

SCUOLA
NORMALE
SUPERIORE

Classe di Scienze

Corso di perfezionamento in Fisica

XXXV ciclo

Charm CP violation measurements with K_S^0 final states at LHCb

Settore Scientifico Disciplinare FIS/01

Candidate:

Lorenzo Pica

Supervisor:

Prof. Giovanni Punzi

Co-Supervisor:

Prof. Michael J. Morello

PhD Coordinator:

Prof. Vittorio Giovannetti

Academic year 2024–2025

Ai miei genitori
To my parents

Abstract

The LHCb collaboration achieved the world’s most precise measurement of the $\mathcal{A}^{CP}(D^0 \rightarrow K_S^0 K_S^0)$ parameter using data collected during Run 2 (2015–2018), reaching a precision of 1.3% and observing compatibility with the absence of CP violation at a significance of 2.4σ . Following a major Upgrade I, LHCb resumed operations in 2022 with a completely overhauled tracking system and a novel triggerless readout architecture, capable of reconstructing events at the full LHC bunch crossing rate of 30 MHz.

This thesis presents the design and optimization of dedicated K_S^0 triggers, aimed at enhancing the LHCb online selection efficiency of $D^0 \rightarrow K_S^0 K_S^0$ decays. The new triggers have been commissioned with early Run 3 data; performed tests confirm the feasibility of this new trigger strategy, never implemented before at a hadron collider. With the availability of larger samples, the impact on signal efficiency has been evaluated, leading to a fourfold enhancement for the target decay $D^0 \rightarrow K_S^0 K_S^0$. These performances motivated an early $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement performed on 6 fb^{-1} of data collected by LHCb in 2024, that already surpassed the precision of the previous measurement. This measurement is a remarkable result, being the most precise determination of this parameter to date, and demonstrating the effectiveness of the new triggers in enhancing the physics reach of the LHCb experiment.

Contents

Introduction	1
1 CP violation in charm decays	5
1.1 Standard Model	6
1.1.1 CPV within SM formalism	6
1.1.2 Neutral mesons phenomenology and mixing	9
1.1.3 CP violation in the decay	10
1.1.4 CP violation in mixing	10
1.1.5 CP violation in the interference between decay and mixing	11
1.2 CP violation in charmed hadron decays	11
1.3 First CP violation observation in charm decays	12
1.3.1 Result interpretation	13
1.4 The role of the $D^0 \rightarrow K_S^0 K_S^0$ mode	14
1.5 Experimental difficulties in measuring $\mathcal{A}^{CP}(K_S^0 K_S^0)$	17
2 Importance of long-lived particles in general	19
2.1 Standard Model LLP	20
2.2 Beyond Standard Model LLP	21
3 LHCb and long-lived particles reconstruction	25
3.1 The Large Hadron Collider	25
3.1.1 LLP decays triggering	26
3.2 LHCb configuration in Run 2	26
3.2.1 LHCb tracking system	27
3.2.2 LHCb particle identification system	35
3.2.3 Run 2 trigger configuration	37
3.3 The LHCb Upgrade I	40
3.3.1 The tracking system upgrades	41
3.3.2 Data acquisition and trigger system for LHCb upgrade	43
4 Measuring $D^0 \rightarrow K_S^0 K_S^0$ at LHCb	47
4.1 Introduction	47
4.2 Variables definition	48
4.3 Trigger selections	50
4.3.1 L0	50
4.3.2 HLT1	50
4.3.3 HLT2	52

4.4	Background sources	53
4.4.1	Peaking backgrounds	54
4.4.2	Combinatorial background	55
4.4.3	Secondary decays	56
4.4.4	Sample splitting	56
4.5	Calibration sample offline preparation	57
4.6	Procedure for cancellation of nuisance asymmetries	57
4.7	Results and systematic effects	59
4.8	Considerations about the Run 2 measurement	61
5	Improving $D^0 \rightarrow K_S^0 K_S^0$ acceptance in Run 3	65
5.1	Triggering on K_S^0 's in HLT1	66
5.1.1	Lessons learned from Run 2	66
5.1.2	The vision for Run 3	67
5.2	Development of a K_S^0 line for HLT1	67
5.2.1	Cuts excluded from optimization	69
5.2.2	Cuts determined by numerical optimization	71
5.2.3	Observables purposely excluded from our trigger line	74
5.2.4	Numerical optimization of thresholds	75
5.2.5	Selection criteria for the TwoTrackKs HLT1 line	76
5.3	Extension to K_S^0 pairs selection	77
5.3.1	Selection criteria for the TwoKs HLT1 line	78
5.4	K_S^0 lines impact on HLT1 throughput	78
5.5	Lines rate cost	79
5.6	Impact on $D^0 \rightarrow K_S^0 K_S^0$ collection efficiency	81
5.7	Run 3 HLT2 $D^0 \rightarrow K_S^0 K_S^0$ lines	82
5.7.1	Rate reduction of $D^0 \rightarrow K_S^0 K_S^0$ HLT2 lines	82
5.7.2	Tuning of HLT2 lines	92
6	Spin-offs of K_S^0 triggers	95
6.1	Impact on efficiency for K_S^0 final state decays	95
6.1.1	$B^0/B_s^0 \rightarrow K_S^0 K_S^0$	95
6.1.2	$D^0 \rightarrow K_S^0 \pi^- \pi^+$	95
6.1.3	$D^+ \rightarrow K_S^0 \pi^+$ and $D_s^+ \rightarrow K_S^0 K^+$	97
6.2	Studies on a Λ -dedicated HLT1 line	98
6.2.1	TwoTrackLambda selections	98
6.2.2	HLT1 Λ line rate cost	99
6.2.3	Impact of TwoTrackLambda line on selected measurements	101
6.2.4	Conclusions Λ -dedicated HLT1 line	102
6.3	Extension to downstream tracks	102
6.3.1	Benefits of trigger-level downstream tracking	103
6.3.2	Efficiency gain on benchmark decay modes and rate cost	105
7	New triggers commissioning with Run 3 data	109
7.1	Tunability of selections	109
7.2	Commissioning with early Run 3 data	112
7.2.1	Commissioning the TwoTrackKs line	112

7.2.2	Commissioning the TwoKs line	116
7.2.3	Rate measurements	118
7.2.4	First attempt at evaluating efficiencies	119
7.3	K_S^0 triggers performance assessment on 2022 data	120
7.4	Analysis of 2023 HLT2 $D^0 \rightarrow K_S^0 K_S^0$ data	124
7.5	Analysis of 2024 HLT2 $D^0 \rightarrow K_S^0 K_S^0$ data	127
8	Measurement of \mathcal{A}^{CP} ($D^0 \rightarrow K_S^0 K_S^0$) in Run 3	133
8.1	Analysis strategy	134
8.2	Data sample	135
8.3	Trigger selection	136
8.3.1	HLT1	136
8.3.2	HLT2	136
8.4	Offline selections	138
8.4.1	Rejection of peaking physical backgrounds	138
8.4.2	D^0 candidate invariant mass selection	141
8.4.3	Secondary decays	142
8.4.4	Combinatorial background	144
8.4.5	Equalization of kinematic region with calibration channel	146
8.4.6	Removal of multiple candidates	146
8.5	Results on charge-integrated sample	147
8.6	The calibration sample	147
8.7	Evaluation of event weights	149
8.7.1	Correction of the intrinsic asymmetry of the calibration sample	152
8.8	Fit procedure	154
8.9	Results	156
8.10	Cross-checks and systematics	158
8.10.1	Check of consistency of \mathcal{A}^{CP} between data-taking periods	158
8.10.2	Closure test of cancellation of spurious asymmetries	159
8.10.3	K_S^0 decays and regeneration effects	159
8.10.4	Systematics on the choice of the k parameter	160
8.10.5	\mathcal{A}^{CP} in $D^0 \rightarrow K_S^0 \pi^+ \pi^-$	161
8.10.6	Fit model uncertainty	161
8.10.7	Finite statistics of the calibration sample and other effects	161
8.11	Conclusions	162
9	Concluding remarks and future prospects	163
A	Complete 2024 fit results	167
A.1	Unweighted fit results	167
A.2	Weighted fit results	176
	References	176
	Acknowledgements	198
	Post Scriptum	200

Introduction

The first evidence of violation of CP in decays of *charm* hadrons has only recently been found (2019), by the LHCb collaboration in $D^0 \rightarrow h^+ h^-$ ($h = K, \pi$) decay modes [1]. While this result amounts to a milestone in particle physics, the picture of CP violation in the charm sector is still far from being clear. Large non-perturbative theoretical uncertainties prevent an unambiguous interpretation of the nature of the observed signal in terms of the Standard Model. Experimental observations in multiple channels, and with greater precision, are much wanted to make progress in this field. An excellent candidate decay mode for these studies is $D^0 \rightarrow K_S^0 K_S^0$, because its CP asymmetry is expected to be determined by a different combinations of amplitudes, and could possibly be larger than that observed in the charged two-body modes.

My thesis is a contribution to the experimental study of this interesting decay mode. The main reason why $\mathcal{A}^{CP}(D^0 \rightarrow K_S^0 K_S^0)$ has not yet been measured precisely, in spite of its reasonably large branching fraction and the huge production of D^0 mesons available at the LHC, is the difficulty of efficiently selecting these decays in real time (*triggering*), leading to only modest samples being available, and a corresponding low precision of the asymmetry measurements.

The current most precise published measurement of this observable has a resolution of 1.3% [2], nearing the upper bound of theoretical predictions, and does not show evidence of CP violation. That measurement, to which I personally contributed in a significant way, includes some innovations over the past, allowing a better exploitation of the data by also including “secondary” D^0 particles coming from beauty hadron decays and other optimizations. I briefly describe the main features of this analysis, that formed the basis of my expertise in the topic, in the first part of this document. Albeit marking a step forward in sensitivity, the precision of that measurement is still much less than what routinely achieved in the charged modes; and even the prospects with its extension to the larger data sample expected from the Run 3, were not extremely promising for an observation of CP asymmetry.

The only way forward to a better measurement was a better trigger methodology. All past data had been taken with a trigger system that could not recognize long-lived particles like K_S^0 before the second level of the trigger (“HLT2”). This caused a very low trigger efficiency for channels like $D^0 \rightarrow K_S^0 K_S^0$, as most of the signal is lost before even reaching the HLT2 level. This was a point on which further improvement was conceivable, at least in principle.

My work developed in a context where LHCb was preparing for the new run (Run 3) with a completely new data acquisition system, capable of processing the full data flow of 30 MHz in real time. This has been an excellent motivation for me to put in a strong effort in pushing the trigger efficiency for $D^0 \rightarrow K_S^0 K_S^0$ decays to the highest possible values,

to make the best from the new data taking period. I therefore worked at improving the capability of LHCb to trigger on objects decaying significantly away from the interaction point, like K_S^0 . Leveraging the high-level reconstruction offered by the novel LHCb HLT1, I designed, optimized, implemented, and tested a new trigger selection capable of recognizing and selecting long-lived particles (LLPs) like the K_S^0 at the HLT1 level. It turned out to be possible to achieve a significant efficiency enhancement for several interesting physics processes with a limited rate cost – with a maximum gain for the $D^0 \rightarrow K_S^0 K_S^0$ mode, the main target of my work, of a factor 2.6.

The promising performances shown by this trigger on simulation motivated the extension of the above trigger methodology to exploit a new heterogeneous tracking device (Downstream Tracker), planned for installation in the following run (Run 4). This is expected to provide reconstruction of “downstream tracks” at the first level, thus extending the LLPs acceptance to much larger decay lengths. I studied the performances of the new triggers when extended to the Run 4 data-taking configuration, that turned out to provide a further boost of the efficiency, and became part of the motivations of the project.

The start of Run 3 marked a major milestone for the LHCb collaboration, as data-taking commenced following an extensive detector upgrade, with the goal of operating at an instantaneous luminosity five times higher than in previous runs. Achieving this ambitious target required an immense collective effort from collaboration members, who worked tirelessly to optimize both the detector and the data acquisition system to reach the expected performance. During my time at CERN in 2022, I was fortunate enough to be part of this endeavor, to which I contributed by promptly and meticulously analyzing the earliest Run 3 data to commission the novel lines, which were under my responsibility.

The initial tests confirmed the proper and robust operation of the K_S^0 triggers, demonstrating their stability on real pp collision data despite the challenging preliminary data-taking conditions. The promising performance observed in simulation was subsequently validated on real data. Although the 2023 data-taking was severely constrained by external factors, leading to a dataset smaller than expected, the robustness of the newly developed K_S^0 -dedicated selections enabled me to observe the first handful of $D^0 \rightarrow K_S^0 K_S^0$ candidates collected by LHCb in Run 3.

In 2024, LHCb achieved a significant milestone, accumulating an integrated luminosity surpassing that of all previous runs combined. The data collected in the early weeks allowed for a real-data validation of the novel K_S^0 trigger lines, confirming a fourfold increase in efficiency for $D^0 \rightarrow K_S^0 K_S^0$ decays. This achievement is remarkable in itself, marking the first successful implementation at a hadron collider of an early-stage trigger capable of selecting a neutral, long-lived particle like the K_S^0 at the full LHC bunch-crossing rate.

These outstanding performances positioned the measurement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$ as a prime candidate for an early 2024 analysis using newly collected data, aiming to provide timely experimental input on a topic of significant interest in the field [3, 4], while at the same time being a kind of global “system test” of the whole upgraded LHCb apparatus. This measurement is now nearing the final approval by the collaboration. This early measurement, exploiting only in part the potential of the new data, already provides a resolution of 1.05%. With its precision, it stands as the most accurate determination of $\mathcal{A}^{CP}(K_S^0 K_S^0)$ to date, promising new insights into charm CP violation in the near future, while demonstrating the enhanced physics reach of LHCb enabled by the new data acquisition system design.

The thesis is structured as follows. Chapter 1 provides a brief overview of CP violation within the Standard Model, with a focus on charm decays and the significance of $D^0 \rightarrow K_S^0 K_S^0$. Chapter 2 expands on the broad physics opportunities offered by the detection of long-lived particles. Chapter 3 describes the LHCb detector’s experimental setup, covering both its Run 2 and Run 3 configurations. Chapter 4 details the methodology used at LHCb to measure $\mathcal{A}^{CP}(K_S^0 K_S^0)$ in Run 2 data. Chapter 5 presents the simulation studies that guided the design of first-level K_S^0 -dedicated trigger selections. Chapter 6 explores potential extensions of these selections, both within the current configuration and for future applications. Chapter 7 documents the commissioning work carried out on the early Run 3 data. Chapter 8 outlines the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement performed using the new triggers on data collected by LHCb in 2024. Finally, Chapter 9 presents some brief conclusions and future prospects of this work.

Chapter 1

CP violation in charm decays

In this chapter, I introduce the phenomenology of CP violation in the Standard Model of particle physics, focusing on the charm quark by describing its unique characteristics and briefly discussing the measurement that led to the first observation of CP violation in charmed hadron decays. I then turn to the $D^0 \rightarrow K_S^0 K_S^0$ decay, highlighting why it is an excellent candidate for confirming CP violation in charmed hadron decays, as well as addressing the challenges associated with determining the CP asymmetry in these decays.

The violation of CP symmetry in elementary interactions is one of the most intensely investigated phenomena in modern physics. Its fundamental source is unknown, but its consequences on the structure of the Universe as we know it are of enormous importance; it is widely believed to be the root cause of the observed huge and pervasive prevalence of matter over antimatter [5]. The CP operator is a transformation defined by the sequential application of spatial parity P and charge conjugation C . These operate, respectively, the inversion of all spatial coordinates ($(x, y, z) \rightarrow (-x, -y, -z)$ for a vector), and the conjugation of all internal quantum numbers of a system ($Q \rightarrow -Q$ for electric charge). Therefore, CP transforms all particles into their respective antiparticles in a given system and vice versa.

The C and P operators were both thought to be universal symmetries of nature at the beginning of the last century. This was true until Wu *et al.* published their findings in 1956 [6]. In their experiment, the study of the decay of polarized Co^{60} nuclei showed that the weak interaction, responsible for the considered decay, only couples with left-handed particles and right-handed antiparticles, completely violating the P symmetry. After this discovery, Lev Landau proposed the CP symmetry to be the actual symmetry between matter and antimatter [7], conserved by weak interactions. This hypothesis was proved wrong as well, by the 1964 experiment by Cronin and Fitch [8]. Here, the CP -violating process $K_L^0 \rightarrow \pi^+ \pi^-$ was observed, demonstrating that the weak interaction also violated the CP symmetry. In contrast, electromagnetic and strong interaction did not show any evidence of P , C or CP violation to date, despite all searches carried out.

After the observation of the CP violation, Sakharov linked this phenomenon to the matter-antimatter imbalance we see in our Universe (*baryogenesis*) indicating the presence of CP -violation during early-universe evolution as one of the necessary conditions for baryogenesis [9]. These findings marked the beginning of a multi-decade exploration of the CP violation phenomenology, conducted by numerous experiments worldwide.

Before proceeding in the discussion, it is useful to briefly summarize the theoretical

framework currently used to model CP violation.

1.1 Standard Model

The Standard Model (SM) of particle Physics is the theory describing the interaction, through electroweak and strong forces, of known elementary particles, schematically reported in Figure 1.1. Quarks and leptons make up the class of *fermions*, both divided into three generations. The remaining particles are *bosons*. The gluon, the photon, the W^\pm and Z are mediators of strong and electroweak forces, respectively, while the Higgs boson, mediates the interaction with the Higgs field. Since only the weak interaction is




Figure 1.1: Schematic view of elementary particles included in the SM.

responsible of CP violation, I will focus my attention on its theoretical formalism.

1.1.1 CPV within SM formalism

CP violation is described within the SM Lagrangian by an irreducible complex phase in the weak interaction term. This can be expressed in the basis of mass eigenstates as:

$$\mathcal{L}_{int}^{CC} = -\frac{g}{\sqrt{2}}(\bar{u}_L, \bar{c}_L, \bar{t}_L)\gamma^\mu V_{CKM} \begin{pmatrix} d_L \\ s_L \\ b_L \end{pmatrix} W_\mu^\dagger + h.c., \quad (1.1)$$

where g is the gauge coupling of the $SU(2)$ weak-isospin group of the SM, \bar{u}_L , \bar{c}_L , \bar{t}_L , d_L , s_L and b_L are left-handed quarks representation, γ^μ are the Dirac matrices, V_{CKM} is the Cabibbo-Kobayashi-Maskawa (CKM) matrix [10, 11] and W_μ is the charged weak gauge boson representation. V_{CKM} describes the coupling of quarks through the weak

interaction, and it can be explicitly expressed as:

$$V_{CKM} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix}. \quad (1.2)$$

Because of the non-correspondence between mass and weak interaction quark eigenstates, it is not diagonal, allowing for transitions between different generations. As a 3×3 unitary matrix, it can be characterized using only four distinct real parameters. Using the three mixing angles θ_{12} , θ_{23} and θ_{13} and a single complex phase δ , V_{CKM} can be written as:

$$V_{CKM} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{13} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}, \quad (1.3)$$

where s_{ij} and c_{ij} stand for $\sin\theta_{ij}$ and $\cos\theta_{ij}$.

