artificial sample obtained from out-of-time GTK tracks: the fraction of such events where any GTK match is found gives a measurement of η_a. This sample has the same size as the one used to evaluate ε_K and η_w.

Both ε_K and η_w measurements suffer from possible contaminations. ε_K is contaminated by the cases where the GTK match and the true K^+ are the same track only by accident: for instance, this can happen if the actual kaon is not reconstructed, but a single, uncorrelated, good GTK track is. η_w is contaminated by cases where the true K^+ track is mistagged by the three pions, and the GTK match is, in fact, a better description of the parent kaon: this happens whenever an accidental track, not matching in time and space with the three pions, has a 3-momentum which is close enough to the total 3π momentum.

It is possible to remove completely the contamination of η_w by using an artificial sample: rather than considering all the in-time tracks for the GTK matching, an artificial set, consisting of out-of-time tracks plus the true K^+ track as tagged by the three pions, is used. In this manner, one makes sure that, if the GTK match and the true K^+ track do not coincide, it is a genuine error of the GTK matching algorithm.

The same method is, in fact, also used for the evaluation of ε_K. This has the effect that the description in terms of $\Delta\rho$ (see equation (3.4)) is cleaner

for events in which two very similar tracks are reconstructed[1]: by using the method as discussed, only the track that was tagged by the three pions is kept, while the other is removed (and possibly replaced with a track taken from the out-of-time sample).

A.4 Choice of the working point

ε_{K}, η_w and η_a depend strongly on the criteria that are applied to ρ_1 in $N_{\mathrm{GTK}} = 1$ events and to $(\rho_1, \Delta\rho)$ in $N_{\mathrm{GTK}} > 1$ events. Figure A.1 shows the distribution of these parameters for the samples used for the measurements of ε_{K} (figure A.1a: events where the GTK match and the true K^+ track coincide), η_w (figure A.1b: events where the GTK match and the true K^+ track do not coincide), and η_a (figure A.1c: events where a GTK match was found among out-of-time tracks).[2]

The events with efficient K^+ identification peak at smaller ρ_1 (namely, smaller probability of having matched a pileup track) compared to events with errors; at the same time, they have higher $\Delta\rho$ (namely larger difference between best and second best match): these characteristics are expected of an event where there is a track which clearly matches better than the others. For $N_{\mathrm{GTK}} = 1$, events with wrong and accidental associations behave similarly: this is a confirmation that, if a single good track is reconstructed which is not the true K^+, it is uncorrelated to the downstream track, and there is no way for the GTK matching algorithm to distinguish between wrong and accidental errors. For $N_{\mathrm{GTK}} > 1$, however, wrong associations tend to cluster towards small ρ_1 and small $\Delta\rho$, meaning that it is easy for the algorithm to get confused between two similarly well-matching tracks; accidental associations, on the other hand, cluster at large ρ_1 and are more spread out in $\Delta\rho$.

In order to choose the *working point*, namely the criteria on the ρ_1 and $\Delta\rho$ variables to accept an event, a preliminary observation must be made. In fact, this task is not as simple as a usual ROC curve, where a threshold on a single variable controls the efficiency and the error probability: in this case, there are three metrics to take into account. Furthermore, there is an extra degree of freedom due to the different criteria for the two different classes of events ($N_{\mathrm{GTK}} = 1$ and $N_{\mathrm{GTK}} > 1$).

In practice, given the selections S_1 and S_2 for the two classes of events,

[1] This can happen if GTK hit clustering fails, resulting in two tracks effectively sharing the same hits.

[2] The threshold effect in ρ_1 for $N_{\mathrm{GTK}} > 1$ is due to the combination of the use of the N_{GTK}-dependent priors and the use of the ρ_0 translation factor.

the final values for the performance parameters are given by

$$\begin{aligned}
\varepsilon_{\mathrm{K}}(S_1, S_2) =& \frac{1}{N_{3\pi}} \big[N(\text{efficient } \mathrm{K}^+ \text{ identification} \cap N_{\mathrm{GTK}} = 1 \cap S_1) \\
&+ N(\text{efficient } \mathrm{K}^+ \text{ identification} \cap N_{\mathrm{GTK}} > 1 \cap S_2) \big] , \\
\eta_w(S_1, S_2) =& \frac{1}{N_{3\pi}} \big[N(\text{wrong association} \cap N_{\mathrm{GTK}} = 1 \cap S_1) \\
&+ N(\text{wrong association} \cap N_{\mathrm{GTK}} > 1 \cap S_2) \big] , \\
\eta_a(S_1, S_2) =& \frac{1}{N_{3\pi}} \big[N(\text{accidental association} \cap N_{\mathrm{GTK}} = 1 \cap S_1) \\
&+ N(\text{accidental association} \cap N_{\mathrm{GTK}} > 1 \cap S_2) \big] .
\end{aligned}$$

Geometrically, this means that, in the $(\varepsilon_{\mathrm{K}}, \eta_w, \eta_a)$ space, scanning over S_1 and S_2 defines a solid region, where the range of, say, ε_{K} values identified by fixed values of η_w and η_a corresponds to different relative proportions of the selected $N_{\mathrm{GTK}} = 1$ and $N_{\mathrm{GTK}} > 1$ classes of events. This degree of freedom is removed by considering only working points which, for every fixed η_w and η_a, have the maximum possible ε_{K}.

At this point, the choice of the working point is a matter of parametrizing the selection S_2 in the $(\rho_1, \Delta\rho)$ plane, and scanning over the parameters (both of the ones of the S_2 selection and the threshold for the S_1 selection). The parametrization has been chosen, after looking at figure A.1, such that the selection boundary is given by two lines with positive slope. The positive slope ensures that events with increasingly higher ρ_1 (namely, likelihood of having matched a pileup track) are selected only for increasingly higher $\Delta\rho$ (namely, relative likelihood of the best match of being the K^+ with respect to the same likelihood for the second best match). The final working point, whose performance as function of intensity is shown in figure 3.6, provides an improvement on the efficiency of 10 % absolute with respect to the algorithm used in Run1 data analysis described in [65], while keeping the same level for the wrong and accidental association errors.

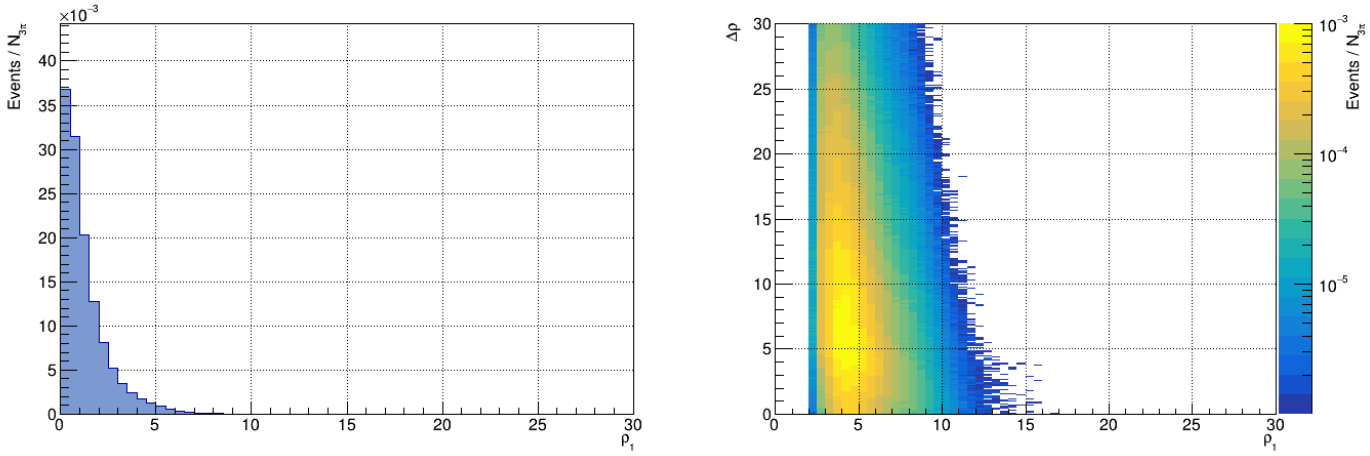

(a) Events with efficient K^+ identification.

