



Hidden sector searches with fixed-target experiments

Doctoral dissertation presented by

Jan Jerhot

in fulfilment of the requirements for the degree of Doctor in Sciences.

November, 2023

Acknowledgements

I acknowledge the funding of this project by the F.R.S.-FNRS, Belgium, through grant FRIA/FC-36305.

I want to express my gratitude to my supervisor, Eduardo Cortina Gil for guiding me through the PhD and even making this project possible at all by polishing the proposal to perfection to secure the funding. I would also like to thank my former supervisors. To Karol Kampf for directing me to this career path and introducing me to the flavor physics and amplitude communities which stand at the forefront of the current HEP research, and to Michal Zamkovsky for introducing me to the NA62 collaboration and explaining to me with patience all the details of the experiment from the analysis software to the last hardware readout board.

I am deeply grateful to Babette Döbrich and Tommaso Spadaro for giving me the opportunity to join the various hidden sector phenomenology projects, which became a major part of my research, and for their guidance over the years and the supervision of the analyses of the NA62 beam-dump data.

I thank Marco Drewes for the discussions about the cosmological aspects of my research and for the opportunity to tutor the QM classes. A warm thanks also to Christophe Delaere and Fabio Maltoni for being a part of my thesis committee, for their comments, which helped to improve the quality of the text of this thesis, and for their suggestions regarding future research.

Thanks to my NA62 colleagues for their time and patience, in particular, Alina, Elisa, and Stefan and the immense work they did on the exotic analyses, Karim, Petre, and Roberto for the help with the NA62 data reconstruction, and to Nicolas for teaching me so much about the RunControl and the NA62 software. Thanks to all my UCLouvain friends and colleagues for making my stay in Belgium full of fun even during the covid times, for teaching me so much about their cultures, cuisines, and languages, and for the “nearly regular” football matches and climbing sessions. I hope we will stay in touch even when the entropy drives us into different corners of the world.

I would like to thank my family for the immense support over the last years despite the distance and my friends for creating an online environment, which made the distances irrelevant even though we were locked up hundreds of kilometers from each other. Foremost, thank you, Eli, for being by my side on this whole journey and for your care and support even during difficult times.

Abstract

Despite its immense predictive power, there are inconsistencies in the Standard Model of particle physics, the most prominent of which are the neutrino masses. The Standard Model also does not address certain phenomena, notably the amount of observed matter and dark matter and the baryon asymmetry of the universe. Still, these phenomena might be explained by Standard Model extensions. Free parameters of the Standard Model might also be derived in these extensions, often predicting new yet unobserved particles almost decoupled from the Standard Model, which can be a component of the dark matter, the so-called hidden sector of particle physics. Fixed-target experiments are among the most promising probes of the hidden sector, thanks to high-intensity beams.

This thesis classifies the hidden sector portals and the various searches for the corresponding processes in the MeV-GeV mass range of the hidden sector mediators. A public software framework for interpreting experimental results in terms of various hidden sector scenarios has been developed during the project. Part of the work is dedicated to the hidden sector searches with the NA62 experiment and the recent results obtained with the data sample collected in the beam-dump mode in 2021. The data analyses performed on this sample are discussed in detail.

Associated publications:

- [1] J. Jerhot, B. Döbrich, F. Ertas, F. Kahlhoefer and T. Spadaro, ALPIN-IST: Axion-Like Particles In Numerous Interactions Simulated and Tabulated. *JHEP*, vol. 07, p. 094, 2022, arXiv:2201.05170 [hep-ph]
- [2] Y. Afik, B. Döbrich, J. Jerhot, Y. Soreq and K. Tobioka, Probing Long-lived Axions at the KOTO Experiment. *Phys. Rev. D*, vol. 108, no. 5, p. 055007. arXiv:2303.01521 [hep-ph]

Contributions to NA62 internal notes:

- [3] B. Döbrich, S. Ghinescu, J. Jerhot, A. Kleimenova, E. Minucci and T. Spadaro, Search for dark photons decaying into a $\mu^+\mu^-$ pair. *NA62-22-07*.
- [4] B. Döbrich, S. Ghinescu, J. Jerhot, A. Kleimenova, E. Minucci and T. Spadaro, Search for dark photons decaying into a e^+e^- pair. *NA62-23-01*.
- [5] B. Döbrich, S. Ghinescu, J. Jerhot, A. Kleimenova, E. Minucci and T. Spadaro, *In preparation*. *NA62-23-xx*.

Contributions to HIKE publications:

- [6] HIKE Collaboration, HIKE, High Intensity Kaon Experiments at the CERN SPS : Letter of Intent. arXiv:2211.16586 [hep-ex]
- [7] HIKE Collaboration, High Intensity Kaon Experiments (HIKE) at the CERN SPS Proposal for Phases 1 and 2. arXiv:2311.08231 [hep-ex]

Thesis Jury

Prof. Fabio Maltoni	President	Université Catholique de Louvain
Prof. Eduardo Cortina Gil	Supervisor	Université Catholique de Louvain
Prof. Christophe Delaere	Secretary	Université Catholique de Louvain
Prof. Marco Drewes	Advisor	Université Catholique de Louvain
Prof. Karol Kampf	Advisor	Charles University
Dr. Babette Döbrich	Member	Max Planck Institute for Physics

Contents

Introduction	1
1 Hidden Sector of Particle Physics	3
1.1 Standard Model of particle physics	4
1.2 Standard Model extensions	14
1.3 Axions	19
1.4 Hidden sector portals	34
2 Search for feebly-interacting particles with fixed-target experiments	43
2.1 FIP production	45
2.2 FIP detection	55
2.3 Fixed-target experiments	66
2.4 Results	74
2.5 Search for exotic particles in kaon decays	82
2.6 Search for axions with KOTO parasitically	86
2.7 Projection for KLEVER	100
3 The NA62 experiment at CERN	103
3.1 NA62 beamline	105
3.2 Upstream region	108
3.3 Downstream region	111
3.4 Trigger and readout	116
3.5 RunControl	124
3.6 NA62 framework	130
4 Search for feebly-interacting particles with NA62 in beam-dump mode	131
4.1 Beam-dump data sample and reconstruction	132
4.2 Exotic signal simulation strategy	143
4.3 Background simulation strategy	144

4.4	Search for di-lepton decay of dark photon	149
4.5	Search for hadronic decays of exotic particles	183
Conclusion		199
Bibliography		201
Appendices		219
A	Treatment of kinematics for the mixing production	221
B	Initial setup of VetoCounter PMTs	225
C	Kaon data analyzers for efficiency studies	229

Introduction

The Standard Model of particle physics (SM) and the Λ CDM model (or what could be called the Standard Model of cosmology) are the two best-tested theories, which together allow us to describe the whole evolution of the universe from about only 10^{-12} s after the Big Bang with immense precision and with just a handful of free parameters. Yet, there are some internal inconsistencies in the SM, such as the amount of matter and anti-matter (baryon asymmetry of the universe), the non-zero mass of neutrinos, and others. A more complete theory should aspire not only to resolve these problems but ideally to do so while reducing the number of free parameters. This reductionist approach drives the research toward associating the additional degrees of freedom in more complete theories. For example, with the amount of cold dark matter (CDM) in the Λ CDM model, which is a free parameter, or with some free parameters of the SM, such as the coupling constants, while unifying the fundamental forces it describes.

The SM extensions often affect the SM processes, which are very precisely measured. The small impact on these well-measured processes suggests that any new physics has to be almost decoupled from the SM at the energies available today to cope with the *non*-observation in interactions with ordinary matter and is referred to as the hypothetical *hidden sector* of particle physics. Putting the typical energy scale of new physics to much higher energies can, on the other hand, affect the cosmology since it alters the evolution of the universe before 10^{-12} s after the Big Bang, possibly during or before the inflation period about which very little is known. This thesis aims to chart the hidden sector with mediators in the $\mathcal{O}(1 \text{ MeV})$ - $\mathcal{O}(1 \text{ GeV})$ mass range and suggests an efficient way for the interpretation of the experimental results which led to the development of a public software framework called ALPINIST.

The second aim of the thesis is to search directly for the hidden sector mediators with data obtained with a fixed-target experiment NA62. NA62 is a multipurpose experiment focused on precision tests of the SM but also allows a search for hidden sector mediators in a well-controlled environment since

the properties of the detector and associated systematical uncertainties can be understood with the SM measurements. The search results are then interpreted using the abovementioned ALPINIST framework.

This thesis consists of two parts: phenomenological and experimental. The first part, consisting of chapters 1 and 2 deals with the phenomenology of the hidden sector of particle physics with more focus on axion-like particles, dark photons, and dark scalars and searches for them with fixed-target experiments. This part follows the author's contributions to the publications [1, 2, 6, 7] and more recent work related to these, which has yet to be published. Chapter 1 introduces the topic and heads towards the description of the ALPINIST framework of which the author is the developer. Chapter 2 is dedicated to the results obtained with ALPINIST framework for the various hidden sector models. This chapter also shows the potential of the framework for studying kinematics in more detail, which proved useful for probing the sensitivity of neutral kaon experiments for hidden sector particles in a less standard manner.

The second part of the thesis, consisting of chapters 3 and 4, describes an experimental search for the hidden sector particles with the NA62 experiment. Chapter 3 describes the NA62 experiment. Although not directly aligned with the thesis's main topic, more space in chapter 3 is given to the VetoCounter detector, of which the author is an expert and has been involved in its construction and installation, and to the RunControl system of NA62 of which the author is a convener and developer. In particular, the author's contributions to the RunControl system update since 2020 are discussed. Chapter 4 is then dedicated to analyses of data collected by the NA62 experiment in 2021 in which the author was directly involved [3–5]. The processes studied in the data analyses are decays of hidden sector mediators into final states with two charged particles of the same flavor or lepton number and photons or neutral pseudoscalars decaying into photons. The author has been responsible for the particle identification, efficiency studies, and interpretation of the analyses with leptons in the final state and the whole analysis procedure for the analyses with hadrons in the final state.

Hidden Sector of Particle Physics

1.1	Standard Model of particle physics	4
1.1.1	Electroweak interactions	8
1.1.2	Quantum Chromodynamics	10
1.2	Standard Model extensions	14
1.2.1	Neutrino masses	16
1.2.2	Baryon asymmetry	16
1.2.3	Mass hierarchy	19
1.3	Axions	19
1.3.1	Strong CP problem	20
1.3.2	QCD axions	22
1.3.3	Cosmological implications	25
1.3.4	Axion-like particles	27
1.3.5	Theory of ALPs at and below the EW scale	30
1.3.6	Low-energy theory of ALPs	31
1.4	Hidden sector portals	34
1.4.1	Dark photon	36
1.4.2	Dark scalar	36
1.4.3	Heavy neutral lepton	38
1.4.4	Model-independent approach	39

Fundamental physics has experienced immense success in giving understanding of nature with the study of isolated interactions. Notably, the study of interactions of constituents typical for a given energy scale at ever-increasing

precision for revealing new phenomena, typically induced by new interactions of new constituents at a larger energy scale, therefore by *indirectly* probing these larger energy scales. At the energies of the current most powerful machines (proton colliders), we are *directly* accessing energy scales of the order of $\mathcal{O}(1)$ TeV to probe indirectly scales several orders of magnitude larger.

The current most precise model of interactions of constituents (elementary particles) at the largest directly accessible energies is called the Standard Model of particle physics (SM). Despite its immense prediction power, we are already witnesses to phenomena the SM cannot explain, which drives both the theoretical and the experimental search for a theory at a higher energy scale from which the SM would emerge. Given the SM's success, much effort is focused on extending the SM with additional interactions and constituents.

This chapter provides an overview of the SM in section 1.1, the most commonly cited phenomena it does not explain, and SM extensions aiming to explain them in section 1.2. More space will be given to axions with section 1.3 since the following chapters mainly address the experimental search for these particles. A more general classification of SM extensions, convenient for experimental searches, will be introduced in section 1.4.

1.1 Standard Model of particle physics

The Standard Model is a gauge field theory, a theory formulated in terms of field operators acting on a Hilbert space, which are arranged in such a way that the whole theory is invariant under local transformations forming a certain symmetry group (gauge group). For each group generator, the theory contains a corresponding (space-time) vector field (called gauge field). The SM gauge group is $SU(3)_C \times SU(2)_L \times U(1)_Y$ and therefore the SM has $8 + 3 + 1$ gauge fields. The four $SU(2)_L \times U(1)_Y$ gauge bosons W_i^μ , $i \in \{1, 2, 3\}$ and B^μ are called *electroweak* (EW) gauge bosons. The 8 fields A_a^μ , $a \in \{1, \dots, 8\}$ belonging to the $SU(3)$ gauge group are called *gluons*.

Besides the gauge bosons, the SM has spin $1/2$ fields, called fermions f , which transform under the electroweak $SU(2)_L \times U(1)_Y$ symmetry. There are three generations of fermionic fields, each of which consists of leptons ℓ that transform only under the $SU(2)_L \times U(1)_Y$ symmetry and quarks q that transform also under the $SU(3)_C$ symmetry. The $SU(2)_L \times U(1)_Y$ sector of the SM is typically referred to as the EW theory (EWT), while the $SU(3)_C$ is referred to as the quantum chromodynamics (QCD). Each generation of fermions (and anti-fermions) has two types of (anti-)quarks and of (anti-)leptons, where

one type has an electric charge by one unit larger (smaller) from the second type. The three quarks with charge $Q = 2/3$ are called *up*-type quarks, and those with $Q = -1/3$ *down*-type quarks, simply by the first generation quarks, which are called u and d . Of the two types of leptons, one is called neutrinos, labeled ν , and they do not have an electric charge $Q = 0$ while the corresponding charged leptons have $Q = -1$. The first generation leptons are called an electron e and its neutrino ν_e . Altogether, there are six flavors of quarks and six flavors of leptons.

Conserved quantities in the SM

As stated by Noether's theorem [8], there is a conserved quantity for each of these symmetries. The conserved charge for $SU(2)_L$ in the EWT is called a *weak isospin*, labeled T , and allows interactions of fermions that have a third component of weak isospin $T_3 = \pm 1/2$, with the $T_3 = \mp 1/2$ fermions via the $T_3 = \pm 1$ charged gauge bosons, which are the electrically charged W bosons. The weak isospin sign convention is such that the sign of T_3 is the sign of the electric charge. Therefore *down*-type quarks and charged leptons have $T_3 = -1/2$ while *up*-type quarks and neutrinos have $T_3 = +1/2$. All $T_3 \neq 0$ fermions have negative chirality while all $T_3 \neq 0$ anti-fermions have positive chirality. Since right-handed fermions and left-handed anti-fermions have $T_3 = 0$, they cannot undergo charged weak interactions. The fact that EWT is chiral led to the conventional labeling of EW gauge symmetry with the following subscripts $SU(2)_L \times U(1)_Y$. The conserved charge of $U(1)_Y$ of the EWT is called the *weak hypercharge*, labeled Y , and together with T_3 defines the electric charge as $Q = T_3 + Y/2$. Since for electroweak gauge bosons $Y = 0$, they cannot directly interact with each other; in other words, they are not affected by the interaction they mediate. Moreover, an accidental global $U(1)_V$ symmetry of EWT is responsible for another conserved charge, called lepton number¹ L , where leptons carry $L = 1$ and antileptons $L = -1$. Several more details on the interactions between fermions will be given in the following subsection 1.1.1.

For QCD, the conserved charge is called *color*, and both quarks and gluons are color-charged. While the eight gluons transform under the adjoint representation of $SU(3)_C$, there are three colors for the quarks and three anti-colors for the antiquarks (fundamental representation of $SU(3)_C$). The self-interaction of gluons leads to color confinement, where only color-neutral

¹Not to confuse with the left-handedness also often labeled with subscript L .

bound states (hadrons) can be observed [9]. Hadrons can be either formed by a quark and antiquark pair in a state called *meson* (for example, pions and kaons) or by three quarks forming a *baryon* (for example, nucleons such as proton and neutron) or three antiquarks forming an *antibaryon*. Similarly to EWT, an accidental global $U(1)_V$ symmetry of QCD is responsible for an additional conserved charge, called baryon (or quark) number B , where quarks carry $B = 1/3$ and antiquarks $B = -1/3$. Baryon number conservation governs the interactions of hadrons [10–12]. More details on the concepts related to QCD and its bound states, which will be helpful for the following chapters of this thesis, will be summarized in the subsection 1.1.2.

BEH mechanism

The SM also has one complex scalar field H transforming under the electroweak $SU(2)_L \times U(1)_Y$ symmetry, so-called Higgs doublet. The imaginary mass μ of this complex scalar doublet is responsible for a spontaneous symmetry breaking (SSB) for three out of four generators of the electroweak symmetry at the energy of about $v = 246$ GeV. Below this energy, a real massive scalar h appears in the spectrum, also called Higgs boson, and the electroweak gauge boson fields mix to give rise to three *weak* vector bosons (W^\pm and Z^0) [13, 14] with masses proportional to v : $m_W = vg/2$, $m_Z = v\sqrt{g^2 + g'^2}/2$ where g and g' are the weak isospin and weak hypercharge coupling constants. This mechanism for gauge bosons acquiring mass from a symmetry breaking is called Higgs or BEH after Brout-Englert-Higgs.

After this $SU(2)_L \times U(1)_Y \rightarrow U(1)_Q$ electroweak symmetry breaking (EWSB) and subsequent electroweak boson mixing

$$\begin{pmatrix} \gamma \\ Z_0 \end{pmatrix} = R(\theta_W) \begin{pmatrix} B_0 \\ W_0 \end{pmatrix}, \quad (1.1)$$

the emergent $U(1)_Q$ symmetry is the one corresponding to the conserved electric charge² Q of the electromagnetism (EM) with gauge boson γ called the photon. The mass of the h boson itself, or in other words, the magnitude of its self-interactions, remains a free parameter of the SM and requires some other mechanism beyond the SM (BSM) to fix the value. Moreover, the measurement of h self-interactions has not yet been performed. Therefore, it may differ from the value fixed by the h mass m_h . In such a case, one could get a hint of

²The electric charge is often expressed in units of electron charges $e = g \sin \theta_W$ where $\sin^2 \theta_W = 0.231$ [15] is the so-called Weinberg or weak angle.

the mechanism that generates the Higgs potential of the form

$$V_H = \mu^2 H^\dagger H + \lambda (H^\dagger H)^2, \quad (1.2)$$

where the free parameters of the model $\mu^2 < 0$ and $\lambda > 0$ determine observed $v = |\mu|/\sqrt{\lambda}$ and $m_h = \sqrt{2}|\mu|$.

Besides the weak bosons, the fermionic fields also acquire mass through the BEH mechanism. In this case, through a Yukawa-type coupling which connects the weak isospin eigenstates with the Higgs doublet and which below the EWSB scale gives rise to interactions of fermions with the h boson and a mass term for the fermions, both equally proportional to the coupling y_f [16]

$$\mathcal{L}_{\text{Yukawa}} = - \sum_f \frac{y_f}{\sqrt{2}} (v + h) \bar{f} f. \quad (1.3)$$

Renormalization

Since the SM is a quantum field theory, the scattering amplitudes encoding the probabilities of the various processes happening (involving in general n particles) are expressed in terms of perturbation expansion of n -point correlation functions $\langle \varphi(x_1) \dots \varphi(x_n) \rangle$, vacuum expectation values (VEV) of time-ordered products of field operators φ , in the various couplings. Since typically only a finite number of terms of the perturbation series are considered in the calculation of the scattering amplitude, the scale invariance requires that the size of the contribution of these terms has to be determined at each scale (a method called renormalization). The physical (renormalized) field $\varphi(\mu)$ at the given scale and the bare fields defining the theory φ_0 are related through the field renormalization Z as $\varphi(\mu) = Z^{-1/2}(\mu)\varphi_0$. The dependence of the strength of the quantum field operators contributing to the given term of the expansion series is captured in the redefinition of the couplings in which the series is expanded. The relation between the way the fields behave at a given energy scale and the evolution of the coupling g with the scale μ is given by the Callan-Symanzik equation (a special case of the more general Renormalization Group (RG) equation)

$$\left[\partial_\mu \frac{\partial}{\partial \mu} + \beta(g) \frac{\partial}{\partial g} + n\gamma(g) \right] \langle \varphi(x_1) \dots \varphi(x_n) \rangle = 0, \quad (1.4)$$

where the β -function characterizes the evolution of g with μ

$$\beta(g) = \mu \frac{\partial g}{\partial \mu}, \quad (1.5)$$

while the γ -function, also called the anomalous dimension, captures the evolution of the field renormalization

$$\gamma = \frac{1}{2} Z^{-1} \mu \frac{\partial Z}{\partial \mu}. \quad (1.6)$$

The sign of the β -function determines whether the value of the coupling grows or decreases with μ . Some β -functions change sign at the so-called infrared fixed point during the evolution while the coupling stays finite. Often, however, when $\beta(g) \rightarrow 0$ also $g \rightarrow \infty$ at so-called *Landau pole*, which is a clear indication that the perturbation expansion is not valid near these scales. While typically the Landau pole lies at very high energy scales, well beyond energies at which measurements can be made, QCD is an example of a theory with a negative β function and Landau pole located at low energies below 1 GeV. This indicates that QCD does not provide a good description at these energies and indeed, the experimental evidence is that quarks and gluons become confined within their bound states as the energy decreases before reaching the Landau pole. On the contrary, as the energy increases, the strong coupling g_s decreases rapidly; this phenomenon to which the color-charged particles are subject is called the *asymptotic freedom*.

1.1.1 Electroweak interactions

In EWSB, besides generating the fermion masses, the BEH mechanism is also responsible for the flavor mixing, as in order to obtain the mass terms for fermions (physical states) of specific flavors, one has to transform the original 3×3 matrices coupling the left- and right-handed fields in the Yukawa terms to the new basis. The two unitary matrices needed to perform this diagonalization can be combined into a single matrix (also unitary), which shows the relation between the weak eigenstates and mass eigenstates. For quarks, this matrix is called V_{CKM} (after Cabibbo-Kobayashi-Maskawa), and by convention, it is chosen to act on the *down*-type quarks. The elements of V_{CKM} are free parameters of the SM, and their size has to be determined experimentally. The current best estimations of the absolute values obtained from the global fit are the following:

$$|V_{\text{CKM}}| = \begin{pmatrix} |V_{ud}| & |V_{us}| & |V_{ub}| \\ |V_{cd}| & |V_{cs}| & |V_{cb}| \\ |V_{td}| & |V_{ts}| & |V_{tb}| \end{pmatrix} = \begin{pmatrix} 0.97435 & 0.2250 & 0.00369 \\ 0.2249 & 0.97349 & 0.0418 \\ 0.00857 & 0.0411 & 0.99912 \end{pmatrix}. \quad (1.7)$$

The corresponding matrix in the leptonic sector is called U_{PMNS} (after Pontecorvo–Maki–Nakagawa–Sakata) and, by convention, is chosen to act on neutrinos. Therefore, these two matrices encode the mixing of *down*-type quarks and neutrinos in weak interactions.

As shown in (1.3), the BEH mechanism connects the fermion fields of both chiralities³ (f_R, f_L) to generate the corresponding mass terms through interactions with the Higgs doublet, resulting in

$$\mathcal{L}_{\text{mass},f} = - \sum_{i \in \{u,d,\ell,\nu\}} m_i \bar{f}_i f_i = -\bar{f} \mathcal{M} f = -\bar{f}_L \mathcal{M} f_R - \bar{f}_R \mathcal{M}^\dagger q_R, \quad (1.8)$$

where the fermion mass is related to the Yukawa coupling as $m_f = y_f v / \sqrt{2}$.

However, only left-handed neutrinos and right-handed antineutrinos are observed in nature; therefore, they can be an exception from the BEH mechanism. Since neutrinos mix, therefore, they are massive even though their mass is tiny compared to other fermions, and exact values have not been measured yet, there has to be some other BSM mechanism that generates their masses to explain their chirality. As neutrinos are electrically neutral, a natural explanation for the observed chiralities could be that they are Majorana rather than Dirac particles, meaning they are their own antiparticles. In such case, the lepton number L would be violated, giving rise to lepton number violating (LNV) processes like, for example, neutrinoless double β decay. No such processes have been, however, observed. Another commonly quoted solution could be that the right-handed neutrinos and left-handed antineutrinos are much heavier than their counterparts. More details on these SM extensions and their consequences will be given in section 1.2.

Symmetries of EWT

Out of the three discrete symmetries P , C and T , corresponding to parity (space reflection), charge conjugation (equality of particles and antiparticles) and time reversal (reversing in- and out-going states in the process) and their

³The chiral projections $f_{R,L}$ are related to f via projection operators $f_{R,L} = P_{R,L} f = (1 \pm \gamma_5) f / 2$.

combinations, only the combination of all three, CPT , is fundamental since the CPT theorem relates the CPT invariance to the Lorentz invariance. It was observed in β -decay of Co^{60} [17] that electroweak interactions maximally violate C and P . This is reflected in the construction of EWT where only (right-)left-handed (anti-)fermions undergo weak interactions

$$\mathcal{L}_{\text{EWT},f} = g\bar{f}_L\gamma^\mu(1 - \gamma_5)f_L W_\mu^i, \quad (1.9)$$

while, for example, electromagnetic force acts on both helicities equally

$$\mathcal{L}_{\text{EM},f} = eQ_f\bar{f}\gamma^\mu f A_\mu. \quad (1.10)$$

The combined CP symmetry, and via CPT theorem also the T symmetry, are also explicitly violated in weak interactions although much less strongly. This is possible thanks to flavor mixing in weak interactions where the V_{CKM} and U_{PMNS} matrices can have a complex phase. CP violation has been experimentally confirmed for the quark sector, while the CP violating phase in the neutrino sector has not been measured yet. If neutrinos were Majorana rather than Dirac particles, the U_{PMNS} matrix would have three CP violating phases instead of one.

1.1.2 Quantum Chromodynamics

QCD is the theory of quarks and their interactions via gluons, gauge bosons of the $\text{SU}(3)_C$ gauge symmetry. As mentioned above, below the EWSB scale, quarks become massive, and unlike in the neutrino sector, the convention is such that the name of the quarks (their flavor) is identified with the mass eigenstates. Therefore, one can express the QCD Lagrangian in the following form

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4}G^{a,\mu\nu}G_{\mu\nu}^a + \bar{q}_f^i(i\gamma^\mu D_\mu - m_f\mathbb{1})_{ij}q_f^j \quad (1.11)$$

where the sum goes over all indices (quark flavors, gluon and quark colors, and space-time indices). The gluon field operators are present in the gluon field strength tensors $G_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a - g_s f^{abc}A_\mu^b A_\nu^c$ and the covariant derivative $(D_\mu)_{ij} = \partial_\mu\delta_{ij} - ig_s(T_a)_{ij}A_\mu^a$ with T and f being the generators of the $\text{SU}(3)_C$ in the fundamental and respectively adjoint representations. The terms proportional to the strong coupling g_s give rise to the quark-gluon and gluon-gluon interactions.

Symmetries of QCD

Unlike in EWT, the QCD Lagrangian has both P and C and combined CP discrete symmetries. Omitting the mass term (1.8), the QCD Lagrangian (1.11) would also have a large global *chiral* $U(N_f)_R \times U(N_f)_L$ symmetry or equivalently rewritten in a different representation as $U(N_f)_V \times U(N_f)_A$ with vectorial ($V = R + L$) and axial $A = R - L$ parts. Moreover, one can explicitly extract the $U(1)$ phases from $q_{R,L}$ fields and decompose $U(N_f)_V \times U(N_f)_A \rightarrow SU(N_f)_V \times SU(N_f)_A \times U(1)_V \times U(1)_A$. While the $SU(N_f)_V \times SU(N_f)_A \times U(1)_A$ symmetry is explicitly broken by quark masses shown in table 1.1, the $U(1)_V$ symmetry of QCD is not, and it is associated with the baryon number B . It turns out that B can be violated within the SM only via chiral anomaly of EWT, [18, 19] same as anomalous breaking of L . The violation of B and L will be elaborated more in the section for SM extensions 1.2 since it is a common prediction of BSM theories.

	1. gen.		2. gen.		3. gen.	
q_f	u	d	c	s	t	b
m_f	2.2 MeV	4.7 MeV	1.28 GeV	96 MeV	173.1 GeV	4.18 GeV

Table 1.1: The measured values of quark masses m_f [20].

As it is apparent from table 1.1, the u , d and potentially also s quarks have relatively small masses compared to the energy scale where QCD becomes non-perturbative.⁴ Therefore, one can study whether at least approximate $SU(2)_V \times SU(2)_A \times U(1)_A$ or $SU(3)_V \times SU(3)_A \times U(1)_A$ symmetries are not manifest in the nature. It turns out that indeed multiplets of hadrons are observed interacting in a way favoring conservation of a quantum number historically called *isospin* corresponding to the vectorial $SU(2)_V$ symmetry (or extended to the so-called *Eightfold way* corresponding to the $SU(3)_V$ symmetry) and the small variations from the conservation can be calculated perturbatively in theories like Chiral Perturbation Theory (χ PT) which study the low-energy dynamics of QCD assuming the P and C and approximate chiral symmetries.

⁴The scale of SSB of chiral symmetry is about $4\pi f_\pi \approx 1161$ MeV, where $f_\pi = 92.42$ MeV is the pion decay constant.

Chiral perturbation theory

Due to the quark confinement, the axial $SU(2)_A$ (or $SU(3)_A$) symmetry is spontaneously broken at the low-energy limit by the quark pair condensate characterized by a non-zero VEV $\langle \bar{q}q \rangle = \langle \bar{q}_L q_R + \bar{q}_R q_L \rangle \neq 0$, where $q^T = (u, d, (s))$ [21]. One should expect the resulting Nambu-Goldstone bosons to have a non-zero mass due to the small explicit breaking of the axial symmetry by quark masses⁵ and indeed, a triplet of pions π (corresponding to SSB of $SU(2)_A$) is observed or eventually also kaons K and the η meson (SSB of $SU(3)_A$) which have a much smaller mass than other bound states composed of the same quark flavors (e.g. ρ meson with $m_\rho = 770$ MeV). Nevertheless, a bound state (η meson for $U(2)_A$ or η' meson for $U(3)_A$) of comparable mass that could be associated with the SSB of the $U(1)_A$ ($m_{\eta'} \gg m_{\pi^0}$ and $m_{\eta'} \gg m_\eta$) is not observed. It was later realized that $U(1)_A$ is also broken by a chiral anomaly, this time of QCD, which is related to the CP violating θ angle of QCD. This topic will be discussed more in section 1.3 in relation to *axions*.

In order to construct the lowest order effective Lagrangian describing the dynamics of pNGBs presented above, $SU(3)$ matrices $U(x) = \exp(i\Phi(x)/f_\pi)$ are defined with

$$\Phi = \Phi_a \lambda_a \begin{pmatrix} \pi^0 + \frac{1}{\sqrt{3}}\eta & \sqrt{2}\pi^+ & \sqrt{2}K^+ \\ \sqrt{2}\pi^- & -\pi^0 + \frac{1}{\sqrt{3}}\eta & \sqrt{2}K^0 \\ \sqrt{2}K^- & \sqrt{2}\bar{K}^0 & -\frac{2}{\sqrt{3}}\eta \end{pmatrix} \quad (1.12)$$

where λ_a are Gell-Mann matrices and Φ_a are Cartesian components forming the individual meson fields [22]. As U transforms under the chiral transform as $U(x) \rightarrow RU(x)L^\dagger$, one can express an effective Lagrangian invariant under the global $SU(3)_R \times SU(3)_L$ with a minimal number of derivatives as

$$\mathcal{L}_{\text{eff}} = \frac{f_\pi^2}{4} \text{Tr} \left(\partial_\mu U \partial^\mu U^\dagger \right), \quad (1.13)$$

a model describing the meson interactions as interactions of NGB with the non-linear realization of chiral symmetry also referred to as the non-linear sigma model [23]. By expanding $U(x)$ in $1/f_\pi$ one obtains in (1.13) the standard kinetic terms for each meson as well as a series of interaction terms among the mesons. An interesting feature in theories with shift symmetry (which is

⁵In such case of relatively small explicit symmetry breaking by an independent mechanism (EWSB in case of quarks) than the SSB mechanism (condensation in case of quarks) the resulting massive modes are also referred to as *pseudo-Nambu-Goldstone* bosons (pNGB).

explicit in (1.13) thanks to the derivative coupling) is that the transition amplitude for processes in which a particle belonging to the field exhibiting the shift symmetry is radiated is zero in momentum soft limit $p \rightarrow 0$. This behavior led to the postulation of numerous *soft pion theorems* [24–26] but it remains a generic feature of derivatively-coupled NGBs and is referred to as Adler’s zero condition [27, 28] and allows construction of the Lagrangian purely from the knowledge of the soft behaviour and of the unbroken symmetry group [29–31].

When the $SU(3)_R \times SU(3)_L$ is explicitly broken by the quark mass term $\mathcal{M} = \text{diag}(m_u, m_d, m_s)$, mesons acquire finite mass. The lowest order effective Lagrangian invariant also under $\mathcal{M} \rightarrow R\mathcal{M}L^\dagger$ analogically to U is

$$\mathcal{L}_{\text{eff}} = \frac{f_\pi^2}{4} \text{Tr} \left(\partial_\mu U \partial^\mu U^\dagger \right) + \frac{f_\pi^2 B_0}{2} \text{Tr} \left(\mathcal{M} U^\dagger + U \mathcal{M}^\dagger \right), \quad (1.14)$$

where B_0 is related to the scalar singlet quark condensate VEV as $\langle \bar{q}q \rangle = -3f_\pi^2 B_0$. By expanding the mass term in U , already at second order $\mathcal{L}_2 = -(B_0/2) \text{Tr}(\Phi^2 \mathcal{M})$ the meson masses can be determined at first order as

$$\begin{aligned} \text{Tr}(\Phi^2 \mathcal{M}) &= 2(m_u + m_d)\pi^+ \pi^- + 2(m_u + m_s)K^+ K^- \\ &\quad + (m_u + m_d)\pi^0 \pi^0 + 2(m_d + m_s)K^0 \bar{K}^0 \\ &\quad + \frac{2}{\sqrt{3}}(m_u - m_d)\pi^0 \eta + \frac{m_u + m_d + 4m_s}{3}\eta^2. \end{aligned} \quad (1.15)$$

In isospin-symmetric limit $\hat{m} = m_u = m_d$ where there is no $\pi^0 - \eta$ mixing, one obtains $m_\pi^2 = 2B_0 \hat{m}$, $m_K^2 = 2B_0(\hat{m} + m_s)$ and $m_\eta^2 = (2/3)B_0(\hat{m} + 2m_s)$ resulting in values remarkably close to those experimentally measured.

Following the construction in Ref. [22], to obtain the χ PT Lagrangian, the QCD Lagrangian $\mathcal{L} = \mathcal{L}_{\text{QCD}} + \mathcal{L}_{\text{ext}}$ is extended by terms involving general external sources (scalar s , pseudoscalar p , vector v^μ and axial-vector a^μ) that comply with the QCD symmetries

$$\mathcal{L}_{\text{ext}} = \bar{q} \gamma_\mu (v^\mu + \gamma_5 a^\mu) q - \bar{q} (s - i\gamma_5 p) q. \quad (1.16)$$

It is clear from the construction that while the vectorial sources do not mix chiralities, the scalar ones do.⁶ It is therefore convenient to introduce also the respective combinations $r^\mu = v^\mu + a^\mu$, $l^\mu = v^\mu - a^\mu$ and $\chi = 2B_0(s + ip)$. Now the χ PT Lagrangian can be constructed at the lowest order with a local

⁶An example of vectorial external source can be electromagnetic four-vector potential $v_\mu = -eA_\mu Q$ with quark charge matrix $Q = \text{diag}(2/3, -1/3, -1/3)$: $\mathcal{L}_{\text{ext}} = -eA_\mu \bar{q} Q \gamma_\mu q$. The quark mass matrix \mathcal{M} , on the other hand, can be expressed as a scalar source.

$SU(3)_R \times SU(3)_L$ symmetry ($U(x)$ transforming as $U(x) \rightarrow V_R(x)U(x)V_L^\dagger(x)$) as

$$\mathcal{L}_{\chi\text{PT}} = \frac{f_\pi^2}{4} \text{Tr} \left(D_\mu U (D^\mu U)^\dagger \right) + \frac{f_\pi^2}{4} \left(\chi U^\dagger + U \chi^\dagger \right), \quad (1.17)$$

where the covariant derivative $D_\mu A \equiv \partial_\mu A - ir_\mu A + il_\mu A$ transforms as $D_\mu A \rightarrow V_R(D_\mu A)V_L^\dagger$. The transition amplitudes for processes involving light pseudoscalar mesons can be computed by introducing additional interactions as external sources (e.g. decay $\pi^+ \rightarrow \mu^+ + \nu_\mu$ by introducing weak interaction as $r_\mu = 0, l_\mu \sim W_\mu$).

An effective Lagrangian for baryons can be constructed in analogy to (1.17).

1.2 Standard Model extensions

The review of the SM presented in the previous sections has pointed out several inconsistencies with the observations, the neutrino masses being the most prominent one. Another point is the mass of the h boson itself, the only dimensionful free parameter of the SM, whose value is susceptible to virtual corrections, which would put the h mass to a much larger energy scale than the one actually measured. The absence of the CP-violating term in QCD is another unexplained feature, which will be addressed in section 1.3. Further, the measurements of the anomalous magnetic dipole moment of the muon result in a 4.2σ tension with the current SM prediction [32]. While this tension is insufficient to consider it an inconsistency in the standard measures (5σ -rule), it may also hint at a possible New Physics contribution to the $(g-2)_\mu$ value.

Other hints for the possible SM extensions come from cosmology. Remarkably, the SM can explain rather well the evolution of the early universe as it is understood today. The three exceptions are the early exponential inflation (in between about the first 10^{-36} and 10^{-33} s after the Big Bang) giving rise to the large structures of the universe, the matter-antimatter asymmetry and the height of the peaks of the oscillations in the Cosmic Microwave Background (CMB) (and in the Baryonic Acoustic Oscillations (BAO) in today's universe). Today's universe, with dominant interactions at the macroscopic level, is thought to be ruled by the General Relativity (GR). If that is the case, several other observations show inconsistencies, among them the kinematics of galaxies and galaxy clusters and the expected amount of matter in them causing the observed gravitational lensing. A popular solution that addresses multiple of these inconsistencies is the postulation of the existence of Dark

Matter (DM), a form of matter that does not interact with the SM or only very weakly. Moreover, observations require that DM is non-relativistic (also called *cold*) for large-scale structures to form. The BAO and CMB oscillations measurements, gravitational lensing, and measurements of the kinematics and evolution of galaxies suggest that cold DM (CDM) is about five times more abundant than the baryonic matter (dominated by protons and neutrons). While radiation (photons and neutrinos) were abundant shortly after the Big Bang, and neutrinos as a DM component would be in agreement with the requirement of feeble interactions with the SM, they would present a relativistic component of DM and, therefore, cannot explain the observations. Another cosmological observation is that the current universe is expanding and that the expansion rate is accelerating. This is reflected in the GR with the positive cosmological constant Λ , associated with a positive vacuum energy density. The current cosmological model, based on the assumptions of CDM and $\Lambda > 0$, is called Λ CDM, and it has six free parameters. While Λ CDM is largely successful, there are tensions between different types of measurements, possibly hinting at underlying mechanisms for how the currently empirical values for DM and Λ are produced. A tension concerning the measurement of Λ , or correspondingly of the Hubble constant H_0 , is called the Hubble tension, and it is a disagreement at a level more than 4σ between the H_0 value obtained from Supernovae Ia and Cepheids measurements and the H_0 value obtained from CMB measurements [33]. Another tension concerns the CDM hypothesis and the behavior of dwarf galaxies. Interestingly, both CDM and Λ can be related to particle physics. If CDM exists and it interacts with the SM, even though very weakly, it can be searched for as a new particle species, possibly also solving one of the problems of the SM. Λ can be associated with the vacuum energy generated by the SM fields. However, in this case, the SM prediction is in disagreement with the cosmological observation by tens to hundreds of orders of magnitude, requiring some additional extremely large contribution of negative vacuum energy almost precisely canceling the SM one, a fine-tuning problem called Cosmological Constant problem to which not many possible solutions are known so far.

Some of the common SM problems will be addressed in the following sections, and a unified model-independent parametrization useful for experimental BSM physics searches will be shown in section 1.4.

1.2.1 Neutrino masses

One of the most compelling solutions for generating the neutrino masses is the so-called Seesaw mechanism, in which the smallness of the SM neutrino masses is implied by existence of some very large mass scale.

The simplest version of a Seesaw mechanism is the Type I seesaw [34, 35], which predicts a right-handed neutrino ν_R $SU(2)_L$ singlet, also called heavy neutral lepton (HNL), with Majorana mass M_R for each SM neutrino flavor and interacts with the SM lepton left-handed and Higgs doublets via Yukawa coupling y_ν in a dimension five operator. The SM neutrino masses come from the diagonalization of the neutrino mass matrix with the Dirac masses m_D obtained from the BEH mechanism as for charged leptons and the Majorana masses of the heavy neutrinos M_R . For $m_D \ll M_R$ the HNL mass eigenstates still have mass $\approx M_R$ while the SM left-handed neutrinos have mass $m_\nu \approx m_D^2/M_R$. While this requires $M_R \approx 10^{14}$ GeV, it can be effectively lowered by tuning the strength of the Yukawa coupling y_ν in the non-minimal scenarios. Alternatively, small m_ν can be obtained from large mass matrix entries if the contributions approximately cancel out in the diagonalization without requiring small y_ν , so-called *low-scale seesaw*.⁷

The Type II seesaw mechanism introduces a new heavy scalar triplet of $SU(2)_L$, which mixes with the Higgs doublet and below the EWSB scale leads to a similar result as Type I seesaw [36]. The Type III seesaw introduces a heavy fermion triplet of $SU(2)_L$, which is also electrically charged and results in a mass matrix similar to Type I seesaw [37].

In practice, many models predicting neutrino masses are variations or combinations of the types above. A popular low-scale seesaw model that can be tested at energies achievable at present accelerators is the neutrino minimal SM (ν MSM), which requires two to three heavy right-handed neutrinos and a scalar singlet of masses below the EWSB scale, which could also present a favorable DM candidate [38].

1.2.2 Baryon asymmetry

The Big Bang nucleosynthesis (BBN), the production of primordial elements from primordial nucleons, can be understood from the SM [39]. Moreover, the observations are in agreement with the predictions [40]. However, it comes

⁷In this scenario, $B - L$ conservation is violated which is not the case in the SM where the chiral anomalies cancel out, as will be shown in the next section 1.2.2. Instead, an equivalent $B - \bar{L}$, where \bar{L} also incorporates the new leptons, would be conserved.

with a requirement of an abundance $\eta = (n_B - n_{\bar{B}})/n_\gamma > 0$ of primordial nucleons n_B over antinucleons $n_{\bar{B}}$ in order to explain the measured⁸ baryon asymmetry of the universe of the order of $\eta \simeq 10^{-10}$, an observation which cannot be explained by the SM and Λ CDM. The baryon asymmetry can be generated from an initial balance between matter and antimatter through baryogenesis in the earlier stages of the universe. A theory aiming to allow baryogenesis and produce the required density of baryons has to satisfy the following set of conditions (called Sakharov conditions) [42]

- B violation;
- C and CP violation;
- Out of thermal equilibrium interactions.

The first two conditions are clearly necessary for processes producing an excess of baryons over antibaryons since, first, a process that generates baryons ($\Delta B \neq 0$) is needed, but secondly, it has to violate C in order to prefer matter over antimatter and similarly CP to ensure the same for given chiralities. The out-of-equilibrium condition is necessary so that the rate of the baryon production processes is not counterbalanced by the equivalent ‘destruction’ processes.

As mentioned in 1.1, the chiral anomaly violates B in the SM with the non-vanishing derivative of the corresponding current given by

$$\partial_\mu j_B^\mu = n \left(\frac{g^2}{32\pi^2} W_{\mu\nu} \tilde{W}^{\mu\nu} - \frac{g'^2}{32\pi^2} B_{\mu\nu} \tilde{B}^{\mu\nu} \right), \quad (1.18)$$

with $\tilde{F}^{\mu\nu} = \varepsilon^{\mu\nu\rho\sigma} F_{\rho\sigma}/2$ dual field strength tensor to $F^{\mu\nu}$ and n number of generations. The existence of anomalous currents of this form is a general feature of gauge theories when coupled to fermions since they correspond to the chiral transformation of the corresponding fermionic fields

$$\psi \rightarrow \exp(-i\alpha\gamma_5)\psi. \quad (1.19)$$

One can verify that such terms correspond to the total derivative and hence do not contribute to the action if the gauge fields vanish at infinity, which is, for

⁸Besides the observed baryon abundance in the current universe, the fact that matter of universe is baryonic while the antibaryonic matter is missing can be also verified from the large structure data, diffuse gamma ray background, or for the early universe with CMB acoustic oscillation measurements [41].

example, the case of Maxwell's theory.⁹ Nevertheless, non-Abelian theories have a non-trivial topological structure of the vacuum with infinitely many periodically repeating minima, thus one meets non-vanishing field configurations in the spatial infinity, called *instantons* [44, 45], where in general it holds [46]

$$\int d^4x F_{\mu\nu} \tilde{F}^{\mu\nu} = \frac{32\pi^2}{g^2} \nu, \quad (1.20)$$

where $\nu \in \mathbb{Z}$ is called instanton number (or also winding or second Chern number).

On the contrary, in chiral non-Abelian theories with Yukawa-coupled fermions, a chiral transformation (1.19) does not change the path integral measure as fields of both chiralities are coupled via the Yukawa matrix, but only one chirality couples to the gauge field, thus in such setup these configurations do not contribute to the action as well. This is the case of EWT, which is an example of chiral non-Abelian theory. Unfortunately, this is not the case of QCD and QCD extensions as will be shown in the following section 1.3.

Furthermore, it has been proven in Ref. [47] that while chiral anomaly violates $B + L$, the opposite, $B - L$ is conserved and equation (1.18) also holds for the lepton number $\partial_\mu j_L^\mu$. Specifically for the SM¹⁰ $n = 3$ and thus $\Delta B = \Delta L = 3$ since fermions of all generations are necessary (therefore, the stability of the proton is ensured in the SM). Such processes are significantly suppressed at current energies (and have never been observed) and could be achieved only by tunneling between equivalent electroweak vacua, so-called *instanton-like* processes between two minima distant by $\nu = 1$. The process could be enhanced at energies available in the early universe or even surpassing the energy barrier.¹¹ Thanks to this conservation of $B - L$ in the SM, a surplus of baryons can also be obtained from the surplus of leptons, which can be generated in an independent process called leptogenesis [48].

⁹By introducing non-trivial boundary conditions, the term can contribute to the action even in Abelian theories, leading, for example, to a so-called Topological Casimir effect in electrodynamics [43].

¹⁰ $B - L$ conservation is also a feature of many BSM theories, but the stability of proton is not necessarily ensured. Current measurements put very stringent limits on proton lifetime $\tau_p > 3.6 \times 10^{29}$ yr and specifically for $\Delta B = \Delta L = 1$ processes as a decay to lepton which give $\tau(p \rightarrow e^+ \pi^0) > 1.6 \times 10^{34}$ yr [20]. An example of $\Delta B = 2$ process would be neutron-antineutron oscillations, which also have not been observed.

¹¹The saddle point of the barrier separating the minima (equivalent vacua in this case) is typically referred to as a *sphaleron*. Surpassing the energy barrier would be, therefore, effectively equivalent to producing excitations of these field configurations (sphalerons).

Similarly, as shown in the previous section 1.1, the SM has both C and CP violations in the EW sector. Unfortunately, the amount of CP violation in the SM is not sufficient to explain the measured baryon asymmetry, indicating further need for a BSM theory which is rather restricted since while it has to introduce additional CP violation, it also has to comply with the current strict measurements of CP violation in flavor physics.

The third condition is also not sufficiently satisfied with the SM as in order for the SM to produce a sufficient amount of baryon asymmetry, the EWSB would have to be a first-order phase transition, which, given the measured mass of h boson $m_h = 125.25$ GeV [20], it is not [49].

Many BSM theories can fulfill the above criteria by adding interactions between leptons and quarks or introducing new heavy states. Those can be theories that extend the SM gauge symmetry group, like Grand Unification Theories, or, for example, supersymmetric theories, which predict the existence of a spacetime symmetry between bosons and fermions above a certain energy scale. These have received large attention also in connection with the Mass hierarchy problem.

1.2.3 Mass hierarchy

The Mass hierarchy problem is another fine-tuning problem of the SM that a more complete theory should aim to solve since the h boson mass is a free parameter of the SM. However, once the value is derived, it suffers from contributions to the mass of the size of the energy scale of the new theory. The dominant contribution comes from the h boson self-coupling, which results in a quadratic divergence of the mass. The supersymmetry has been a popular candidate for the solution thanks to the cancellations from the nearly equivalent bosonic and fermionic contributions. Nevertheless, no supersymmetric particle candidates have been found so far and the energy scales at which they are often expected are too large for such theory to be tested.

1.3 Axions

The conserved $U(1)_A$ current in QCD has originally posed a problem unlike the vectorial $U(1)_V$ counterpart associated with the baryon number since it is not an observable symmetry, neither a corresponding pNGB with mass $\sim m_\pi$ [50] has been found. Nevertheless, it has been later shown that as $U(1)_V$, also the $U(1)_A$ symmetry is broken anomalously [47] with the non-vanishing

derivative of the corresponding current

$$\partial_\mu j_5^\mu = \partial_\mu (\bar{q} \gamma^\mu \gamma_5 q) = N_f \frac{g_s^2}{16\pi^2} G_{\mu\nu} \tilde{G}^{\mu\nu} \quad (1.21)$$

with N_f number of flavors. However, since QCD is not a chiral theory, such a term cannot be rotated away by chiral transformation (1.19) as was the case of EWT and an equivalent term, which is in general parametrized by a phase parameter $\theta \in [0, 2\pi)$, should be present in the QCD Lagrangian.

Moreover, in massive QCD, such rotation changes the phase $\arg(\det \mathcal{M})$ of the quark mass matrix \mathcal{M} from equation (1.8) by θ . Therefore, given (1.21), the QCD Lagrangian should contain a term of the form

$$\mathcal{L}_{\text{QCD}} \supset \bar{\theta} \frac{g_s^2}{32\pi^2} G_{\mu\nu} \tilde{G}^{\mu\nu}, \quad (1.22)$$

which breaks the $U(1)_A$ symmetry, where the angle $\bar{\theta}$ is a free parameter of the SM. When one rotates away all the γ_5 -dependency of quark mass matrix \mathcal{M} , the total angle $\bar{\theta}$ can be expressed as $\bar{\theta} = \theta + \arg(\det \mathcal{M})$. As will be shown in the following section, this term, however violates CP symmetry in QCD contrary to what is observed, giving rise to the so-called Strong CP problem.

1.3.1 Strong CP problem

The term (1.22) can be rewritten as a scalar product of vector and axial-vector fields¹² $G_{\mu\nu}^a \tilde{G}^{a,\mu\nu} = 4\vec{E}_\mu^a \vec{B}^{a,\mu}$ hence it conserves C , but it violates P thus also CP .

In massive QCD, by performing a chiral rotation in the quark fields (1.19), the equation (1.21) becomes [22]

$$\partial_\mu (\bar{q} \gamma^\mu \gamma_5 q) = 2i\bar{q} \gamma_5 \mathcal{M} q + N_f \frac{g_s^2}{16\pi^2} G_{\mu\nu} \tilde{G}^{\mu\nu} \quad (1.23)$$

therefore by performing a chiral rotation by $\bar{\theta}$ one can introduce an equivalent term to (1.22) in quark fields.

In the effective theory of pions π^a , $a \in \{1, 2, 3\}$ and nucleons $N^T = (pn)$, this terms creates additional CP violating contribution to the πNN interaction Lagrangian

$$\mathcal{L}_{\pi NN} = \pi^a \bar{N} \tau^a (i\gamma_5 g_{\pi NN} + \bar{g}_{\pi NN}) N, \quad (1.24)$$

¹²Equivalent of electric and magnetic fields in the Maxwell theory.

where τ^a are Pauli matrices, $g_{\pi NN}$ is the standard CP conserving coupling determined by Goldberger-Trieman relation [51] $g_{\pi NN} = g_A m_N / f_\pi$ and $\bar{g}_{\pi NN}$ is the new CP violating coupling [52]

$$\bar{g}_{\pi NN} = \bar{\theta} g_A \frac{\mu}{f_\pi} \frac{m_\Xi - m_N}{2m_s - m_u - m_d}. \quad (1.25)$$

At one-loop order (1.24) together with pion-photon interaction a triangle diagram generates a neutron electric dipole moment¹³ (EDM) [53]

$$d_n = e \frac{\bar{g}_{\pi NN} g_{\pi NN}}{8\pi^2 m_N} \log\left(\frac{m_N^2}{m_\pi^2}\right) \sim 3 \times 10^{-16} \bar{\theta} \cdot e \cdot \text{cm}. \quad (1.26)$$

The most recent measurements of neutron EDM give an upper limit (UL) $|d_n| < 1.8 \times 10^{-26} e \cdot \text{cm}$ [54], therefore $|\bar{\theta}| \lesssim 6 \times 10^{-11}$ meaning that QCD is in fact CP conserving given the current precision. Similarly to other fine-tuning problems mentioned in the previous section, there is no explanation within the SM why QCD should be CP conserving, which in this case became known as the Strong CP problem.¹⁴

It is clear that even one massless quark would allow to rotate the CP violating term away. However, the quark masses are known to a very good precision today, and they all have non-zero values (see table 1.1).

Many of the fine-tuning problems, and notably those mentioned above, are connected to the evolution of the universe in a way that if given parameters had different values, the current universe would look very differently and could not support life in the way we know it, a reasoning called the anthropic principle. On the contrary, introducing non-zero neutron EDM does not pose a problem for the evolution of the universe, so naively, the reasoning above cannot be used.¹⁵

A common BSM solution setting $\bar{\theta} = 0$ is the introduction of an additional global $U(1)$ symmetry to the SM which is spontaneously broken [58]. While the original proposal predicting a new pseudoscalar particle called *axion* [59–61] has been ruled out experimentally, new $U(1)$ symmetries appear frequently

¹³In order to evaluate the size of the contribution numerically, the axial-vector coupling in the chiral limit $g_A = 1.2694$ can be measured from the neutron β decay, $m_N = (m_p + m_n)/2 = 938.92$ MeV and $\mu = m_u m_d / (m_u + m_d) = 1.55$ MeV.

¹⁴Some attempts were made to resolve the Strong CP problem within the SM e.g. [55]. A recent work [56] is suggesting that the θ contribution from the instanton effects can be aligned with the chiral phases of \mathcal{M} effectively canceling $\bar{\theta}$.

¹⁵Nevertheless, for example, Ref. [57] suggests that $\bar{\theta}$ can be linked to the Cosmological Constant problem.

in BSM theories and thus axions and axion-like particles (ALPs) remain a popular extension of the SM.

1.3.2 QCD axions

By introducing a new global chiral $U(1)_{\text{PQ}}$ symmetry to the SM which is spontaneously broken at some scale f_a , a new NGB a an *axion* appears which transforms under $U(1)_{\text{PQ}}$ as $a \rightarrow a + \alpha f_a$. Coupling a to the SM analogous to the θ contribution as

$$\mathcal{L}_{\text{SM+axion}} \supset \bar{\theta} \frac{g_s^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a,\mu\nu} - \frac{1}{2} (\partial_\mu a)(\partial^\mu a) + \xi \frac{a}{f_a} \frac{g_s^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a,\mu\nu}, \quad (1.27)$$

with model-dependent coupling parameter ξ , results in an anomalous current [62]

$$\partial_\mu j^\mu = \xi \frac{g_s^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a,\mu\nu}, \quad (1.28)$$

which, together with the QCD anomaly, generates a periodic effective potential for the axion field $V_{\text{eff}} \sim \cos(\bar{\theta} + \xi \langle a \rangle / f_a)$ with the minimum for the VEV of a being $\langle a \rangle = -\bar{\theta} f_a / \xi$. At the minimum of the effective potential the $\bar{\theta}$ -term is cancelled out

$$\left\langle \frac{\partial V_{\text{eff}}}{\partial a} \right\rangle = -\frac{\xi}{f_a} \frac{g_s^2}{32\pi^2} \left\langle G_{\mu\nu}^a \tilde{G}^{a,\mu\nu} \right\rangle \Big|_{\langle a \rangle = -\bar{\theta} f_a / \xi} = 0. \quad (1.29)$$

Thus, the Strong CP problem is dynamically solved, and the $\bar{\theta}$ -term is absorbed in a physical axion field $a_{\text{phys}} = a - \langle a \rangle$. Now (1.27) can be written as

$$\mathcal{L}_{\text{SM+axion}} \supset -\frac{1}{2} (\partial_\mu a_{\text{phys}})(\partial^\mu a_{\text{phys}}) + \xi \frac{a_{\text{phys}}}{f_a} \frac{g_s^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a,\mu\nu}. \quad (1.30)$$

Based on the review in Ref. [63], several examples of the first axion models will be shown in the following.

Peccei-Quinn-Weinberg-Wilczek (PQWW) model

The original model [58–61] associated f_a to the EWSB scale v by assuming two Higgs doublets

$$\phi_1 = \frac{v_1}{\sqrt{2}} e^{i\frac{a}{v}x} \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \phi_2 = \frac{v_2}{\sqrt{2}} e^{i\frac{a}{v}\frac{1}{x}} \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad (1.31)$$

where $x = v_2/v_1$ and $v = \sqrt{v_1^2 + v_2^2}$. The following Yukawa-type interaction Lagrangian

$$\mathcal{L}_{\text{Yukawa}} = \bar{q}_L \phi_1 u_R + \bar{q}_L \phi_2 d_R + \bar{L}_L \phi_1 l_R + \text{h.c.}, \quad (1.32)$$

is then also invariant under $U(1)_{\text{PQ}}$

$$\begin{aligned} a &\rightarrow a + \alpha v; \\ u_R &\rightarrow e^{-i\alpha x} u_R; \\ d_R &\rightarrow e^{-i\alpha/x} d_R; \\ l_R &\rightarrow e^{-i\alpha/x} l_R. \end{aligned} \quad (1.33)$$

If one omits the lepton part, the associated Noether current will be

$$j_{\text{PQ}}^\mu = v \partial^\mu a + x \bar{u}_R \gamma^\mu u_R + \frac{1}{x} \bar{d}_R \gamma^\mu d_R \quad (1.34)$$

giving the coefficient ξ for this model

$$\xi = \frac{N_f}{2} \left(x + \frac{1}{x} \right). \quad (1.35)$$

In the following will be proceeded using Lagrangian (1.14) but in the limit with two light quarks. The quark operators will be replaced by their condensates $\langle \bar{q}q \rangle$, giving rise to π and η mesons

$$U = \exp \left(i \frac{\tau_i \pi_i + \eta}{f_\pi} \right); \tau_i \pi_i \equiv \begin{pmatrix} \pi^0 & \sqrt{2} \pi^+ \\ \sqrt{2} \pi^- & \pi^0 \end{pmatrix} \quad (1.36)$$

τ_i here are the Pauli matrices.