The coupling between different families is not equally probable, and the size of different V_{CKM} elements can differ by orders of magnitude. This can be made manifest adopting the Wolfenstein parametrisation [12], where four parameters λ , A , ρ and η are exploited:

$$V_{CKM} = \begin{pmatrix} 1 - \lambda^2/2 & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \lambda^2/2 & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix} + O(\lambda^4), \quad (1.4)$$

where λ is the expansion parameter, historically known as the sine of the Cabibbo angle ($\sin\theta_C \simeq 0.22$). In terms of this parameter, it can be more easily seen how different quark generations mix. Coupling between the same generation is significantly favored, while the transition is suppressed by a factor λ , λ^2 and λ^3 for the first and second, second and third, first and third generations, respectively. Current experimental knowledge of CKM matrix elements is [13]:

$$V_{CKM} = \begin{pmatrix} 0.97410 \pm 0.00011 & 0.22650 \pm 0.00048 & 0.00361^{+0.00011}_{-0.00009} \\ 0.22636 \pm 0.00048 & 0.97320 \pm 0.000011 & 0.04053^{+0.00083}_{-0.00061} \\ 0.00854^{+0.00023}_{-0.00016} & 0.03978^{+0.00082}_{-0.00060} & 0.999172^{+0.000024}_{-0.000035} \end{pmatrix}. \quad (1.5)$$

V_{CKM} elements directly enter into the computation of the decay amplitudes, and weak decay processes are usually classified based on the lower power of λ appearing in the contributing matrix element. Different categories are referred to as Cabibbo favored (CF), singly Cabibbo suppressed (SCS) and doubly Cabibbo suppressed (DCS), corresponding to λ^0 , λ^1 and λ^2 , respectively.

In order to better understand how CP violation arises from the presence of a single complex phase in the CKM matrix, it is useful to consider a minimal example of a hadron decay from an initial state $|i\rangle$ to a final state $|f\rangle$:

$$\mathcal{A}_f = \langle f | \mathcal{H} | i \rangle, \quad (1.6)$$

and its CP -conjugate process:

$$\bar{\mathcal{A}}_{\bar{f}} = \langle \bar{f} | \mathcal{H} | \bar{i} \rangle, \quad (1.7)$$

where \mathcal{H} is the Hamiltonian describing the decay process. Decay amplitudes can generally be expressed as:

$$\mathcal{A}_f = \sum_{i=1}^n |\mathcal{A}_n| e^{i(\phi_n + \delta_n)}, \quad \bar{\mathcal{A}}_{\bar{f}} = \sum_{i=1}^n |\mathcal{A}_n| e^{i(-\phi_n + \delta_n)}, \quad (1.8)$$

where i runs over different processes that can contribute to the decay, $|\mathcal{A}_n|$ is the amplitude of each involved process, δ_n and ϕ_n are the so-called *strong* and *weak* phases, respectively. The two contributions to the phase are kept separate because of their different behavior under the CP transformation. δ_n are defined as the phases that do not change sign under CP transformation, and are named like this since they are usually arising from the strong interactions involved in the decay. ϕ_n are, on the opposite, defined as the ones that change sign under the same transformation, and are referred to as weak phases, since the only source of CP -odd phases in the SM is the δ phase contained in the V_{CKM} matrix.

The most common observable exploited in the search for CP -violating effects is the so-called CP -asymmetry (\mathcal{A}^{CP}), hence the asymmetry between the amplitude of a process and of its CP -conjugate. This can be expressed as:

$$\frac{|\mathcal{A}_f|^2 - |\bar{\mathcal{A}}_{\bar{f}}|^2}{|\mathcal{A}_f|^2 + |\bar{\mathcal{A}}_{\bar{f}}|^2} = \frac{\sum_{n \neq m} |\mathcal{A}_n| |\mathcal{A}_m| \sin(\phi_n - \phi_m) \sin(\delta_n - \delta_m)}{\sum_n |\mathcal{A}_n|^2 + \sum_{n \neq m} |\mathcal{A}_n| |\mathcal{A}_m| \cos(\phi_n - \phi_m) \cos(\delta_n - \delta_m)}. \quad (1.9)$$

The reported expression indicates that to have a non-zero \mathcal{A}^{CP} , the decay should be influenced by two distinct processes that exhibit different strong and weak phases. A graphical visualization of CP violation within the SM is possible considering the six equations arising from the CKM matrix unitarity requirement $V_{CKM}^\dagger V_{CKM} = \mathbb{I}$. These can be geometrically represented as triangles in the $(\bar{\rho} - \bar{\eta})$ plane and are usually called *unitarity triangles*. A visual representation of these is shown in Figure 1.2. All the triangles




Figure 1.2: Visual representation of the six unitarity triangles. Taken from [14].

have the same area $|J|/2$, where J is the Jarlskog invariant [15], defined as:

$$\text{Im}[V_{ij} V_{kl} V_{il}^* V_{kj}^*] = J \sum_{n,m \in (d,s,b)} \epsilon_{ikm} \epsilon_{jln}. \quad (1.10)$$

Exploiting Wolfenstein parametrization, value of J can be approximated as $J \simeq \lambda^6 A^2 \eta$, and current experimental knowledge of its value is $J = (3.00)_{-0.09}^{+0.15} \times 10^{-5}$ [13]. In the SM, all the CP -violating observables are proportional to J . Therefore, the non-null observed value for J confirms the presence of CP violation. This can be visualized graphically as having triangles that do not degenerate into lines.

1.1.2 Neutral mesons phenomenology and mixing

The neutral meson system provides an excellent framework for investigating CP violation, as it offers multiple distinct manifestations of this phenomenon. Because of this, I will briefly introduce the phenomenology of these particles. A unique feature of such particles, namely K_S^0 , D^0 , B^0 and \bar{B}_s^0 , is their nonzero probability of oscillating into their own antiparticle, through a transition where the flavor quantum number of the meson changes by two ($\Delta F = 2$ transitions). This phenomenon is referred to as *mixing* and it happens because of the non-correspondence between weak and mass eigenstates.

Let's consider the time evolution of a given meson M^0 . For times significantly larger than the typical strong interaction, it can be described by the following Schrödinger equation:

$$i \frac{d}{dt} \begin{pmatrix} M^0(t) \\ \bar{M}^0(t) \end{pmatrix} = \left[\mathbf{M} - \frac{i}{2} \mathbf{\Gamma} \right] \begin{pmatrix} M^0(t) \\ \bar{M}^0(t) \end{pmatrix} \quad (1.11)$$

where \mathbf{M} and $\mathbf{\Gamma}$ being 2x2 Hermitian matrices,

$$\mathbf{M} = \begin{pmatrix} M_{11} & M_{12} \\ M_{12}^* & M_{22} \end{pmatrix} \quad \mathbf{\Gamma} = \begin{pmatrix} \Gamma_{11} & \Gamma_{12} \\ \Gamma_{12}^* & \Gamma_{22} \end{pmatrix}. \quad (1.12)$$

The $\mathbf{M} - i\mathbf{\Gamma}/2$ term is the effective Hamiltonian describing the time evolution of the system (\mathcal{H}_{eff}). The flavor-conserving transitions ($\Delta F = 0$) $M^0 \rightarrow M^0$ and $\bar{M}^0 \rightarrow \bar{M}^0$ are associated with diagonal elements, and they are constrained by the system invariance under CPT transformation to be $M_{11} = M_{22}$ and $\Gamma_{11} = \Gamma_{22}$. The off-diagonal elements are associated with flavor-changing transitions ($\Delta F = 2$), through which mixing can occur. When the probability of mixing is non-zero, M_{12} and Γ_{12} are non-vanishing, making \mathcal{H}_{eff} non-diagonal. As a result, its eigenstates are not identical to $|M^0\rangle$ and $|\bar{M}^0\rangle$, but instead are superpositions of these states. These can be expressed as:

$$|M_{L,H}^0\rangle = p|M^0\rangle + q|\bar{M}^0\rangle, \quad (1.13)$$

where L, R stands for light and heavy, referring to the masses of \mathcal{H}_{eff} eigenstates and p and q are the complex coefficients describing the relation between the two eigenstate sets. The complex coefficients are related by the following expressions:

$$|p|^2 + |q|^2 = 1 \quad \text{and} \quad \frac{q}{p} = \sqrt{\frac{M_{12}^* - i\Gamma_{12}^*/2}{M_{12} - i\Gamma_{12}/2}} = \left| \frac{q}{p} \right| e^{i\Phi}. \quad (1.14)$$

Mass and width difference between light and heavy states is usually expressed exploiting two dimensionless parameters x and y :

$$x = \frac{\Delta m}{\Gamma} \quad y = \frac{\Delta \Gamma}{2\Gamma}. \quad (1.15)$$

These dictate the time evolution of a particle created in a specific flavor state as:

$$\begin{aligned} |M^0(t)\rangle &= g_+(t)|M^0\rangle + \frac{q}{p}g_-(t)|\bar{M}^0\rangle, \\ |\bar{M}^0(t)\rangle &= g_+(t)|\bar{M}^0\rangle + \frac{p}{q}g_-(t)|M^0\rangle, \end{aligned} \quad (1.16)$$

where $g_{\pm}(t)$ are defined as:

$$|g_{\pm}(t)|^2 = \frac{1}{2}e^{-t/\tau} \left[\cos\left(\frac{xt}{\tau}\right) \pm \cosh\left(\frac{yt}{\tau}\right) \right]. \quad (1.17)$$

1.1.3 CP violation in the decay

Following this introduction to the phenomenology of neutral meson decays, it is important to examine the possible manifestations of CP violation within this system. One such manifestation is CP violation occurring in their decay, usually referred to as *direct* CP -violation. Let us consider the case where the same final state f is accessible from both CP -conjugated states M^0 and \bar{M}^0 . In this case, the golden observable to measure CP -violating effects is the previously defined \mathcal{A}^{CP} . Given this specific initial and final states, it can be expressed as:

$$\mathcal{A}^{CP}(f) = \frac{\Gamma(M^0 \rightarrow f) - \Gamma(\bar{M}^0 \rightarrow f)}{\Gamma(M^0 \rightarrow f) + \Gamma(\bar{M}^0 \rightarrow f)}, \quad (1.18)$$

with Γ being the time-integrated decay width of the considered decay. Γ is proportional to the squared amplitude of the decay process and $\mathcal{A}^{CP}(f)$ can be expressed as:

$$\mathcal{A}^{CP}(f) = \frac{|\mathcal{A}_f|^2 - |\bar{\mathcal{A}}_f|^2}{|\mathcal{A}_f|^2 + |\bar{\mathcal{A}}_f|^2} = \frac{1 - \mathcal{R}_f^2}{1 + \mathcal{R}_f^2}, \quad (1.19)$$

where:

$$\mathcal{R}_f = \left| \frac{\bar{\mathcal{A}}_f}{\mathcal{A}_f} \right|. \quad (1.20)$$

From this expression it is clear how direct CP -violation arises in case \mathcal{R}_f differs from 1, hence when the two CP -conjugated states have a different probability to go into the considered final state f .

1.1.4 CP violation in mixing

Another manifestation of CP violation can be observed in the phenomenon of mixing, where neutral mesons oscillate between their particle and antiparticle states. This type of CP violation arises when the mass eigenstates $|M_L^0\rangle$ (light) and $|M_H^0\rangle$ (heavy) are not composed of the same mixture of the flavor eigenstates $|M^0\rangle$ and $|\bar{M}^0\rangle$.

To quantify CP violation in mixing, the dimensionless parameter \mathcal{R}_m is introduced, defined as:

$$\mathcal{R}_m = \left| \frac{q}{p} \right|.$$

In the absence of CP violation, the coefficients p and q are equal in magnitude, yielding $\mathcal{R}_m = 1$. However, if $\mathcal{R}_m \neq 1$, this signals the presence of CP violation in mixing. Physically, this inequality implies that the probabilities for a neutral meson $|M^0\rangle$ to transition into its antiparticle state $|\bar{M}^0\rangle$ and vice versa are not identical.

1.1.5 CP violation in the interference between decay and mixing

The last manifestation of CP violation we consider arises from the interference between decay and mixing. This occurs when both CP -conjugate states $|M^0\rangle$ and $|\bar{M}^0\rangle$ can decay into the same final state $|f\rangle$, and mixing between $|M^0\rangle$ and $|\bar{M}^0\rangle$ is also possible. In this scenario, $|M^0\rangle$ can decay into the final state f both directly ($|M^0\rangle \rightarrow f$) and through mixing ($|M^0\rangle \rightarrow |\bar{M}^0\rangle \rightarrow f$), leading to interference between these two decay paths. The same occurs for the CP -conjugate state $|\bar{M}^0\rangle$.

Considering the expression reported in Eq. (1.16) for an initial state of pure $|M^0\rangle$, the time-dependent amplitude for it to go into the final state $|f\rangle$ can be expressed as:

$$\langle f|\mathcal{H}|M^0(t)\rangle = \mathcal{A}_f g_+(t) + \bar{\mathcal{A}}_f \frac{q}{p} g_-(t). \quad (1.21)$$

From this expression, it is possible to compute the actual decay rate, that is:

$$\begin{aligned} \frac{d\Gamma}{dt}(M^0(t) \rightarrow f) \propto |\mathcal{A}_f|^2 \Bigg[& (1 - |\lambda_f|^2) \cos\left(\frac{xt}{\tau}\right) + (1 + |\lambda_f|^2) \cosh\left(\frac{yt}{\tau}\right) \\ & - 2 \cdot \text{Im}(\lambda_f) \sin\left(\frac{xt}{\tau}\right) + 2 \cdot \text{Re}(\lambda_f) \sinh\left(\frac{yt}{\tau}\right) \Bigg], \end{aligned} \quad (1.22)$$

where:

$$\lambda_f = \frac{q\bar{\mathcal{A}}_f}{p\mathcal{A}_f}. \quad (1.23)$$

A similar expression can be extracted for the case of $|\bar{M}^0\rangle$ decay and it will not be reported here because of brevity reasons. The asymmetry between the two happens to be different from zero when the complex phase of the λ_f parameter is non-null, hence:

$$\text{Im}(\lambda_f) \neq 0. \quad (1.24)$$

1.2 CP violation in charmed hadron decays

The role of charmed hadron decays in the study of CP violation is unique. They provide the only opportunity to measure CP violation effects associated with the decay of up-type quarks, since up mesons are CP eigenstates (*e.g.* π^0) and the top quark decays before reaching a bound state in a meson, due to its very large mass. As a result, they allow probing a different set of interactions compared to B and K mesons, where there is the decay of a down-type quark. The study of CP violation in charmed hadron decays is so interesting also because of the expected size of the CP violation effects. This arises from the smaller mass of the beauty quark compared to the top quark involved in the relevant loop diagrams, as well as the small CKM matrix elements that couple the first two generations of quarks to the third. CP violation is directly linked to this specific combination of CKM matrix elements:

$$\text{Im}(V_{cb}V_{ub}^*/V_{cs}V_{us}^*) \approx -6 \times 10^{-4}, \quad (1.25)$$

which leads to charm CP asymmetries typically in the range of 10^{-4} to 10^{-3} . The size of expected CP violation effects provides an excellent null test for SM predictions, making

any precise measurement of \mathcal{A}^{CP} in charmed hadron decays a sensitive probe for the influence of New Physics (NP).

The mentioned arguments motivate the efforts for a detailed investigation of CP violation in charmed hadron decays. However, the small expected CP violation size poses several experimental challenges that must be faced, such as the collection of very large statistics samples and comprehending experimental effects at a sub-per-mille scale. As a consequence, even if the search for charm CP violation has been widely investigated, any evidence for an effect lacked until very recently. In contrast, evidence was found decades ago for the systems of K and B decays, in 1964 by Cronin and Fitch and in 2001 by Belle [16] and BaBar [17] experiments, respectively.

Providing a clear theoretical interpretation of the experimental results is challenging due to non-negligible contributions from non-perturbative strong interactions involving strange and down quarks [18], which are present at the energy scales of the *charm* quark ($\simeq 2 \text{ GeV}/c^2$). These contributions face additional complications because they are subject to large cancellations, vanishing in the $SU(3)_F$ limit where the masses of the strange, down, and up quarks are considered negligible compared to the typical hadronic scale of charm decays, Λ_{QCD} , making precise predictions even more challenging.

A rigorous comparison between experimental measurements and theoretical predictions will only be possible with improvements on both of these aspects.

1.3 First CP violation observation in charm decays

A significant step toward this direction has been represented by the first observation of CP violation in charm decays. This pivotal result has been achieved in 2019 by the LHCb experiment [1], by measuring the difference between the time-integrated \mathcal{A}^{CP} of $D^0 \rightarrow K^+ K^-$ and $D^0 \rightarrow \pi^+ \pi^-$ decays, defined as $\Delta\mathcal{A}^{CP} = \mathcal{A}^{CP}(K^+ K^-) - \mathcal{A}^{CP}(\pi^+ \pi^-)$. The result, achieved using a sample of 5.9 fb^{-1} collected at a center-of-mass energy of 13 TeV, is:

$$\Delta\mathcal{A}^{CP} = (-15.4 \pm 2.9) \times 10^{-4}. \quad (1.26)$$

This result is a milestone in experimental particle physics, being the endpoint of a decades-long research, carried out by various experiments, such as Belle, BaBar and LHCb itself. The achievement of such a precise result has been made possible by the very large samples collected by the LHCb experiment, corresponding to 54×10^6 $D^0 \rightarrow K^+ K^-$ and 17×10^6 $D^0 \rightarrow \pi^+ \pi^-$ decays. The accumulation of such extensive samples was possible thanks to the very large production cross section of charmed hadrons provided by the proton collisions happening in the LHC and to the capabilities of the LHCb trigger system, designed to effectively collect heavy-flavor decays.