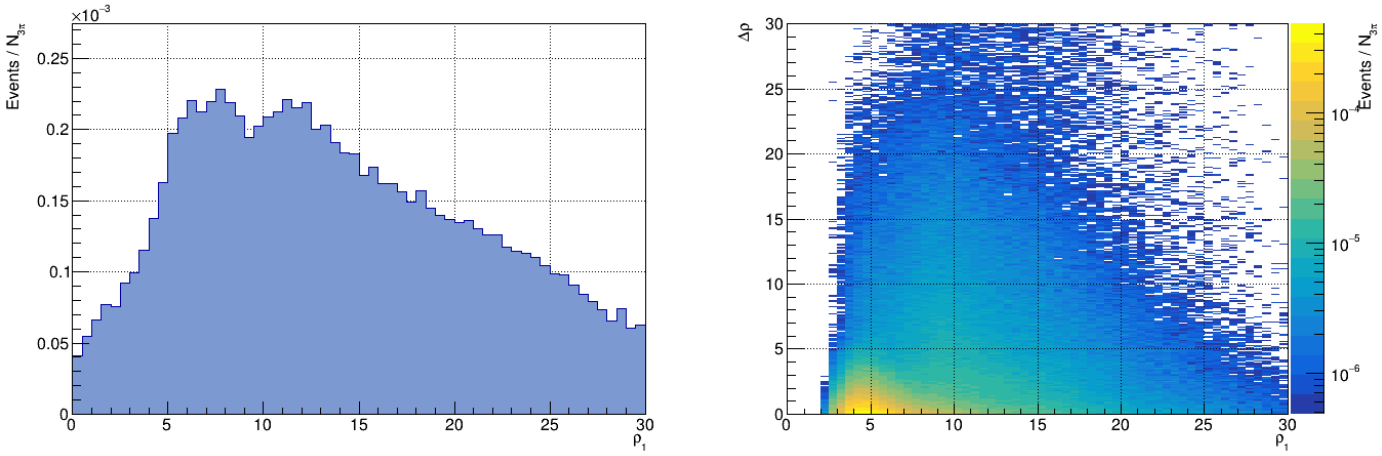

(b) Events with wrong associations.

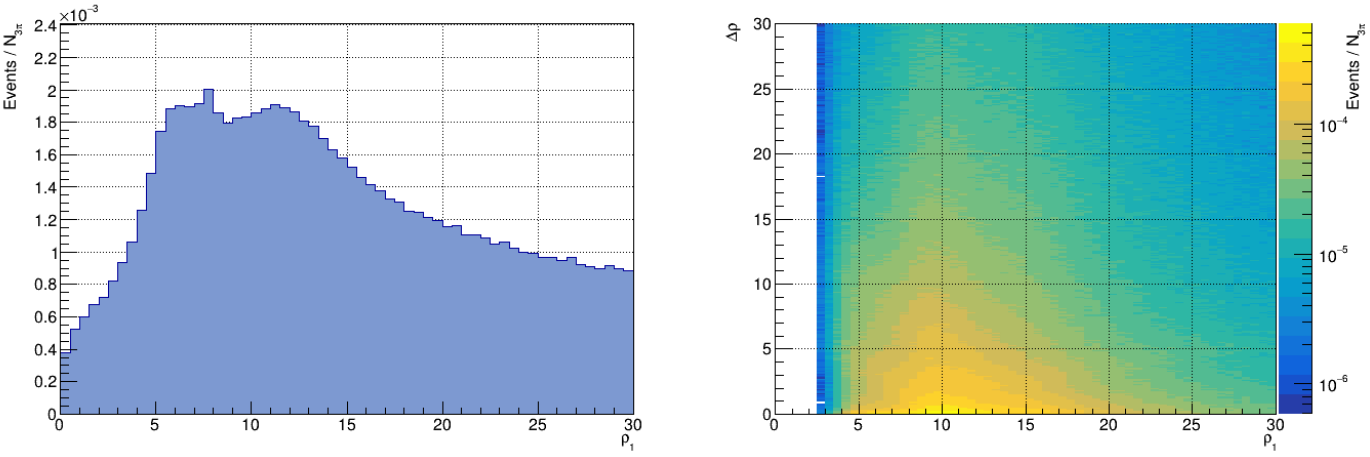

(c) Events with accidental associations.

Figure A.1: Distribution of ρ_1 in $N_{\mathrm{GTK}} = 1$ events (left) and of $(\rho_1, \Delta\rho)$ in $N_{\mathrm{GTK}} > 1$ events (right), for different classes of events.

Appendix B

LKr energy fine calibration

The LKr energy fine calibration is a procedure to measure an average calibration factor R as function of the LKr reconstructed energy E_{reco}, so that reconstructed energies can be corrected to calibrated energies E_{LKr} via

$$E_{\text{LKr}} = \frac{E_{\text{reco}}}{R(E_{\text{reco}})} \,.$$

In order to do so, a "standard candle" has to be identified, whose energy E can be calculated independently of the LKr: then, an evaluation of $R(E_{\text{reco}})$ is given by E_{reco}/E. It is important that this procedure is applied over small batches of data (also known as *runs*), so that the calibration reabsorbs possible differences in the LKr conditions across the data sample.

During my PhD course, I implemented a new calibration method, which has then been used for the $K^+ \to \pi^+\nu\bar{\nu}$ analysis of 2021–2022 data, and is currently used as default for all other analyses.

Two calibration methods were existing before. One, called *π^0-1D,* is described in [132], and is based on a sample $\pi^0 \to \gamma\gamma$, where the invariant mass of the diphoton is constrained to be the π^0 mass as provided by the PDG: $m_{\pi^0} = 134.9768\,\text{MeV}$ [27]. This method, however, tries to correct both photons at the same time, which results in complications and in a computationally expensive iterative procedure. The other (π^0_D) is based on π^0 Dalitz decays, $\pi^0 \to e^+e^-\gamma$: it is very clean, because the energy of the e^+e^- pair can be measured by the STRAW, but it is very much limited by statistics, as the Dalitz decay is pretty rare (BR $\approx 1.2\,\%$ [27]).

The new calibration method uses $\pi^0 \to \gamma\gamma$ decays as the π^0-1D method, but exploits each of the two photons separately, by reconstructing the π^0 from the $K^+ \to \pi^+\pi^0$ decay that produced it. For this reason it has been named *Single Photon.*

Concretely, a $K^+ \to \pi^+\pi^0$ sample is selected: in each event, based on the 3-momenta of K^+ and π^+ as measured by GTK and STRAW, respectively,

the 4-momentum of the π^0 is calculated as

$$\mathbf{p}_{\pi^0} = \mathbf{p}_{\mathrm{K}^+} - \mathbf{p}_{\pi^+} ,$$
$$E_{\pi^0} = \sqrt{m_{\pi^0}^2 + \left|\mathbf{p}_{\pi^0}\right|^2} .$$

Assuming no bias in the photon position as measured by the LKr, the direction of flight of each single photon can be calculated as the direction from the K^+–π^+ vertex to the photon position at the LKr plane. This direction can be used to calculate the angle θ between the directions of flight of the π^0 and of the photon. Then, by simple kinematics considerations, the energy of the photon is calculated:

$$E_\gamma = \frac{1}{2} \frac{m_{\pi^0}^2}{E_{\pi^0} - \left|\mathbf{p}_{\pi^0}\right| \cos\theta} .$$

This energy allows to calculate $R = E_{\mathrm{reco}}/E_\gamma$. Figure B.1 shows the distribution of R as function of E_{reco} for one run.

For each (1 GeV wide) bin of E_{reco}, the average R, with its statistical uncertainty, is calculated from a gaussian fit. From the average R as function of E_{reco}, the calibration $R(E_{\mathrm{reco}})$ is extracted via a fit with a high degree polynomial. The result is shown in figure B.2, together with the comparison with the previous LKr calibration methods. The limited statistics for the π^0_{D} method lead to the impossibility of fitting a trend of R as function of E_{reco}. The Single Photon method gives results that are consistent with the π^0-1D method, but without needing iterative procedures. Furthermore, the calibration looks more stable at high energies than the one provided by the π^0-1D method, so that both the fit range, and the range of validity for which to apply the calibration, is enlarged.

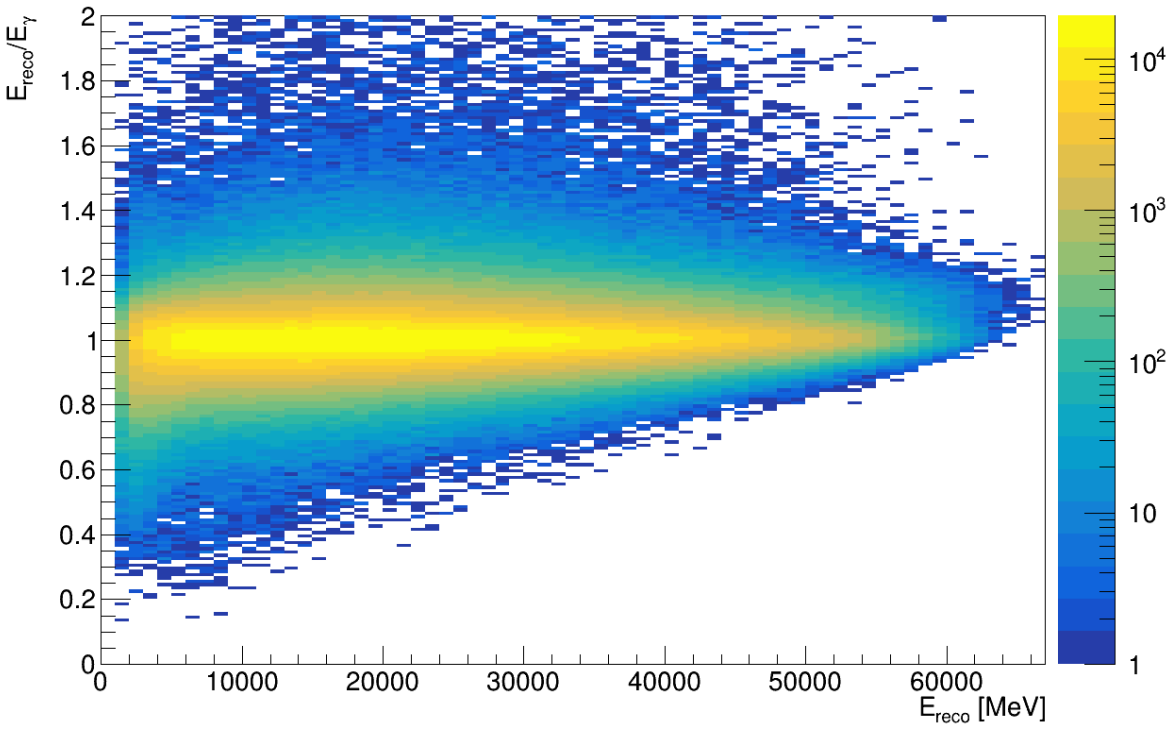

Figure B.1: Distribution of E_{reco}/E_γ as function of E_{reco}.