The U transforms under $U(1)_{\text{PQ}}$ as

$$U \rightarrow U \begin{pmatrix} e^{i\alpha x} & 0 \\ 0 & e^{i\alpha/x} \end{pmatrix} \quad (1.37)$$

and hence only the $U(1)_A$ flavour symmetry is broken by the quark mass term of (1.14), which is symmetric under $U(1)_{\text{PQ}}$ and gives

$$\begin{aligned} \mathcal{L}_{\mathcal{M}} &= \frac{f_\pi^2 B_0^2}{2} \text{Tr} \left(U A \mathcal{M} + (U A \mathcal{M})^\dagger \right), \text{ with} \\ A &= \text{diag} \left(e^{-i \frac{a}{v} x}, e^{-i \frac{a}{v} \frac{1}{x}} \right), \mathcal{M} = \text{diag}(m_u, m_d). \end{aligned} \quad (1.38)$$

The anomalous terms breaking both $U(1)_A$ in equation (1.21) and $U(1)_{\text{PQ}}$ in equations (1.28) are responsible for the correction of the mass of the η meson

as well as giving a mass to our axion and the a - π and a - η mixing

$$\mathcal{L}_{\text{anomaly}} = -\frac{m_\eta^2}{2} \left[\eta + \frac{f_\pi}{v} \frac{1}{2} \left(\frac{N_f}{2} - 1 \right) \left(x + \frac{1}{x} \right) a \right]^2. \quad (1.39)$$

Combining both previous contributions, one gets the axion mass as

$$m_a = \frac{N_f}{2} \left(x + \frac{1}{x} \right) \frac{f_\pi}{v} m_\pi \frac{\sqrt{m_u m_d}}{m_u + m_d} \quad (1.40)$$

and from mixing with π and η also the coupling to two photons $g_{a\gamma\gamma}$ from the EM anomaly

$$\partial_\mu j_{\text{PQ}}^\mu = \frac{e^2}{24\pi^2} N_f \left(x + \frac{1}{x} \right) F_{\mu\nu} \tilde{F}^{\mu\nu} \quad (1.41)$$

$$\mathcal{L}_{a\gamma\gamma} = g_{a\gamma\gamma} a F_{\mu\nu} \tilde{F}^{\mu\nu}; \quad g_{a\gamma\gamma} = \frac{e^2}{32\pi^2} N_f \left(x + \frac{1}{x} \right) \frac{1}{v} \frac{m_u}{m_u + m_d}. \quad (1.42)$$

As this model gives the SSB scale of $U(1)_{\text{PQ}}$ coincident with the electroweak SSB scale $f_a = v$, all the above variables have fixed values. However, these were quickly excluded by experiments, and other models came to the fore as, in principle, the SSB scale might be at much higher energies $f_a \gg v$. Then the coupling would be much weaker, and axions would be very light, so-called *invisible axions*, which was later realized that makes them also DM particle candidates.

Kim-Shifman-Vainshtein-Zakharov (KSVZ) model

There are, in principle, two mechanisms to extend the original PQWW model to shift the f_a at higher energies, which can be collectively called the *hadronic models* and the *non-hadronic models*.

The Kim-Shifman-Vainshtein-Zakharov [64,65] (KSVZ) model was the first hadronic model proposal. It expands the SM with a heavy quark Q and an additional complex Higgs scalar σ both are weak-interaction-singlets and are symmetric under $U(1)_{\text{PQ}}$. The describing Lagrangian can be expressed in a form

$$\mathcal{L}_{\text{KSVZ}} = \bar{Q} i \not{D} Q + \partial^\mu \sigma^\dagger \partial_\mu \sigma - h_\sigma (\bar{Q}_L \sigma Q_R + \text{h.c.}) + m_\sigma^2 \sigma^\dagger \sigma - \lambda_\sigma (\sigma^\dagger \sigma)^2. \quad (1.43)$$

The $U(1)_{\text{PQ}}$ is spontaneously broken by the $f_a = \langle \sigma \rangle = m_\sigma / \sqrt{2\lambda_\sigma} \gg v$, so the quark Q gets a mass $m_Q = h_\sigma f_a$ and the complex field σ in (1.43) gives a rise to a scalar state from its radial part with a mass $m_\sigma \sqrt{2}$ and a pseudoscalar state a from its phase

$$\sigma = f_a \exp\left(i \frac{a}{\sqrt{2} f_a}\right). \quad (1.44)$$

The heavy quark loop interactions give rise to

$$\mathcal{L}_{\text{KSVZ}}^{\text{gauge}} = \frac{a}{f_a} \left(\frac{g_s^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a,\mu\nu} + 3e_Q^2 \frac{e^2}{16\pi^2} F^{\mu\nu} \tilde{F}_{\mu\nu} \right), \quad (1.45)$$

where e_Q is the EM charge of the quark Q . The anomalous mass term is then

$$\mathcal{L}_{\text{anomaly}} = -\frac{m_\eta^2}{2} \left(\eta + \frac{f_\pi}{2f_a} a \right)^2 \quad (1.46)$$

and together with the standard meson mass term (1.15) determine the a - π^0 and a - η mixing, resulting in an effective photon coupling

$$g_{a\gamma\gamma} = \frac{e^2}{32\pi^2} \frac{2}{f_a} \left(3e_Q^2 - \frac{4m_d + m_u}{3(m_u + m_d)} \right). \quad (1.47)$$

Diagonalizing (1.39) and (1.15) give axion mass inversely proportional to f_a which remain a free parameter of the model

$$m_a = m_\pi \frac{f_\pi}{f_a} \frac{\sqrt{m_u m_d}}{m_u + m_d} = 6.3 \text{ eV} \left(\frac{10^6 \text{ GeV}}{f_a} \right). \quad (1.48)$$

The Dine-Fischler-Srednicki-Zhitnisky [66, 67] (DFSZ) model was the first non-hadronic axion model. It adds to the PQWW model (i.e., to the two scalars φ_1 and φ_2) a new scalar φ transforming under $U(1)_{\text{PQ}}$, which $\langle \varphi \rangle = f_a \gg v$. One gets similar results to the KSVZ model, when one rescales $f_a \rightarrow f_a/N_f$. The axion mass is given by the same expression as in the former case and the coupling to photons is given by

$$g_{a\gamma\gamma} = \frac{e^2}{32\pi^2} \frac{2}{f_a} \left(\frac{4}{3} - \frac{4m_d + m_u}{3(m_d + m_u)} \right). \quad (1.49)$$

1.3.3 Cosmological implications

Axions, or axion-like particles (ALPs) in general, which will be presented in the next section, present favorable DM candidates, and thus, one has to con-

sider the corresponding cosmological phenomena, most of which are model-dependent. This presents an extensive topic, reviewed, e.g., in Ref. [68], while only a few generic instances derived for QCD axions will be presented here. The case of hot (relativistic) axions created and annihilated in the primordial soup influences also the amount of observable matter as in the early universe they couple predominantly to gluons. Moreover, the population of hot axions is wiped out by inflation unless they are reheated, and the hot QCD axion becomes less motivated as a DM candidate. Thus, the case of cold axions will be mainly discussed further.

There are two general scenarios differing in the production of cold DM (CDM) axions in the early universe, as discussed in Ref. [69]. In the first case, the $U(1)_{PQ}$ symmetry was spontaneously broken at energy f_a after the inflation and in the second one before or during the inflation.

In the post-inflation scenario, a network of one-dimensional topological defects (cosmic strings) creates and consequently starts disintegrating by the emission of massless axions, which at temperatures below the QCD scale obtain mass by the instanton (anomalous) term (1.39) and the axion potential takes a periodic form

$$V_{\text{eff}}(a) = m_a^2 f_a^2 [1 - \cos(a/f_a)]. \quad (1.50)$$

Since the potential (1.50) has infinitely many minima, the axion field settles down to different minima around the universe as it cools down and gets divided into domains separated by quasi-stable domain walls attached between cosmic strings, creating a string-domain network. The different θ which leave the QCD Lagrangian up to a multiple of $2\pi/N$ are a manifest of explicit anomalous breaking of $U(1)_{PQ}$ symmetry into a discrete subgroup $Z(N)$. The spontaneous breaking of the $Z(N)$ symmetry implies N degenerate vacua, and the further development of the string-domain network varies for $N = 1$ and $N > 1$.

- For $N = 1$, the domain walls are quantum-mechanically unstable and decay into axions, giving a significant contribution to the matter density of the universe.
- For $N > 1$, the domain walls are stable, giving rise to extra acceleration to the universe expansion. One could identify this expansion with the observed one. Nevertheless, such a universe would be less homogeneous than is observed, which became known as the *domain wall problem*.

This problem is solved for $N = 1$ or if the $Z(N)$ symmetry is somehow explicitly broken or in the pre-inflation scenario discussed further.

In the pre-inflation scenario, the inflation expands ‘our’ domain into a larger size than is the present observable universe, and thus, there is no domain wall problem, and the current axion density is given by coherent oscillation of the axion field.

1.3.4 Axion-like particles

In principle, the whole (C_a, m_a) parameter space can be accessed in certain models, as the nature of the axion-SM coupling can be more general. The class of pseudoscalar NGBs from SSB of BSM shift symmetries is typically referred to as axion-like particles since they share a similar nature to the axions solving the strong CP problem (also referred to as QCD axions). In order to solve the strong CP problem, QCD axions have to be very light ($m_a < 1$ keV) to achieve large suppression by f_a . ALPs, on the other hand, allow much larger masses $m_a > 1$ MeV, which are the main aim of this work, and in certain heavy QCD axion models [70–77] also to solve the strong CP problem. On the contrary, the requirement $f_a \lesssim 10$ TeV, which is in reach of current laboratory probes, is supported by the so-called axion quality problem [72].

The most general ALP Lagrangian at the ultraviolet (UV) scale¹⁶ $\Lambda \gg v$ with operators up to dimension 5 is

$$\begin{aligned} \mathcal{L}_{\text{ALP}} = & \frac{1}{2}(\partial^\mu a)(\partial_\mu a) - \frac{m_{a,0}^2}{2}a^2 + \sum_f \frac{C_{ff}}{2\Lambda}(\partial^\mu a)\bar{f}\gamma_\mu\gamma_5 f \\ & + g^2 \frac{C_{WW}}{\Lambda} a W^{\mu\nu} \tilde{W}_{\mu\nu} + g'^2 \frac{C_{BB}}{\Lambda} a B^{\mu\nu} \tilde{B}_{\mu\nu} + g_s^2 \frac{C_{GG}}{\Lambda} a G^{\mu\nu} \tilde{G}_{\mu\nu} , \end{aligned} \quad (1.51)$$

where only flavor-conserving ALP-fermion couplings are considered at the UV scale and flavor-changing couplings $C_{f_i f_j}$ are present at low energy scales as a consequence of RG evolution of the couplings to the EW scale and matching conditions derived in Ref. [78]. The results in Ref. [1] were derived for FCNC ALP couplings for C_{WW} -only coupling at UV using the one-loop calculation from Ref. [79] while the two-loop calculation for C_{GG} -only coupling at UV

¹⁶UV scale stands here for a general high energy scale at which \mathcal{L}_{SM} and \mathcal{L}_{ALP} are matched onto the Lagrangian of the full theory.

scale is from Ref. [80]. By comparing the ALP Lagrangian (1.51) with the QCD axion one (1.30), the scale Λ is related to f_a as $32\pi^2 C_{GG}/\Lambda = -1/f_a$ and the physical ALP mass is given by the sum of the potential $m_{a,0}$ contribution of (1.51) which would explicitly break the shift symmetry and the low-energy QCD contribution of (1.48). The ALP-Higgs couplings at dimension five can be removed by ALP field re-definitions using the five global U(1) symmetries of the SM [81, 82].

In the fermion mass basis chosen in [82] and coupling definition used in (1.51), the gauge field couplings become

$$\begin{aligned}\tilde{C}_{GG}(\Lambda) &= C_{GG} + \frac{1}{32\pi^2} \sum_q C_{qq}(\Lambda) \\ \tilde{C}_{WW}(\Lambda) &= C_{WW} - \frac{1}{32\pi^2} \text{Tr}[3\mathbf{k}_U + \mathbf{k}_E] \\ \tilde{C}_{BB}(\Lambda) &= C_{BB} + \frac{1}{32\pi^2} N_c^f Q_f^2 \sum_f C_{ff} + \frac{1}{32\pi^2} \text{Tr}[3\mathbf{k}_U + \mathbf{k}_E],\end{aligned}\tag{1.52}$$

where N_c^f and Q_f denote the color multiplicity and electric charge. The ALP-fermion coupling matrices \mathbf{k}_F and \mathbf{k}_f correspond to the coupling to the left-handed doublets and right-handed singlets respectively and $C_{f_i f_i} = [\mathbf{k}_F]_{ii} - [\mathbf{k}_f]_{ii}$.

Motivation for generalization

The breaking of U(1) symmetries is a common phenomenon in many theories beyond the SM. It is, therefore, convenient to search for these particles in a model-independent manner. Examples of symmetries already present in the SM, which were presented in section 1.1 which can be spontaneously broken are, e.g., lepton number, resulting in pNGBs called *majorons* or *familons* which connected with the breaking of the family number. One of the first ideas of majorons came in Ref. [83], where neutrinos are considered Majorana particles and hence break the lepton number, but unlike in the usual approach, the lepton number is broken spontaneously. In the case of familons, where the flavor symmetries are spontaneously broken, a general approach was discussed in Ref. [84].

Many of the unified field theories follow the idea behind the Kaluza-Klein theory, where the space-time in which the theory is postulated is made of a tensor product of the 4-dimensional Minkowski space-time (or a pseudo-Riemannian manifold in general) and a manifold geometrically reflecting given

symmetries. At low energies, this manifold can be seen as compactified, and the original action in the higher dimensional space-time gives rise to gauge fields in four dimensions obeying the symmetries of the compactified manifold¹⁷ and extra modes which decouple from these gauge fields are not observable in four dimensions.

In the late 90s, it was realized that the size of some dimensions need not be as small as Planck length ($\sim 10^{-35}$ m) and instead could have the right size to solve the hierarchy problem [86]. This new approach has also influenced other theories beyond the SM, even if they have a different original motivation. Axions, which in models built in the previous section 1.3.2 seek for a mechanism putting the SSB scale f_a into high values, are one such case. Thus, the solution could be similar to the solutions for gravity as axions would propagate into our four dimensions as Kaluza-Klein particles [87].

Meanwhile, a different strategy, much more similar to the original KK one, was followed during the years, and it was shown that by reducing more than one dimension, non-Abelian gauges could be restored [88]. Most of the unification schemes of today try to restore the whole $SU(3) \times SU(2) \times U(1)$ SM gauge group, which corresponds to a manifold of at least seven compact dimensions. In such theories, the cases of ALPs as presented above (majorons, familons, etc.) have a much more general character, which supports the idea that many ALP modes could appear in the spectrum independently. A similar dimensional reduction scheme is also followed in string theories, where ALP modes (besides the axion one [89]) appear very naturally as KK zero modes of antisymmetric tensor fields or more similarly to the previous case as open strings on D-branes. The plenitude of ALPs in string theory is recently referred to as a *string axiverse* [90] and predicts ALPs in experimentally reachable regions.

¹⁷The original KK theory assumed the simplest topological manifold, the circle, and a 5-dimensional metric incorporating the 4-dimensional metric, the EM 4-potential and an additional scalar field. As a result of the compactification, it restored the general relativity equations and the QED from the unified equation, where the $U(1)_{EM}$ symmetry corresponded to the transformations in the compactified dimension. The theory, however, failed in the prediction as, when incorporating the fermionic field corresponding to the electron with the given electric charge, it predicted an incorrect value of the electron mass by many orders of magnitude [85].

1.3.5 Theory of ALPs at and below the EW scale

At energy scales below the EW symmetry breaking scale, the interactions of ALPs with EM and weak bosons take the following form

$$\mathcal{L}_{\text{ALP}} \supset e^2 \frac{C_{\gamma\gamma}}{\Lambda} a F^{\mu\nu} \tilde{F}_{\mu\nu} + 2e^2 \frac{C_{\gamma Z}}{\Lambda} a F^{\mu\nu} \tilde{Z}_{\mu\nu} + 4g^2 \frac{C_{WW}}{\Lambda} a \varepsilon^{\mu\nu\alpha\beta} \partial_\mu W_\nu^+ \partial_\alpha W_\beta^-, \quad (1.53)$$

with couplings given by

$$C_{\gamma\gamma} = C_{WW} + C_{BB} \quad (1.54)$$

$$C_{\gamma Z} = \frac{C_{WW} \cos^2 \theta_W - C_{BB} \sin^2 \theta_W}{\sin \theta_W \cos \theta_W}. \quad (1.55)$$

At the one-loop level, $C_{\gamma\gamma}$ receives additional contributions from fermion and gauge boson loops, leading to

$$C_{\gamma\gamma} = C_{WW} + C_{BB} + \sum_f \frac{N_c^f Q_f^2}{16\pi^2} C_{ff} B_1 \left(\frac{4m_f}{m_a} \right) + \frac{2\alpha}{\pi} \frac{C_{WW}}{\sin^2 \theta_W} B_2 \left(\frac{4m_W}{m_a} \right) \quad (1.56)$$

where the loop functions $B_{1,2}(\tau)$ are defined in eq. (14) of Ref. [91].

While the gauge boson couplings do not run from the UV scale Λ to the weak scale $\mu \simeq m_t$, the fermion couplings undergo a nontrivial development, and they receive an effective contribution from the gauge couplings even if $C_{ff}(\Lambda) = 0$. In this work the numerical results derived in [78] for $\Lambda = 1$ TeV are used.¹⁸ For the flavor-conserving interaction

$$\mathcal{L}_{\text{fermion}}(\mu) = \frac{\partial^\mu a}{2\Lambda} \sum_{f \neq t} C_{ff}(\mu) \bar{f} \gamma_\mu \gamma_5 f \quad (1.57)$$

the used numerical values from RG evolution and matching are for $\Delta C_{ff} = C_{ff}(m_t) - C_{ff}(\Lambda)$ and $C_{VV} = C_{VV}(\Lambda)$

$$\begin{aligned} \Delta C_{tt} &\simeq -[6.17\tilde{C}_{GG} + 0.23\tilde{C}_{WW} + 0.02\tilde{C}_{BB}] \times 10^{-3} \\ \Delta C_{bb} &\simeq 0.097C_{tt} - [7.02\tilde{C}_{GG} + 0.19\tilde{C}_{WW} + 0.005\tilde{C}_{BB}] \times 10^{-3} \\ \Delta C_{uu,cc} &\simeq -0.116C_{tt} - [6.35\tilde{C}_{GG} + 0.19\tilde{C}_{WW} + 0.02\tilde{C}_{BB}] \times 10^{-3} \\ \Delta C_{dd,ss} &\simeq 0.116C_{tt} - [7.08\tilde{C}_{GG} + 0.22\tilde{C}_{WW} + 0.005\tilde{C}_{BB}] \times 10^{-3} \\ \Delta C_{ee} &\simeq 0.116C_{tt} - [0.37\tilde{C}_{GG} + 0.22\tilde{C}_{WW} + 0.05\tilde{C}_{BB}] \times 10^{-3}. \end{aligned} \quad (1.58)$$

¹⁸The full implementation of the fermion coupling RG evolution is planned in the future.

Under the assumption of minimal flavor violation [92], the numerical result for the FCNC couplings is

$$[k_U(m_t)]_{ij}^{\text{MFV}} \simeq V_{ti}^* V_{tj} [\Delta k_U(\Lambda) + 1.9 \times 10^{-2} C_{tt}(\Lambda) - 6.1 \times 10^{-5} \tilde{C}_{GG}(\Lambda) - 2.8 \times 10^{-5} \tilde{C}_{WW}(\Lambda) - 1.8 \times 10^{-7} \tilde{C}_{BB}(\Lambda)], \quad (1.59)$$

where $\Delta k_U(\Lambda)$ is the contribution from matching at UV, and $[k_F(m_t)]_{ij}^{\text{MFV}} = [k_f(m_t)]_{ij}^{\text{MFV}} = 0$ for all other contributions. Below the EW scale, the flavor-changing couplings do not evolve, while the flavor-conserving couplings are given by equation (2.35) in Ref. [78].

The term (1.57) is related to

$$\mathcal{L}_{\text{Yukawa}} = i \sum_f y_f a \bar{f} \gamma_5 f \quad (1.60)$$

via equations of motion. Note, however, that they are not equivalent and can be transformed from one to the other only in a free-field theory. For example, (1.57) exhibits the shift symmetry and Adler's zero condition as was the case of pions in (1.13) while the amplitudes of (1.60) do not vanish for $p \rightarrow 0$. Nevertheless, this case of Yukawa-coupled pseudoscalar will also be probed in the following chapter for $y_f = g_Y m_f / v$ as it also leads to FCNC transitions at one-loop level

$$\mathcal{L}_{a q_i \bar{q}_j} = C_{q_i q_j, a}^{R, L} (a \bar{q}_{L, i} q_{R, j} + \text{h.c.}) \quad (1.61)$$

with (depending if q_i, q_j are up- or down-type) [93, 94]

$$C_{q_i q_j, a}^{R, L} = i g_Y \frac{m_{q_i, q_j}}{v} \frac{g^2}{16\pi^2} \log\left(\frac{\Lambda^2}{m_t^2}\right) \sum_{q=u, c, t(d, s, b)} \xi_q V_{tq_i} V_{tq_j}^* ; \xi_q = \frac{m_q^2}{m_W^2}, \quad (1.62)$$

where $\log(\Lambda/m)$ regularizes the otherwise divergent contribution. Exploiting this coupling can be useful in theories extending the Higgs sector to solve the hierarchy problem.

1.3.6 Low-energy theory of ALPs

At energies below $\simeq 2 \text{ GeV}$, the Lagrangians (1.53) and (1.57) are matched with the chiral effective Lagrangian (1.14). To do so, first, the gluon-coupling term in (1.51) can be absorbed in the quark field definition by performing a

chiral rotation

$$q \rightarrow \exp[i(a/f_a)\kappa_q\gamma_5]q \quad (1.63)$$

under the assumption that $\text{Tr}[\kappa] = 1$. With this choice the quark mass matrix is modified as $m_q(a) = \exp[i(a/f_a)\kappa_q\gamma_5]m_q\exp[i(a/f_a)\kappa_q\gamma_5]$ and the quark- and photon-ALP couplings as

$$\begin{aligned} \tilde{C}_{\gamma\gamma} &= \tilde{C}_{\gamma\gamma} - 2N_c\text{Tr}[\kappa\mathbf{Q}\mathbf{Q}]C_{GG} \\ \tilde{C}_{qq} &= C_{qq} + 64\pi^2\kappa_qC_{GG}, \end{aligned} \quad (1.64)$$

where $N_c = 3$ is the number of colors. The effective Lagrangian then becomes

$$\begin{aligned} \mathcal{L}_{\text{ALP,eff}} &= \bar{q}[i\not{D} - m_q(a)]q + \frac{1}{2}(\partial^\mu a)(\partial_\mu a) - \frac{m_{a,0}^2}{2}a^2 \\ &+ e^2\frac{\tilde{C}_{\gamma\gamma}}{\Lambda}aF^{\mu\nu}\tilde{F}_{\mu\nu} + \sum_f\frac{\tilde{C}_{ff}}{2\Lambda}(\partial^\mu a)\bar{f}\gamma_\mu\gamma_5f, \end{aligned} \quad (1.65)$$

which is matched to

$$\begin{aligned} \mathcal{L}_{\text{ALP,eff}} &= \frac{f_\pi^2}{4}\text{Tr}\left(D_\mu U D^\mu U^\dagger\right) + \frac{f_\pi^2 B_0}{2}\text{Tr}\left(\mathcal{M}U^\dagger + U\mathcal{M}^\dagger\right) \\ &+ \frac{1}{2}(\partial^\mu a)(\partial_\mu a) - \frac{m_{a,0}^2}{2}a^2 + e^2\frac{\tilde{C}_{\gamma\gamma}}{\Lambda}aF^{\mu\nu}\tilde{F}_{\mu\nu} \\ &+ i\frac{f_\pi}{2\Lambda}(\partial^\mu a)\text{Tr}[\tilde{C}_q(U^\dagger D_\mu U - U D_\mu U^\dagger)] \end{aligned} \quad (1.66)$$

where with EM field as the external source, the covariant derivative from (1.17) is $D_\mu U = \partial_\mu U + ieA_\mu[\mathbf{Q}, U]$.

The chiral rotation (1.63) induces mass m_{aP} and kinetic K_{aP} mixing between the ALP and the neutral pseudoscalar mesons $P = \{\pi^0, \eta, \eta'\}$. After diagonalization of the kinetic and mass matrices, the obtained physical fields are given by

$$P = P_{\text{phys}} - \delta_I \sum_{P'} \frac{m_{PP'}^2}{m_P^2 - m_{P'}^2} P'_{\text{phys}} + \theta_{aP} a_{\text{phys}}. \quad (1.67)$$

The ALP field can then be expressed as a U(3) matrix

$$a = \theta_{a\pi^0}\pi^0 + \theta_{a\eta}\eta + \theta_{a\eta'}\eta' \quad (1.68)$$

with individual pseudoscalar fields as in (1.12) and the mixing angles in the leading order in isospin breaking [95]

$$\theta_{aP} = -\frac{f_\pi K_{aP} m_a^2 + m_{aP}^2}{\Lambda (m_a^2 - m_P^2)}. \quad (1.69)$$

A convenient choice of $\kappa_q = m_q^{-1}/\text{Tr}[\mathbf{m}^{-1}]$ with $\mathbf{m} = \text{diag}\{m_u, m_d, m_s\}$ results in $m_{a\pi^0} = m_{a\eta_8} = 0$ [81]. The ALP mass mixing is then given solely by $m_{a\eta_8} \neq 0$ and η - η' mixing $\theta_{\eta\eta'} \simeq -13^\circ$ [15] as

$$\begin{aligned} m_{a\eta}^2 &= -64\pi^2 C_{GG} \frac{\sqrt{6} B_0 \sin \theta_{\eta\eta'}}{\text{Tr}[\mathbf{m}^{-1}]} \\ m_{a\eta'}^2 &= 64\pi^2 C_{GG} \frac{\sqrt{6} B_0 \cos \theta_{\eta\eta'}}{\text{Tr}[\mathbf{m}^{-1}]} \end{aligned} \quad (1.70)$$

The kinetic mixing terms are given by

$$\begin{aligned} K_{a\pi^0} &= -\frac{1}{2}(\tilde{C}_{uu} - \tilde{C}_{dd}) \\ K_{a\eta} &= -\frac{1}{\sqrt{12}}[(\tilde{C}_{uu} + \tilde{C}_{dd})(\cos \theta_{\eta\eta'} - \sqrt{2} \sin \theta_{\eta\eta'}) \\ &\quad - \tilde{C}_{ss}(2 \cos \theta_{\eta\eta'} + \sqrt{2} \sin \theta_{\eta\eta'})] \\ K_{a\eta'} &= -\frac{1}{\sqrt{12}}[(\tilde{C}_{uu} + \tilde{C}_{dd})(\cos \theta_{\eta\eta'} + \sqrt{2} \sin \theta_{\eta\eta'}) \\ &\quad - \tilde{C}_{ss}(2 \cos \theta_{\eta\eta'} - \sqrt{2} \sin \theta_{\eta\eta'})]. \end{aligned} \quad (1.71)$$

The effective ALP-photon coupling receives additional contributions from the a - P mixing

$$\tilde{C}_{\gamma\gamma} = C_{\gamma\gamma} - 1.92 C_{GG} + \tilde{\alpha}_s(m_a^2) \mathcal{F}_{\text{VMD}} \sum_P \frac{\Lambda}{16\pi^2 f_P} \theta_{aP} \quad (1.72)$$

where the C_{GG} contribution from (1.64) is evaluated to one-loop [96] and $f_{\eta'} \approx 73$ MeV [15]. The functions \mathcal{F}_{VMD} and $\tilde{\alpha}_s$, will be defined in (2.4) and (2.5) respectively, allow to extrapolate the mixing effects for $m_a > m_{\eta'}$.

The ALP-meson interactions are calculated in this work using the vector meson dominance (VMD) [97] approach charted in Ref. [95] while keeping $C_{qq} \neq 0$.¹⁹

¹⁹An independent study [98] of $C_{qq} \neq 0$ case has been published when this work was being finalized. The computations in Ref. [98] include additional interactions on top of those included in Ref. [95], also used in this work.

1.4 Hidden sector portals

The origin of Dark Matter remains one of the most significant unknowns of modern physics, with no indication of which solution is more probable, even though some solutions have been excluded, such as the SM neutrinos. Theory models aiming to explain the DM existence are often based on convenience rather than likelihood, where this convenience is typically the aim to simultaneously address some other unexplained phenomena to reduce the number of free parameters of our current theory. However, the most pessimistic scenario, that the problems are unrelated, cannot be excluded. If the DM can be explained by the existence of some unknown quantum fields, the most general prediction that can be made is the postulation of a so-called *Hidden* (or *Dark*) *sector* (HS), which is a collection of quantum fields independent of the SM. If the HS interacts with the SM only gravitationally, any further analysis of its internal structure becomes very challenging in today's universe, no matter if it has a rich or very simple gauge structure.

Lorentz invariance requires that the *elementary* quantum fields corresponding to observed particles have spin ≤ 2 and that the only spin 2 field is the massless graviton coupled to the stress-energy tensor [99, 100]. Spin 3/2 elementary fields are possible only as superpartners to bosonic fields in supersymmetry. If the HS is coupled to the SM via other interaction than gravity and one does not assume supersymmetry, then by dropping the restriction on specific models, so not restricting oneself to solving any of the problems mentioned in previous sections, one can probe experimentally the couplings of general spin 0, 1/2 and 1 HS fields to all possible SM fields, forming so-called *hidden sector portals*. Without any prior knowledge of the HS internal structure, this model-independent classification might be the most efficient way to search for any NP beyond the SM.

The interpretation of experimental searches for these new interactions can be complicated because of a typically large number of possible production and decay (if unstable) mechanisms of the corresponding NP (or as referred to later on *exotic*) particles where the production cross-sections σ and decay widths Γ vary depending on the model or more specifically in our approach, on the coupling setup. From the Effective Field Theory (EFT) point of view, one is extending the SM Lagrangian with additional operators \mathcal{O}_i which can be formed by both SM and HS fields. In the case of the above-mentioned HS portals, one talks about operators containing both terms

$$\mathcal{L} \supset \sum_i \frac{\mathbf{C}_i}{\Lambda_i^{d-4}} \mathcal{O}_i^d = \sum_{j,k} \frac{\mathbf{C}_{j,k}}{\Lambda_{j,k}^{d-4}} \mathcal{O}_{\text{SM},j} \times \mathcal{O}_{\text{HS},k}, \quad (1.73)$$

where i, j, k are arbitrary indices labeling all the contributions, d is the dimension of the operator, and Λ is the “New Physics scale”, a scale at which the fields forming \mathcal{O}_{SM} and \mathcal{O}_{HS} are replaced by fields of a BSM theory which possibly unifies these two distinct sectors. Λ and the dimensionless coupling constant²⁰ \mathbf{C} determine the strength of the interaction involving given fields and, therefore, the size of the production and decay cross-sections.

The various contributions i to the searched process may not be trivially disentangled as very different NP models can leave very similar experimental signatures and vice versa. This has led to the definition of a number of benchmark scenarios, which are a focus of the experimental searches. For simplicity, these benchmark scenarios often focus on contributions driven by one sole coupling constant \mathbf{C} . This setup typically does not correspond to the predicted models since the coupling strength is scale-dependent and the same as for SM fields, as explained in section 1.1; coupling an NP field to an SM field introduces effectively also a coupling to other SM fields at a loop level. While these contributions are typically subdominant, a cancellation of tree-level contributions may occur in some cases. Alternatively, even a suppressed channel can become important if favored by a specific experimental setup or if studied at a specific energy. Therefore, data analyses of experimentally collected data that assume very specific models, such as the benchmark scenarios, are becoming disfavored as it can be challenging to generalize or reinterpret them for different models. Redoing the analysis and corresponding simulations can be very computationally demanding, which calls for a more model-independent approach already at the analysis level.

The benchmark models can still be a very useful probe, however, if the portal and coupling are carefully selected, as often different setups lead to similar experimental signatures as will be shown in chapters 2 and 4. The commonly used benchmark models for HS searches are those selected in Ref. [101, 102], often denoted as *BCI-BCII*. In the following subsections, several of these benchmark models will be introduced, and a model-independent treatment for all of them following Ref. [1] will be shown in section 1.4.4.

²⁰Note that often instead of Λ and \mathbf{C} a dimensionful coupling is defined.

1.4.1 Dark photon

The Dark photon (DP) A' is a gauge boson of a new $U(1)_{A'}$ symmetry interacting with the SM hypercharge through kinetic mixing term

$$\mathcal{L} \supset \mathcal{L}_{\text{HS}} - \frac{\varepsilon}{2 \cos \theta_W} F'_{\mu\nu} B^{\mu\nu}, \quad (1.74)$$

where $F'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$, $B_{\mu\nu}$ is the SM hypercharge field strength tensor, θ_W is the Weinberg angle and ε is the coupling constant. The main theoretical motivation for this SM extension was the explanation of the discrepancy between the SM-predicted and the measured [103] value of the anomalous magnetic dipole moment, so-called ($g-2$) anomaly [104]. A possible extension of the dark photon portal is to dark vectors with coupling to the $B-L$ current or separately lepto- and proto-phobic vectors [105].

In the minimal scenario, the hidden sector consists only of a fermion field χ coupled to A' via g_D

$$\mathcal{L}_{\text{HS}} \supset -\frac{1}{4}(F'_{\mu\nu})^2 + \frac{1}{2}m_{A'}^2(A'_\mu)^2 + |(\partial_\mu + ig_D A'_\mu)\chi|^2. \quad (1.75)$$

Three DP benchmark scenarios set in Ref. [101] ($BC1$ - $BC3$) correspond to the limiting cases

- $g_D \rightarrow 0$, in which case when DP is produced, it can only decay back to the SM particles ($BC1$);
- g_D dominates, so when DP is produced, it decays predominantly to the DM particles ($BC2$);
- $m_{A'} \rightarrow 0$, then $\chi, \bar{\chi}$ become so-called *milicharged particles* as they get an effective electric charge $|Q_\chi| = |\varepsilon g_D e|$ ($BC3$).

1.4.2 Dark scalar

The Dark scalar (DS) S presents the minimal extension of the SM scalar sector and a natural solution to the hierarchy problem [106]. The DS is a scalar singlet, therefore, without any direct coupling to the SM fermions with which it interacts only via mixing with the SM h boson. Above the EWSB scale, a scalar field Φ is coupled to the Higgs doublet H in the following way

$$\mathcal{L} \supset \mathcal{L}_{\text{HS}} - (\mu_S \Phi + \lambda_S \Phi^2) H^\dagger H. \quad (1.76)$$

Similarly to the DP, the HS is assumed to consist of a Dirac fermion field χ

$$\mathcal{L}_{\text{HS}} \supset \frac{1}{2}m_0^2\Phi^2 + g_\chi\Phi\chi\bar{\chi}. \quad (1.77)$$

Below the EWSB scale when the H field settles in $H^T = (0, (v+h)/\sqrt{2})$ with the EW VEV $v = 246$ GeV, the non-zero μ_S leads to a mass mixing between H and Φ via matrix

$$\mathcal{M}^2 = \begin{pmatrix} 2\lambda v^2 & \mu_S v \\ \mu_S v & m_0^2 \end{pmatrix} \quad (1.78)$$

which can be diagonalized by rotation

$$\begin{pmatrix} H \\ \Phi \end{pmatrix} = R(\theta) \begin{pmatrix} h \\ S \end{pmatrix} \quad (1.79)$$

where S is the physical DS and the mixing angle θ [107]

$$\tan 2\theta = \frac{2\mu_S v}{m_0^2 - 2\lambda v^2}. \quad (1.80)$$

For small values of the angle, the mixing can be approximated by $\sin \theta$ or θ and mass eigenvalues m_h and m_S

$$\sin \theta = \frac{\mu_S v}{m_S^2 - m_h^2}. \quad (1.81)$$

The h - S mixing leads to an effective Yukawa-type coupling of S to the SM fermions

$$\mathcal{L}_{Sf\bar{f}} = \sum_f C_{ff,S} S \bar{f} f, \quad (1.82)$$

where $C_{ff,S} = m_f \sin \theta / v$. At a one-loop level, it also induces flavor-changing neutral current (FCNC) transitions such as $b \rightarrow s + S$ or $s \rightarrow d + S$ leading to decays $B \rightarrow KS$ or $K \rightarrow \pi S$

$$\mathcal{L}_{Sq_i\bar{q}_j} = C_{q_i q_j, S} (S \bar{q}_{L,i} q_{R,j} + \text{h.c.}), \quad (1.83)$$

where

$$C_{q_i q_j, S} = \frac{m_{q_i} \sin \theta}{v} \frac{3g^2}{64\pi^2} \sum_{q=u,c,t(d,s,b)} \xi_q V_{qq_i} V_{qq_j}^*. \quad (1.84)$$

For non-zero g_χ , SM particle decays into $\chi\bar{\chi}$ pair such as $h \rightarrow \chi\bar{\chi}$ are possible if kinematically allowed, similarly as $S \rightarrow \chi\bar{\chi}$. This scenario of large g_χ and invisibly decaying mediator S is however largely excluded by $h \rightarrow \chi\bar{\chi}$ and FCNC $B \rightarrow K + \text{inv}$ and $K \rightarrow \pi + \text{inv}$ measurements [107]. Visibly decaying mediator scenarios are therefore considered.

The coupling $\lambda_S \neq 0$ does not lead to additional decays of S but enhances the production as S can be produced in pairs. Directly from (1.76) one obtains $\mathcal{L}_{\text{int}} \supset \lambda_S S^2 (vh + h^2/2)$ and therefore $h \rightarrow SS$ decay

$$\Gamma_{h \rightarrow SS} = \frac{\lambda_S^2 v^2}{8\pi m_h} \sqrt{1 - \frac{4m_S^2}{m_h^2}}. \quad (1.85)$$

Also, S can be produced in pairs in FCNC $K \rightarrow \pi SS$, $B \rightarrow KSS$, etc. decays [108]

$$\mathcal{L}_{SSq_i\bar{q}_j} = C_{q_i q_j, SS} (S^2 \bar{q}_{L,i} q_{R,j} + \text{h.c.}), \quad (1.86)$$

with

$$C_{q_i q_j, SS} = \frac{m_{q_i} \lambda_S}{m_h^2} \frac{3g^2}{64\pi^2} \sum_{q=u,c,t(d,s,b)} \xi_q V_{qq_i} V_{qq_j}^*. \quad (1.87)$$

Two DS benchmarks were set in Ref. [101] (*BC4-BC5*)

- $\lambda_S \rightarrow 0$, in which case the S production is controlled by $\sin \theta$ (*BC4*);
- enhanced pair production ($\lambda_S \neq 0$): The current UL of $\text{BR}_{h \rightarrow SS}$ is assumed to comply with observation and enhance production in $B \rightarrow KSS$, etc. decays. The statistical combination of the LHC Run 1 and 2 data, has put $\text{BR}_{h \rightarrow \text{invisible}} < 0.093$ at 90% CL [109]. Given equation (1.85) and assuming $m_S \ll m_h$, this case corresponds to the extension of *BC4* with $\lambda_S \simeq 3.9 \times 10^{-3}$ (*BC5*).

1.4.3 Heavy neutral lepton

The HNLs, including several popular models, were already introduced in section 1.2.1. Here, the main focus will be again on simplified benchmark scenarios characteristic by couplings to a single generation of leptons, which can give a rough idea of the impact of these couplings on experimental searches.

The general form of the HNL portal is

$$\mathcal{L} \supset \mathcal{L}_{\text{HS}} + \sum_{\alpha, I} F_{\alpha I} (\bar{L}_\alpha H) N_I, \quad (1.88)$$

where L_α is the left-handed doublet of the SM neutrino of flavor α , N_I is the I -th HNL field and $F_{\alpha I}$ are the corresponding couplings. The form of \mathcal{L}_{HS} is model-dependent, but for the Type I seesaw, it would contain the Majorana mass terms M_R for the HNLs, but in general, they can also be of the Dirac type.

As described in section 1.2.1, below the EWSB scale, the SM neutrino fields obtain the Dirac mass m_D , and after diagonalizing the mass matrix, the SM neutrinos mass eigenstates have mass $m_\nu \approx m_D^2/M_R$. Upon the diagonalization, one obtains mixing between the SM neutrino fields ν and the HNL fields N , typically parametrized by elements of matrix U for the respective flavors. Processes involving HNLs can then be calculated from the neutrino processes by an exchange $\nu_\alpha \rightarrow U_{\alpha I} N_I$.

The minimal benchmark models (*BC6*, *BC7* and *BC8*) from Ref. [101] always assume one Majorana HNL with dominant mixing with one SM neutrino flavor (e , μ and τ).

1.4.4 Model-independent approach

When a signal is observed at an experiment, the number of events in the detector N_{det} depends at first sight in a complicated non-linear way on the couplings \mathbf{C} from (1.73) as they determine not only the production cross sections and decay ratios into various final states but also the kinematic distributions and particle lifetimes $\tau = 1/\Gamma$. Upon a closer look, for a fixed set of couplings \mathbf{C} , the quantity can be factorized into an exotic particle X production and a decay part

$$N_{\text{det}} = \int d\theta_X dE_X \frac{d^2 N(m_X, \mathbf{C})}{d\theta_X dE_X} p_{\text{det}}(m_X, \Gamma_X, \theta_X, E_X, \mathbf{C}). \quad (1.89)$$

This factorization has been used initially in Ref. [1] on the example of fixed-target experiments, which will be elaborated more in the following chapter 2, but is applicable in general for any laboratory and cosmological searches. Variables θ_X and E_X correspond to the angle under which X is produced with respect to some reference axis and the particle energy. $d^2 N/d\theta_X dE_X$ denotes the spectrum of produced particles, or later referred to as a *differential yield*,

and p_{det} denotes the probability that an exotic particle of given properties will lead to an observable signal in the detector.

The exotic particle production can happen through many different processes and up to some rare cases where an interference among them occurs²¹, these processes are independent, and one can write the differential yield simply as a sum over all the contributions

$$\frac{d^2N(m_X, \mathbf{C})}{d\theta_X dE_X} = \sum_i \frac{d^2N_i(m_X, \mathbf{C})}{d\theta_X dE_X}. \quad (1.90)$$

Each contribution is driven by some, in general unique, model-dependent parameters which at the first order affect the overall normalization but do not affect the spectrum. The model-dependent parameter can be a branching ratio (BR) or an SM-HS coupling itself.²² It is therefore possible to calculate the spectrum for a reference coupling \mathbf{C}_{ref} and then perform an appropriate rescaling

$$\frac{d^2N_i(m_X, \mathbf{C})}{d\theta_X dE_X} \simeq f_i(\mathbf{C}, \mathbf{C}_{\text{ref}}) \frac{d^2N_i(m_X, \mathbf{C}_{\text{ref}})}{d\theta_X dE_X}. \quad (1.91)$$

On the contrary, the spectrum depends on the exotic mass m_X in a non-trivial way, and it is therefore necessary to calculate $d^2N_i/d\theta_X dE_X$ separately for each m_X considered.

Similarly to (1.91), the detection probability can be split into different final states f as

$$p_{\text{det}}(m_X, \theta_X, E_X, \mathbf{C}) = \sum_f \text{BR}_{X \rightarrow f}(m_X, \mathbf{C}) p_{\text{det},f}(m_X, \Gamma_X, \theta_X, E_X), \quad (1.92)$$

where $\text{BR}_{X \rightarrow f} = \Gamma_{X \rightarrow f}/\Gamma_X$ denotes the branching ratio into the final state f . While these branching ratios may depend in a non-trivial way on the couplings \mathbf{C} , the detection probability $p_{\text{det},f}(m_X, \Gamma_X, \theta_X, E_X)$ for a given final

²¹This can be the case of the *mixing* production processes when exotic particle mixing with multiple SM particles occurs. In these cases, the relative phases of the corresponding amplitudes are, unfortunately, often not known. However, in mixing production, the differential yield is typically strongest close to the mass of the SM particle the exotic particle mixes with, and therefore, if mixing with multiple SM particles occurs, these channels can be considered independent to a good approximation.

²²The coupling typically enters in second power, but also in more complicated cases, the model-dependent contribution can be factorized out even if only as a branching ratio even though such use cases are less universal.

state depends on the couplings only through the exotic particle lifetime, which determines the probability of the X to decay at a given position. Again, this makes it possible to calculate the detection probabilities $p_{\text{det},f}$ in advance and then apply an appropriate rescaling through the branching ratios $\text{BR}_{X \rightarrow f}$.

Putting everything together

$$N_{\text{det}} = \sum_{i,f} f_i(\mathbf{C}, \mathbf{C}_{\text{ref}}) \text{BR}_{X \rightarrow f}(m_X, \mathbf{C}) \times \int d\theta_X dE_X \frac{d^2 N_i(m_X, \mathbf{C}_{\text{ref}})}{d\theta_X dE_X} p_{\text{det},f}(m_X, \Gamma_X, \theta_X, E_X) \quad (1.93)$$

$$= \sum_{i,f} M_{if}(m_X, \mathbf{C}) \times E_{\text{det},if}(m_X, \Gamma_X), \quad (1.94)$$

where M_{if} denotes the part that depends on the NP model (i.e., on the couplings \mathbf{C}), while $E_{\text{det},if}$ encapsulates all the experimental details.

The calculation of $E_{\text{det},if}$ for a given experiment can be challenging, typically requiring Monte Carlo (MC) simulations of both the particle production and the decay. However, these simulations only have to be performed once for given m_a and Γ_a and a choice of \mathbf{C}_{ref} and the (tabulated) functions $E_{\text{det},if}$ can easily be made publicly available. This makes it possible to perform fast analyses of any model for which the model-dependent functions M_{if} have been calculated without redoing the analysis for every change in the searched model.

This factorization of the calculation of the number of expected events on a model-dependent and model-independent part stood behind the development of the ALPINIST²³ framework, an open-source tool available at [110]. The tool was originally designed for the interpretation of ALP models in which the entries of $M_{if}(m_a, \mathbf{C})$ are masses and couplings from the lagrangian (1.51), but has now been extended to interpret also a vector $BC1$ and a scalar $BC4$ and $BC5$ models from the previous section. Extension of $BC1$ to a generic vector is under development, as well as the implementation of HNLs with generic couplings, which will allow to interpret not only $BC6-8$ but any combinations of the three couplings. This interpretation is handled by the `ALPrescale` module of the framework.

Since the framework was originally used to interpret proton beam fixed-target experiment searches, it also includes a module that handles the exotic

²³An acronym for *Axion-like particles in numerous interactions simulated and tabulated*.

particle production in processes that take place at energies available in fixed-target facilities. The `ALPproduction` module uses analytic expressions for perturbative processes or performs an MC simulation of exotic particle production in SM particle decays. It can use `PYTHIA` [111] or other external generator for the SM particle production. In between the two stands the `ALPdecay` module, which is a simplified MC containing the basic geometry of fixed-target experiments. Tabulated datasets are passed among the modules to avoid rerunning all the steps after modifying some parameters. A schematic outline of the framework is shown in figure 1.1 where box elements represent modules and triangle elements represent tabulated datasets.

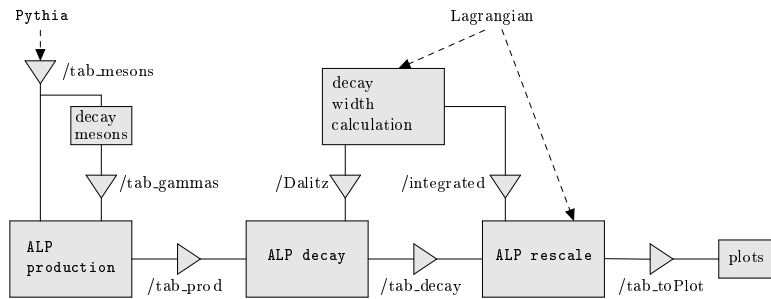


Figure 1.1: Diagram representing the internal structure of the ALPINIST framework.

The `ALPrescale` module for reinterpretation does not ask for the experimental details from the previous steps and allows the interpretation of data from different sources, such as electron or muon fixed-target experiments, collider experiments, or cosmology. A more complete simulator instead of the `ALPdecay` module can also be used. This option has been used in this work for the data analysis with the NA62 experiment and will be discussed more in chapters 3 and 4.

The `ALPdecay` module has been used in Refs. [1, 2, 6, 7] for estimation of the sensitivity of several past and future experiments, a work that will be discussed in the following chapters 2. Since the code of the framework is publicly available, simulation upgrades, including new production mechanisms or decay channels, or more realistic experiment modeling, can be easily included by the users from the public, allowing them straightforward estimates of the related phenomenological consequences.²⁴

²⁴Note that the `ALPdecay` module can also be helpful for optimization of experimental geometry of proposed experiments and crude optimization of analysis cuts for existing experiments and study of their impact on specific physics models.

Search for feebly-interacting particles with fixed-target experiments

2.1	FIP production	45
2.1.1	Direct production	45
2.1.2	Meson mixing production	48
2.1.3	Meson decay production	50
2.2	FIP detection	55
2.2.1	General principles of the ALP Monte Carlo simulation	56
2.2.2	Decay channels	58
2.2.3	Propagation of decay products	65
2.3	Fixed-target experiments	66
2.3.1	Results from past experiments	66
2.3.2	Acceptance of current experiments	69
2.3.3	Acceptance of future experiments	70
2.4	Results	74
2.4.1	Axions	74
2.4.2	Yukawa-coupled pseudoscalars	80
2.4.3	Dark photons	80
2.4.4	Dark scalars	81
2.5	Search for exotic particles in kaon decays	82
2.5.1	Search for ALPs and dark scalars in kaon decays . . .	82
2.5.2	Search for HNLs	85
2.6	Search for axions with KOTO parasitically	86

2.6.1	KOTO setup and data-taking	87
2.6.2	Axion production and detection	91
2.6.3	Bounds and projections for ALPs	98
2.7	Projection for KLEVER	100

Fixed-target experiments are among the most promising probes of feebly-interacting particles in the MeV-GeV mass range, thanks to typically large beam intensities. While the mass range that can be probed is determined mainly by the energy available in the primary beam interaction and the experiment sensitivity for the various decay channels, the range of the couplings of the particle to the SM depends on the detector geometry and efficiency and the search strategy.

There are two¹ main strategies for direct searches for *exotic* particles at fixed-target experiments:

1. Search for a decay of an observed SM particle into other observed SM particles and an exotic particle. The exotic particle mass, momenta and other properties can be reconstructed from the knowledge of the properties of the SM initial and final states.
2. Search for a decay of an exotic particle into observed SM particles. The exotic particle is then reconstructed using the SM final states.

The first strategy can be very sensitive in terms of the coupling as the exotic particle does not have to decay back to the SM final states, and therefore, the BR in equation (1.92) can be set to 1. The leading order contribution to matrix M_{ij} from equation (1.94) is then dependent on the coupling \mathbf{C} only in second power. This strategy can be limited by the beam intensity, which has to be such that the individual SM particle decays can be clearly separated in time to allow correct reconstruction of the events. Also, the mass of the SM particles limits the mass of the exotic particle.

The second strategy does not require the knowledge of the exotic particle production process, and the particle properties are reconstructed fully from its decay products. This allows shielding of the exotic particle production point instead of placing a detector around it, which suppresses the SM backgrounds

¹A third strategy can be a search for an interaction of the exotic particle in the detector material. This strategy may prove to be challenging in terms of background because the exotic particle cannot be tracked. Nevertheless, it can address scenarios in which the exotic particle cannot decay back into the SM final states.

and enables much higher beam intensity as the particle flux through the detector is very low. This technique is used in the so-called *beam-dump experiments*, which will be the main focus of this chapter. The disadvantage of this strategy is that at the leading order $M_{ij} \propto \mathbf{C}^4$ for relatively long-lived particles.

The two strategies can also be combined, in which case the kinematics is fully under control, and the backgrounds can be further suppressed. However, it carries the disadvantages of both strategies (low intensity and $M_{ij} \propto \mathbf{C}^4$) and, therefore, typically provides the lowest discovery potential. Specific examples of the first strategy or the combination will be given later in this chapter in section 2.5 on the case of the NA62 experiment, which can be operated in several data-taking modes, and a proposed HIKE experiment.

The first four sections of the chapter will be dedicated to the estimation of the sensitivity of fixed-target experiments in a beam-dump-like regime using the results for ALPs from [1] and extending the framework to other particle types, such as dark photons and dark scalars. Section 2.5 will show a summary of NA62 searches for exotic particles in charged kaon decays, projection for HIKE experiment from Ref. [6, 7] and the comparison with beam-dump searches. Sections 2.6 and 2.7 will show the results from Ref. [2] on *parasitic* search for axions in K_L fixed-target experiments.

2.1 FIP production

The fixed-target experiments studied in this work are provided with proton beams. The exotic particle can be produced in the target either directly in the interaction of the beam proton with the target nuclei or nucleons, in interactions of the secondary particles from the initial collision with the target nuclei, or in decays of the secondary particles. Following [1], the production modes will be introduced for ALPs; the implementation for other particles will be analogous.

2.1.1 Direct production

The ALPs with non-zero effective coupling to the SM photons $C_{\gamma\gamma}$ can be produced in the target in the so-called *Primakoff* process, in the interaction of photons with the target EM fields of the target nuclei described by $d\sigma_{\gamma N}$, calculated in [112] using the Weizsäcker-Williams approximation. The differential

yield per one beam proton interaction is then given by

$$\frac{d^2 N}{d\theta_a dE_a} = \frac{C_{\gamma\gamma}^2}{\Lambda^2} \frac{8e^4}{\sigma_{pN}} \frac{\sin\theta_a}{\pi E_{\text{beam}}} \int dp_t^2 d\phi_a \frac{d\sigma_{\gamma N}}{d\cos\theta_a} \gamma(E_a, p_t^2), \quad (2.1)$$

where the total cross-section for the incident proton on the nuclei with given atomic mass number A is $\sigma_{pN} \simeq \sigma_{pp} A^{2/3}$. $\gamma(E_a, p_t^2)$ is the photon distribution, and there are two sources of photons considered

1. off-shell photons in the EM scattering of the beam proton itself, a production calculated in [112];²
2. on-shell photons from 2γ decays of π^0 , η , and η' mesons produced in the collision of the beam proton with a target nucleon.

In order to obtain the $\gamma(E_a, p_t^2)$ distribution for on-shell photons from π^0, η, η' meson di-photon decays, the production of these secondary mesons has to be simulated. In this work the meson yields from p - p collisions at experimentally relevant energy values³ were simulated using PYTHIA 8.2 [111] with a `SoftQCD:all` flag and a pomeron flux parametrization parameter set as `SigmaDiffractive:PomFlux(5)`, a setup for which the light meson yield has been validated in [114] including the meson kinematic distributions. The corresponding photon distribution is then obtained by interpolating the dataset resulting from the flat phase-space decay of these mesons into photon pairs. The decay products are boosted to the laboratory frame, and the events are weighted with the π^0, η, η' di-photon branching fractions.

In figure 2.1, the predicted distribution of ALPs in the E_a - θ_a plane for the Primakoff production channel is shown.⁴ An anti-correlation between the two parameters in the sense that higher ALP energies imply smaller production angles and vice versa can be observed.

²It has been pointed out in Ref. [113] that this approximation becomes inaccurate for small ALP masses as well as for low beam energies. Given that elastic scattering only gives a subdominant contribution compared to ALP production from photons produced in meson decays, this approximation is nevertheless fully sufficient for the purposes of this work.

³For 400, 120 and 70 GeV proton beams the respective p - p center-of-mass energies \sqrt{s} are approximately 27.4, 15 and 11.5 GeV.

⁴The choice of parameters in this example is purely instructive and corresponds to multiple setups discussed in section 2.3.




Figure 2.1: Production differential yield as a function of energy and angle for a *photon-from-meson* induced ALP production at a copper target with a 400 GeV proton beam for fixed $C_{\gamma\gamma}/\Lambda = 10^{-4} \text{ GeV}^{-1}$ and $m_a = 0.5 \text{ GeV}$.

Analogously, the DP that mixes with the SM photon can be produced in the *dark photon bremsstrahlung*. The interactions for both direct productions are depicted in figure 2.2.




Figure 2.2: Feynman diagrams for the direct production. Left: Primakoff production of ALP a in interaction of beam proton p with either external EM field γ^* from target nuclei or on-shell photons γ from decays of secondary particles. Right: Bremsstrahlung production of dark photon A' in beam proton p scattering in external EM field.

This work employs the bremsstrahlung cross-section calculated in [115] also using the Weizsäcker-Williams approximation and therefore having the

same limitations as in the ALP production. The differential yield is given by

$$\frac{dN}{d\theta_{A'} dE_{A'}} = \varepsilon^2 \frac{e^2 E_{A'}^2 \sin(2\theta_{A'})}{8\pi^2 E_{\text{beam}}} \frac{\sigma_{pN}(s')}{\sigma_{pN}(s)} \omega(E_{A'}, \theta_{A'}), \quad (2.2)$$

with the weight function $\omega(E_{A'}, \theta_{A'})$ [GeV⁻²] given by equations (13) and (14) in [115] while factorizing out explicitly the coupling and transforming $(z, p_t^2) \rightarrow (E_{A'}, \theta_{A'})$. The reduced center-of-mass energy is $s' \simeq (E_{\text{beam}} + 2m_p)(E_{\text{beam}} - E_{A'})$.

2.1.2 Meson mixing production

If the ALP field is coupled to the SM quark or gluon fields, it naturally enters the kinetic and mass matrices of the effective Lagrangian of pseudoscalar mesons. Upon diagonalization of these matrices in order to obtain the mass eigenstates, the mixing angle between ALPs and mesons $P = \pi^0, \eta, \eta'$ emerges in the form of (1.69). It should be emphasized that this expression is valid for values m_a such that $|\theta_{aP}|^2 \ll 1$.

In the presence of such a mixing, it is possible to produce an ALP instead of the meson P via SM processes in a small fraction of cases proportional to $|\theta_{aP}|^2$. Note that this approach is an approximation as it does not account for the mutual dependence of the mixing angles θ_{aP} . This is manifest in the explicit dependence of θ_{aP} on matrix κ whose elements are arbitrary, while in the full calculation of the corresponding processes, the resulting amplitudes should be κ -independent.