In order to extract \mathcal{A}^{CP} for considered decays, the flavour of the D^0 has to be determined. This is done by requiring the D^0 to come from either a $D^{*+} \rightarrow D^0 \pi^+$ and $\bar{B} \rightarrow D^0 \mu^- \bar{\nu}_\mu X$ decay. These strong decays allow the determination of the D^0 flavor by the sign of the accompanying pion (π -tagged sample) or muon (μ -tagged sample). The analysis was carried out separately for each of the two samples. The asymmetry between

the number of D^0 and \bar{D}^0 is measured on data as:

$$\begin{aligned}\mathcal{A}_{\text{raw}}^{\pi\text{-tagged}}(f) &= \frac{N(D^{*+} \rightarrow D^0 \pi^+) - N(D^{*-} \rightarrow \bar{D}^0 \pi^-)}{N(D^{*+}(f) \rightarrow D^0 \pi^+) + N(D^{*-}(f) \rightarrow \bar{D}^0 \pi^-)}, \\ \mathcal{A}_{\text{raw}}^{\mu\text{-tagged}}(f) &= \frac{N(\bar{B} \rightarrow D^0 \mu^- \bar{\nu}_\mu X) - N(B \rightarrow \bar{D}^0 \mu^+ \nu_\mu X)}{N(\bar{B} \rightarrow D^0 \mu^- \bar{\nu}_\mu X) + N(B \rightarrow \bar{D}^0 \mu^+ \nu_\mu X)},\end{aligned}\quad (1.27)$$

where $f = \pi^+ \pi^-, K^+ K^-$ and N is the measured signal yield for each given decay. These asymmetries are referred to as *raw*, because they do not directly correspond to the physical asymmetry \mathcal{A}^{CP} , but to the sum of this with the so-called nuisance asymmetries, hence detection and production asymmetry \mathcal{A}_{det} and $\mathcal{A}_{\text{prod}}$. Raw asymmetries can be approximated as:

$$\begin{aligned}\mathcal{A}_{\text{raw}}^{\pi\text{-tagged}}(f) &\simeq \mathcal{A}^{CP}(f) + \mathcal{A}_{\text{det}}(\pi) + \mathcal{A}_{\text{prod}}(D^*), \\ \mathcal{A}_{\text{raw}}^{\mu\text{-tagged}}(f) &\simeq \mathcal{A}^{CP}(f) + \mathcal{A}_{\text{det}}(\mu) + \mathcal{A}_{\text{prod}}(B).\end{aligned}\quad (1.28)$$

Given this expansion, $\Delta\mathcal{A}^{CP}$ can be directly extracted from the difference of raw asymmetries as:

$$\Delta\mathcal{A}_{\text{raw}} = \mathcal{A}_{\text{raw}}(K^+ K^-) - \mathcal{A}_{\text{raw}}(\pi^+ \pi^-) = \mathcal{A}^{CP}(K^+ K^-) - \mathcal{A}^{CP}(\pi^+ \pi^-) = \Delta\mathcal{A}^{CP}. \quad (1.29)$$

From the previous equation, it can be seen how powerful this observable is, as it ensures the cancellation of all spurious effects without the exploitation of a calibration channel. Additionally, $\mathcal{A}^{CP}(K^+ K^-)$ and $\mathcal{A}^{CP}(\pi^+ \pi^-)$ are expected to be of the same magnitude, but opposite in sign, increasing the size of the expected effect and making $\Delta\mathcal{A}^{CP}$ more sensitive to CP violation effects with respect to the single observables.

Results obtained from the separate analysis of the two differently tagged samples are:

$$\begin{aligned}\Delta\mathcal{A}^{CP\pi\text{-tagged}} &= [-18.2 \pm 3.2(\text{stat.}) \pm 0.9(\text{syst.})] \times 10^{-4}, \\ \Delta\mathcal{A}^{CP\mu\text{-tagged}} &= [-9 \pm 8(\text{stat.}) \pm 5(\text{syst.})] \times 10^{-4},\end{aligned}\quad (1.30)$$

both in agreement both with previous LHCb results [19, 20] and world average [21]. The combination of these results with all previous LHCb measurements corresponds to:

$$\Delta\mathcal{A}^{CP} = (-15.4 \pm 2.9) \times 10^{-4}, \quad (1.31)$$

resulting in a 5σ -significance on the measured parameter, and the observation of CP violation in charmed hadron decays.

1.3.1 Result interpretation

In order to interpret this result, it is useful to expand on the theoretical meaning of the measured observable. The $\mathcal{A}^{CP}(K^+ K^-)$, $\mathcal{A}^{CP}(\pi^+ \pi^-)$ considered in the reported measurement are the so-called time integrated asymmetries, obtained integrating over all possible D^0 decay time values. They correspond to a mixture of two effects, as:

$$\mathcal{A}^{CP}(f) = a_{\text{dec}}^{CP}(f) - \frac{\langle t(f) \rangle}{\tau(D^0)} A_\Gamma(f). \quad (1.32)$$

The first term is the component due to the CP violation in the decay, independent of time. The second is the time-dependent CP violation component, where $\langle t(f) \rangle$ is the $D^0 \rightarrow f$ mean decay time in the analyzed sample, $\tau(D^0)$ is the D^0 lifetime and $A_\Gamma(f)$ is the $D^0/\bar{D}^0 \rightarrow f$ effective decay width asymmetry and represents the CP asymmetry given by mixing and interference between decay and mixing [22–24]. Making the time-dependent and time-independent terms explicit in the $\Delta\mathcal{A}^{CP}$ expression, one obtains:

$$\Delta\mathcal{A}^{CP} = \Delta a_{dec}^{CP} - \frac{\Delta\langle t \rangle}{\tau(D^0)} \Delta A_\Gamma \simeq \Delta a_{dec}^{CP} - \frac{\Delta\langle t \rangle}{\tau(D^0)} A_\Gamma, \quad (1.33)$$

where in the last equivalence A_Γ has been approximated to be equal for both final states [25]. The second term is expected to be of the order of 10^{-5} , given that the two factors composing it are measured to be $A_\Gamma = (-2.8 \pm 2.8) \times 10^{-4}$ and $\Delta\langle t \rangle/\tau(D^0) = 0.115 \pm 0.002$. Therefore, it is clear that the impact of time-dependent CP violation given by mixing and interference is almost negligible and the measured effect is given with good approximation by CP violation in the decay. The direct CP violation contribution in $\Delta\mathcal{A}^{CP}$ obtained by subtracting the time-dependent effect, is:

$$\Delta\mathcal{A}_{dec}^{CP} = (-15.7 \pm 2.9) \times 10^{-4}. \quad (1.34)$$

As for other charm CP violation parameters, a theoretical prediction with a precision similar to the experimental one is not yet available. At the current state of the art, predictions rely in the $10^{-4} - 10^{-3}$ range [26–34], making an exhaustive comparison between theory and experiment impossible. All we can say is that the experimental result is compatible with the theoretically expected range, albeit close to its high end.

A promising approach for the estimation of expected value of $\Delta\mathcal{A}^{CP}$ is the so-called topological-diagram approach [35]. This alternative approach extracts magnitude and phase of weak annihilation amplitudes arising from final state rescattering directly from experimental data, avoiding QCD calculations and possibly achieving a more precise prediction. Therefore, the addition of further measurements of charm-related CP violation quantities can be crucial to clarify the theoretical interpretation.

1.4 The role of the $D^0 \rightarrow K_S^0 K_S^0$ mode

The $D^0 \rightarrow K_S^0 K_S^0$ channel is an excellent candidate for a second, confirmatory charm CP violation observation. The mode is closely related to the mode of first observation, but accidental effects conjure to give a significantly larger predicted CP asymmetry, that can be up to 1.1% according to some predictions [36, 37]. $\mathcal{A}^{CP}(K_S^0 K_S^0)$ is usually measured as a time-integrated CP asymmetry, and it can be expressed as:

$$\mathcal{A}^{CP}(K_S^0 K_S^0) = a_{dec}^{CP}(K_S^0 K_S^0) - \frac{\langle t(K_S^0 K_S^0) \rangle}{\tau(D^0)} A_\Gamma, \quad (1.35)$$

similarly as done in 1.32. Given the values of $A_\Gamma = (-2.8 \pm 2.8) \times 10^{-4}$ and $\langle t(K_S^0 K_S^0) \rangle/\tau(D^0) \sim 1$, the second term in Eq. (1.35) can be neglected, given the current sensitivity on this parameter is of the order of $\sim 1\%$. Therefore the time-integrated measurement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$ is with good approximation a measurement of direct CP violation:

$$\mathcal{A}^{CP}(K_S^0 K_S^0) \simeq a_{dec}^{CP}(K_S^0 K_S^0). \quad (1.36)$$

In order to better understand what determines the possible enhancement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$ size, it is convenient to exploit the mentioned topological-diagram approach, expanding further on it. All involved processes are classified and grouped according to their decay topology and in particular two groups are identified. The first one includes tree and penguin amplitudes, that are: color-allowed (T) and color-suppressed (C) tree amplitudes (where an external and internal emission of a W boson is present, respectively), QCD-penguin amplitude (P), color-favored (P_{EW}) and color-suppressed (P_{EW}^C) electro-weak penguin amplitudes, and the singlet QCD-penguin amplitude (S), which involves SU(3)-singlet mesons like η , ω and ϕ . The second group includes weak annihilation amplitudes, classified as: W -exchange (E) and W -annihilation (A) amplitudes, QCD-penguin exchange (PE) and QCD-penguin annihilation (PA) amplitudes, electro-weak penguin exchange (PE_{EW}) and electro-weak penguin annihilation (PA_{EW}) amplitudes. Visual representation of these diagrams is displayed in Figure 1.3. It is important to notice how these are not Feynman




Figure 1.3: Graphical representation of possible flavor diagrams in the topological approach. Taken from [35].

diagrams, since they are meant to implicitly include all possible QCD effects (*e.g.* gluon lines, or quark loops). Only weak interactions and flavour flows are exploited in the classification.

The $D^0 \rightarrow K_S^0 K_S^0$ decay is classified as singly Cabibbo suppressed. The usual decomposition in terms of CKM matrix elements for these decays is:

$$\mathcal{A} = \frac{1}{2} [(\lambda_s - \lambda_d)\mathcal{A}_{sd} + \lambda_b\mathcal{A}_b], \quad (1.37)$$

where $\lambda_q \equiv V_{cq}^* V_{uq}$ and \mathcal{A}_q represent the involved CKM matrix elements and the relative involved amplitudes, respectively. The size of direct CP violation can be derived from the reported expression for the decay amplitude, as:

$$a_{dec}^{CP}(K_S^0 K_S^0) = \frac{|\mathcal{A}|^2 - |\bar{\mathcal{A}}|^2}{|\mathcal{A}|^2 + |\bar{\mathcal{A}}|^2} = \text{Im} \frac{\lambda_b}{\lambda_s - \lambda_d} \text{Im} \frac{\mathcal{A}_b}{\mathcal{A}_{sd}}. \quad (1.38)$$

The first term is given by the relative magnitude of CKM matrix elements for contributing decays, at first and second order, and has a value of:

$$\text{Im} \frac{\lambda_b}{\lambda_s - \lambda_d} \sim 6 \times 10^{-4}. \quad (1.39)$$

In order to determine the expected size for the second term it is useful to decompose \mathcal{A}_{sd} and \mathcal{A}_b in terms of topological amplitudes. These can be combined making use of the so-called *Wilson coefficients*, that in case of the $D^0 \rightarrow K_S^0 K_S^0$ corresponds to [37, 38]:

$$\mathcal{A}_{sd} = \frac{E_1 + E_2 - E_3}{\sqrt{2}} \quad (1.40)$$

and

$$\mathcal{A}_b = \frac{2E + E_1 + E_2 + E_3 + PA}{\sqrt{2}}, \quad (1.41)$$

where E_1 , E_2 , E_3 , E and PA are the topological amplitudes involved in the decay, reported in Figures 1.4 and 1.5. E_1 , E_2 and E_3 are tree level amplitudes and they are non-vanishing




Figure 1.4: Graphical representation of SU(3)-breaking topological amplitudes. Taken from [39].




Figure 1.5: Representation of exchange (E) and penguin annihilation (PA) topological amplitudes involved in the $D^0 \rightarrow K_S^0 K_S^0$ decay. Taken from [40].

only in case the SU(3) symmetry is broken. The SU(3) symmetry assumes the three light quarks to be exactly exchangeable, with $m_u = m_d = m_s$, and it's exploited in the theoretical calculations to describe the decay modes amplitudes with a reduced number of free parameters. SU(3) is known to be just slightly broken ($\sim O(30\%)$ [39]), and this could suppress \mathcal{A}_{sd} , since only E_1 , E_2 and E_3 contribute to it. On the contrary, \mathcal{A}_b receives also contributions from E and PA diagrams, not affected by the SU(3) limit. The combination of these two effects could cause an enhancement of the $\mathcal{A}_{sd}/\mathcal{A}_b$ ratio, resulting in the mentioned possible large $\mathcal{A}^{CP}(K_S^0 K_S^0)$ size. However, it should be mentioned that not all predictions concur on such a large value for $\mathcal{A}^{CP}(K_S^0 K_S^0)$, some being closer to more typical charm \mathcal{A}^{CP} size of 10^{-3} [30, 41, 42]. This makes the experimental measurement interesting also from the viewpoint of discriminating between calculation approaches.

$D^0 \rightarrow K_S^0 K_S^0$ is quite interesting for being sensitive to a different mix of amplitudes with respect to decays involved in $\Delta\mathcal{A}^{CP}$ measurement. As a result, evidence of a non-zero

CP violation effect in $D^0 \rightarrow K_S^0 K_S^0$ decays could also contribute to refining theoretical predictions for other measurements, with $\Delta\mathcal{A}^{CP}$ being the most significant.

1.5 Experimental difficulties in measuring $\mathcal{A}^{CP}(K_S^0 K_S^0)$

The previous section highlighted the strong motivations currently standing for a search for CP violation effects in $D^0 \rightarrow K_S^0 K_S^0$ decays. Because of this, several measurements of $\mathcal{A}^{CP}(K_S^0 K_S^0)$ have been performed at different experiments and facilities. The current status is summarized in Table 1.1. As it can be seen, this topic is garnering significant

Collaboration	Year	Yield	$\mathcal{A}^{CP}(K_S^0 K_S^0)$ [%]
Cleo [43]	2008	65 ± 14	$-23. \pm 19.$
LHCb Run 1 [44]	2015	635 ± 74	$-2.9 \pm 5.2 \pm 2.2$
LHCb Run 2 [2]	2021	8102 ± 136	$-3.1 \pm 1.2 \pm 0.4 \pm 0.2$
CMS [3]	2024	2046 ± 64	$6.2 \pm 3.0 \pm 0.2 \pm 0.8$
Belle + Belle II [4]	2024	7078 ± 93	$-1.4 \pm 1.3 \pm 0.1$
Average	-	-	-1.7 ± 0.9

Table 1.1: Recent measurements of $\mathcal{A}^{CP}(K_S^0 K_S^0)$. The first reported uncertainty is the statistical one, the second, and the third where present, is systematic.

attention in the field – two new results have been published just last year, by CMS and Belle and Belle II collaborations. The most precise determination is still the one I took part in, published by LHCb in 2021 - although the latest (still unpublished) combined measurement by Belle and Belle II comes very close. It is worth noting that the world average precision has now dropped below the percent level, entering the range of potential predictions. This makes the prospect of a further improvement in the precision of this measurement particularly compelling.

However, achieving a more precise measurement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$ poses significant challenges at both hadron and lepton colliders. The Belle and Belle II experiments, for instance, operate in an environment particularly well-suited for measurements involving fully neutral final states like $\mathcal{A}^{CP}(K_S^0 K_S^0)$. They benefit from the low track multiplicity characteristic of electron collisions, enabling the implementation of a trigger with an overall efficiency exceeding 99% [45], and the collection of high-purity samples. However, the biggest challenge for the collection of a sizeable $D^0 \rightarrow K_S^0 K_S^0$ sample at e^+e^- colliders working at the $\Upsilon(4S)$ comes from the available cross section for D^0 production, much lower than that of hadron colliders.

For example, considering the D^0 production, the rate achieved at LHCb is:

$$R(D^0)_{\text{LHCb}} = \mathcal{L}_{\text{LHC}} \times \sigma(pp \rightarrow D^0 X) = 4 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1} \times 2072 \text{ } \mu\text{b} \simeq 1 \text{ MHz}, \quad (1.42)$$

where \mathcal{L}_{LHC} is the Run 2 LHC peak instantaneous luminosity, $\sigma(pp \rightarrow D^0 X)$ is the inclusive D^0 production cross section at 13 TeV in the LHCb acceptance [46]. This means that at the LHC non-empty bunch crossing rate of 30 MHz, 3% of collisions produce a D^0 meson, on average. A similar estimate can be made for e^+e^- collisions; for instance, considering the rate achieved by KEKB for the Belle experiment, that is:

$$R(D^0)_{\text{Belle}} = \mathcal{L}_{\text{KEKB}} \times \sigma(e^+e^- \rightarrow D^0 X @ 10.6 \text{ GeV}) \simeq 760 \text{ Hz}, \quad (1.43)$$

where $\mathcal{L}_{\text{KEKB}}$ is the peak instantaneous luminosity of KEKB, $\sigma(e^+e^- \rightarrow D^0 X @ 10.6 \text{ GeV})$ [47] is the D^0 inclusive production cross section at the Belle running conditions. Three orders of magnitude separate the two rates, showing how a much larger integrated luminosity needs to be collected at a lepton collider in order to accumulate a sample of similar size. This is confirmed by considering the integrated luminosity of 1.4 ab^{-1} exploited in the Belle and Belle II measurement, larger by a factor 240 with respect to the 5.9 fb^{-1} exploited in the LHCb measurement, even if the achieved sample size and precision is of the same order.

On the other hand, experiments collecting data at hadron colliders, such as LHCb, face significant challenges in efficiently identifying signal decays within a much harsher environment characterized by a higher track multiplicity. This challenge is particularly pronounced in the online identification and selection of fully neutral, long-lived final states like $D^0 \rightarrow K_S^0 K_S^0$, which have historically suffered from lower efficiencies compared to charged final states such as $K^+ K^-$ and $\pi^+ \pi^-$. For example, let's consider the most recent $\Delta\mathcal{A}^{CP}$ and $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurements [1, 2]. The number of signal candidates in the samples used for these measurements is:

$$\begin{aligned} N(D^0 \rightarrow K^+ K^-)/N(D^0 \rightarrow K_S^0 K_S^0) &\sim 6.7 \times 10^3, \\ N(D^0 \rightarrow \pi^+ \pi^-)/N(D^0 \rightarrow K_S^0 K_S^0) &\sim 2 \times 10^3, \end{aligned} \tag{1.44}$$

where N is the event yield available in the indicated channel. The different branching ratios can only partially account for the observed yield difference, since $BR(D^0 \rightarrow K^+ K^-)/BR(D^0 \rightarrow K_S^0 K_S^0, K_S^0 \rightarrow \pi^+ \pi^-) \sim 60$ and $BR(D^0 \rightarrow \pi^+ \pi^-)/BR(D^0 \rightarrow K_S^0 K_S^0, K_S^0 \rightarrow \pi^+ \pi^-) \sim 20$, meaning that the selection efficiency for the $K_S^0 K_S^0$ channel was a *factor 100 lower* than $K^+ K^-$ or $\pi^+ \pi^-$. Moreover, one is led to expect that suppression of the background in the offline analysis should not be the reason for a significant loss of efficiency, as correctly identifying a K_S^0 particle is much easier than identifying a charged kaon or a pion exploiting an offline-level reconstruction, and K_S^0 particles are not more abundant in the background than kaons or pions. Therefore, trigger selection must be addressed for the observed yield difference, as current triggers do not leverage the K_S^0 decay characteristics. This results in limited efficiency for K_S^0 selection at the online level, significantly constraining the overall efficiency for decays involving K_S^0 final states.

Overcoming this limitation is the quest I have been pursuing during the years of my PhD, that I describe in the chapters that follow.

Chapter 2

Importance of long-lived particles in general

In the previous chapter, the significance of the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement was discussed, along with the challenges associated with selecting a K_S^0 final state at a hadron collider. However, that measurement is only one example within an extensive set of high-energy physics measurements, that are crucially dependent on the experimental detection of long-lived particles. Given the importance of that field, before proceeding with the discussion of my specific target, I make a digression here to provide a brief overview of the physics case for LLPs in general. This serves as an additional motivation for the work described in the chapters that follow, showing its potential benefits in a wider perspective.