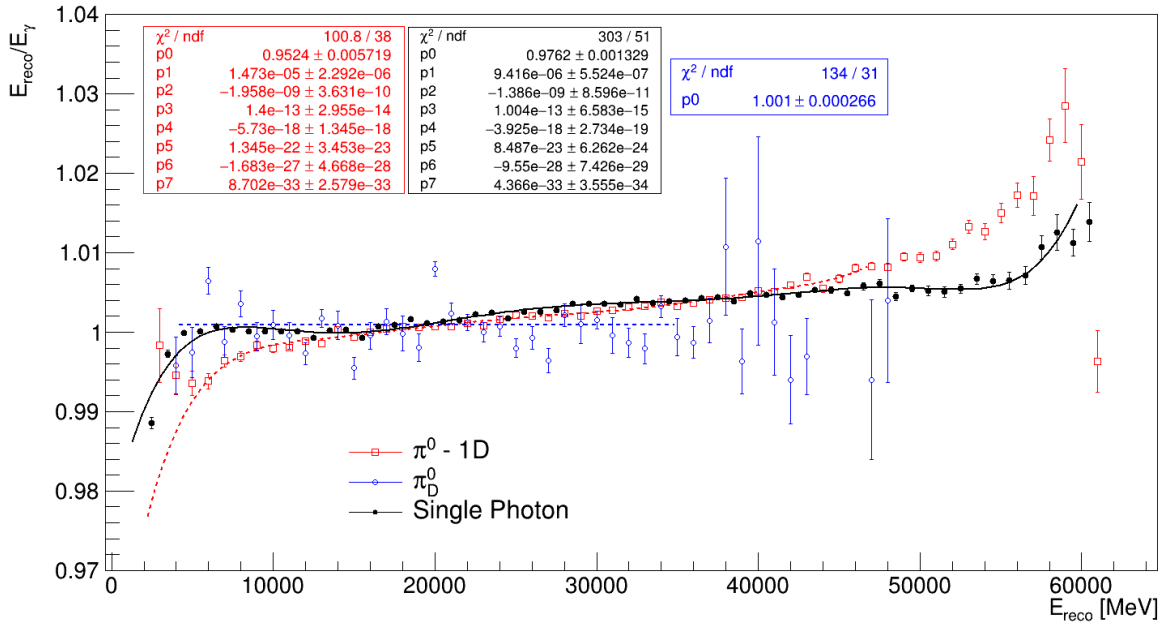

Figure B.2: Mean R as function of E_{reco}, compared between fine calibration methods.

Appendix C

CHANTI studies

During my PhD, as part of the CHANTI working group, I have performed multiple studies for the maintenance and improvement of the CHANTI detector. A selection of the most important ones is reported in this appendix.

The CHANTI and its electronics chain of the CHANTI are described in detail in [117]. The CHANTI is composed of 6 stations of 48 scintillator bars each, and each bar is read out by a SiPM. The 288 channels are connected to 9 front-end boards (FEE), managing 32 channels each, which apply the bias voltage to the SiPMs, monitor the SiPM currents and temperatures, and transfer signals from the SiPMs to the next electronics stage through trans-resistance amplifiers. The next electronic stage (*ToT boards*) is adapted from the readout electronics of the LAV [138], and outputs LVDS signals corresponding to leading and trailing times of the analog SiPM signal at two configurable thresholds. These LVDS signals are fed to TEL62 boards [125], which digitize them and sends them to the PC farm in response to L0 triggers.

During the 2021 data taking, most likely owing to the high intensity flares (see section 3.1), an increase of the SiPM dark currents of about 60 to 80 nA/day was observed (the nominal dark current was $O(10\,\text{nA})$). As a consequence, most channels had reached or exceeded 2.5 µA, which was the maximum readable current for the FEE, meaning that no monitoring of the dark currents was possible. SiPM replacement before the 2022 data taking was also not possible due to logistics, and no sizable deterioration of the CHANTI performance was seen in the 2021 data; the FEE boards were therefore modified to change the accuracy of the current reading from around 1 nA to around 10 nA, so that the maximum readable current could correspondingly increase to 25 µA. Since 2022, the dark current increase was much lower (typically 20 nA/day); SiPMs were finally replaced before the 2023 data taking.

C.1 Calibration of frontend electronics boards

On the FEE boards, each channel the reads current and voltage via AD7708 16 bit ADC, while the applied voltage is controlled via a LTC2620 DAC. After the modification of the FEE boards, each channel of each FEE was tested by connecting it to a variable resistance, and reading the current I through the resistance R via a Keithley™ 2400 programmable source measure unit, which provides both accurate power supply and precision low current measurements. Each channel was tested with resistances of 10 MΩ, 20 MΩ, and 100 MΩ; for each resistance value, the FEE channels were set to provide voltages corresponding to DAC count values from 0 to 2400, with spacing of 200. The calibration of the set voltage consisted in a linear fit of the points in the plane (DAC counts, $R \times I$); similarly, the read voltage was calibrated by fitting the points in the plane (ADC counts, $R \times I$), and the read current was calibrated by fitting the points in the plane (ADC counts, I).

Figure C.1 shows the calibration of a typical FEE channel. All channels were found to have an excellent linearity, for all three calibrations (χ^2/ndf typically less than 1/37), and the slope of the read current is found to be between 9.8 and 10.2 nA/ADC count for 96 % of the channels.

C.2 Characterization of SiPMs

After the 2022 data taking, before replacing the SiPMs, 384 new SiPMs were characterized. Each SiPM was powered via the Keithley™ 2400 source measure unit at different bias voltages V to reconstruct the I–V curve (5 measurements of I were done per V setting); I–V curves were measured at several temperatures (generally 25 °C, 30 °C and 40 °C). Each I–V curve was fitted with the function

$$I(V) = I_0 + \alpha(V - V_{bd})^2 H(V - V_{bd})$$

(V_{bd} is the breakdown voltage and H is the Heaviside function) to extract the I_0, V_{bd} and α parameters. The I–V curves and fit results of a typical SiPM are reported in figure C.2. The V_{bd} versus temperature T points were then fitted with $V_{bd}(T) = mT + q$: a typical result is shown in figure C.3.

The values of the dark current (from the $I(V)$ fit) at 25 °C were all found to be compatible with the nominal values (typically between 30 and 70 nA) within 15 %. All values of $V_{bd}(25\,°\text{C})$, as extracted from the $V_{bd}(T)$ fit, were found to be compatible with their respective nominal values within the experimental uncertainties. Most SiPMs were found to have $m \approx 50\,\text{mV}/°\text{C}$; the ones that were not compatible with this value were not used in the final SiPM replacement (only 288 SiPMs, out of the 384 tested, are needed to instrument the CHANTI completely). The value of m is especially important to validate because the during the data taking the SiPM bias voltage is adjusted with the measured temperature to keep them at a constant overvoltage of 3 V.