The corresponding differential yield can be written as

$$\frac{d^2 N}{d\theta_a dE_a} = |\mathcal{F}_{\text{VMD}}(m_a)|^2 \sum_P \frac{\hat{\alpha}_s^2(m_a)}{\hat{\alpha}_s^2(m_P)} \frac{d^2 N_P}{d\theta_P dE_P} \Big|_{\substack{E_P \rightarrow E_a \\ \theta_P \rightarrow \theta_a}} |\theta_{aP}|^2, \quad (2.3)$$

where $d^2 N_P/d\theta_P dE_P$ denotes the differential yield of the meson P . To consider ALP masses up to 3 GeV and capture the suppression of ALP production for $m_a \gg m_P$ the VMD form factor

$$\mathcal{F}_{\text{VMD}}(m) = \begin{cases} 1, & \text{for } m \leq 1.4 \text{ GeV} \\ \sum_{i=0}^3 a_i m^i, & \text{for } 1.4 \text{ GeV} < m \leq 2 \text{ GeV} \\ (1.4 \text{ GeV}/m)^4, & \text{for } 2 \text{ GeV} < m \leq 3 \text{ GeV} \\ 0, & \text{for } m > 3 \text{ GeV}, \end{cases} \quad (2.4)$$

is included as well as the running strong coupling $\hat{\alpha}_s(m)$ in the form

$$\hat{\alpha}_s(m) = \begin{cases} 1, & \text{for } m \leq 1 \text{ GeV} \\ \sum_{i=0}^3 b_i m^i, & \text{for } 1 \text{ GeV} < m \leq 1.5 \text{ GeV} \\ 4\pi/7 \log(m/0.34 \text{ GeV})^2, & \text{for } m > 1.5 \text{ GeV}, \end{cases} \quad (2.5)$$

where parameters a_i and b_i are determined by requiring continuity of the resulting function and its first derivative.

For obtaining the distribution of $d^2N_P/d\theta_P dE_P$, the validated datasets from the simulation described in section 2.1.1 can be conveniently re-used. There is, however, an important subtlety: Since the pseudoscalar meson masses are generally different from the ALP mass, it is not possible to simply replace a pseudoscalar with an ALP with the same energy and momentum. Even if the replacement is made so that the three-momentum is conserved in a specific frame, it may not be conserved in a different frame. This introduces some degree of arbitrariness in the precise prescription used for the replacement. In this work, it is required that the three-momentum of the ALP is equal to the three-momentum of the pseudoscalar meson in the p - p center-of-mass (cm) frame. The corresponding ALP energy and angle are then calculated and then boosted result into the laboratory frame. For further details on the kinematic adjustment, including the calculation of the corresponding systematic uncertainty, see appendix A.




Figure 2.3: Production differential yield for an ALP with $m_a = 0.5 \text{ GeV}$ produced via a - η mixing assuming a fixed mixing angle $\theta_{a\eta} = 10^{-4}$ and 400 GeV proton beam.

The differential yield for production via ALP- η mixing is shown in figure 2.3. For the meson mixing contribution, the differential yield peaks at larger ALP angles than for Primakoff production shown in figure 2.1, but still gives a sizeable contribution in typical on-axis beam-dump experiments with an angular acceptance of a few mrad.

In a complete analogy, dark photons can be produced via mixing with light vector mesons $V = \rho, \omega$ and ϕ . The fitted values of

$$\mathcal{R}_\mu^V = \frac{\sigma(e^+e^- \rightarrow V \rightarrow \text{hadrons})}{\sigma(e^+e^- \rightarrow \mu^+\mu^-)} \quad (2.6)$$

from [105] based on the measured $\sigma(e^+e^- \rightarrow \text{hadrons})$ cross sections are used to parametrize the $A'-V$ mixing. The strategy for the kinematic adjustment is the same as for ALPs.

2.1.3 Meson decay production

The ALPs can have a very rich set of various couplings to the SM triggering production in FCNC transitions. Flavor-changing ALP couplings arise not only from interactions of ALPs with quarks [93] and electroweak gauge bosons [79] but also in higher-order processes from ALP-gluon interactions [80]. In all of these cases, ALPs can be produced in FCNC processes and can therefore be probed through rare meson decays in fixed-target experiments [94, 116]. However, in order for the produced exotic particles to be focused in the forward direction, it is essential that the parent meson decays before interacting with the target material. B and D mesons can thus be of particular importance since when produced at the beam energies of $\mathcal{O}(100)$ GeV their decay length is significantly smaller than the interaction length of the target material.⁵ The same holds for light vector mesons $V = \rho, \omega, \phi$ and neutral pseudoscalars $P = \pi^0, \eta, \eta'$ which were already shown to be a potentially interesting source of ALPs and DPs.

Light mesons

Electromagnetic decay of light mesons can be a significant source of dark photons that mix with the SM photon. Two-body decays of neutral pseudoscalar $P \rightarrow \gamma + A'$ and vector meson $V \rightarrow P + A'$ are considered. Since the decay

⁵The nuclear collision length for most of the target materials is ≈ 10 cm [15]. The decay length for a 100 GeV momentum B meson is < 1 cm and ≈ 1 cm for a D meson, while for K it is > 500 cm.

is isotropic in the rest frame of the original meson, the differential production yield depends only on the distribution of mesons and the branching ratios for the decays involving the dark photons. The decay width is given by the decay width of the corresponding SM process $\Gamma_{P \rightarrow \gamma\gamma}$ and $\Gamma_{V \rightarrow P\gamma}$ and accounting for the finite mass of the outgoing dark photon.

For the production in pseudoscalar decay, one obtains

$$\Gamma_{P \rightarrow \gamma + A'} = 2\varepsilon^2 \left(1 - \frac{m_{A'}^2}{m_P^2}\right)^3 \Gamma_{P \rightarrow \gamma\gamma} \quad (2.7)$$

while for the vector meson decay

$$\Gamma_{V \rightarrow P + A'} = \varepsilon^2 \frac{(m_V^2 - m_P^2 - m_{A'}^2)^2}{(m_V^2 - m_{A'}^2)^3} \lambda^{1/2}(m_V, m_P, m_{A'}) \Gamma_{V \rightarrow P\gamma}, \quad (2.8)$$

where $\lambda(x, y, z) = [x^2 - (y + z)^2][x^2 - (y - z)^2]$ is the Källén function.

Heavy mesons

This work includes the exotic (pseudo)scalar production from $B \rightarrow K^{(*)}X$ decays and dark scalar production in $B \rightarrow K^{(*)}SS$ (both charged and neutral). As for the light mesons, the differential production yield depends only on the distribution of B mesons and the branching ratios for the decays involving the exotic particles.

The B meson distributions are again obtained using simulations of p - p collisions with PYTHIA 8.2. The PYTHIA setup of Ref. [94] is followed, i.e., only the bottom quark production hard-QCD processes are allowed in order to have sufficient statistics. Then, the final distributions are reweighted according to the bottom production cross section at the given beam energies. The value reported by PYTHIA is $\sigma_{pp} = 39.85 \times 10^9$ pb and $\sigma_{bb} = 1.866 \times 10^3$ pb for 400 GeV beam, which has already been validated in Ref. [94] using the values reported in [117]. For 120 (and 70 GeV), the values are $\sigma_{pp} = 38.54 \times 10^9$ pb and $\sigma_{bb} = 3.8$ pb ($\sigma_{pp} = 38.38 \times 10^9$ pb and $\sigma_{bb} = 1.15 \times 10^{-7}$ pb respectively), for which no measurement has been found in the existing literature for validation. In any case, for experiments operating with such small beam energies, the ALP production via B meson decays is found to be negligible. A separate simulation involving the two-body decay kinematics of $B \rightarrow K^{(*)} + a$ is performed to obtain the resulting ALP distribution as a function of E_a and

θ_a . The corresponding decay widths are given by

$$\begin{aligned}\Gamma_{B \rightarrow K+a} &= \frac{|C_{bs}|^2}{\Lambda^2} \frac{1}{64\pi m_B^3} (m_B^2 - m_K^2)^2 f_0^2(m_a) \lambda^{1/2}(m_B, m_K, m_a) \\ \Gamma_{B \rightarrow K^*+a} &= \frac{|C_{bs}|^2}{\Lambda^2} \frac{1}{64\pi m_B^3} A_0^2(m_a) \lambda^{3/2}(m_B, m_{K^*}, m_a) .\end{aligned}\tag{2.9}$$

The values for FCNC couplings from (1.59) evaluated at $\Lambda = 1$ TeV are used. Since this result is applied to a broad range of ALP masses, the following parametrization is used for the form factors [118, 119]:

$$f_0(q) = \frac{0.330}{1 - q^2/37.46} ,\tag{2.10}$$

$$A_0(q) = \frac{1.364}{1 - q^2/m_B^2} - \frac{0.990}{1 - q^2/36.78 \text{ GeV}^2} ,\tag{2.11}$$

which depend on the momentum transfer q . An example of the obtained differential yield for a fixed value of the branching ratio $\text{BR}_{B^{\pm,0} \rightarrow K^{\pm,0}+a} = \Gamma_{B \rightarrow K+a}/\Gamma_B = 10^{-10}$ is shown in figure 2.4.




Figure 2.4: Production differential yield of an $m_a = 0.5$ GeV ALP produced in decays $B^{\pm,0} \rightarrow K^{\pm,0} + a$ of B produced at a copper target with 400 GeV proton beam assuming a referential value $\text{BR}_{B \rightarrow K+a} = 10^{-10}$.

An analogous expression to (2.9) is for the Yukawa-coupled pseudoscalars

$$\begin{aligned}\Gamma_{B \rightarrow K+a} &= \frac{|C_{bs,Y}^S|^2 (m_B^2 - m_K^2)^2}{16\pi m_B^3 (m_b - m_s)^2} f_0^2(m_a) \lambda^{1/2}(m_B, m_K, m_a) \\ \Gamma_{B \rightarrow K^*+a} &= \frac{|C_{bs,Y}^P|^2}{16\pi m_B^3} \frac{1}{(m_b + m_s)^2} A_0^2(m_a) \lambda^{3/2}(m_B, m_{K^*}, m_a)\end{aligned}\quad (2.12)$$

with effective couplings $C_{bs,Y}^{S,P} = (C_{bs,Y}^R \pm C_{bs,Y}^L)/2$ given by equation (1.62).

As mentioned, also the dark scalar can be produced in FCNC $b \rightarrow s$ transitions with decay widths of $B \rightarrow K + S$ and $B \rightarrow K^* + S$ processes also given by equation (2.9) equivalent

$$\begin{aligned}\Gamma_{B \rightarrow K+S} &= \frac{|C_{bs,S}|^2 (m_B^2 - m_K^2)^2}{16\pi m_B^3 (m_b - m_s)^2} f_0^2(m_S) \lambda^{1/2}(m_B, m_K, m_S) \\ \Gamma_{B \rightarrow K^*+S} &= \frac{|C_{bs,S}|^2}{16\pi m_B^3} \frac{1}{(m_b + m_s)^2} A_0^2(m_S) \lambda^{3/2}(m_B, m_{K^*}, m_S).\end{aligned}\quad (2.13)$$

but driven by the corresponding FCNC effective coupling from equation (1.84).

As mentioned in section 1.4.2, non-zero λ_S gives rise to an effective coupling $C_{bs,SS}$ from equation (1.87). Dark scalars can then be also produced in pairs in 3-body⁶ $B \rightarrow K^*SS$ decays with differential widths [108]

$$\begin{aligned}d\Gamma_{B \rightarrow KSS} &= \frac{|C_{bs,SS}|^2 (m_B^2 - m_K^2)^2}{128\pi^3 m_B^3 (m_b - m_s)^2} f_0^2(m_{12}^2) I^{1/2}(m_{12}) dm_{12}^2 \\ d\Gamma_{B \rightarrow K^*SS} &= \frac{|C_{bs,SS}|^2 \lambda(m_B, m_{K^*}, m_{12})}{128\pi^3 m_B^3 (m_b + m_s)^2} A_0^2(m_{12}^2) I^{1/2}(m_{12}) dm_{12}^2,\end{aligned}\quad (2.14)$$

where $m_{12}^2 = (p_1 + p_2)^2$ is the invariant mass of the two dark scalars and

$$I(m_{12}) = [m_{12}^4 - 2m_{12}^2(m_B^2 + m_K^2) + (m_B^2 - m_K^2)^2] [1 - 4m_S^2/m_{12}^2]. \quad (2.15)$$

Analogously, the production of D mesons is simulated with

$$\sigma_{cc} = \begin{cases} 3.601 \times 10^6 \text{ pb}, & \text{for } E_{\text{beam}} = 400 \text{ GeV} \\ 0.518 \times 10^6 \text{ pb}, & \text{for } E_{\text{beam}} = 120 \text{ GeV} \\ 1.551 \times 10^5 \text{ pb}, & \text{for } E_{\text{beam}} = 70 \text{ GeV} \end{cases} \quad (2.16)$$

⁶The `TGenPhaseSpace` class from ROOT framework is used for handling the multi-body decay kinematics. Section 2.2.2 will give more details about the procedure.

and a subsequent decay $D \rightarrow \pi + a$ with decay width

$$\Gamma_{D \rightarrow \pi + a} = \frac{|C_{cu}|^2}{\Lambda^2} \frac{1}{64\pi m_D^3} (m_D^2 - m_\pi^2)^2 f_0^2(m_a) \lambda^{1/2}(m_D, m_\pi, m_a), \quad (2.17)$$

where $f_0(0) = 0.612$ [120] is used (see also [121]) and the couplings is again given by (1.59). To account for the q -dependence, an analogous parametrization as for B mesons is used:

$$f_0(q) = \frac{f_0(0)}{1 - q^2/m_{\text{fit}}^2} \quad (2.18)$$

For fitted $m_{\text{fit}}^2 = 6.46 \text{ GeV}^2$, the mean values of [120] are reproduced with a relative uncertainty $< 1\%$, which is sufficient for the purpose of this work since, as will be pointed out, there are potentially more significant contributions to the total uncertainty.

Despite a larger production cross-section of D mesons, for the ALP models considered in this work the branching ratio $\text{BR}_{D \rightarrow \pi + a}$ is many orders of magnitude smaller than $\text{BR}_{B \rightarrow K^{(*)} + a}$ because of the CKM suppression in (1.59). Nevertheless, this type of production was also simulated in this work since, in specific models, the up-type quark transition can be enhanced (see e.g. [122]). Note that the D meson spectra generated by PYTHIA need to be validated by comparing to the experimentally measured spectra to give a realistic estimate as was done for the light mesons and B mesons. A direct comparison of the total cross-section reported by PYTHIA in (2.16) with the experimentally measured cross-sections [123] suggests that the PYTHIA result is rather conservative.

The obtained ALP distribution from $D \rightarrow \pi + a$ decay is shown in figure 2.5. Compared to the spectra shown in figures 2.1 and 2.3, one can observe that the distributions of exotic particles produced in rare meson decays peak at even larger angles and energies. In particular, particles produced from B mesons have very high energies, making this production mode particularly promising for off-axis experiments searching for relatively heavy particles with short lifetimes.




Figure 2.5: Production differential yield of an $m_a = 0.5$ GeV ALP produced in decays $D^{\pm,0} \rightarrow \pi^{\pm,0} + a$ of D mesons produced at a copper target with 400 GeV proton beam assuming a referential value $\text{BR}_{D \rightarrow \pi+a} = 10^{-10}$.

As shown in [124], rare meson decays also present the main production mechanism for HNLs at fixed targets. The implementation of this interpretation in the ALPINIST framework is in progress. While it is a straightforward extension of the current work, it relies on the abovementioned validation of the D meson yield and implementation of the numerous production and decay channels.

2.2 FIP detection

To determine the probability that an exotic particle decay induces an observable signal, the various fixed-target experiments under consideration have been modeled using a simplified⁷ and generic MC simulation of exotic particle propagation and decays. The layout consists of a decay volume, a spectrometer, analyzing magnets, and a forward calorimeter (see schematic drawing in figure 2.6). The specific setup for each experiment is described in more detail in section 2.3.

⁷For example, not all veto detectors are considered in the simulation. Also, efficiencies of individual detectors and possible scattering or absorption of final state particles in the material of the experiment before detection are either not accounted for at all or only as an additional global inefficiency of the detector for given final states if the value is known.




Figure 2.6: Schematic view of a simplified experimental setup of fixed-target experiment used in the toy Monte-Carlo. An example of ALP production and $a \rightarrow \pi^+ \pi^- \gamma$ decay is depicted.

The outcome of the simulation is the function $E_{\text{det},if}(m_X, \Gamma_X)$ defined in eq. (1.94). Together with this work, tabulated data sets of $E_{\text{det},if}$ are provided based on a simulation with $N = 10^6$ trials on a grid of Γ_X and m_X values for each of the production and decay channels i, f , and for each of the experiments described in section 2.3. Note that the function $E_{\text{det},if}$ includes the efficiency $\epsilon_{\text{det},f}$ of the experiment for the detection of a given final state f . For past experiments, the efficiencies as determined by the experiment collaborations are used (see section 2.3), while for future experiments without any prior knowledge $\epsilon_{\text{det},f} = 1$ is used. However, it is straight-forward in this approach to redo this analysis for different values of $\epsilon_{\text{det},f}$, e.g., using directly the measured distributions, making the ALPINIST framework a powerful tool to assess the impact of the experimental design on sensitivity projections.

2.2.1 General principles of the ALP Monte Carlo simulation

The probability $p_{\text{det},f}$ that an exotic particle with given m_X, Γ_X, θ_X and E_X decaying to channel f will be detected at the experiment is a potentially complicated function of the distance l that the particle travels before decaying and the momenta of the final-state particles, which can be characterized by a set of kinematical variables a_1, \dots, a_n . The conditions imposed on the decay products by experimental cuts can be summarized by a function $\Theta_{\text{det},f}(m_X, \theta_X, E_X, l, a_1, \dots, a_n)$, which is equal to 1 if all experimental conditions are met and 0 otherwise. Since it is typically impossible to perform the integration over all of these variables, it is useful to define $p_{\text{det},f}$ instead as a sum over randomly

generated decay processes

$$p_{\text{det},f}(m_X, \Gamma_X, \theta_X, E_X) = \sum_{k=1}^N \epsilon_{\text{det},f} \frac{\Theta_{\text{det},f}(m_X, \Gamma_X, \theta_X, E_X; \{l, a_1, \dots, a_n\}_{\text{ran}}^k)}{N}, \quad (2.19)$$

where $\{l, a_1, \dots, a_n\}_{\text{ran}}^k$ denotes the k -th simulated event of the exotic particle decay.

To carry out these simulations the exotic particle decay length is first calculated in the laboratory frame $l_X = \gamma_X \beta_X / \Gamma_X$ with $\gamma_X = E_X / m_X$ and $\beta_X = p_X / m_X$ and using the distribution $p(l) = l_X^{-1} \exp(-l/l_X)$ to choose a random value of l . Alternatively, instead of sampling from the $p(l)$ distribution, one can sample from a flat distribution and re-weight the sampled events after the simulation.

The four-momenta of the second generation (SM) particles are generated based on the exotic particle four-momentum and potentially also on additional dynamics which can be added externally, e.g., the specific distributions used to simulate the ALP decays will be discussed in section 2.2.2. The same holds for other generations of particles if the previous generation particle is unstable, leading to a chain of subsequent decays, e.g., $a \rightarrow 2\rho \rightarrow 2\pi + 2\pi^0 \rightarrow 2\pi + 4\gamma$. As an approximation, unstable SM particles are decayed immediately, while relatively long-lived particles are considered detector-stable. Given the magnetic fields of the experiment, one can use these four-momenta to calculate the trajectories of the final-state particles and check the experimental conditions in order to evaluate $\Theta_{\text{det},f}$.

In practice, in order to obtain the distribution $E_{\text{det},if}(m_X, \Gamma_X)$ for given values of m_X and Γ_X , the binned distribution $\Delta^2 N_i / \Delta\theta_X \Delta E_X$ for the production mode i from the previous section is read into a θ_X - E_X histogram using ROOT [125]. At each iteration a random bin $\{\theta_X, E_X\}_{\text{ran}}^k$ is sampled according to the histogram. The exotic particle is then propagated, and subsequent decay is simulated. By summing up a sufficiently large sample N , a reliable estimate of $E_{\text{det},if}(m_X, \Gamma_X)$ is obtained in the form

$$E_{\text{det},if}(m_X, \Gamma_X) = \sum_{k=1}^N \epsilon_{\text{det},f} \frac{\Theta_{\text{det},f}(m_X; \{\theta_X, E_X, l, a_1, \dots, a_n\}_{\text{ran}}^k)}{N}. \quad (2.20)$$

Note that the right-hand side depends implicitly on Γ_X through the distribution used to obtain l .

2.2.2 Decay channels

Since fixed-target experiments can probe exotic particles of masses up to $\mathcal{O}(1)$ GeV, there are several kinematically allowed decay channels, listed in table 2.1. Up to 4-body decays are probed, accounting for the dominant contributions to the total decay width for the exotic particle masses and coupling scenarios under consideration.

	ALP	DP	DS
<i>neutral</i>	$\gamma\gamma$ $\pi^0\pi^0\pi^0$ $\pi^0\pi^0\eta^{(\prime)}$	-	-
<i>charged</i>	e^+e^- $\mu^+\mu^-$ - $\pi^+\pi^-\gamma$ $\pi^+\pi^-\pi^0$ $\rho\rho(\pi^+\pi^-\pi^0\pi^0)$ $\pi^+\pi^-\eta^{(\prime)}$ - $K^+K^-\pi^0$	e^+e^- $\mu^+\mu^-$ $\pi^+\pi^-$ - $\pi^+\pi^-\pi^0$ $\rho\rho(\pi^+\pi^-\pi^0\pi^0)$ - K^+K^- $K^+K^-\pi^0$	e^+e^- $\mu^+\mu^-$ $\pi^+\pi^-$ - - $\rho\rho(\pi^+\pi^-\pi^0\pi^0)$ - K^+K^- -

Table 2.1: Summary of exotic particles and resulting final states probed in this work.

The leptonic channels $X \rightarrow \ell\ell$ ($\ell = e, \mu$) are common to all considered exotic particles with the decay width for scalars

$$\Gamma_{S \rightarrow \ell\ell} = C_{\ell\ell, \text{eff}}^2 \frac{m_\ell^2}{8\pi} \sqrt{m_S^2 - 4m_\ell^2}, \quad (2.21)$$

and for vectors

$$\Gamma_{V \rightarrow \ell\ell} = \frac{C_{\ell\ell, \text{eff}}^2}{3} \left(1 + 2\frac{m_\ell^2}{m_V^2}\right) \sqrt{m_V^2 - 4m_\ell^2}, \quad (2.22)$$

where m_ℓ denotes the lepton mass and $C_{\ell\ell, \text{eff}}$ the effective leptonic coupling specific for each exotic particle scenario.⁸

The rest of the decay channels will be discussed in the respective sections for each exotic particle type.

⁸Note that for scalars $C_{\ell\ell, \text{eff}}$ has dimension $[\text{eV}^{-1}]$.

Axions

A di-gamma decay is a generic feature of ALPs [91, 116, 126], and it presents the most studied decay mode. This is especially due to the large amount of ultra-light ALP and cosmology searches. An ALP coupling to any SM particle often leads to a non-zero branching ratio for the $a \rightarrow \gamma\gamma$ channel unless a very precise cancellation among the couplings occurs. The corresponding decay width is simply

$$\Gamma_{a \rightarrow \gamma\gamma} = \frac{C_{\gamma\gamma, \text{eff}}^2 e^4 m_a^3}{\Lambda^2 4\pi} \quad (2.23)$$

with an effective ALP-photon coupling given by (1.72).

The leptonic channels described by equation (2.21) are most interesting for fermion- or electroweak-coupled ALPs. The effective coupling used is given by equation (26) in Ref. [91] calculated at one loop.

For ALPs coupled to gluons and quarks, such as the original axions, a plethora of decay channels open in the $\mathcal{O}(1)$ mass range. Two-body pseudoscalar ALP decays are CP-violating; therefore, the first kinematically accessible channels are three- and four-body ones with decay widths

$$\Gamma_{a \rightarrow 3P} = \frac{k}{2Sm_a} \int |\mathcal{M}_{a \rightarrow 3P}|^2 d\Phi_3; \quad (2.24)$$

$$\Gamma_{a \rightarrow 4P} = \frac{k}{2Sm_a} \int |\mathcal{M}_{a \rightarrow 4P}|^2 d\Phi_4, \quad (2.25)$$

where S is the symmetry factor, k is a correcting factor to match the experimentally measured width, and \mathcal{M} is the amplitude of the corresponding process.

While two-body decays are treated with a flat phase space, for three-body decays a more careful treatment is necessary. In a three-body decay, the invariant masses of first and second and second and third particles, m_{12} and m_{23} , can be used to define a Dalitz plot.⁹ If the transition amplitude square $|\mathcal{M}|^2$ of the process depends on the momenta of outgoing particles, the resulting Dalitz plot density is generally non-uniform. Such a momentum dependence arises in a wide range of models, for example, if the decay proceeds via a resonance, as in the case of $a \rightarrow \pi\pi\gamma$, which receives a contribution from virtual ρ meson exchange.

⁹In all the decays considered, there are at least two particles with the same mass. These are taken to calculate m_{12} .

For the simulation of hadronic decays, the framework derived in [95] using χ PT and vector meson dominance (VMD) is used.




Figure 2.7: Decay widths for ALPs coupled to gluons for $C_{GG}/\Lambda = 1 \text{ TeV}^{-1}$ calculated using matrix elements calculated in Ref. [95]. Decays into heavy meson resonances are included for the calculation of the total decay width.

These results are derived under the assumption that ALPs are coupled only to gluons, but they remain valid also for ALPs with couplings to quarks when the corresponding contributions to the mixing angles are taken into account. For models with non-universal couplings to light quarks, as considered, for example, in [127], there may be significant differences in both the individual branching ratios and the corresponding distributions of the decay products. In this case, one has to determine the correct form of the ALP field before calculating explicitly the $\mathcal{M}_{a \rightarrow \text{hadrons}}$ amplitudes derived in [95]. The resulting ALP decay widths for gluon coupling are shown in figure 2.7, for flavor-universal fermion coupling in figure 2.8.




Figure 2.8: Decay widths for ALPs coupled to fermions for $C_{ff}/\Lambda = 1 \text{ TeV}^{-1}$ calculated using the same strategy as for the gluon coupling but using full form of the kinetic mixing terms from equation (1.71).

For a given ALP mass m_a , we use $|\mathcal{M}|_{a \rightarrow \text{hadrons}}^2 / \Gamma_{a \rightarrow \text{hadrons}}$ to weight a flatly distributed Dalitz plot density which is obtained using the `TGenPhaseSpace` class of `ROOT` framework. The procedure is repeated for each m_a bin. As an example, the density for the decay $a \rightarrow \pi^+ \pi^- \eta$ is shown in figure 2.9.

Note that the total width mainly impacts the experiment sensitivity estimation, while the non-uniformity of the Dalitz plot affects the sensitivity only marginally. However, it can be a very useful analysis tool in case of signal observation, which might help to determine the precise setup of the various ALP couplings to give guidance for the search for the corresponding ALP model.

An important remark is the Dalitz plot normalization. `TGenPhaseSpace` does not sample the events uniformly, but the extreme m_{12}^2 and m_{23}^2 values are more preferred, as shown in the top left histogram of figure 2.10 on the example of $a \rightarrow \pi^+ \pi^- \gamma$ decay of 386 MeV ALP. The event weight of `TGenPhaseSpace` then gives the expected flat distribution. However, the normalization is not one as expected but rather corresponds to the square of the portion of the whole m_{12}^2, m_{23}^2 surface occupied by the kinematically allowed states as shown in the top right histogram of figure 2.10.




Figure 2.9: Dalitz plot density for six values of ALP masses used for the $a \rightarrow \pi^+ \pi^- \eta$ decay.

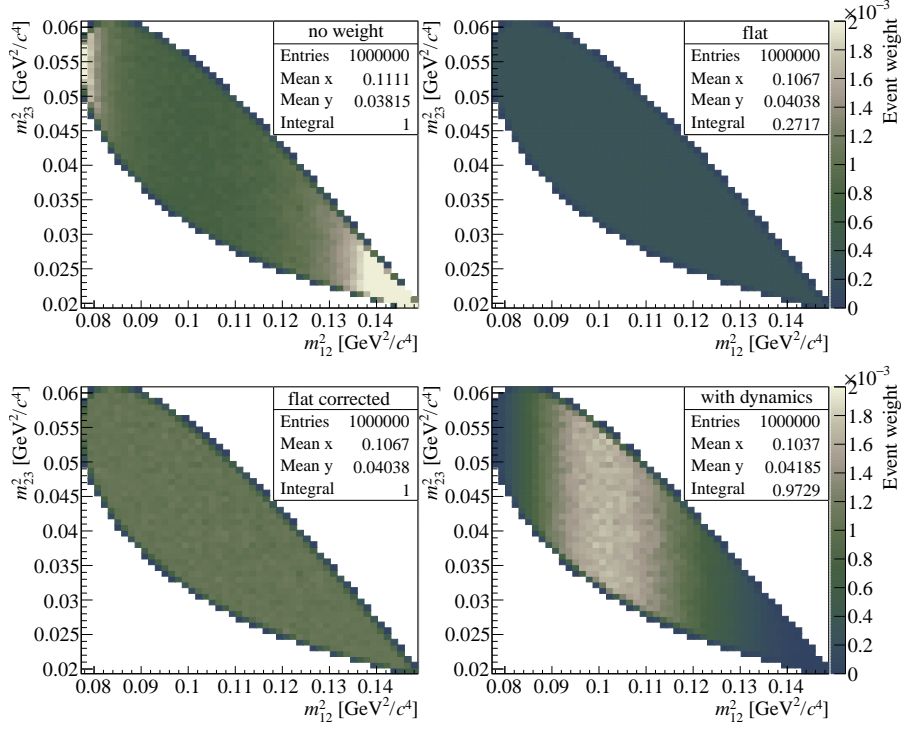


Figure 2.10: Dalitz plot for $a \rightarrow \pi^+\pi^-\gamma$ decay of 386 MeV ALP obtained using TGenPhaseSpace, shown with linear scaling for the z-axis. Top left: distribution of the sampled events. Top right: flat distribution with wrong normalization when using the TGenPhaseSpace event weights. Bottom left: corrected flat distribution. Bottom right: distribution with corrected normalization when also applying dynamics from $|\mathcal{M}|_{a \rightarrow \pi^+\pi^-\gamma}^2$ calculated for gluon-coupled ALP.

The correction is therefore first derived numerically for each final state configuration and applied on the TGenPhaseSpace weights, which results in the correct normalization as shown in the bottom left histogram of figure 2.10. On the right is the expected distribution when the abovementioned dynamics given by $|\mathcal{M}|^2$ is taken into account for the event weight calculation.

Dark photons

The decay width for hadronic decays of dark photons is given by

$$\Gamma_{A' \rightarrow \text{hadrons}} = \Gamma_{A' \rightarrow \mu\mu} \mathcal{R}_\mu^V, \quad (2.26)$$

where $\Gamma_{A' \rightarrow \mu\mu}$ is given by equation (2.22) and \mathcal{R}_μ^V by equation (2.6). The resulting widths are shown in figure 2.11. Unlike in the case of ALPs, the Dalitz plot rescaling is not applied in this simulation as the corresponding densities have not been calculated. For simplicity, all the DP decays are taken to be uniform. It has been demonstrated on the case of ALPs that such choice does not affect the sensitivity estimation for highly boosted exotic particles, such as those detected at beam-dump experiments, which are fully reconstructed.¹⁰




Figure 2.11: Decay widths for dark photons for a fixed value $\varepsilon = 1$ calculated using the \mathcal{R}_μ^V values fitted in Ref. [105].

Dark scalars

The results from [106] are used for the calculation of the hadronic decay widths, notably the $S \rightarrow \pi\pi$ and $S \rightarrow KK$. Note that for $S \rightarrow \rho\rho(4\pi)$, etc. the estimated result from [106] given by $\Gamma_{S \rightarrow \rho\rho} = C(\sin\theta)^2 m_S^2 \sqrt{m_S^2 - 4m_\pi^2}$ is oversimplifying and a more precise determination of this width should be implemented in the future. Used decay widths are plotted in figure 2.12 as a function of DS mass.

¹⁰This situation may differ for HNLs with neutrinos in the final state escaping detection. For example, suppose the dynamics prefers events with neutrinos with large momenta. In that case, the overall sensitivity estimation might be biased by assuming uniform decays due to a larger sensitivity of the event acceptance to the detector resolution.




Figure 2.12: Decay widths for dark scalars for a fixed value $\sin \theta = 1$. Partonic widths are assumed for masses above 2 GeV. The $\pi\pi$ and KK widths are calculated following Ref. [106] and a combined 4π , etc. decay width is such that the total width is continuous at 2 GeV.

2.2.3 Propagation of decay products

The photons, either emitted as daughter particles of the ALP or as decay products of SM particles, are propagated in a straight line up to the upstream face of the forward calorimeter. When extrapolating any charged particle, if present, the bending induced by the spectrometer analyzing magnets is accounted for. The corresponding magnetic field is treated as that of an ideal solenoid, and effects due to detailed field maps, stray fields, etc., are not modeled.

Photons and e^\pm are assumed to release all their energy in the calorimeter. No resolution effects are simulated, and any experimental condition is directly applied to the photon or e^\pm energy. A minimum separation between photon or e^\pm showers is required for “cluster counting”: a pair of energy releases is merged into a single “cluster” if the relative distance of the calorimeter impact points is below an experiment-dependent radius. In this case, the cluster energy is the energy sum of the two merged particles. For cluster counting, the particle impact point is required to be within the calorimeter-sensitive region, and an experiment-dependent minimum energy is required. Muons are extrapolated forward up to the front face of a muon detector if it is present. For muon detection, the muon impact point at the front face of the muon detector is required to be in the sensitive region. The presence of a passive shield in front

of a muon detector is only accounted for by applying a loose minimum-energy threshold.

It should be emphasized that decays with all neutral final state particles are considered not only in the decay volume but also in the spectrometer area, effectively enlarging the decay volume up to the calorimeter location.

2.3 Fixed-target experiments

This section briefly discusses modeling the considered experimental setups and the related acceptance conditions. Table 2.2 provides an overview of the basic parameters used for each experiment. A beam of protons with energy E_{beam} is made to interact with an absorber based on the element listed in the column “Target”. The total number of protons on target (POT) is denoted by N_{POT} . The experiment should be sensitive to ALPs produced in the target and reaching a decay volume (DV) located z_{DV} meters downstream. The longitudinal axis of the DV has an angle θ_{off} with respect to the beam axis, and Ω_{cov} denotes the solid angle covered by the calorimeter. Neutral (charged) ALP-daughter particles can be detected if the ALP decays within $l_{n(c),\text{DV}}$ meters from the DV entrance. Details are given in the following subsections.

2.3.1 Results from past experiments

To explore future opportunities, first, the results of the past CHARM and Nu-Cal experiments need to be recasted in terms of their sensitivity for exotic processes probed in this work.

CHARM

For CHARM (using a copper target), an off-axis search for two photons in the final state was presented in [128], with an average efficiency of 0.51. The search required at least one photon detected at the calorimeter. The sensitive area of the calorimeter has transverse dimensions of $3\text{ m} \times 3\text{ m}$. In the analysis quoted, the shower energy was required to be between 5 and 50 GeV. The same condition has been implemented here.

The same reference includes results for the detection of final states with two daughter muons.¹¹ The detection efficiency is reported to be 0.85. In the

¹¹This is in contrast to the simulation performed in [94], where it was mistakenly assumed that CHARM would be sensitive also to final states with just *one* muon.

Experiment	Status	E_{beam} [GeV]	N_{POT} [10^{18}]	Target	$l_{\text{n,DV}}$ [m]	$l_{\text{c,DV}}$ [m]	z_{DV} [m]	θ_{off} [mrad]	Ω_{cov} [μsr]
CHARM	completed	400	2.4	Cu	35	35	480	10	34
NuCal	completed	70	1.7	Fe	23	23	64	0	700
NA62	running	400	1	Cu	139	81	82	0	84
HIKE	proposed	400	50	Cu	139	81	82	0	84
DUNE ND	proposed	120	1100	C	10	10	575	0	36
DarkQuest	proposed	120	1.44	Fe	13.5	1	5	0	12000
SHADOWS	proposed	400	50	Cu	26	23	10	75	9200
SHiP	proposed	400	200	Mo	50	28	45	0	4500

Table 2.2: Overview of the basic parameters for the experimental setups considered in this work.

ALPINIST simulation, the muons are required to have at least 1 GeV at the final detector plane in order to be seen as minimum ionizing particles throughout the detector material.

Regarding hadronic final states, only on-axis searches are known, which go hand in hand with a larger background rate. Around 80 muon-less events are reported in Ref. [129] based on statistics of 7×10^{17} POT [130], after requiring a shower-like energy release of more than 20 GeV. According to the quoted analysis, 4 ± 21 events can be attributed to the interaction of other neutrinos or neutrino-like particles. Due to the limitation in statistics and due to the uncertainty induced by the subtraction of the background, the strongest exclusion by CHARM on hadronic decay modes is again given by a re-interpretation of the search for decays to final states with emitted photons: final states such as $\pi^+\pi^-\gamma$, $\pi^+\pi^-\pi^0$, and $\pi^+\pi^-\eta$ are considered, and at least one and at most two photons are required to be detected in the calorimeter, while all of the charged hadrons are required to escape the volume.

NuCal

The sensitivity of NuCal to final states with photons was evaluated in [131], requiring the detection of one electromagnetic shower. While NuCal operated with a 70 GeV proton beam, it benefited from a relatively small distance between target and detector (64 m) and from a comparably large decay volume length (23 m). After rejecting hadronic-like showers, NuCal observed 1 event compared to a background expectation of 0.3 events. ALP signals corresponding to a prediction of more than 3.6 events are excluded with a confidence level of 90%. For the detection of a two-track final state, the same procedure from [115] is followed, requiring both tracks to reach a circle of 2.6 m diameter at the end of the decay volume with a combined energy of at least 10 GeV. As in Ref. [115], an overall detection efficiency of 0.7 is assumed for the photonic final state and 0.8 for the di-muon final state.

By contrast, no corresponding results for hadronic final states in NuCal are known. Thus, by proceeding similarly as for CHARM, the search can be reinterpreted as a search for final states with photons as appropriate (i.e., if the final state contains at least one γ , π^0 , or η) and requiring that the charged hadrons escape detection.

2.3.2 Acceptance of current experiments

The only operational proton beam experiment that is considered in this part is NA62, which is collecting data until 2025; a successor experiment called HIKE with a similar setup is proposed for the beyond 2025 period [6]. Some first data taken in beam-dump mode also exist from the KOTO experiment [132]. In section 2.6, it will be shown that while KOTO cannot compete with past experiments in this mode, there exists an interesting possibility in the kaon mode [2]. The present KOTO setup, as well as the future update will be therefore described in details there, including the validation of the used soft QCD PYTHIA dataset.

NA62/HIKE

The primary goal of NA62 is the precise measurement of the rare decay of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$, but the experiment is also sensitive to a variety of exotics scenarios. An overview of exotic particle searches with NA62 in kaon mode will be given in section 2.5 while here, the focus will be on the beam-dump mode.

For a dump-mode run, the standard beryllium target of NA62 must be removed, and the movable collimators (TAX) located 22 m downstream of the target must be closed and used to dump the proton beam. This operation has been validated during the 2016–2018 data-taking, and about 10^{16} POT have been collected in beam-dump mode. Triggers sensitive to decay channels with neutral and charged daughter particles have been deployed. During the current 2021–2025 run about 10^{18} POT will be collected in beam-dump mode. Around 1.4×10^{17} POT have been collected in 2021 [133] and the detailed analysis of the 2021 results will be shown in chapter 4. The sensitivity projection for NA62 estimated in this chapter correspond to 10^{18} POT while the projection for HIKE correspond to 5×10^{19} POT. In line with the present knowledge based on 2021 data, there is no background limitation for given statistics [6].

NA62 is modeled in the ALPINIST toy MC as follows: the beam-defining collimator for ‘regular’ data-taking is used to dump the beam, and its distance to the start of the fiducial volume is 82 m. A decay region in a vacuum contains a spectrometer with a first station 81 m downstream of the fiducial volume entrance. The spectrometer hosts a 0.7 T solenoidal magnet 95 m downstream of the entrance of the decay volume. The Liquid Krypton Calorimeter (LKr) is located about 139 m downstream of the entrance of the decay volume, and it is modeled as an octagon [134]. A detailed description of the NA62 detector

and its various subsystems will be given in chapter 3, which is needed for the proper analysis performed in chapter 4, simplified cuts used in the ALPINIST toy MC are described below.

For the neutral final states, the following acceptance conditions are required: Both photons from the ALP decay must be detected at a minimum mutual distance of 10 cm at the LKr plane. In addition, the energy of each photon must be above 1 GeV and each photon impact point at the LKr must be at least 15 cm away from the LKr central hole through which the beam-pipe passes. Finally, the energy sum of the two photons must exceed 3 GeV. The NA62 target material (i.e., the material of the upstream section of the final collimator) is copper.

For the di-muon decay, both muons have to be in the acceptance of the NA62's first and last tracking stations. The acceptance loss due to the central hole allowing the passage of the beam pipe is modeled, cf. also Ref. [94]. Each of the muons should have at least 5 GeV in order to be tracked efficiently.

Finally, for hadronic final states that include photons, not only the charged hadrons must be tracked, but also the additional requirement is made that all photons reach the calorimeter and deposit at least 3 GeV of combined energy.

2.3.3 Acceptance of future experiments

To evaluate the sensitivity of future experiments, the analysis is based on the available proposals. Some of the considered endeavors are more defined, while others might be subject even to significant change in the future. Moreover, some of the proposed experiments may face non-negligible backgrounds, which are not estimated in this work.

SHiP

The Search for Hidden Particles experiment (SHiP) is a dedicated proposal at CERN SPS for the FIPs search. Since the envisaged geometry for SHiP has changed since the publication of the SHiP proposal [135], the layout of [136] is followed. Note that a proposal for placing the detector in a different location with slightly modified geometry has been published [137] after publishing the results of this work. The impact on the detector sensitivity is, however, negligible.

The prospects for the detection of photons from ALP decays in the SHiP calorimeter is modeled as follows. The fiducial region is taken to be 45 m downstream of the production point of the ALPs and contains a spectrometer

with the first spectrometer station located 50.76 m after the entrance of the decay volume and a 0.15 T magnet. The fourth spectrometer station is 60.16 m downstream of the entrance of the decay volume, and it is followed by the calorimeter. The distance between the spectrometer and the calorimeter inferred from the schematics is 0.5 m.

It is required that both photons of di-gamma decay are in the acceptance area of $5 \text{ m} \times 10 \text{ m}$. The energy of each photon must be above 1 GeV. The sum of the photon energies should be above 3 GeV. The photon impact points at the calorimeter should be at least 10 cm apart. The target material is molybdenum, and the POT is 2×10^{20} . Note that the proposed SHiP calorimeter has the potential of reconstructing the photon direction, which allows for an ALP mass reconstruction.

The steps provided in [138] are mimicked to model the detection of charged particles in SHiP. The decay vertex must lay in the decay volume, and the two tracks must hit all spectrometer chambers in their sensitive volume, with a minimum distance of 5 cm away from the decay vessel walls. Each particle must have a minimum energy of 5 GeV. Finally, for hadronic final states that also include photons, all photons should reach the calorimeter and deposit at least 3 GeV of combined energy.

DarkQuest

To make projections for DarkQuest, the proposed extension of the SeaQuest experiment is considered [139], particularly the proposed phase-I parameters. Details of the DarkQuest setup can be found in [140]. The target material is iron. Following [141], for the di-photon final state, the detection of 10 signal events to claim a signal beyond the background fluctuations is required.

It is assumed that the experiment is upgraded with a calorimeter placed between tracking stations 3 and 4 (at $\sim 18.5 \text{ m}$ downstream of the target). The fiducial volume entrance is 5 meters downstream of the target. An independent study of the DarkQuest sensitivity to ALPs [142] uses 7-8 m as fiducial volume (following [141], whereas the fiducial volume definition in this work follows [140]). The first spectrometer station is located about 6 m from the target, and a geometric acceptance of $2 \text{ m} \times 2 \text{ m}$ is assumed in the transverse direction. The “phase-I statistics” of 1.44×10^{18} POT is assumed.

As with the other setups, a minimum energy of 1 GeV is required for each detected photon and a total energy of at least 3 GeV and – given the photon shower Moliere radius – a minimum mutual distance of 10 cm at the calorime-

ter plane to avoid shower overlap. The abovementioned alternative study [142] has a slightly different requirement regarding the overall energy and the photon separation.

For the di-muon final state, the tracking in the simulation uses the KMAG magnet and imposes that the particles are in the acceptance of the first and third tracking stations. In addition, 1 GeV of minimum energy for each muon is required so that it is not stopped in the iron absorber.¹² Following [139], muon track pairs should be detected with a $\sim 60\%$ efficiency for di-muon masses above 4.2 GeV. Here, the possibility of detecting also hadronic final states is not considered, as the K_L background in the phase-I setup is expected to be a limiting factor [140].

DUNE ND

The results presented in Ref. [143] motivated the inclusion of the DUNE near detector [144] (ND) in this study. DUNE ND uses a neutrino beam produced at the Long-Baseline Neutrino Facility (LBNF) at Fermilab by a 120 GeV proton beam impacting on a graphite target with integrated intensity of 1.1×10^{21} POT per year [145]. Note that the nominal intensity can increase with further updates, and the target material can be changed to beryllium [145]. Since DUNE ND is located underground, separated by approximately 0.5 km of earth from the target, it can also effectively serve as an on-axis¹³ proton beam-dump experiment.

The essential parts of DUNE ND used in this analysis are the Liquid Argon detector (ND-LAr) and a multipurpose detector based on gaseous argon (ND-GAr). The ND-LAr is a 5 m long time projection chamber with height of 3 m and width of 7 m located 574 m downstream of the target. The ND-GAr located right downstream of the ND-LAr is composed of a 5 m long cylindrical high-pressure gaseous time projection chamber of 5.2 m diameter surrounded by an electromagnetic calorimeter, in a 0.5 T magnetic field and a muon system. In this analysis, the ND-LAr and ND-GAr are treated as two independent

¹²Note that this requirement goes slightly beyond that made in Ref. [140], where no minimum energy requirement is made. There, instead, only the transverse p_T kick induced by KMAG and the third tracking station are considered.

¹³The DUNE ND will also have the opportunity to take data up to 60 mrad off-axis as so-called DUNE-PRISM. This option is not considered in this work since the amount of time allocated to off-axis operation is unknown. It should be stressed, however, that since in certain scenarios, exotic particles in beam-dumps are dominantly produced off-axis, DUNE-PRISM can have a large potential for this type of hidden sector physics searches even if operated for a short period of time.

decay volumes. Following the reconstruction in [144], exotic particle decays are allowed to happen in $6\text{ m} \times 3\text{ m} \times 2\text{ m}$ volume of ND-LAr, excluding 50 cm from the sides and upstream and 150 cm downstream corresponding to roughly $5\text{ cm} \times 9\text{ cm}$ Moliere radius to allow a full reconstruction of the showers. For ND-GAr a 4 m long, 2 m high and 4.8 m wide block located 580 m from the target is assumed while excluding the outer layers following the exclusion in the active volume calculation of [144]. The photons in the 2γ decays in the ND-GAr must have a minimum 0.8 deg (14 mrad) separation angle and a minimum 20 MeV energy each. Charged hadrons from hadronic ALP decays must have a minimum kinetic energy of 5 MeV, based on the measured resolution [144]. For decays taking place in the volume of ND-LAr, no conditions on the minimum energy are applied until the expected performance of the detector based on the ProtoDUNE [146] results is known.

For muons in DUNE, following Figure 2.30 of [144], it is assumed that muons with energies below 1 GeV are reconstructed with full efficiency. Above those energies, the angle w.r.t. to the beam axis must be smaller than 40 deg (700 mrad).

SHADOWS

SHADOWS is a proposed off-axis experiment in the NA62 experimental cavern, which can take data concurrently with HIKE operated in a beam-dump mode. It would use a 400 GeV proton beam dumped in the copper NA62 TAXes, and it would collect the same statistics as HIKE operated in beam-dump mode 5×10^{19} POT. The geometry of the detector that was employed in [1] was based on the 2021 PBC proposal [147], the final layout proposed in Ref. [148] includes slight modifications which mainly constitute of the detector shielding which requires a shift of the detector downstream by 4 m. The decay volume starts 14 m downstream of the target and accommodates a 12-14 m long spectrometer¹⁴ located 34 m downstream of the target. The spectrometer has four tracking stations with a gap between the second and the third station foreseen to accommodate a 1 T dipole magnet. A calorimeter is located about 50 cm behind the fourth spectrometer station. The calorimeter sensitive area has a rectangular shape of $2.5\text{ m} \times 2.5\text{ m}$ and in our implementation, the detector is off-axis by $\theta_{\text{off}} = 75\text{ mrad}$, so that the calorimeter center is shifted by 2.25 m from the beam axis. The calorimeter is likely to be succeeded by a

¹⁴Compared to 2.5 m assumed originally in Ref. [1].

muon detector¹⁵ since muons generated in the TAXes present the main background component based on the preliminary studies in [147]. The request on muons is similar to that of NA62 and SHiP: Two spectrometer chambers should be hit as well as a minimum energy of 5 GeV for each particle.

2.4 Results

Having obtained the model-independent functions E_{if} by combining the differential feebly-interacting particle yields from section 2.1 with the various particle decay modes discussed in section 2.2 and the experimental setups considered in section 2.3, the predicted number of FIPs in a given experiment can be calculated for any model for which the model-dependent functions $M_{if} = f_i(\mathbf{C}, \mathbf{C}_{\text{ref}})\text{BR}_{X \rightarrow f}(m_X, \mathbf{C})$ are known.

In the following, the results for data obtained with the beam-dump toy MC for various FIPs will be shown.¹⁶

2.4.1 Axions

The necessary formalism has been worked out in great detail for a broad class of ALP effective theories that include interactions of ALPs with SM quarks and leptons, as well as gauge and Higgs bosons. The ALPINIST framework allows a custom setup of couplings. For comparison with the benchmark scenarios presented in Ref. [101], several cases will be shown in the following, first for ALPs coupled dominantly to SM gauge bosons and second for dominant coupling to SM fermions.

ALPs coupled to SM gauge bosons

The corresponding Lagrangian for this case is

$$\mathcal{L}_{a,\text{int}} = g'^2 \frac{C_{BB}}{\Lambda} a B^{\mu\nu} \tilde{B}_{\mu\nu} + g^2 \frac{C_{WW}}{\Lambda} a W^{\mu\nu} \tilde{W}_{\mu\nu} + g_s^2 \frac{C_{GG}}{\Lambda} a G^{\mu\nu} \tilde{G}_{\mu\nu}, \quad (2.27)$$

where g' , g and g_s denote the hypercharge, weak and strong gauge couplings respectively, $X^\mu = B^\mu, W^\mu, G^\mu$ denote the corresponding gauge fields (with the group generator index suppressed), $X^{\mu\nu}$ denotes the field strength tensor and $\tilde{X}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\rho\sigma} X_{\rho\sigma}$ denotes its dual.

¹⁵The muon detector plane is assumed to be 1 m behind the calorimeter plane in this work.

¹⁶In order to keep the results clear, the exclusions from other sources, such as cosmology, colliders, rare decays, etc., will be shown separately in the following sections.

Rather than varying all three couplings independently, four different benchmark scenarios are considered:

- (i) **B dominance:** $C_{BB} \neq 0$; $C_{WW} = C_{GG} = 0$
- (ii) **W dominance:** $C_{WW} \neq 0$; $C_{BB} = C_{GG} = 0$
- (iii) **Gluon dominance:** $C_{GG} \neq 0$; $C_{BB} = C_{WW} = 0$
- (iv) **Co-dominance:** $C_{BB} = C_{WW} = C_{GG} \neq 0$

The first three scenarios are chosen to highlight the different production channels. In particular, the first scenario is similar to the frequently studied case of *photon dominance* (see e.g. [114]), except that it includes additional interactions between ALPs and Z bosons, which are, however, of no relevance for fixed-target experiments. This roughly corresponds to the *BC9* scenario defined in Ref. [101]. The second scenario additionally predicts FCNC processes involving ALPs, effectively enhancing ALP production. The third scenario features ALP-meson mixing and decays of ALPs into hadronic final states, also denoted as *BC11* in the literature. Finally, the fourth scenario, first proposed in [116] and explored further in [143] investigates the potential interplay between the different couplings.

In order to predict ALP signals in beam-dump experiments, one first needs to obtain explicit expressions for the functions $f_i(\mathbf{C}, \mathbf{C}_{\text{ref}})$, which describe how the contributions from the different production channels depend on the fundamental interactions in eq. (2.27). The second step is to calculate the ALP decay length and branching ratios as a function of the couplings and the ALP mass. In most cases, the dominant decay mode will be the one into a pair of photons, with decays into leptons only appearing at the one-loop level and, therefore giving a negligible contribution. This conclusion changes when the direct coupling of ALPs to leptons is considered, which will be discussed later in this section. For scenarios with $C_{GG} \neq 0$ and $m_a > 3m_\pi$, there will also be a relevant contribution from three-body decays into hadronic final states.

Once the functions M_{if} have been obtained, it is straightforward to calculate the predicted number of ALP events N_{det} for each experiment as a function of the fundamental couplings and the ALP mass. For convenience, these numbers are provided in tabulated form as a part of the ALPINIST framework. In the following, the exclusion limits (projected sensitivities) for past (future) experiments will be shown at 90% confidence level.

For both scenarios (i) and (ii) one finds $\text{BR}(a \rightarrow \gamma\gamma) \approx 1$ since the one-loop processes $a \rightarrow \ell\ell$ are suppressed by many orders of magnitude. Moreover, since in both cases $C_{GG} = 0$, the production via meson mixing vanishes

($f_P = 0$), while $f_{\gamma\gamma} \neq 0$ thanks to the contribution from C_{BB} and C_{WW} to the effective ALP-photon coupling. The only difference between the two scenarios is that for (i), there are no FCNCs, and hence $f_{BK^{(*)}} = 0$, whereas they give an important contribution for (ii). Likewise, the contribution of $f_{D\pi}$ vanishes for (i) and gives a marginal contribution for (ii).

The existing exclusion limits and projected sensitivities for the two scenarios are compared in figure 2.13.¹⁷ As expected, all experiments under consideration follow a similar pattern, with the shape of the (projected) exclusions dictated by the ALP decay length (which decreases for large couplings and large ALP masses) and the production cross section (which decreases for small couplings and large ALP masses). Another notable feature in these plots is the low-mass cut-off in sensitivity for most future experiments. This cut-off results from the required separation distance between the two photons in the final state, which decreases for smaller ALP masses and correspondingly larger boost factors.




Figure 2.13: 90% CL exclusion limits for scenarios (i) (left) and (ii) (right). The exclusion from past experiments, which accounts for signal and background studies, is indicated by a full contour, while the future experiments, assuming a zero background hypothesis, are indicated by a corresponding line.

One finds that in order to extend the reach towards smaller couplings, the decisive quantity is the assumed POT, which is the largest for DUNE and SHiP. To make improvements towards larger couplings and masses, the beam energy and detector geometry are decisive. In particular, one can observe that the

¹⁷Note that the slight differences in reach compared to Ref. [114] can be mainly attributed to the inclusion of η' decays and of a wider range of meson momenta.

contribution from rare B meson decays in (ii) is especially relevant for experiments with $E_{\text{beam}} = 400 \text{ GeV}$ and clearly favors off-axis experiments like SHADOWS as well as on-axis experiments with large angular coverage like SHiP.

In scenarios (iii) and (iv), the gluon coupling is non-zero, and therefore, all production channels considered in this work become relevant. Moreover, one finds $\text{BR}(a \rightarrow \gamma\gamma) \ll 1$ for $m_a > 1 \text{ GeV}$ where hadronic decays dominate. The effects of ALP-meson mixing in both production and decay are clearly visible in figure 2.14 and imprint a ‘‘pole structure’’ in the exclusion plots whenever the ALP mass approaches one of the pseudoscalar meson masses. Note that if the two masses become very close, the condition $|\theta_{aP}| \ll 1$ no longer holds. These regions are indicated by a dark shading and masked in this analysis.

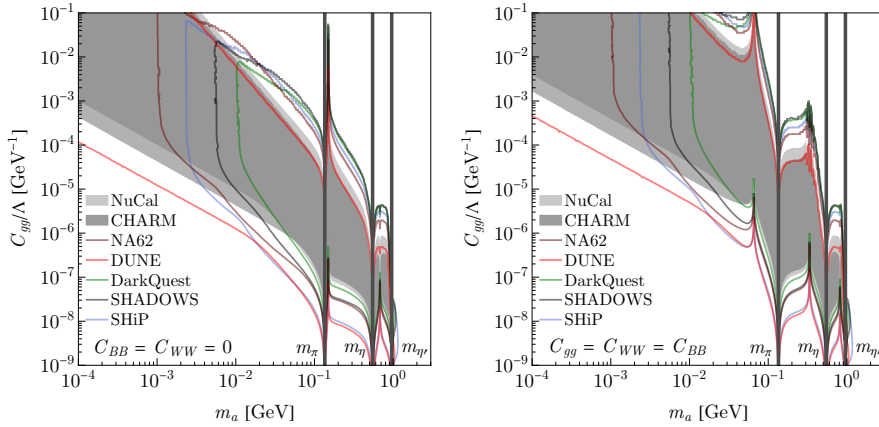


Figure 2.14: 90% CL exclusions limits for scenarios (iii) (left) and (iv) (right) assuming a di-gamma search only.

In scenario (iv) there furthermore occurs a partial cancellation in the effective ALP-photon coupling, see equations (1.54) and (1.72), which has two effects: First of all, for small ALP masses ($m_a \lesssim 100 \text{ MeV}$) it shifts all exclusion contours to somewhat larger couplings, which are necessary to achieve comparable ALP production yields. Second, it leads to additional ‘‘poles’’ outside of the masked regions whenever the different contributions to $C_{\gamma\gamma}$ cancel accidentally.

For scenarios with ALP-gluon couplings, the branching ratio into photons is strongly suppressed for heavy ALP masses and whenever there is a cancellation in the effective ALP-photon coupling. The incorporation of hadronic

decays is, therefore, essential to probe these parameter regions and utilize the full potential of beam-dump experiments for ALP searches. This can be seen in figure 2.15, which compares the exclusions and sensitivities for the di-photon channel only and the combination of all channels.

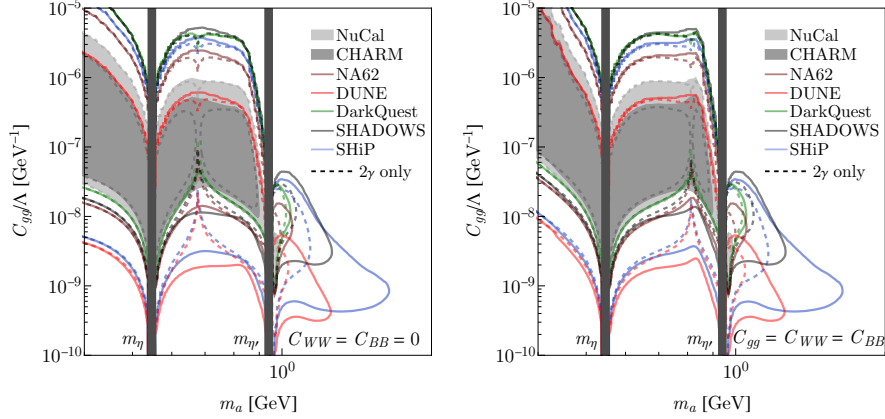


Figure 2.15: 90% CL exclusion limits for scenarios (iii) (left) and (iv) (right) given by a di-gamma decay only search (dashed line), compared with the exclusion when all mentioned hadronic channels are also taken into account (full line).