The majority of high-energy physics (HEP) measurements focus on the properties of particles with short lifetimes. Prominent examples include the W , Z , and Higgs bosons, which decay almost instantaneously after production. Even for beauty and charm hadrons, which have slightly longer lifetimes, their flight distances are typically limited to a few millimeters at most. However, a rich set of interesting processes exists involving states with longer lifetimes, and flight distances ranging from a few millimeters to several meters or more. These particles, commonly referred to as long-lived particles (LLPs), include both well-known states described by the Standard Model and hypothetical new particles predicted by various beyond-Standard Model theories. The study of these particles has a long-standing history in high-energy physics, tracing back to experiments such as those conducted at LEP by the ALEPH [48] and DELPHI [49] collaborations. Significant efforts continue to be dedicated to such measurements by a wide range of contemporary collaborations, including the major LHC experiments. Moreover, the interest in LLP decays extends well into the future, as evidenced by extensive studies exploring potential LLP measurements at upcoming facilities and the development of dedicated experiments, such as MATHUSLA [50] and FASER 2 [51] at the HL-LHC, or FCC-ee [52].

The extensive potential offered by the study of LLP will be demonstrated by reporting some notable examples where these are involved. The list does not intend to serve as a comprehensive report; for a complete overview, please refer to Ref. [53].

2.1 Standard Model LLP

A wide array of intriguing measurements can be performed through the detection of Standard Model LLPs. Among these, the strange hadrons K_S^0 and Λ , commonly referred to as V^0 s, offer exceptional versatility for investigation. These particles are identified as two-track vertices, with displacements of the order of meters, given by lifetimes of 9×10^{-11} s (K_S^0) and 2.6×10^{-10} s (Λ), and boosts larger than 50. It is instructive to consider a few examples of measurements enabled by the detection of these states.

$$D^0 \rightarrow K_S^0 K_S^0$$

The measurement of $\mathcal{A}^{CP}(D^0 \rightarrow K_S^0 K_S^0)$ serves as a clear example of a compelling analysis enabled exclusively by the detection of K_S^0 final states. A detailed discussion of the significance of this measurement has been provided in the previous chapter and will not be reiterated here.

$$B_s^0 \rightarrow K_S^0 K_S^0$$

The measurement of the branching fraction and CP violating parameters for the $B_s^0 \rightarrow K_S^0 K_S^0$ decay is an excellent probe for the presence of NP influences [54–56], showing how also beauty physics can be investigated through the detection of a K_S^0 final state. Both LHCb and Belle collaborations performed a measurement of such parameters in the last decade [57, 58]. It is worth noting how the sample exploited by both collaborations is made of just a few tens of signal candidates (~ 30) despite the markedly different experimental setup, given the small branching ratio of the considered decay. Therefore, any technical improvement that could possibly enhance available sample size would have a major impact on the precision of such measurements.

$$B^0 \rightarrow J/\psi K_S^0$$

The system of neutral beauty mesons includes another particularly compelling decay channel, $B^0 \rightarrow J/\psi K_S^0$. Due to the dominance of tree-level contributions in the determination of the decay amplitude, the measurement of CP -violating parameters in this decay provides a theoretically clean approach to determining the CKM angle β [59, 60]. This characteristic has led to multiple measurements of CP violation in $B^0 \rightarrow J/\psi K_S^0$ by various collaborations, including BaBar [61], Belle [62, 63], Belle II [64], and LHCb [65–67]. Notably, the latest result from LHCb [67] represents the most precise determination of the CKM angle β to date. In this measurement, the detection of particles with very long flight distances plays a critical role, as approximately 60% of the sample involves a K_S^0 with a flight distance on the order of meters, escaping the LHCb vertex detector. The ability to reconstruct such highly displaced decays is essential for maximizing the available statistics, without which the achievable precision would be severely limited.

$$K_S^0 \rightarrow \mu^- \mu^+$$

The measurement of the branching fraction for the $K_S^0 \rightarrow \mu^- \mu^+$ decay provides a stringent null test of the SM. The SM predicts an extremely small branching fraction of approximately

5×10^{-12} [68]. Therefore, any BSM influence that enhances this branching fraction to a detectable level would serve as a clear indication of new physics.

The most recent $BR(K_S^0 \rightarrow \mu^+ \mu^-)$ measurements have been performed by the LHCb collaboration [69–71], exploiting full statistics collected during Run 1 and Run 2 of the LHC. The combined result set the upper limit that is $BR(K_S^0 \rightarrow \mu^+ \mu^-) < 2.1 \times 10^{-10}$ @90% CL. In the latest analysis, the detection of a displaced vertex is necessary for the selection of both signal and calibration sample, as $BR(K_S^0 \rightarrow \mu^+ \mu^-)$ has been normalized to the number of observed $K_S^0 \rightarrow \pi^+ \pi^-$ decays. The measurement is still statistically limited, despite the huge production cross-section of K_S^0 at the LHC. Therefore, in order to set more stringent limits, the ability of the system to collect these decays must be improved.

$J/\psi \rightarrow \Lambda \bar{\Lambda}$

The study of Λ baryons also offers valuable opportunities for interesting measurements. One notable example is the determination of the branching fraction for the decay $J/\psi \rightarrow \Lambda \bar{\Lambda}$. In this process, the two final-state baryons are produced via the annihilation of a $c\bar{c}$ pair, either through a virtual photon or the emission of three gluons. The dynamics of this process, combined with the complexities of QCD effects, make an accurate theoretical prediction for the branching fraction $BR(J/\psi \rightarrow \Lambda \bar{\Lambda})$ particularly challenging, as reflected in the wide range of existing predictions [72–74]. In this context, precise experimental measurements of this parameter are crucial, as they can provide essential inputs for theory, enabling a deeper understanding of the underlying effects.

Consequently, several collaborations have performed measurements of this decay, including MARKII [75], DM2 [76], BES [77, 78], BESII [79, 80], and BaBar [81]. Despite the substantial body of results, the world average for the $J/\psi \rightarrow \Lambda \bar{\Lambda}$ branching fraction has only recently reached a precision at the 10% level [82]. This motivates continued efforts to improve the measurement precision, and LHCb is well-positioned to make significant contributions in this area.

$\Lambda_c^+ \rightarrow \Lambda \pi^+$

A precise determination of the SCS $\Lambda_c^+ \rightarrow \Lambda \pi^+$ decays branching fraction can provide a crucial input for theory as well. In particular, it can help shed light on the heavy-to-light baryonic transitions models describing the considered Λ_c^+ decay [83–85]. The most recent $BR(\Lambda_c^+ \rightarrow \Lambda \pi^+)$ measurement has been performed by the BESIII collaboration with the result published in 2022 [86], whose result is $BR(\Lambda_c^+ \rightarrow \Lambda \pi^+) = (1.31 \pm 0.08(stat) \pm 0.05(syst))\%$.

2.2 Beyond Standard Model LLP

Several particles predicted by beyond SM (BSM) theories are characterized by a long lifetime, either because of a small coupling between the BSM and detectable SM particles where the new state would decay, or a limited phase space for the same state. Because of this, a significant number of searches for BSM states require the detection of LLPs, with different masses, final states, and decay topologies.

Supersymmetry

Supersymmetry (SUSY) is one of the most compelling BSM theories of the last decades [87–90]. It predicts an entire set of SM states’ supersymmetric partners, typically heavier than the SM ones. The potential coincidence between these and dark matter particle candidates makes their search even more appealing.

The LHC serves as an optimal facility for detecting evidence of long-lived SUSY particles. Its exceptionally high instantaneous luminosity and center-of-mass collision energy significantly enhance the likelihood of producing such states compared to other accelerators. Consequently, numerous searches have been conducted by various LHC experiments. For example, the ATLAS collaboration explored various decay signatures and possible final states for long-lived SUSY particles, as decays into a final state composed of a fully-hadronic displaced vertex associated with a muon [91], or the appearance of a single high-impact-parameter lepton in the final state [92]. Another interesting phenomenology investigated by the ATLAS collaboration is the production of particles with a very long lifetime that - produced in the collisions - come to rest in the detector material. This decay signature is the presence of a displaced vertex when collisions are absent [93]. So-called disappearing tracks can be another manifestation of SUSY LLP decays. These are identified as charged tracks that leave a minimum of four hits within the pixel detector, yet they lack any association with further hits in the outer tracker region or with energy deposits in the calorimeters [94]. It is noteworthy that all of the aforementioned examples are characterized by the presence of displaced charged tracks in their final states.

Heavy neutral leptons

In the search for BSM particles, neutrinos gather significant attention, as they have already shown laboratory evidence for a BSM phenomenon, with the observation of their oscillatory nature between different flavors [95–100]. The Seesaw mechanism is one of the most credited theoretical interpretations describing this phenomenon [101–105]. This SM extension predicts the existence of a number of electroweak-singlet (“sterile”) neutrinos, with masses ranging from MeV/c^2 to TeV/c^2 scale. In the latter case, these are usually referred to as heavy neutral leptons (HNLs).

The observation of BSM phenomena in neutrino behavior has provided strong motivation for the search for their heavier counterparts. Several of these have been conducted at LHC experiments, as this facility is again representing an extremely suitable place for the search for such states. For example, the CMS collaboration recently published a search for HNLs performed exploiting the full Run 2 statistics of 138 fb^{-1} [106]. The investigated final state is composed of three leptons: one compatible with coming from the primary pp interaction, the other two forming a displaced secondary vertex within the CMS tracking system. This work successfully set a limit on the coupling between muon-neutrinos and HNLs with masses in the $1 - 20 \text{ GeV}$ range, both for Dirac and Majorana hypotheses.

This is just an example of the various searches performed for HNLs manifestations. A comprehensive review for these can be found in Ref. [107].

Higgs portal

An alternative series of investigations aimed at discovering BSM particles examines the potential for their production at the LHC via the decay process of a Higgs boson.

This hypothesis is permitted by the prevailing understanding of the Higgs boson’s total branching fraction, as existing experimental constraints continue to permit decay into BSM particles with branching ratios reaching up to 30% [108]. An example for these searches is the result published by LHCb collaboration, where a search is performed for a *hidden-valley* pion π_ν . This represents the lightest state of the “hidden sector”, accessible through the decay of a SM-like Higgs [109]. The investigated final state is represented by a pair of hadronic jets originating from a displaced vertex. The analysis is performed exploiting the integrated luminosity of 2.0 fb^{-1} collected during Run 1 of the LHC. No evidence for π_ν -like particles has been found, and limits are set for particles with a mass between 25 and 50 GeV, and a lifetime between 2 and 500 ps. The presented LHCb result represents the latest investigation within the scope of analogous searches undertaken by other collaborations, including the D0 [110], CDF [111], ATLAS [112–114], and CMS [115] collaborations.

A multitude of theoretical predictions and experimental explorations have been conducted on various other scenarios, including those involving dark photons [116–119] and axion-like particles [120, 121]. While a detailed account of these specific investigations is omitted here for reasons of brevity, it should nonetheless be evident that the diversity and significance of potential beyond Standard Model long-lived particle (BSM LLP) scenarios are thoroughly recognized and understood within the field.

Chapter 3

LHCb and long-lived particles reconstruction

In this chapter I briefly describe the LHCb detector and trigger system, for both Run 2 and Run 3 configurations. The description will mainly be focused on the aspects regarding the detection and online selection of LLPs.

Having presented the broad physics case offered by the detection of LLPs, it is now interesting to consider the most suitable experimental setups for an effective study of such decays.

3.1 The Large Hadron Collider

LHCb is hosted at the CERN (*Conseil Européen pour la Recherche Nucléaire*) laboratories, and it is located in one at the collision points along the ring of the Large Hadron Collider (LHC) [122]. The LHC is a circular hadron collider with a ring length of around 27 km, the largest of CERN's accelerators. It can produce collisions between both protons and heavy ions. Beams circulate along the ring in opposite directions, colliding at four interaction points where the major experiments are placed, hence ALICE, ATLAS, CMS and LHCb. Hadrons are kept in a stable orbit by superconducting NbTi dipole magnets, capable of providing two opposite-direction magnetic fields up to 8 T within the same yoke. Protons do not circulate continuously in the accelerator, but are organized in *bunches* of around 10^{11} protons, that are time-spaced for a multiple of 25 ns, determining a bunch-crossing rate of 40 MHz. However, some empty bunches are present for technical reasons, leading to an actual maximum non-empty average bunch-crossing rate of 30 MHz. LHC is designed to be able to produce proton-proton collisions at a center-of-mass energy up to 14 TeV, at a maximum instantaneous luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. However, it has been set to 7 TeV during the 2010–2011 data taking, 8 TeV during 2012 and 13 TeV during 2015–2018.

The LHC provides an ideal environment for the study of long-lived particles, as evidenced by the numerous searches and measurements in which LHC experiments have played a leading role mentioned in Chapter 2. This pivotal role is given by the combination of the very large cross-section offered by the collision of protons, together with the high instantaneous luminosity and multi-TeV center of mass energy at which collisions take place. These allow for the collection of large samples in a relatively short timescale, for

a wide range of phenomena. Additionally, some of the most advanced detectors in the world are hosted along the LHC ring. These exploit a wide range of detection geometry and techniques, allowing the achievement of unprecedented precision measurements.

The long-term research perspective offered by CERN is an additional key factor that makes the LHC an ideal environment for the study of long-lived particles. The LHC is expected to continue providing collisions for at least another 20 years, with future facilities already under evaluation by CERN. Among these, the Future Circular Collider (FCC) stands out as the most promising, as outlined in the 2020 update by the European Strategy for Particle Physics [123].

3.1.1 LLP decays triggering

Typical LLP decay signatures include the presence of significantly displaced tracks and vertices, often accompanied by relatively low momentum and transverse momentum, as highlighted in examples of Chapter 2. To leverage these signatures at the trigger level in the identification of signal decays, at least a partial event reconstruction is necessary during online processing, as reconstruction of tracks and vertices, identification of the primary pp vertices, and determination of track displacements. These tasks are computationally intensive and are typically executed only in the later stages of the trigger system, where more processing time is available. As a result, implementing LLP-dedicated selections at the early trigger stages has traditionally been infeasible, leading to limited trigger efficiency for LLP events. Early trigger stages predominantly rely on high- p_T tracks and muons as their primary selection criteria, which are not optimal for capturing LLP decays. To enhance trigger efficiency for LLPs, it is essential to advance reconstruction tasks to earlier stages of the trigger system. This shift allows LLP decay signatures to be effectively utilized, enabling the development and implementation of dedicated trigger strategies tailored to LLP events. The only LHC experiment taking a significant step forward in this kind of approach within the LHC landscape is LHCb, with its Run 3 upgrade. It is now able to reconstruct tracks and vertices at the full non-empty bunch-crossing collision rate of 30 MHz (more details will be given later), making it the perfect testbed to develop and test novel trigger techniques aimed at early LLP selection.

I will provide an overview of the LHCb detector and data acquisition systems in both the Run 2 and Run 3 configurations.

3.2 LHCb configuration in Run 2

The LHCb detector [124] is a single-arm forward spectrometer designed for the study of hadrons containing beauty and charm quarks. Its forward angular coverage spans from 10 mrad to 300 mrad (250 mrad in the non-bending plane), corresponding to a pseudorapidity range of $1.8 < \eta < 4.9$, where η is defined as $\eta = -\log[\tan(\theta/2)]$ and θ is the polar angle relative to the beam direction. This forward configuration enables LHCb acceptance to encompass a region where 27% of produced b quarks and 24% of $b\bar{b}$ pairs from pp collisions are located, despite covering only about 4% of the solid angle, as illustrated in Figure 3.1. Heavy-flavor hadrons produced in these collisions typically have momenta on the order of tens of GeV/c. Combined with their lifetimes ranging from 0.1 to 1 ps, this results in flight distances of up to several millimeters. The LHCb detector has been designed




Figure 3.1: Distribution for the production cross-section of $b\bar{b}$ pairs at a center-of-mass energy of 14 TeV, as a function of the polar angle of the produced particles. Red area represents the LHCb acceptance. Plot taken from [125].

to leverage these displacements and the specific invariant mass of the decaying hadron to identify signal decays among tracks originating from the primary pp interaction. The general structure of the detector is represented in Figure 3.2.

3.2.1 LHCb tracking system

The sub-detector group aimed at the reconstruction of particle trajectories (referred to as *tracks*) is composed of:

- the *vertex locator* (VELO) - a silicon-strip detector surrounding the pp interaction zone, designed to determine the positions of both the primary vertices where proton-proton interactions occur and the secondary vertices associated with the decay of heavy-flavored hadrons. It has an acceptance that spans the z -direction range from -17.5 cm to 75 cm relative to the nominal interaction point [127];
- the *tracker turicensis* (TT) - a silicon strip tracking detector placed upstream the magnet. It is positioned at a distance of 230 cm [128] from the primary interaction point along the z direction. One of the main purposes of the TT is the detection of particles escaping the VELO detector acceptance, mainly K_S^0 and Λ ;
- the LHCb magnet - a room-temperature magnet producing a dipole vertical magnetic field with integrated bending power of 4 Tm;
- the three tracking stations (T1 - T3) placed downstream the magnet.

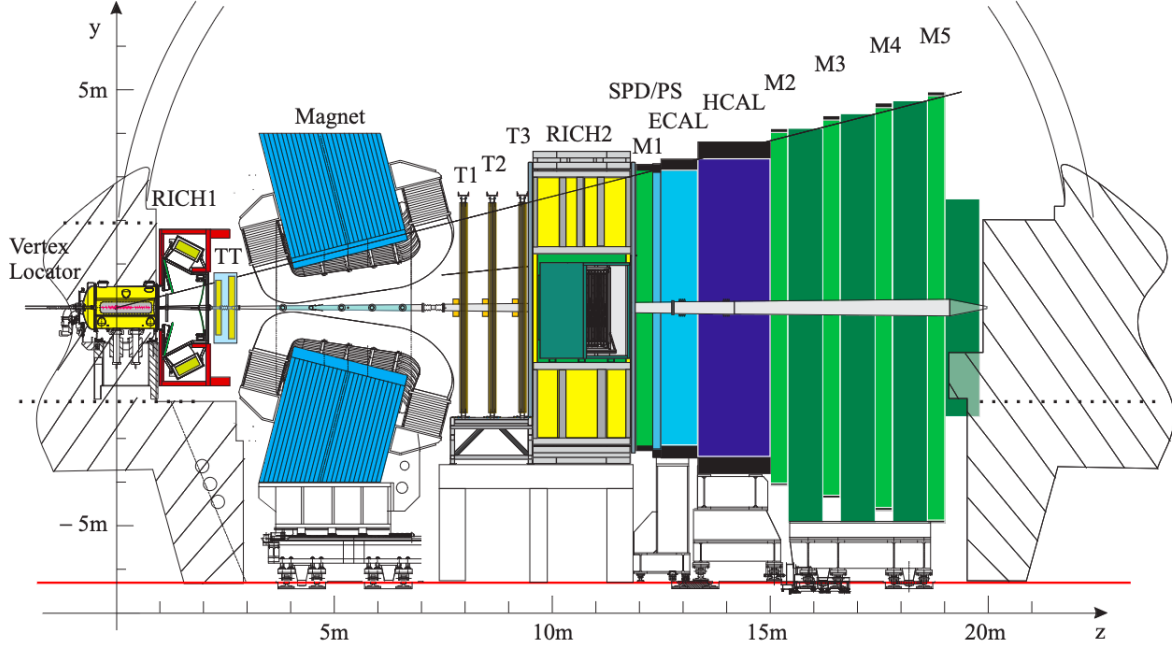


Figure 3.2: A schematic representation of the Run 2 LHCb detector configuration. Taken from [126].