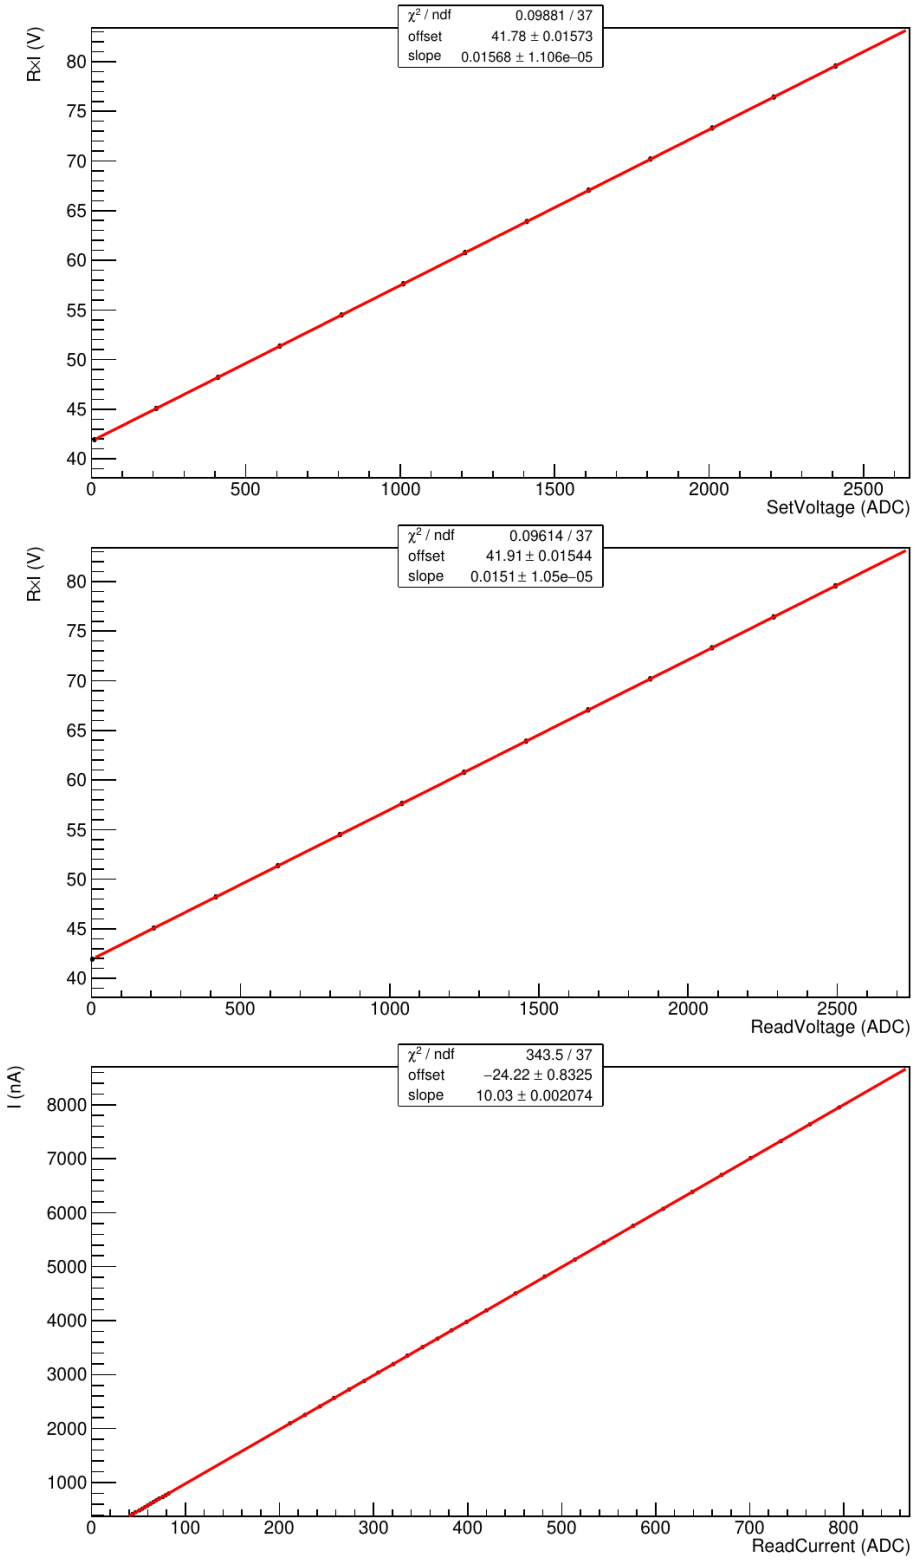

Figure C.1: Calibration of a typical FEE channel: set voltage (top), read voltage (middle), read current (bottom).

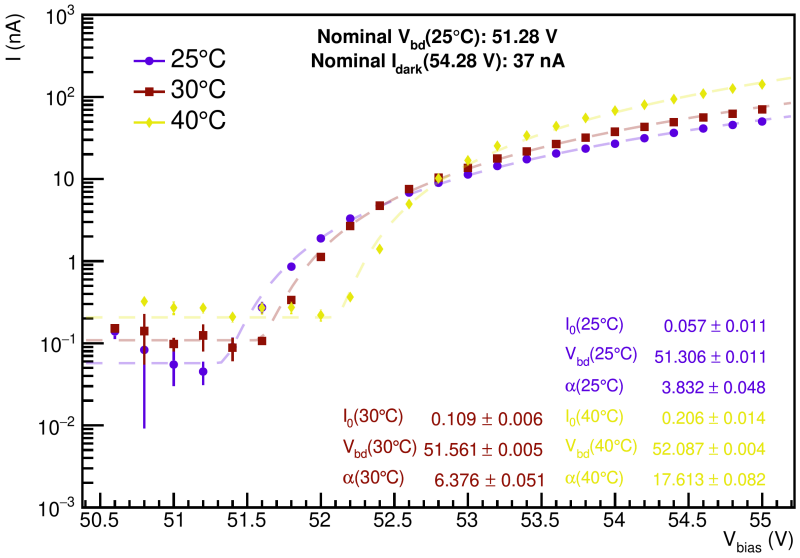

Figure C.2: Typical measured and fitted I–V curves of a SiPM. I_0 is provided in nA, V_{bd} is provided in volts, α in nA/V^2.

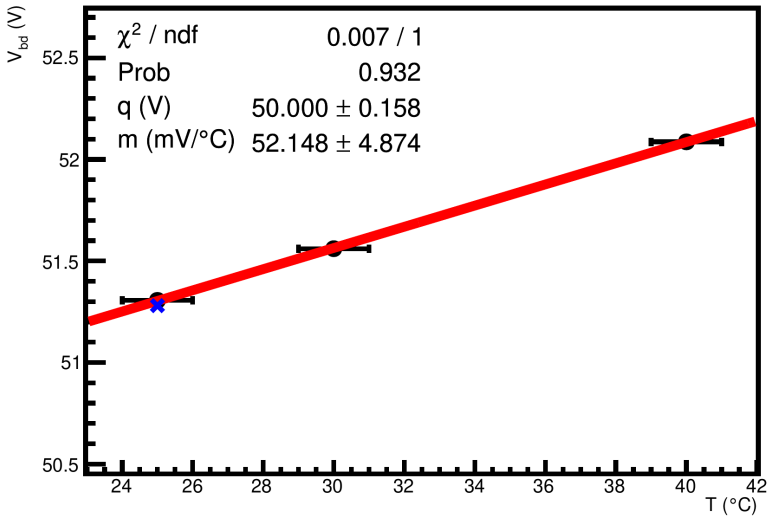

Figure C.3: Typical measured and fitted V_{bd} dependence on temperature. The blue cross denotes the nominal value.

C.3 Study of slewing corrections and time resolution

ToT boards provide leading and trailing times of two thresholds: this allows the study of slewing corrections. In particular, from the two leading times of signals that cross both thresholds, an extrapolation of the zero-crossing time can be made; for signals that cross only the low threshold, the correlation

between the time-over-threshold (ToT) and the time-walk correction can be exploited. A sketch is shown in figure C.4.

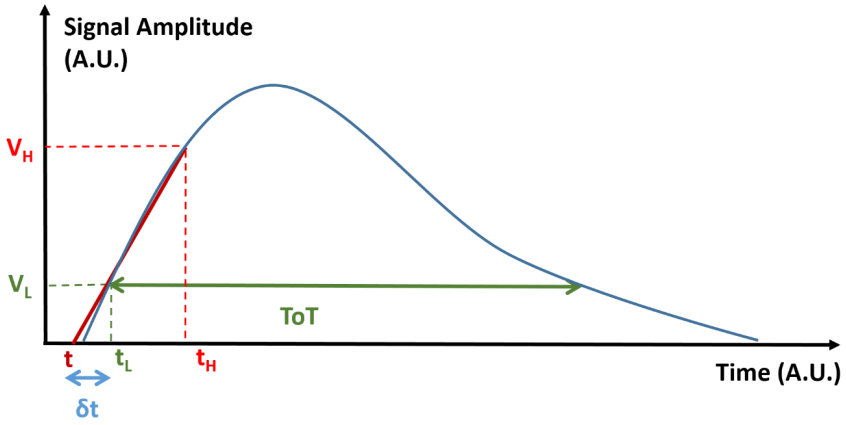

Figure C.4: Sketch of slewing corrections. [117]

The high threshold should be tuned according a trade-off: on one hand, a higher value means a more precise extrapolation for the zero-crossing time; on the other hand, it also means a smaller fraction of signals will cross it, thus increasing the fraction of signals that will have to be corrected with an average, ToT-based, correction.

At the beginning of the 2022 data-taking, data have been collected in the muon configuration (see section 2.3) where the illumination of the detector is more spread out, and composed mainly of muons.[1] During this muon run, the high threshold has been varied in about a dozen settings from 70 mV to 220 mV (while the low threshold was fixed at 60 mV), collecting about 100 spills per high threshold setting.

The comparison of the average slewing corrections as function of ToT for the different high threshold settings is shown in figure C.5. They were all found to be compatible as expected (the time walk as function of ToT should not depend on the high threshold that measures it). These average slewing corrections were therefore fitted, in order to extract the ToT-based correction for signals that do not cross the high threshold.