It should be emphasized that even experiments vetoing hadronic final states can partially cover regions with suppressed decays into photons thanks to hadronic decays with only photons in the final state and with charged hadrons escaping detection (provided that a partial reconstruction of the event is allowed). For experiments sensitive to a range of different final states, on the other hand, this analysis highlights the exciting possibility that beam-dump experiments may not only discover ALPs but also infer their dominant branching ratios, giving crucial clues regarding the underlying model.

ALPs coupled to SM fermions

For ALPs coupled to quarks and leptons $f = q, \ell$ at the UV scale as

$$\mathcal{L}_{a,\text{int}} \supset \sum_{f=q,\ell} \frac{C_{ff}}{2\Lambda} \partial^\mu a \bar{f} \gamma_5 \gamma_\mu f, \quad (2.28)$$

we show the results for two benchmark models:

- (v) **quark dominance:** $C_{qq} \neq 0; \quad C_{\ell\ell} = 0$
- (vi) **fermion dominance:** $C_{qq} = C_{\ell\ell} \neq 0,$

where the former induces hadronic ALP decays as in the case of gluon-coupled ALPs while the latter gives rise to the non-negligible contribution of di-lepton decays. Purely leptonic coupling, typical for leptophilic models, such as [149], can also be set within the framework. Note, however, that additional production channels, such as secondary lepton scattering in the target, which are not yet implemented in the framework, might become dominant. Similarly, ALP decays via the Bethe-Heitler process [150] are not considered in this work. The result under these considerations is shown in figure 2.16.

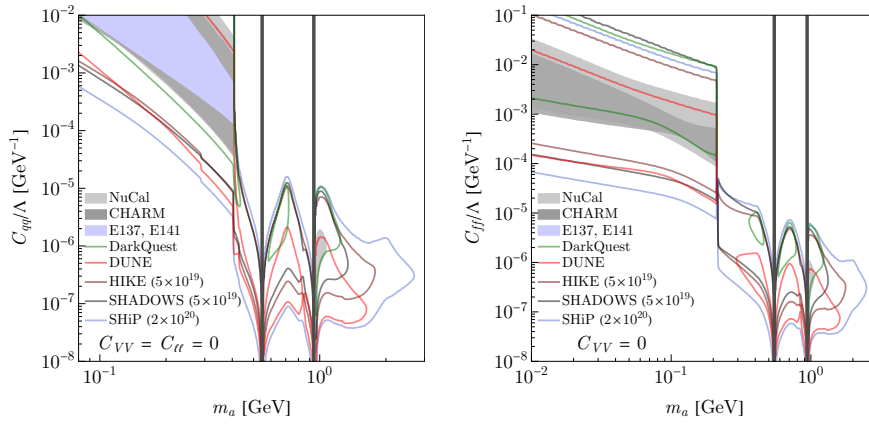


Figure 2.16: 90% CL exclusion limits for scenarios (v) (left) and (vi) (right) assuming a search for all final states from the ALP column of table 2.1. The exclusion for electron beam-dumps E137 and E141 is recasted using the ALPINIST framework using the data provided in [93].

To the author’s knowledge, this is the first estimation of the experimental sensitivity for fermion-coupled ALPs for beam-dump experiments in this mass range. A dedicated study on flavor physics bounds has been done in Ref. [151] separately for the ALP coupling to up- and down-type quarks.¹⁸ In Ref. [151], the beam-dump CHARM bound corresponds to the bound of gluon-coupled ALPs from Ref. [1] shown above under the assumption that the decay widths for fermion coupling are similar to those for the gluon-coupled ALPs. Figure 2.8 indicates that while this is true for the hadronic decay widths, the $a \rightarrow \gamma\gamma$ decay width does not dominate the total width anymore and, therefore, the CHARM bound is much weaker.

¹⁸The result of Ref. [151] is also shown here in section 2.5 in the context of ALPs produced in kaon decays.

2.4.2 Yukawa-coupled pseudoscalars

As mentioned in section 1.3, a scenario similar (but not equivalent) to (vi) are Yukawa-coupled exotic pseudoscalars, also called *BC10*,

(vii) **Yukawa coupling:** $g_Y \neq 0$.

The results for this scenario assuming production in $b \rightarrow s + a$ transitions and probing decays to leptons are shown in figure 2.17. The total decay width, including the hadronic final states calculated in Ref. [126] is used. The hadronic contribution is small in the mass range shown. While the sensitivity of the experiments goes clearly beyond the shown mass range, the total width suffers from large hadronic uncertainties at larger masses, and the sensitivity cannot be reliably estimated.




Figure 2.17: 90% CL exclusions limits for scenario (vii) assuming di-lepton final states. The exclusion bound from the LHCb searches with di-muon final state [152, 153] reinterpreted in Ref. [94] is included.

2.4.3 Dark photons

For the dark photon model given by lagrangian (1.74) the *BCI* scenario of dominant DP-SM is considered. Two results are shown in the following figure 2.18. One result considers the standard DP production in bremsstrahlung and light meson decays, while the other also includes mixing production. In the low-mass limit, the small opening angle of the $A' \rightarrow e^+e^-$ decay affects the sensitivity of experiments with a spectrometer magnet close to the calorimeter, similar to what was observed for photons in section 2.4.1 due to the cluster separation.




Figure 2.18: 90% CL exclusion limits for dark photons produced via bremsstrahlung and light meson decays (left) and when also including mixing with vector mesons (right).

2.4.4 Dark scalars

Both μ_S and λ_S dark scalar couplings are implemented in the framework and can be varied independently. The exclusions and projections corresponding to the $BC4$ and $BC5$ scenarios with fixed $\lambda_S = 5 \times 10^{-4}$ and $\lambda_S = 3.9 \times 10^{-3}$ are shown in 2.19 as a function of $\sin^2 \theta$.

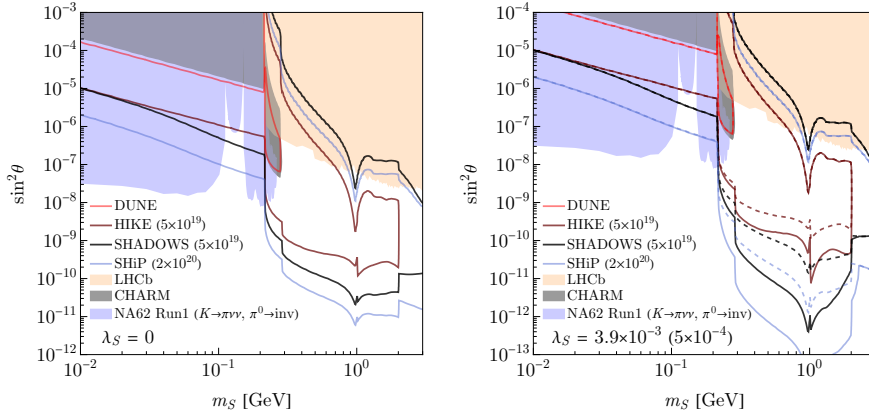


Figure 2.19: 90% CL exclusion limits for dark scalars in $BC4$ (left) and two cases of $BC5$ scenarios (right).

The values of λ_S correspond to the current upper limit at 90% CL of the $h \rightarrow SS$ BR [109] and the projection of the value used in Ref. [102]. Note that in the case of $BC5$ scenario, the restriction coming from the branching ratio on

$h \rightarrow SS$ is rather strict, and together with the phase-space suppression of the $B \rightarrow KSS$ 3-body decay, the contribution from the pair production becomes effective only at very low $\sin \theta$, when the single S branching ratio becomes very small.

The sharp drop in the sensitivity at 2 GeV is because the $\pi\pi$ and KK decay widths are conservatively calculated only up to 2 GeV while beyond 2 GeV the sensitivity is driven solely by the di-lepton final states.

The plot also includes the exclusions coming from B physics from the LHCb experiment [152, 153] and from kaon physics from the NA62 experiment, more on which will be shown in the following section.

2.5 Search for exotic particles in kaon decays

The exotic particles can be also search for in decays of relatively long-lived particles. The NA62 is a notable example of an experiment which can perform such searches in kaon decays in its kaon data-taking mode.¹⁹ In this data-taking mode, a kaon beam is formed using the beryllium target and the TAX. The kaons then decay in flight in the NA62 detector fiducial volume.

2.5.1 Search for ALPs and dark scalars in kaon decays

As in the case of B meson decays, the dark scalars or if one of the ALP-SM couplings C_{qq} , C_{GG} , C_{WW} is non-zero, the DS/ALP can be produced in an FCNC decay $K \rightarrow \pi S(S)$ or $K \rightarrow \pi a(a)$.

The main goal of NA62 is the measurement of the SM $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decay. Since the decay $K^+ \rightarrow \pi^+ X$ with X escaping detector would have the same signature as the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decay, it would appear as a peak above the continuous missing invariant mass spectrum of the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decay (see the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ spectrum in 2018 NA62 data in figure 2.20 left). The obtained limit from the full NA62 2016-2018 dataset on the $K^+ \rightarrow \pi^+ X$ BR at 90% CL as a function of exotic particle mass and lifetime is shown in figure 2.20 right.

¹⁹The author has not been directly involved in exotic particle searches using kaon data at NA62. The following two short subsections are included for completeness and to give context to the corresponding projections for the HIKE experiment.

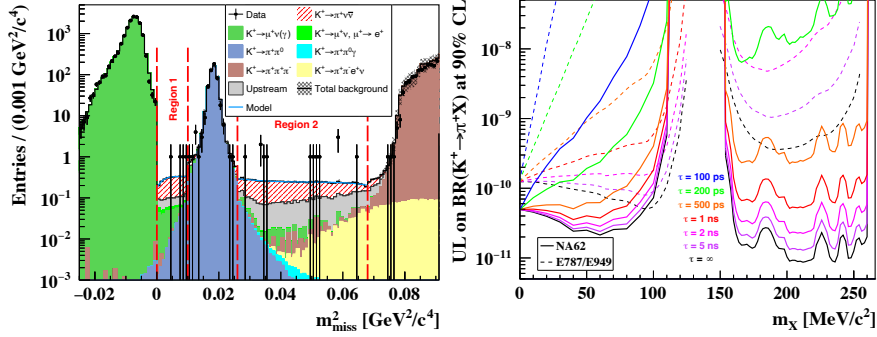


Figure 2.20: Left: Expected and observed number of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ events as a function of the reconstructed m_{miss}^2 for the 2018 data set [154]. Right: Bounds on the $K^+ \rightarrow \pi^+ X$ BR at 90% CL as a function of exotic particle mass and lifetime.

The interpretation of the limit as an exclusion in the parametric space $(m_S, \sin \theta^2)$ for the DS scenario is shown in figure 2.21 left, while for the case of an ALP coupling to SM up-quarks a parametric space (m_a, C_{uu}) is shown on the right.




Figure 2.21: Left: Bound on the dark scalar mixing angle [154]. Right: Bounds on flavor-diagonal pseudoscalar quark couplings: coupling of ALP to up quarks [151].

The results of these studies, combined with the result of the previous section for the beam-dump mode operation, are also projected to the sensitivity of the proposed HIKE experiment. The estimated sensitivity for dark scalars and gluon-coupled ALPs published in [6, 7] is shown in figure 2.22.

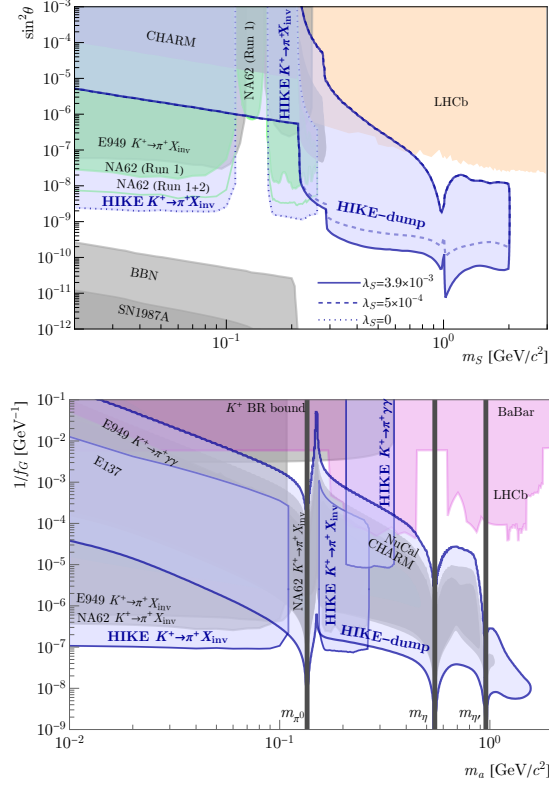


Figure 2.22: Left: HIKE sensitivity for dark scalars. Sensitivity in the dump mode is evaluated for two reference values of the λ_S parameter, $\lambda_S = 3.9 \times 10^{-3}$ and $\lambda_S = 5 \times 10^{-4}$. HIKE sensitivity in the kaon mode ($K^+ \rightarrow \pi^+ X_{inv}$) is shown for $\lambda_S = 0$, which corresponds to the scenario *BC4* from Ref. [101]. Right: HIKE sensitivity for gluon-coupled ALPs (*BC11* scenario from Ref. [101]).

A measurement of the $K^+ \rightarrow \pi^+ e^+ e^- e^+ e^-$ BR can be interpreted as decay to ALPs $K^+ \rightarrow \pi^+ a a$ with subsequent $a \rightarrow e^+ e^-$ or decay to dark scalar coupled to dark photons with a decay chain $K^+ \rightarrow \pi^+ S$, $S \rightarrow A' A'$, $A' \rightarrow e^+ e^-$ [155]. No signal is observed in the 2017-2018 NA62 data set [156] and the obtained upper limits on the BRs are shown in figure 2.23, excluding the QCD axion interpretation of the 17 MeV anomaly [157–159].




Figure 2.23: Left: Bound on the axion $\text{BR}_{K^+ \rightarrow \pi^+ a a}$ as a function of m_a . Right: Bound on the BR of the $K^+ \rightarrow \pi^+ S, S \rightarrow A' A', A' \rightarrow e^+ e^-$ decay chain as a function of the dark scalar and dark photon masses [156].

2.5.2 Search for HNLs

As explained in section 1.4.3, processes involving HNLs can be calculated from the neutrino processes by an exchange $\nu_\alpha \rightarrow U_{\alpha I} L_I$. In the minimal scenario with one HNL, one can express the branching ratio of the HNL production in a decay of a kaon K^+ to a charged lepton l^+ as

$$\text{BR}_{K^+ \rightarrow \ell^+ N} = \text{BR}_{K^+ \rightarrow \ell^+ \nu} \cdot \rho_\ell(m_N) \cdot |U_{l4}|^2, \quad (2.29)$$

where ρ_ℓ is a kinematic factor.

A search for HNL production in K^+ decays has been performed at NA62 in the 2016-2018 dataset in two independent analyses: with a muon [160] and with a positron [161] in the final state and with the HNL escaping detection. As in the $K^+ \rightarrow \pi^+ X$ analysis, the strategy is to search for a spike in the missing mass spectrum $m_{\text{miss}}^2 = (P_K - P_\ell)^2$ that would correspond to the HNL mass m_N . The scan in m_N is performed in steps of $\mathcal{O}(1)$ MeV/c^2 in mass ranges: 144–462 MeV/c^2 for the $K^+ \rightarrow e^+ N$ decay and 200–384 MeV/c^2 for the $K^+ \rightarrow \mu^+ N$ decay. The upper limits on $|U_{e4}|^2$ and $|U_{\mu4}|^2$ obtained are plotted in figure 2.24, left.




Figure 2.24: Left: UL at 90% CL on $|U_{\ell 4}|^2$ from production searches, red: $|U_{e4}|^2$, blue: $|U_{\mu 4}|^2$. Right UL on $\mathcal{B}(K^+ \rightarrow \mu^+ \nu X)$, where X is scalar or vector.

The analysis is also reinterpreted as a search for decay $K^+ \rightarrow \mu^+ \nu X$, where X is a scalar or vector. The corresponding upper limits on the branching ratios of these two cases, given the mass of particle X , are plotted in figure 2.24, right.

2.6 Search for axions with KOTO parasitically

While the primary goal of the J-PARC KOTO experiment is to measure the rare decay $K_L \rightarrow \pi^0 \nu \bar{\nu}$, as for NA62, the unique setup of KOTO raises the possibility to search for physics beyond the Standard Model [162]. In particular, in this section, the sensitivity of the current KOTO setup, as well as KOTO Step 2 for various benchmark scenarios of ALP coupling to the Standard Model gauge bosons, is probed. Like NA62, KOTO can effectively serve as a beam-dump experiment searching for NP without relying on kaon decays since NP particles can already be produced in the target. This beam-dump potential was pointed out as a possible explanation for the three candidate events in the KOTO 2019 data [163].

Following Ref. [2], the potential of the KOTO experiment to serve as a proton beam-dump for sub-GeV NP searches is studied in this section. Unlike NA62, KOTO can probe long-lived new particles in a digamma final state during its kaon physics running without needing a dedicated trigger or run, in a completely parasitic scheme; see figure 2.25. This limitation is caused by the NA62 trigger requiring the presence of a kaon in the standard data-taking.




Figure 2.25: A schematic illustration of the KOTO layout for Step 1 and Step 2. The parameters illustrating the detector size and the relative angle with respect to the primary proton beam are inscribed in the figure. The values for Step 2 are in brackets.

In the following, the KOTO experimental setup and data-taking modes are discussed in section 2.6.1. In section 2.6.2, the ALP production and decay are described, and the quantitative impact of this analysis is shown in section 2.6.3, where the bounds from current data and estimate the projection for future data-taking are derived.

2.6.1 KOTO setup and data-taking

Several past and future setups of the KOTO experiment are exploited while accounting for the available information on the experimental conditions. These setups fall into two independent categories: one regarding the experiment layout and one regarding the data-taking mode.

Two experimental layouts are considered as follows:

- Step 1: The present 2022 layout, which was also used in the 2015 data-taking [164].
- Step 2: The proposed setup for the future as described in Ref. [165].

A schematic view of the setups is found in figure 2.25. For both steps, two data-taking modes are considered:

- Kaon mode: The standard mode with a K_L beam.
- Beam-dump mode: A special run in a beam-dump mode, which includes a shield that blocks the beam (beam plug) and different selection cuts.

The experimental setups

In both setups, the experiment uses a primary 30 GeV proton beam from the J-PARC main ring. The proton beam impinges on a golden target T1 and generates a secondary hadronic beam, which, besides other particles, consists of K_L . In the present setup, Step 1 [164], the experiment axis is under a 16° angle with respect to the primary proton beam, and its front end is located 21 m from the T1 target with a set of collimators, sweeping magnets in between for forming the neutral K_L beam, and veto detectors for upstream background suppression. The CsI calorimeter (ECAL), located 27 m from the T1 target, has a 2 m diameter with a $15\text{ cm} \times 15\text{ cm}$ central hole for the beam. The decay volume is 2.9 m long and precedes the ECAL.

The proposed KOTO Step 2 setup [165] with a higher intensity beam assumes a 5° angle between the detector and the primary beam. The beginning of the decay volume is expected to be 45.75 m away from the target. The calorimeter size is increased to a 3 m diameter with a $20\text{ cm} \times 20\text{ cm}$ central hole for the beam, and it is located 64 m from the target.

In a special beam-dump mode [166] during the operation with the Step 1 setup, a beam plug was placed to close the K_L beamline, however, the sweeping magnet was not functional. The dataset of this run is smaller than in the kaon mode. It is not known if the backgrounds stated in Ref. [166] could be further reduced at the analysis level, but without a functioning sweeping magnet, a 0-background setting with the acquired data seems unlikely. With a functional sweeping system or a dedicated run with optimized magnet sweeping, one may be optimistic that a small background can be achieved similarly to NA62. The above-described experimental layouts and modes are implemented in the ALPINIST framework together with specific selection conditions described in the following section.

KOTO data-taking modes and their interpretation

While operation in the beam-dump mode can potentially allow a direct search for particles beyond the SM in a background-clean environment, the majority of the data are collected in the kaon mode, with the main aim to measure the extremely rare $K_L \rightarrow \pi^0 \nu \bar{\nu}$ decay.

In the kaon mode, searching for particles of different origin than in the K_L decay may prove to be challenging due to the backgrounds originating both from the beam and from upstream. Nevertheless, since $K_L \rightarrow \pi^0 \nu \bar{\nu}$ with a consequent $\pi^0 \rightarrow \gamma\gamma$ decay has the same signature as the axion $a \rightarrow \gamma\gamma$

decay, this provides an opportunity to reinterpret the (non)observation of the $\pi^0\nu\bar{\nu}$ signal to constrain also the axion parameter space. Unlike SM particles, axions can propagate through the beamline elements and upstream detectors, and they can decay off-axis from the neutral K_L beam. Obviously, the different kinematics of these two options would render a dedicated analysis for the case of the axion possibly more sensitive than what is stated below.

For the reinterpretation of the $K_L \rightarrow \pi^0\nu\bar{\nu}$ search as an axion search, the $a \rightarrow \gamma\gamma$ decays simulated and reanalyzed using the ALPINIST toy MC, assuming the $K_L \rightarrow \pi^0\nu\bar{\nu}$ event reconstruction and selection conditions for Step 1 analysis of 2015 data [164], for the future Step 1 run [167] and for Step 2 [165]. These are summarized in Table 2.3. In particular, cuts on the following kinematic variables are implemented: the photon cluster coordinates in the plane perpendicular to the K_L beam axis, $x_{\gamma_{1,2}}$ and $y_{\gamma_{1,2}}$; and r , which is the separation distance between the photon clusters. R_{COE} is the center-of-energy-deposited distance from the beam, based on the photon position at the calorimeter $(x_{\gamma_{1,2}}, y_{\gamma_{1,2}})$ and the final photon energies, $E_{\gamma_{1,2}}$. The photon separation angle projection on the calorimeter plane and the angle between the beam axis and the $\pi\nu\nu$ -hypothesis-reconstructed photon momenta are denoted as $\theta_{\gamma,\text{calo}}$ and $\theta_{\gamma,\text{beam}}$, respectively. The z_{vtx} position is calculated assuming that an on-axis $\pi^0 \rightarrow \gamma\gamma$ decay is being reconstructed, where $z_{\text{vtx}} = 0$ corresponds to $z = 21$ m from the T1 target for Step 1 and 44 m for Step 2. Finally, \mathcal{A}_{add} , the selection efficiency of additional shape-related cuts (cluster shape, pulse shape, and shower depth), is quoted. Since more detailed information about \mathcal{A}_{add} is not publicly accessible, a uniform distribution over the whole signal region is assumed in this work, using the numbers quoted in Refs. [164, 165]. As can be seen from Table 2.3, the main difference between the current Step 1 data and the future planned run is the collected statistics in terms of the number of protons on target N_{pOT} and the selection efficiency of the shape-related cut algorithms which has improved considerably while keeping a good rejection power for hadronic backgrounds.

In the search for $a \rightarrow \gamma\gamma$ signal events, one is not limited to the $K_L \rightarrow \pi^0\nu\bar{\nu}$ signal region since the $a \rightarrow \gamma\gamma$ events have different kinematics; justification will be given in section 2.6.2. Therefore, the expected N_{exp} and the observed N_{obs} number of SM events has to be estimated in the whole $z_{\text{vtx}}-p_{T,\pi^0}$ plane for the various KOTO datasets. For the Step 1 2015 dataset, estimations are made in all regions in the $z_{\text{vtx}}-p_{T,\pi^0}$ plane, using the N_{exp} and N_{obs} shown in figure 3 of Ref. [164]. In addition, the sensitivity of the region of $p_{T,\pi^0} > 0.5$ GeV that is out of the range of the referential figure is estimated.

Since the pion transverse momentum is expected to be smaller than 0.5 GeV for most of the known physics processes [168] and it is below the kinematic threshold for the production in the K_L beam, it is assumed that there is no SM background in this region, i.e., $N_{\text{exp}} = 0$.

	Step 1	Step 2
$\sqrt{x_{\gamma_{1,2}}^2 + y_{\gamma_{1,2}}^2}$	< 0.85 m	< 1.35 m
$\min(x_{\gamma_{1,2}} , y_{\gamma_{1,2}})$	> 0.15 m	> 0.175 m
R_{COE}	> 0.2 m	–
r	> 0.3 m	> 0.3 m
$\theta_{\gamma, \text{calo}}$	< 150°	< 150°
$E_{\gamma_1} + E_{\gamma_2}$	> 0.65 GeV	> 0.5 GeV
$E_{\gamma_{1,2}}$	$\in [0.1, 2.0]$ GeV	> 0.1 GeV
$E_{\gamma_1}/E_{\gamma_2}$	> 0.2	–
z_{vtx}	$\in [2.9, 6.0]$ m	$\in [1.75, 15]$ m
$E_{\gamma}\theta_{\gamma, \text{beam}}$	> 2.5° GeV	–
\mathcal{A}_{add}	0.52 (0.9)	0.73
$N_{\text{PoT}} \times 10^{19}$	2.2 (14)	60

Table 2.3: The selection conditions for $K_L \rightarrow \pi^0 \nu \bar{\nu}$ for Step 1 [164] and Step 2 [165]. The N_{PoT} and shape-related cut efficiency (in parenthesis) are for the future Step 1 run [167].

In order to project the KOTO sensitivity for the future Step 1 dataset, the N_{exp} backgrounds are estimated in the various $z_{\text{vtx}}-p_{T, \pi^0}$ regions which were presented in Ref. [169] for the $N_{\text{PoT}} = 3.05 \times 10^{19}$ statistics and rescaled to the expected statistics $N_{\text{PoT}} = 14 \times 10^{19}$ [167]. When excluding the $\pi \nu \bar{\nu}$ and the surrounding region ($z_{\text{vtx}, \pi \nu \bar{\nu}} < 5.1$ m, and $p_{T, \pi^0} < 0.26$ GeV), where the number of background events is large compared to the expected number of $a \rightarrow \gamma \gamma$ events,²⁰ one gets $N_{\text{exp}} = 3.35$ events with $N_{\text{PoT}} = 14 \times 10^{19}$.

For KOTO Step 2, the information from Ref. [168] is used for the background estimation in the $1.75 \text{ m} < z_{\text{vtx}} < 15 \text{ m}$ and $p_{T, \pi^0} > 0.4 \text{ GeV}$ regions. Approximately $N_{\text{exp}} \approx 1.38$ events are found with $N_{\text{PoT}} = 6 \times 10^{20}$ statistics (assuming again that there is no background in the $p_{T, \pi^0} > 0.5 \text{ GeV}$ region for $z_{\text{vtx}} < 15 \text{ m}$).

In the beam-dump mode, so far, KOTO has collected data corresponding to $N_{\text{PoT}} = 2.2 \times 10^{17}$ [166], which is about two orders of magnitude less than in the kaon mode. For the projection of the beam-dump mode, it is assumed

²⁰The SM $K_L \rightarrow \pi^0 \nu \bar{\nu}$ is considered to be a background for the search for $a \rightarrow \gamma \gamma$ decay.

that 10 times more data can be collected in this mode (i.e., $N_{\text{PoT}} = 2.2 \times 10^{18}$) while keeping the background under control, i.e., a background-free search is considered. This case is explored for both KOTO Step 1 and KOTO Step 2 layouts,²¹ assuming $\mathcal{A}_{\text{add}} = 1$. The selection conditions are simply both photons being in the calorimeter acceptance with cluster distance > 0.3 m (as used in the $\pi^0 \nu \bar{\nu}$ analysis). For KOTO Step 1 dump-mode, it is required that at least 50 MeV is deposited on the calorimeter per photon; and for KOTO Step 2, at least 100 MeV per photon and at least 500 MeV in total deposited on the calorimeter.

2.6.2 Axion production and detection

The conventional axion production at the K_L experiments is from $K_L \rightarrow \pi^0 a$ decay, which is CP -violating and thereby suppressed. Here, the relevant production to probe the long-lived axion occurs at the fixed target T1 where K_L is produced i.e., in the proton-gold collisions.

Axion production in the target

The Primakoff production and axion-meson mixing, described before for the case of other beam-dump experiments, are probed here. Additional production could be from flavor-changing kaon decays near the fixed target. With the gluon coupling C_{GG} , the $K^\pm (K_S) \rightarrow \pi^\pm (\pi^0) a$ is not suppressed by a loop or CP violation [81, 151, 170]. The C_{WW} coupling induces the same processes at one loop [79]. These production rates could be sizable because the total width of kaons is small; i.e., the BR is enhanced. However, most kaons are removed by the collimators or deflected by the magnetic field. Including these effects, the $K^+ \rightarrow \pi^+ a$ contribution is subdominant. Still, the production from K_S is potentially interesting due to the shorter lifetime, but it requires a simulation of K_S transport, which is beyond the scope of this work; therefore, the ALP production from K^+ and K_S decays is neglected, and the ALP yield at T1 target is approximated to be

$$N_a \approx N_a^{\text{mix}} + N_a^{\text{Prim}}, \quad (2.30)$$

²¹The choice of \mathcal{A}_{add} is because not enough information is available in Ref. [166] to derive it. Also, it is later shown that even with this optimistic choice of \mathcal{A}_{add} , the sensitivity at the dump mode is not competitive with that of the kaon mode.

where the exact relative contributions of the two mechanisms depend on the specific model.

In the following, the axion production yields for these mechanisms are shown for the 30 GeV proton beam, including their validation with the detector under the 16° angle. As shown in Ref. [114], the yield and the momentum spectrum are well described by PYTHIA 8 also for angles smaller than 16° and higher beam energy.

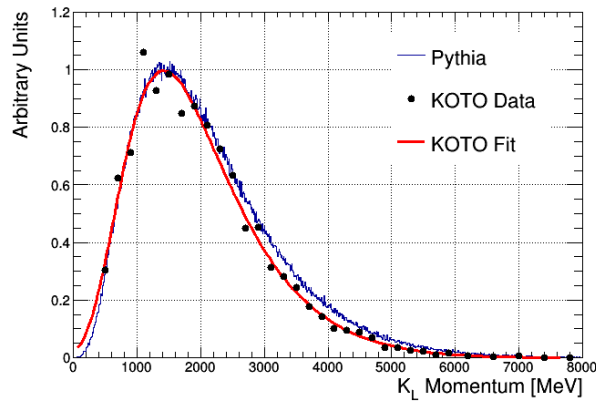


Figure 2.26: Comparison of the K_L total momentum distribution from the simulation used in this work (PYTHIA 8) and the data derived by KOTO [171] (KOTO Data). The red curve presents a fit of the data done by KOTO (KOTO Fit).

For the estimation of the axion flux from both the Primakoff and the mixing production, the flux of the π^0 , η , and η' mesons has to be estimated first. For this purpose, PYTHIA 8 [111] is used for the generation of the pp interactions. Since the meson flux simulation cannot be validated directly as the detector is far from the target, the measured K_L flux at KOTO [171] is used instead to normalize the meson yields. The measured number of K_L collimated into $8 \times 8 \text{ cm}^2$ [164] at the end of a 20 m beamline is $(4.2 \pm 0.02_{\text{stat}} \pm 0.06_{\text{sys}}) \times 10^7$ K_L per 2×10^{14} protons on target, which corresponds to the measured number of K_L per proton on target $N_{K_L}^{\text{KOTO}} \sim 2.1 \times 10^{-7}$. Therefore, the number of mesons simulated by PYTHIA 8 is normalized to the measured number as

$$N_P = \frac{N_P^{\text{sim}}}{N_{K_L}^{\text{sim}}} \times N_{K_L}^{\text{KOTO}}, \quad (2.31)$$

where N_P^{sim} is the number of (π^0, η, η') mesons produced per pp interaction in the simulation. For 10^8 simulated pp interactions, after accounting for K_L decays assuming the peak K_L momentum 1.4 GeV [164] and 60% loss of K_L due to absorption in the beamline material [171], the obtained number of K_L 's simulated per pp interaction is $N_{K_L}^{\text{sim}} \sim 2.0 \times 10^{-7}$. Therefore, $N_{K_L}^{\text{KOTO}}/N_{K_L}^{\text{sim}} \sim 1.05$, showing a good agreement between the total number of K_L 's measured and simulated using PYTHIA 8. Based on the simulation the following multiplicities of (π^0, η, η') mesons are obtained:

$$\frac{N_{\pi^0}^{\text{sim}}}{N_{K_L}^{\text{sim}}} \sim 21, \quad \frac{N_{\eta}^{\text{sim}}}{N_{K_L}^{\text{sim}}} \sim 2.2, \quad \frac{N_{\eta'}^{\text{sim}}}{N_{K_L}^{\text{sim}}} \sim 0.17. \quad (2.32)$$

Furthermore, since ALPs interact rarely with ordinary matter, the surface of the axion flux potentially entering the detector $S_{\text{axion}}^{\text{sim}}$ is much larger, occupying the whole decay volume plane. When compared to the K_L flux, which is collimated to the $S_{K_L}^{\text{KOTO}} = 8 \times 8 \text{ cm}^2$ profile at the end of the beamline, the ratio of the two surfaces is about $S_{\text{axion}}^{\text{sim}}/S_{K_L}^{\text{KOTO}} \sim 490$.

Finally, in order to validate the kinematic distributions, the distributions of K_L obtained with PYTHIA 8 are compared with the distributions measured by KOTO [171]. As shown in figure 2.26, a good agreement is observed between the shapes of the distribution of the K_L total momentum from the simulation used in this paper and the data measured by KOTO. This validation for K_L gives some credibility to the simulated (π^0, η, η') distributions at 30 GeV, and therefore to the validity of the expected axion distributions which are used in this work.

The obtained ALP differential yield distributions with respect to the axion energy (E_a) and its production angle to the proton beam (θ_a , see figure 2.25) are shown in figure 2.27 for two benchmarks with axion masses of 40 MeV and 400 MeV and the couplings being dominated by the gluon, $f_a/C_{GG} = 200 \text{ GeV}$ and 20 TeV, respectively. In this coupling benchmark for the $m_a = 400 \text{ MeV}$ case, the overall production yield from the mixing with π^0 and η exceeds significantly the Primakoff production, while for smaller axion masses, the Primakoff production becomes more relevant, similarly to what has been observed for the experiments operating with higher beam energies described in section 2.1.




Figure 2.27: Expected axion production yield N_a at the target as a function of axion energy, E_a , and its angle to the incident proton beam for two cases of axion masses and couplings. The respective angles at which the detector is placed in Step 1 and Step 2 are shown by solid lines with dashed lines indicating the part of the distribution of interest given the ECAL edges.

Axion detection mechanism

Given the beam energy and the distance between the axion production location and the detector, the KOTO experiment can search for long-lived ALPs for which, at the relevant masses, the dominant decay channel is $a \rightarrow \gamma\gamma$. For an axion heavier than 1 GeV with non-zero C_{GG} coupling, the width of hadronic decay modes, such as $a \rightarrow \pi\pi\eta$, dominates the total width, resulting in reduced sensitivity, as will be discussed later on.

Long-lived ALPs produced at the fixed target can reach the distant decay volume, and a diphoton decay leaves a characteristic signal; see a schematic picture of the axion decay event in Fig. 2.28. When axions enter the decay volume, they are almost parallel to the K_L beam axis because the distance to the ECAL is larger than the ECAL size but away from the axis with the distance ρ_a . The two photons from an $a \rightarrow \gamma\gamma$ decay in the decay volume can then hit the ECAL, mimicking the signal of $K_L \rightarrow \pi^0\nu\bar{\nu}$. If the standard reconstruction algorithm for $K_L \rightarrow \pi^0\nu\bar{\nu}$ is applied, the reconstructed position (dz) will be different from the true distance between the axion decay point and the ECAL (dz_a), but the event is not discarded. In this sense, this signal is similar to the halo $K_L \rightarrow \gamma\gamma$ background that the KOTO Collaboration found in the earlier data, but the reconstructed distributions are typically different. See figure 29 of Ref. [168] for the distribution.




Figure 2.28: A schematic picture of the axion event in the kaon mode. The disk represents the ECAL, and the axis is the K_L beamline. The distance from the ECAL to the upstream is dz . A typical axion trajectory is almost parallel to the beam axis in the distance ρ_a (see Fig. 2.25). If the final-state photons from the axion decay at dz_a leave the energy of $E_{1,2}$ with a separation of D_{12} on the ECAL, the vertex position dz is reconstructed on the beam axis.

Contrary to the kaon mode, the detection mechanism in the beam-dump mode is relatively straightforward, as it is a dedicated run to search for long-lived particles decaying to photons.

In the kaon mode, the $a \rightarrow \gamma\gamma$ decay passes the event selections described in section 2.6.1 when $\rho_a > 0$, which introduces factitious transverse momentum of the diphoton system. Some selection criteria are universal, often limited by the experimental resolution, but the remaining cuts assume the topology of $K_L \rightarrow \pi^0(\rightarrow \gamma\gamma)\nu\bar{\nu}$. The KOTO experiment could implement a dedicated analysis for the long-lived particles, but here the standard $K_L \rightarrow \pi^0\nu\bar{\nu}$ analysis is adopted because the backgrounds are well investigated. Assuming $K_L \rightarrow \pi^0\nu\bar{\nu}$ topology for the long-lived axion decays leads to several nontrivial characteristics of the axion signal event distribution N_{sig} in the $z_{\text{vtx}}-p_{T,\pi^0}$ plane. Example distributions for parameters for which we expect interesting sensitivities of Step 2 are shown in figure 2.29. In the following, analytic understandings of the characteristics based on several simplifications are given.




Figure 2.29: Axion signal event distributions N_{sig} as could be found in the $\pi\nu\bar{\nu}$ analysis at KOTO Step 2 for specific axion models. For the analytic understanding of the distributions, the dashed magenta lines corresponding to equation (2.36) with $\rho_a = 0.175, 1.35$ m and a specific E_a are shown.

Given the distance from the target, it can be assumed that the axion enters the decay volume in parallel to the beamline with distance ρ_a . The axion invariant mass is given by the photon energies, $E_{1,2}$, and the opening angle, θ_{12} ,

$$m_a^2 = 2E_1E_2(1 - \cos\theta_{12}) \simeq E_1E_2\theta_{12}^2. \quad (2.33)$$

Using this, the separation of the two photons on the ECAL, D_{12} , is approximately

$$D_{12} \simeq dz_a\theta_{12} \simeq dz_a \frac{m_a}{\sqrt{E_1E_2}}. \quad (2.34)$$

Then, the *reconstructed* vertex position, assuming the event topology of $K_L \rightarrow \pi^0\nu\bar{\nu}$, is given by

$$dz \equiv D_{12} \frac{\sqrt{E_1E_2}}{m_{\pi^0}} \simeq dz_a \frac{m_a}{m_{\pi^0}}. \quad (2.35)$$

The distance from the ECAL (dz) is translated to the standard coordinate system by $z_{\text{vtx}} = L - dz$ where $L = 6$ (20) m in Step 1 (Step 2). Because only the axion decays in the decay volume are accepted, there exists a limitation of $dz_a < dz_a^{\text{max}} = 2.9$ (18.25) m in Step 1 (Step 2). This limitation, together with Eq. (2.35), gives a condition on the maximum spread of the distribution over the reconstructed vertex position as $z_{\text{vtx}} > L - dz_a^{\text{max}}(m_a/m_{\pi^0})$. The boundary $z_{\text{vtx}}^{\text{min}}$ is shown as a blue vertical line in the left panel of Fig. 2.29.

Another feature can be seen as a correlation with both $p_T^{\pi^0}$ and z_{vtx} . In the case of $K_L \rightarrow \pi^0 \nu \bar{\nu}$, the transverse kick is from K_L decay, and hence, $p_T^{\pi^0} < m_{K_L}$. However, the transverse asymmetry of the ECAL hits is merely from the transverse position of the incident axion. Supposing the distance from the axion to the beamline is ρ_a , the reconstructed $p_T^{\pi^0}$ is roughly

$$p_T^{\pi^0} \simeq \frac{E_a \rho_a}{\sqrt{\rho_a^2 + dz^2}} \simeq \frac{E_a}{\sqrt{1 + (L - z_{\text{vtx}})^2 / \rho_a^2}}. \quad (2.36)$$

where $E_a = E_1 + E_2$ is the axion energy. The correlation between $p_T^{\pi^0}$ and z_{vtx} is explained by the above formula. Note that $p_T^{\pi^0}$ can easily exceed m_{K_L} , because $E_a \sim \mathcal{O}(1)$ GeV.

The last feature that can be observed is the vanishing of the distribution at large $p_T^{\pi^0}$ and z_{vtx} , which is a consequence of the two-photon separation cut, $D_{12} \leq 0.3$ m. This can be understood by combining Eqs. (2.34), (2.35), and (2.36) with a simplification of $E_{1,2} \approx E_a/2$ as

$$p_T^{\pi^0} \lesssim \frac{2m_{\pi^0}}{0.3 \text{ m} \sqrt{(L - z_{\text{vtx}})^{-2} + \rho_a^{-2}}}. \quad (2.37)$$

Therefore, the distribution between the two dashed lines from Eq. (2.36) vanishes at large $p_T^{\pi^0}$ and z_{vtx} at an approximate bound corresponding to Eq. (2.37) with $\rho_a = 1.35$ m.

Since $p_T^{\pi^0}$ of the long-lived axion events can be significantly larger compared to the K_L events or background, the high $p_T^{\pi^0}$ region is used.

What concerns the searches for hadronic decays of gluon-coupled ALPs, charged decays are rejected, but KOTO could be in principle sensitive to decays $a \rightarrow \pi^0 \pi^0 \eta$ or $a \rightarrow 3\pi^0$ with a subsequent $\pi^0(\eta) \rightarrow \gamma\gamma$ decay resulting in a six-cluster event. While studying hadronic decays in the kaon mode would require a dedicated analysis to address the various backgrounds, in the case of beam-dump mode, it is stated in Ref. [166] that no 6-cluster events have been found in the collected sample, indicating that the K_L background is kept under control in this case.

After running a simulation with ALPINIST for $a \rightarrow 3\pi^0$ and $a \rightarrow \pi^0 \pi^0 \eta$ decays with simple selection criteria on minimal cluster energy and cluster separation as for the $a \rightarrow \gamma\gamma$ decay in the beam-dump mode, the KOTO sensitivity with hadronic decays was not found to surpass the sensitivity using $a \rightarrow \gamma\gamma$ even for larger axion masses. As can be seen in Fig. 2.30 for KOTO Step 2, the difference in the number of observable signal events N_{sig} is by

several orders of magnitude smaller for a combined search for $a \rightarrow 3\pi^0$ and $a \rightarrow \pi^0\pi^0\eta$ decays compared to $a \rightarrow \gamma\gamma$.

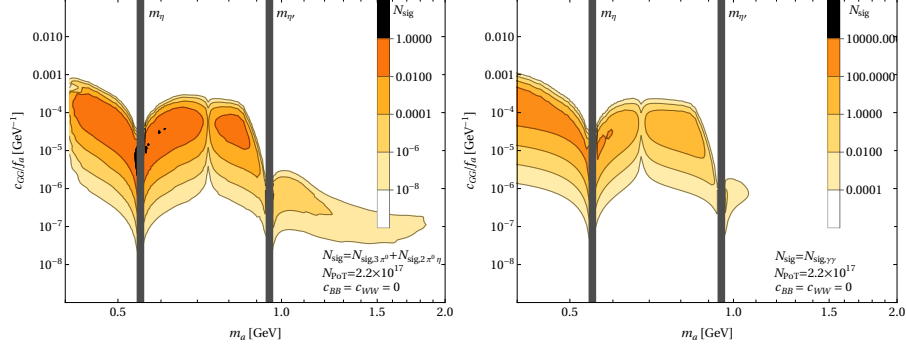


Figure 2.30: Expected number of signal events N_{sig} for given values of axion mass m_a and coupling C_{GG} for scenario (i) with $N_{\text{PoT}} = 2.2 \times 10^{17}$ in KOTO Step 2 in beam-dump mode. Comparison between N_{sig} from $a \rightarrow \pi^0\pi^0\eta$ and $a \rightarrow 3\pi^0$ decays (left) and that from $a \rightarrow \gamma\gamma$ decay (right).

2.6.3 Bounds and projections for ALPs

The N_{sig} distributions obtained using the simulation²² were reinterpreted using the ALPINIST ALP_{rescale} module for the i, iii and iv ALP benchmark models from section 2.4. The resulting current bounds and future projections are shown in the left panel of figure 2.31. In the right panel are shown the results for several fixed masses for variable C_{GG} vs C_{BB} couplings, which is expected to be similar to the case of C_{GG} vs C_{WW} , up to the FCNC production.

The obtained KOTO bounds and projections are compared to the existing bounds from different experiments. The electron beam-dumps E137 [172] and E141 [173] are considered, where the ALPINIST framework is used for the interpretation of the data provided in Ref. [174], implementation which has been already done for Ref. [6]. The proton beam-dump bounds based on CHARM [128] and NuCal [131] (bounds in gray shade) come from results of section 2.4.

As the ALP can be produced by flavor-changing meson decays, especially in the presence of C_{WW} and C_{GG} (the corresponding bounds are shown in the blue shade), the results of $K^+ \rightarrow \pi^+ a$ from Ref. [170] are applied to obtain

²²For increasing the precision of the estimated number of observed events, each mass and decay width bin is evaluated with 2×10^6 ALP decay events.

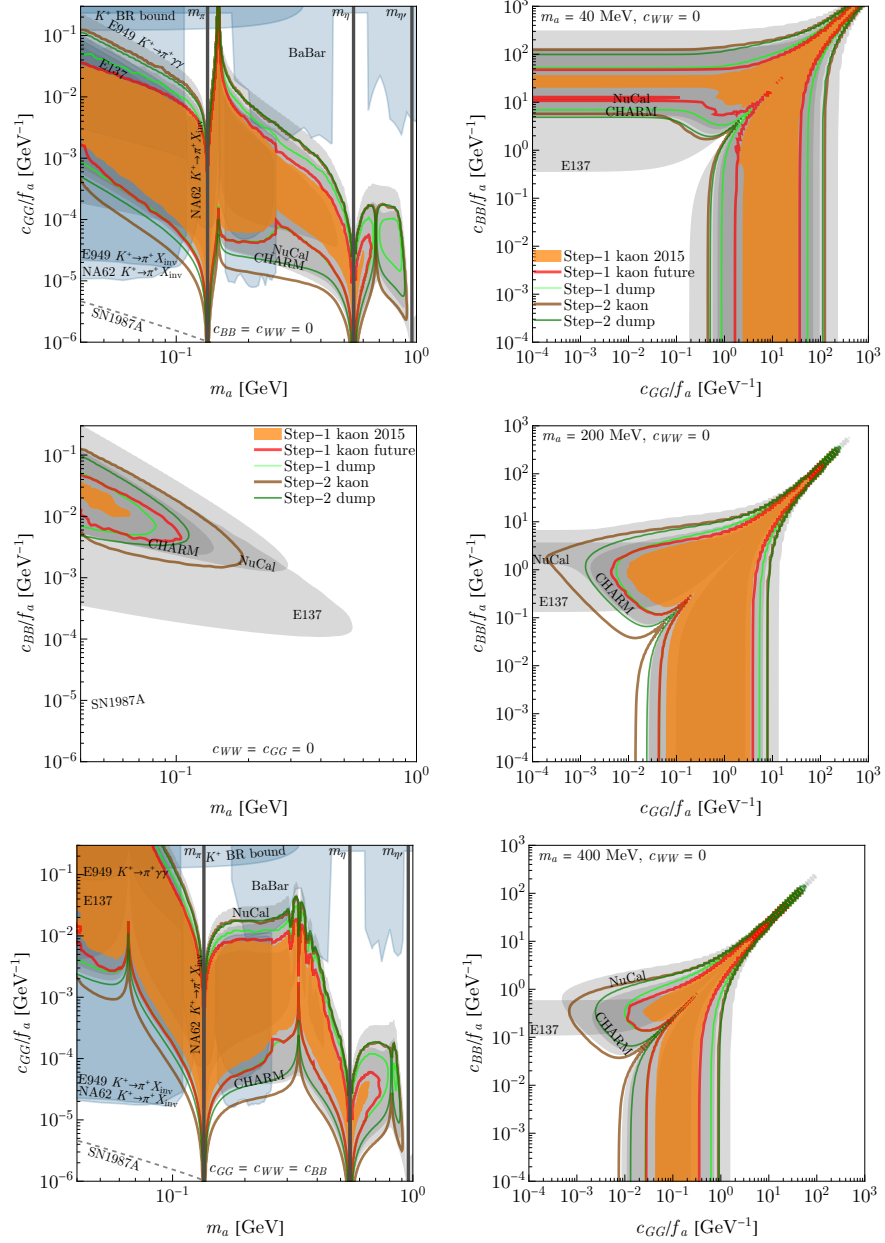


Figure 2.31: Left panel: 90% C.L. exclusion bounds (filled contours) and projected limits (empty contours) for scenarios (i)-(iii) for all KOTO setups considered compared to the exclusions from beam-dump experiments, exclusion from $B \rightarrow Ka$ and $K \rightarrow \pi a$ decays and the bound from the supernova SN1987A (shown as a dashed line as it is affected by significant uncertainties see, e.g., Refs. [116, 175]). Right panel: 90% C.L. exclusion bounds and projected limits for fixed axion mass and variable c_{BB} and c_{GG} couplings.

the bounds of E949 [176, 177] and NA62 [154, 178] whose upper limits on BRs were shown in section 2.5. The scheme of the recast is found in Ref. [162]. For $B \rightarrow Ka(\rightarrow \gamma\gamma)$, the BABAR bound can be used [179, 180]. The two-loop production calculation with C_{GG} is found in Ref. [80]. The contribution from C_{WW} at one loop is calculated in [79], but it is numerically subdominant for the benchmark iv, so for simplicity, $B \rightarrow K + a$ is approximated by the C_{GG} contribution. The total width K^+ would be modified significantly for low f_a/C_{GG} , which leads to a relevant bound at $m_a \sim m_{\pi^0}$. Requiring $\text{BR}(K^+ \rightarrow \pi^+ a) < 3 \times 10^{-3}$, based on Sec. 2.2.2 of Ref. [151] results in the bound $f_a/C_{GG} \lesssim 5 \text{ GeV}$. The C_{BB} -only scenario is not significantly constrained by the meson decays since the production originates from electroweak two-loop diagrams. Therefore, in Fig. 2.31, the meson decay bounds are omitted because the corresponding bounds in the limit of the benchmark i are unknown. Finally, the shown SN1987A bounds are those derived in Ref. [116] but are plotted with a dashed line as their robustness is under debate [175].

It is found that at Step 1, KOTO cannot probe new regions in the parameter space, and it is sensitive only in regions already covered by other proton beam-dump experiments for all coupling scenarios. The *non*-observation of additional signal on top of the expected backgrounds in the past KOTO $K_L \rightarrow \pi^0 \nu \bar{\nu}$ analyses only confirms the results of these past experiments. While for scenarios with photon coupling domination in the future, KOTO Step 2 cannot compete with the past electron beam-dump experiments E137 and E141 either, it can probe new regions of parameter space for larger masses ($m_a \gtrsim m_{\pi^0}$) for scenarios with gluonic coupling thanks to enhanced axion production through mixing with other neutral pseudoscalars.

2.7 Projection for KLEVER

KLEVER is the proposed stage 3 of the HIKE experiment supplied by K_L beam with a mean momentum of 40 GeV [6]. The main goal is, as for the KOTO, the measurement of $K_L \rightarrow \pi^0 \nu \bar{\nu}$ with the aim to collect 60 signal events to measure the BR with a 20% uncertainty. In this section, a similar study as for KOTO is performed on the reinterpretation of this future analysis as a search for an off-axis $a \rightarrow \gamma\gamma$ decay of an ALP produced upstream in the target. This interpretation is completely complementary to the beam-dump search of the HIKE experiment since it does not require beam-dump operation and it is a ubiquitous byproduct of the $K_L \rightarrow \pi^0 \nu \bar{\nu}$ search, which does not even need a dedicated trigger line.

The geometry of the KLEVER detector is very different from the HIKE stages 1 and 2. The spectrometer is completely removed to minimize the interactions in the fiducial volume of the experiment while many photon vetoes are added to hermetically cover the decay volume, which is ended by an electromagnetic calorimeter with a radius of 1.25 m and central hole 20×20 cm. Two decay volume lengths are under study. The original proposal with decay volume starting at 120 m, which is given by the position of the upstream veto, and calorimeter located 241 m from the target as for NA62. Another proposal is to extend the decay volume by 150 m to enhance background suppression. In this case, the decay volume starts 270 m downstream of the target, and the calorimeter is also shifted by 150 m. Several options for the target material are under consideration: Be, Cu, and W.

The search strategy is analogous to the KOTO analysis. The $K_L \rightarrow \pi^0 \nu \bar{\nu}$ decay is assumed to happen on the K_L beam axis, and the vertex z position is reconstructed using the opening angle of the $\pi^0 \rightarrow \gamma \gamma$ decay. The $K_L \rightarrow \pi^0 \nu \bar{\nu}$ signal box is then defined in the z_{vtx} and p_T plane. Since the decay of an upstream ALP $a \rightarrow \gamma \gamma$ would be reconstructed as the signal $K_L \rightarrow \pi^0 \nu \bar{\nu}$, the obtained distribution is spread into an unphysically high p_T region. The requested cuts mirror the ones for the SM search in Ref. [6]. The radius of the photons on the calorimeter has to be > 0.35 cm, and the minimum energy per photon has to be more than 2 GeV. The expected ALP distributions in the signal plane for a toy ALP model are shown in the left panel of figure 2.32 in the case of the standard decay volume and in the right panel for the extended decay volume.




Figure 2.32: Signal event distribution N_{sig} for ALP with mass $m_a = 513$ MeV for the (i) coupling scenario with $C_{\gamma\gamma} = 2 \times 10^{-6} \text{ GeV}^{-1}$ as could be found in the $K_L \rightarrow \pi^0 \nu \bar{\nu}$ analysis at KLEVER (left) and extended KLEVER (right) with Be target with the full dataset 6×10^{19} POT. The grey boxes correspond to the $\pi \nu \bar{\nu}$ signal regions as defined in Ref. [6].

By comparing the ALP distributions with the background (and $K_L \rightarrow \pi^0 \nu \bar{\nu}$ signal) distributions, shown in figures 22–24 of Ref. [6], one can see that similarly to KOTO, the ALPs can be searched for in an almost background-free regime even for the setups for which the $K_L \rightarrow \pi^0 \nu \bar{\nu}$ search itself can be highly affected by the SM backgrounds. The obtained exclusions for KLEVER and the extended versions are shown in figure 2.33 for ALP benchmark scenarios (i) in the left panel and (iii) in the right panel. In order to avoid the SM backgrounds, the requirement for KLEVER is $p_T > 0.3 \text{ GeV}$ for all z_{vtx} , while for the extended KLEVER $p_T > 0.1 \text{ GeV}$ is required for $z_{\text{vtx}} < 260 \text{ m}$ and $p_T > 0.3 \text{ GeV}$ for $z_{\text{vtx}} > 260 \text{ m}$. The reconstructed z_{vtx} is required to be in the range $120 \text{ m} < z_{\text{vtx}} < 220 \text{ m}$ for KLEVER and in the range $120 \text{ m} < z_{\text{vtx}} < 370 \text{ m}$ for the extended version.

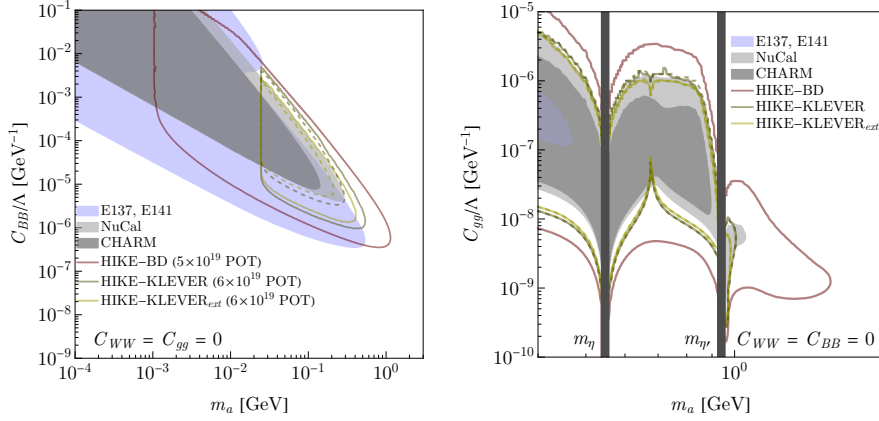


Figure 2.33: 90% CL exclusions limits for ALP scenarios (i) (left) and (iii) (right). Comparison of KLEVER with HIKE in beam-dump mode and past experiments. Full contours are assuming the Be target, dashed contours are assuming the Cu target.

As in the case of KOTO, it should be emphasized that this result is only a byproduct of the $K_L \rightarrow \pi^0 \nu \bar{\nu}$ search even without a dedicated beam-dump run and can be searched for in a *parasitic regime*. A dedicated analysis might result in better sensitivity. It can be observed that in spite of having a much longer decay volume, the acceptance of the extended version for ALPs is lower, which is given by the lower overall angular coverage. It is also found that despite lower statistics, HIKE operated in a dedicated beam-dump mode has much larger sensitivity.

Chapter 3

The NA62 experiment at CERN

3.1	NA62 beamline	105
3.2	Upstream region	108
3.2.1	CEDAR/KTAG	108
3.2.2	GigaTracker	108
3.2.3	VetoCounter	109
3.2.4	CHANTI	111
3.2.5	ANTI0	111
3.3	Downstream region	111
3.3.1	LAV	112
3.3.2	STRAW	112
3.3.3	RICH	113
3.3.4	CHOD/NewCHOD	113
3.3.5	LKr	114
3.3.6	MUV	115
3.3.7	Additional veto calorimeters	115
3.4	Trigger and readout	116
3.4.1	Readout	117
3.4.2	L0 trigger	119
3.4.3	HLT	122
3.4.4	PCFarm	123
3.5	RunControl	124
3.5.1	New devices	126
3.5.2	New control and monitoring tools and features	129
3.6	NA62 framework	130

The NA62 is a multipurpose fixed-target experiment situated in the so-called North Area of the CERN laboratory located in Preveessin (France). Its main aim is the study of an ultra-rare $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decay, but it also serves for precision measurements in the kaon sector and probe of SM-forbidden kaon decays. Moreover, it can search for BSM particles produced in kaon decays, but it can also be operated in a beam-dump mode to search for new particles of masses higher than m_K in a background-clean environment.

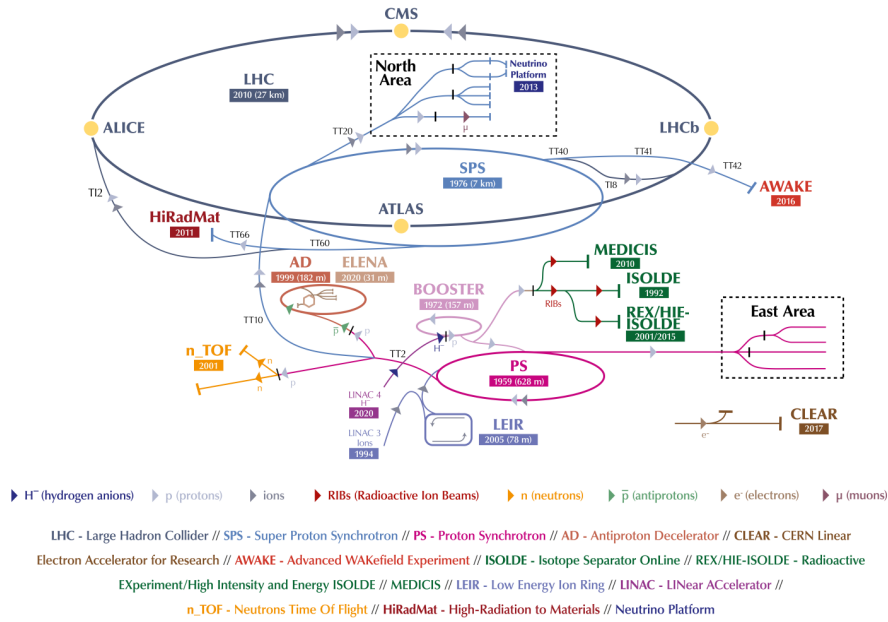


Figure 3.1: CERN's accelerators complex [181].