LHCb tracks reconstruction

Having described the structure of the LHCb tracking system, it is useful to describe how charged particles are detected when passing through the detector. This understanding is particularly important in the context of LLP decay detection.

Five different track categories are defined, depending on which sub-detector the track crossed, producing hits. These are represented in Figure 3.3, and are:

- *Long* tracks - these tracks are reconstructed exploiting entire tracking system. Tracks are based on VELO and T-stations information and TT hits are added when available;
- *Downstream* tracks - these tracks are reconstructed using only TT and T-stations hits;
- *Upstream* tracks - these tracks are reconstructed exploiting VELO and TT hits only. Upstream tracks are usually associated with low-momentum particles swept away by the magnetic field before reaching the T-stations;
- *VELO* tracks - these tracks are reconstructed exploiting only VELO hits. They are used for the reconstruction of primary pp interaction vertex, and as seeds for the reconstruction of Long and Upstream tracks;
- *T-tracks* - these tracks are reconstructed exploiting only T-stations hits.

It is important to highlight the differing longitudinal acceptance of tracks, as this plays a crucial role in the reconstruction of long-lived particles. A track can be reconstructed in the VELO only if it produces at least three hits in the VELO layers. Consequently,




Figure 3.3: A visual representation of the five different LHCb charged particles reconstruction categories. Taken from [126].

only particles decaying within approximately 70 cm of the primary interaction point can be reconstructed as long tracks, while a significant fraction of LLP decays can only be reconstructed as downstream tracks. For instance, the K_S^0 candidates analyzed in the Run 3 $\mathcal{A}^{CP}(D^0 \rightarrow K_S^0 K_S^0)$ measurement discussed in Chapter 8, have an average momentum of $p \sim 40$ GeV/ c and an average flight distance of 2 meters. As a result, approximately two-thirds of these decays can only be reconstructed using downstream tracks.

VELO

The Vertex Locator (VELO) [127] is a silicon-strip tracking detector positioned around the pp interaction point, designed to reconstruct particle trajectories and identify both the pp primary interaction vertices and the secondary vertices associated with b - and c -hadron decays. It consists of 42 semicircular silicon modules, each containing two overlapping sensors that cover a radial distance from 8.2 mm to 41.9 mm relative to the beamline. Each module is equipped with two distinct sensor types, R and ϕ , providing complementary information. The structure of a module pair is illustrated in Figure 3.4. The R sensors are designed for the measurement of the radial distance from the beam using semicircular-shaped strips. Their pitch is not uniform, going from 38 μm to 102 μm , linearly increasing with the radial distance from the beam axis. The ϕ sensors provide measurement of the azimuthal angle, exploiting radially-oriented strips, again exploiting a non-uniform pitch arrangement. Each ϕ sensor is divided into two regions, inner and outer. The border between the two is placed at a distance of 17.25 mm from the center. The strip pitch increases linearly from 38 to 78 μm and from 39 to 97 μm , for the inner and outer region, respectively. The strips have an angle of approximately 20° and 10° to the radial direction for the inner and outer regions, respectively.

Modules are arranged along the beam direction, covering both the upstream and downstream regions of the nominal interaction point, as shown in Figure 3.5. Two VELO halves are installed on retractable supports. This allows placing the detector at a distance of 8.2 mm from the interaction point in closed configuration, while ensuring the detector's safety during LHC beam injection, when retracted. Achieving the closest feasible distance from the beamline is important to maximize the vertex resolution, which is proportional to the distance from the interaction region and the sensitive area. No beam pipe is present




Figure 3.4: A schematic representation of structure of each VELO station, composed by two different modules. Taken from [127].




Figure 3.5: A schematic representation of modules arrangement in the VELO. Taken from [127].

within the VELO region and the VELO vessel is kept separate from the beam vacuum through a $300\ \mu\text{m}$ [129] aluminum foil, usually referred to as RF-foil. This shields the VELO from electromagnetic effects induced by the high-frequency circulating beams. It has a corrugated structure, as shown in Figure 3.6, allowing for an overlap between modules when VELO is positioned in closed configuration. The VELO raw hit resolution




Figure 3.6: A schematic representation of the RF-foil structure, surrounding the VELO detector. Taken from [127].

is in the $10 - 25 \mu\text{m}$ range.

The VELO, with the described design, achieves a primary vertex resolution of $13 \mu\text{m}$ in the transverse plane and $71 \mu\text{m}$ along the beam direction, along with an impact parameter¹ resolution of $(15 + 29/p_T[\text{GeV}/c]) \mu\text{m}$. These performances enable LHCb to effectively distinguish heavy-flavored hadron decays from the abundant background processes produced in pp collisions. However, only long tracks can leverage the resolution offered by the VELO detector, and in case of LLPs these can be a small fraction of the total, as reported in Section 3.2.1.

The dipole magnet

LHCb uses a room-temperature dipole magnet for the measurement of particle momentum, through their bending angles. The magnet is made of two coils, in which a current of 5.85 kA circulates in nominal conditions, dissipating 4.2 MW of power. The produced magnetic field's main component is along the y direction. It reaches an intensity of 1.1 T, with an overall bending power of 4 Tm. The x and z components of the magnetic field are negligible and $x - y$ can be considered the actual bending plane with good approximation. A schematic representation of the magnet, together with the plot showing the field intensity as a function of z , can be found in Figure 3.7. The y component of the magnetic field is measured with Hall probes before the start of data-taking periods, with a relative precision of 4×10^{-4} . This is important to establish a good momentum and invariant mass scale calibration.

Because of the LHCb detector and magnet configuration, opposite sign particles have very different acceptance, since they are swept in opposite directions by the dipole

¹The impact parameter is defined as the distance of closest approach between the particle trajectory and the considered vertex; see Section 4.2.




Figure 3.7: Representation of the LHCb dipole magnet (left) and visualization of the y component of the magnetic field. Taken from [130].

magnetic field. In order to mitigate the effect of this acceptance bias on the charge detection asymmetry, the magnetic field is periodically reversed, defining two different field configurations (MagUp and MagDown). When data samples collected with the two different configurations of roughly similar size are combined, most of the charge asymmetry is expected to cancel out. While this helps reduce the overall charge asymmetries, analyses do not rely solely on this cancellation to mitigate detection asymmetries. A common practice in LHCb analyses is to perform the measurement separately on data collected with each configuration and verify the consistency of the results between the two.

Tracker Turicensis

The Tracker Turicensis (TT) is a silicon microstrip sensors located before the magnet, covering the full LHCb acceptance. It is placed 230 cm after the nominal interaction point along the z direction, and its main purpose is to allow the detection of both low-momentum particles, swept away from LHCb acceptance by magnetic field, and longer-lived particles (like K_S^0 and Λ) that don't produce hits in the VELO because of their large flight distance. It is composed of four different layers, grouped in two stations, separated by 30 cm. Layers are arranged in a “ x - u - v - x ” pattern, where x layers have vertical strips, and u and v ones are rotated by $\pm 5^\circ$, respectively. This configuration allows for a bi-dimensional measurement of hit positions, avoiding at the same time the ambiguity that would arise in case horizontal orientation would be exploited for these layers. A schematic representation of TT can be found in Figure 3.8. Downstream tracks, reconstructed with TT and T-stations only, are crucial to enhance the acceptance for long-lived decays, but offer a lower resolution with respect to long tracks, that can also leverage VELO hits. In particular, mass resolution is worse by a factor of 2 for downstream tracks with respect to long ($3.5 \text{ MeV}/c^2$ vs $7 \text{ MeV}/c^2$ [132]), and it can also be observed for the K_S^0 candidates of LHCb Run 2 $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement reported in Chapter 4, whose invariant mass for the long and downstream-reconstructed case is shown in Figure 3.9.




Figure 3.8: A schematic representation of the Tracker Turicensis and its layers arrangement. Taken from [131].




Figure 3.9: Invariant mass distribution for K_S^0 candidates reconstructed with long (left) and downstream (right) tracks exploited in the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ Run 2 measurement [2].

T-stations

Three T-stations, each consisting of four layers, are positioned approximately 7.5 meters downstream of the nominal interaction point along the z direction, after the magnet. These stations enable the measurement of particle trajectories by detecting the curvature of particles passing through the magnetic field. The T-stations cover the full LHCb angular acceptance with a dimension of 6 meters horizontally and 4.9 meters vertically. Two different acceptance regions are identified in the T-stations, covering the inner and outer regions. These are respectively identified as the one exceeding and not exceeding a 10% occupancy at the nominal Run 2 luminosity of $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$. As a result, the two regions are instrumented differently. This choice significantly lowers the cost of the detector placed downstream the magnet, without suffering a significant hit-resolution loss.

Inner Tracker

The Inner Tracker (IT) covers the innermost region of T-stations. The same $x-u-v-x$ layers arrangement and microstrip technology of the TT is exploited. A schematic representation of the IT can be found in Figure 3.10.




Figure 3.10: A schematic representation of the inner tracker. Taken from [131].

Outer tracker

The outer tracker (OT) [133] covers the outer region of the three T-stations, outside the IT acceptance. It is a gaseous ionisation detector, composed of straw tubes operating as proportional counters. Each drift tube is 2.4 m long, with an inner diameter of 4.9 mm. Tubes are arranged in layers of a single T-station with the same $x-u-v-x$ configuration of TT and IT. A schematic representation of the OT is reported in Figure 3.11

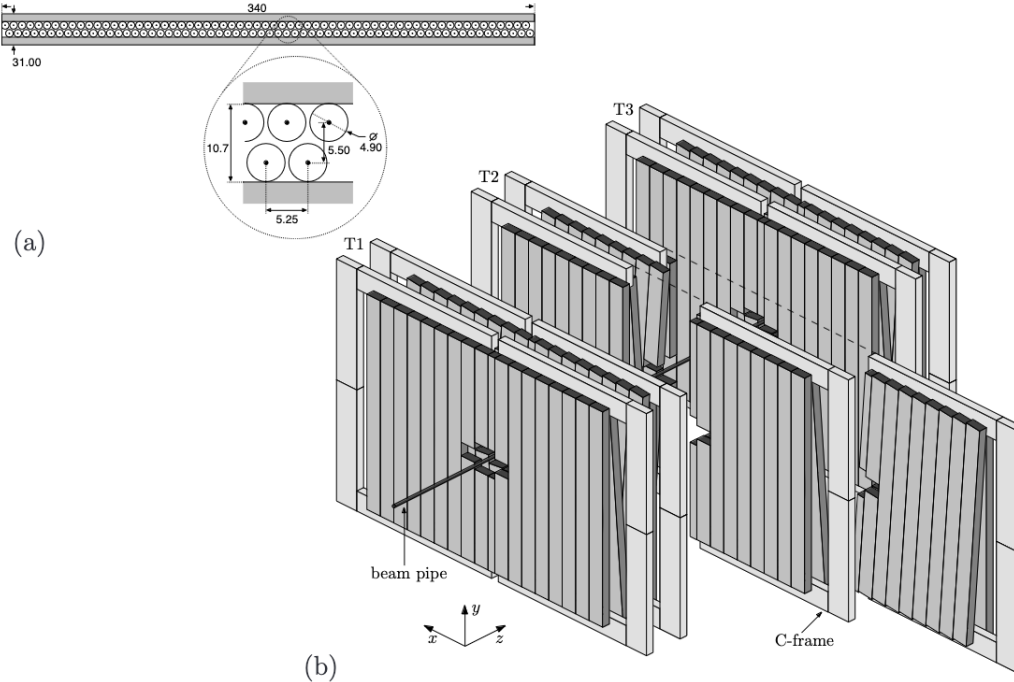


Figure 3.11: A schematic representation of the outer tracker. Taken from [134].

3.2.2 LHCb particle identification system

The particle identification system (PID) enables the distinction of hadron species. It is composed of:

- two *ring imaging Cherenkov* detectors (RICH1 and RICH2) - these allow the identification of hadron species by exploiting the Cherenkov radiation they emit. The RICH1 and RICH2 provide identification for particles in the $1 - 60 \text{ GeV}/c$ and $15 - 100 \text{ GeV}/c$ momentum regimes, respectively;
- the *scintillating pad* and the *preshower* detectors (SPD/PS) - these are employed for the discrimination of electrons from photons and hadrons, respectively;
- the *electromagnetic calorimeter* (ECAL) - it is exploited for electrons and photons identification, together with energy measurement of the latter;
- the *hadronic calorimeter* (HCAL) - it provides the measurement of hadrons energy deposit;
- the five *muon stations* (M1-M5) - they are composed by alternating layers of iron and multi-wire proportional chambers and are used to identify muons.

Cherenkov detectors

Two ring-imaging Cherenkov detectors (RICH1 and RICH2) are used to discriminate between different hadron species in LHCb [135] by measuring the radius of the emitted Cherenkov radiation.

The RICH1 is placed between the VELO and the TT. It covers the full LHCb angular acceptance and it is designed to provide PID information for particles within a momentum range of $1 - 60 \text{ GeV}/c$. The RICH2 is placed farther from the interaction point, after the last T-station. It is designed to provide PID information for particles with larger momentum, in the $15 - 100 \text{ GeV}/c$ range. It covers only partially the LHCb acceptance, corresponding to around $[15, 120]$ mrad in the horizontal plane and up to 100 mrad in the vertical plane. A schematic view of RICH1 and RICH2 is reported in Figure 3.12. With the adopted configuration the two sub-detectors are able to provide particle identification in a wide momentum range ($1 - 100 \text{ GeV}/c$) as can be seen from the plot reported in Figure 3.13. In both detectors, the Cherenkov light is detected by a lattice of hybrid photon detectors (HPDs), placed outside of the LHCb acceptance to avoid increasing the material budget. A system of flat and spherical mirrors is used to direct the Cherenkov light toward the HPDs, see Figure 3.12.

Calorimetric system

The LHCb calorimetric system is composed of four different modules, placed between the first and second muon stations, covering the angular acceptance from 25-300 (250) mrad in the bending (non-bending) plane. The calorimetric modules are:

- Scintillator Pad Detector (SPD) and PreShower detector (PS) - a pair of polystyrene-based scintillating planes whose purpose is to distinguish between photons and charged hadrons, for SPD, and electromagnetic showers from hadronic ones, in case of PS;

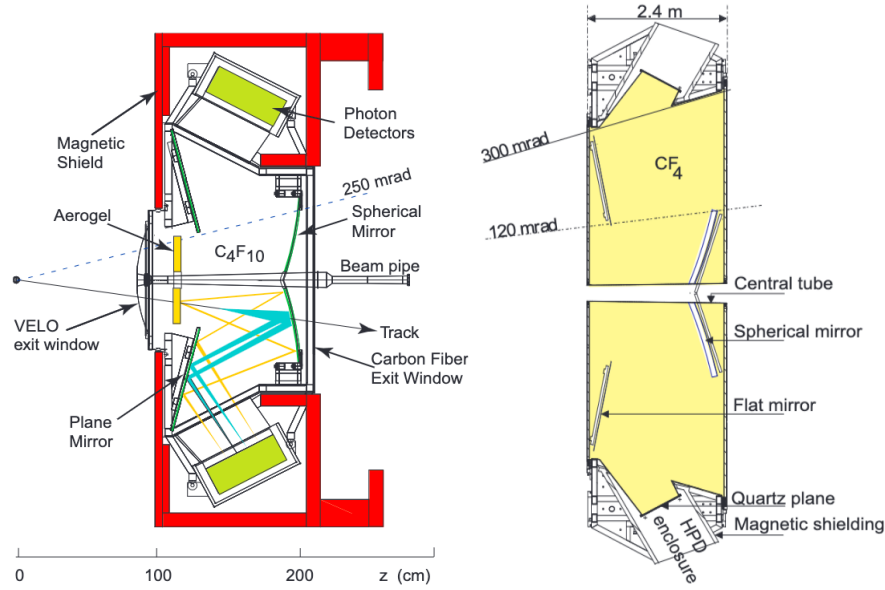


Figure 3.12: A schematic representation of RICH1 (left) and RICH2 (right) detectors. Taken from [130].




Figure 3.13: Cherenkov angle as a function of particle momentum for isolated tracks (tracks whose Cherenkov ring does not overlap with any other ring) in the C4F10 radiator. Taken from [136].

- Electromagnetic calorimeter (ECAL) - it is made up of 66 scintillator-lead alternating layers, whose thickness is of 4 mm and 2 mm, respectively. It is exploited for the measurement of electromagnetic showers energy. Achieved resolution is $\sigma(E)/E = 1\% \oplus 10\%/\sqrt{E/\text{GeV}}$ [137];
- Hadronic calorimeter (HCAL) - placed farther from the interaction point, consists of scintillator-iron alternating layers, of thickness 4 mm / 16 mm, respectively. The HCAL total thickness limited by the available space inside the cavern. This condition makes it not sufficient for the total containment of produced hadronic showers. The achieved energy resolution is $\sigma(E)/E = 9\% \oplus 69\%/\sqrt{E/\text{GeV}}$ [137]. The HCAL is mostly used for trigger purposes.

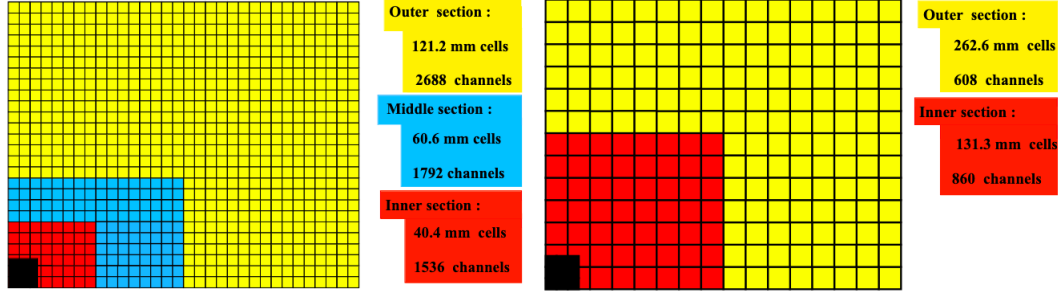


Figure 3.14: Representation of the SPD, PS and ECAL (left) and HCAL (right) segmentation scheme. Just one quadrant is reported. Black square represent the beam pipe region. Taken from [130].

Muon system

The LHCb muon system consists of five stations (M1–M5), designed to differentiate muons from hadrons while also enabling a rapid measurement of the muons’ transverse momentum with a resolution of approximately $\Delta p_T/p_T = 20\%$. This information is exploited for triggering. The first station (M1) is placed upstream of the calorimeters, while M2–M5 are placed downstream of it. This configuration allows for a better momentum resolution, since M1 is less affected by multiple scattering effects. M2–M5 are interleaved by 80 cm iron absorbers, in order to stop hadrons not contained in the HCAL. A minimum of 6 GeV/c is required for a muon to traverse the entire system. Muon stations cover the angular acceptance from 20 (16) to 306 (258) mrad in the bending (non-bending) plane. Each station is split in four quadrants, and each quadrant in four concentric regions (R1–R4), where R1 is the closest to the beam pipe.