A study of the time resolution as function of the high threshold was then performed. The time resolution is measured from a gaussian fit of the distribution of the time differences between the corrected times and the trigger time (which is given by the RICH, and therefore has an intrinsic time resolution which is negligible compared to the CHANTI). The result of the time resolutions for a typical CHANTI view is shown in figure C.6, and for a typical CHANTI channel in figure C.7.

No minimum of the time resolution is found: the time resolution seems to saturate at the highest settings of the high threshold. Figure C.8 shows

[1] A few muon runs are usually collected at the beginning of each data-taking year, and they are used to calibrate many sub-detectors.

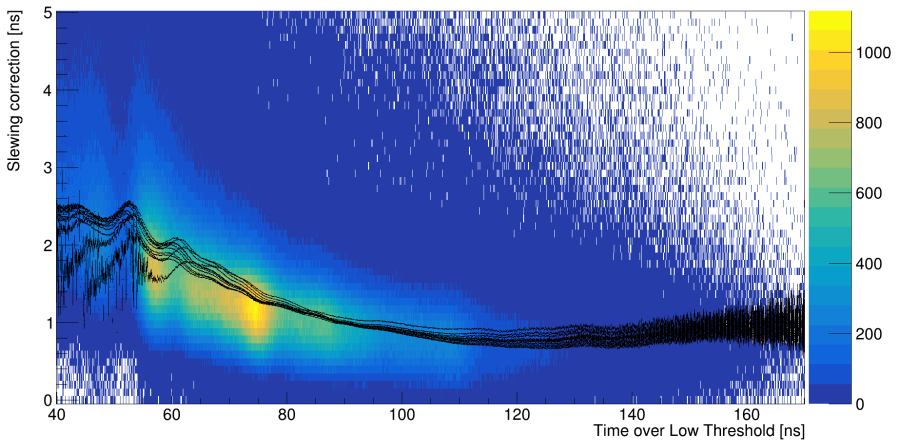

Figure C.5: Slewing corrections as function of the ToT. Color represents the slewing corrections in a spill of 2021 data. Black lines (one per high threshold settings) represent the average slewing corrections as function of the ToT.

the fraction of signals that pass the low threshold only, the high threshold only[2], and both thresholds, as function of the high threshold. As a result of this study, the high threshold was moved from 120 mV to 200 mV at the beginning of the 2022, resulting in an improved time resolution from around 1 ns to about 0.7 ns.

[2]This category of signals is constituted by signals generated by pileup particles that hit the detector right after the trailing time of the high threshold of a previous signal.

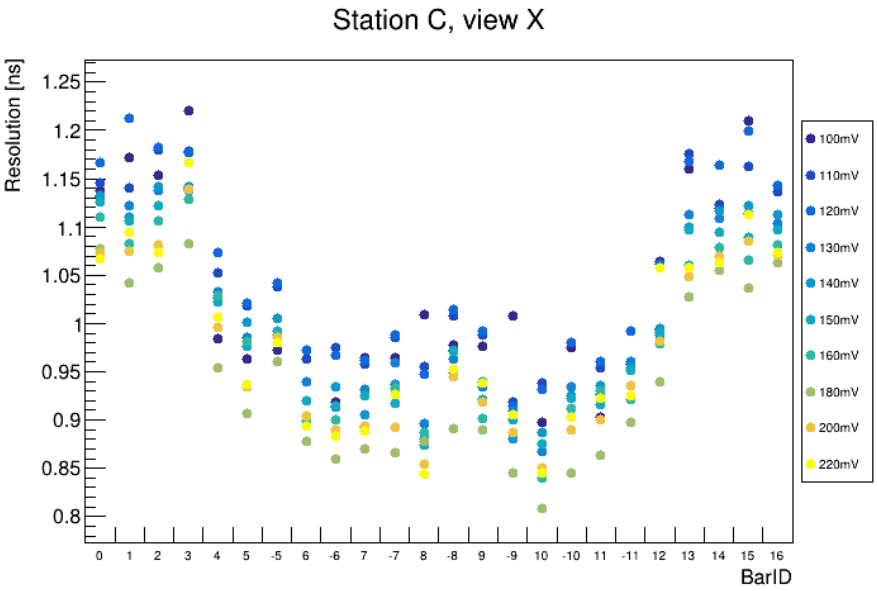

Figure C.6: Time resolution, for each channel of a typical CHANTI view, as function of the high threshold (color). CHANTI channels of each view are labelled by BarIDs: their absolute value is linearly related to the coordinate (in this case, the X coordinate); opposite-signed BarIDs correspond to scintillator bars that are on either side of the beam hole. Central bars are shorter, and therefore have an intrinsically smaller spread for their times.

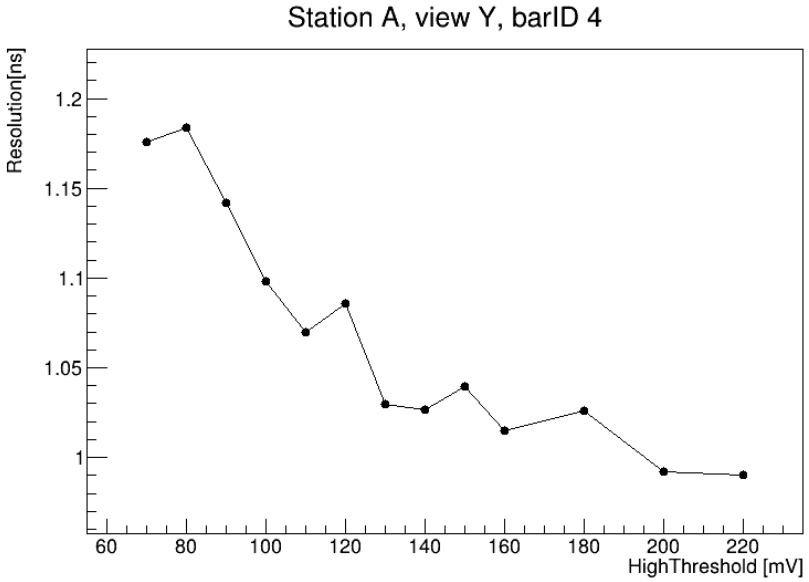

Figure C.7: Time resolution, for a typical CHANTI channel, as function of the high threshold.

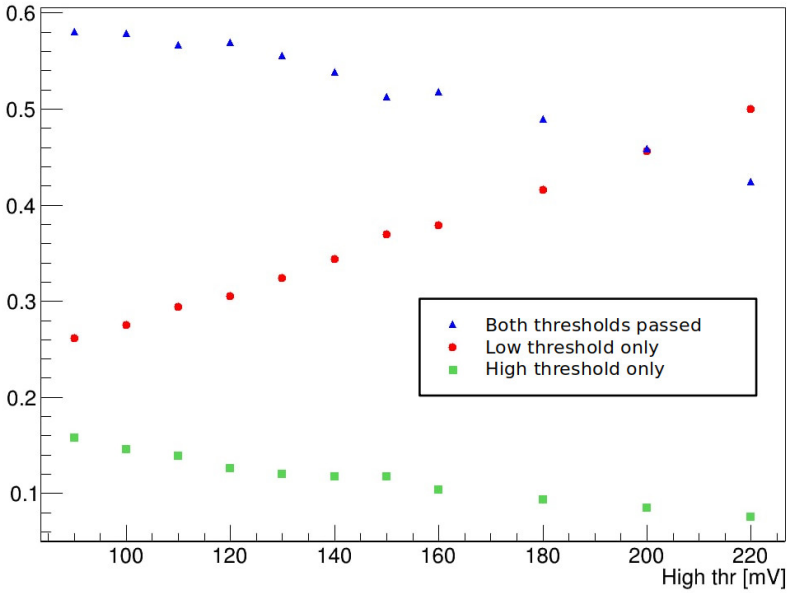

Figure C.8: Fraction of signals that pass the low threshold only, the high threshold only, and both thresholds, as function of the high threshold.

Bibliography

[1] S.L. Glashow, *Partial symmetries of weak interactions*, *Nucl. Phys.* **22** (1961) 579.

[2] F. Englert and R. Brout, *Broken symmetry and the mass of gauge vector mesons*, *Phys. Rev. Lett.* **13** (1964) 321.

[3] P.W. Higgs, *Broken symmetries and the masses of gauge bosons*, *Phys. Rev. Lett.* **13** (1964) 508.

[4] M. Gell-Mann, *A Schematic Model of Baryons and Mesons*, *Phys. Lett.* **8** (1964) 214.

[5] G. Zweig, *An SU_3 model for strong interaction symmetry and its breaking*, CDS Report CERN-TH-401 (Jan. 17, 1964), DOI: 10.17181/CERN-TH-401.

[6] G. Zweig, *An SU_3 model for strong interaction symmetry and its breaking II*, CDS Report CERN-TH-412 (Feb. 21, 1964), DOI: 10.17181/CERN-TH-412. Published in *Developments in the Quark Theory of Hadrons*, Volume 1. Edited by D. Lichtenberg and S. Rosen. Nonantum, Mass., Hadronic Press, 1980. pp. 22-101.