After a successful data-taking period 2016-18, called Run 1, NA62 restarted collecting data in 2021 (Run 2) after many software and hardware updates. It will keep operating until CERN Long-Shutdown 3¹. The layout of the detector in Run 2 is briefly discussed in this chapter, following the outline of [182], with a focus on experimental changes with respect to Run 1.

This chapter is organized as follows: the beamline layout and its various elements are described in section 3.1, the detectors preceding the decay volume and detectors in and downstream the decay volume are discussed in sections 3.2 and 3.3 respectively and the detector readout and trigger streams are briefly discussed in section 3.4. The different detector utilization and changes to the setup in the beam-dump mode with respect to the standard data-taking is highlighted in each section.

¹The Long-Shutdown 3 is currently scheduled between December 2025 and February 2029.

3.1 NA62 beamline

The NA62 is provided with a 400 GeV/ c proton beam by the CERN SPS accelerator with an intensity of the order of 10^{12} protons/s with protons delivered in spills with an effective duration of about 3 s.

The NA62 beamline (K12) starts with a 200 mm long beryllium target (T10). When T10 is impinged by the primary proton beam, a high-intensity flux of secondary particles is produced, from which a secondary beam is formed using a triplet of quadrupole magnets. These magnets collect large angle of particles produced at T10 and focus them into an achromat, a system that selects a specific momentum component of a beam, located about 23 m downstream from the target, as shown in the left panel of figure 3.2. The achromat is composed of two pairs of dipole magnets and movable collimators (TAX) composed of 800 mm of copper and 2400 mm of iron, each of which has a set of apertures. The two dipole magnets B1A and B1B preceding TAX displace the beam vertically proportionally to the particle momentum. The momentum component is selected by choosing the corresponding TAX aperture position, and the second pair of dipoles B1C and B2 shifts the beam back into the original vertical coordinate. Another set of three quadrupole magnets focuses the beam, and muons in the beam are swept out of the acceptance of the experiment by three dipole magnets (B3). The secondary beam then enters the experimental area. During the standard data-taking, the formed positively charged secondary beam has 750 MHz particle rate, 75 GeV/ c momentum with about 1% (rms) momentum spread and it consists of about 70% of π^+ , 23% of p^+ and 6% of K^+ .

In the area upstream of the decay region, the kaons in the beam are tagged, and the momenta of single beam particles are measured. The beam is further collimated with a final collimator TCX and deflected by TRIM5 magnet² before entering at about 105 m from T10 a 117 m long tank with 10^{-6} mbar vacuum. In the first, about 60 m of the vacuum tank, about 18% of the beam kaons decay in flight. This region is further referred to as a fiducial volume (FV). The FV is followed by a spectrometer measuring the momenta of the decay products and other detectors downstream, allowing the decay product particle identification. Together, the upstream and downstream detectors allow the full reconstruction of a kaon decay event. The leftover beam consisting of

²The beam is deflected before entering the decay volume to compensate for a subsequent deflection in the opposite direction by a downstream spectrometer magnet MNP33.

undecayed particles propagates further through the downstream region and is deflected by a magnet at the end of the experimental hall and dumped.

The experiment can also be operated in a beam-dump mode, in which the T10 target is lifted from the beamline and the primary proton beam is propagated to the TAX, as shown in the right panel of figure 3.2. The beam is made to interact in the TAX collimators, which are operated in a "closed" position: the positions of TAX11 and TAX12 are set at the maximum value of 140 cm. In this position, the beam cannot pass through the apertures, and the beam is effectively dumped in the TAX material, corresponding to approximately 19.6 interaction lengths. Dipoles B1A and B1B, on the other hand, are operated as in the kaon mode: the magnetic induction vectors are opposite, and the beam is made to interact along the Z direction at a nominal impact point (IP)

$$P_0 = (0, -22, 23070) \text{ mm} , \quad (3.1)$$

with the spread in the transverse plane [183]

$$\sigma_{P_{0,\perp}} = (4.7, 3.2) \text{ mm} . \quad (3.2)$$




Figure 3.2: The TAX area in the standard (left) and beam-dump (right) configuration. Blue (red) line indicates the $75 \text{ GeV}/c$ ($400 \text{ GeV}/c$) beam component.

Unlike in the standard data-taking, the currents of the B1C and B2 dipoles in the beam-dump mode are set to produce magnetic fields in the same direction, nominally along the $-Y$ axis. The magnetic induction generated by B1C is approximately -1.82 T ; that of B2 is approximately a factor of three lower. Such a setup has been suggested by Monte Carlo studies to minimize the flux of "halo" muons produced by pion decays within the TAX [184].




Figure 3.3: NA62 setup in 2021.

3.2 Upstream region

The detectors in the upstream region serve mainly for beam tracking and kaon tagging. Several veto detectors are also situated here for detecting events in which beam particle has decayed or interacted inelastically upstream, possibly resulting in final states entering the downstream region and mimicking a signal of decay inside the decay region.

While these detectors are needed for the operation with the kaon beam, the data from them are also collected during the beam-dump mode, even if no secondary beam is present. Notably, the veto detectors, but also the beam-tracking detectors described in the following sections operated as a veto, can help study the remaining muon halo in the beam-dump mode.

3.2.1 CEDAR/KTAG

A Cherenkov light detector (CEDAR) equipped with eight arrays of photodetectors (KTAG) is used in the beamline to identify the kaon component of the beam.

The CEDAR is a vessel filled with nitrogen at 1.75 bar at room temperature. Since the nitrogen-filled CEDAR presented the largest fraction of the material budget in the beamline (about 3.7×10^{-2} interaction lengths), in 2023, a new CEDAR vessel filled with hydrogen at 3.6 bar pressure has been installed, reducing the CEDAR material budget to 8×10^{-3} interaction lengths, while reaching a comparable detection efficiency [185].

In the CEDAR, the cone of radiated Cherenkov light is reflected by a mirror at the end of the vessel back to the front face where KTAG is located. At a given pressure, only the light with an angle corresponding to a particle of a given mass can pass through a set of diaphragms and enter KTAG. The light is focused on a set of eight spherical mirrors, which reflect the light into eight sections with photomultipliers (PMTs).

The kaon identification efficiency of KTAG exceeds 98% with the time resolution below 100 ps, also providing a very precise time measurement for the kaons.

3.2.2 GigaTracker

The GigaTracker (GTK) is a beam spectrometer that measures the momenta, time, and direction of passing beam particles.

The GTK comprises four stations (GTK0-4)³, located within an achromat (stations GTK0 and GTK1 placed before the B4A and B4B magnets, GTK2 station placed in the position of vertically shifted beam and GTK3 station situated after the B5 and B6 magnets) where the particle momentum is measured from its trajectory vertical displacement.

Each station is a silicon pixel detector with a sensor with a matrix of 18000 $300 \times 300 \mu\text{m}^2$ pixels located on an area of $63.1 \times 29.3 \text{mm}^2$. The sensor for each station is read out by ten bump-bonded chips (5 chips in 2 rows), so-called TDCpix. The substantial beam rate results in radiation damage to the silicon structure, accumulating over time and leading to degradation of the signal yield and time resolution. The effects of radiation damage can be reduced by increasing the bias voltage and by active cooling, which is also necessary for removing the heat generated by the TDCpix electronics.

The high beam rate ($\sim 750 \text{MHz}$ at nominal intensity) requires a single hit time resolution below 200 ps, which has been achieved with a measured time resolution of 115 ps [186].

3.2.3 VetoCounter

The VetoCounter is a new scintillating-tile detector commissioned in 2021, serving for detection of charged pions and photons from upstream kaon decays (mainly $K^+ \rightarrow \pi^+\pi^0$ and $K^+ \rightarrow \pi^+\pi^+\pi^-$) before GTK 3 station which has proven to be a significant background source for the $K^+ \rightarrow \pi^+\nu\bar{\nu}$ analysis in Run 1 [154].

It comprises three stations: stations F1 and F2 are located upstream, and station B is located downstream of the final collimator, before the GTK station 3, see left panel of figure 3.5. Stations F1 and F2 are contained in the same structure with a 25 mm thick lead sheet placed between the stations for photon conversion. Each station consists of 11 scintillating tiles, 9 tiles having the size $40 \times 120 \text{mm}^2$ and 2 tiles having half the size, $20 \times 120 \text{mm}^2$. Each tile is made of a polyvinyltoluene scintillator glued on the sides to two light guides with a PMT attached at the end.⁴ Three tiles are located above the beam pipe and eight below, always with one smaller tile located just above and below the beam pipe so that the detector readout can sustain the higher particle rate (the

³Formerly, the GTK consisted of 3 stations. The fourth station, GTK0, has been added in 2021 for the NA62 Run 2.

⁴The PMTs used and their initial settings after the VetoCounter installation are described more in appendix B.

particle rate in each tile in this configuration is below 1 MHz). The PMTs are read out by four constant fraction discriminators (CFD) developed at CERN, each giving a low- and high-threshold output channel for the time-to-digital converters and the readout electronics described in section 3.4.1.

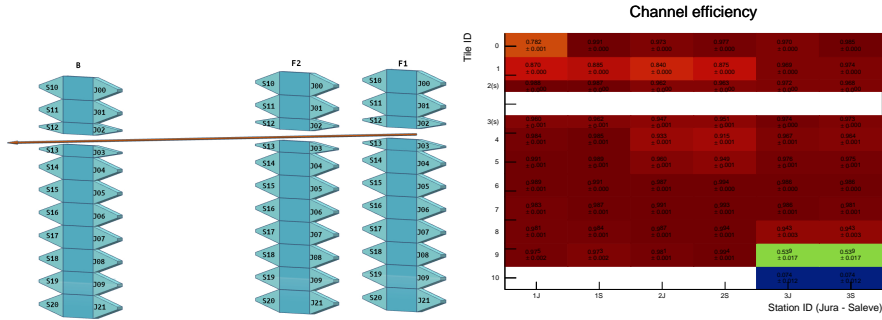


Figure 3.4: Left: Sketch of the VetoCounter tile layout with stations F1 and F2 located before the final collimator and station B located after. Right: Channel halo efficiency measured in kaon run 13669 using `VetoCounterDQMonitor` analyzer.

The VetoCounter efficiency can be measured using reconstructed downstream tracks of halo particles, propagated upstream to the VetoCounter stations and matching with VetoCounter hits, see right panel of figure 3.4 for the measurement during a kaon run. A more precise measurement can be performed during a muon run in which the kaon beam is not present, and a uniform flux of muon halo passes through the experiment. The highest efficiency is observed in station B, while for the upstream stations, the muon halo traverses the final collimator and undergoes non-trivial multiple scattering; therefore, the observed efficiency is typically lower.




Figure 3.5: Gaussian-fit parameters for the distribution of the difference between the times of the Jura and Saleve channels obtained with `VetoCounterDQTELMonitor` analyzer. Left: all hits with overall measured $\sigma_{\text{measured}} = 0.456$ ns. Right: fitted values for each tile individually.

The time resolution is measured for each tile using the Jura and Saleve channels on the mutual time difference $\Delta t = (t_{\text{Jura}} - t_{\text{Saleve}})$, see left panel of figure 3.5, as $\sigma_{\text{res}} = \sigma_{\text{measured}}/\sqrt{2}$ for σ_{measured} obtained by a Gaussian fit of Δt . See right panel of figure 3.5 for the obtained fit results.

3.2.4 CHANTI

The Charged Anti-Coincidence detector (CHANTI) is a veto detector placed between the GTK station 3 and the beginning of the decay region. Its purpose is to detect inelastic particle interactions in the last GTK station since the final states originating from this process are not swept away by the magnetic field as in the other GTK stations. The CHANTI consists of six $300 \times 300 \text{ mm}^2$ hodoscope stations, each made of 48 bars triangular in cross-section arranged in two perpendicular planes: one for X and one for Y coordinate. Each bar has a polystyrene scintillator with an optical fiber in the middle connected at the ends to a PMT.

3.2.5 ANTI0

The ANTI0 is a new hodoscope commissioned in 2021, situated just before the vacuum tank end cap. It has an octagon shape circumscribing a circle with radius 1080 mm to cover the whole geometrical acceptance of the first section of the vacuum vessel, except for a $232 \times 232 \text{ mm}^2$ gap in the center for the beam pipe. It is used for detecting events with upstream halo muon entering the decay volume, possibly mimicking a signal of $K^+ \rightarrow \mu^+ \nu$ decay when passing close to the beam kaon trajectory within the decay volume. It can also help with the trigger rate reduction in the standard data-taking with the kaon beam. The ANTI0 gets even more significance in the beam-dump mode, where the muon halo presents the only background source.

It is composed of 280 scintillating tiles arranged in a chessboard-like layout. Each tile has a $124 \times 124 \text{ mm}^2$ polystyrene scintillator glued with light guides read out by silicon photomultipliers (SiPM).

3.3 Downstream region

The detectors located in and downstream of the 117 m long vacuum tank serve for identifying the decay products of kaon decays and determining their properties. To account for the spread of the decay products as they propagate down-

stream, the vacuum tank diameter widens from 1.92 m at the beginning up to 2.8 m at the end.

3.3.1 LAV

In order to achieve the needed background suppression for the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$, a complete hermetical coverage of the downstream region for photons escaping the detection by the calorimeters is needed. For vetoing photons radiated at large angles between 8.5 and 50 mrad at which they would escape the vacuum tank, a Large Angle Veto (LAV) system consisting of 12 ring-shaped stations is placed along the whole downstream region. 11 LAV stations encircle the vacuum tank and a twelfth station operates in the air just about 3 m upstream of the calorimetric system. The diameter of the LAV stations gradually increases with the diameter of the vacuum tank.

Each station has several rows of lead glass blocks encircling the perimeter of the inner volume. The rows are mutually shifted for the blocks to cover the whole circumference. A particle passing through the lead glass leaves an electromagnetic shower, which is detected using a PMT via the Cherenkov light produced.

3.3.2 STRAW

The STRAW spectrometer, occupying the last 35 m of the 117 m long vacuum tank, serves for tracking and momentum measurement of the final state particles. It is composed of four straw chambers, two located before and two after a large aperture dipole magnet (MNP33) with an integrated magnetic field of $0.9 \text{ T} \cdot \text{m}$. Each straw chamber has four views perpendicular to the beam axis under different angles (X (0°), Y (90°) and U (-45°), V ($+45^\circ$)) to provide a more precise measurement of the particle coordinate in the chamber and to prevent ambiguities in cases with multiple in-time hits. Each view consists of 448 straws, 2160 mm long about 9.82 mm in diameter and a 12 cm gap for the beam is left near the center.

The straws are staggered in four layers in each view to cover all possible angles under which the particle can cross it and they are tensioned at 1.5 kg to decrease the deflection, especially of the horizontally positioned straws due to gravity. To minimize the scattering, the straws are made of a very thin ($36 \mu\text{m}$) PET foil. The straw chambers are operated in the vacuum tank with the single straws filled with a mixture of 30% of CO_2 and 70% of argon at atmospheric pressure. The PET foil is coated from the inside by two metal layers (50 nm

of copper and 20 nm of gold) as a cathode, and a gold-plated tungsten anode wire tensioned at 80 g with an electric field ~ 1 kV between the anode and the cathode goes through the center of the straw. The passing particle ionizes the filling gas, and the leftover charge is collected and processed by the front-end electronics.

The charged particle trajectory is reconstructed by combining the hits from different chambers, and the particle momentum is calculated from the momentum kick in the horizontal plane provided by MNP33. The momentum resolution of the spectrometer is $\sigma(p)/p = 0.30\% \oplus 0.005\% \cdot p$, where p is in GeV/ c .

3.3.3 RICH

The Ring Imaging Cherenkov Counter (RICH) is a downstream particle identification detector designed to separate pions from muons between 15 and 35 GeV/ c momentum. With a very good intrinsic event time resolution of about 70 ps, it also provides a referential time for charged downstream particles.

The RICH is a 17.5 m long vessel with a conical shape with a diameter of 3.2 m at the downstream end and an increasing diameter backward against the direction of the beam. It is filled with a neon gas at a constant pressure of 990 mbar. The entrance and the exit windows are made of 2 and 4 mm respectively thick aluminum, where the entrance window separates the vessel from the vacuum of the decay region. To increase the resolution and to keep the photomultipliers out of the active area, the downstream end of the vessel is filled with a mosaic of 20 spherical mirrors, which reflect the cone of Cherenkov light backward onto the photomultipliers. Half of the mirrors point left and half right of the beam pipe. The inner side of the vessel is covered with black epoxy paint to avoid reflection of the Cherenkov light from it. The measured efficiency in the 15-35 GeV/ c range for pion selection was 86% with a 1.3% muon survival probability.

3.3.4 CHOD/NewCHOD

A system of charged particle hodoscopes (CHOD) is located between RICH and the calorimeters to trigger events with at least one charged particle passing through and to provide the impact point position and timing. There are two CHOD detectors, one refurbished from the previous NA48 experiment (NA48-CHOD) and a new detector (NewCHOD).

The NA48-CHOD is a fast scintillator detector with $\mathcal{O}(200)$ ps time resolution composed of two planes of scintillator slabs to provide the X and Y coordinates of the impact point. Each plane has 64 slabs read out at one end by a PMT. The NA48-CHOD has lower efficiency for high-intensity beams than the NewCHOD but is used as a control detector in the experiment.

The NewCHOD is a single plane of 152 tiles made by a polystyren scintillator read out by optical fibers connected to SiPMs. The tiles are typically 108 mm high and 134 mm or 268 mm wide with 1 mm overlap between the tiles.

3.3.5 LKr

The Liquid Krypton calorimeter (LKr) is an electromagnetic calorimeter reused from NA48. It is a quasi-homogeneous calorimeter placed inside a cryostat filled with 9 m^3 of liquid krypton at a temperature of 130 K. The outer radius of the LKr is 128 cm, and the inner 8 cm are left for the beam pipe. The active area of LKr is composed of a single plane about 127 cm in depth, corresponding to $27X_0$, allowing for full development of electromagnetic showers. It enables a precise photon detection with an inefficiency of the order 10^{-5} . The LKr contains 13248 cells of about $2 \times 2\text{ cm}^2$ with thin Cu-Be ribbons as electrodes. The alternating electrodes filling the plane are oriented parallel to the beam axis but slightly deflected, changing the direction of deflection several times in a ± 5 mrad zig-zag fashion to avoid inefficiency for very shallow tracks close to the anodes. Electrodes are connected directly to the preamplifiers, sending a signal to the readout electronics.




Figure 3.6: Left: One LKr quadrant layout, Right: LKr cells ‘zig-zag’ arrangement detail [134].

3.3.6 MUV

The muon veto (MUV) is a system of three independent detectors succeeding the LKr: hadronic calorimeters MUV1 and MUV2 and a fast muon veto MUV3. MUV1 and MUV2 are iron-scintillator sandwich calorimeters of twelve vertical and twelve horizontal layers of iron alternating with scintillator strips (44 per plane for MUV1 and 22 per plane for MUV2) covering a plane of $260 \times 260 \text{ cm}^2$ and corresponding to about $\sim 8X_0$. Each scintillator strip is 260×6 (or 4) $\times 1 \text{ cm}$ and is connected to photomultipliers through the WLS fibers in the case of MUV1. The MUV2 is a refurbished original NA48 hadron calorimeter, very similar to the new MUV1. It has readout photomultipliers connected directly to scintillators. The iron wall of 80 cm is placed between the MUV2 calorimeter and the MUV3 muon veto and is thick enough to stop all particles but muons. MUV3 itself is composed of $22 \times 22 \text{ cm}$ tiles with two photomultipliers placed 21 cm behind each tile. They cover the angle for scintillation photons, but as they also detect Cherenkov radiation if a muon passes through the photomultiplier windows, two photomultipliers are necessary for better timing.




Figure 3.7: MUV detectors layout [134].

3.3.7 Additional veto calorimeters

Additional veto detectors are placed in the area downstream of the vacuum tank to study decays with final states radiated at very large or very small angles. The Small angle veto (SAV) is a calorimetric system serving for the detection of photons emitted at small angles out of the acceptance of the LKr. It is composed of the Small-angle calorimeter (SAC) and the Intermediate ring calorimeter (IRC). Both detectors have a similar construction and are made of

an alternating sequence of lead plates and scintillators read out by PMTs coupled through WLS fibers. The IRC is located between the CHOD and LKr, covering a radius between 7 cm and 14 cm, and the SAC is placed inside the beam vacuum towards the beamline when the charged part of the beam is swept to the beam-dump, to detect the leftover photons in the beamline.

A second pair of detectors serves for the detection of charged pions in the $K^+ \rightarrow \pi^+\pi^+\pi^- (K_{3\pi})$ decays and consists of the peripheral MUV0 detector and the hadronic sampling calorimeter (HASC). The MUV0 is a single plane of two layers of scintillator tiles read out by PMTs via WLS and detects π^- with momenta below 10 GeV/c deflected by the spectrometer magnet out of the acceptance of other downstream detectors. The HASC is a sandwich calorimeter consisting of 9 identical modules and is located slightly off-axis after the MUV3 before the SAC. The HASC can detect π^+ above 50 GeV/c propagating through the beamline as these π^+ are deflected off-axis more than the charged (75 GeV/c) component of the beam which is dumped. In 2021, a second HASC station was placed symmetrically on the other side of the beamline for detecting electrons produced by high-energy photons interacting with the RICH beam pipe. This will allow further reduction of the $K^+ \rightarrow \pi^+\pi^0$ background, necessary for a precise $K^+ \rightarrow \pi^+\nu\bar{\nu}$ measurement.

3.4 Trigger and readout

As a consequence of the high beam rate, the rate of downstream decays reaches up to 10 MHz, resulting in a data rate too high for the storage. In order to reduce the data rate while maximizing high-quality data collection, a unified high-performance trigger and data acquisition (TDAQ) system is used at NA62. The TDAQ system is composed of the readout, a single hardware trigger level (L0), and high-level software triggers (HLT).

The L0 trigger decision is stated by the L0 Trigger Processor (L0TP), and it is based on the trigger primitives generated by the readout of single detectors. The link between the L0TP and the readout electronics is provided by a Local Trigger Unit (LTU) with roughly one LTU per subsystem.⁵ The HLTs are performed at an online computing farm (PCFarm). Both levels reduce the data rate to 10 - 100 kHz, which meets the bandwidth of the data flow for permanent

⁵The link mode of operation of L0TP is called the *global* mode. The LTU can also be operated in a *standalone* mode in which the LTU emulates the L0TP protocol and allows running the subsystem without L0TP for performing tests and calibration and development tasks [187].

storage. The data passing all the trigger levels are written on the disk by the Central Data Recording (CDR) service.

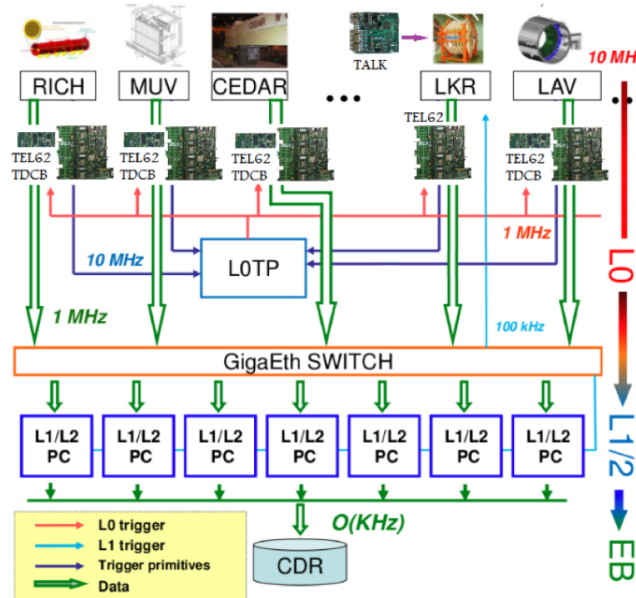


Figure 3.8: An overview of the NA62 trigger system and the corresponding data flow reduction [188].

A common Timing, Trigger and Control (TTC) system, adopted by all LHC experiments, is also used at NA62 for clock and trigger distribution. The LTU sends signals via two multiplexed channels encoded using 40 MHz clock and transmitted at 160 MHz rate using a unidirectional optical fiber. One channel carries the L0 trigger signal, and the other has information about resets and trigger types. All the synchronous elements of the TDAQ system run on the TTC clock and are synchronously reset by a Start of Burst (SOB) command and similarly stop on an End of Burst (EOB) signal delivered after a given number of 25 ns clock periods.

3.4.1 Readout

Since most of the detector sub-systems at NA62 have similar requirements on timing precision and readout capabilities, a common system based on high-precision Time-to-Digital converters (TDCs) has been designed. In order for the readout to cope with the increasing beam intensity, a new system is being tested as a readout of the VetoCounter detector in parallel with the common

readout and will be deployed at first for the highest-rate detectors, KTAG and CHANTI. The spectrometers use their own dedicated systems made suitable for a high number of channels. GTK has the highest number of channels, which are coupled by optical fibers to meet the necessary time resolution. In contrast, for the Straw spectrometer with a lower intrinsic time resolution, a lower time resolution readout is sufficient. Calorimeters use flash analog-to-digital converters (FADC) and have their own readout system.

TEL62

A common TDC-based system, TEL62, provides a 100 ps time resolution and can be reprogrammed, allowing scalability and the flexibility to implement different trigger conditions. One TEL62 board is able to host up to four mezzanine TDC boards (TDCB) for input data channels controlled by field-programmable gate array (FPGA) with up to 2 GB of DDR2 RAM. Each TDCB houses four high-performance TDC (HPTDC) chips and can digitize times of leading and trailing edges for 128 detector channels. Therefore, one TEL62 can read out and trigger on 512 detector channels. The TEL62 can also host an output quad-gigabit Ethernet board driven by a fifth identical FPGA, which sends L0 trigger primitives to the L0TP and detector data blocks to the PCFarm when the L0 trigger decision is positive.

GTK-RO

Each GTK TDCpix is read out by a GTK off-detector readout card (GTK-RO) via a 200 m long optical fiber. Five GTK-RO cards are housed by one computer in the PCFarm, with two computers per GTK station to balance the amount of data. The GTK-RO card is controlled by an FPGA with 2 GB of DDR2 RAM, and it is interfaced with TTC by a mezzanine card. The design's modularity allows reading the GTK data either at the L0 or at the L1 trigger.

Straw readout

The Straw readout is located in the experimental cavern and is composed of an on-detector front-end part and a back-end part positioned about 10-15 m from the detector. Each front-end board, also used as a cover for the gas operation (FE-cover), serves 16 straws and houses 2 readout chips (each processing leading and trailing edges of 8 channels), one chip for setting the discriminator thresholds via 16 DACs and an FPGA with 32 TDCs. The back-end electronics

is built on Straw Readout Boards (SRB) and serves for communication with the TTC system and sending the collected data to the PCFarm upon selection by the trigger. Each SRB manages up to 16 FE-cover boards.

CREAM

The Calorimeter Readout Modules (CREAM) are used for LKr, MUV1, MUV2, IRC and SAC calorimeters. Each CREAM board can read up to 32 analog channels via a daughter board, which shapes the input signals and further digitizes them using a 14-bit FADC [189]. Serialized data are then sent to an FPGA [190] on the motherboard and copied to a DDR3 RAM with 8 GB capacity [191].

FELIX

The Front-End Link eXchange (FELIX) system, originally designed for the ATLAS upgrade, is a new FPGA-based system insensitive to radiation and exhibiting no rate limitation. Output from the high-voltage and low-voltage threshold channels from the detector front-end TDCs is sent via a radiation-hard optical link to the FELIX board. It is an FPGA-based PCIe card housed in a PCFarm PC; the complete spill data are loaded into the PC's memory and sorted for matching with the trigger. This also allows to see a detailed fast spill profile, which can be useful for beam-quality monitoring. The system allows operation at both L0 and L1. Currently, it is implemented parasitically as a VetoCounter readout in parallel with the TEL62 readout for testing. Further, it will be deployed for CHANTI and KTAG detectors.

3.4.2 L0 trigger

The L0 is a hardware trigger designed to filter the events given the inputs from a small set of fast detectors. It is based on L0 trigger primitives, 64-bit data blocks containing identifiers indicating which conditions are satisfied at a given time. The L0 primitives are generated directly by TEL62s; for the calorimeters, a dedicated system of TEL62s reading out first the information from CREAMs (L0-Calo) is used. The maximum output rate of the L0 trigger is 1 MHz.

TDC-based L0

The TDC-based L0 trigger is based on the identification of hits belonging to the same event. This operation first requires hit sorting which is performed by the TEL62 since the trigger algorithms are significantly more efficient if fed by already sorted data. Each detector then has a custom set of conditions for generating the trigger primitives.

The CHOD L0 provides a selection of just single-track events since impact time corrections are required due to the length of NA48 CHOD scintillator bars and these time corrections become time-demanding in the case of multiple hits. It is used in L0 as an independent stream (control trigger). The New-CHOD L0 does not require such corrections and offers more selective conditions, including multiple hits. Since RICH has many channels distributed over several TEL62s, additional communication between the boards is necessary for primitive evaluation. Also, an additional TEL62 is used, reading channels by groups of 8 to have the information from the whole detector available on one board. The RICH L0 is then based on the multiplicity of clusters of hits. Since LAV, similarly to RICH, is read out by multiple TEL62 boards, a similar inter-communication among the boards is needed. A single station can also be used at L0, typically the LAV12 L0. The MUV3 L0 can be used either as a positive trigger or a fast muon veto. Its channels are connected to a single TEL62 board equipped with three TDC boards.

L0-Calo

The L0-Calo identifies clusters in data read out from CREAMs and sorts them in time. Different conditions can then be put on the number of clusters of given energies or the total deposited energy. The system is composed of 37 TEL62 boards arranged in three stages: Front-End, Merger, and Concentrator. 29 Front-End TEL62s (currently 28 for LKr) use simplified information from CREAMs for blocks of 4×4 cells and perform a search for energy peaks for which time, position, and energy are calculated. The Merger TEL62s merge peaks from Front-End TEL62s to single clusters. A single Concentrator board counts the number of clusters, computes total electromagnetic energies, and generates trigger primitives. Due to a significant data rate of calorimeters, which does not allow readout at the L0 rate, CREAMs store data in local temporary buffers after receiving the L0 trigger, and the data are transmitted only when receiving the L1 trigger at a lower rate.

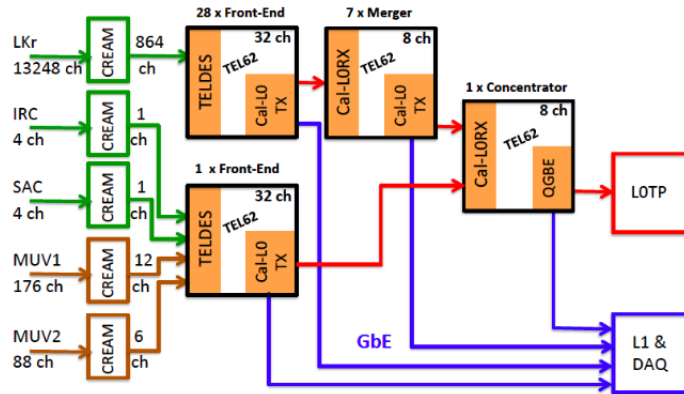


Figure 3.9: A scheme of the L0-Calo system [134].

L0TP

The L0TP acquires the L0 trigger primitives from the TEL62s via Gigabit Ethernet links using a UDP internet protocol, sorts them in time, and checks if the combination of primitives within the time window meets the preset configuration, so-called trigger mask. If it does, the L0TP sends a signal for the readout to send the data to the PCFarm with a preset downscaling for given trigger masks. The L0TP is based on a commercial board hosting a single FPGA.

The L0TP firmware is built of several blocks. Seven Ethernet interfaces receive primitives from detectors, which are aligned by a Delay Generator and stored in 7 circular buffers. The synchronization of the internal clock with the TTC is managed by a Phase Locked Loop (PLL), which automatically adjusts the phase of the internal signal with the input one. There are two clock domains; the 40 MHz one is common with the TTC, and the 125 MHz clock is common with the trigger-algorithm logic and the Ethernet communication. The Memory Management Unit (MMU) allows driving the two uncorrelated streams of Physics and Control trigger corresponding to the same events. The time alignment of these two uncorrelated flows is critical as the output RAMs are read in parallel, so the addresses of RAM are using the output of MMU, which is the content of merged FIFO (First-In-First-Out type buffer) filled with sorted primitives of one $6.4 \mu\text{s}$ frame from the reference and control FIFOs. The matching of the primitives with trigger masks is managed by the Associative Memory Module (AMM). Up to 8 trigger masks are available for up to 7 detectors at L0. Each mask is defined by a sequence of 7 bits, each of which can be in one of 3 states: 0, 1, “don’t care” depending on whether the primitives from the detector are ignored or whether they are used as a trigger or as

a veto. The AMM also generates a so-called trigger word sent to the detectors via LTU to tell them which action to take. Besides Physics and Control triggers, there are also special trigger words used, for example, for calibration, but also the already mentioned SOB and EOB, which enable and disable the data-taking and trigger-sending SOB- and EOB-specific data. There are also two special trigger words for anomalous activities that are raised by the detectors: the CHOKE and the ERROR trigger.

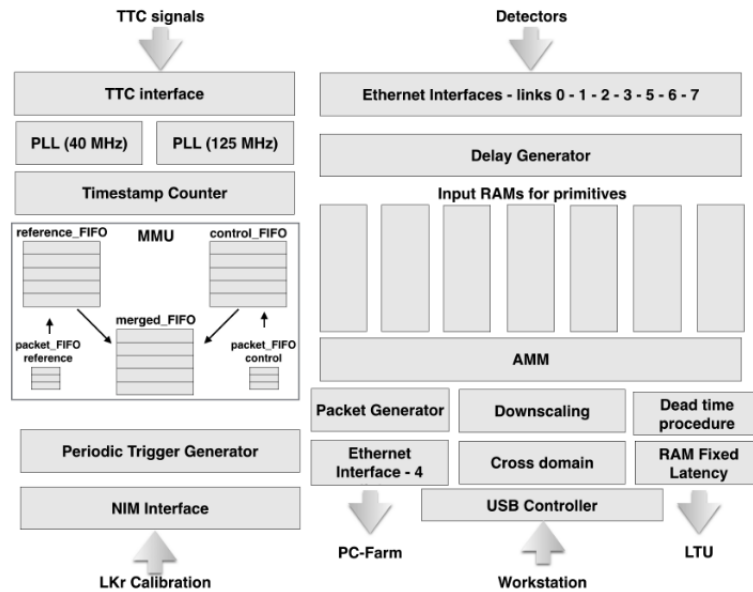


Figure 3.10: A scheme of the L0-Calo system [192].

The 1Gbit Ethernet link of the L0TP presented a bottleneck for data collected at high beam intensity. In 2021, a new L0TP, called L0TP+, with a 10 Gbit UDP port, was developed to overcome this issue. The L0TP+ also has a more powerful FPGA, allowing more complex trigger lines, possibly even using machine learning for triggering in the future. The L0TP+ has been operated parasitically with L0TP in 2022 and fully deployed in 2023 [185].

3.4.3 HLT

While the maximum rate of the L0 trigger is 1 MHz, a considerable data reduction has yet to be made to meet the requirements of the permanent data storage, and hence, a high-level software trigger (HLT) is used. Currently, several HLT algorithms are operated at the PCFarm as an L1 trigger. Standalone information from any sub-detector except the calorimeters can be used.

KTAG L1

The KTAG L1 is based on the multiplicity of hits in the sectors of KTAG in order to identify the presence of kaon in the event. Typically, a signal in > 4 sectors from 8 is required.

LAV L1

The LAV L1 is a veto with a limit on the hit multiplicity in selected LAV stations. Typically, < 3 hits in stations LAV2-LAV11 are required in order to suppress the $K^+ \rightarrow \pi^+\pi^0$ background.

STRAW L1

In STRAW L1, multi-track events are identified, and track charges are determined.

CHOD L1

The CHOD L1 can also be used for identifying multi-track events utilizing the multiplicity of hits in the slabs.

3.4.4 PCFarm

The PCFarm is a set of 30 computers that receive the L0-triggered data in multi-event packets (MEP) from the detectors via optical links using a UDP internet protocol. The MEPs are distributed over the PCFarm computers using a round-robin algorithm, which is reset on the SOB signal. Standalone information from each subdetector is processed via L1 algorithms issuing L1 trigger.⁶ After passing all trigger stages, the data from all PCs are sent to one of the merger PCs⁷ which orders the events in time, merges them in a single file, and sends the file to the permanent storage.

⁶L2 triggers using information from multiple subdetectors can also be implemented in the PCFarm.

⁷Three merger PCs were operated during Run 1. This number has been increased to five for Run 2, and a second hard drive has been installed in Merger 2 PC.

3.5 RunControl

The RunControl is a SCADA-based⁸ system built using the WinCC OA system and CERN-made JCOP [193] and UNICOS [194] frameworks.⁹ It controls and monitors all NA62 TDAQ systems and provides a common front-end for the various types of devices. The current layout of the RunControl graphical user interface can be seen in figure 3.11. The network communication between the RunControl and the devices is mediated by the Distributed Information Management System (DIM) [195] developed at CERN. The DIM system uses a client-server model where a server can publish a list of services while a client can subscribe to these. The DIM services can be of a *value*-type or a *command*-type where the former is used for the server→client information propagation, while the latter goes in the opposite direction. A comprehensive description of the RunControl system, including technical details and implementation, is given in Ref. [196]; only a general description will be given with more focus on the modifications of the system compared to [196].

As described in the previous section 3.4, the nature of the TDAQ components can be very diverse. From the front-end point of view, however, the devices can be described by a limited number of states and operated by a limited number of commands, which are common to all devices. Some devices present sub-components of a larger component, like, for example, a detector. To allow efficient control and monitoring of a large number of TDAQ devices, the system follows a tree-like hierarchy where the various devices (leaves) are arranged in common logical nodes (branches) connected to the root node. Each node presents an abstract device, called a finite state machine (FSM), whose purpose is to propagate the commands down the FSM tree to the leaves, which then send the command via DIM to the device and, in the opposite direction, to propagate the states received from the device DIM services from the leaves up the FSM-tree to the root. Initially, each branch had its own RunManager [196], which allowed disconnecting it from the main FSM tree and operating it independently under a different user. This feature can be useful, for example, for subsystem development purposes, allowing multiple experts to operate in parallel. It has, however, never been used and in made the work with the RunControl more complicated for the operators. After the CERN Long Shutdown

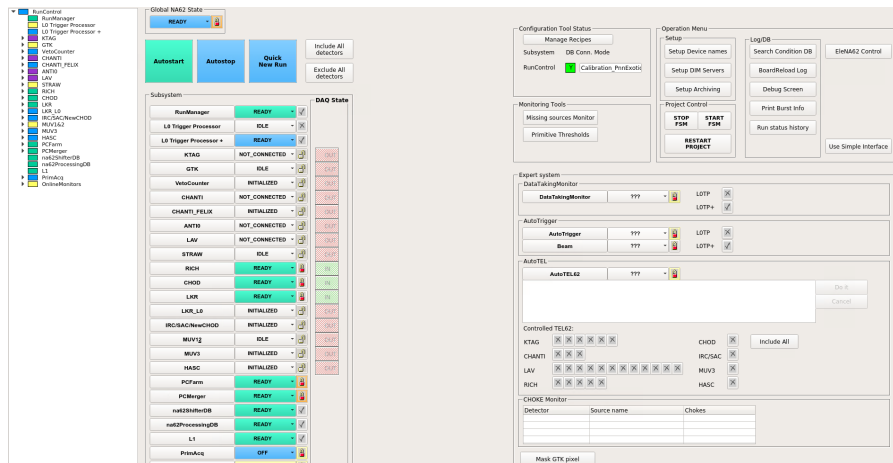


Figure 3.11: The RunControl graphical user interface. The NA62 root node and its branches are shown in the FSM tree on the left. Various additional control and monitoring tools are available to the shifters and the experts on the right.

2 (LS2), this feature has been deprecated and the main tree is operated by a single RunManager.

For the software and hardware nodes of the PCFarm and mergers, the states are simply **ON**, **OFF** and **ERROR**, while the commands are **STOP**, **START** and **RESTART**. For all other devices, the states are **READY**, **IDLE**, **INITIALIZED**, **ERROR** and **NOT_CONNECTED**¹⁰ and the commands are **START_RUN**, **END_RUN**, **RESET** and **INITIALIZE**.¹¹ The LOTP has an additional state **RUNNING** and additional commands **ENABLE_TRIGGERS** and **PAUSE_TRIGGERS** for starting and pausing the data acquisition. The **ENABLE-** and **PAUSE_TRIGGERS** and **START-** and **END_RUN** commands are also issued automatically when the burst counter reaches 1500 in order to have a consistent length of the run but also to save the data-taking time by performing the task for the shifters. An additional **LOAD_RECIPE** FSM command, which does not result in sending a command to the devices via DIM, loads the configuration for given devices from the configuration database into the RunControl. The JCOP recipe mechanism is used for the management of

⁸Supervisory Control And Data Acquisition.

⁹Joint Controls Project and Unified Industrial Control System.

¹⁰The device is in **NOT_CONNECTED** state when the connection to the corresponding DIM server is off.

¹¹Additional, newly implemented **BOARD_RELOAD** command allows reloading of the firmware of given board in case it has been corrupted for example due to the radiation. This operation can be performed by the shifters, and it is automatically logged in the electronic log-book, which notifies the corresponding expert about the event.

the device configuration.¹² When issuing the **START_RUN** or the **INITIALIZE** commands,¹³ the configuration specific to the given device is stored in an XML file and loaded by the devices into their internal memory.

3.5.1 New devices

During the LS2, several new TDAQ devices have been installed at NA62. While the expansion of the GTK and HASC detectors or the installation of the new ANTI0 detector concerned only the addition of new devices of already implemented types, such as TEL62 or GTK-RO boards, the installation of the VetoCounter detector required the implementation of an interface for the FELIX readout. The requirements on the interface stem mainly from the structure and the information available in the configuration files for a given device. Even though the content is very different, a similar data structure allowed for deriving the FELIX interface from the one used for SRBs. In 2023, the FELIX readout was also implemented for CHANTI testing. In order to allow the comparison between the data collected with TEL62 and with FELIX, the VetoCounter and CHANTI FELIX are implemented in the RunControl as independent detectors on the TEL62-readout ones. This information is then also propagated to the data processing, and it is accounted for in the data reconstruction.

The installation of L0TP+ also required a dedicated interface since it has a unique configuration file structure. Moreover, since the installation of L0TP+, both the L0TP and L0TP+ devices are present, and during a testing period, they were exchanged frequently. This required changes in the logic of how the L0TP FSM type is handled since the presence of two L0TP devices was not foreseen during the RunControl development. The L0TP DIM server publishes numerous services that provide real-time statistics of the number of trigger primitives generated and processed. The monitoring of these values is central for the data-taking shifters as they are the first indicator if the TDAQ is running correctly, and they are displayed in the RunControl *BigScreen* panel, see figure 3.12. However, exactly one L0TP device can be used for the triggering

¹²A cross-check between the loaded recipe and the recipe in the DB is newly done in the configuration panel and notifies the expert if the most up-to-date configuration has not been committed to the DB, preventing overwriting of the loaded configuration with an outdated recipe.

¹³Except for the **RESET** command, the commands can be followed by parameters, allowing to specify the action without the need of reloading the configuration file.

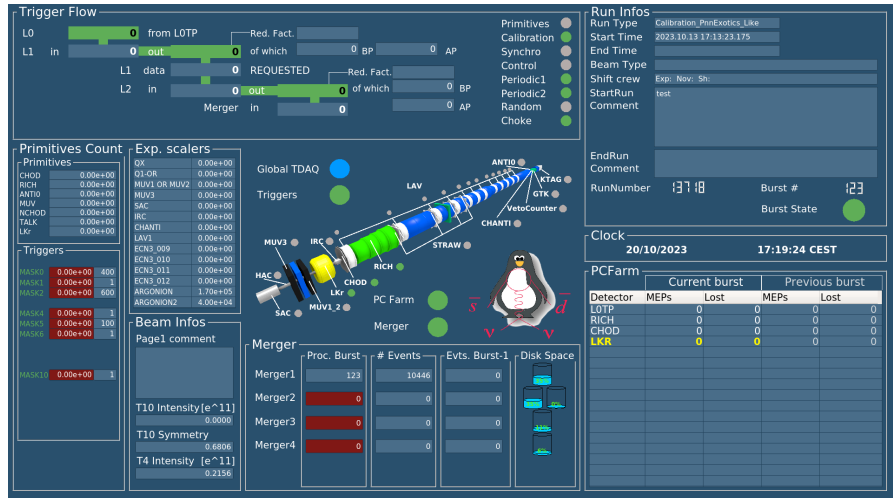


Figure 3.12: The RunControl BigScreen panel which is displayed on one of the wall screens in the NA62 Control Room.

while the other can be operated only in a parasitic mode,¹⁴ therefore, the possibility to change easily the LOTP device used for triggering and monitoring in the RunControl had to be implemented.

A slight change also concerns the LOTP communication with the RunControl. LOTP+ has an additional JTAG board, independent of the LOTP+ FPGA, which allows external access to the LOTP+ for debugging but also, for example, for reloading the LOTP+ firmware which cannot be done from the LOTP+ itself. Even though they are a single device from the RunControl point of view, the JTAG board is provided with an additional dedicated DIM server. Therefore, when commands are issued to the LOTP+ device, the RunControl has to propagate them to the correct DIM server.¹⁵ Another modification concerns the trigger configuration, which is communicated differently to the LOTP+, while for the user, the operation should be the same as for LOTP. For LOTP+, the enabled trigger types are specified in the configuration file sent to the board at the start of each run, unlike for LOTP, where the list of enabled triggers is sent as a parameter of the **ENABLE_TRIGGERS** command. Originally, this feature permitted LOTP to change the trigger configuration during the run, which can be useful for debugging by using the number of trigger primitives generated.

¹⁴This does not prohibit sending commands to the LOTP device which is not used for the data acquisition.

¹⁵For example **BOARD_RELOAD** command followed by the firmware version as a parameter is addressed to the JTAG DIM server.

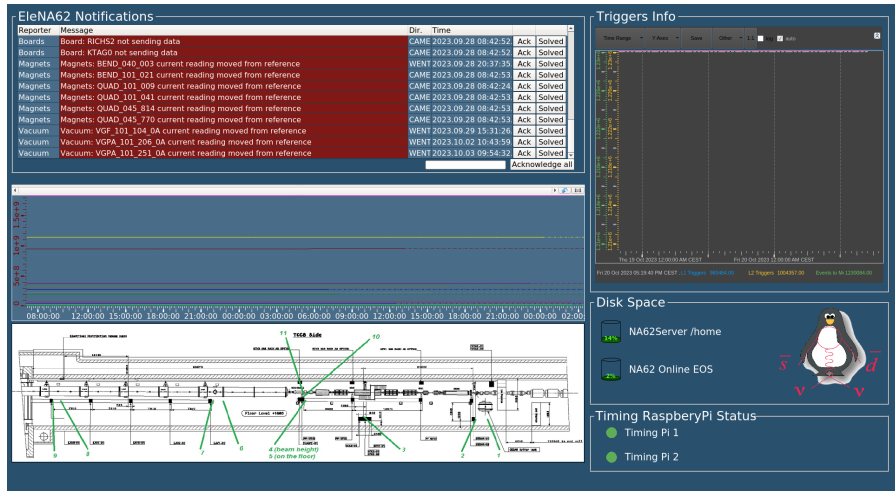


Figure 3.13: The RunControl SideScreen panel which is displayed on one of the wall screens in the NA62 Control Room.

However, it contradicts the definition of the run as a collection of bursts with the same data-taking conditions, and, as a consequence, it complicates the data processing. Therefore, it has been disabled for LOTP, and the modification is allowed only out-of-run, the same as for LOTP+.

Slight modification concerns also TEL62. The PCFarm expects all readout boards included in the DAQ to send data. The possibility of using a TEL62 board for triggering data collected by another TEL62 board has been implemented. This additional TEL62 board does not send the data to the PCFarm; therefore, in the RunControl, the user can now specify for each board whether it is used as a data source for the PCFarm. A different implementation is for the L0Calo, which is viewed as a single device from the RunControl even though it contains many TEL62 boards. The user can now specify the number of TEL62 boards for L0Calo used as a source, similar to the case of SRBs for the spectrometer, each of which manages multiple FE-cover boards.

New merger PCs were implemented in the PCFarm for Run 2 and had to be also added in the RunControl. The hardware node of the PCFarm computers uses the Farm Monitoring and Control (FMC) component of the JCOP framework. The FMC development has been halted, and the last version supports WinCC OA v3.16. With the planned update of WinCC OA to v3.19 in 2023/2024, an alternative PCFarm hardware control and monitoring tool is under development, and it will be implemented in the RunControl.

3.5.2 New control and monitoring tools and features

A new general-purpose monitoring tool, called EleNA62 Notifications, has been implemented during LS2. It is a manager that collects information from various reporters, and based on these, it creates or updates notifications which are shown to the user in a table on the SideScreen, an additional screen shown in the Control Room added after LS2, see figure 3.13. An action is expected from the shifter to resolve the issue reported. The currently implemented reporters are:

RunControl: Notifications that regard RunControl itself, such as, for example, a wrong shifter logged in the RunControl.

Magnets: Notify the shifter that a magnet current has moved from reference beyond a certain threshold. If it is a case that can be harmful to the experiment (for example, a beam passing through the detectors), an alarm sound is also played, and an additional notification is shown on the BigScreen.

Vacuum: Similar notifications and logic as in the case of magnets.

Boards: Can notify shifters if some TEL62 is not sending data.

CREAMs: Can notify shifters if some CREAM is not sending data or if a *swap* has been spotted by a swap monitor; more will be described below.

The information from each type of reporter is processed in a dedicated module, allowing straightforward scalability of the system in case additional reporters should be implemented.

The SideScreen also provides the trend of the number of L1 and L2 triggers and the number of events sent to mergers as a function of time, giving an estimate of the data-taking efficiency. It further shows the disk usage of the NA62Server machine and the NA62 EOS and the status of the timing RaspberryPi boards. These two boards are running independently so that in case one malfunctions, the data-taking can continue until access to the experimental area is possible.

The RunControl is interfaced with the NA62 OnlineMonitor, which reconstructs the collected data in real time and allows spotting problems that can compromise the data quality. The OnlineMonitor informs the RunControl if there are critical errors present, in which case the RunControl puts the node

of a given device in a **WARNING** state to notify the shifter. It also runs a so-called *Swap Monitor*, which reports a corruption of CREAM firmware, which leads to swapping in the mapping of readout channels.

Several additional monitoring tools were implemented directly in the graphical user interface and the BigScreen:

Choke monitor: Displays a list of boards that are sending choke signals, indicating that they are overwhelmed with data that they cannot process in real-time. This feature is currently implemented for TEL62s and CREAMs.

Missing source monitor: Displays a list of boards from which the PCFarm expects to receive data, but the data did not arrive.

Primitive monitor: Expected number of primitives for given devices for a reference T10 rate is set by the experts. The primitive monitor scales the numbers with the instantaneous beam rate and notifies the shifters if the primitive rate is out-of-threshold. This can, for example, indicate that a given board is malfunctioning and generating fewer primitives or that there is a noisy channel giving a surplus of primitives.

3.6 NA62 framework

The NA62 framework (NA62FW) is a collection of software modules developed by NA62 Collaboration for NA62 data analysis. The NA62FW is mainly written in C++ and Python and comprises four base modules:

NA62Tools Contains all the classes common to all other modules, such as the persistency libraries. It also contains the access code to the conditions database or external generators for the NA62MC module.

NA62MC A GEANT4-based module for the NA62 detector and processes simulations.

NA62Reconstruction A ROOT-based module which allows to reconstruct both the NA62 raw data and the output of NA62MC. It contains libraries for reconstructing hits and candidates for all subdetectors.

NA62Analysis A ROOT-based module containing NA62 analysis tools and a framework for new analyzer development.

Chapter 4

Search for feebly-interacting particles with NA62 in beam-dump mode

4.1	Beam-dump data sample and reconstruction	132
4.1.1	Uncertainty on the number of protons on tax	134
4.1.2	Reconstruction	135
4.1.3	T0 evaluation and additional calibrations	140
4.1.4	Bad burst definition	141
4.1.5	Trigger monitoring and efficiency	142
4.2	Exotic signal simulation strategy	143
4.3	Background simulation strategy	144
4.3.1	Prompt background	146
4.3.2	Combinatorial background	147
4.4	Search for di-lepton decay of dark photon	149
4.4.1	Dark photon MC	149
4.4.2	Event selection	152
4.4.3	Background determination for $\mu\mu$	162
4.4.4	Background determination for ee	166
4.4.5	Efficiency determination	172
4.4.6	Determination of expected sensitivity	179
4.4.7	Control and signal region opening	181
4.5	Search for hadronic decays of exotic particles	183
4.5.1	Event selection	185
4.5.2	Background determination	191

4.5.3	Efficiency determination	195
4.5.4	Determination of the expected sensitivity	197

In the beam-dump mode, the intense P42 400 GeV proton beam can be exploited to search for the production and decay of exotic particles at NA62. The beam-dump beamline setup described in 3.1 is used to minimize the flux of halo muons. The rates of MUV3 hits measured at the beginning of the beam-dump data-taking fully confirmed the prediction of Ref. [184]: as shown in figure 4.1, the optimal current is approximately -250 A, which for an MTR-type magnet corresponds to -0.6 T, i.e., one third of the nominal value used for B1C.




Figure 4.1: Rate of muons measured at the MUV3 during a current scan of the B2 magnet performed by the experts at the beginning of the beam-dump data-taking. The reference point, +770 A, corresponds to 1.82 T.

4.1 Beam-dump data sample and reconstruction

The data sample considered in this analysis was collected in 2021, with the experiment running in beam-dump mode. The sample is composed of 22 good runs,¹ collected in a period of almost ten days and corresponds to 13.9×10^{16}

¹An optimal length of one run in terms of size for the storage, the time it takes to process it, and for how long stable data-taking conditions can be kept, is about 1500 bursts.

integrated number of protons on TAX (POT). The full list of runs, with the corresponding number of bursts and POT, is shown in table 4.1.

Run ID	N. of bursts	POT $\times 10^{16}$	Average T10 $\times 10^{11}$
11069	1174	0.78	67.1
11070	420	0.29	70.7
11071	1390	0.97	69.6
11072	591	0.46	70.0
11074	495	0.33	68.8
11075	1415	0.98	69.6
11076	1115	0.77	69.3
11077	187	0.13	69.8
11078	1409	0.96	68.8
11079	633	0.43	67.7
11080	415	0.25	59.3
11082	398	0.26	65.0
11083	1444	0.93	64.6
11084	1297	0.82	63.5
11085	1416	0.92	65.0
11086	1126	0.72	65.2
11088	1343	0.85	63.3
11089	1082	0.67	64.2
11090	730	0.46	63.2
11092	1369	0.85	62.2
11093	1439	0.91	63.4
11094	281	0.17	62.2
22 runs	21185	13.9	65.7

Table 4.1: List of all good runs with the corresponding number of processed bursts and cumulative POT

Three trigger lines have been implemented:

- A control-trigger "neutral" mask, based on the LKr total energy release. It requires a total energy above 1 GeV, with one or more reconstructed clusters.
- Two physics-trigger masks:
 1. Q1/20, triggered by events with at least one hit in the NewCHOD, downscaled by a factor of 20;
 2. H2, triggered by events with two in-time hits in two different tiles of the NewCHOD.

The intensity of the proton beam has steadily increased soon after the beginning of the data-taking, so most of the data has been acquired with approximately 66×10^{11} protons per pulse, equivalent to more than 1.5 times the nominal intensity. At such intensity, the approximate rates of control, downscaled-Q1, and H2 triggers are 4, 14, and 16 KHz, respectively. The trigger rates as a function of the beam intensity are shown in figure 4.2. A fraction of 2% of the bursts are shorter than the standard, most often due to a CHOKE signal. This is accounted for while evaluating the total POT. The trigger rate has a good linearity with the measured T10 intensity. The error on the T10 evaluation is discussed in section 4.1.1.




Figure 4.2: Trigger rate as a function of the beam intensity for run 11078: (a) Q1/20 trigger, (b) H2 trigger. The trigger rate is computed using an effective spill length of 4.8 s. The 2% shorter bursts are due to a CHOKE signal at the L0TP and deviate from the linear trend.

4.1.1 Uncertainty on the number of protons on tax

The number of protons on TAX (POT) is evaluated from the measurement of beam flux, the so-called T10. T10 is provided by the titanium-foil secondary-emission monitor placed inside the target station (target box instrumentation upstream, TBIU). The expected uncertainty on the T10 measurement is at the level of 15–20% [197]. The T10 measurement is registered for each spill. The presence of spills shorter than the standard is accounted for by proportionally reducing the corresponding estimated T10.

The K12 beam intensity of K^+ mesons, measured by the NA62 detector, has been exploited to improve on the quoted uncertainty using both the kaon rate measured out-of-time from the KTAG and the rate of $K^+ \rightarrow \pi^+\pi^+\pi^-$ measured from control triggers. For comparison, the measurement provided by the ionization chamber close to the K12 beam-dump (the so-called Ar-

gonion) has been considered as well. Large systematic variations along the data-taking are observed in all the cases; see figure 4.3. A number of effects can qualitatively explain the observed differences. On one hand, the T10 measurement might have undergone a progressive degradation. On the other hand, the trigger rate was not entirely linear with the instantaneous intensity in 2018, while the so-called "bump" intensity peak at the beginning of the spill has been harming the kaon data-taking in 2021. Unlike part of the 2021 kaon data, the beam-dump data has been acquired without any rejection of the initial part of the spill. In conclusion, the $\pm 20\%$ uncertainty on T10, as quoted originally, is assumed.




Figure 4.3: Left: ratio of the POT evaluated from $K^+ \rightarrow \pi^+\pi^+\pi^-$ decays to that from the T10 measurement vs the run number, including 2017, 2018 and 2021 data-taking periods. Right: POT from the KTAG rate, the Argonion measurement and the $K_{3\pi}$ counts vs the run ID in 2021.

4.1.2 Reconstruction

A dedicated configuration of NA62Reconstruction for beam-dump data-taking has been used, which allows the reconstruction of tracks with momentum up to 400 GeV. Initially, the data set has been centrally reconstructed with the NA62FW, revision v3.1.5. A hardware limitation of the calorimeter read-out did not allow efficient reconstruction of high energetic clusters in v3.1.5., which affected signal acceptance for analyses with electrons and photons in the final state and had to be treated at the analysis level for the $A' \rightarrow ee$ search described in section 4.4.2. The impact of this limitation is best seen in the distribution of E/p of electrons with energy given by the reconstructed deposited energy E in an LKr cluster and the momentum p measured by the spectrometer and cluster time distribution with respect to the track time, see figure 4.4.