3.2.3 Run 2 trigger configuration

The trigger system employed by LHCb during Run 2 was specifically designed for selecting and identifying heavy-flavor decays with a high efficiency [138]. It processes pp collision data produced by the LHC at an initial rate of 40 MHz (or approximately 30 MHz if only non-empty bunch crossings are considered). The needed rejection factor is obtained through a two-stage sequential trigger system that progressively reduces the event rate. The trigger scheme employed in Run 2 is shown in Figure 3.15.

The Level-0 trigger

The first stage of the LHCb trigger system is the Level-0 (L0) trigger, tuned to reduce the input event rate to 1.1 MHz, corresponding to the maximum readout frequency of the full detector. Operating synchronously with the LHC bunch-crossing rate of 40 MHz, the L0 must deliver its decision within a fixed latency of 4 μ s. To achieve this, L0 is implemented in custom-built FPGA boards optimized for low-latency processing. The trigger leverages information from the detectors capable of high-frequency readout, specifically the calorimeters and the muon system. L0 selections are based on three main signatures: *L0 muon*, *L0 calorimeter*, and *L0 hadron*.

The *L0 muon* trigger is the L0 thread aimed at identifying events containing muons in the final state. This is motivated by the fact that muons are a distinctive signature

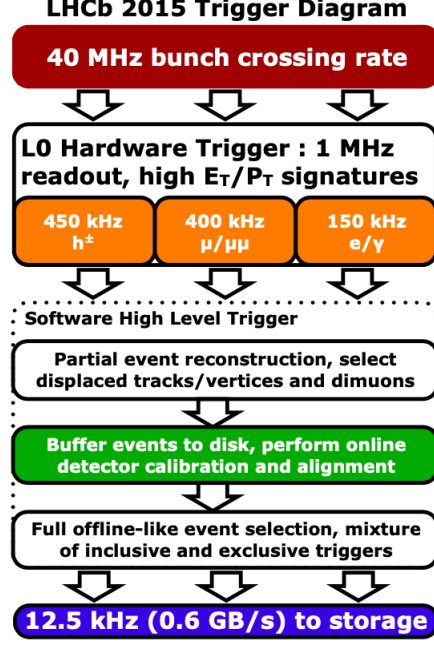


Figure 3.15: The LHCb trigger structure for Run 2 data taking.

of heavy-flavor decays, particularly helpful in hadronic collisions, as they are frequently produced in the weak decays of heavy quarks. The L0 *muon* decision is based exclusively on data from the five muon stations, activated when one of two criteria is satisfied: either the presence of a single muon with transverse momentum (p_T) exceeding a defined threshold (typically $1.5 \text{ GeV}/c$), or when the geometric mean of the transverse momenta of two muons, $\sqrt{p_T(\mu_1)p_T(\mu_2)}$, exceeds a predefined threshold (typically $1.3 \text{ GeV}/c$).

The detection of electrons, high-energy photons, and hadrons is also utilized at the L0 level to identify heavy-flavored hadron decays. The L0 calorimeter and L0 hadron triggers are specifically designed to respond to these signatures. A coarse estimate of particle transverse momentum (p_T) is employed in their decision-making process, leveraging the projective design of the calorimeter system. In this design, the calorimeter cells in different layers cover the same solid angle with respect to the nominal collision point, enabling the use of energy deposit positions to rapidly estimate transverse energy (E_T) for hadrons, photons, and electrons. Energy deposits are grouped into clusters of 2×2 cells, that typically fully contain the energy shower from a single particle while minimizing contributions from nearby particles. The transverse energy is then computed as:

$$E_T = \sum_{i=0}^{i=3} E_i \sin(\theta_i) \quad (3.1)$$

where E_i represents the energy recorded in the i -th cell, and θ_i is the angle between the beam direction and the line connecting the nominal interaction point to the center of the cell. The energy deposit pattern in different calorimeters allows a partial particle identification. Electrons and photons are linked to an ECAL energy deposit respectively with and without the presence of a SPD deposit. Hadrons are linked to a HCAL energy deposit, not matched to any ECAL one. This allows the application of separate thresholds for the three species. In particular, taking into account only the most energetic electron,

photon, and hadron, the event is retained if their transverse energy is respectively larger than 3.7, 2.4 and 2.8 GeV.

The L0 thread typically exploited to trigger most of the $D^0 \rightarrow K_S^0 K_S^0$ candidates is L0 hadron, activated by the energy deposit of final state pions (this is the case for the decays used for the measurement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$ using Run 2 data of Chapter 4).

High Level Trigger

The second trigger stage is the software-based High-Level Trigger (HLT). It runs on a dedicated cluster of CPUs made of 1.7k nodes and equipped with a total of 27k physical cores. The HLT processes the events selected by L0, performing more detailed event reconstruction and applying channel-specific selections. This stage reduces the event rate further, from 1.1 MHz to 12.5 kHz, which matches the maximum rate at which data can be written to permanent storage. This rate reduction is achieved with the execution of two subsequent steps, HLT1 and HLT2, both implemented as a C++ executable and run on the farm nodes.

HLT1 reduces the event rate from 1.1 MHz to 30 kHz, performing a partial event reconstruction. VELO tracks are fully reconstructed and exploited for the identification of primary vertices. Only long tracks are reconstructed at this stage, applying minimal displacement and transverse momentum thresholds, in order to keep combinatorics and processing time under control. Selection algorithms (referred to as *trigger lines*) are then run on the reconstructed objects. These are based on the presence of at least one good-quality, large p_T and large displacement reconstructed track or muon. Charmed hadrons decaying into hadronic final states are typically selected based on the presence of one, or two, high- p_T displaced tracks. When at least one of the trigger lines has a positive response, the event is retained and passed to HLT2.

HLT2 represents the final step in the LHCb trigger system, reducing the event rate to the 12.5 kHz that can be written to permanent storage. At this stage, an offline-like event reconstruction is performed, including entire tracks reconstruction and full PID information. The advanced reconstruction capabilities of HLT2 enable the application of specific decay selections. A few hundred trigger lines are implemented, combining inclusive criteria with exclusive, channel-specific selections tailored to particular decay modes. The HLT2 output event rate is constrained not by the expected signal rate, but by the available bandwidth. This creates a bottleneck for certain types of events, such as charm decays, that occur at significantly higher rates compared to beauty decays. To address this limitation, LHCb upgraded both the computing power and disk space of the event filter farm during 2013–14. This enhancement allowed HLT2 to achieve offline-level precision in its reconstruction, enabling direct physics analyses on trigger-reconstructed objects without requiring further offline event reconstruction. This advancement was paired with the development of a novel data-taking paradigm known as *Turbo* [139, 140]. In this approach, only the reconstructed objects directly associated with the candidate decay selected by the trigger are saved for offline analysis, while the remaining event data are discarded. This paradigm shift, from an event-oriented data collection to a candidate-oriented approach, led to a reduction of the size of stored events by nearly an order of magnitude, allowing for a significant increase in the event rate that could be saved to disk.

In this context, it is important to underline that HLT1 did not reconstruct any particles

decaying outside the acceptance of the VELO detector, such as downstream tracks. This restriction was primarily due to the limited time and processing power available. This configuration, although well-suited for selecting short-lived heavy hadron decays through inclusive selection strategies, imposes a significant limitation on the online selection efficiency for LLPs. Given that the reconstruction of downstream tracks is only performed in the final trigger stage (HLT2), any LLPs with a flight distance large enough to not leave enough hits in the VELO detector cannot be directly triggered at L0 nor HLT1. Those decays can only be retained if the same event also contains some long tracks satisfying some of the HLT1 trigger criteria; this clearly results in a significantly limited efficiency.

3.3 The LHCb Upgrade I

The LHCb experiment recorded a dataset of 9 fb^{-1} during Run 1 and Run 2 of the LHC, which corresponds to several million decays for certain channels, such as $D^0 \rightarrow K^+ K^-$ and $D^0 \rightarrow \pi^+ \pi^-$ [1]. Despite the availability of such large datasets, most physics measurements conducted so far remain statistically limited, highlighting the importance of collecting larger datasets to achieve higher sensitivities. However, acquiring the necessary statistics to enable significant precision improvements within a reasonable timescale requires an increase in the data-taking rate compared to Run 1 and 2. To meet this goal, the LHCb collaboration decided to raise the instantaneous luminosity during Run 3 from the $4 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ used in Run 2 to $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. However, taking data at a higher instantaneous luminosity while exploiting the same Run 2 system would not lead to a substantial increase in the number of collected signal events. This limitation arises from the saturation of the efficiency suffered by the Run 2 detector and trigger system for hadronic decay modes with increasing instantaneous luminosities, as can be seen in Figure 3.16. The main bottleneck is given by L0. As described earlier, L0 selects events based




Figure 3.16: B meson trigger yield decays for different instantaneous luminosities when L0 trigger is exploited. Values are normalized to the yield expected in design conditions for $\mathcal{L} = 2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. Taken from [141].

on transverse energy deposits exceeding several GeV in the detector, targeting hadrons, muons, electrons, or photons. This mechanism efficiently selects di-muon final states, but

requires the application of a rather tight transverse energy threshold for hadronic modes, leading to an inefficiency of $\sim 50\%$ for hadronic final states, already in Run 2 conditions. At higher luminosities, the significant increase in background processes would require even tighter thresholds for hadronic channels. This would further compromise the signal efficiency, ultimately leading to the observed saturation in signal yield.

In order to fully exploit the larger instantaneous luminosity, LHCb went under a major upgrade during the Long Shutdown 2 of the LHC [142], with the target of collecting 50 fb^{-1} by 2030. Several sub-detectors have been replaced in order to cope with the higher occupancy, the L0 trigger step has been removed, and the DAQ system has been upgraded so that is now able to perform a trigger-less readout of the entire detector, and a pretty advanced event reconstruction at the full bunch-crossing rate of the LHC.

3.3.1 The tracking system upgrades

All sub-detectors within the LHCb tracking system were replaced to handle the increased detector occupancy at higher luminosities. Additionally, the front-end electronics were entirely upgraded to support a readout rate of 30 MHz, the design read-out rate for Run 3.

The VELO Upgrade

The Run 3 VELO detector [129] has been redesigned to accommodate the increased luminosity and higher particle densities expected in Run 3 data-taking. It features a fully silicon pixel design, replacing the previous silicon-strip technology. This change allows for enhanced spatial resolution, higher radiation tolerance, and improved capability to handle pattern recognition at the higher track densities. The geometrical arrangement of the detector follows the Run 2 one. It is made of 26 stations, each one composed of two modules, placed on the two sides of the beam. 19 stations (7) are placed in the forward (backward) region with respect to the nominal interaction point. On each module, four different sensors are mounted, two on each side. As in Run 2, the two VELO sides are installed on retractable halves and placed inside a vacuum tank directly connected with the LHC beam pipe. In the upgraded system, these can be placed closer to the beamline (5 mm instead of 8 mm in Run 2), increasing acceptance for displaced vertices from heavy-flavor decays and improving the impact parameter resolution. The thickness of the RF-foil has also been decreased in Run 3 from 300 to 250 μm [129], to decrease the material interaction effects. The system incorporates a fast, lightweight cooling system using evaporative CO_2 to manage the increased heat load while minimizing material interactions that could degrade tracking performance. A schematic representation of LHCb upgraded VELO is reported in Figure 3.17. The novel VELO detector is able to provide the Run 2 precision on vertex position, and a long track impact parameter resolution of $(16 + 17/p_T[\text{GeV}/c]) \mu\text{m}$, improving the precision for particles with lower momentum.

The Upstream Tracker

The Tracker Turicensis (TT) detector was replaced in Run 3 by the Upstream Tracker (UT) [143], which fulfills the same primary roles as its predecessor. The UT is a silicon micro-strip detector consisting of two pairs of layers separated by approximately 30 cm. These four layers follow the same x - u - v - x configuration as the TT, with the overall




Figure 3.17: A schematic representation of the upgraded VELO structure. Taken from [129].

structure shown in Figure 3.18, but the UT represents a significant upgrade over the




Figure 3.18: A schematic representation of the Upstream Tracker structure. Taken from [143].

TT, featuring thinner sensors with finer segmentation. Specifically, the UT employs strip pitches ranging from $95 \mu\text{m}$ to $190 \mu\text{m}$, compared to the TT's uniform pitch of $183 \mu\text{m}$. Furthermore, the UT's geometry was designed to offer greater acceptance coverage than the TT, enhancing the efficiency for detecting low-momentum particles. These improvements ensure better track reconstruction performance and reduced background contamination in the higher luminosity environment of Run 3. The UT plays a strategical role also in the trigger, as the first estimate of particle momentum provided by VELO-UT tracks ($\Delta p/p \sim 15\%$) is exploited to speed up matching of hits in the forward tracker [144].

The Scintillating Fibre detector

The Scintillating Fibre (SciFi) detector [143] is the upgraded replacement for the T-stations used in Run 2, for both the Inner Tracker (IT) and Outer Tracker (OT) detectors. The SciFi detector consists of 12 detection layers arranged in the x - u - v - x configuration. Each layer comprises 12 modules, each 5 m tall and 52 cm wide, with special accommodations for the beam pipe in the form of a hole in the two central modules. It uses scintillating fibers that are 2.5 m long and $250\ \mu\text{m}$ in diameter. These fibers feature a polystyrene core with 1% organic fluorescent dye by weight, which enhances the scintillation process. The light produced in the fibers is read out by silicon photomultipliers located at the top and bottom of the detector. This design achieves a raw hit resolution of $42\ \mu\text{m}$. The structure of the SciFi detector is illustrated in Figure 3.19.

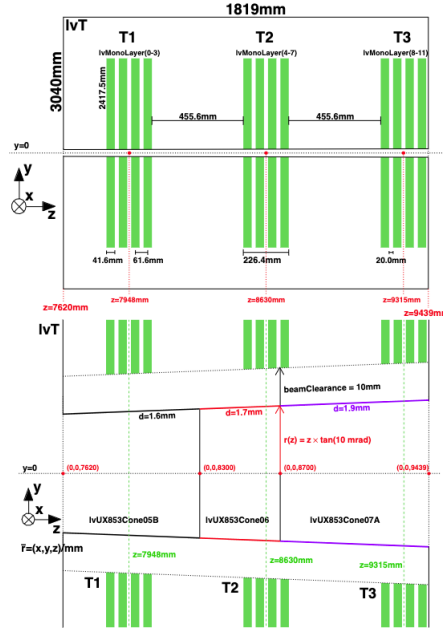


Figure 3.19: A schematic representation of the Scintillating Fibre detector structure. Taken from [143].

3.3.2 Data acquisition and trigger system for LHCb upgrade

The upgrade of the data acquisition and trigger systems is the most innovative and challenging part of the LHCb Upgrade I. The exploitation of a trigger-less readout system, performing a complex event reconstruction at the full bunch-crossing rate of the LHC, is a first-time at a hadron collider, and required the complete replacement of the readout electronics, online and trigger systems. A scheme displaying the structure of the one exploited in Run 3 is shown in Figure 3.20.

The Online system

In the upgraded system, detector information is read out at the full maximum LHC bunch-crossing rate of 40 MHz and it is handled by the so-called Online system. The first

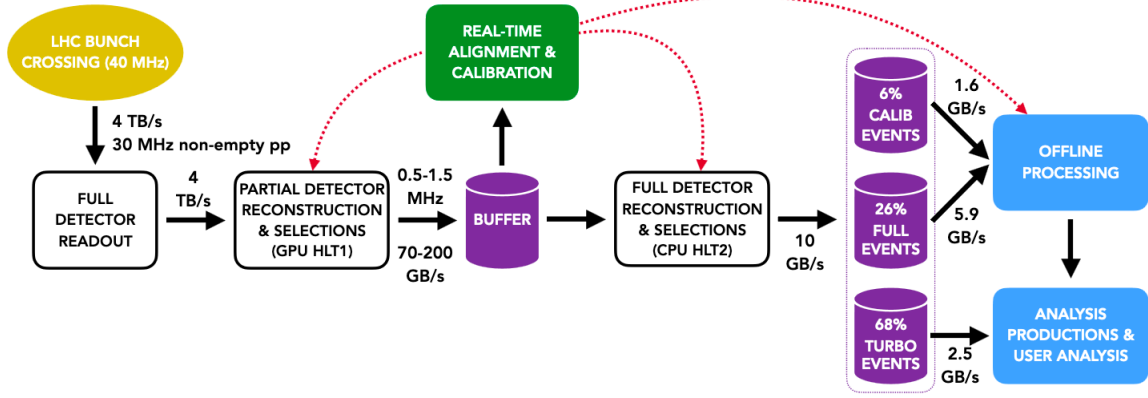


Figure 3.20: Diagram reporting the LHCb dataflow exploited in Run 3 after the Upgrade. Taken from [145].

data acquisition step is event building, and is performed by a farm of 162 event builder (EB) servers. Each server hosts an FPGA board (TELL40), working as a back-end receiver for the server. The overall input stream corresponds to a bandwidth of 4 TB/s. Each EB server receives data from sub-detectors' front-end electronics in the form of Multi-Fragment Packets (MFPs), and information from a single packet of events is aggregated into Multi Event Packets (MEP), each containing 1000 events, that are then passed to HLT1. A schematic representation of the Online system structure is reported in Figure 3.21.

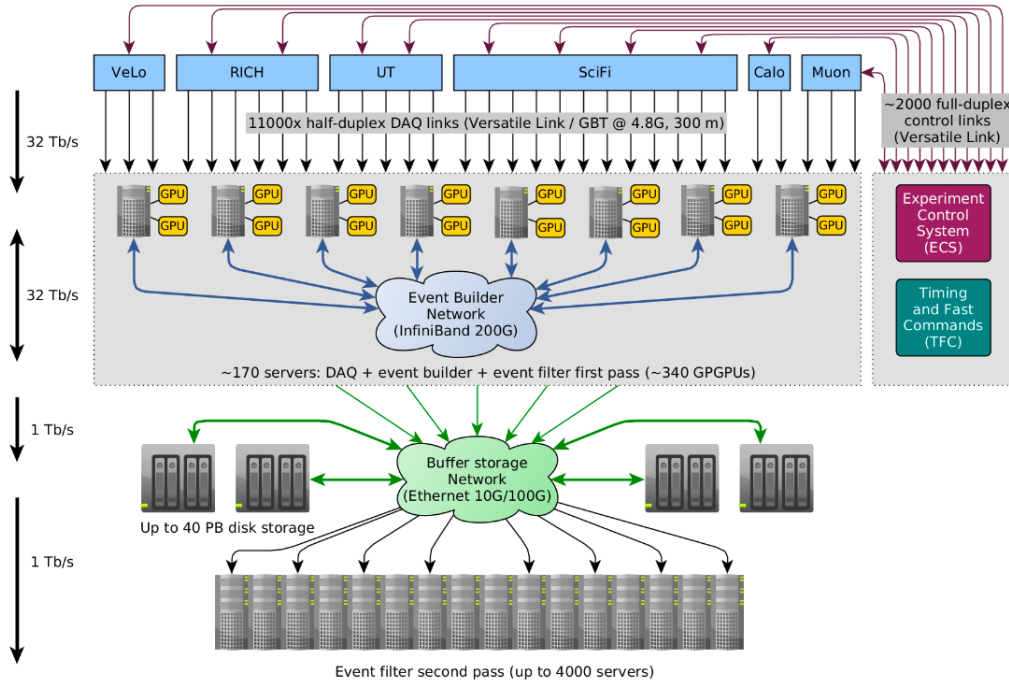


Figure 3.21: A schematic representation of the Online system. Taken from [144].