[7] S. Weinberg, *A model of leptons*, *Phys. Rev. Lett.* **19** (1967) 1264.

[8] A. Salam, *Weak and electromagnetic interactions*, *Conf. Proc. C* **680519** (1968) 367.

[9] S.L. Glashow, J. Iliopoulos and L. Maiani, *Weak Interactions with Lepton-Hadron Symmetry*, *Phys. Rev. D* **2** (1970) 1285.

[10] D.J. Gross and F. Wilczek, *Ultraviolet Behavior of Nonabelian Gauge Theories*, *Phys. Rev. Lett.* **30** (1973) 1343.

[11] S. Weinberg, *The Making of the Standard Model*, *Eur. Phys. J. C* **34** (2004) 5, arXiv:hep-ph/0401010.

[12] J.F. Donoghue, E. Golowich and B.R. Holstein, *Dynamics of the Standard Model*, Cambridge University Press, 2nd ed. (2022). ISBN: 978-1-009-29100-2, DOI: 10.1017/9781009291033.

[13] GARGAMELLE NEUTRINO collaboration, *Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment*, *Phys. Lett. B* **46** (1973) 138.

[14] E598 collaboration, *Experimental Observation of a Heavy Particle J*, *Phys. Rev. Lett.* **33** (1974) 1404.

[15] SLAC-SP-017 collaboration, *Discovery of a Narrow Resonance in e^+e^- Annihilation*, *Phys. Rev. Lett.* **33** (1974) 1406.

[16] M.L. Perl et al., *Evidence for Anomalous Lepton Production in e^+ – e^- Annihilation*, *Phys. Rev. Lett.* **35** (1975) 1489.

[17] E288 collaboration, *Observation of a Dimuon Resonance at* 9.5 GeV *in* 400 GeV *Proton-Nucleus Collisions*, *Phys. Rev. Lett.* **39** (1977) 252.

[18] UA1 collaboration, *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at* $\sqrt{s} = 540\,\text{GeV}$, *Phys. Lett. B* **122** (1983) 103.

[19] UA2 collaboration, *Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN $\bar{p}p$ Collider*, *Phys. Lett. B* **122** (1983) 476.

[20] UA1 collaboration, *Experimental Observation of Lepton Pairs of Invariant Mass Around* $95\,\text{GeV}/c^2$ *at the CERN SPS Collider*, *Phys. Lett. B* **126** (1983) 398.

[21] UA2 collaboration, *Evidence for $Z^0 \to e^+e^-$ at the CERN $\bar{p}p$ Collider*, *Phys. Lett. B* **129** (1983) 130.

[22] CDF collaboration, *Observation of top quark production in $\bar{p}p$ collisions*, *Phys. Rev. Lett.* **74** (1995) 2626, arXiv:hep-ex/9503002.

[23] D0 collaboration, *Observation of the top quark*, *Phys. Rev. Lett.* **74** (1995) 2632, arXiv:hep-ex/9503003.

[24] DONUT collaboration, *Observation of tau neutrino interactions*, *Phys. Lett. B* **504** (2001) 218, arXiv:hep-ex/0012035.

[25] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1, arXiv:1207.7214 [hep-ex].

[26] CMS collaboration, *Observation of a New Boson at a Mass of* 125 GeV *with the CMS Experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30, arXiv:1207.7235 [hep-ex].

[27] S. Navas et al. (Particle Data Group), *Review of Particle Physics*, *Phys. Rev. D* **110** (2024) 030001.

[28] A. Ceccucci et al., *Proposal to measure the rare decay $K^+ \to \pi^+ \nu\bar{\nu}$ at the CERN SPS*, Tech. Rep. CERN-SPSC-2005-013, SPSC-P-326, CERN, Geneva (2005), DOI: 10.17181/CERN.HVBQ.34HD.

[29] 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*, *JHEP* **11** (2015) 033, arXiv:1503.02693 [hep-ph].

[30] J. Brod, M. Gorbahn and E. Stamou, *Updated Standard Model Prediction for $K \to \pi\nu\bar{\nu}$ and ϵ_K*, *PoS* **BEAUTY2020** (2021) 056, arXiv:2105.02868 [hep-ph].

[31] G. D'Ambrosio, A.M. Iyer, F. Mahmoudi and S. Neshatpour, *Anatomy of kaon decays and prospects for lepton flavour universality violation*, *JHEP* **09** (2022) 148, arXiv:2206.14748 [hep-ph].

[32] A.J. Buras and E. Venturini, *The exclusive vision of rare K and B decays and of the quark mixing in the standard model*, *Eur. Phys. J. C* **82** (2022) 615, arXiv:2203.11960 [hep-ph].

[33] L. Allwicher, M. Bordone, G. Isidori, G. Piazza and A. Stanzione, *Probing third-generation New Physics with $K \to \pi\nu\bar{\nu}$ and $B \to K^{(*)}\nu\bar{\nu}$* , arXiv:2410.21444 [hep-ph].

[34] NA62 collaboration, *Measurement of the very rare $K^+ \to \pi^+ \nu\bar{\nu}$ decay*, *JHEP* **06** (2021) 093, arXiv:2103.15389 [hep-ex].

[35] J. Swallow (NA62 collaboration), *A new measurement of $\mathrm{BR}(K^+ \to \pi^+ \nu\bar{\nu})$ by the NA62 experiment*, CERN EP Seminar (Sept. 24, 2024), https://indico.cern.ch/event/1447422/.

[36] R. Fiorenza (NA62 collaboration), *New measurement of $\mathrm{BR}(K^+ \to \pi^+ \nu\bar{\nu})$*, contribution to *43rd International Symposium on Physics in Collision* (Oct. 22, 2024), https://indico.cern.ch/event/1414470/contributions/6136109/.

[37] G.D. Rochester and C.C. Butler, *Evidence for the Existence of New Unstable Elementary Particles*, *Nature* **160** (1947) 855.

[38] T.D. Lee and C.N. Yang, *Question of Parity Conservation in Weak Interactions*, *Phys. Rev.* **104** (1956) 254.

[39] C.S. Wu, E. Ambler, R.W. Hayward, D.D. Hoppes and R.P. Hudson, *Experimental Test of Parity Conservation in Beta Decay*, *Phys. Rev.* **105** (1957) 1413.

[40] 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** (1964) 138.

[41] NA48 collaboration, *A new measurement of direct CP violation in two pion decays of the neutral kaon*, *Phys. Lett. B* **465** (1999) 335, arXiv:hep-ex/9909022.

[42] KTeV collaboration, *Observation of direct CP violation in $K_{S,L} \to \pi\pi$ decays*, *Phys. Rev. Lett.* **83** (1999) 22, arXiv:hep-ex/9905060.

[43] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, *Phys. Rev. Lett.* **10** (1963) 531.

[44] M. Kobayashi and T. Maskawa, *CP-Violation in the Renormalizable Theory of Weak Interaction*, *Prog. Theor. Phys.* **49** (1973) 652.

[45] L. Wolfenstein, *Parametrization of the Kobayashi-Maskawa Matrix*, *Phys. Rev. Lett.* **51** (1983) 1945.

[46] A.J. Buras, *Weak Hamiltonian, CP Violation and Rare Decays*, contribution to *Les Houches Summer School in Theoretical Physics, Session 68: Probing the Standard Model of Particle Interactions*, p. 281 (1998), arXiv:hep-ph/9806471.

[47] F. Mescia and C. Smith, *Improved estimates of rare K decay matrix elements from $K_{\ell 3}$ decays*, *Phys. Rev. D* **76** (2007) 034017, arXiv:0705.2025 [hep-ph].

[48] G. Isidori, F. Mescia and C. Smith, *Light-quark loops in $K^+ \to \pi^+ \nu\bar{\nu}$*, *Nucl. Phys. B* **718** (2005) 319, arXiv:hep-ph/0503107.

[49] A.J. Buras and E. Venturini, *Searching for New Physics in Rare K and B Decays without $|V_{cb}|$ and $|V_{ub}|$ Uncertainties*, *Acta Phys. Polon. B* **53.6** (2021) A1, arXiv:2109.11032 [hep-ph].

[50] A.J. Buras, S. Uhlig and F. Schwab, *Waiting for precise measurements of $K^+ \to \pi^+ \nu\bar{\nu}$ and $K_L \to \pi^0 \nu\bar{\nu}$*, *Rev. Mod. Phys.* **80** (2008) 965, arXiv:hep-ph/0405132.

[51] Y. Grossman and Y. Nir, *$K_L \to \pi^0 \nu\bar{\nu}$ beyond the standard model*, *Phys. Lett. B* **398** (1997) 163, arXiv:hep-ph/9701313.

[52] A.J. Buras, D. Buttazzo and R. Knegjens, *$K \to \pi\nu\bar{\nu}$ and ε'/ε in simplified new physics models*, *JHEP* **11** (2015) 166, arXiv:1507.08672 [hep-ph].