Figure 4.4: The E/p distribution of single electrons generated with a single electron MC (left) and the corresponding time distribution of the LKr clusters with respect to the track time (right). Dashed lines indicate the cuts typically used in the analyses $E/p(e) \leq 1 \pm 0.05$ and $|\Delta t| \leq 2$ ns.

Cluster size

The first limitation of the reconstruction comes at the software level and concerns the maximum number of cells $\max(N_{\text{cells}})$ per cluster taken into account, which can be seen in the left panel of figure 4.5 and is capped at about 100 in v3.1.5. This is given by the limit on the cluster radius $r_{\text{clus}} < 110$ mm. Given the size of the LKr cell $S_{\text{cell}} \simeq 2 \times 2 \text{ cm}^2$ this indeed gives $N_{\text{cells}} = \pi r_{\text{clus}}^2 / S_{\text{cell}} \simeq 95$. This choice is optimized for the kaon run and the maximum energy of 75 GeV [198]. However, as seen in the right panel of figure 4.5, it is insufficient for the beam-dump run.




Figure 4.5: The number of cluster cells as a function of momentum of the corresponding track in v3.1.5 (left) and when removing the limit (right) in a single electron MC sample.

Parametrization of the maximum radius of the reconstructed cluster as a function of the deposited energy E_{clus} as $r_{\text{clus,max}} = r_0 - \exp[(E_0 - E_{\text{clus}})/E_1]$ allows to capture better the shape of EM clusters. Values of $r_0 = 160$ mm, $E_0 = 350$ GeV and $E_1 = 75$ GeV allow to contain the majority of the EM cluster, see left panel of figure 4.6, and allow a better separation of EM clusters from hadronic clusters² based on E/p information than a flat cut, see right panel of figure 4.6. This positively impacts $e-\pi$ particle identification (PID) separation in analyses using cuts on E/p using data reconstructed with version v3.5.0 or newer.³ This modification, however, has no effect on the behavior observed in figure 4.4 as a much more significant change in the deposited energy is needed to achieve such a shift. As the Molière radius of liquid krypton is about $r_M \simeq 4.7$ cm [199], the majority of the cluster energy is contained only in several central cells and the energy mis-reconstruction happens already at the cell-level.

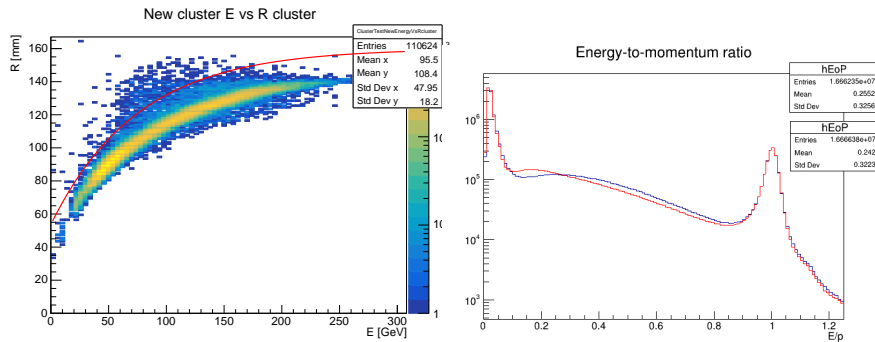


Figure 4.6: Left: the LKr cluster radius as a function of cluster energy in single electron MC. The red curve indicates the new cut applied at the reconstruction level for chosen values of r_0 , E_0 , and E_1 . Right: E/p spectrum in kaon mode before v3.5.0 (blue) and after (red). The peak around $E/p = 1$ corresponds to the electrons, around 0 to muons, and the continuous spectrum in between are pions. The change clearly results in a shift of the pion spectrum away from the electron peak.

LKr cell saturation

The *seed* of the cluster is defined as the most energetic hit of the cluster and defines the location of the cluster in the LKr plane and its time. When plotting the energy deposited in the seed per hit as a function of the momentum of

²Hadronic clusters are typically more spread than EM clusters.

³The spectrum shift towards the muon peak has no impact on the PID $\pi \rightarrow \mu$ mis-identification as MUV3 is used for the μ -tagging.

the corresponding track, the distribution in left panel of figure 4.7 is obtained indicating two features deviating the distribution from a linear dependence: capping of the maximum energy at about 70 GeV and about -60 GeV offset when $p_{\text{track}} > 150$ GeV/ c . The cause has been identified with the limitation of LKr CREAMs whose ADC can count up to $2^{14} = 16384$ per sample with eight samples per one readout time window. With the setting of about 3.5 MeV per count, the maximum energy read out per sample is about 57 GeV, depending on the cell gain. When the deposited energy is higher, the output value of the ADC remains the highest possible one, leading to an energy saturation of the cell.

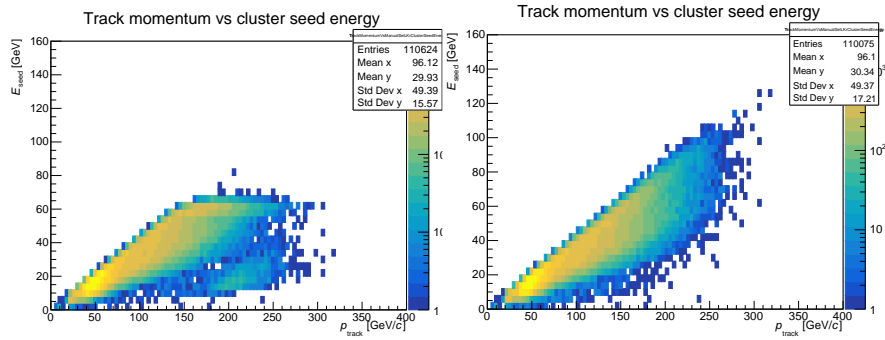


Figure 4.7: Seed energy as a function of track momentum of single electrons before (left) and after (right) v3.7.0.

When looking at the pulse structure of saturated cells at the sample level, one obtains the distributions shown in figure 4.8 in the 8 sample window where hits with one, two, and three saturated samples per pulse are shown. The hit reconstruction algorithm can then malfunction in two ways. Either it identifies the mean of the pulse correctly and returns approximately correct hit time and underestimates the hit energy which went above the maximum possible per sample, or when the search for the mean fails, the time distribution is shifted, and some samples in the pulse are possibly not counted in the hit, leading to an energy offset.




Figure 4.8: Examples of three pulses with one (left), two (center), and three (right) saturated samples. The pulses correspond to the seed hits from electrons in a beam-dump single-track MC. The dashed curve indicates what would be the expected distribution without saturation

A cell saturation reconstruction algorithm has been introduced in v3.7.0 to treat this hardware limitation at the software level. The strategy is to identify the saturated samples in each pulse and prepare a list of up to two samples adjacent from each side to the saturated samples. Another list of adjacent samples satisfying two conditions is prepared: the samples have to be located symmetrically on both sides; the energy has to decrease as one goes away from the saturation. This condition is set to allow the potential reconstruction of a scenario when two saturated pulses are identified in one 8-sample window.⁴ A fitting and correcting procedure that uses the saturated and adjacent samples follows:

- expected peak location t_{peak} is taken as the mean of the list of saturated and adjacent samples;
- expected peak energy E_{peak} is calculated from the adjacent sample energies E_{adjacent} and their distance from the peak $t_{\text{adjacent}} - t_{\text{peak}}$, giving a larger weight to the adjacent samples that are closer to the peak;
- energy of saturated samples E_{sat} is calculated⁵ from E_{peak} using their distance $t_{\text{sat}} - t_{\text{peak}}$;
- χ^2 is calculated for the adjacent and newly calculated saturated sample energies for different phases;
- procedure above (calculating t_{peak} , E_{peak} and correcting E_{sat}) is repeated until a phase for which χ^2 converges is found;

⁴This scenario has never been observed when testing the algorithm's performance.

⁵The standard `LKrRecoTools::PulseHeight()` method is used for this purpose.

- if χ^2 does not converge, the procedure is repeated with small correction $k(t_{\text{peak}}, t_{\text{sat}})$: $E_{\text{sat,corr}} = E_{\text{sat}}(1 + k)$ where k is varied minimizing χ^2 ; k and phase for which χ^2 is minimized is then used for calculating E_{sat} and t_{sat} .

The effect of the reconstruction procedure on the seed energy can be seen in the right panel of figure 4.7. The impact on the E/p distribution and the cluster timing can be seen in figure 4.9.




Figure 4.9: The E/p distribution after v3.7.0 of single electrons generated with a single electron MC (left) and the corresponding time distribution of the LKr clusters with respect to the track time (right)

4.1.3 T0 evaluation and additional calibrations

The detector time calibration (T0 correction evaluation) has been performed following the standard procedure implemented within the NA62TZ. Because of the relatively low rate of the triggers in the beam-dump configuration, at least 1000 bursts need to be processed to have enough statistics.⁶ For this reason, the time corrections cannot be evaluated for each run, as 40% of the runs collected have less than 1000 bursts. Four reference runs have been identified according to differences in the data-taking conditions: LOTP granularity and time-zero used for LOCalo, as defined in table 4.2.

⁶In kaon mode, 100 bursts per run are sufficient for the T0 correction evaluation.

Run ID	Reference run	LOTP granularity
11069-11076	11071	6.25 ns
11077-11080	11078	12.5 ns
11082-11090	11083	25.0 ns
11092-11094	11092	25.0 ns + New L0Calo T0

Table 4.2: Reference runs

All detectors have been enabled except the KTAG and the GTK, for which there were not enough statistics due to the data-taking conditions. The RICH, aligned to the trigger, is used as a reference detector for the T0 determination. The stability of the time response obtained has been verified with this procedure, and it is at the level of or better than 300 ps.

Additional calibrations, such as the LKr calibration or the GTK spatial alignment, cannot be evaluated from beam-dump data; the results from the last run taken in standard conditions before the beam-dump data-taking (Run 011048) are used.

4.1.4 Bad burst definition

The machinery to flag bad bursts has been implemented. Several analyzers used for kaon runs have also been adapted for beam-dump data, setting various input parameters according to the beam mode. The list of adapted analyzers is the following:

- ANTI0Efficiency: based on single muon tracks backward extrapolated at the ANTI0 plane;
- GigatrackerDQMonitor (only worth to address hardware malfunctioning);
- CHODEfficiency: based on single muon tracks;
- CHANTIMuonEfficiency: completely standalone efficiency for penetrating tracks;
- NewCHODEfficiency: based on single muon tracks;
- MUV3Efficiency: based on single muon tracks.

Two specific analyzers have been prepared: one to check the data quality from the trigger efficiency point of view called `BeamDumpMonitoring` (see section 4.1.5 for details) and one to check the RICH response using halo muons

called `RICHBeamDumpMonitor` which derives the efficiency for charged particles reconstructed with STRAW and (New)CHOD.

For this analysis, the following systems are checked: LAV, STRAW, CHOD, NewCHOD, RICH, MUV3, L0Calo and L0NewCHOD. All other systems either cannot be checked as the corresponding analyzer relies on the selection of kaon decay modes or are not of interest for the present analysis. In total, 0.4% of all bursts have been rejected because they were flagged as bad, of which 25% due to problems in the L0-trigger and 75% due to problems in the RICH, CHOD, or STRAW.

4.1.5 Trigger monitoring and efficiency

A monitor of the trigger response and a first estimate of its efficiency is obtained using a dedicated data quality analyzer based on a simple selection. Tracks are selected based on quality conditions described in 4.4.2. The track time is defined as the time of the CHOD-associated candidate if it exists; otherwise, the time of the associated NewCHOD candidate is used. For the study of the Q1 (H2) trigger efficiency, single-track (two-track) events are selected, and the presence of isolated NewCHOD tiles associated with each track is required, using time and spatial information. The efficiency is measured on control triggers, so a further condition is applied, asking for an associated LKr cluster for each track with energy above 1.5 GeV.

The Q1 and H2 trigger efficiencies as a function of the tile in the X-Y plane are shown in figure 4.10.




Figure 4.10: Q1 (left) and H2 (right) trigger efficiency from run 11078 as a function of the Y and X coordinates of the tile (or tiles) under consideration.

4.2 Exotic signal simulation strategy

The basic set-up of the signal MC for exotic particles has been documented in Ref. [200]. Since then, however, several updates and improvements have been integrated into NA62MC. Regarding the exotic MC itself, namely 3- and 4-body exotic particle decays have been implemented, following the implementation in Ref. [1] described in section 2.2.2, allowing also non-flat Dalitz distributions. All production and decay channels from sections 2.1 and 2.2.2 have been implemented with additional custom settings such as choice of flat and non-flat 3-body decays, scanning through exotic particle masses with a logarithmic step or choosing between derivative and Yukawa coupling for exotic particle production via FCNC transitions.

The kinematics for exotic particle production in meson decays is handled within the NA62MC, using meson distributions obtained with PYTHIA as described in section 2.1.3 as an input. For direct and mixing production, the exotic particle spectra obtained from an external generator described in sections 2.1.1 and 2.1.2 are used and located in NA62Tools.

As illustrated in section 2.2.2, each decay channel is characterized by at least two parameters: the exotic particle mass and the decay width. A Monte Carlo production spanning such bi-dimensional parameter space would be extremely demanding regarding CPU time and storage space. At the generation level, while scanning on fewer exotic particle mass points, an infinite exotic particle lifetime has been assumed instead, and a decay within the Z range 102–180(182) m has been forced.⁷ At the analysis level, for each value of coupling considered, the events generated are re-weighted to account for the probability of reaching and decaying within the chosen Z range to match the distribution of flight paths to the expected exotic particle decay length. The full weight used to estimate yield for each mass and coupling can be expressed as:

$$N_{\text{exp}} = N_{\text{POT}} \cdot \chi(pp \rightarrow X) \cdot \text{BR}(X \rightarrow \text{SM}) \cdot P_{\text{RD}}(\Gamma) \cdot \varepsilon_{\text{acc}} \quad (4.1)$$

where

- POT is the number of protons on TAX;
- $\chi(pp \rightarrow X)$ is the exotic particle production probability;

⁷Range 102–180 m was used for the leptonic dark photon decays while 102–182 m is used for the hadronic channels.

- $\text{BR}(X \rightarrow \text{SM})$ is the branching fraction of the exotic particle decay into given observable final states;
- $P_{\text{RD}}(\Gamma)$ is the probability of the exotic particle reaching the NA62 fiducial volume and decaying therein;
- ε_{acc} is the selection efficiency defined as

$$\varepsilon_{\text{acc}} = \frac{\sum_{i=1}^{N_{\text{sel}}} w_i}{\sum_{i=1}^{N_{\text{gen}}} w_i} \quad (4.2)$$

where N_{sel} is the number of all selected MC events, N_{gen} is the number of all generated exotic particles, and w_i is a weight that is used to correct the initially infinite lifetime as

$$w_i = \frac{1}{\gamma_i \beta_i c \tau} \times e^{-\frac{Z_i^{\text{decay}} - Z_i^{\text{prod}}}{\gamma_i \beta_i c \tau}} \quad (4.3)$$

Variables $P_{\text{RD}}(\Gamma)$ and ε_{acc} provide equivalent information to E_{det} of equation 1.94 which also allows interfacing the output of NA62MC to the interpretation module `ALPrescale` of ALPINIST framework.

4.3 Background simulation strategy

A solid estimate of the expected background is paramount for a credible search for a new physics signature. Two methods of first-principle Monte Carlo evaluations were tested [184, 201]; however, in order to reach the statistics equivalent of the order of 10^{17} POT, large oversampling of the order of $\mathcal{O}(10^3)$ is required, thus inducing non-physical correlations. The techniques used stem instead from a mixture of data-driven and Monte Carlo methods.

A number of data "side-bands" within the selection algorithm, which will be discussed in section 4.4.2, have been exploited to give indications about the expected background in dump mode:

- Opposite-charge events with e^+e^- or $\mu^+\mu^-$ PID assignments with reconstructed impact point away from the actual impact point at TAX (the signal and control regions in the CDA_{TAX} vs Z_{TAX} plane are defined properly later in section 4.4.2);

- Events with same-charge vertices, $\mu^+\mu^+$ or $\mu^-\mu^-$ ⁸;
- Events with $e\text{-}\mu$ PID assignments (with same or opposite charges).




Figure 4.11: Time difference of the two selected tracks. The tracks are sorted according to the charge and index within the STRAW reconstruction. The panels refer to: (a) $\mu^+\mu^-$, (b) e^+e^- , (c) μ^+e^- , (d) μ^-e^+ , (e) $\mu^+\mu^+$, (f) $\mu^-\mu^-$

The time difference between the two selected tracks for the various control samples is shown in figure 4.11. The following conclusions can be derived:

1. Events with at least one electron or positron are dominated by in-time tracks, mostly promptly emitted by secondary interactions of halo particles within traversed material upstream of the decay volume;

⁸The statistics of the corresponding e^+e^+ or e^-e^- final states is too poor to derive any meaningful information.

2. Di-muon events are approximately flatly distributed in time and result mostly from a combinatorial random pairing of incoming halo particles.

The two background sources, combinatorial pairing, and prompt emission, have been evaluated separately.

4.3.1 Prompt background

This analysis uses the data-driven backward MC concept for simulating the in-time background to overcome the abovementioned issues with a first-principle background MC. A backward simulation is performed starting from an observed muon halo at a given plane, conceptually inverting the time flow. As a result, one retrieves the corresponding halo at another plane (more upstream), where measurement is impossible. Practically, this is achieved with PUMAS [202], a standalone tool often used in muography studies, interfaced with GEANT4. The input consists of the halo observed in data in the form of single-track kinematics (x, y, P_x, P_y, P_z and charge) at the spectrometer reference plane ($Z = 180\text{m}$). PUMAS then propagates backward each muon, increasing its energy depending on the amount of material traversed. The backward simulation is stopped at the entry of BEND5 ($Z = 92\text{m}$), where particles are scored in a G4BL-like tree. The resulting halo is then used as a particle gun for a forward MC, which uses GEANT4, and scoring at the spectrometer reference plane.

The obtained distribution has discrepancies: the simulated halo is more spread than the original one, and the high momentum component is enhanced in the simulation. The increased spread can be attributed to the different treatment of muon multiple scattering in PUMAS and GEANT4. The increase in average momentum in MC is due to the asymmetric shape of the energy loss/gain differential cross-section as a function of the energy. In order to cure these issues, an unfolding technique was applied using `ROOUnfold`. The unfolded quantities are the momentum magnitude (P) and transverse radius (R) distributions of the halo in the form of a 2D histogram ($PvsR$). The transfer matrix is constructed with $PvsR$ at BEND5 as input and the $PvsR$ of the simulated halo at spectrometer entry as output for each charge. `ROOUnfold` inverts the transfer matrix and applies it on the $PvsR$ of the data halo at spectrometer entry. The result is the $PvsR_{\text{Unfold}}$ distribution at BEND5 that, when used as input for NA62MC, in the absence of other relevant variables, would reproduce the distributions observed in data. Each muon is then re-weighted based

on individual P and R as

$$w = w_i \times \frac{PvsR_{\text{Unfold}}}{PvsR_{\mu}} \quad (4.4)$$

where w_i is the weight after the backward transport and $PvsR_{\mu}$ is the population in the respective bin of the 2D histogram. The data/MC comparison of other quantities (transverse positions, direction of tracks in the XZ and YZ plane) shows an agreement to within a relative deviation of $\pm 50\%$, which is considered an absolute systematic error for the prompt MC estimates.

The background simulation is then performed in two stages:

- Stage 1, from BEND5 to the spectrometer exit plane. The input consists of the backward-propagated muon halo, and each muon is oversampled four times. A filtering condition is applied to each event. Two separate filter conditions for the particles to be scored at the spectrometer exit plane are applied:
 - 1a** requiring at least one $e^+/e^-/\gamma/\pi^+/\pi^-/p/n/K^+/K^-/K_L^0$ with momentum at least 1 GeV/c;
 - 1b** requiring two muons independent of their charge.
- Stage 2, from the spectrometer exit plane to the end of ECN3 with the beam consisting of all particles scored at the spectrometer exit plane in the previous stage.

The division in stages is performed because a shower developed at $Z < 180\text{m}$ does not necessarily imply an event with final states in the downstream detector acceptance, while the complexity of the simulation grows with the increasing number of particles. The MC for each stage is reconstructed individually, and the result is merged.

4.3.2 Combinatorial background

To evaluate the expected combinatorial background, events with single tracks in time with the trigger are selected. The trigger condition is established by the presence of a Q1 primitive and the absence of the H2 primitive. No condition is applied related to the LKr L0 primitive. In this way, the starting sample is unrelated to the signal sample, and no bias is introduced by the potential presence of the LKr primitive. For the single-track selection, the STRAW reconstruction, the association to downstream detectors, and the PID assignments

are performed exactly as for the signal selection, the only difference being the apparent absence of a vertex condition. The track reconstructed charge and the PID assignment are saved in a dedicated output tree, as well as the LAV-veto and ANTI0 acceptance conditions.

Pairs of the selected single tracks are then built, simulating a random superposition within ± 5 ns in the same event. The vertex condition is then applied, and the total momentum reconstructed at the vertex is used to evaluate the backward extrapolation at the TAX. Vertex of any charge combination is reconstructed.

Nominally, the starting single-track, Q1-based sample is supposed to be downscaled at the L0 trigger by $D = 20$. Therefore, in each spill, the number of "simulated" pairs considered is given by:

$$N_{\text{pair}}^{\text{sim}} = \sum_{i,j} \frac{(R_i \times T_{\text{spill}})(R_j \times T_{\text{spill}})}{2 \times D^2} \times \varepsilon_{\text{Q1}}^2, \quad (4.5)$$

where the indices i and j run over the spill time, $R_{i,j}$ are the instantaneous single-track rates, T_{spill} is the total spill duration and ε_{Q1} is the Q1-trigger efficiency. The number of expected track pairs within a time window ΔT is given by:

$$N_{\text{pair}}^{\text{exp}} = \sum_{i,j} \frac{(R_i \times T_{\text{spill}})(R_j \times \Delta T)}{2} \times \varepsilon_{\text{H2}}, \quad (4.6)$$

where ε_{H2} is the H2-trigger efficiency. By comparing the two equations above, the simulated pair should be weighted by a factor

$$w = \frac{\Delta T \times D^2}{T_{\text{spill}}} \frac{\varepsilon_{\text{H2}}}{\varepsilon_{\text{Q1}}^2}. \quad (4.7)$$

With an error well below the per-mil level, the Q1 trigger is always fired, given an H2 condition. Therefore, the trigger efficiency factor reduces to the probability that the H2 condition is fired, given that two tracks firing Q1 are within ΔT .

A complication arises due to the downscale factor in equation (4.7). In order to avoid any induced bias, the single-track events with both Q1 and LKr primitives are kept. Since LKr-based triggers are not downscaled, using $D = 20$ in the above equation leads to overestimating the number of events. Quite symmetrically, since events with Q1 and H2 primitives are rejected (which again are not downscaled at L0), the tendency is to underestimate the number

of expected pairs. Studying the trigger correlations in the four reference runs, in total, although the two effects are opposite in sign, the LKr-induced one is larger than that due to H2. The bias induced on D by Q1 !H2 LKr events is $B_{\text{LKr}} = +(15.5 \pm 0.5)\%$, while that induced by the rejection of Q1 !H2 events is $B_{\text{H2}} = -(5.7 \pm 0.3)\%$ (the errors quoted are statistical). To correct for these effects, the effective downscale is redefined as

$$D_{\text{eff}} = D \times \left(\frac{1}{1 + K \times B_{\text{LKr}} + B_{\text{H2}}} \pm (1 - K) \times B_{\text{LKr}} \right), K = 0.5, \quad (4.8)$$

to be used in place of D in equation (4.7). Only half of the LKr bias is corrected while quoting half of the LKr bias as a systematic error on D_{eff} since a fraction of the events with LKr and Q1 primitives does not lead to LKr-triggered events due to the large time jitter of the LKr at L0. Independent acquisition of the trigger primitive information has been disabled before the beginning of the beam-dump data-taking; therefore, a more precise estimation is not possible. Since Q1 and H2 are both physics masks, no further consideration is needed concerning B_{H2} . With this approach, the relative systematic error is approximately 7.5% and 15% on D_{eff} and w , respectively. The statistical error on the expected number of background events is negligible with respect to the systematic error.

4.4 Search for di-lepton decay of dark photon

The first two analyses performed on the data sample collected in 2021 were searches for exotic particle di-lepton decays with the focus on dark photons produced in light meson decays and via bremsstrahlung for which NA62 has the largest sensitivity [3,4]. This result also allows the reinterpretation for other exotic particles, such as ALPs or dark scalars produced in B meson decays. The di-muon result has been published in [133].

4.4.1 Dark photon MC

Generally, the considered DP emission can proceed via a bremsstrahlung process or from the decay of secondary mesons from the proton shower in the dump. The two production mechanisms are simulated and studied separately, as the kinematics, and thus the acceptance probability of the Dark Photon decay products depends on the production channel. This point is illustrated in figure 4.12: For a fixed set of exotic particle masses, it shows the spectra of

longitudinal and transverse momenta. Thanks to the hardness of the spectra and partly to the overall yield, *direct* DP production is the most relevant channel in this analysis at the POT at hand, as will be shown below. For the production in meson decays, the relevant mesons are π^0 , η , η' , Φ , ω and ρ . For comparison, figure 4.12 also includes the spectra for ALP emission in B meson decay.




Figure 4.12: Illustration of kinematics for fermionic-coupled axion-like particles (ALPs) and Dark Photons for two different masses and selected production modes.

The validation of the implementation of the DP signal production and decay in NA62MC has been done against a toy MC. The latter is documented in [203]. The validation of NA62MC against the toy has been presented in [204].

For the present study, signal samples have been generated from revision v3.1.6 in NA62 with a central production using the computing grid of NA62. As will be shown below, with the available POT, the expected sensitivity extends up to dark photon masses of approximately 700 MeV, within a range of coupling constants $2 \times 10^{-7} < \varepsilon < 2 \times 10^{-5}$. The expected mass sensitiv-

Decay mode	Production mode	N_{mass}	N_{evts}
e^+e^-	bremsstrahlung-like	140	$\simeq 14 \text{ M}$
	meson-mediated	140	$\simeq 14 \text{ M}$
$\mu^+\mu^-$	bremsstrahlung-like	98	$\simeq 13 \text{ M}$
	meson-mediated	98	$\simeq 13 \text{ M}$

Table 4.3: Summary of grid-based, TAX-generated, dark photon signal Monte Carlo productions.

ity region is such that the contribution of hadron resonances to the production channels is modest, which is a significant simplification. For the MC generation, fixed-width (5 MeV) mass steps were used. The two decay modes, e^+e^- and $\mu^+\mu^-$ are generated separately for each production mode. For the di-electron channel, dark photon masses in the range 5–700 MeV are considered, corresponding to $N_{\text{mass}} = 140$ mass points. For the di-muon channel, the mass range is 215–700 MeV, corresponding to $N_{\text{mass}} = 98$. A first grid-based production has been performed allowing the dark photon to be produced either at the TAXes (as is the case with a beam-dump setup) or at the T10 target (as is the case parasitically to the standard-mode setup, not of interest for the present document)⁹ A second grid-based production only generated dark photons at the TAXes. For dark photons generated at the TAXes, the number of events per mass point available is $N_{\text{evts}}/N_{\text{mass}} \gtrsim 120\,000$. The statistics are summarized in table 4.3.

With a simple geometrical selection ensuring that the DP decays in the fiducial volume and its daughters are in the LKr acceptance and 1.4×10^{17} protons on TAX are acquired, the MC generation suggests a sensitivity (90% confidence level in absence of background) extending up to 700 MeV and including a coupling interval $3 \times 10^{-7} < \epsilon < 2 \times 10^{-5}$, see figure 4.13.

In the case of signal observation, the di-lepton invariant mass can be used as an estimate of the DP mass. The experimental resolution on the di-lepton invariant mass varies with the DP mass and production mode. For di-muon decays, the invariant mass resolution is 1 MeV at the kinematic threshold and raises to more than 5 MeV (3 MeV) for the direct (meson-mediated) production at 700 MeV, which is well understood in terms of DP kinematics. An interpolation has been performed to continuously explore the mass range with the sensitivity needed; see section 4.4.6.

⁹Obviously, the Monte Carlo truth related to the production point can be exploited to keep only TAX-generated DPs, as needed for beam-dump-mode signal studies.




Figure 4.13: Sensitivity of NA62 to dark photon production in meson decays and via a bremsstrahlung process separated by decay mode (left) and by production mode (right), 90% CL and 0-background assumption. The sensitivity accounts only for the effects of geometrical acceptance of NA62 and 1.4×10^{17} protons on TAX. The grey underlying exclusion is the one adapted by the PBC and originally based on Ref. [205].

4.4.2 Event selection

Considering a $A' \rightarrow \ell\ell$ decay in the NA62 FV, the signal signature is defined by a lepton-antilepton reconstructed (secondary) vertex within this volume and a primary vertex between the reconstructed momentum of the lepton-antilepton pair and the proton beam, close to the TAXes region. The event selection basis is common for both ee and $\mu\mu$ channels, with the criteria that are different for the former from the latter described at the end of each section.

Event selection algorithm

The first part of the event selection consists of selecting good tracks reconstructed by the STRAW spectrometer. Each track must satisfy the following conditions:

- it is reconstructed with good-quality ($\chi^2 < 20$ and a difference between momentum before and after the track fit $< 20 \text{ GeV}/c$);
- it has reconstructed hits in at least three STRAW chambers;
- it is in the geometrical acceptance of the (New)CHOD, LKr, MUV1, MUV2, MUV3 detectors and outside the geometrical acceptance of LAV12;
- it is away from any other track by at least 20 mm at the first STRAW chamber to reduce the probability of sharing the same STRAW hits, as is typical of fake tracks.
- each track extrapolated at the LKr front plane must be at least 200 mm distant from any other track, so to avoid possible cluster overlapping,

- at least one NewCHOD candidate must be associated with the track to match the trigger condition (described in 4.1).

The track time is defined as the time of the associated CHOD candidate, if it exists, or otherwise of the associated NewCHOD candidate (figure 4.14). Only tracks passing the above conditions and within a time window of ± 5 ns with respect to the trigger time are selected.




Figure 4.14: CHOD candidate (left) and NewCHOD candidate (right) time with respect to the trigger time.

The association of each spectrometer track with candidates in other detectors ((New)CHOD, LKr, MUV3) is done using the tools present in the NA62FW accessed via the `DownstreamTrack` class. Only events with two or more tracks passing the above criteria are kept.

The vertex reconstruction is performed using the `VertexLSF` tool [206] based on a least-square vertex fitting method. Inputs to this algorithm are the reconstructed STRAW track fit parameters (track slopes, positions, and momenta at the first STRAW station) and their covariance matrices. Outputs of the tool are the vertex position and the fitted tracks momenta. The tool is used, enabling the particle extrapolation through the residual field in the blue tube. Events with exactly one reconstructed vertex with $\chi^2 < 30$ and a vertex Z coordinate in the range [105,180] m are selected. No condition on the total charge at the vertex is applied at this point to allow all possible combinations.

The mean time of the two selected tracks forming the vertex defines the reference time.

Particle identification for $\mu\mu$ analysis The event selection proceeds with PID conditions, using information from the STRAW, LKr, and MUV3 detectors, to distinguish μ^\pm , e^\pm , and π^\pm tracks. For each track, the ratio E/p is calculated using the track momentum p and the energy E of the geometrically closest LKr cluster, within 50 mm of the track impact point at the LKr front plane and 10 ns of the track time. Further, the following PID requirements are applied.

Muon (μ^\pm): Exactly one LKr cluster is allowed within a 50 mm radius around the track impact point and within $\pm 5 \sigma_{LKr}^\mu$ of the track time, where σ_{LKr}^μ is the time resolution of a muon-like cluster (figure 4.15 (right)). A correction for the residual time misalignment is also considered according to the mean values shown in figure 4.15 (left). If there is an associated cluster, it is required that $E/p < 0.2$. Exactly one signal must be geometrically associated with the track in any MUV3 tile and within 5 ns of the track time.

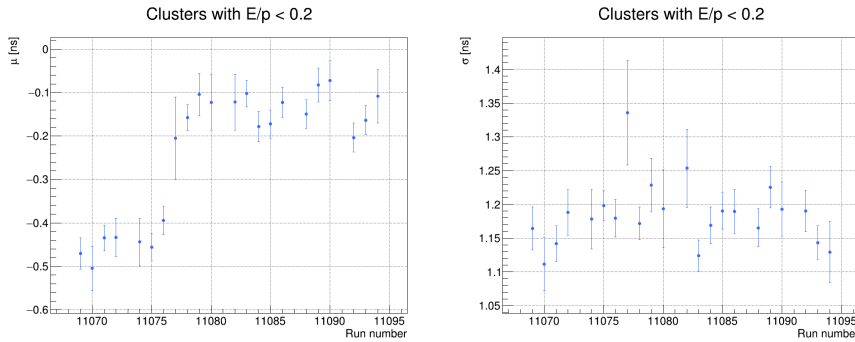


Figure 4.15: Gaussian-fit parameters for the distribution of the difference between the times of the STRAW track and of the associated LKr cluster as a function of the run number: Mean value (left) and standard deviation (right). The track and the associated LKr cluster satisfy the condition $E/p < 0.2$.

Positron/electron (e^\pm): Exactly one LKr cluster is required within a 100 mm radius around the track impact point within $\pm 5 \sigma_{LKr}^e$ of the track time, where σ_{LKr}^e is the time resolution of an electron-like cluster (figure 4.16 (right)). A correction for the residual time misalignment (figure 4.16 (left)) is also taken into account. The associated cluster must have $0.95 < E/p < 1.05$. There must be no signal in MUV3 geometrically associated with the track and within 5 ns of the track time.

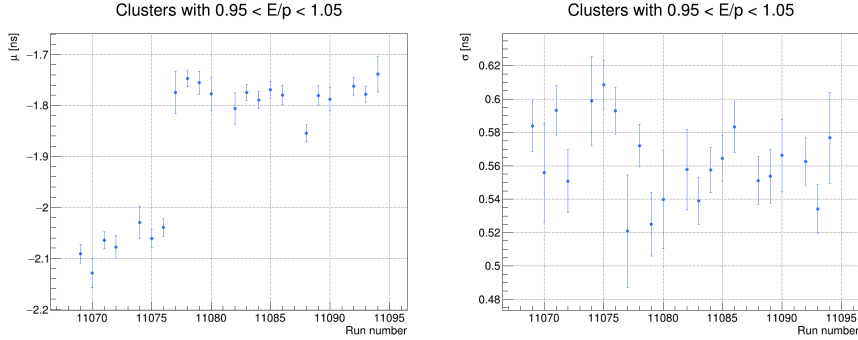


Figure 4.16: Gaussian-fit parameters for the distribution of the difference between the times of the STRAW track and of the associated LKr cluster as a function of the run number: Mean value (left) and standard deviation (right). The track and the associated LKr cluster satisfy the condition $0.95 < E/p < 1.05$.

Pions (π^\pm): The conditions are the same as for e^\pm , but the associated cluster must have $0.2 < E/p < 0.9$.

For the $\mu^+\mu^-$ analysis, no distance cut with respect to the LKr dead cells is requested for the LKr associations, while the two muons have to have a different MUV3 tile associated with each track. On the contrary, at least 30 mm distance of the track projection on the LKr plane from a dead cell is required.

Particle identification for ee analysis The ee analysis uses a custom LKr association instead of the `DownstreamTrack` association tool in order to increase signal acceptance. A momentum-dependent search radius is defined as

$$R_{\text{direct}} [\text{mm}] = \frac{1600}{p [\text{GeV}/c]} + 20. \quad (4.9)$$

All clusters within R_{direct} and within 10 ns from the track time are labeled *direct* matches. Note that for $\mu\mu$, the radius was fixed at 50 mm, and at most, one cluster was allowed in the circle around the track impact point. Such a requirement introduces inefficiencies for e^\pm because e^\pm traversing material before the MNP-33 magnet can radiate photons that can reach the LKr. High-momentum electrons are only slightly bent by the magnet; hence, the photon cluster can appear in the search radius of the track. This happens often for $p > 150 \text{ GeV}/c$, as shown in figure 4.17. To mitigate this issue, another circle is defined, this time around the projection of the track at the LKr front face

using the position and direction before the MNP-33 magnet. All clusters with $E > 1$ GeV within the radius

$$R_{\text{radiative}} [\text{mm}] = \frac{1600}{p [\text{GeV}/c]} + 40 \quad (4.10)$$

are labeled *radiative* matches. Tracks having more than one *direct* matches that are not *radiative* are discarded. Finally, the associated cluster is the one with the highest energy among the geometrically closest two *direct* matches.

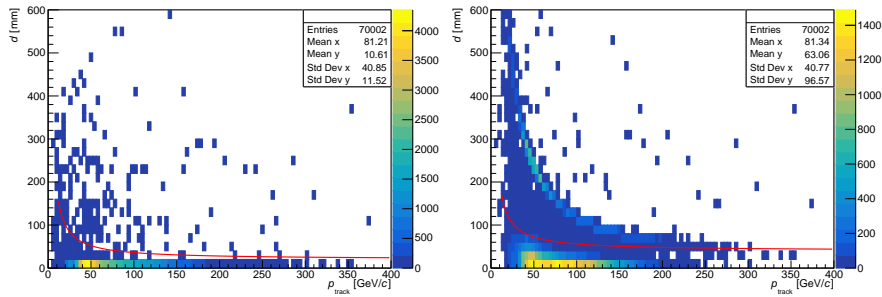


Figure 4.17: Distance from the projection of the track after (left) and before (right) magnet trajectory onto the LKr to the geometrically closest cluster as a function of the track momentum. Red curves show the corresponding search radii (left equation (4.9) and right equation (4.10)) for associating the *direct* and *radiative* clusters to the track.

Another potential limitation of the original PID algorithm comes from the cluster mis-reconstruction for high momentum electrons in data reconstructed with v3.1.5, as described in section 4.1.2. This issue can be largely mitigated at the analysis level. The clusters with wrongly calculated pulse time often have seed energies lower than the energies of the neighboring cells; see left panel of figure 4.18. Clusters identified this way can be cured by applying a correction derived from the fraction of the electron energy expected to be in cells other than the seed. The expected behavior is then restored for most clusters, as seen in figure 4.18 (right). Moreover, the time of these clusters can also be corrected using the time of the cell with the highest energy instead of the seed time. A tuning of the e and π PID condition allows recovering most of the clusters not perfectly corrected:

Positron/electron (e^\pm): It is required that exactly one *direct* cluster is found and that $Q_e < E/p < 1.05$ with

$$Q_e = 1 - 0.0013 \times \max(p [\text{GeV}/c] - 100, 50). \quad (4.11)$$

There must be no signal in MUV3 geometrically associated with the track and within 5 ns of the track time.

Pions (π^\pm): It is required that exactly one *direct* cluster is found and that $0.2 < E/p < Q_\pi$ with

$$Q_\pi = 0.95 - 0.013 \times \max(p [\text{GeV}/c] - 100, 50). \quad (4.12)$$

There must be no signal in MUV3 geometrically associated to the track and within 5 ns of the track time.




Figure 4.18: Difference in the ratio E/p of the associated cluster seed and the highest energetic neighbor as a function of track momentum as obtained with single-electron MC before applying the cluster energy correction (left). E/p for the treated clusters (right). The modified E/p conditions for the ee analysis are indicated by red lines. The distribution and the PID conditions can be compared to those of figure 4.4 (left).

Additional vetoes

LAV veto: To reduce the possible selection of vertices derived by the interaction of incoming particles, mostly muons, with the material in the LAV stations, an LAV veto condition is applied: no signal from any LAV station must be present within 10 ns from the reference time (the LAVMatching tool is used for this purpose). The distribution of the

difference between the LAV hits and the reference time is shown in figure 4.19 left. For the time window used, a random veto $\epsilon_{RV}^{LAV} = (2.5 \pm 0.1)\%$ is expected.




Figure 4.19: Left: Time distribution of the LAV hits for the reference time (data from run 011092). Right: Time distribution of the ANTI0 candidates for the reference time (data from run 011092)

ANTI0: Each track must be in the ANTI0 acceptance to efficiently exploit the in-time activity in the ANTI0. A track has an ANTI0 match if the ANTI0 hit position is within 200 mm from the track extrapolated position at the ANTI0 downstream face and the ANTI0 hit time is within 5 ns from the track time. During the beam-dump data-taking, one channel of the ANTI0 was time-misaligned. To correct this, the times of hits in this channel (ID=131) are shifted back by 6.647 ns. This algorithm's efficiency and random veto have been evaluated on single tracks (muons) in the data sample. They amount to 98% and 0.3% respectively. An efficiency-tile map has been obtained, as the probability that a minimum ionizing track traversing a given tile satisfies the ANTI0-veto condition, see figure 4.20. The ANTI0 vetoes a vertex if at least one track has an ANTI0 hit matched.¹⁰

¹⁰The ANTI0 veto is used only in ee analysis to suppress the prompt background defined in section 4.3.1.




Figure 4.20: Efficiency map for the ANTI0-veto condition from a sample of single-muon data events triggered by the Q1 condition.

Signal and control region definition Signal events are due to the decay of exotic particles produced by the interaction of protons in the TAX. To identify this signature, the total momentum evaluated by the sum of the two-track momenta at the vertex is backward extrapolated in a straight line. The point of minimum approach between the beam direction at the TAX entrance given by 3.1 and the total momentum is calculated: its longitudinal coordinate, Z_{TAX} , is plotted against the distance of closest approach CDA_{TAX} in figure 4.21 for signal Monte Carlo events (see section 4.2 for details related to the MC production).




Figure 4.21: Distance of closest approach between the beam at the TAX entrance and the total momentum of the two leptons with respect to the longitudinal position, for signal Monte Carlo events. The grey box represents the defined signal region.

For the signal, the mean value of Z_{TAX} is approximately 23 m, and the related standard deviation is 5.5 m. The standard deviation on CDA_{TAX} is 7 mm. The signal (SR) and control regions (CR) are defined around the expected pro-

duction point given by the point of smallest distance of closest approach. The analyses are performed as *blinded analyses*; therefore, both the SR and CR are kept masked until the analysis is finalized.

For the $\mu\mu$ analysis, the SR, including ± 3 standard deviations in each of the plot axes, is defined as

$$\text{SR} : 6 < Z_{\text{TAX}} < 40 \text{ m} \ \& \ \text{CDA}_{\text{TAX}} < 20 \text{ mm}. \quad (4.13)$$

Analogously, the CR to be exploited for validation of the background estimate is defined as

$$\text{CR} : -4 < Z_{\text{TAX}} < 50 \text{ m} \ \& \ \text{CDA}_{\text{TAX}} < 150 \text{ mm} \ \text{and not in SR}. \quad (4.14)$$

An improved version of the SR definition is used in the ee analysis to better account for the signal distribution. The region is obtained by evaluating the lines equally populated by the signal distribution. The black solid line of figure 4.22 is a constant-population contour corresponding to a fraction of 10^{-3} of the distribution maximum and includes $\sim 98\%$ of the signal events. The red dashed line is an ellipse defined to surround the constant-population contour and is used to define the SR for the ee analysis

$$\text{SR} : \left(\frac{Z_{\text{TAX}}[\text{m}] - 23}{12} \right)^2 + \left(\frac{\text{CDA}_{\text{TAX}}[\text{mm}]}{30} \right)^2 < 1, \quad (4.15)$$

which contains $\gtrsim 98\%$ of the signal events for both production modes. The CR is kept the same.

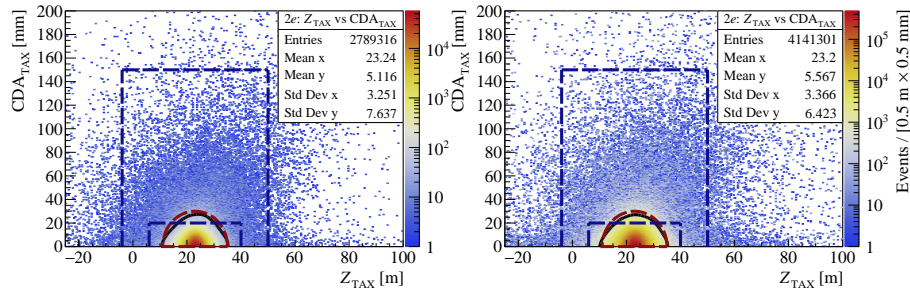


Figure 4.22: Distance of closest approach between the beam and the total momentum of the two leptons (y-axis) as a function of the longitudinal position (x-axis), for signal MC events using all A' masses: left (right) refers to bremsstrahlung (meson-mediated) production modes. The blue dashed contours represent the signal and control regions defined in equations (4.13) and (4.14). The black solid line is a constant-population contour, and the red dashed contour is an ellipse defined to surround it.

Modified FV definition

While in the $\mu\mu$ analysis, the FV is given simply by the range $102 < Z_{\text{VTX}} < 180$ m, the ee uses a modified FV defined to minimize the background coming from muon interactions in the LAV 1-8 stations and to reject as little signal as possible. This is approximately achieved by the union of two cones along the z -axis, both with the vertex at the beginning of the TAX given by (3.1) but bases at the downstream face of the LAV5 and LAV8 stations with the respective inner radii ($Z_{\text{LAV5}} = 152.393$ m, $R_{\text{LAV5}} = 536.5$ mm and $Z_{\text{LAV8}} = 180.923$ m, $R_{\text{LAV8}} = 767.5$ mm). Pairs created by muon interactions in LAV5 can be reconstructed in this volume due to resolution effects. Consequently, an optimized version of the above is used, namely

- Cone 1, with vertex position given by (3.1) and the base at $Z_1 = 157$ m with radius $R_1 = 520.5$ mm. This cone is truncated longitudinally, requiring $105 \text{ m} < Z < 157 \text{ m}$.
- Cone 2, with vertex position given by (3.1) and the base at $Z_2 = 180.92$ m with inner radius $R_2 = 767.5$ mm. This cone is also truncated longitudinally, requiring $157 \text{ m} < Z < 180 \text{ m}$.

A 2-track vertex is kept if it is reconstructed within the volume defined by Cone 1 \cup Cone 2. An *upstream* side-band is also defined as the region outside of the FV of the longitudinal cylinder with $105 \text{ m} < Z < 157 \text{ m}$ and radius $R_{\text{upstream}} = 536.5$ mm. Events with vertices reconstructed in this region are also kept for further analysis. Figure 4.23 shows the defined fiducial volumes.




Figure 4.23: Transverse position of all secondary vertices selected in the data sample as a function of the longitudinal position. LAV and ANTI0 veto conditions are not applied. The black lines are the edges of the fiducial volume. The upstream region is the red-hatched area. CR/SR are masked for e^+e^- vertices.

4.4.3 Background determination for $\mu\mu$

Prompt background for di-muon events An equivalent of 8.37×10^{15} POT has been generated for the $\mu\mu$ analysis using Stage-1b as described in section 4.3.1. The analysis of the MC truth for these events revealed that they do not originate from an actual $\mu^+\mu^-$ prompt production but rather are the effect of a secondary interaction of the incoming "beam" muon. The interaction most frequently happens in the LAV, creates another charged particle, and the spectrometer reconstructs both. A secondary meson can then decay into a muon.

A total of $141 \pm 66_{\text{stat}} \pm 70_{\text{syst}}$ events are expected from the prompt background before the LAV veto. The first error is purely statistical, and the second is the MC systematic error. The background distribution expected for the prompt background before the LAV veto in the CDA_{TAX} -vs- Z_{TAX} plane is shown in figure 4.24. No event is found in the CR.




Figure 4.24: Events selected as $\mu^+\mu^-$ from the prompt-background MC before the LAV veto request. Distribution of the distance of closest approach of the total momentum at the vertex to the beamline at the TAX impact point (CDA_{TAX}) vs the Z coordinate of the minimum approach. The grey rectangles represent the signal and control regions.

This MC prediction has been compared with the data. The distribution of the track-track time difference of $\mu^+\mu^-$ data events before the LAV veto is shown in figure 4.25 where the prompt component is clearly visible on top of the combinatorial background and amounts to 270 ± 27 events. The data/MC ratio, $R = 1.91 \pm 0.91_{\text{stat}} \pm 0.95_{\text{syst}}$, differs from one by 0.7 standard deviations. The obtained R is used as a scale factor for the MC prediction: after the LAV veto condition $0.28 \pm 0.13_{\text{stat}} \pm 0.14_{\text{syst}}$ prompt-background events are expected. Note that the reduction factor of the CR (SR) request in data is below

1.6% (0.16%) at a 90% confidence level, making the expected contribution of the prompt background completely negligible.




Figure 4.25: Events selected as $\mu^+\mu^-$ before the LAV veto request. Time difference between the two selected tracks.

Combinatorial background for di-muon events After weighting the events according to equation 4.7, the distributions of CDA_{TAX} vs Z_{TAX} for events with $\mu\mu$ PID assignments are shown in figure 4.26 superimposed with data events as full dots. Three additional control regions, labeled as CR3, CR2, and CR1, are defined for $\mu^+\mu^-$ progressively closer to the actual CR. A qualitative agreement between the data and predicted combinatorial background can be observed.




Figure 4.26: Distributions of CDA_{TAX} vs Z_{TAX} for (left) $\mu^+\mu^+$, (right) $\mu^-\mu^-$ and (bottom) $\mu^+\mu^-$ events selected from the combinatorial sample. Data events are superimposed as full black dots. Possible data within the control or signal regions (CR, SR) of $\mu^+\mu^-$ distribution is masked.

A quantitative comparison between the observed and expected number of events is shown in table 4.4. A good agreement is observed in all the cases.

Final state	region	$N_{\text{exp}} \pm \delta N_{\text{exp}}$	N_{obs}	$p(L \leq L_{\text{obs}})$
$\mu^+ \mu^+$	Outside CR	62.5 ± 9.4	53	0.46
	CR	0.46 ± 0.07	0	1.0
	SR	0.040 ± 0.006	0	1.0
$\mu^- \mu^-$	Outside CR	9.1 ± 1.4	8	0.88
	CR	0.050 ± 0.007	0	1.0
	SR	0.005 ± 0.0007	0	1.0
$\mu^+ \mu^-$	Outside CR	30.9 ± 4.6	28	0.78
	CR ₃	2.00 ± 0.30	2	1.0
	CR ₂	0.68 ± 0.10	1	0.48
	CR ₁	0.34 ± 0.05	1	0.29
	CR ₁₊₂₊₃	3.02 ± 0.45	4	0.56
	CR	0.20 ± 0.04	–	–
	SR	0.019 ± 0.004	–	–

Table 4.4: Summary of expected di-muon events from combinatorial background (N_{exp}), the related systematic error (δN_{exp}), the observed events in data (N_{obs}) and the p -values representing the probability to obtain a likelihood L for data-MC compatibility equal or worse than that corresponding to N_{obs} ($p(L \leq L_{\text{obs}})$).

For the $\mu^+ \mu^-$ final state, an evaluation of the combinatorial background alternative to the use of equation (4.8) can be obtained by determining D_{eff} from the data/MC ratio from same-sign events. The central value of D_{eff} is 7.7% lower than that of equation (4.8) and its fractional uncertainty is 6.4%, dominated by the data statistics. The results of this procedure are shown in table 4.5, considered the final estimate of the combinatorial background.

Region in CDA _{TAX} vs Z _{TAX}	$N_{\text{exp}} \pm \delta N_{\text{exp}}$	N_{obs}	$p(L \leq L_{\text{obs}})$
Outside CR	26.3 ± 3.4	28	0.74
CR ₃	1.70 ± 0.22	2	0.68
CR ₂	0.58 ± 0.07	1	0.44
CR ₁	0.29 ± 0.04	1	0.25
CR ₁₊₂₊₃	2.57 ± 0.33	4	0.34
CR	0.17 ± 0.02	–	–
SR	0.016 ± 0.002	–	–

Table 4.5: Summary of expected $\mu^+ \mu^-$ events from combinatorial background (N_{exp}), the related uncertainty (δN_{exp}), the observed events in data (N_{obs}) and the p -values representing the probability to obtain a likelihood L for data-MC compatibility equal or worse than that corresponding to N_{obs} ($p(L \leq L_{\text{obs}})$). The data/MC ratio for same-sign events determines the effective downscale factor.

Expected background in the control and signal region

The expected background is shown as a function of the di-muon invariant mass in the left panel of figure 4.27. The total expected number of background events is 0.016 ± 0.002 in the range of invariant masses (215, 1200) MeV and 0.009 ± 0.001 in the range (215, 700) MeV, where the errors are due to the uncertainty in the scaling of the combinatorial background. Consequently, with a 90% CL coverage and assuming no signal, no observed events are expected in the data SR in the full accessible range. The di-muon background invariant mass distribution has been fitted with a fourth-order polynomial to account for the asymmetric concave shape of the distribution. The PDF obtained is integrated in mass-sliding windows with a width of ± 2 and ± 3 DP-signal standard deviations, and the expected number of background events in each window is shown in the right panel of figure 4.27.

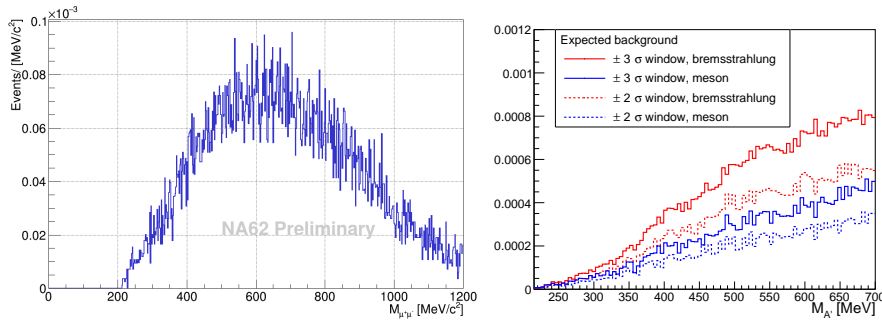


Figure 4.27: Expected number of background events for the search of $A' \rightarrow \mu^+ \mu^-$ as a function of the di-muon invariant mass in 1 MeV bins (left) and in ± 2 (dashed line) and ± 3 (solid line) signal standard deviations, for bremsstrahlung (red) and meson-mediated (blue) production modes (right).

To determine the discovery potential, the ± 3 sigma windows related to the direct production from above are used. The number of observed events necessary in data to infer a five-sigma level discovery has been evaluated: if the Feldman-Cousins unified approach is used, with a confidence level of 0.9999994, the lower limit on the number of signal events is away from zero by more than 10^{-5} . The number of necessary observed events is two throughout the entire mass range. This evaluation suffers from the so-called look-elsewhere effect since it is based solely on the event counts inside a given mass window, and it does not account for the compatibility of the entire distribution of observed events with the signal-plus-background or background-only

hypotheses. The evaluation of p -values for the observed events to derive from a signal in the presence of background can be performed using an unbinned, extended maximum-likelihood fit [207]. It turns out that the probability of finding two background events within $\pm 3\sigma$ is approximately 5×10^{-6} . Therefore, the minimum number of events to be observed in a three-sigma window to claim a discovery is three.

4.4.4 Background determination for ee

As suggested by the distribution of the track-track time difference, vertices with at least one e^\pm in the final state are dominated by in-time tracks. The combinatorial background due to random pairing of e^+ and e^- tracks, evaluated the same way as in the case of $\mu\mu$, is less than 10^{-3} events, hence completely negligible with respect to the prompt background as will be shown below.

The size of the prompt background MC sample generated for the ee analysis using the backward MC is equivalent to 1.53×10^{17} POT. The numerical evaluation of the agreement between data and MC before applying the ANTI0 and LAV veto is shown in table 4.6. The number of events observed in data is fully compatible with the MC expectation in all regions, a remarkable result considering statistics of more than 10^{17} POT.

Region	$N_{\text{exp}} \pm \delta N_{\text{exp}}$	N_{obs}	$p(L \leq L_{\text{obs}})$
Fiducial volume	18.8 ± 10.2	25	0.59
Upstream side-band	40.1 ± 20.8	56	0.50
Total	58.9 ± 30.2	81	0.51

Table 4.6: Summary of expected e^+e^- events from the MC sample (N_{exp}), the related total uncertainty (δN_{exp}), the observed events in data (N_{obs}) and the p -values representing the probability to obtain a likelihood L for data-MC compatibility equal or worse than that corresponding to N_{obs} ($p(L \leq L_{\text{obs}})$).

The cross-check for the data-MC agreement for the three outer control regions in the $(Z_{\text{TAX}}, CDA_{\text{TAX}})$ plane labeled as OCR_{3,2,1} in figure 4.28 is summarized in table 4.7. A good agreement is observed in all cases.




Figure 4.28: Distance of closest approach of the vertex momentum to the beamline at the TAX impact point (CDA_{TAX}) for the Z coordinate of the point of minimum approach in MC (left) and data (right). The distributions are separated by the position of the secondary vertex. The LAV and ANTI0 vetoes are not applied. The solid black contours are the CR and SR, as described in section 4.4.2. CR and SR are masked in the data sample. Control regions additional to CR are also shown by dashed gray contours.

Condition	$N_{\text{exp}} \pm \delta N_{\text{exp}}$	N_{obs}	$p(L \leq L_{\text{obs}})$
ee PID	58.9 ± 30.2	81	0.50
ee PID, OCR ₃	0.7 ± 0.7	4	0.40
ee PID, OCR ₂	2.5 ± 1.8	2	0.65*
ee PID, OCR ₁	2.9 ± 2.0	2	0.65*
ee PID, CR	0.5 ± 0.5	0	0.99*

Table 4.7: Summary of the expected number of e^+e^- vertices from the MC sample (N_{exp}) before the requirement of the LAV-ANTI0 veto conditions, the related total uncertainty (δN_{exp}), the observed events in data (N_{obs}) and the p-values representing the probability of obtaining a likelihood L for data-MC compatibility equal or worse than that corresponding to N_{obs} ($p(L \leq L_{\text{obs}})$). Values with * are obtained using Fisher's exact test for proportions.

The background evaluation in the SR and CR is divided into four steps: i) definition of the control sample before LAV and ANTI0 veto conditions are applied; ii) evaluation of the rejection power for LAV and ANTI0 veto conditions; iii) evaluation of the rejection power for the CR/SR request; iv) combination of the rejection powers to evaluate the expected background. Each rejection power indicated as η in the following is defined as the fraction of events rejected by the condition applied.

Rejection power of the LAV and ANTI0 veto The next step is to evaluate the rejection factor of the ANTI0 and LAV veto conditions combined. The data counts are compared to the expected number of events before and after

requiring the LAV-ANTI0 veto conditions, keeping CR and SR masked in both samples, in table 4.8.