HLT1

With the removal of L0, the high-level trigger does all the processing and selection, based just on software processes. It is divided into two sequential steps as in Run 2 - HLT1 and HLT2. HLT1 receives the full unfiltered detector information at the non-empty pp average bunch crossing rate of 30 MHz. Its purpose is to reduce the event rate to 1-1.5 MHz. This is achieved by performing a partial event reconstruction, and applying inclusive selections on reconstructed objects.

The HLT1 is run on a farm of $O(500)$ NVIDIA RTX A5000 GPU cards and the processing is implemented in the so-called Allen framework [146, 147]. The GPU cards are hosted in the EB servers themselves, as this allows for a fast communication between EB and HLT1, exploiting the server's internal communication. This design is also a cheaper solution for the communication between EB servers and the LHCb CPU farm, now responsible for the execution of HLT2 only. In fact, since HLT1 is expected to reduce the event rate by a factor ~ 30 , a smaller bandwidth connection is needed between the two farms.

Despite the stringent throughput constraints under which HLT1 must operate, it performs reconstruction of long tracks, primary and secondary vertices. The PID information is partially processed, with the exploitation of ECAL and muon stations for the identification of muons and electrons (while RICH information is not processed at this level).

The *Looking Forward* algorithm [146, 148] (commonly referred to as the *forward* algorithm) is the HLT1 long track reconstruction algorithm, according to its original design. The first reconstruction is run in the VELO tracker. In this region, particle trajectories are reconstructed as straight lines, since no magnetic field is present, and then exploited in the reconstruction of primary vertices. UT hits are then attached to VELO tracks to form VELO-UT tracks. This is done by extrapolating VELO tracks to the UT. Search windows are defined in the UT around the predicted x track position, and found hits are matched to the track. The VELO-UT tracks allow the reconstruction of low-momentum tracks that are swept away by the magnetic field, and they also provide an initial momentum estimate ($\sim 15\%$ resolution) to be exploited in subsequent long tracks reconstruction. The long tracks reconstruction is then performed by attaching SciFi hits to the VELO-UT tracks. Each track is extrapolated to the SciFi and a window is opened in the SciFi layers around the predicted position, where possible hits are searched. The pattern recognition needed for the hits matching is a computationally heavy task, and the exploitation of the momentum estimate provided by the VELO-UT tracks allows defining a significantly narrower search window in the SciFi. Despite this strategy, a limitation on the window size has to be applied anyway, with the net effect of excluding long tracks below a certain p_T , typically $\sim 450 \text{ MeV}/c^2$ [144, 147]. This can limit the achievable trigger efficiency for relatively low momentum decays, like *charm* or *strange* ones; implications will be discussed later. Once the pattern recognition is solved, the identified tracks are fitted to extract track parameters.

Following the identification of long tracks, they are utilized to reconstruct secondary vertices. To minimize computing time and reduce the number of possible combinations, specific selection criteria are applied to the tracks. In particular, only long tracks with $p > 0.2 \text{ GeV}/c$, $\chi^2_{track} < 2.5$ and $\chi^2_{IP} > 4^2$ are exploited for the identification of secondary

²See definition of this variables in Section 4.2.

vertices. After the application of those selections, tracks are matched in pairs and a least-squares fit is performed on each identified vertex. Vertices are retained for further selection if they exhibit a positive-definite covariance matrix and satisfy quality criteria, including a good χ^2 value.

Then, a set of selections is applied to reconstructed objects, in order to achieve the desired event rate reduction. These are based on the presence of at least one track with high- p_T and large displacement, similarly to Run 2. More details on the HLT1 selections in use will be given in the next chapter.

Before being handed to HLT2, retained events are stored in an intermediate disk buffer, whose size is 40 PB. It provides temporary storage for HLT1 filtered data, allowing running HLT2 at a lower instantaneous rate, exploiting the periods without collisions to process stored data. Additionally, data stored in the buffer are exploited to compute alignment and calibration constants, that are updated and exploited in reconstruction algorithms.

HLT2

The second trigger level reads out data from the buffer, working asynchronously with respect to the LHC bunch-crossing. Its purpose is to reduce the rate to around 100 kHz, corresponding to a bandwidth of 10 GB/s written to tape. It performs a complete reconstruction of the event, exploiting the full detector information. In particular, the reconstruction of all track categories is performed and full PID information is exploited at this stage. At the HLT2 level, both inclusive and exclusive selections are applied to the reconstructed objects, utilizing a total of over a thousand distinct selection lines. The implementation of such a large number of lines is feasible due to the increased processing time available at this stage and the flexibility of general-purpose CPUs. This extended processing capacity provides significant flexibility, enabling the use of channel-specific selections that take into account the unique characteristics of entire decay chains. As a result, it is possible to achieve both high signal efficiency and minimal background acceptance simultaneously. The exploitation of the *Turbo* paradigm is continued in Run 3, as it demonstrated in Run 2 to be able to save bandwidth for those decays that can rely only on trigger candidates during offline analysis and do not require information from the rest of the event. This is the case for several charmed hadrons decays, as $D^0 \rightarrow K^+ K^- / \pi^+ \pi^-$ exploited in the $\Delta\mathcal{A}^{CP}$ measurement, or the $D^0 \rightarrow K_S^0 K_S^0$ decay, as described in Chapter 8.

Chapter 4

Measuring $D^0 \rightarrow K_S^0 K_S^0$ at LHCb

In this chapter I will describe the baseline LHCb analysis procedure for $\mathcal{A}^{CP}(K_S^0 K_S^0)$, based on the Run 2 measurement, that I directly contributed to finalize and publish during the initial period of my PhD. This will serve to illustrate the primary limitations to the precision of the measurement, that are the key motivations for the subsequent work detailed in this thesis.

4.1 Introduction

The LHCb collaboration published the most precise $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement to date in 2021 [2], obtaining the following result:

$$\mathcal{A}^{CP}(K_S^0 K_S^0) = (-3.1 \pm 1.2 \pm 0.4 \pm 0.2)\%.$$

The relevance of this measurement has been detailed in Chapter 1 and will not be reiterated here.

The measurement was conducted using data collected during the 2015–18 period at a center-of-mass energy of $\sqrt{s} = 13$ TeV, corresponding to an integrated luminosity of approximately 5.9 fb^{-1} . A portion of this dataset, specifically the data collected in 2015 and 2016, had already been used to measure the same observable in a previous publication [149], but was included in the updated measurement to benefit from several improvements in the analysis methodology, providing a significant improvement in sensitivity.

The measured quantity is the $\mathcal{A}^{CP}(K_S^0 K_S^0)$, defined as:

$$\mathcal{A}^{CP}(K_S^0 K_S^0) = \frac{\Gamma(D^0 \rightarrow K_S^0 K_S^0) - \Gamma(\bar{D}^0 \rightarrow K_S^0 K_S^0)}{\Gamma(D^0 \rightarrow K_S^0 K_S^0) + \Gamma(\bar{D}^0 \rightarrow K_S^0 K_S^0)}, \quad (4.1)$$

where Γ is the decay width of the D^0 or \bar{D}^0 meson. This quantity can only be computed knowing the flavor of the D^0 decay at production. This is achieved by exploiting a sample of D^0 originating from the $D^{*+} \rightarrow D^0 \pi^+$ decay (implying charge-conjugate process from here on). The sign of the pion present in the final state (usually referred to as tagging pion, π_{tag}) allows for the identification of the accompanying D^0 meson flavor. The K_S^0 are reconstructed in the $\pi^+ \pi^-$ final state.

The considered decay widths Γ can't be directly observed in data, as the only measurable quantity is the number of observed candidates $N(D^0 \rightarrow K_S^0 K_S^0)$. This quantity is

extracted from data through a simultaneous 3D fit to the $\Delta m = m(D^{*+}) - m(D^0)$, $m(D^0)$ and the two K_S^0 invariant mass distributions, where the $D^0 \rightarrow K_S^0 K_S^0$ signal component appears as a peak in the Δm distribution around $145.5 \text{ MeV}/c^2$ and around $497.6 \text{ MeV}/c^2$ in $m(K_S^0)$ distributions. Calculating the asymmetry directly from the $N(D^0 \rightarrow K_S^0 K_S^0)$ yields obtained from the fit, without any prior corrections, results in the raw asymmetry. This value encompasses not only the physical \mathcal{A}^{CP} but also contributions from production and detection asymmetries, collectively referred to as nuisance asymmetries. The relation between $N(D^0 \rightarrow K_S^0 K_S^0)$ and $\Gamma(D^0 \rightarrow K_S^0 K_S^0)$ is:

$$N(D^0 \rightarrow K_S^0 K_S^0) \propto \sigma(D^{*+}) \cdot \varepsilon^+(D^{*+}) \cdot \Gamma(D^0 \rightarrow K_S^0 K_S^0), \quad (4.2)$$

where $\sigma(D^{*+})$ are the production cross sections and ε^+ the detection probabilities for D^{*+} decays. The production asymmetry arises from the different production rates for D^{*+} and D^{*-} candidates, given by the flavor-asymmetric pp initial state in collisions. The detection asymmetry arises solely from the tagging pion in this channel, as the D^0 meson decays into a charge-symmetric final state. This asymmetry is expected to reverse sign with opposite magnet polarities and is significantly reduced when integrating data from both MagUp and MagDown configurations. However, complete cancellation of this effect cannot be guaranteed. In LHCb measurements, a *calibration* channel is employed to subtract nuisance asymmetries and extract \mathcal{A}^{CP} . This calibration channel shares the same production and detection asymmetries as the signal channel and has an \mathcal{A}^{CP} value known with much greater precision than the expected result of the measurement. In this analysis, a large sample of $D^0 \rightarrow K^+ K^-$ decays is utilized. It offers significantly higher statistics compared to the signal sample and a much more precise \mathcal{A}^{CP} value than that of $D^0 \rightarrow K_S^0 K_S^0$ [150], allowing for an accurate determination and subtraction of nuisance asymmetries.

4.2 Variables definition

Before proceeding further, it is helpful to define a set of variables that will be used throughout the remainder of this document. While some of these variables have already been introduced earlier, their definitions are gathered in one place for clarity and ease of reference. The list of variables is:

- *Primary Vertex (PV)*: spatial point of the primary pp interaction;
- *Secondary Vertex (SV)*: spatial point of the decay of the particle produced in the interaction. Typically a heavy flavor hadron as a charmed or beauty hadron;
- χ^2/ndf : χ^2 obtained from the fit of a track revealed in the detector, normalized to its degrees of freedom;
- χ_{vvd}^2 : squared significance of the distance between primary and decay vertex calculated using the positions and covariance matrices of the two fitted vertexes.
- p : momentum of the particle;
- p_T : transverse component of the particle momentum with respect to beams flight direction;

- τ : lifetime of the particle;
- τ_{BPV} : defined as $\tau_{BPV} = \text{FD}/\gamma\beta c = \text{FD}m/pc$, where FD is the particle flight distance between the PV and the decay vertex; it represents the time necessary for the particle to travel from the PV to its decay vertex, normalized to the particle boost;
- Pseudorapidity (η): defined as $-\log(\tan(\theta/2))$, where θ is the angle between the beams flight direction and the momentum of the particle;
- Helicity angle θ_H : the helicity angle for a decaying candidate, *i.e.* the angle of between the momentum direction of the second daughter particle in the rest frame of the mother particle with respect to the direction of the boost from the laboratory frame to the rest-frame of mother particle;
- θ_{DIRA} (*direction angle*): angle between the momentum of a particle and the direction identified by the PV and the decay vertex of that particle. This angle tends to be larger in case of a partially reconstructed prompt particle, or a particle not produced in the primary pp interaction;
- $DOCA$: distance of closest approach between the tracks of interest;
- \mathcal{P}_{ghost} : probability for the track to be a misidentified track (*ghost*); this variable is returned by a Neural Network, that exploits various variables which describe track reconstruction and global event properties in order to separate ghost tracks, which are spurious combination of hits, from real tracks;
- $Impact\ Parameter(IP)$: distance of closest approach between the PV and the direction identified by the momentum of the particle. For a particle not produced in the PV due to the pp interaction, as K_S^0 in the considered decay chain, this variable is more likely to assume larger values;
- χ_{IP}^2 : difference between χ^2 obtained from PV fitting including or not including a particle in the fit. For a particle does not come from the PV this variable assumes larger value on average;
- χ_{vtx}^2/ndf : χ^2 of the fit to the vertex of origin of two or more tracks, normalized to its degrees of freedom.

It is also useful to introduce two definitions largely used within the LHCb collaboration, regarding how the signal candidate has been selected by the trigger selections. As outlined in the introduction, the LHCb trigger system is organized into three sequential levels: L0, HLT1, and HLT2. At each level, events are categorized into two distinct classes based on the method of their selection:

- *Trigger On-Signal* (TOS): events in this category are triggered directly by the presence of the signal decay chain because one or more tracks originating from the signal decay directly satisfy the selection criteria. Such events are recorded even if the signal process is the only one present in the event;

- *Trigger Independent-of-Signal* (TIS): events in this class are recorded due to the presence of another process occurring in the same interaction other than the signal decay chain. In this case, tracks originating from the signal decay do not contribute to satisfying the trigger criteria.

4.3 Trigger selections

4.3.1 L0

The events exploited in this measurement are required to be selected at the L0 level either by the presence of an HCAL deposit linked to the D^0 candidate (D^0 is TOS on L0Hadron) or by any L0 signal not related to the D^* decay chain (D^* is TIS on L0Global). This configuration minimizes the L0 trigger requirements on the tagging pion and reduces the collected data charge asymmetry. The same selection is applied for both signal and calibration channel.

4.3.2 HLT1

At the HLT1 level, $D^0 \rightarrow K_S^0 K_S^0$ events can be triggered by any of the existing lines. However, due to the final state consisting of charged tracks, the most relevant lines are **TrackMVA** and **TwoTrackMVA**. These lines respectively select events with a single long track or a combination of two long tracks, provided the tracks have sufficient transverse momentum and displacement with respect to reconstructed primary vertices. The **TrackMVA** line selects a single track which satisfies the following requirements:

- $\chi^2/\text{ndf} < 2.5$
- $\mathcal{P}_{\text{ghost}} < 0.2$
- $\{(p_T > 25.0) \wedge (\chi_{\text{IP}}^2 > 7.4)\} \vee$

$$\left\{ [1.0 < p_T < 25.0] \wedge \left[\ln \chi_{\text{IP}}^2 > \ln(7.4) + \frac{1.0}{(p_T - 1.0)^2} + \alpha \left(1 - \frac{p_T}{25} \right) \right] \right\}$$

where p_T is the transverse momentum expressed in GeV/ c and α varied during the Run 2 data taking, as shown in Fig. 4.1. The **TwoTrackMVA** line benefits from the requirement of two tracks, enabling it to apply less stringent p_T and displacement criteria, making it generally more suitable for selecting charm decays. This line fires if a pair of long tracks satisfies the following selection:

- $p_T > 500 \text{ MeV}/c$
- $p > 5000 \text{ MeV}/c$
- $\chi^2/\text{ndf} < 2.5$
- $\chi_{\text{IP}}^2 > 4$

The two-track combination must then satisfy the following requirements:




Figure 4.1: p_T - χ_{IP}^2 region selected by the **TrackMVA** line. The shaded area represents the excluded region. The parameter α takes the value of 1.1 during the 2017 and 2018 data-taking.

- $\chi^2 < 10$
- $2 < \eta < 5$
- $1 < m_{corr} < 10^6 \text{ GeV}/c^2$
- $\cos \theta_{DIRA} > 0$
- selection on the output of a BDT classifier calculated from the following variables: (1) χ^2 of the two-track vertex, (2) distance between the primary vertex (PV) and the two-track vertex, (3) sum of the p_T of the two tracks, (4) number of tracks with $\chi_{IP}^2 > 16$.

However, **TwoTrackMVA** is highly inefficient for selecting K_S^0 mesons due to the selection criterion imposed on the corrected mass (m_{corr}) variable, defined as:

$$m_{corr} = \sqrt{m^2(h_1 h_2) + p_\perp^2(h_1 h_2) + p_\perp(h_1 h_2)}, \quad (4.3)$$

where h_1 and h_2 are the two input particles for the trigger line and p_\perp is the momentum component perpendicular to the direction identified by the PV and the secondary vertex. This variable provides an approximate estimate of the parent particle's mass, accounting for the possibility of missing daughters. In fully reconstructed decays the corrected mass closely approximates the true mass and the applied selection of $m_{corr} > 1 \text{ GeV}/c^2$ rejects the majority of real K_S^0 candidates, making the line extremely inefficient for $D^0 \rightarrow K_S^0 K_S^0$ decays. Because of the limited efficiency of **TwoTrackMVA** line, in this analysis the D^0 is only required to be TOS on **TrackMVA**. In order to have a larger sample size, also events triggered by other decays are considered, including also events for which the D^* is TIS on any of the HLT1 physics lines. Summarizing, the adopted HLT1 selection is: D^0 candidate to be TOS on **TrackMVA** line, or the D^{*+} candidate to be TIS on any HLT1 physics line.

It is important to notice how K_S^0 decaying outside of the VELO acceptance can only be reconstructed in a TIS approach, since downstream tracking and triggering is absent at the HLT1 level.

The calibration channel events are triggered at HLT1 level by the dedicated line, whose selections are summarized in Table 4.1. This line was chosen because it introduces minimal

Variable	
$p_T(K^\pm)$	> 800 MeV
$p_{Tmax}(K^\pm)$	> 900 MeV
$p(K^\pm)$	> 4000 MeV
$\chi_{track}^2(K^\pm)$	< 2
$\sum_{p_T}(K^\pm)$	> 1500 MeV
$ m(K^+K^-) - m(D^0) $	< 100 MeV
$DOCA(K^+K^-)$	< 0.2 mm
$\tau(D^0)$	> 0.2 ps
$\chi_{vtx}^2/ndf(D^0)$	< 10

Table 4.1: Summary of selection in HLT1 on the $D^0 \rightarrow K^+K^-$ decay channel.

biases on the D^0 lifetime by avoiding any displacement-based selection on the kaons. Only a rectangular selection is directly applied on D^0 lifetime ($\tau(D^0)$), making its effect easier to interpret and reproduce. Due to these triggering conditions, the $D^0 \rightarrow K^+K^-$ sample selected by this line is commonly referred to as *lifetime-unbiased*. This sample has been chosen as the applied trigger selections determine also a minimal alteration of the fraction of D^0 coming from beauty decays. This approach is different to what is usually done in other charm analyses and it is a crucial point for the developed analysis strategy, and it will be explained later in detail.