[53] J. Aebischer, A.J. Buras and J. Kumar, *Another SMEFT story: Z' facing new results on ϵ'/ϵ, ΔM_K and $K \to \pi\nu\bar{\nu}$*, *JHEP* **12** (2020) 097, arXiv:2006.01138 [hep-ph].

[54] F.F. Deppisch, K. Fridell and J. Harz, *Constraining lepton number violating interactions in rare kaon decays*, *JHEP* **12** (2020) 186, arXiv:2009.04494 [hep-ph].

[55] M. Gorbahn, U. Moldanazarova, K.H. Sieja, E. Stamou and M. Tabet, *The anatomy of $K^+ \to \pi^+\nu\bar{\nu}$ distributions*, *Eur. Phys. J. C* **84** (2024) 680, arXiv:2312.06494 [hep-ph].

[56] C. Bobeth and A.J. Buras, *Leptoquarks meet ε'/ε and rare Kaon processes*, *JHEP* **02** (2018) 101, arXiv:1712.01295 [hep-ph].

[57] D. Marzocca, S. Trifinopoulos and E. Venturini, *From B-meson anomalies to Kaon physics with scalar leptoquarks*, *Eur. Phys. J. C* **82** (2022) 320, arXiv:2106.15630 [hep-ph].

[58] Ò.L. Crosas, G. Isidori, J.M. Lizana, N. Selimović and B.A. Stefanek, *Flavor non-universal vector leptoquark imprints in $K \to \pi\nu\bar{\nu}$ and $\Delta F = 2$ transitions*, *Phys. Lett. B* **835** (2022) 137525, arXiv:2207.00018 [hep-ph].

[59] M. Bordone, D. Buttazzo, G. Isidori and J. Monnard, *Probing Lepton Flavour Universality with $K \to \pi\nu\bar{\nu}$ decays*, *Eur. Phys. J. C* **77** (2017) 618, arXiv:1705.10729 [hep-ph].

[60] S. Descotes-Genon, S. Fajfer, J.F. Kamenik and M. Novoa-Brunet, *Implications of $b \to s\mu\mu$ anomalies for future measurements of $B \to K^{(*)}\nu\bar{\nu}$ and $K \to \pi\nu\bar{\nu}$*, *Phys. Lett. B* **809** (2020) 135769, arXiv:2005.03734 [hep-ph]. Addendum: *Phys. Lett. B* **840** (2023) 137830.

[61] A.J. Buras, J. Harz and M.A. Mojahed, *Disentangling new physics in $K \to \pi\nu\bar{\nu}$ and $B \to K(K^*)\nu\bar{\nu}$ observables*, arXiv:2405.06742 [hep-ph].

[62] Y. Asano et al., *Search for a rare decay mode $K^+ \to \pi^+\nu\bar{\nu}$ and axion*, *Phys. Lett. B* **107** (1981) 159.

[63] E949 collaboration, *Study of the decay $K^+ \to \pi^+\nu\bar{\nu}$ in the momentum region $140 < P_\pi < 199\,\mathrm{MeV}/c$*, *Phys. Rev. D* **79** (2009) 092004, arXiv:0903.0030 [hep-ex].

[64] NA62 collaboration, *First search for $K^+ \to \pi^+\nu\bar{\nu}$ using the decay-in-flight technique*, *Phys. Lett. B* **791** (2019) 156, arXiv:1811.08508 [hep-ex].

[65] NA62 collaboration, *An investigation of the very rare $K^+ \to \pi^+\nu\bar{\nu}$ decay*, *JHEP* **11** (2020) 042, arXiv:2007.08218 [hep-ex].

[66] U. Camerini, D. Ljung, M. Sheaff and D. Cline, *Experimental search for semileptonic neutrino neutral currents*, *Phys. Rev. Lett.* **23** (1969) 326.

[67] J.H. Klems, R.H. Hildebrand and R. Stiening, *Limits on the $K^+ \to \pi^+ + \nu + \bar{\nu}$ and $K^+ \to \pi^+ + n\gamma$ decay rates*, *Phys. Rev. D* **4** (1971) 66.

[68] G.D. Cable, R.H. Hildebrand, C.Y. Pang and R. Stiening, *Search for rare K^+ decays. II. $K^+ \to \pi^+\nu\bar{\nu}$*, *Phys. Rev. D* **8** (1973) 3807.

[69] M.K. Gaillard and B.W. Lee, *Rare Decay Modes of the K-Mesons in Gauge Theories*, *Phys. Rev. D* **10** (1974) 897.

[70] J.R. Ellis, J.S. Hagelin and S. Rudaz, *Reexamination of the Standard Model in the Light of B Meson Mixing*, *Phys. Lett. B* **192** (1987) 201.

[71] E787 collaboration, *Search for the Decay $K^+ \to \pi^+\nu\bar{\nu}$*, *Phys. Rev. Lett.* **64** (1990) 21.

[72] E787 collaboration, *Search for the Decay $K^+ \to \pi^+\nu\bar{\nu}$*, *Phys. Rev. Lett.* **70** (1993) 2521. Erratum: *Phys.Rev.Lett.* **71** (1993) 305.

[73] E787 collaboration, *Search for the Decay $K^+ \to \pi^+\nu\bar{\nu}$*, *Phys. Rev. Lett.* **76** (1996) 1421, arXiv:hep-ex/9510006.

[74] G. Buchalla and A.J. Buras, *$K \to \pi\nu\bar{\nu}$ and High Precision Determinations of the CKM Matrix*, *Phys. Rev. D* **54** (1996) 6782, arXiv:hep-ph/9607447.

[75] J. Brod, M. Gorbahn and E. Stamou, *Two-Loop Electroweak Corrections for the $K \to \pi\nu\bar{\nu}$ Decays*, *Phys. Rev. D* **83** (2011) 034030, arXiv:1009.0947 [hep-ph].

[76] *Minutes of the 237th meeting of the Research Board, held on 9 June 2021*, Tech. Rep. CERN-DG-RB-2021-505, M-237, CERN, Geneva (2021).

[77] NA62 collaboration, *2021 NA62 Status Report to the CERN SPSC*, Tech. Rep. CERN-SPSC-2021-009, SPSC-SR-286, CERN, Geneva (2021).

[78] NA62 collaboration, *2022 NA62 Status Report to the CERN SPSC*, Tech. Rep. CERN-SPSC-2022-012, SPSC-SR-306, CERN, Geneva (2022).

[79] NA62 collaboration, *2023 NA62 Status Report to the CERN SPSC*, Tech. Rep. CERN-SPSC-2023-013, SPSC-SR-326, CERN, Geneva (2023).

[80] NA62 collaboration, *2024 NA62 Status Report to the CERN SPSC*, Tech. Rep. CERN-SPSC-2024-012, SPSC-SR-345, CERN, Geneva (2024).

[81] NA62 collaboration, *ADDENDUM I TO P326: Continuation of the physics programme of the NA62 experiment*, Tech. Rep. CERN-SPSC-2019-039, SPSC-P-326-ADD-1, CERN, Geneva (2019).

[82] E391A collaboration, *Experimental study of the decay* $K_L^0 \to \pi^0\nu\bar{\nu}$, *Phys. Rev. D* **81** (2010) 072004, arXiv:0911.4789 [hep-ex].

[83] KOTO collaboration, *Search for the* $K_L \to \pi^0\nu\bar{\nu}$ *and* $K_L \to \pi^0 X^0$ *decays at the J-PARC KOTO experiment*, *Phys. Rev. Lett.* **122** (2019) 021802, arXiv:1810.09655 [hep-ex].

[84] KOTO collaboration, *Study of the* $K_L \to \pi^0\nu\bar{\nu}$ *Decay at the J-PARC KOTO Experiment*, *Phys. Rev. Lett.* **126** (2021) 121801, arXiv:2012.07571 [hep-ex].

[85] KOTO collaboration, *Search for the* $K_L \to \pi^0\nu\bar{\nu}$ *Decay at the J-PARC KOTO Experiment*, arXiv:2411.11237 [hep-ex].

[86] NA62 collaboration, *Search for a feebly interacting particle* X *in the decay* $K^+ \to \pi^+ X$, *JHEP* **03** (2021) 058, arXiv:2011.11329 [hep-ex].

[87] NA62 collaboration, *A measurement of the* $K^+ \to \pi^+\mu^+\mu^-$ *decay*, *JHEP* **11** (2022) 011, arXiv:2209.05076 [hep-ex].

[88] NA62 collaboration, *A study of the* $K^+ \to \pi^0 e^+\nu\gamma$ *decay*, *JHEP* **09** (2023) 040, arXiv:2304.12271 [hep-ex].

[89] NA62 collaboration, *Search for* K^+ *decays into the* $\pi^+e^+e^-e^+e^-$ *final state*, *Phys. Lett. B* **846** (2023) 138193, arXiv:2307.04579 [hep-ex].

[90] NA62 collaboration, *Measurement of the* $K^+ \to \pi^+\gamma\gamma$ *decay*, *Phys. Lett. B* **850** (2024) 138513, arXiv:2311.01837 [hep-ex].