Condition	$N_{\text{exp}} \pm \delta N_{\text{exp}}$	N_{obs}	$p(L \leq L_{\text{obs}})$
ee PID	58.9 ± 30.2	81	0.50
ee PID, ANTI0	0.7 ± 0.7	1	0.45
ee PID, LAV	11.9 ± 6.7	5	0.36
ee PID, LAV-ANTI0	0.7 ± 0.7	0	1.0

Table 4.8: Summary of the expected number of e^+e^- vertices from the MC sample (N_{exp}) before and after requiring the ANTI0, LAV, and LAV-ANTI0 veto conditions, the related total uncertainty (δN_{exp}), the observed events in data (N_{obs}) and the p-values representing the probability of obtaining a likelihood L for data-MC compatibility equal or worse than that corresponding to N_{obs} ($p(L \leq L_{\text{obs}})$). Values with * are obtained using Fisher's exact test for proportions.

Data and MC observations are compatible. The rejection power of the LAV-ANTI0 veto condition is then evaluated from MC, using e^+e^- vertices in the fiducial volume and upstream, but this time not masking the CR/SR. The result is shown in the second row of the summary table 4.11.

Rejection power of the CR and SR requests In complete analogy, using the weights of events entering CR and SR in the MC sample and considering the fiducial volume and upstream side-band as a single sample, the evaluated CR and SR rejections without applying LAV and ANTI0 veto are shown table 4.11.

The rejection powers of the CR and SR requests are determined, assuming each to be independent of the LAV-ANTI0 veto condition. This ansatz is verified using a control sample obtained from $e\mu$ vertices, which are produced via the same processes as e^+e^- but are much more abundant. The data counts and expected number of events from the MC for the $e\mu$ final states before and after the LAV-ANTI0 veto conditions and within the SR and CR are shown in table 4.9. The event counts agree between data and MC within the uncertainties.

Condition	$N_{\text{exp}} \pm \delta N_{\text{exp}}$	N_{obs}	$p(L \leq L_{\text{obs}})$
$e\mu$ PID	2905 ± 1455	2896	0.97
$e\mu$ PID, ANTI0	8.6 ± 6.1	12	0.61
$e\mu$ PID, LAV	728 ± 365	645	0.94
$e\mu$ PID, LAV-ANTI0	0	2	0.25*
$e\mu$ PID, CR	50 ± 26	49	0.98
$e\mu$ PID, SR	2.5 ± 1.8	3	0.83
$e\mu$ PID, LAV-ANTI0, CR	0	0	–
$e\mu$ PID, LAV-ANTI0, SR	0	0	–

Table 4.9: Summary of the expected number of $e\mu$ vertices from the MC sample (N_{exp}) before and after requiring the ANTI0, LAV, and LAV-ANTI0 veto conditions, the related total uncertainty (δN_{exp}), the observed events in data (N_{obs}) and the p-values representing the probability to obtain a likelihood L for data-MC compatibility equal or worse than that corresponding to N_{obs} ($p(L \leq L_{\text{obs}})$). The entry with * is obtained from a Fisher’s test on proportions assuming integer counts.

For the combined LAV-ANTI0 condition, only two events remain in data and no events in MC, preventing the ansatz validation. To validate the independence of the requests, an inverse sample is used by using events which are vetoed by ANTI0 and re-weighted by $w = \prod_i (1 - \varepsilon_i) / \varepsilon_i$, where ε_i are the measured tile efficiencies, shown in figure 4.20, and the index runs over tiles vetoing given event. The LAV veto is then applied, and, using the surviving events, the rejection powers of the CR and SR requests are evaluated. The CR and SR rejection factors in data with and without the ANTI0-LAV veto condition agree within the uncertainties, as shown in table 4.10, validating the factorization ansatz.

	Condition	$1 - \eta_{\text{CR}}$	$1 - \eta_{\text{SR}}$
MC	$e\mu$ PID	$0.0172^{+0.0043}_{-0.0020}$	$0.0009^{+0.0010}_{-0.0004}$
Data	$e\mu$ PID	$0.0169^{+0.0048}_{-0.0022}$	$0.0010^{+0.0012}_{-0.0004}$
Data	$e\mu$ PID, LAV-ANTI0	$0.0107^{+0.0042}_{-0.0030}$	$0.00008^{+0.00127}_{-0.00007}$

Table 4.10: Rejection power of the CR/SR requests from a sample of vertices identified as $e\mu$ PID. The third row is obtained after mimicking the ANTI0 veto request using the tile-efficiency map.

Determination of the expected prompt background Using the derived rejection factors η for the LAV-ANTI0, SR, and CR conditions, shown in table 4.11, the expected number of events in CR and SR after the full selection

can be computed as

$$N_{\text{bkg}}^{\text{CR,SR}} = N_{\text{bkg}}^{\text{inFV}} \times (1 - \eta_{\text{LAV-ANTI0}}) \times (1 - \eta_{\text{CR,SR}}). \quad (4.16)$$

Condition	$N_{\text{exp}} \pm \delta N_{\text{exp}}$	$1 - \eta$
e^+e^- PID	59.9 ± 6.7	–
e^+e^- PID, LAV-ANTI0	0.72 ± 0.72	$0.012^{+0.020}_{-0.008}$
e^+e^- PID, CR	0.51 ± 0.51	$0.009^{+0.018}_{-0.006}$
e^+e^- PID, SR	0.47 ± 0.47	$0.008^{+0.018}_{-0.006}$

Table 4.11: Summary of the expected number of e^+e^- vertices from the MC sample (N_{exp}) surviving various stages of the selection, the related statistical uncertainty (δN_{exp}), the rejection power and the related statistical uncertainty ($\eta \pm \delta\eta$) of the LAV-ANTI0 veto conditions and of the CR/SR requirement. The $1 - \eta$ factors are relative to the PID request.

Two independent methods are used to combine the uncertainties in equation (4.16). In both, $N_{\text{bkg}}^{\text{inFV}}$ is assumed to fluctuate with a gamma distribution. In a fully frequentist scenario, all the rejection powers η are assumed to be flatly distributed within the related 90% confidence interval. In a bayesian scenario, the distributions of the η factors are obtained when evaluating them using a beta distribution. For both scenarios, numerical simulations are performed to combine the PDF (10^6 experiments were simulated). The following 68% CL intervals are obtained:

$$\begin{aligned} 0.0021 &\leq N_{\text{bkg}}^{\text{CR}} \leq 0.0243 @ 68\% \text{CL (frequentist)} \\ 0.0012 &\leq N_{\text{bkg}}^{\text{CR}} \leq 0.0194 @ 68\% \text{CL (bayesian)} \\ 0.0021 &\leq N_{\text{bkg}}^{\text{SR}} \leq 0.0237 @ 68\% \text{CL (frequentist)} \\ 0.0012 &\leq N_{\text{bkg}}^{\text{SR}} \leq 0.0188 @ 68\% \text{CL (bayesian)}. \end{aligned} \quad (4.17)$$

An upgraded bayesian evaluation can be obtained starting from the number of events in the FV, this time masking the CR and SR, using

$$N_{\text{bkg}}^{\text{CR(SR)}} = \frac{N_{\text{bkg}}^{\text{inFV}} \Big|_{\text{CR,SR masked}}}{\eta_{\text{CR}} + \eta_{\text{SR}} - 1} \times (1 - \eta_{\text{LAV-ANTI0}}) \times (1 - \eta_{\text{CR(SR)}}), \quad (4.18)$$

where the denominator is needed to obtain the corresponding $N_{\text{bkg}}^{\text{inFV}}$. While $N_{\text{bkg}}^{\text{inFV}} \Big|_{\text{CR,SR masked}} = 18.8 \pm 3.9 \pm 9.4$ in the MC as well, the number of events observed in data, 25, allows reducing the expected spread of the true value.

The obtained distribution of $N_{\text{bkg}}^{\text{SR}}$ for 10^6 simulated outcomes in this bayesian scenario is shown in figure 4.29. The results of this method at 90% and 68% CL are¹¹

$$\begin{aligned}
 0.0008 &\leq N_{\text{bkg}}^{\text{CR}} \leq 0.059 @ 90\% \text{CL} \\
 0.0007 &\leq N_{\text{bkg}}^{\text{SR}} \leq 0.058 @ 90\% \text{CL} \\
 0.0024 &\leq N_{\text{bkg}}^{\text{CR}} \leq 0.031 @ 68\% \text{CL} \\
 0.0022 &\leq N_{\text{bkg}}^{\text{SR}} \leq 0.030 @ 68\% \text{CL},
 \end{aligned} \tag{4.19}$$

which are also expressed using the intervals around the median:

$$\begin{aligned}
 N_{\text{bkg}}^{\text{SR}} &= 0.0094_{-0.009}^{+0.049} @ 90\% \text{CL} \\
 N_{\text{bkg}}^{\text{SR}} &= 0.0094_{-0.007}^{+0.021} @ 68\% \text{CL}.
 \end{aligned} \tag{4.20}$$




Figure 4.29: Distribution of the expected number of background events in CR, SR for 1M simulated experiments in the bayesian approach of eq. (4.18).

With a total expected number of background events below 0.06 with a 90% CL coverage and assuming no signal, no observed events are expected in the data CR/SR. A numerical simulation has been performed to explore the discovery potential implied by the background estimate. In a bayesian approach, the number of background events in the SR has been sampled from a Poisson distribution with a true value distributed as in figure 4.29. The background events are then distributed in invariant mass. The expected background distribution has been taken from the e^+e^- events before the LAV-ANTI0 veto

¹¹It is expected that the result in (4.19) is larger than (4.17) as the data observation exceeds the central value of the MC expectation by more than 30%.

conditions are applied, modeled with a Landau fit (see the left panel of figure 4.30). The probability of observing one, two, three, and four background events within ± 3 signal standard deviations centered around an assumed A' mass is shown in the right panel of figure 4.30. In conclusion, observing two events within $\pm 3\sigma_M$ is enough for a discovery for $M_{A'} > 50$ MeV. For masses below 50 MeV, three events are needed.




Figure 4.30: Left: fit with a Landau PDF of the invariant mass distribution of the background to the e^+e^- search. Right: expected probability of observing one, two, three, and four events within $\pm 3\sigma_M$ around a given A' mass. The black horizontal line indicates the probability corresponding to a five-sigma level (discovery).

4.4.5 Efficiency determination

The selection efficiency is determined by Monte Carlo as a function of the unknown dark photon mass and coupling, separately for the bremsstrahlung-like and the meson-mediated production processes. The results are shown in figure 4.31. The statistics generated by MC are sufficient to guarantee a relative statistical error of a few percent in the relevant parameter space (except for the region at the very border of the sensitivity).

The selection algorithm has been split into five steps: track and vertex reconstruction, association to downstream detectors, extrapolation within the signal region, trigger, LAV, and ANTI0 veto conditions, and the PID. The systematic errors are determined separately for each component.




Figure 4.31: Selection efficiency for the $A' \rightarrow \mu^+ \mu^-$ (top) and $A' \rightarrow e^+ e^-$ (bottom) signal, as a function of the DP mass and coupling. Left (right) panel refers to the bremsstrahlung-like (meson-mediated) production mode.

Tracking

The quality of the track reconstruction has been checked using the `SpectrometerDataQualityMonitor` analyzer. The results show that the track residuals are below $50 \mu\text{m}$ throughout the entire detector, and the distributions of first leading time, last trailing time, and the sum of the radii of two consecutive planes are compliant with the expectation, as cross-checked with the relevant experts. These results are not entirely trivial since the distribution of momentum, position, and angles of the incoming charged particles greatly differ from that of a standard kaon run. It must be underlined that the tracking efficiency depends on the instantaneous intensity, as thoroughly discussed in [208]. The beam-dump data corresponds to a total single-track rate well below 1 MHz. Moreover, the distribution of the transverse position of tracks at the first STRAW chamber is smooth, so the pile-up effects are expected to be irrelevant. In these conditions, the absolute inefficiency is seen to be at the level of a few per mil. The systematic error related to the tracking efficiency is considered negligible. Systematic errors on the quality of the track reconstruction are discussed in Sec. 4.4.5.

Association to downstream detectors

For each di-lepton mode, association with NewCHOD candidates for both tracks is mandatory to ensure a well-controlled trigger efficiency for the H2 condition. Therefore, an accurate evaluation is performed of the probability for a track in the acceptance of the NewCHOD sensitive region to be associated with a NewCHOD tile in time. The NewCHOD efficiency is evaluated at the last step of the centralized data-quality process for each data run and for the MC. The data inefficiency is always below the percent level i) throughout all the interesting momentum range (track charges are separately treated); ii) throughout the entire burst duration; iii) along the bursts of each run; iv) throughout most of the NewCHOD tiles. The MC inefficiency is lower than the data one, and the data-MC difference averages to 0.6%. The association efficiency has been corrected for such discrepancy, leading to a -1.2% relative correction. A systematic error of $\pm 0.6\%$, equal to the single-track correction, is conservatively assumed.

Similar comparisons have been made for the CHOD association and (for muons) for the MUV3 association, where the inefficiency of the related data-quality analyzers is at or below the few per-mil level.¹² No correction or systematic error is derived in this case.

STRAW resolution effects

A study related to the data-MC agreement on the resolution of kinematic variables [209], done within the context of the $K^+ \rightarrow \pi^+ \mu^+ \mu^- (\gamma)$ analysis, estimates a maximum data-MC difference of about 5% on the $K \rightarrow 3\pi$ mass resolution. This would be a reassuring figure of merit if one also considers it to be applied to the invariant mass of the DP signal. As said, in 2021, data set track residuals smaller than $50 \mu\text{m}$ are observed throughout the entire STRAW detector. While not signaling any intrinsic calibration problem, it should be stressed that this is the effect of the missing run-by-run fine calibrations, which have not been applied to the beam-dump data. An effective treatment of MC signal events is by using the `SpectrometerRecoAnalyzer` pre-analyzer with settings corresponding to a 1 ns smearing of the STRAW times and comparing the analysis results with the standard treatment of STRAW data. Once the smearing is applied, the standard deviation of the Z_{TAX} distribution increases by 2% for the signal bremsstrahlung production and by 3% for the

¹²The `MUV3Efficiency` analyzer considers a track associated to the MUV3 if one or more MUV3 tiles are geometrically associated within ± 15 ns from the trigger time.

meson-mediated production. These variations have negligible effect in terms of the fraction of signal events possibly passing from the signal region to the control region. Once the smearing is applied, the number of two-track vertices selected decreases by 1.5%, and this is considered a figure of merit of the systematic error related to the reconstruction and resolution on spectrometer tracks and vertices.

Trigger

The trigger efficiency has been continuously measured along the data set. The efficiency for the H2 primitive has been measured starting from control trigger events, with a two-track vertex reconstructed and two NewCHOD tiles associated. The inefficiency averages to 2%, with a 0.5% overall run-by-run variation. A 2% correction is therefore applied to the MC efficiency, while a systematic error of 0.5% is considered.

Particle ID efficiency for $\mu\mu$

The efficiency of the particle identification (PID) conditions described in section 4.4.2 has been determined on a single-particle MC for the whole beam-dump momentum range (up to 400 GeV/ c) and verified on kaon data and MC for the low momentum range (up to 60 GeV/ c).

The single-particle MC sample serves only for the study of the detector response for high momentum tracks of different particle species and does not aim to reproduce the observed momentum spectra. A data-driven MC production is used for generating the sample, using the reconstructed tracks from the data and propagating them backward to the TRIM5 plane. These tracks are then used as a gun for the MC with a fixed particle type (μ^\pm).

The $K^+ \rightarrow \mu^+ \nu_\mu$ decay ($K_{\mu 2}$) is used for the validation of the PID efficiency for μ . The $K_{\mu 2}$ sample is extracted from the data using a kinematic selection with an additional veto on in-time activity in LAV, IRC, and SAC and requiring at most one reconstructed candidate in MUV1 and 2 and exactly one reconstructed track with an associated candidate in NewCHOD. Appendix C provides more details on the kaon data analyzers. Data from run 11001 selected with the trigger mask 2 is used (primitives based on RICH & NewCHOD, downscaled by a factor of 600).




Figure 4.32: Inefficiency for μ PID identification as a function of track momenta based on a single-muon beam-dump MC, $K_{\mu 2}$ MC (2021 conditions) and $K_{\mu 2}$ 2021 data (top left) and the fraction of events in 2021 data selected as not muons by the PID used in this analysis (top right). Difference of the inefficiency estimated for $K_{\mu 2}$ data and $K_{\mu 2}$ MC (bottom left) and single-muon MC vs 2021 data difference (bottom right).

The observed PID inefficiencies in kaon data and MC¹³ and in the single-muon beam-dump MC are compared in the left panel of figure 4.32 where the total PID inefficiency for muons in the beam-dump MC is calculated to be $(3.546 \pm 0.005)\%$ while the value calculated on the $K_{\mu 2}$ data yields $(7.4 \pm 0.1)\%$. The observed data-MC discrepancy might be due to an improper simulation of the muon multiple scattering or other effects producing additional signals in MUV3. The overall fitted value of the data-MC difference yields $\varepsilon_{\text{MC}, K_{\mu 2}} - \varepsilon_{\text{data}, K_{\mu 2}} = (3.8 \pm 0.1)\%$. The data-MC efficiency difference shows a momentum dependence with a decreasing trend towards higher momenta. In order to confirm this trend for the whole range, the inefficiency for the single-particle MC with the fraction of data single tracks not identified as muons is compared in the right panel of figure 4.32: the contamination of the data sample by particles other than muons is estimated to be below 1% using more precise PID tools, namely the boosted decision trees [210] (BDT) and convolutional neural network [211] (CNN) algorithms (see figure 4.33). The conclusion is that the observation from the $K_{\mu 2}$ sample (figure 4.32 left) is in agreement with the beam-dump one (figure 4.32 right). The total correction to the signal selection efficiency is estimated by weighting each single muon

¹³Overlaid and non-overlaid MC simulations yield comparable results.

track in the MC-signal events with the data-MC discrepancy observed for the beam-dump (figure 4.32 right). The impact on the overall efficiency is 8%, and a systematic error of $\pm 1\%$ has been evaluated from the uncertainty on the data-MC inefficiency difference.




Figure 4.33: Fraction of events selected as not muons in one beam-dump run (11084) by the BDT and CNN PID algorithms.

Particle ID efficiency for ee

Similarly to $\mu\mu$, the PID efficiency has been studied both in data and MC control samples. A single-electron “beam-dump” MC sample has been produced starting from data events with single tracks from muon-halo: the track momenta and positions are evaluated for $Z = 102$ m, and single-particle MC gun is built while enforcing electron/positron ID. A control sample of data and MC electron tracks is obtained from the selection of K_{e3} decays (2021 data are used). The K_{e3} analyzer is described in more detail in appendix C.

The left panel of figure 4.34 shows the efficiency of the electron-PID condition originally used in the $\mu\mu$ analysis, in which the cut $0.95 < E/p < 1.05$ is applied, compared to the new PID definition after applying cluster corrections in the ee analysis. In the single-electron MC, the sizeable inefficiency of the original PID observed for momenta above 100 GeV/c is primarily due to saturation effects, as discussed in section 4.1.2. The consequence of such inefficiency would be a decreased sensitivity for the $X \rightarrow e^+e^-$ decays in the region of high coupling values. After applying energy corrections to the clusters identified as saturated, and with the modified PID conditions described in section 4.4.2, a significant improvement in the PID efficiency is achieved, as shown in figure 4.34 (left). The evaluation of the systematic error on the improved PID conditions would require a data sample in which the saturation

effects can be investigated. This is not available using e^+e^- or $e\mu$ control samples, as the e^\pm momenta exceed 100 GeV/ c with insufficient statistics. However, two methods correcting for the saturation at the analysis level were compared. In the present analysis, the energy for saturation clusters is corrected using the expected fraction of energy deposited in cells other than the seed, and the saturated-cluster time is obtained from the cell adjacent to the seed with the highest energy. An alternative method consists in correcting the saturated-cluster energies by a fixed offset (70 GeV) and avoiding any modification of the time. Data counts and MC expected events for the e^+e^- final state are independent of which of the two methods is applied, suggesting a negligible systematic error from the saturation correction.




Figure 4.34: Inefficiency for e PID as a function of the track momentum. Left: comparison of e PID inefficiencies using the original conditions of $\mu\mu$ analysis (black) and the conditions upgraded for the ee analysis as explained in section 4.4.2 (blue). Right: results from a single-electron beam-dump MC (black), a K_{e3} MC in 2021 conditions (blue), and a K_{e3} 2021 data sample (red) are shown. The bottom left plot is the data/MC inefficiency ratio for K_{e3} events.

The right panel of figure 4.34 compares the performance of the new PID on single-electron MC with positrons selected in a K_{e3} sample from both 2021 data and MC. Data and MC agree at the percent level, although only a region of momenta below 55 GeV/ c is explored. The PID inefficiency from MC K_{e3} concurs with that from the beam-dump single-electron MC.

Summary of signal selection efficiency

A summary of uncertainties and corrections for the signal selection efficiency is given in table 4.12. The entry for the statistical error represents a nominal value since the error slightly varies with the exotic particle mass and coupling. The total relative uncertainty on the efficiency is below 3%. The MC efficiency needs to be corrected by -13.7% for $\mu\mu$ analysis and by -6.3% for ee analysis to account for the various data/MC disagreements discussed.

Source	Uncertainty	Relative correction
Statistics	2.1%	—
NewCHOD association	0.6%	-1.2%
Straw quality and resolution	1.5%	—
Trigger	0.5%	-2.0%
LAV random veto	0.1%	-2.5%
ANTI0 random veto ($A' \rightarrow e^+e^-$)	0.1%	-0.6%
PID ($A' \rightarrow \mu^+\mu^-$)	1.0%	-8.0%
PID ($A' \rightarrow e^+e^-$)	1.1%	—
Total ($A' \rightarrow \mu^+\mu^-$)	2.8%	-13.7%
Total ($A' \rightarrow e^+e^-$)	2.9%	-6.3%

Table 4.12: Summary of uncertainties and corrections to the signal selection efficiency (alternative efficiency).

4.4.6 Determination of expected sensitivity

After accounting for the selection efficiency, the expected signal yield is shown as a function of the unknown DP coupling and mass in the top panel of figure 4.35 for the $\mu\mu$ analysis and in the bottom panel for the ee analysis. Contributions from direct and meson-mediated productions are separately shown. The first dominates over the second for most of the parameter space. Direct and meson productions are characterized by different resolution of the invariant-mass peaks, with the former larger than the latter for most of the mass range for a given coupling value (top panel of figure 4.36 for $\mu\mu$ and bottom panel for ee). Both the expected signal yield and the expected standard deviations of the signal invariant-mass peaks have been parameterized as two-dimensional analytic functions of the DP coupling and mass.




Figure 4.35: Expected number of events for the DP decay to $\mu^+\mu^-$ (top) and to e^+e^- (bottom) in the ϵ vs DP mass parameter space. The left (right) panel refers to direct (meson-mediated) production. A 2.3-event contour has been superimposed in red to guide the eye.




Figure 4.36: Standard deviation (MeV) of the invariant-mass peak for the DP decay to $\mu^+\mu^-$ (top) and to e^+e^- (bottom) in the ϵ vs DP mass plane. The left (right) panel refers to direct (meson-mediated) production.

Given the expected background, in the absence of a signal, no data events are expected to be observed in the SR at 90% CL. The uncertainty on the contours in figure 4.35 is dominated by the 20% relative error on the POT.

4.4.7 Control and signal region opening

$A' \rightarrow \mu^+ \mu^-$ analysis

In the $\mu\mu$ analysis, 0.17 ± 0.02 background events are expected in the CR, so the probability for a non-zero observation is approximately 15%. The control region was opened, and 0 events were observed in it.

In the entire di-muon invariant mass range (215–1200 MeV), 0.016 ± 0.002 events are expected in the SR, with the corresponding probability of non-zero observation of approximately 1.6%. After the SR has been opened, one event has been found with a two-track invariant mass of approximately 411 MeV. The corresponding observed 90% CL upper limit is represented by the region enclosed within the black contour in the left panel of figure 4.37. In the same picture, the color-filled area represents the expected contour uncertainty in the absence of an A' signal with a one-sigma (green) and two-sigma (yellow) statistical coverage.

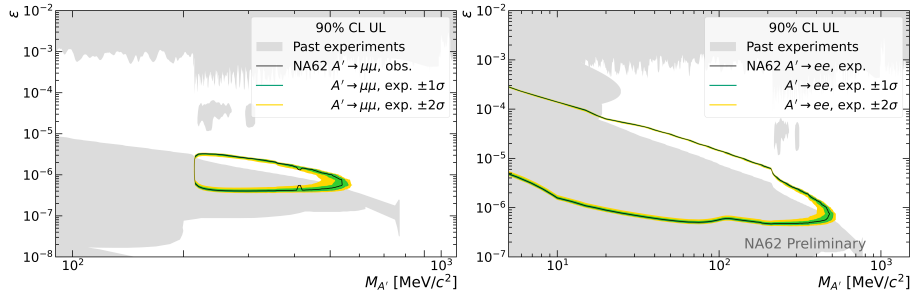


Figure 4.37: Search for $A' \rightarrow \mu^+ \mu^-$ (left) and $A' \rightarrow e^+ e^-$ (right): the region of the parameter space within the solid line is excluded at 90% CL. The color-filled area represents the expected contour uncertainty in the absence of a signal: green (yellow) corresponds to one (two) sigma statistical coverage.

The event found corresponds to a $2.4\text{-}\sigma$ global significance. Only two STRAW tracks are present in it: the positive has a momentum of 99.5 GeV, and the negative has a momentum of 39.5 GeV. The vertex would be found both with the `VertexLSF` algorithm and using the point of closest approach of the two tracks. The geometrical track-track distance at minimum is 2.71 mm.

The vertex position is well within the fiducial volume: the longitudinal coordinate is 158 m, and the transverse radial distance is 382 mm. Only two CHOD/NewCHOD candidates are present, and both are time-associated. Only two MUV3 tiles are fired in time: the tile IDs are 31 and 39, corresponding to non-adjacent tiles. The event might be due to the combinatorial background since the track time difference is -1.69 ns, away from zero by approximately two times the standard deviation expected from in-time events. Moreover, the extrapolation to the TAX impact point is barely within the SR, in a corner populated by far resolution tails of the expected signal, see figure 4.38.




Figure 4.38: Distance of closest approach between the beam direction at the TAX entrance and the total momentum of the two tracks (CDA_{TAX}) vs the longitudinal position of the minimum approach (Z_{TAX}). Left: data (dots) vs expected background (color density histogram). Right: data (dots) vs expected fraction of signal Monte Carlo events (color density). Bins of $2 \text{ mm} \times 1 \text{ m}$ size are used for the color density plots.

$A' \rightarrow e^+e^-$ analysis

The e^+e^- control region was opened, and 0 events were observed. Given the distribution of the expected number of background events in figure 4.29, the probability for a non-zero observation was approximately 1.7%.

After opening the signal region, 0 events were also observed, with the probability for a non-zero observation of approximately 1.6%. Note that there were no events in the CR and SR even before the LAV-ANTI0 veto conditions. Similarly, there were no events in the SR from the upstream data side-band.

The corresponding expected 90% CL upper limit is represented by the region enclosed within the black contour in the right panel of figure 4.37. In the same picture, the color-filled area represents the expected contour uncertainty in the absence of an A' signal with a one-sigma (green) and two-sigma (yellow) statistical coverage.

Model-independent interpretation for ALPs and dark scalars

The result can be interpreted as the search for an axion-like particle or a dark scalar emitted in a B meson decay. Model-independent interpretation is used here, allowing three parameters free to vary: the exotic particle mass, the rest lifetime, and the branching ratio product of the exotic particle production and decay [94]. The results for the selected value of the exotic particle mass are shown in figure 4.39 for the $\mu\mu$ analysis. For exotic particle masses below 300 MeV, the result from NA62 allows covering a region previously unexplored.




Figure 4.39: Search for an exotic particle a produced from decay of B mesons and decaying to $\mu^+ \mu^-$. Four values of the exotic particle mass are considered. The region of the parameter space within the green area is excluded at 90% CL. For comparison, the solid line represents the expected exclusion power of NA62 from a toy-MC simulation used in [94] assuming 1.4×10^{17} POT. The excluded regions from LHCb and CHARM measurements (see [94] for details) are superimposed as grey-filled areas. The shown updated CHARM bound has been obtained using the toy MC of the ALPINIST framework [1].

4.5 Search for hadronic decays of exotic particles

This analysis is an extension of the $A' \rightarrow \ell^+ \ell^-$ analyses and concludes the search for new physics, which can be achieved with a two-charged track analysis with the statistics collected by NA62 in the beam-dump mode in 2021. The analysis also paves the way for analyzing the new dataset collected in 2023, equivalent to about 2.5×10^{17} POT, which can be used to probe an even larger class of new physics models, such as heavy neutral leptons [212]. Given the larger variety of final states in the hadronic channels compared to the leptonic ones, the analysis result can be interpreted as a search for multiple

possible new physics particles. For this reason, the analysis is performed in a model-independent way, avoiding reference to any specific coupling setup and parametrizing the observation using only kinematic parameters. The result can then be interpreted as a search for a specific new physics model by re-weighting the output using the kinematic parameters according to the selected coupling setup by using the ALPINIST tool. The result of this analysis can be interpreted in at least three of the PBC benchmark models (*BC1*, *BC4* and *BC11*). An effective coupling to quarks arises in all of these models, resulting in coupling to hadrons in the low energy limit, but other coupling setups can also be probed.

The exotic particle masses kinematically accessible at NA62 are in the range of MeV to several GeVs, which allow exotic particle decays to light pseudoscalar mesons such as π , η , K and η' . The lowest-point amplitudes are naturally favored unless suppressed or forbidden by some *well-established* SM symmetry. Therefore, two-body hadronic decays are allowed for the scalar and vector exotic particles, while for the pseudoscalar ALP, three-body decays are favored. Two-body decays to vector mesons such as ρ , ω , and ϕ can be an interesting probe above the kinematic threshold as the branching ratios (BR) can be large for these channels. However, the signal yield is very low already for the lightest vector meson final states $\rho\rho$ even though the $\text{BR}(\rho \rightarrow \pi\pi) = 1$, due to the lower acceptance. For this reason, the decays to vector mesons are not simulated for this analysis, while the case of $\rho\rho$ is partially covered by the 4π decay, which is simulated for the $2\pi 2\pi^0$ final states. A summary of the final states of exotic particle decays probed by this analysis is shown in table 4.13.

DP	DS	ALP
$\pi^+\pi^-$	$\pi^+\pi^-$	$\pi^+\pi^-\gamma$
$\pi^+\pi^-\pi^0$		$\pi^+\pi^-\pi^0$
$\rho\rho(\pi^+\pi^-\pi^0\pi^0)$	$\rho\rho(\pi^+\pi^-\pi^0\pi^0)$	$\rho\rho(\pi^+\pi^-\pi^0\pi^0)$
		$\pi^+\pi^-\eta$
K^+K^-	K^+K^-	
$K^+K^-\pi^0$		$K^+K^-\pi^0$

Table 4.13: Summary of exotic particles and resulting final states probed by this analysis.

A simulation of NA62 sensitivity for the three benchmark models given the table 4.13 final states and assuming 1.4×10^{17} protons on TAX obtained with a toy Monte Carlo module of the ALPINIST framework [1, 2, 6] is shown in

figure 4.40. For the dark photon search, the addition of hadronic channels can extend the reach of the original leptonic analyses from section 4.4, especially around the mass of vector resonances. By including the hadronic decays, this analysis allows probing new regions beyond the CHARM and LHCb exclusion also in the minimal dark scalar scenario. For the ALP, this analysis is the first search in the $\mathcal{O}(100 \text{ MeV}) - \mathcal{O}(1 \text{ GeV})$ region since 1990 [128, 131].

As this analysis shares most of the selection with the $\ell^+\ell^-$ ones, the following sections will describe in detail only the modified selection criteria and the final state classification according to table 4.13 in section 4.5.1 and the expected backgrounds for these channels in section 4.5.2.

In order to allow the reinterpretation of the result in any new physics model, the selection efficiency and the NA62 sensitivity will be evaluated separately for each production and decay channel as a function of exotic particle mass m_X and decay width Γ_X in sections 4.5.3 and 4.5.4. The sensitivity for the three abovementioned benchmark models will be shown by combining the separate results.

4.5.1 Event selection

The signal searched by this analysis is characteristic by the presence of two charged tracks (h^+h^-) belonging to a hadron and an antihadron and up to four additional photons. The two tracks allow the reconstruction of the decay vertex and the association of the additional photons with the event. The selection conditions for the two tracks concern the track acceptance and association to the downstream detectors where the cuts applied are the same as were used for $A' \rightarrow e^+e^-$. The same holds for the vertex, which is reconstructed using the `VertexLSF` tool. The vertex time, given by the average of the corresponding track times, defines the reference time of the event. The vertex's location has to be within the modified FV as defined in section 4.4.2 as a union of two cones avoiding the LAV stations. Events reconstructed in the *upstream region* of the modified FV are kept for background study.

Particle Identification

The particle identification (PID) used in $\ell^+\ell^-$ was based on the E/p information from the spectrometer and LKr and the association with MUV3. The signal acceptance proves to be reasonable ($\gtrsim 90\%$) in these two borderline cases with $E/p \sim 1$ for e and $E/p \sim 0$ for μ , as shown in section 4.4.5.




Figure 4.40: NA62 sensitivity estimated using [1] for the *BCI* benchmark with meson decay and bremsstrahlung production (top left), when adding also mixing production (top right), *BC4* benchmark (bottom left) and *BC11* (bottom right). The dotted contour shows the sensitivity of this analysis as a standalone search. Full contour shows the combined sensitivity of all decay modes simulated. All contours correspond to the 90% confidence level (CL) exclusion.

While for hadrons, the efficiency can be as low as 50%. In order to increase the signal acceptance for hadronic decays, this analysis uses the BDT-based `SpectrometerCalorimetersAssociation` tool [210] which is used for the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ analysis and which is optimized for maximizing the pion signal acceptance and minimizing the contamination by other particles. The conditions used for selecting $e/h/\mu$ are the following:

Positron/electron (e^\pm): BDT electron hypothesis probability $p_e > 0.8$, and there must be no signal in MUV3 geometrically associated with the track and within 5 ns of the track time.

Hadron (h^\pm): BDT pion hypothesis probability $p_\pi > 0.8$, and there must be no signal in MUV3 geometrically associated with the track and within 5 ns of the track time.

Muon (μ^\pm): BDT muon hypothesis probability $p_\mu > 0.8$, and there must be exactly one signal geometrically associated with the track in any MUV3 tile and within 5 ns of the track time.

Exotic signal events typically have tracks of high momenta. This allows using RICH for the positively charged track to further distinguish π^+ and K^+ . As hadron-antihadron pairs are searched, the BDT-selected h^+h^- events with RICH most likely kaon hypothesis are labeled as K^+K^- while the rest is labeled as $\pi^+\pi^-$. Note that RICH information is not used here for vetoing any events in order not to decrease signal acceptance but only to provide additional information for the event classification. As will be shown in section 4.5.4, RICH can perform very well for distinguishing K^+K^- and $\pi^+\pi^-$.

Event classification and signal region definition

The information from possible additional neutral final states is needed in order to reconstruct the exotic particle correctly. The final state is assumed to be either a single photon or an unstable particle decaying instantly to a photon pair. Except for the track-associated LKr clusters,¹⁴ all clusters with $E_{\text{cluster}} > 5 \text{ GeV}$ are assumed to be belonging to photons with time $t_\gamma = t_{\text{cluster}}$ originating from the vertex. Their momentum \vec{p}_γ is then defined by E_{cluster} and direction given by the difference of the vertex and the LKr cluster position.

¹⁴The association from the $A' \rightarrow e^+e^-$ analysis is used for excluding the track clusters in this analysis.

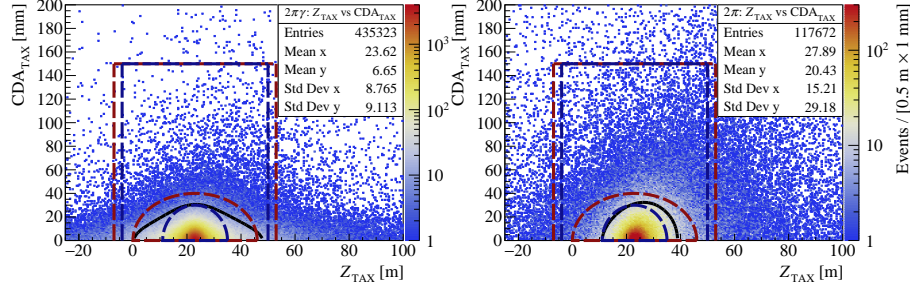


Figure 4.41: Distance of closest approach between the primary proton beam and the reconstructed exotic particle from a $\pi^+\pi^-\gamma$ event simulated with NA62 Monte Carlo. The red ellipse and box define the signal and control regions used in this analysis, compared to the blue signal and control regions used in $A' \rightarrow e^+e^-$. Left: correctly reconstructed $\pi^+\pi^-\gamma$. Right: $\pi^+\pi^-\gamma$ event mis-reconstructed as $\pi^+\pi^-$ event. The black solid line is a constant-population contour containing $\sim 93\%$ of the $a \rightarrow \pi^+\pi^-\gamma$ signal events. The constant-population contour corresponds to a fraction of 5×10^{-3} of the $a \rightarrow \pi^+\pi^-\gamma$ distribution maximum.

In case that two photons with $|t_{\gamma_1} - t_{\gamma_2}| < 5 \text{ ns}$ are found, it is checked whether their 4-momenta give together π^0 or η meson mass $|m_{\gamma_1\gamma_2} - m_{\pi^0,\eta}| < 15 \text{ MeV}/c^2$. If π^0 or η is reconstructed, the time $t_{\pi^0,\eta} = (t_{\gamma_1} + t_{\gamma_2})/2$ is assigned to it. When four photons are found, by combining their momenta, it is checked if two π^0 are reconstructed by this procedure.

This algorithm sorts the events in 5 *final state topologies*:¹⁵

1. h^+h^- ;
2. $h^+h^-\gamma$;
3. $h^+h^-\pi^0$;
4. $h^+h^-\eta$;
5. $h^+h^-2\pi^0$.

Events not belonging to any topology above are rejected. Suppose events from $h^+h^-2\pi^0$ topology are identified by the PID as $\pi^+\pi^-2\pi^0$. In that case, an additional check is performed by comparing the reconstructed invariant mass of each two pions matching in time $|t_{\pi_i} - t_{\pi_j}| < 5 \text{ ns}$ with the ρ meson mass.

¹⁵The particle identification is applied independently on the tracks, h^\pm can therefore be also an electron or muon at this stage.

When $|m_{\pi_i\pi_j} - m_\rho| < 80 \text{ MeV}/c^2$ a ρ meson with $t_\rho = (t_{\pi_i} + t_{\pi_j})/2$ is reconstructed. Events with reconstructed $\rho^+\rho^-$ or $\rho^0\rho^0$ are labeled as 2ρ ; otherwise, they are assumed to be $\pi^+\pi^-2\pi^0$. This additional classification is taken as complementary information rather than a sixth *topology* as the signal acceptance for a 2ρ event is very low. For all practical purposes, 2ρ events are assumed to belong to the $h^+h^-2\pi^0$ topology.

The combination of event topology and the PID of individual tracks allows the reconstruction of an exotic particle decaying to all hadronic final states from table 4.13 from their momenta. The exotic particle is produced in the TAX area via one of the mechanisms described in section 2.1.

By projecting the reconstructed exotic particle to the point of smallest distance of closest approach CDA_{TAX} to the impact point (3.1) one can select the signal events by requiring the proximity of the interaction point same as was done for the $\ell^+\ell^-$ analyses. The expected signal distribution of $\pi^+\pi^-\gamma$ events produced in the interaction point, simulated with NA62 Monte Carlo, is shown in left panel of figure 4.41. The signal region (SR) is defined as an ellipse centered at $Z_{\text{TAX}} = 23070 \text{ mm}$ with $r_z = 23 \text{ m}$ and $r_{\text{CDA}} = 40 \text{ mm}$ while the control region (CR) is a box with $\text{CDA}_{\text{TAX}} < 150 \text{ mm}$ and $-7 \text{ m} < Z_{\text{TAX}} < 53 \text{ m}$. Both regions are shown by red dashed lines in figure 4.41. The regions are larger compared to the ones defined for the $A' \rightarrow e^+e^-$ analysis (blue dashed lines) to accommodate a larger spread of the signal indicated by black solid contour which contains 93% of the $\pi^+\pi^-\gamma$ signal events. Both regions are kept masked until the analysis is finalized.

Additional radiated photons in the event could affect the event reconstruction as it can *promote* an event of simpler topology to an event of a more complex one. Moreover, the reconstructed exotic particle momentum gets modified by the additional contribution, which can shift the projection to the interaction point from the SR to the CR. Radiated photons are typically soft, as shown in figure 4.42 for simulated $\pi^+\pi^-$ events, therefore in the majority of cases, this mis-reconstruction is prevented by the $E_{\text{cluster}} > 5 \text{ GeV}$ requirement.




Figure 4.42: Left: Energy of additional LKr clusters in $A' \rightarrow \pi^+\pi^-$ decay which can promote the event to an $h^+h^-\gamma$ topology. Right: Energy of LKr clusters not associated to π in $a \rightarrow \pi^+\pi^-\gamma$ decay. The energy limit on photon association is indicated by the red line.

However, an opposite scenario is possible if a signal photon is lost, but LAV or SAV does not veto the event, the event is *demoted* to a simpler topology. For example, for an $X \rightarrow \pi^+\pi^-\gamma$ event, this can occur in as many as 20% of cases. Right panel of figure 4.41 depicts the shift of units of percent of $X \rightarrow \pi^+\pi^-\gamma$ events from the SR of $X \rightarrow \pi^+\pi^-$ to its CR and $\sim 1\%$ even outside of it.

The mis-reconstruction of events of more complex topologies as simpler ones occur for all topologies. More complex topologies are also linked to a higher mass of the exotic particle and, therefore, to a lower signal yield, as will be shown in section 4.5.4, as these exotic particles are less boosted. This suggests a strategy for unmasking the CRs and SRs to open the regions individually from more complex topologies to simpler ones as a non-observation of signal in SR of a more complex topology ensures that the CR of a simpler topology is not polluted by the signal of a more complex one. It will be shown in section 4.5.2 that this strategy also goes in line with the SM background expectation, which is lower for more complex topologies.

LAV, ANTI0 and SAV vetoes

The LAV and ANTI0 vetoes are applied with the same cuts as before: The LAV veto using the LAVMatching tool with a 10 ns time window with respect to the reference time results in a random veto of $\epsilon_{RV}^{LAV} = (2.5 \pm 0.1)\%$ and the ANTI0 veto applied when any of the tracks has a hit in an associated ANTI0 tile within 5 ns with respect to the track time. The ANTI0 veto results in a random veto of $\epsilon_{RV}^{ANTI0} = 0.3\%$.

Beyond LAV and ANTI0, this analysis also applies a veto on SAV in order to suppress the mis-reconstruction of signal events when a photon is lost in the center of LKr. The SAVMatching tool with a 20 ns time window with respect to the reference time is used. The wider time window is chosen to reflect the worse time resolution. This choice does not impact the random veto which is completely negligible.¹⁶

4.5.2 Background determination

Since about one π is observed per ten e in the single-track spectrum in the data, the combinatorial background, already negligible for e^+e^- , is completely absent in the sample obtained for $\pi^+\pi^-$. The $\pi^+\pi^-$ background is dominated by the *prompt component* - in-time tracks produced by a muon halo traversing the material, while additional radiated photons can mimic the signal of other decay channels. Potential upstream $K_S \rightarrow \pi^+\pi^-$ events are rejected by the ANTI0 acceptance requirement as the outgoing pions from this process have to pass through the ANTI0 central hole.

The prompt background simulation is based on a backward Monte Carlo technique, and the sample corresponding to 1.53×10^{17} POT used for the leptonic analyses will be employed for the background estimation also for this analysis.

In the following subsection 4.5.2, the background estimation outside the CR for all investigated decay channels will be determined and compared to the observation in the data. A background estimation in the CR and SR after applying the LAV, ANTI0, and SAV vetoes will be shown in subsection 4.5.2.

Estimation outside CR before applying vetoes

The LAV or ANTI0 veto removes all h^+h^- and h^+h^-X events in both data and MC. Also, the SR or CR requirement vetoes all SM background events in the MC. Therefore, the background estimation will be determined while lifting these cuts and employing the rejection factors $1 - \eta$ of SR/CR and LAV/ANTI0 cuts calculated for the prompt background in e^+e^- , shown in table 4.11. The table 4.14 shows a good data-MC agreement for $\pi^+\pi^-$ while for other final states, a more detailed study of the MC sample can help to determine N_{exp} .

¹⁶The random veto estimated with single tracks before applying ANTI0 and LAV veto is 3.8×10^{-5} . After applying ANTI0 and LAV veto, the statistics reduction is such that no out-of-time SAV hits are observed.

Channel	$N_{\text{exp}} \pm \delta N_{\text{exp}}$	N_{obs}	$p(L \leq L_{\text{obs}})$
$\pi^+\pi^-$	1.95 ± 1.24	2	0.76
$\pi^+\pi^- + \text{anything}$	0	0	-
$K^+K^- (+ \text{anything})$	0	0	-

Table 4.14: Expected number of $h^+h^- + \text{anything}$ vertices from the MC sample (N_{exp}) before requiring the LAV-ANTI0-SAV veto conditions, the related total uncertainty (δN_{exp}), the observed events in data (N_{obs}) and the p-values representing the probability to obtain a likelihood L for data-MC compatibility equal or worse than that corresponding to N_{obs} ($p(L \leq L_{\text{obs}})$).

The K^+K^- estimation can be calculated from the number of selected K^+ by RICH per hadron¹⁷ selected by the BDT PID n_{K^+/h^+} . The obtained ratio in a 68% CL score interval is:

$$n_{K^+/h^+} = 0.0082^{+0.0017}_{-0.0015}. \quad (4.21)$$

By using n_{K^+/h^+} and N_{exp} for $\pi^+\pi^-$, the number of K^+K^- outside CR is determined as

$$\begin{aligned} N_{\text{exp}}^{K^+K^-} &= 0.013^{+0.015}_{-0.008} @ 68\% \text{CL (frequentist)} \\ N_{\text{exp}}^{K^+K^-} &= 0.013^{+0.016}_{-0.008} @ 68\% \text{CL (bayesian)}, \end{aligned} \quad (4.22)$$

where a flat prior in the interval (4.21) is used in the frequentist approach while a β -distribution is used to model the PDF of (4.21) in the bayesian case.

In order to estimate the background for channels with additional photons in the final state, the particle identification requirement has to be lifted. By studying the spectra of radiated photons in two-track events, we can estimate the number of photons per two-track event $n_\gamma/n_{2\text{-track}}$ and project for the case of an h^+h^- event. Note that this estimation is conservative as, for example, an electron penetrating material radiates photons more easily than a hadron.

The obtained $n_\gamma/n_{2\text{-track}}$ at 68% CL is $n_\gamma/n_{2\text{-track}} = 0.077 \pm 0.005$. In both frequentist and bayesian scenarios, this gives

$$N_{\text{exp}}^{\pi^+\pi^-\gamma} = 0.11^{+0.13}_{-0.07}. \quad (4.23)$$

Since energy losses of a particle traversing material are well-described by the Landau distribution, as validated in left panel of figure 4.43 where the observed photon spectrum is fitted by the Landau distribution, the estimation

¹⁷This includes true kaons as well as pions mis-tagged as kaons by RICH.

of $n_\gamma/n_{2\text{-track}}$ can be improved further by accounting for the shape. As the estimation using a flat prior and a realistic $n_\gamma/n_{2\text{-track}}$ distribution already gives the same result within the decimal point, the same result is expected when using the Landau distribution.




Figure 4.43: Left: Energies of additional photons in $X^+X^-\gamma$ events in background MC where X can be any charged particle. The red line indicates the 5 GeV cut on event classification as an $h^+h^-\gamma$ topology. Right: Reconstructed m_{inv} of pairs of photons selected in X^+X^- events in prompt MC. Red lines indicate the π^0 and η mass windows used for classification as $X^+X^-\pi^0(\eta)$ topology. The green curve is a fit with Landau distribution. Both distributions are fitted with Landau distribution (blue).

For estimating backgrounds for more complex topologies with π^0 and η mesons, one can probe the spectrum of 2γ invariant masses m_{inv} reconstructed from two-track events with additional two photons. The spectrum obtained in the prompt MC fitted with the Landau distribution is plotted in right panel of figure 4.43 with indicated windows for the π^0 and η selection. The number of two-track events with π^0 or η in the final state $n_{\pi^0,\eta}/n_{2\text{-track}}$ is estimated using the fit from right panel of figure 4.43. At 68% CL, the values obtained are

$$\begin{aligned} n_{\pi^0}/n_{2\text{-track}} &= (8_{-4}^{+8}) \times 10^{-4} \\ n_{\eta}/n_{2\text{-track}} &= (5.3_{-4.7}^{+44.7}) \times 10^{-5}. \end{aligned} \quad (4.24)$$

Combined with the PID requirement, the expected background for the $\pi^+\pi^-\pi^0$ decay channel is

$$\begin{aligned} N_{\text{exp}}^{\pi^+\pi^-\pi^0} &= 0.0014_{-0.0009}^{+0.0020} @ 68\% \text{CL (frequentist)} \\ N_{\text{exp}}^{\pi^+\pi^-\pi^0} &= 0.0015_{-0.0011}^{+0.0026} @ 68\% \text{CL (bayesian)}, \end{aligned} \quad (4.25)$$

while for $\pi^+\pi^-\eta$:

$$\begin{aligned} N_{\text{exp}}^{\pi^+\pi^-\eta} &= 0.00032_{-0.00024}^{+0.00060} @ 68\% \text{CL (frequentist)} \\ N_{\text{exp}}^{\pi^+\pi^-\eta} &= 0.00044_{-0.00036}^{+0.00117} @ 68\% \text{CL (bayesian)}. \end{aligned} \quad (4.26)$$

The $K^+K^-\pi^0$ background estimation is given by the convolution of n_{K^+/h^+} and $n_{\pi^0}/n_{2\text{-track}}$ and yields

$$\begin{aligned} N_{\text{exp}}^{K^+K^-\pi^0} &= (1.2_{-0.8}^{+1.7}) \times 10^{-5} @ 68\% \text{CL (frequentist)} \\ N_{\text{exp}}^{K^+K^-\pi^0} &= (1.3_{-0.9}^{+2.3}) \times 10^{-5} @ 68\% \text{CL (bayesian)}. \end{aligned} \quad (4.27)$$

For $\pi^+\pi^-2\pi^0$, N_{exp} is driven by the probability to have 4γ accidentally forming two π^0 . The corresponding estimation is

$$\begin{aligned} N_{\text{exp}}^{\pi^+\pi^-2\pi^0} &= 0.0007_{-0.0006}^{+0.0014} @ 68\% \text{CL (frequentist)} \\ N_{\text{exp}}^{\pi^+\pi^-2\pi^0} &= 0.0004_{-0.0003}^{+0.0011} @ 68\% \text{CL (bayesian)}. \end{aligned} \quad (4.28)$$

There are not enough statistics of reconstructed π to estimate the shape of the $\pi\pi$ invariant mass distribution to calculate the expected number of on-shell 2ρ events but given that based on the accidental matching of π necessarily $N_{\text{exp}}(2\rho) \ll N_{\text{exp}}(4\pi)$, the expectation falls in the confidence interval of 4π .

Expected background in the signal region

The rejection factors of the LAV-ANTI0 and CR(SR) cuts ($1 - \eta_{\text{LAV,ANTI0}}$ and $1 - \eta_{\text{CR(SR)}}$) have been calculated for the $A' \rightarrow e^+e^-$ analysis, giving a combined rejection of about a factor of 10^{-4} . The extension of the CR and SR compared to the one used in $A' \rightarrow e^+e^-$ does not modify $1 - \eta_{\text{CR(SR)}}$ as the e^+e^- events observed in data and MC are still contained within the same CR and SR and their number does not change. After applying the full selection, the expected backgrounds for all channels studied are summarized in table 4.15. As advertised in section 4.5.1 the expected background is lowest for the most complex topologies and highest for the simple h^+h^- decay. Even for this simplest topology, the expected background is much lower than the expectation in the $A' \rightarrow \ell^+\ell^-$ analyses.

Channel	$N_{\text{exp,CR}} \pm \delta N_{\text{exp,CR}}$	$N_{\text{exp,SR}} \pm \delta N_{\text{exp,SR}}$
$\pi^+ \pi^-$	$0.0005_{-0.0004}^{+0.0017}$	$0.0005_{-0.0004}^{+0.0017}$
$\pi^+ \pi^- \gamma$	$(2.5_{-2.1}^{+10.1}) \times 10^{-5}$	$(2.4_{-2.1}^{+9.8}) \times 10^{-5}$
$\pi^+ \pi^- \pi^0$	$(3.3_{-2.9}^{+13.8}) \times 10^{-7}$	$(3.2_{-2.8}^{+13.4}) \times 10^{-7}$
$\pi^+ \pi^- \pi^0 \pi^0 (\rho\rho)$	$(1.5_{-1.3}^{+6.4}) \times 10^{-7}$	$(1.4_{-1.4}^{+6.3}) \times 10^{-7}$
$\pi^+ \pi^- \eta$	$(9.3_{-8.1}^{+40.2}) \times 10^{-8}$	$(8.9_{-7.8}^{+39.1}) \times 10^{-8}$
$K^+ K^-$	$(2.9_{-2.5}^{+12.0}) \times 10^{-6}$	$(2.8_{-2.4}^{+11.6}) \times 10^{-6}$
$K^+ K^- \pi^0$	$(3.8_{-3.2}^{+10.9}) \times 10^{-9}$	$(3.1_{-3.1}^{+10.6}) \times 10^{-9}$

Table 4.15: Summary of expected number of background events at 68% CL for all studied decay channels from table 4.13 in CR and SR after full selection.

4.5.3 Efficiency determination

The signal selection efficiency is calculated as a function of mass and decay width for each exotic particle. The efficiency is calculated for each production and decay channel in order to allow combining the various channels for different new physics models. An MC sample with the exotic particle mass range starting at the kinematic threshold and going up to $m_X = 3 \text{ GeV}$ with logarithmic step is produced for each production and decay channel with 2×10^5 events per mass bin. The simulated exotic particles are decayed uniformly along their path in the range of Z_{VTX} positions $102 \text{ m} < Z_{\text{VTX}} < 182 \text{ m}$, covering the NA62 decay volume with a crossover of about 2 m on both sides to account for the resolution of the vertex position reconstruction in the modified FV.

Except for the particle identification, the particle-type-independent selection applied in this analysis is consistent with the one used in 4.4.2, and therefore, the systematic errors and the corrections to the efficiencies derived from the MC are also applicable here. The PID efficiency for this analysis is derived in the following section.

Particle ID efficiency

To determine the PID efficiency and the mis-identification probability, a single track pion sample has been produced with NA62 MC. Same as for the μ and e samples, single track data events forming the beam-dump muon halo at the $z = 102 \text{ m}$ plane are stored and used as a gun for the MC with enforced particle ID of a pion. The generated sample is compared with kaon $K_{2\pi}$ data and both overlaid and non-overlaid samples.

In figure 4.44 left, it is observed that the non-overlaid sample distribution follows as expected, the single pion sample in the low momentum region. The total observed inefficiency of the BDT pion PID yields 4.7%, confirming the much better performance of the BDT PID for hadrons compared to the E/p -based PID. To estimate the correction applied for the MC, the overlaid MC is compared with the data to account for inefficiencies caused by $K_{2\pi}$ event mis-reconstruction. Except for the very low momentum region, negligible data-MC difference is observed (well below 1% level), implicating that no additional correction to the signal MC is needed. Systematic PID uncertainty of 0.4% is assigned to the MC based on the average data-MC discrepancy observed.

The larger data-MC difference of BDT-based PID at low momentum, however, calls for cross-checking the data-MC difference for the mis-identification of electrons as pions since the electron background is typically soft, which can potentially cause a surplus of π -identified background events in data compared to the MC. A kaon $K_{2\pi D}$ sample is probed both in data and MC. As shown in figure 4.44 right, no deviation of the $e \rightarrow \pi$ mis-identification is observed in the comparison of overlaid MC to the data, underlining the credibility of the MC-derived expected backgrounds from the previous section.




Figure 4.44: Left: comparison of π PID inefficiencies as a function of the track momentum for single track MC (black), non-overlaid $K_{2\pi}$ MC (blue), overlaid MC (green) and data (red). Right: comparison of $e \rightarrow \pi$ mis-ID probabilities for $K_{2\pi D}$ data and MC samples.

Source	Uncertainty	Relative correction
Statistics	2.1%	–
NewCHOD association	0.6%	–1.2%
Straw quality and resolution	1.5%	–
Trigger	0.5%	–2.0%
LAV random veto	0.1%	–2.5%
ANTIO random veto	0.1%	–0.6%
PID	0.4%	–
Total	2.7%	–6.3%

Table 4.16: Summary of uncertainties and corrections to the signal selection efficiency.

Summary of signal selection efficiency

A summary of uncertainties and corrections for the signal selection efficiency is given in table 4.16. The total relative uncertainty on the efficiency is about 2.7%. The MC efficiency needs to be corrected by -6.3% to account for data/MC disagreements.

4.5.4 Determination of the expected sensitivity

After accounting for the selection efficiency, the expected signal yield in the 2021 data sample is shown for an ALP produced in a Primakoff process and decaying into $\pi^+\pi^-\gamma$ for a fixed value of the production coupling $C_{\gamma\gamma}/\Lambda = 10^{-4}$ in left panel of figure 4.45.




Figure 4.45: Expected number of events (left) and standard deviation (MeV) of the invariant-mass peak (right) for the $X \rightarrow \pi^+\pi^-\gamma$ decay after full selection in Γ_X - m_X plane for a Primakoff-produced ALP with a fixed value of the production coupling $C_{\gamma\gamma}/\Lambda = 10^{-4}$ (note that Γ_X is kept variable, while $C_{\gamma\gamma}$ represents only the normalization on the number of events on the production side). The events from signal regions of mis-reconstructed final state topologies are taken into account. Right:

The resolution of the reconstructed exotic particle mass m_X varies from $\mathcal{O}(1 \text{ MeV})$ at low masses to $\mathcal{O}(10 \text{ MeV})$ for heavy exotic particles. In the low-mass limit near the kinematic threshold for given decay, worsening of the mass resolution can also occur, notably in events with additional photons; see right panel of figure 4.45 for the abovementioned case of a Primakoff-produced ALP.

The calculation for other exotic production and decay processes is completely analogical, except for the mixing production, where additional correction to the production yield is done as a function of the exotic particle mass. The correction and the related systematic uncertainty are calculated in appendix A.

To finalize this analysis, the evaluation of the number of expected events using the full MC samples and statistical combination of the various channels is in progress.

Conclusion

The thesis discusses searches for hidden sector mediators in the MeV–GeV mass range with the focus on dark photon, axion-like and Higgs-like portals, first in general at fixed-target experiments and then in the specific case of NA62.

The aim of the phenomenology part of the project was to develop a tool which would allow to re-interpret experimental data for a custom hidden sector model where, upon plugging in the experimental data and choosing the setup of the couplings at the UV scale, the user would obtain the expected number of observed events for custom set of parameters.¹⁸ For this purpose numerous production and decay mechanisms effective in the MeV–GeV mass range were implemented. The tool also includes a toy Monte Carlo of a custom setup of proton fixed-target experiment with several setups for past and future experiments pre-included. This allows production of test data in case that true experimental data is not available. Combining different parts of the tool resulted in a unified software framework which is now available in a public repository. Extension of the framework for heavy neutral leptons is under development, as well as the inclusion of additional hidden sector particle production and decay modes.