[91] NA62 collaboration, *Searches for lepton number violating* K^+ *decays*, *Phys. Lett. B* **797** (2019) 134794, arXiv:1905.07770 [hep-ex].

[92] NA62 collaboration, *Search for lepton number and flavor violation in* K^+ *and* π^0 *decays*, *Phys. Rev. Lett.* **127** (2021) 131802, arXiv:2105.06759 [hep-ex].

[93] NA62 collaboration, *Searches for lepton number violating $K^+ \to \pi^-(\pi^0)e^+e^+$ decays*, *Phys. Lett. B* **830** (2022) 132172, arXiv:2202.00331 [hep-ex].

[94] NA62 collaboration, *A search for the $K^+ \to \mu^+\nu e^+e^+$ decay*, *Phys. Lett. B* **838** (2022) 137679, arXiv:2211.04818 [hep-ex].

[95] NA62 collaboration, *First search for $K^+ \to \pi^0\pi\mu e$ decays*, arXiv:2409.12981 [hep-ex].

[96] NA62 collaboration, *Search for production of an invisible dark photon in π^0 decays*, *JHEP* **05** (2019) 182, arXiv:1903.08767 [hep-ex].

[97] NA62 collaboration, *Search for heavy neutral lepton production in K^+ decays to positrons*, *Phys. Lett. B* **807** (2020) 135599, arXiv:2005.09575 [hep-ex].

[98] NA62 collaboration, *Search for π^0 decays to invisible particles*, *JHEP* **02** (2021) 201, arXiv:2010.07644 [hep-ex].

[99] NA62 collaboration, *Search for K^+ decays to a muon and invisible particles*, *Phys. Lett. B* **816** (2021) 136259, arXiv:2101.12304 [hep-ex].

[100] NA62 collaboration, *Search for dark photon decays to $\mu^+\mu^-$ at NA62*, *JHEP* **09** (2023) 035, arXiv:2303.08666 [hep-ex].

[101] NA62 collaboration, *Search for Leptonic Decays of Dark Photons at NA62*, *Phys. Rev. Lett.* **133** (2024) 111802, arXiv:2312.12055 [hep-ex].

[102] F. Hahn et al. (NA62 collaboration), *NA62: Technical Design Document*, Tech. Rep. NA62-10-07, CERN, Geneva (2010).

[103] NA62 collaboration, *The beam and detector of the NA62 experiment at CERN*, *JINST* **12** (2017) P05025, arXiv:1703.08501 [physics.ins-det].

[104] C. Parkinson, M. Corvino, N. Lurkin, A. Bragadireanu, S. Ghinescu and S. Kholodenko, *Transforming NA62*, in *Newsletter of the EP department*, CERN (Dec. 13, 2021), https://ep-news.web.cern.ch/content/transforming-na62. URL consulted in Sept. 2024.

[105] H. Danielsson et al., *New veto hodoscope ANTI-0 for the NA62 experiment at CERN*, *JINST* **15** (2020) C07007, arXiv:2004.09344 [physics.ins-det].

[106] NA62 collaboration, *2019 NA62 Status Report to the CERN SPSC*, Tech. Rep. CERN-SPSC-2019-012, SPSC-SR-249, CERN, Geneva (2019).

[107] C. Bovet, R. Maleyran, L. Piemontese, A. Placci and M. Placidi, *The CEDAR counters for particle identification in the SPS secondary beams: a description and an operation manual*, Tech. Rep. CERN-82-13, CERN, Geneva (1982), DOI: 10.5170/CERN-1982-013.

[108] E. Goudzovski et al., *Development of the kaon tagging system for the NA62 experiment at CERN*, *Nucl. Instrum. Meth. A* **801** (2015) 86, arXiv:1509.03773 [physics.ins-det].

[109] F. Anghinolfi et al., *NINO: An ultra-fast and low-power front-end amplifier/discriminator ASIC designed for the multigap resistive plate chamber*, *Nucl. Instrum. Meth. A* **533** (2004) 183.

[110] NA62 collaboration, *Development of a new CEDAR for kaon identification at the NA62 experiment at CERN*, *JINST* **19** (2024) P05005, arXiv:2312.17188 [hep-ex].

[111] G. Aglieri Rinella et al., *The NA62 GigaTracKer: a low mass high intensity beam 4D tracker with 65 ps time resolution on tracks*, *JINST* **14** (2019) P07010, arXiv:1904.12837 [physics.ins-det].

[112] G.A. Rinella et al., *TDCpix pixel detector ASIC with* 100 ps *time stamping*, *Nucl. Instrum. Meth. A* **1053** (2023) 168331.

[113] A. Kleimenova (GigaTracKer working group), *Operation and performance of the NA62 GigaTracKer*, *PoS* **VERTEX2023** (2024) 008.

[114] M. Ceoletta, *The NA62 VetoCounter sub-detector readout system upgrade and performance evaluation*, Master's thesis, Università degli Studi di Padova (2023), https://hdl.handle.net/20.500.12608/45502.

[115] NA62 collaboration, *2020 NA62 Status Report to the CERN SPSC*, Tech. Rep. CERN-SPSC-2020-007, SPSC-SR-266, CERN, Geneva (2020).

[116] W. Wu (ATLAS TDAQ collaboration), *FELIX: the New Detector Interface for the ATLAS Experiment*, *IEEE Trans. Nucl. Sci.* **66** (2019) 986, arXiv:1806.10667 [physics.ins-det].

[117] F. Ambrosino et al., *CHANTI: a Fast and Efficient Charged Particle Veto Detector for the NA62 Experiment at CERN*, *JINST* **11** (2016) P03029, arXiv:1512.00244 [physics.ins-det].

[118] D. Aisa et al., *Mirror system of the RICH detector of the NA62 experiment*, *JINST* **12** (2017) P12017.

[119] G. Anzivino et al., *Precise mirror alignment and basic performance of the RICH detector of the NA62 experiment at CERN*, *JINST* **13** (2018) P07012, arXiv:1809.04026 [physics.ins-det].

[120] G. Anzivino et al., *Light Detection System and Time Resolution of the NA62 RICH*, *JINST* **15** (2020) P10025, arXiv:2009.07581 [physics.ins-det].

[121] NA48 collaboration, *The beam and detector for the NA48 neutral kaon CP violations experiment at CERN*, *Nucl. Instrum. Meth. A* **574** (2007) 433.

[122] A. Ceccucci, R. Fantechi, P. Farthouat, G. Lamanna and V. Ryjov, *The NA62 Liquid Krypton calorimeter readout module*, *JINST* **6** (2011) C12017.

[123] R. Aliberti, *Particle identification with calorimeters for the measurement of the rare recay $K^+ \to \pi^+ \nu\bar{\nu}$ at NA62*, Ph.D. thesis, Johannes Gutenberg University Mainz (2018), http://cds.cern.ch/record/2671696.

[124] OPAL collaboration, *The OPAL detector at LEP*, *Nucl. Instrum. Meth. A* **305** (1991) 275.

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

[126] 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** (2015) 1050, arXiv:1407.2456 [physics.ins-det].

[127] R. Ammendola et al., *The integrated low-level trigger and readout system of the CERN NA62 experiment*, *Nucl. Instrum. Meth. A* **929** (2019) 1, arXiv:1903.10200 [physics.ins-det].

[128] D. Soldi and S. Chiozzi, *Level Zero Trigger Processor for the NA62 experiment*, *JINST* **13** (2018) P05004, arXiv:1802.06548 [physics.ins-det].

[129] NA62 collaboration, *Performance of the NA62 trigger system*, *JHEP* **03** (2023) 122, arXiv:2208.00897 [hep-ex].

[130] 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*, Tech. Rep. CERN-74-02, CERN, Geneva (1974), DOI: 10.5170/CERN-1974-002.

[131] J. Allison et al., *Recent developments in Geant4*, *Nucl. Instrum. Meth. A* **835** (2016) 186.

[132] F. Brizioli, *Measurement of* BR($K^+ \to \pi^+ \nu\bar{\nu}$) *with the NA62 experiment at CERN*, Ph.D. thesis, Università degli Studi di Perugia (2021), https://cds.cern.ch/record/2765463.

[133] U. Müller et al., *Particle identification with the RICH detector in experiment WA89 at CERN*, *Nucl. Instrum. Meth. A* **343** (1994) 279.

[134] S.S. Wilks, *The large-sample distribution of the likelihood ratio for testing composite hypotheses*, *Ann. Math. Statist.* **9** (1938) 60.

[135] HIKE collaboration, *HIKE, High Intensity Kaon Experiments at the CERN SPS: Letter of Intent*, arXiv:2211.16586 [hep-ex].

[136] HIKE collaboration, *High Intensity Kaon Experiments (HIKE) at the CERN SPS: Proposal for Phases 1 and 2*, arXiv:2311.08231 [hep-ex].

[137] G. Anzivino et al., *Workshop summary: Kaons@CERN 2023*, *Eur. Phys. J. C* **84** (2024) 377, arXiv:2311.02923 [hep-ph].

[138] A. Antonelli et al., *The NA62 LAV front-end electronics*, *JINST* **7** (2012) C01097, arXiv:1111.5768 [physics.ins-det].