Using the framework, several exclusion bounds of past experiments were re-calculated, including processes not accounted for before. By comparing with the past experiments, the discovery potential of future prospects has been determined, each profiting from a specific geometric setup. The framework also showed the potential for discovering axion-like particles in a Standard Model search for a $K_L \rightarrow \pi^0 \nu \bar{\nu}$ decay originating from a completely independent process which allows for axion masses higher than what would be achievable with kaon decays. Besides the reinterpretation of the results for different models, the toy Monte Carlo of the framework has proven to be use-

¹⁸The parameters can be of purely kinematical nature, such as mass and lifetime or decay width, or it can be the branching ratios for the production and decay processes or the couplings determining these variables.

ful for kinematics studies, for emulating analysis cuts but also for optimization of the design of future detectors.

A simulation using the framework toy Monte Carlo has shown that the currently operating NA62 experiment, which studies a kaon beam, also has a discovery potential for hidden sector mediators too heavy to be produced in kaon decays. For this purpose, the data at NA62 have to be collected in a beam-dump mode, to ensure a background-free environment. In 2021, the experiment has collected 1.4×10^{17} protons on target in the beam-dump mode. The experimental part of this project aimed to probe several hidden sector scenarios for which NA62 has a discovery potential with the 2021 statistics.

All the analyses performed require a hidden sector mediator decay into Standard Model final states, out of which two are electrically charged and are of the same type. These searches share the same type of background, namely combinatorial background from random pairing of uncorrelated halo particles and in-time background from interaction of the halo particles in the material of the detector. The electrically charged final states allow the kinematic reconstruction of the original particle and the location of its production, keeping effectively the background under control. These common characteristics resulted in the development of an analyzer whose base is the same for the various types of final states, mainly differing in the final state reconstruction and slight changes to the signal region. A high efficiency of the final state reconstruction was possible thanks to the knowledge of the detector response studied thoroughly using the kaon data collected in the standard data-taking.

The results of the first two analyzes, probing decays into a pair of leptons, have been already shown publicly. No decisive signal was observed and an upper bound on the observation has been put. A third analysis, probing decays into hadrons, is being finalized.

Another data sample, an equivalent of 2.5×10^{17} protons on target, has been collected at NA62 in 2023. The reconstruction of the sample is in progress. Besides extending the reach of the abovementioned analyses, with the combined statistics, NA62 also has the potential to probe semi-leptonic decays of hidden sector mediators and decays with purely neutral final states if the background can be kept under control. A long-term plan at NA62 until the end of 2025 is the collection of 1×10^{18} protons on target. A successor experiment to NA62 with even better sensitivity in kaon mode and planned collection of 5×10^{19} protons on target in beam dump mode is proposed beyond 2025.

Bibliography

- [1] J. Jerhot, B. Döbrich, F. Ertas, F. Kahlhoefer, and T. Spadaro, “ALPIN-IST: Axion-Like Particles In Numerous Interactions Simulated and Tabulated,” *JHEP*, vol. 07, p. 094, 2022, 2201.05170.
- [2] Y. Afik, B. Döbrich, J. Jerhot, Y. Soreq, and K. Tobioka, “Probing long-lived axions at the KOTO experiment,” *Phys. Rev. D*, vol. 108, no. 5, p. 055007, 2023, 2303.01521.
- [3] B. Döbrich, S. Ghinescu, J. Jerhot, A. Kleimenova, E. Minucci, and T. Spadaro, “Search for dark photons decaying into a $\mu^+\mu^-$ pair,” *Internal note: NA62-22-07*, 2022.
- [4] B. Döbrich, S. Ghinescu, J. Jerhot, A. Kleimenova, E. Minucci, and T. Spadaro, “Search for dark photons decaying into a e^+e^- pair,” *Internal note: NA62-23-01*, 2023.
- [5] B. Döbrich, S. Ghinescu, J. Jerhot, A. Kleimenova, E. Minucci, and T. Spadaro, “In preparation,” *Internal note: NA62-23-xx*, 2023.
- [6] E. Cortina Gil *et al.*, “HIKE, High Intensity Kaon Experiments at the CERN SPS: Letter of Intent,” 11 2022, 2211.16586.
- [7] M. U. Ashraf *et al.*, “High Intensity Kaon Experiments (HIKE) at the CERN SPS Proposal for Phases 1 and 2,” 11 2023, 2311.08231.
- [8] E. Noether, “Invariant Variation Problems,” *Gott. Nachr.*, vol. 1918, pp. 235–257, 1918, physics/0503066.
- [9] K. G. Wilson, “Confinement of Quarks,” *Phys. Rev. D*, vol. 10, pp. 2445–2459, 1974.
- [10] T. Nakano and K. Nishijima, “Charge Independence for V-particles,” *Prog. Theor. Phys.*, vol. 10, pp. 581–582, 1953.
- [11] K. Nishijima, “Charge Independence Theory of V Particles,” *Prog. Theor. Phys.*, vol. 13, no. 3, pp. 285–304, 1955.

- [12] M. Gell-Mann, “The interpretation of the new particles as displaced charge multiplets,” *Nuovo Cim.*, vol. 4, no. S2, pp. 848–866, 1956.
- [13] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons,” *Phys. Rev. Lett.*, vol. 13, pp. 321–323, 1964.
- [14] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons,” *Phys. Rev. Lett.*, vol. 13, pp. 508–509, 1964.
- [15] P. A. Zyla *et al.*, “Review of Particle Physics,” *PTEP*, vol. 2020, no. 8, p. 083C01, 2020.
- [16] S. Weinberg, “A Model of Leptons,” *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 1967.
- [17] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, “Experimental Test of Parity Conservation in β Decay,” *Phys. Rev.*, vol. 105, pp. 1413–1414, 1957.
- [18] S. L. Adler, “Axial-vector vertex in spinor electrodynamics,” *Phys. Rev.*, vol. 177, pp. 2426–2438, Jan 1969.
- [19] J. S. Bell and R. Jackiw, “A PCAC puzzle: $\pi^0 \rightarrow \gamma\gamma$ in the σ model,” *Nuovo Cim. A*, vol. 60, pp. 47–61, 1969.
- [20] R. L. Workman *et al.*, “Review of Particle Physics,” *PTEP*, vol. 2022, p. 083C01, 2022.
- [21] M. Peskin and D. Schroeder, *An Introduction To Quantum Field Theory*. Frontiers in Physics, Avalon Publishing, 1995.
- [22] S. Scherer and M. Schindler, *A Primer for Chiral Perturbation Theory*. Lecture Notes in Physics, Springer Berlin Heidelberg, 2011.
- [23] M. Gell-Mann and M. Levy, “The axial vector current in beta decay,” *Nuovo Cim.*, vol. 16, p. 705, 1960.
- [24] Y. Nambu and D. Lurie, “Chirality conservation and soft pion production,” *Phys. Rev.*, vol. 125, pp. 1429–1436, 1962.
- [25] S. L. Adler, “Tests of the Conserved Vector Current and Partially Conserved Axial-Vector Current Hypotheses in High-Energy Neutrino Reactions,” *Phys. Rev.*, vol. 135, pp. B963–B966, 1964.

- [26] S. Weinberg, “Current Commutator Calculation of the $K_{\ell 4}$ form Factors,” *Phys. Rev. Lett.*, vol. 17, pp. 336–340, 1966. [Erratum: *Phys.Rev.Lett.* 18, 1178 (1967)].
- [27] S. L. Adler, “Consistency conditions on the strong interactions implied by a partially conserved axial vector current,” *Phys. Rev.*, vol. 137, pp. B1022–B1033, 1965.
- [28] S. L. Adler, “Consistency conditions on the strong interactions implied by a partially conserved axial-vector current. II,” *Phys. Rev.*, vol. 139, pp. B1638–B1643, 1965.
- [29] S. R. Coleman, J. Wess, and B. Zumino, “Structure of phenomenological Lagrangians. 1.,” *Phys. Rev.*, vol. 177, pp. 2239–2247, 1969.
- [30] C. G. Callan, Jr., S. R. Coleman, J. Wess, and B. Zumino, “Structure of phenomenological Lagrangians. 2.,” *Phys. Rev.*, vol. 177, pp. 2247–2250, 1969.
- [31] I. Low, “Adler’s zero and effective Lagrangians for nonlinearly realized symmetry,” *Phys. Rev. D*, vol. 91, no. 10, p. 105017, 2015, 1412.2145.
- [32] B. Abi *et al.*, “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm,” *Phys. Rev. Lett.*, vol. 126, no. 14, p. 141801, 2021, 2104.03281.
- [33] M. Dainotti, B. De Simone, G. Montani, T. Schiavone, and G. Lambiase, “The Hubble constant tension: current status and future perspectives through new cosmological probes,” in *CORFU2022: 22th Hellenic School and Workshops on Elementary Particle Physics and Gravity*, 1 2023, 2301.10572.
- [34] P. Minkowski, “ $\mu \rightarrow e\gamma$ at a Rate of One Out of 10^9 Muon Decays?,” *Phys. Lett. B*, vol. 67, pp. 421–428, 1977.
- [35] T. Yanagida, “Horizontal gauge symmetry and masses of neutrinos,” *Conf. Proc. C*, vol. 7902131, pp. 95–99, 1979.
- [36] M. Magg and C. Wetterich, “Neutrino Mass Problem and Gauge Hierarchy,” *Phys. Lett. B*, vol. 94, pp. 61–64, 1980.
- [37] R. Foot, H. Lew, X. G. He, and G. C. Joshi, “Seesaw Neutrino Masses Induced by a Triplet of Leptons,” *Z. Phys. C*, vol. 44, p. 441, 1989.

- [38] T. Asaka and M. Shaposhnikov, “The ν MSM, dark matter and baryon asymmetry of the universe,” *Phys. Lett. B*, vol. 620, pp. 17–26, 2005, hep-ph/0505013.
- [39] R. V. Wagoner, W. A. Fowler, and F. Hoyle, “On the Synthesis of elements at very high temperatures,” *Astrophys. J.*, vol. 148, pp. 3–49, 1967.
- [40] D. N. Schramm and M. S. Turner, “Big Bang Nucleosynthesis Enters the Precision Era,” *Rev. Mod. Phys.*, vol. 70, pp. 303–318, 1998, astro-ph/9706069.
- [41] L. Canetti, M. Drewes, and M. Shaposhnikov, “Matter and Antimatter in the Universe,” *New J. Phys.*, vol. 14, p. 095012, 2012, 1204.4186.
- [42] A. D. Sakharov, “Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe,” *Pisma Zh. Eksp. Teor. Fiz.*, vol. 5, pp. 32–35, 1967.
- [43] C. Cao, M. van Caspel, and A. R. Zhitnitsky, “Topological Casimir effect in Maxwell Electrodynamics on a Compact Manifold,” *Phys. Rev. D*, vol. 87, no. 10, p. 105012, 2013, 1301.1706.
- [44] R. Jackiw and C. Rebbi, “Vacuum periodicity in a yang-mills quantum theory,” *Phys. Rev. Lett.*, vol. 37, pp. 172–175, Jul 1976.
- [45] C. Callan, R. Dashen, and D. Gross, “The structure of the gauge theory vacuum,” *Physics Letters B*, vol. 63, no. 3, pp. 334–340, 1976.
- [46] A. Bilal, “Lectures on Anomalies,” 2008, 0802.0634.
- [47] G. 't Hooft, “Symmetry Breaking Through Bell-Jackiw Anomalies,” *Phys. Rev. Lett.*, vol. 37, pp. 8–11, 1976.
- [48] M. Fukugita and T. Yanagida, “Baryogenesis Without Grand Unification,” *Phys. Lett. B*, vol. 174, pp. 45–47, 1986.
- [49] K. Kajantie, M. Laine, K. Rummukainen, and M. E. Shaposhnikov, “A Nonperturbative analysis of the finite T phase transition in SU(2) x U(1) electroweak theory,” *Nucl. Phys. B*, vol. 493, pp. 413–438, 1997, hep-lat/9612006.
- [50] S. Weinberg, “The U(1) Problem,” *Phys. Rev. D*, vol. 11, pp. 3583–3593, 1975.

- [51] M. L. Goldberger and S. B. Treiman, “Decay of the pi meson,” *Phys. Rev.*, vol. 110, pp. 1178–1184, 1958.
- [52] R. Crewther, P. Di Vecchia, G. Veneziano, and E. Witten, “Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics,” *Physics Letters B*, vol. 88, no. 1, pp. 123–127, 1979.
- [53] A. Hook, “TASI Lectures on the Strong CP Problem and Axions,” *PoS*, vol. TASI2018, p. 004, 2019, 1812.02669.
- [54] C. Abel *et al.*, “Measurement of the Permanent Electric Dipole Moment of the Neutron,” *Phys. Rev. Lett.*, vol. 124, no. 8, p. 081803, 2020, 2001.11966.
- [55] S. Khlebnikov and M. Shaposhnikov, “Extra space-time dimensions: Towards a solution to the strong cp-problem,” *Physics Letters B*, vol. 203, no. 1, pp. 121–124, 1988.
- [56] W.-Y. Ai, J. S. Cruz, B. Garbrecht, and C. Tamarit, “Consequences of the order of the limit of infinite spacetime volume and the sum over topological sectors for CP violation in the strong interactions,” *Phys. Lett. B*, vol. 822, p. 136616, 2021, 2001.07152.
- [57] N. Kaloper and J. Terning, “Landscaping the Strong CP Problem,” *JHEP*, vol. 03, p. 032, 2019, 1710.01740.
- [58] R. D. Peccei and H. R. Quinn, “CP Conservation in the Presence of Instantons,” *Phys. Rev. Lett.*, vol. 38, pp. 1440–1443, 1977.
- [59] R. D. Peccei and H. R. Quinn, “Constraints Imposed by CP Conservation in the Presence of Instantons,” *Phys. Rev. D*, vol. 16, pp. 1791–1797, 1977.
- [60] S. Weinberg, “A New Light Boson?,” *Phys. Rev. Lett.*, vol. 40, pp. 223–226, 1978.
- [61] F. Wilczek, “Problem of Strong P and T Invariance in the Presence of Instantons,” *Phys. Rev. Lett.*, vol. 40, pp. 279–282, 1978.
- [62] M. Kuster, G. Raffelt, and B. Beltrán, *Axions: Theory, Cosmology, and Experimental Searches*. Lecture Notes in Physics, Springer Berlin Heidelberg, 2007.

- [63] R. D. Peccei, “The Strong CP problem and axions,” *Lect. Notes Phys.*, vol. 741, pp. 3–17, 2008, hep-ph/0607268.
- [64] J. E. Kim, “Weak Interaction Singlet and Strong CP Invariance,” *Phys. Rev. Lett.*, vol. 43, p. 103, 1979.
- [65] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, “Can Confinement Ensure Natural CP Invariance of Strong Interactions?,” *Nucl. Phys. B*, vol. 166, pp. 493–506, 1980.
- [66] M. Dine, W. Fischler, and M. Srednicki, “A Simple Solution to the Strong CP Problem with a Harmless Axion,” *Phys. Lett. B*, vol. 104, pp. 199–202, 1981.
- [67] A. R. Zhitnitsky, “On Possible Suppression of the Axion Hadron Interactions. (In Russian),” *Sov. J. Nucl. Phys.*, vol. 31, p. 260, 1980.
- [68] D. J. E. Marsh, “Axion Cosmology,” *Phys. Rept.*, vol. 643, pp. 1–79, 2016, 1510.07633.
- [69] M. Kawasaki and K. Nakayama, “Axions: Theory and Cosmological Role,” *Ann. Rev. Nucl. Part. Sci.*, vol. 63, pp. 69–95, 2013, 1301.1123.
- [70] H. Fukuda, K. Harigaya, M. Ibe, and T. T. Yanagida, “Model of visible QCD axion,” *Phys. Rev. D*, vol. 92, no. 1, p. 015021, 2015, 1504.06084.
- [71] P. Agrawal, G. Marques-Tavares, and W. Xue, “Opening up the QCD axion window,” *JHEP*, vol. 03, p. 049, 2018, 1708.05008.
- [72] P. Agrawal and K. Howe, “Factoring the Strong CP Problem,” *JHEP*, vol. 12, p. 029, 2018, 1710.04213.
- [73] M. K. Gaillard, M. B. Gavela, R. Houtz, P. Quilez, and R. Del Rey, “Color unified dynamical axion,” *Eur. Phys. J. C*, vol. 78, no. 11, p. 972, 2018, 1805.06465.
- [74] T. Gherghetta, V. V. Khoze, A. Pomarol, and Y. Shirman, “The Axion Mass from 5D Small Instantons,” *JHEP*, vol. 03, p. 063, 2020, 2001.05610.
- [75] R. S. Gupta, V. V. Khoze, and M. Spannowsky, “Small instantons and the strong CP problem in composite Higgs models,” *Phys. Rev. D*, vol. 104, no. 7, p. 075011, 2021, 2012.00017.

- [76] T. Gherghetta and M. D. Nguyen, “A Composite Higgs with a Heavy Composite Axion,” *JHEP*, vol. 12, p. 094, 2020, 2007.10875.
- [77] A. Valenti, L. Vecchi, and L.-X. Xu, “Grand Color axion,” *JHEP*, vol. 10, p. 025, 2022, 2206.04077.
- [78] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, “Flavor probes of axion-like particles,” *JHEP*, vol. 09, p. 056, 2022, 2110.10698.
- [79] E. Izaguirre, T. Lin, and B. Shuve, “Searching for Axionlike Particles in Flavor-Changing Neutral Current Processes,” *Phys. Rev. Lett.*, vol. 118, no. 11, p. 111802, 2017, 1611.09355.
- [80] S. Chakraborty, M. Kraus, V. Loladze, T. Okui, and K. Tobioka, “Heavy QCD axion in $b \rightarrow s$ transition: Enhanced limits and projections,” *Phys. Rev. D*, vol. 104, no. 5, p. 055036, 2021, 2102.04474.
- [81] H. Georgi, D. B. Kaplan, and L. Randall, “Manifesting the Invisible Axion at Low-energies,” *Phys. Lett. B*, vol. 169, pp. 73–78, 1986.
- [82] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, “The Low-Energy Effective Theory of Axions and ALPs,” *JHEP*, vol. 04, p. 063, 2021, 2012.12272.
- [83] Y. Chikashige, R. N. Mohapatra, and R. D. Peccei, “Are There Real Goldstone Bosons Associated with Broken Lepton Number?,” *Phys. Lett. B*, vol. 98, pp. 265–268, 1981.
- [84] F. Wilczek, “Axions and Family Symmetry Breaking,” *Phys. Rev. Lett.*, vol. 49, pp. 1549–1552, 1982.
- [85] R. Coquereaux and G. Esposito-Farese, “The Theory of Kaluza-Klein-Jordan-Thiry revisited,” *Ann. Inst. H. Poincaré Phys. Theor.*, vol. 52, pp. 113–150, 1990.
- [86] N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, “The Hierarchy problem and new dimensions at a millimeter,” *Phys. Lett. B*, vol. 429, pp. 263–272, 1998, hep-ph/9803315.
- [87] K. R. Dienes, E. Dudas, and T. Gherghetta, “Invisible axions and large radius compactifications,” *Phys. Rev. D*, vol. 62, p. 105023, 2000, hep-ph/9912455.

- [88] E. Witten, “Search for a Realistic Kaluza-Klein Theory,” *Nucl. Phys. B*, vol. 186, p. 412, 1981.
- [89] P. Svrcek and E. Witten, “Axions In String Theory,” *JHEP*, vol. 06, p. 051, 2006, hep-th/0605206.
- [90] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, “String Axiverse,” *Phys. Rev. D*, vol. 81, p. 123530, 2010, 0905.4720.
- [91] M. Bauer, M. Neubert, and A. Thamm, “Collider Probes of Axion-Like Particles,” *JHEP*, vol. 12, p. 044, 2017, 1708.00443.
- [92] G. D’Ambrosio, G. F. Giudice, G. Isidori, and A. Strumia, “Minimal flavor violation: An Effective field theory approach,” *Nucl. Phys. B*, vol. 645, pp. 155–187, 2002, hep-ph/0207036.
- [93] M. J. Dolan, F. Kahlhoefer, C. McCabe, and K. Schmidt-Hoberg, “A taste of dark matter: Flavour constraints on pseudoscalar mediators,” *JHEP*, vol. 03, p. 171, 2015, 1412.5174. [Erratum: *JHEP* 07, 103 (2015)].
- [94] B. Döbrich, F. Ertas, F. Kahlhoefer, and T. Spadaro, “Model-independent bounds on light pseudoscalars from rare B-meson decays,” *Phys. Lett. B*, vol. 790, pp. 537–544, 2019, 1810.11336.
- [95] D. Aloni, Y. Soreq, and M. Williams, “Coupling QCD-Scale Axionlike Particles to Gluons,” *Phys. Rev. Lett.*, vol. 123, no. 3, p. 031803, 2019, 1811.03474.
- [96] G. Grilli di Cortona, E. Hardy, J. Pardo Vega, and G. Villadoro, “The QCD axion, precisely,” *JHEP*, vol. 01, p. 034, 2016, 1511.02867.
- [97] T. Fujiwara, T. Kugo, H. Terao, S. Uehara, and K. Yamawaki, “Non-abelian Anomaly and Vector Mesons as Dynamical Gauge Bosons of Hidden Local Symmetries,” *Prog. Theor. Phys.*, vol. 73, p. 926, 1985.
- [98] G. Dalla Valle Garcia, F. Kahlhoefer, M. Ovchinnikov, and A. Zaporozhchenko, “Phenomenology of axion-like particles with universal fermion couplings – revisited,” 10 2023, 2310.03524.

- [99] S. Weinberg, “Photons and Gravitons in S -Matrix Theory: Derivation of Charge Conservation and Equality of Gravitational and Inertial Mass,” *Phys. Rev.*, vol. 135, pp. B1049–B1056, 1964.
- [100] S. Weinberg and E. Witten, “Limits on Massless Particles,” *Phys. Lett. B*, vol. 96, pp. 59–62, 1980.
- [101] J. Beacham *et al.*, “Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report,” *J. Phys. G*, vol. 47, no. 1, p. 010501, 2020, 1901.09966.
- [102] C. Antel *et al.*, “Feebly Interacting Particles: FIPs 2022 workshop report,” in *Workshop on Feebly-Interacting Particles*, 5 2023, 2305.01715.
- [103] D. P. Aguillard *et al.*, “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.20 ppm,” 8 2023, 2308.06230.
- [104] M. Pospelov, “Secluded $U(1)$ below the weak scale,” *Phys. Rev. D*, vol. 80, p. 095002, 2009, 0811.1030.
- [105] P. Ilten, Y. Soreq, M. Williams, and W. Xue, “Serendipity in dark photon searches,” *JHEP*, vol. 06, p. 004, 2018, 1801.04847.
- [106] M. W. Winkler, “Decay and detection of a light scalar boson mixing with the Higgs boson,” *Phys. Rev. D*, vol. 99, no. 1, p. 015018, 2019, 1809.01876.
- [107] G. Krnjaic, “Probing Light Thermal Dark-Matter With a Higgs Portal Mediator,” *Phys. Rev. D*, vol. 94, no. 7, p. 073009, 2016, 1512.04119.
- [108] C. Bird, P. Jackson, R. V. Kowalewski, and M. Pospelov, “Search for dark matter in $b \rightarrow s$ transitions with missing energy,” *Phys. Rev. Lett.*, vol. 93, p. 201803, 2004, hep-ph/0401195.
- [109] “Combination of searches for invisible decays of the Higgs boson using 139 fb^{-1} of proton-proton collision data at $\sqrt{s}=13 \text{ TeV}$ collected with the ATLAS experiment,” *Phys. Lett. B*, vol. 842, p. 137963, 2023, 2301.10731.
- [110] J. Jerhot, B. Döbrich, E. Ertas, F. Kahlhoefer, and T. Spadaro, “ALPINIST: v1.3.0,” Jan. 2022. 10.5281/zenodo.5844011.

- [111] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, “An introduction to PYTHIA 8.2,” *Comput. Phys. Commun.*, vol. 191, pp. 159–177, 2015, 1410.3012.
- [112] B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald, and K. Schmidt-Hoberg, “ALPtraum: ALP production in proton beam dump experiments,” *JHEP*, vol. 02, p. 018, 2016, 1512.03069.
- [113] L. Harland-Lang, J. Jaeckel, and M. Spannowsky, “A fresh look at ALP searches in fixed target experiments,” *Phys. Lett. B*, vol. 793, pp. 281–289, 2019, 1902.04878.
- [114] B. Döbrich, J. Jaeckel, and T. Spadaro, “Light in the beam dump - ALP production from decay photons in proton beam-dumps,” *JHEP*, vol. 05, p. 213, 2019, 1904.02091. [Erratum: *JHEP* 10, 046 (2020)].
- [115] J. Blümlein and J. Brunner, “New Exclusion Limits on Dark Gauge Forces from Proton Bremsstrahlung in Beam-Dump Data,” *Phys. Lett. B*, vol. 731, pp. 320–326, 2014, 1311.3870.
- [116] F. Ertas and F. Kahlhoefer, “On the interplay between astrophysical and laboratory probes of MeV-scale axion-like particles,” *JHEP*, vol. 07, p. 050, 2020, 2004.01193.
- [117] C. Lourenco and H. K. Wohri, “Heavy flavour hadro-production from fixed-target to collider energies,” *Phys. Rept.*, vol. 433, pp. 127–180, 2006, hep-ph/0609101.
- [118] P. Ball and R. Zwicky, “New results on $B \rightarrow \pi, K, \eta$ decay formfactors from light-cone sum rules,” *Phys. Rev. D*, vol. 71, p. 014015, 2005, hep-ph/0406232.
- [119] P. Ball and R. Zwicky, “ $B_{d,s} \rightarrow \rho, \omega, K^*, \phi$ decay form-factors from light-cone sum rules revisited,” *Phys. Rev. D*, vol. 71, p. 014029, 2005, hep-ph/0412079.
- [120] V. Lubicz, L. Riggio, G. Salerno, S. Simula, and C. Tarantino, “Scalar and vector form factors of $D \rightarrow \pi(K)\ell\nu$ decays with $N_f = 2 + 1 + 1$ twisted fermions,” *Phys. Rev. D*, vol. 96, no. 5, p. 054514, 2017, 1706.03017. [Erratum: *Phys.Rev.D* 99, 099902 (2019), Erratum: *Phys.Rev.D* 100, 079901 (2019)].

- [121] J. Martin Camalich, M. Pospelov, P. N. H. Vuong, R. Ziegler, and J. Zupan, “Quark Flavor Phenomenology of the QCD Axion,” *Phys. Rev. D*, vol. 102, no. 1, p. 015023, 2020, 2002.04623.
- [122] A. Carmona, C. Scherb, and P. Schwaller, “Charming ALPs,” *JHEP*, vol. 08, p. 121, 2021, 2101.07803.
- [123] B. Vonck, “Charm production in 400 GeV/c pp interactions,” in *22nd Rencontres de Moriond: Hadrons, Quarks and Gluons*, pp. 77–84, 1987.
- [124] M. Drewes, J. Hajer, J. Klaric, and G. Lanfranchi, “NA62 sensitivity to heavy neutral leptons in the low scale seesaw model,” *JHEP*, vol. 07, p. 105, 2018, 1801.04207.
- [125] R. Brun and F. Rademakers, “Root - an object oriented data analysis framework,” in *AIHENP’96 Workshop, Lausanne*, vol. 389, pp. 81–86, 1996.
- [126] F. Domingo, “Decays of a NMSSM CP-odd Higgs in the low-mass region,” *JHEP*, vol. 03, p. 052, 2017, 1612.06538.
- [127] H.-C. Cheng, L. Li, and E. Salvioni, “A theory of dark pions,” *JHEP*, vol. 01, p. 122, 2022, 2110.10691.
- [128] F. Bergsma *et al.*, “Search for Axion Like Particle Production in 400-GeV Proton - Copper Interactions,” *Phys. Lett. B*, vol. 157, pp. 458–462, 1985.
- [129] F. Bergsma *et al.*, “Bounds on Supersymmetric Particles From a Proton Beam Dump Experiment,” *Phys. Lett. B*, vol. 121, p. 429, 1983.
- [130] M. Jonker *et al.*, “Experimental Study of Prompt Neutrino Production in 400-GeV Proton - Nucleus Collisions,” *Phys. Lett. B*, vol. 96, p. 435, 1980.
- [131] J. Blumlein *et al.*, “Limits on neutral light scalar and pseudoscalar particles in a proton beam dump experiment,” *Z. Phys. C*, vol. 51, pp. 341–350, 1991.
- [132] H. Nanjo, “Dark Sector search at KOTO experiment,” 2021. Contribution to “3rd meeting on Searches for Hidden Sectors at Kaon and Hyperon Factories”.

- [133] E. Cortina Gil *et al.*, “Search for dark photon decays to $\mu^+\mu^-$ at NA62,” *JHEP*, vol. 09, p. 035, 2023, 2303.08666.
- [134] E. Cortina Gil *et al.*, “The Beam and detector of the NA62 experiment at CERN,” *JINST*, vol. 12, no. 05, p. P05025, 2017, 1703.08501.
- [135] M. Anelli *et al.*, “A facility to Search for Hidden Particles (SHiP) at the CERN SPS,” 4 2015, 1504.04956.
- [136] C. Ahdida *et al.*, “SHiP Experiment - Progress Report,” tech. rep., CERN, Geneva, Jan 2019.
- [137] O. Aberle *et al.*, “BDF/SHiP at the ECN3 high-intensity beam facility,” tech. rep., CERN, Geneva, 2022.
- [138] C. Ahdida *et al.*, “Sensitivity of the SHiP experiment to dark photons decaying to a pair of charged particles,” *Eur. Phys. J. C*, vol. 81, no. 5, p. 451, 2021, 2011.05115.
- [139] C. A. Aidala *et al.*, “The SeaQuest Spectrometer at Fermilab,” *Nucl. Instrum. Meth. A*, vol. 930, pp. 49–63, 2019, 1706.09990.
- [140] B. Batell, J. A. Evans, S. Gori, and M. Rai, “Dark Scalars and Heavy Neutral Leptons at DarkQuest,” *JHEP*, vol. 05, p. 049, 2021, 2008.08108.
- [141] A. Berlin, S. Gori, P. Schuster, and N. Toro, “Dark Sectors at the Fermilab SeaQuest Experiment,” *Phys. Rev. D*, vol. 98, no. 3, p. 035011, 2018, 1804.00661.
- [142] N. Blinov, E. Kowalczyk, and M. Wynne, “Axion-like particle searches at DarkQuest,” *JHEP*, vol. 02, p. 036, 2022, 2112.09814.
- [143] K. J. Kelly, S. Kumar, and Z. Liu, “Heavy axion opportunities at the DUNE near detector,” *Phys. Rev. D*, vol. 103, no. 9, p. 095002, 2021, 2011.05995.
- [144] V. Hewes *et al.*, “Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report,” *Instruments*, vol. 5, no. 4, p. 31, 2021, 2103.13910.
- [145] R. Acciarri *et al.*, “Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 1: The LBNF and DUNE Projects,” 1 2016, 1601.05471.

- [146] B. Abi *et al.*, “The Single-Phase ProtoDUNE Technical Design Report,” 6 2017, 1706.07081.
- [147] W. Baldini *et al.*, “SHADOWS (Search for Hidden And Dark Objects With the SPS),” 10 2021, 2110.08025.
- [148] M. Alviggi *et al.*, “SHADOWS Letter of Intent,” tech. rep., CERN, Geneva, 2022.
- [149] C. Rella, B. Döbrich, and T.-T. Yu, “Searching for muonphilic dark sectors with proton beams,” *Phys. Rev. D*, vol. 106, no. 3, p. 035023, 2022, 2205.09870.
- [150] J. Blumlein *et al.*, “Limits on the mass of light (pseudo)scalar particles from Bethe-Heitler e^+e^- and $\mu^+\mu^-$ pair production in a proton - iron beam dump experiment,” *Int. J. Mod. Phys. A*, vol. 7, pp. 3835–3850, 1992.
- [151] E. Goudzovski *et al.*, “New physics searches at kaon and hyperon factories,” *Rept. Prog. Phys.*, vol. 86, no. 1, p. 016201, 2023, 2201.07805.
- [152] R. Aaij *et al.*, “Search for hidden-sector bosons in $B^0 \rightarrow K^{*0}\mu^+\mu^-$ decays,” *Phys. Rev. Lett.*, vol. 115, no. 16, p. 161802, 2015, 1508.04094.
- [153] R. Aaij *et al.*, “Search for long-lived scalar particles in $B^+ \rightarrow K^+\chi(\mu^+\mu^-)$ decays,” *Phys. Rev. D*, vol. 95, no. 7, p. 071101, 2017, 1612.07818.
- [154] E. Cortina Gil *et al.*, “Measurement of the very rare $K^+ \rightarrow \pi^+\nu\bar{\nu}$ decay,” *JHEP*, vol. 06, p. 093, 2021, 2103.15389.
- [155] M. Hostert and M. Pospelov, “Novel multilepton signatures of dark sectors in light meson decays,” *Phys. Rev. D*, vol. 105, no. 1, p. 015017, 2022, 2012.02142.
- [156] E. Cortina Gil *et al.*, “Search for K^+ decays into the $\pi^+e^+e^-e^+e^-$ final state,” 7 2023, 2307.04579.
- [157] A. J. Krasznahorkay *et al.*, “Observation of Anomalous Internal Pair Creation in Be8 : A Possible Indication of a Light, Neutral Boson,” *Phys. Rev. Lett.*, vol. 116, no. 4, p. 042501, 2016, 1504.01527.

- [158] A. J. Krasznahorkay, M. Csatlós, L. Csige, J. Gulyás, A. Krasznahorkay, B. M. Nyakó, I. Rajta, J. Timár, I. Vajda, and N. J. Sas, “New anomaly observed in He4 supports the existence of the hypothetical X17 particle,” *Phys. Rev. C*, vol. 104, no. 4, p. 044003, 2021, 2104.10075.
- [159] A. J. Krasznahorkay *et al.*, “New anomaly observed in C12 supports the existence and the vector character of the hypothetical X17 boson,” *Phys. Rev. C*, vol. 106, no. 6, p. L061601, 2022, 2209.10795.
- [160] E. Cortina Gil *et al.*, “Search for K^+ decays to a muon and invisible particles,” *Phys. Lett. B*, vol. 816, p. 136259, 2021, 2101.12304.
- [161] E. Cortina Gil *et al.*, “Search for heavy neutral lepton production in K^+ decays to positrons,” *Phys. Lett. B*, vol. 807, p. 135599, 2020, 2005.09575.
- [162] S. Gori, G. Perez, and K. Tobioka, “KOTO vs. NA62 Dark Scalar Searches,” *JHEP*, vol. 08, p. 110, 2020, 2005.05170.
- [163] T. Kitahara, T. Okui, G. Perez, Y. Soreq, and K. Tobioka, “New physics implications of recent search for $K_L \rightarrow \pi^0 \nu \bar{\nu}$ at KOTO,” *Phys. Rev. Lett.*, vol. 124, no. 7, p. 071801, 2020, 1909.11111.
- [164] J. K. Ahn *et al.*, “Search for the $K_L \rightarrow \pi^0 \nu \bar{\nu}$ and $K_L \rightarrow \pi^0 X^0$ decays at the J-PARC KOTO experiment,” *Phys. Rev. Lett.*, vol. 122, no. 2, p. 021802, 2019, 1810.09655.
- [165] K. Aoki *et al.*, “Extension of the J-PARC Hadron Experimental Facility: Third White Paper,” 10 2021, 2110.04462.
- [166] H. Nanjo, “Dark Sector search at KOTO experiment,” 2021. Contribution to “3rd meeting on Searches for Hidden Sectors at Kaon and Hyperon Factories”.
- [167] Y.-C. Tung, “Status of $K_L \rightarrow \pi^0 \nu \bar{\nu}$ Study at J-PARC KOTO Experiment,” 2022. Contribution to “41st International Conference on High Energy Physics, Bologna, Italy”.
- [168] “Searches for new physics with high-intensity kaon beams,” in *Snowmass 2021*, 4 2022, 2204.13394.

- [169] J. K. Ahn *et al.*, “Study of the $K_L \rightarrow \pi^0 \nu \bar{\nu}$ Decay at the J-PARC KOTO Experiment,” *Phys. Rev. Lett.*, vol. 126, no. 12, p. 121801, 2021, 2012.07571.
- [170] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, “Consistent Treatment of Axions in the Weak Chiral Lagrangian,” *Phys. Rev. Lett.*, vol. 127, no. 8, p. 081803, 2021, 2102.13112.
- [171] T. Masuda *et al.*, “Long-lived neutral-kaon flux measurement for the KOTO experiment,” *PTEP*, vol. 2016, no. 1, p. 013C03, 2016, 1509.03386.
- [172] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu, L. W. Mo, T. A. Nunamaker, and P. Rassmann, “Search for Neutral Metastable Penetrating Particles Produced in the SLAC Beam Dump,” *Phys. Rev. D*, vol. 38, p. 3375, 1988.
- [173] E. M. Riordan *et al.*, “A Search for Short Lived Axions in an Electron Beam Dump Experiment,” *Phys. Rev. Lett.*, vol. 59, p. 755, 1987.
- [174] M. J. Dolan, T. Ferber, C. Hearty, F. Kahlhoefer, and K. Schmidt-Hoberg, “Revised constraints and Belle II sensitivity for visible and invisible axion-like particles,” *JHEP*, vol. 12, p. 094, 2017, 1709.00009. [Erratum: *JHEP* 03, 190 (2021)].
- [175] N. Bar, K. Blum, and G. D’Amico, “Is there a supernova bound on axions?,” *Phys. Rev. D*, vol. 101, no. 12, p. 123025, 2020, 1907.05020.
- [176] A. V. Artamonov *et al.*, “Study of the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ in the momentum region $140 < P_\pi < 199$ MeV/c,” *Phys. Rev. D*, vol. 79, p. 092004, 2009, 0903.0030.
- [177] A. V. Artamonov *et al.*, “Search for the decay $K^+ \rightarrow \pi^+ \gamma \gamma$ in the π^+ momentum region $P > 213$ MeV/c,” *Phys. Lett. B*, vol. 623, pp. 192–199, 2005, hep-ex/0505069.
- [178] E. Cortina Gil *et al.*, “Search for π^0 decays to invisible particles,” *JHEP*, vol. 02, p. 201, 2021, 2010.07644.
- [179] J. P. Lees *et al.*, “Search for an Axionlike Particle in B Meson Decays,” *Phys. Rev. Lett.*, vol. 128, no. 13, p. 131802, 2022, 2111.01800.

- [180] E. Bertholet, S. Chakraborty, V. Loladze, T. Okui, A. Soffer, and K. Tobioka, “Heavy QCD axion at Belle II: Displaced and prompt signals,” *Phys. Rev. D*, vol. 105, no. 7, p. L071701, 2022, 2108.10331.
- [181] E. Lopienska, “The CERN accelerator complex,” Feb. 2022. Available at: <https://home.cern/science/accelerators/accelerator-complex>, accessed on 1/9/2023.
- [182] J. Jerhot, “Search for Axion-Like Particles at the NA62 experiment,” Master’s thesis, Charles U., 2019.
- [183] L. Gagnon, D. Banerjee, J. Bernhard, M. Brugger, N. Charitonidis, G. L. D’Alessandro, N. Doble, M. Van Dijk, A. Gerbershagen, E. Montbarbon, and M. Rosenthal, “Report from the Conventional Beams Working Group to the Physics Beyond Collider Study and to the European Strategy for Particle Physics,” tech. rep., CERN, Geneva, 2022.
- [184] M. Rosenthal *et al.*, “Single-muon rate reduction for beam dump operation of the K12 beam line at CERN,” *Int. J. Mod. Phys. A*, vol. 34, no. 36, p. 1942026, 2019.
- [185] C. NA62, “2023 NA62 Status Report to the CERN SPSC,” tech. rep., CERN, Geneva, 2023.
- [186] F. Hahn, F. Ambrosino, A. Ceccucci, H. Danielsson, N. Doble, F. Fantechi, A. Kluge, C. Lazzeroni, M. Lenti, G. Ruggiero, M. Sozzi, P. Valente, and R. Wanke, “NA62: Technical Design Document,” tech. rep., CERN, Geneva, 2010.
- [187] T. Blažek, V. Černý, R. Lietava¹, M. Koval’, and M. Krivda, “NA62 LTU software User Guide,” *Internal note: NA62-12-11*, 2012.
- [188] D. Soldi and S. Chiozzi, “Level Zero Trigger Processor for the NA62 experiment,” *JINST*, vol. 13, no. 05, p. P05004, 2018, 1802.06548.
- [189] Analog Devices, Inc, *Octal, 14-Bit, 40/65 MSPS Serial LVDS 1.8 V Analog-to-Digital Converter (Datasheet)*. Rev. B.
- [190] Altera Corporation, *Stratix IV FPGA (Datasheet)*, 2011.
- [191] A. Ceccucci, R. Fantechi, P. Farthouat, G. Lamanna, J. Rouet, V. Ryjov, and S. Venditti, “The New Readout System of the NA62 LKr Calorimeter,” *IEEE Trans. Nucl. Sci.*, vol. 62, no. 5, pp. 2134–2140, 2015.

- [192] D. Soldi, *The FPGA-based first level trigger for the NA62 Experiment at CERN SPS*. PhD thesis, Turin U., 2017.
- [193] M. Gonzalez-Berges, “The Joint CONTROLS Project Framework,” *eConf C*, vol. 0303241, p. THGT006, 2003. On behalf of the Framework Team.
- [194] H. Milcent, E. Blanco, F. Bernard, and P. Gayet, “UNICOS: AN OPEN FRAMEWORK,” 2009.
- [195] C. Gaspar and M. Dönszelmann, “DIM - A distributed information management system for the DELPHI experiment at CERN,” 1994.
- [196] N. Lurkin, *Neutral Pion Transition Form Factor Measurement and Run Control at the NA62 experiment*. PhD thesis, Birmingham U., 2017. Presented 2017.
- [197] L. Gatignon and J. Bernhard, “BSI Test,” 2021. NA62 run meeting 15/11/2021.
- [198] G. Unal, “Reconstruction program for the LKr,” *Internal note: NA48-98-01*, 1998.
- [199] V. Fanti *et al.*, “Performance of an electromagnetic liquid krypton calorimeter,” *Nucl. Instrum. Meth. A*, vol. 344, pp. 507–520, 1994.
- [200] C. Cesarotti, B. Döbrich, and T. Spadaro, “Monte Carlo Implementation for Exotic Particle Searches at NA62,” *Internal note: NA62-17-09*, 2017.
- [201] S. Ghinescu, B. Döbrich, E. Minucci, and T. Spadaro, “A biased MC for muon production for beam-dump experiments,” *Eur. Phys. J. C*, vol. 81, no. 8, p. 767, 2021, 2106.01932.
- [202] V. Niess, A. Barnoud, C. Cârloganu, and E. Le Méneüdeu, “Backward Monte-Carlo applied to muon transport,” *Comput. Phys. Commun.*, vol. 229, pp. 54–67, 2018, 1705.05636.
- [203] M. Raggi and T. Spadaro, “Search for visible decays of long-lived dark photons at the NA62 experiment: a feasibility study,” *Internal note: NA62-16-01*, 2016.
- [204] A. Kleimenova, “Exotic production modes for the dark photon: validation and perspectives,” 2022. MC validation meeting 16/2/22.

- [205] Y.-D. Tsai, P. deNiverville, and M. X. Liu, “Dark Photon and Muon $g - 2$ Inspired Inelastic Dark Matter Models at the High-Energy Intensity Frontier,” *Phys. Rev. Lett.*, vol. 126, no. 18, p. 181801, 2021, 1908.07525.
- [206] E. Minucci and P. Petrov, “Method for vertex fit in NA62,” *Internal note: NA62-16-05*, 2016.
- [207] R. J. Barlow, “Extended maximum likelihood,” *Nucl. Instrum. Meth. A*, vol. 297, pp. 496–506, 1990.
- [208] Z. Kučerová and M. Koval, “Spectrometer reconstruction efficiency for single track events,” *Internal note: NA62-20-05*, 2020.
- [209] v. Bičian, “New Measurement of $K^+ \rightarrow \pi^+ \mu^+ \mu^- (\gamma)$ Decay with NA62 Run 1 Data,” *Internal note: NA62-22-02*, 2022.
- [210] R. Aliberti, *Particle Identification with Calorimeters for the Measurement of the Rare Decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ at NA62*. PhD thesis, Mainz U., 2018. Presented 21 Jan 2019.
- [211] D. Bryman, W. Fedorko, B. Velghe, and M. Yu, “Particle identification in NA62 using a Light Gradient Boosting Machine and Convolutional Neural Networks,” *Internal note: NA62-21-02*, 2021.
- [212] L. Iacobuzio, “Full Monte Carlo simulation for Heavy Neutral Lepton decay searches at NA62,” *Internal note: NA62-19-11*, 2019.
- [213] Hamamatsu Photonics K.K., *Compact Metal Package PMT (Datasheet)*, 2022.
- [214] J. C. Swallow, “Particle Identification Studies for the $K \rightarrow \pi \mu e$ Analysis,” *Internal note: NA62-19-03*, 2019.

Appendices

Appendix A

Treatment of kinematics for the mixing production

The mixing production serves as an approximation to the various processes in which an ALP can be produced. The total yield of produced ALPs can be derived from the yield of the neutral pseudoscalar mesons and the respective mixing angles with the ALP. This approach, however, does not say anything about the change in the kinematics when the ALP is produced instead of the meson. The general minimal adjustment, which can be applied irrespective of the specific process that leads to the ALP production, is the kinematics change when the original pseudoscalar's mass is changed. Since the center-of-mass frame of the processes that lead to the meson production is unknown, we choose a common frame for the mass adjustment to be the p - p collision cm frame. Note that this choice has a certain level of arbitrariness, with the obvious alternative being the laboratory frame. However, we expect that the p - p cm frame leads to a more accurate description for highly-energetic ALPs, where only a smaller portion of the available energy is spent on the ALP production, and the majority of the energy is carried in the ALP boost. These highly boosted ALPs are particularly important for the sensitivity of beam-dump experiment searches.

The cm momentum of a particle with mass m is then simply found as

$$p_{cm}(p_{lab}, m) = \gamma_{beam} p_{lab} - \gamma_{beam} \beta_{beam} \sqrt{p_{lab}^2 + m^2} \quad (\text{A.1})$$

and analogous relations can be used for the E_{cm} and, by inversion, for p_{lab} and E_{lab} . By adjusting the particle mass, one necessarily violates the momentum or energy conservation (or both) in the original process. In this simulation, the momentum of the ALP is fixed in the p - p cm frame to be the same as the momentum of the original meson. The energy and angle of the ALP a in the laboratory frame then translate to the energy and angle of the original meson

P in the laboratory frame as

$$\begin{aligned} E_P|_{p_{\text{fixed}}} &= \sqrt{p_{\text{lab}}(p_{\text{cm}}(p_a, m_a), m_P) + m_P^2}, \\ \theta_P|_{p_{\text{fixed}}} &= \arcsin \left[\sin(\theta_a) \frac{p_a}{p_{\text{lab}}(p_{\text{cm}}(p_a, m_a), m_P)} \right]. \end{aligned} \quad (\text{A.2})$$

The plotted distributions for the case of $P = \eta'$ can be found in the upper row of figure A.1.




Figure A.1: Dependence of the energy (left) and the angle of the ALP (right) on the ALP mass in the laboratory frame for the case of the fixed momentum (upper panel) or fixed energy (lower panel) in the p - p cm frame when the original particle was the η' meson, produced by a 400 GeV proton beam. The angular distribution is evaluated for the case $E_a = 100$ GeV.

If one chooses to fix the ALP energy to the meson energy in the p - p cm frame, one obtains in complete analogy the following relations:

$$\begin{aligned} E_P|_{E_{\text{fixed}}} &= E_{\text{lab}}(E_{\text{cm}}(p_a, m_a), m_P) \\ \theta_P|_{E_{\text{fixed}}} &= \arcsin \left[\sin(\theta_a) \frac{p_a}{\sqrt{E_{\text{lab}}(E_{\text{cm}}(E_a, m_a), m_P)^2 - m_P^2}} \right]. \end{aligned} \quad (\text{A.3})$$

The resulting distributions are plotted in the bottom row of figure A.1. By comparing the two distributions, we can see that even in the borderline case of choosing the p - p collision as a universal cm frame, we can get a significantly different E_a - θ_a distribution solely by choosing a different method of fixing the kinematics in the cm frame. This choice, however, impacts mostly ALPs with a small boost (large masses and low energies), which are of limited relevance for beam-dump experiments.

The impact of different choices on the yield in the acceptance of the NA62 experiment will be discussed in the following section.

Mixing production systematic uncertainty

Fixing the exotic particle momentum or the exotic particle energy in the cm frame presents two borderline cases. In NA62 data analysis, the average yield from these two cases is used as the central value. The assigned systematic uncertainty corresponds to the difference between the central value and the yield obtained while fixing the energy or momentum respectively in the cm frame. Figure A.2 shows the relative correction to yield obtained while fixing the energy in the cm frame. Besides studying the correction and uncertainty of the ALP mixing with neutral pseudoscalars, the correction to the dark photon yield is also calculated, accounting for the contributions of ρ , ω , and ϕ mesons to the yield.




Figure A.2: Correction to the DP and ALP yields at NA62 from mixing production and the related systematic uncertainty. Top left: dark photon (combination of all vector meson contributions), Top right: ALP mixing with π^0 , Bottom left: ALP mixing with η , Bottom right: ALP mixing with η' .

Appendix **B**

Initial setup of VetoCounter PMTs

Hamamatsu series R9880U-110 are used as VetoCounter PMTs [213]. The PMT high voltages are supplied from two 32-channel CAEN HV boards. The initial course voltage setup of the PMTs was chosen so that each PMT has an equal gain in order to get a uniform behavior for the whole detector. The gain G is related to the voltage V as

$$\frac{dG}{G} = n \frac{dV}{V}. \quad (\text{B.1})$$

The slope n is obtained from fitting the Hamamatsu-measured characteristics as $n \simeq 8.7$. The gain G is obtained from the anode S_a and cathode S_c sensitivities measured by Hamamatsu for each PMT at $V = 1\text{kV}$, and the initial voltage is then obtained using (B.1). See tables B.1, B.2, B.3 for the values of V and G obtained when using PMT BAG0068 as a reference.

The final *fine* settings of the VetoCounter voltages were obtained during a muon run after VetoCounter installation by comparing the gain of Jura and Saleve channels for each tile and equalizing those.

Channel	PMT	S_c [μ A/lm]	S_a [A/lm]	$G \times 10^6$	V for $G = 10^6$	V for $S_a = 799$ A/lm
101J	BAG0068	132	799	6.05	813.1	1000.0
101S	BAG0069	134	628	4.69	837.3	1028.1
102J	BAG0070	134	616	4.60	839.2	1030.3
102S	BAG0071	132	489	3.70	860.3	1058.1
103J	BHD2645	96	470	4.90	833.1	1062.9
103S	BHD2649	105	481	4.58	839.5	1060.1
104J	BHD2655	106	597	5.63	819.8	1034.1
104S	BHD2656	104	375	3.61	862.9	1090.8
105J	BHD2657	101	414	4.10	850.3	1078.5
105S	BHD2658	102	480	4.71	836.9	1060.3
106J	BHD2659	102	438	4.29	845.8	1071.5
106S	BHD2660	104	469	4.51	841.0	1063.2
107J	BHD2661	104	429	4.13	849.7	1074.1
107S	BHD2663	103	329	3.19	875.0	1107.4
108J	BHD2664	101	418	4.14	849.4	1077.3
108S	BHD2702	102	501	4.91	832.8	1055.1
109J	BHD2705	103	481	4.67	837.7	1060.1
109S	BHD2707	101	439	4.35	844.6	1071.3
110J	BHD2709	103	244	2.37	905.6	1146.1
110S	BHD2710	103	317	3.08	878.8	1112.1
111J	BHD2711	102	408	4.00	852.7	1080.3
111S	BHD2715	103	432	4.19	848.1	1073.2

Table B.1: List of station F1 VetoCounter PMTs and their initial settings.

Channel	PMT	S_c [$\mu\text{A}/\text{lm}$]	S_a [A/lm]	$G \times 10^6$	\bar{V} for $G = 10^6$	\bar{V} for $S_a = 799 \text{ A}/\text{lm}$
201J	BHD2716	102	393	3.85	856.4	1085.0
201S	BHD2718	101	480	4.75	836.0	1060.3
202J	BHD2719	100	533	5.33	825.0	1047.6
202S	BHD2720	97	460	4.74	836.2	1065.5
203J	BHD2736	92	349	3.79	857.9	1099.9
203S	BPE5762	110	381	3.46	866.9	1088.8
204J	BPE5763	110	301	2.74	890.7	1118.8
204S	BPE5764	111	286	2.58	896.9	1125.3
205J	BPE5767	113	285	2.52	899.1	1125.8
205S	BPE5768	112	294	2.63	895.0	1121.8
206J	BPE5771	116	371	3.20	874.9	1092.2
206S	BPE5774	112	306	2.73	890.9	1116.6
207J	BPE5775	110	361	3.28	872.3	1095.6
207S	BPE5781	111	349	3.14	876.6	1099.9
208J	BPE5785	114	298	2.61	895.4	1120.0
208S	BPE5786	110	342	3.11	877.8	1102.4
209J	BPE5798	106	580	5.47	822.5	1037.5
209S	BPE5799	106	554	5.23	826.9	1043.0
210J	BPE5800	110	545	4.95	832.0	1045.0
210S	BPE5801	108	706	6.54	805.9	1014.3
211J	BPE5803	113	478	4.23	847.2	1060.8
211S	BPE5805	113	610	5.40	823.8	1031.5

Table B.2: List of station F2 VetoCounter PMTs and their initial settings.

Channel	PMT	S_c [$\mu\text{A}/\text{lm}$]	S_a [A/lm]	$G \times 10^6$	V for $G = 10^6$	V for $S_a = 799 \text{ A}/\text{lm}$
301J	BPE5807	113	697	6.17	811.3	1015.8
301S	BPE5810	115	792	6.89	801.1	1001.0
302J	BPE5812	111	686	6.18	811.1	1017.7
302S	BPE5814	107	547	5.11	829.0	1044.5
303J	BHG0488	121	588	4.86	833.8	1035.9
303S	BHG0489	120	720	6.00	813.9	1012.0
304J	BHG0490	123	504	4.10	850.3	1054.4
304S	BHG0491	123	473	3.85	856.6	1062.1
305J	BHG0492	123	354	2.88	885.6	1098.1
305S	BHG0493	124	421	3.40	868.9	1076.4
306J	BHG0494	123	396	3.22	874.2	1084.0
306S	BHG0495	122	410	3.36	869.9	1079.7
307J	BHG0496	105	392	3.73	859.5	1085.3
307S	BHG0498	111	461	4.15	849.0	1065.3
308J	BHG0499	113	417	3.69	860.6	1077.6
308S	BHG0512	110	347	3.15	876.3	1100.6
309J	BHG0524	115	337	2.93	883.8	1104.3
309S	BHG0526	117	300	2.56	897.4	1119.2
310J	BHG0555	115	241	2.10	918.5	1147.7
310S	BHG0561	113	320	2.83	887.2	1110.9
311J	BHG0568	107	713	6.66	804.1	1013.2
311S	BHG0570	109	741	6.80	802.3	1008.7

Table B.3: List of station B VetoCounter PMTs and their initial settings.

Kaon data analyzers for efficiency studies

Four analyzers for the following kaon decays are used $K^+ \rightarrow \mu^+ \nu_\mu$ ($K_{\mu 2}$), $K^+ \rightarrow \pi^0 e^+ \nu_e$ ($K_{e 3}$), $K^+ \rightarrow \pi^+ \pi^0$ ($K_{2\pi}$) and $K^+ \rightarrow \pi^+ \pi^0, \pi^0 \rightarrow e^+ e^- \gamma$ ($K_{2\pi D}$). Each analyzer has a specific kinematical selection for given decay while keeping the additional cuts on detector acceptances and event reconstruction the same as those of the beam-dump analyses as possible.

The data samples used for each analysis are chosen to avoid any bias in the PID, therefore, using only triggers not employing any cuts on MUV3 or LKr energy. In particular, for single-track analysis, the largely downscaled ($D = 10^4$) control sample is used, and for three-track analysis, the three-track filtered sample is used. For the same reason, all PID steps in the analyzers are omitted, and the same PID criteria as in beam-dump analyses are applied to the filtered events to estimate the PID efficiency.

$K_{\mu 2}$ selection

The $K_{\mu 2}$ analyzer is derived by modifying the `Kmu2Selection` tool from `NA62Analysis`, omitting any PID criteria. The GTK K^+ candidates are matched to downstream π^+ candidates using `MatchingRG` tool from `NA62-Analysis`. The track time association uses the time of the `KTAG` K^+ candidate, and the reconstructed vertex must have $CDA_{VTX} < 4$ cm. The final kinematical selection is done by requiring $(p_K - p_\mu)^2 < 0.01 \text{ GeV}^2$ with K^+ selected by the `KTAG` and GTK and assuming a μ^+ mass hypothesis for the downstream track.

$K_{e 3}$ selection

The $K_{e 3}$ analyzer also uses the GTK for the event reconstruction with upstream and downstream tracks matched using the `SpectrometerGigaTracker-`

Matching tool with the reference time of the KTAG selected K^+ candidate and a time window ± 5 ns for the matching with the downstream track. $CDA_{VTX} < 7$ mm is required for the track matching. The kinematic cuts follow the setup from Ref. [214], requiring

- $|(p_K - p_e)^2 - m_{\pi^0}^2| > 5\sigma_{m_{K_{2\pi}}}^2$;
- $(p_K - p_e)^2 - m_{\pi^0}^2 > 5\sigma_{m_{K_{2\pi}}}^2$;
- $(p_K - p_\mu)^2 > 5\sigma_{m_{K_{\mu 2}}}^2$;
- $(p_K - p_\pi)^2 < 0.065 \text{ GeV}^2$,

where $p_{e,\pi,\mu}$ are calculated assuming the respective mass hypothesis, $\sigma_{m_{K_{2\pi}}}^2 = 2890 \text{ MeV}^2$ and $\sigma_{m_{K_{\mu 2}}}^2 = 3830 \text{ MeV}^2$. Next, the π^0 is reconstructed from the reconstructed vertex position and deposited energy of the two photons in LKr. The final selection is by combining the reconstructed p_e and p_{π^0} and requiring large transverse momentum and small m_ν

- $(p_{\pi^0} + p_e - p_K)^T > 30 \text{ MeV}$;
- $|p_{\pi^0} + p_e - p_K| < 20 \text{ MeV}$.

$K_{2\pi}$ and $K_{2\pi D}$ selection

The $K_{2\pi}$ and $K_{2\pi D}$ analyzers use simply the `K2piSelection` and `K2piD-Selection` tools of `NA62Analysis` while using the option for event reconstruction with `GTK`. The additional vetoes on `LAV` and `SAV` matching are applied in analogy with the beam-dump analyses.