4.3.3 HLT2

The K_S^0 candidates are reconstructed through the $K_S^0 \rightarrow \pi^+\pi^-$ decay mode, where the pions can be reconstructed using either long or downstream tracks. This classification results in three distinct categories in which $D^0 \rightarrow K_S^0 K_S^0$ are classified: LL, LD, and DD. These correspond to cases where both K_S^0 candidates are reconstructed with long tracks, one K_S^0 is reconstructed with long tracks and the other with downstream tracks, and both K_S^0 candidates are reconstructed with downstream tracks, respectively. The candidates of each category are selected by a dedicated HLT2 Turbo line whose selections have been specifically tuned for the selection of $D^0 \rightarrow K_S^0 K_S^0$ candidates. Because of the different resolution of the categories, these are separately analyzed and results are only combined at the end. The selections adopted for K_S^0 candidates are reported in Table 4.2. Retained K_S^0 candidates are then combined to form D^0 candidates. The applied selections varied during data-taking and are reported in Table 4.3 and Table 4.4, respectively, for the 2015-2016 and 2017-2018 periods. A dedicated Turbo line is exploited for the selection of the calibration sample at the HLT2 level. This is again refraining from the application of any displacement cut on the kaons, similarly to what has been done at the HLT1 level. The selections applied by this line are summarized in Table 4.5

Variable	$K_{SL}^0 \rightarrow \pi_L^+ \pi_L^-$	$K_{SD}^0 \rightarrow \pi_D^+ \pi_D^-$
$\chi^2/\text{ndf}(\pi)$	< 3	< 4
$\mathcal{P}_{\text{ghost}}(\pi)$	< 0.4	< 0.4
$\chi_{\text{IP}}^2(\pi)$	> 36	—
$p_T(\pi)$	—	$> 175 \text{ MeV}/c$
$p(\pi)$	—	$> 3000 \text{ MeV}/c$
$ m(\pi^+ \pi^-) - m(K_S^0) $	$< 35 \text{ MeV}/c^2$	$< 64 \text{ MeV}/c^2$
$\chi_{\text{vtx}}^2/\text{ndf}(K_S^0)$	< 30	< 30
$\tau(K_S^0)$	$> 2 \text{ ps}$	$> 0.5 \text{ ps}$
$z(K_S^0) - z(PV)$	—	$> 400 \text{ mm}$
$z(K_S^0)$	$\in [-100, 500] \text{ mm}$	$\in [300, 2275] \text{ mm}$

Table 4.2: Summary of K_S^0 selection in HLT2: these candidates are combined to form D^0 candidates.

Variable	$D^0 \rightarrow K_{SL}^0 K_{SL}^0$	$D^0 \rightarrow K_{SL}^0 K_{SD}^0$	$D^0 \rightarrow K_{SD}^0 K_{SD}^0$
$\sum_{K_S^0} p_T$	$> 1500 \text{ MeV}/c$	$> 2000 \text{ MeV}/c$	$> 2000 \text{ MeV}/c$
$p_T(K_S^0)$	$> 500 \text{ MeV}/c$	$> 750 \text{ MeV}/c$	$> 750 \text{ MeV}/c$
$\chi_{\text{IP}}^2(K_S^0)$	> 9	> 4	> 4
$\chi_{VVD}^2(D^0)$	> 20	> 10	> 10
$m(K_S^0 K_S^0)$	$\in [1789, 1949] \text{ MeV}/c^2$		
$\chi_{\text{vtx}}^2/\text{ndf}(D^0)$	< 10		
$\theta_{\text{DIRA}}(D^0)$	$< 34.6 \text{ mrad}$		
$\tau(D^0)$	$> 0.2 \text{ ps}$		
$m(D^0 \pi_{\text{tag}}) - m(K_S^0 K_S^0)$	$\in [130.0, 160] \text{ MeV}/c^2$		
$p_T(\pi_{\text{tag}})$	$> 100 \text{ MeV}/c$		
$\mathcal{P}_{\text{ghost}}(\pi_{\text{tag}})$	< 0.4		
$\chi^2/\text{ndf}(\pi_{\text{tag}})$	< 3		
$\chi_{\text{vtx}}^2/\text{ndf}(D^*)$	< 25		

Table 4.3: Summary of HLT2 trigger selection on D^0 candidates for 2015 and 2016 $D^0 \rightarrow K_S^0 K_S^0$ samples.

4.4 Background sources

The $D^0 \rightarrow K_S^0 K_S^0$ sample collected using the described trigger selections is subject to contamination from various background sources. These backgrounds can introduce biases into the measurement and reduce its precision. Consequently, addressing and mitigating these backgrounds is a critical aspect of performing the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement at LHCb.

Variable	$D^0 \rightarrow K_{SL}^0 K_{SL}^0$	$D^0 \rightarrow K_{SL}^0 K_{SD}^0$	$D^0 \rightarrow K_{SD}^0 K_{SD}^0$
$\sum_{K_S^0} p_T$	$> 1500 \text{ MeV}/c$		
$p_T(K_S^0)$	$> 500 \text{ MeV}/c$		
$\chi_{IP}^2(K_S^0)$	> 4		
$\chi_{VVD}^2(D^0)$	> 5		
$m(K_S^0 K_S^0)$	$\in [1775, 1955] \text{ MeV}/c^2$		
$\chi_{\text{vtx}}^2/\text{ndf}(D^0)$	< 10		
$\theta_{DIRA}(D^0)$	$< 34.6 \text{ mrad}$		
$m(D^0 \pi_{tag}) - m(K_S^0 K_S^0)$	$\in [-75, 170] \text{ MeV}/c^2$		
$p_T(\pi_{tag})$	$> 200 \text{ MeV}/c$		
$\mathcal{P}_{\text{ghost}}(\pi_{tag})$	< 0.25		
$\chi^2/\text{ndf}(\pi_{tag})$	< 3		
$\chi_{\text{vtx}}^2/\text{ndf}(D^*)$	< 25		

 Table 4.4: Summary of HLT2 trigger selection on D^0 candidates for 2017 and 2018 $D^0 \rightarrow K_S^0 K_S^0$ samples.

Variable	
$p_T(K^\pm)$	$> 800 \text{ MeV}$
$p_{Tmax}(K^\pm)$	$> 1200 \text{ MeV}$
$\chi_{\text{vtx}}^2/\text{ndf}(D^0)$	< 10
$\tau(D^0)$	$> 0.25 \text{ ps}$
$DOCA(K^+ K^-)$	$< 0.1 \text{ mm}$
$\theta_{DIRA}(D^0)$	$< 141.5 \text{ mrad}$
$p_T(D^0)$	$> 2000 \text{ MeV}$
$m(D^0 \pi_{tag}) - m(K^+ K^-)$	$\in [130, 160] \text{ MeV}$
$p_T(\pi_{tag})$	$> 200 \text{ MeV} (> 100 \text{ MeV})$
$\mathcal{P}_{\text{ghost}}(\pi_{tag})$	$< 0.25 (< 0.4)$
$\chi^2/\text{ndf}(\pi_{tag})$	< 3
$\chi_{\text{vtx}}^2/\text{ndf}(D^*)$	< 25
$\cos(\theta_{DIRA}(D^0))$	> 0.9

 Table 4.5: Summary of selection in HLT2 on the $D^0 \rightarrow K^+ K^-$ decay channel. In brackets (), 2015 and 2016 selections are listed (if different from 2017 and 2018 ones).

4.4.1 Peaking backgrounds

A set of insidious backgrounds are the ones simulating the same peaking structure as the signal in the Δm and $m(K_S^0)$ distributions. If these are not rejected before the result extraction, they can determine an overestimate of the $D^0 \rightarrow K_S^0 K_S^0$ statistics, together with a bias on the \mathcal{A}^{CP} value.

Some peaking backgrounds are due to partially reconstructed D^0 decays, as $D^0 \rightarrow K_S^0 K_S^0 \pi^0$, or non- D^0 decays, as $D_s^+ \rightarrow K_S^0 K_S^0 \pi^+$. These are effectively rejected in the analysis by applying an offline selection on the invariant mass of the D^0 candidate.

In particular, the selection applied for the LL/LD(DD) sample is:

$$|m(K_S^0 K_S^0) - m(D^0)| < 20(30) \text{ MeV}/c^2, \quad (4.4)$$

where $m(D^0)1864.84 \text{ MeV}/c^2$ is the PDG value for the D^0 mass [82].

Another insidious peaking background is due to the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays. These are not completely rejected by the $D^0 \rightarrow K_S^0 K_S^0$ HLT2 selections because of the finite width of the mass window exploited to select K_S^0 candidates allowing for the $\pi^+ \pi^-$ pair is interpreted as coming from a K_S^0 decay. This decay is completely reconstructed and it is not rejected by the selection applied on the D^0 invariant mass. It does not however bias our measurement, as it is disentangled from the signal component in the 3D fit, since only one of the K_S^0 candidates comes from a real K_S^0 decay. In order to improve the purity of the signal to maximize resolution, a selection based on the K_S^0 flight distance is applied to further reject this background:

- LL sample: $\log(\chi_{FD}^2(K_S^0)) > 4$, applied on both K_S^0
- LD sample: $\log(\chi_{FD}^2(K_S^0)) > 2$, applied only to the long K_S^0

This background is negligible in the DD sample.

4.4.2 Combinatorial background

The primary source of background in the $D^0 \rightarrow K_S^0 K_S^0$ sample is the so-called *combinatorial* background. This term refers to any random combination of tracks and/or candidates that mimic the expected signal. Such combinations may include: unrelated track pairs falsely reconstructing a K_S^0 candidate; fake or genuine K_S^0 candidates that combine to form a fake D^0 candidate; or fake/real D^0 candidates paired with a random pion to produce a fake D^* candidate. The combinatorial background as it appears as a flat distribution in the $m(\pi^+ \pi^-)$, $m(K_S^0 K_S^0)$ and $m(D^0 \pi)$ mass spectra and can be disentangled from signal in the fit. However, it significantly lowers the sample's purity, limiting the achievable precision on $\mathcal{A}^{CP}(K_S^0 K_S^0)$ and it's partially rejected before the yields extraction. The output of a k-nearest-neighbours (kNN) classifier [151] has been exploited to reject the combinatorial background. A kNN classifier is a simple, non-parametric machine learning algorithm that classifies data points based on the majority class of their k nearest neighbors in the feature space, as measured by a distance metric (e.g., Euclidean distance). It is trained exploiting a simulated $D^0 \rightarrow K_S^0 K_S^0$ sample as signal and candidates from D^0 sidebands from data as background. The feature is defined by several track-related variables, as track and vertex quality, K_S^0 and D^0 candidates p_T , K_S^0 and D^0 decay helicity angles, and particle identification parameters of the D^0 final state particles. Any variable related to the position and displacement of D^0 and D^* candidates is carefully avoided, in order to prevent any bias in the secondaries fraction. To maximize the use of the available statistics, rather than applying a single threshold on the classifier output, the output has been divided into bins, and the analysis has been conducted across these individual bins. The optimized figure of merit is then determined as the value of σ_S/N_S obtained by combining the contributions from all bins.

4.4.3 Secondary decays

As previously mentioned, a fraction of the D^{*+} candidates present in the signal sample are not produced in a primary pp collision (*prompt* decays), but they come from the decay of a *beauty* hadron. These are referred to as *secondary decays*. They are affected by a different production asymmetry w.r.t. prompt decays and their presence in the signal sample can potentially bias the \mathcal{A}^{CP} measurement if not correctly taken into account. Because of this, they are customarily suppressed in LHCb \mathcal{A}^{CP} measurements, through selections based on D^0 vertex displacement. However, for $D^0 \rightarrow K_S^0 K_S^0$ decays, this approach would result in a substantial loss of statistics. The D^0 vertex position in this decay mode is determined with significantly lower precision compared to decays involving charged particles, such as $D^0 \rightarrow K^+ K^-$, primarily due to the long flight distance of the K_S^0 . Consequently, the separation between prompt and secondary decays is poor, and applying stringent selection criteria to reject secondary decays would lead to a significant reduction in the signal yield. This is an even greater issue when downstream tracks are involved, as the lower resolution of the T-stations and the increased projection distance further degrade vertex precision.

Because of this, it was decided not to apply such selections, and to accept secondary decays as part of the signal sample instead. The bias in the \mathcal{A}^{CP} measurement can be prevented by using a calibration sample containing the same proportion of secondary decays (f) as the signal sample. This ensures that both samples exhibit identical spurious asymmetries, enabling their proper cancellation when measuring $\mathcal{A}^{CP}(K_S^0 K_S^0)$. While the $D^0 \rightarrow K_S^0 K_S^0$ and $D^0 \rightarrow K^+ K^-$ samples initially share the same f at production, applied selections — both at the trigger and offline levels — could alter this fraction. Consequently, any selection criteria based on variables with differing resolutions for the D^0 vertex between the signal and calibration samples must be avoided, as such selections could affect the secondary-to-prompt ratio differently in the two samples.

4.4.4 Sample splitting

A highly effective tool widely used in charm analyses for improving the precision of the mass measurement is the decay tree fitter (DTF) [152]. This tool performs a simultaneous fit to the entire decay chain while allowing for the inclusion of specific constraints during the fit. One such constraint is the requirement for the D^* decay vertex to coincide with the PV (PV constraint). Since the PV position is determined with much greater precision, this constraint leads to a substantial enhancement in the invariant mass resolution, and the signal-to-background ratio (S/B) of analyzed samples, improving the precision of measurements.

For secondary decays, the D^{*+} exhibits a non-zero displacement from the PV and applying the PV constraint causes a distortion of the Δm distribution, complicating its fit. To enable the exploitation of this tool also in this measurement where secondary decays are included in the signal sample, the dataset is divided into two subsets: D^{*+} candidates with decay vertices either compatible (PV-compatible) or not compatible (PV-incompatible) with the PV. The variable exploited to separate the two samples is defined as:

$$\Delta\chi_{DTF}^2 = (\chi_{DTF, PV}^2 - \chi_{DTF, NO-PV}^2), \quad (4.5)$$

where $\chi_{DTF, PV}^2$ ($\chi_{DTF, NO-PV}^2$) is the χ^2 of the DTF applying (not applying) the PV constraint. The chosen threshold separating the two sub-samples is $\Delta\chi_{DTF}^2 = 15$. This

corresponds to a 5σ deviation, since the $\Delta\chi_{DTF}^2$ is distributed as a χ^2 with 3 degrees of freedom. The two sub-samples are separately analyzed and the results are only combined at the end, having different invariant mass resolution. This approach allows the PV constraint to be applied to the compatible subset, fully leveraging the resolution improvement, while omitting the constraint for the incompatible subset. Since the PV-compatible subset contains the majority of candidates — owing to the limited resolution on the D^0 vertex position — this strategy still achieves a substantial enhancement in the precision of the \mathcal{A}^{CP} measurement, without the need for removing secondaries.

4.5 Calibration sample offline preparation

Some offline selections are applied to the calibration sample in order to equalize selections with the signal sample and improve its purity. The main offline selections applied to $D^0 \rightarrow K^+ K^-$ control sample are:

- $|m(K^+ K^-) - m(D^0)| < 20 \text{ MeV}/c^2$;
- identification of PV-compatible and PV-incompatible samples according to $\Delta\chi_{DTF}^2$;
- $|\Delta m - 145.5 \text{ MeV}/c^2| < 1.5 \text{ MeV}/c^2$.

No further offline selection needs to be applied to the $D^0 \rightarrow K^+ K^-$ sample, as its statistics allow the determination of $A_{raw}(K^+ K^-)$ with a precision that is much larger than the one expected on $\mathcal{A}^{CP}(K_S^0 K_S^0)$. The Δm distributions for the control sample after the application of reported selections, except for the Δm one, are shown in Figure 4.2.

To allow utilizing this sample to cancel nuisance asymmetries, it is still necessary to correct for the impact on the secondary decays fraction f of the $\tau(D^0)$ selection applied at trigger level. The selection applied on $K^+ K^-$ decays is $\tau(D^0) > 0.25 \text{ ps}$. Applying the same selection to the $K_S^0 K_S^0$ sample does not equalize f between the two datasets due to the differing resolution of the variable in each sample. The strategy adopted in Run 2 to correct for this effect was to assign a weight of 2 to events with $0.25 < \tau < 0.5 \text{ ps}$, in order to compensate for the lack of candidates at lower lifetimes. This works under the assumption that candidates with $0 < \tau < 0.5 \text{ ps}$ have the same f , since the lifetime of a b -hadron is larger than a c -hadron, and we are looking at very low decay times [153]. This procedure is not repeated in the analysis of Run 3 data (Chapter 8), as the $\tau(D^0)$ selection has been removed from trigger selections.

4.6 Procedure for cancellation of nuisance asymmetries

In LHCb asymmetry measurements, nuisance asymmetries are usually canceled by subtracting the asymmetry measured on a calibration sample from that observed in the signal sample. This cancellation procedure is valid in the approximation of small asymmetries (typically below $\sim 20\%$). This requires excluding from analysis any acceptance areas affected by larger asymmetries. There are indeed in the LHCb acceptance some regions, close to the beam pipe and to the edge of acceptance, where particles of opposite sign are swept outside, or inside, the LHCb acceptance by the magnet. The selections applied to




Figure 4.2: Δm distribution for the $D^0 \rightarrow K^+ K^-$ sample. 1/4 of the statistics is shown. Left (right) plot shows distribution for the D^0 (\bar{D}^0) candidates. Top (bottom) plot shows PV-compatible (PV-incompatible) sample. Taken from [153]. A $|\Delta m - 145.5 \text{ MeV}/c^2| < 1.5 \text{ MeV}/c^2$ selection is applied on this sample before its exploitation in the nuisance asymmetry cancellation procedure.

reject these regions will be referred to in the following as *fiducial acceptance cuts*. Applying those fiducial cuts causes a statistics loss of the order of 30% in the case of $D^0 \rightarrow K_S^0 K_S^0$. In order to avoid this statistics loss, a novel method for nuisance asymmetries cancellation has been developed within this analysis, allowing also the exploitation of regions affected by larger asymmetries.

The novel technique achieves cancellation of nuisance asymmetries by assigning a weight to each signal candidate [153]. This weight is derived from the asymmetry measured in the calibration sample at the corresponding point in the kinematic phase space occupied by the signal candidate. The procedure is structured as follows:

- for each signal candidate the position in the kinematic phase space is identified. This is operatively done by using the D^0 3D momentum \vec{p}_0 ;
- the local particle density of $D^{*\pm}$ candidates $n_C^\pm(\vec{p}_0)$ is estimated in a local neighborhood of the considered \vec{p}_0 . This correspond to measuring the nuisance asymmetry in the kinematic phase space point identified by \vec{p}_0 ;
- the weight applied on the signal candidate is

$$w^\pm(\vec{p}_0) = \frac{n_C^+(\vec{p}_0) + n_C^-(\vec{p}_0)}{2n_C^\pm(\vec{p}_0)} [1 \pm \mathcal{A}^{CP}(K^+ K^-)] .$$

This procedure applies a local correction of the nuisance asymmetries in the signal sample. The $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement can then be done by performing a weighted fit to the $K_S^0 K_S^0$ sample.

A kNN classifier is used to estimate the local particle density $n_C^\pm(\vec{p}_0)$. The $D^0 \rightarrow K^+ K^-$ sample is split into two halves, according to the sign of the tagging pion. These are used as input to the kNN algorithm as the two classes to be distinguished, exploiting $\vec{p}_0(D^0)$ as feature space. In this way, the kNN output is

$$out = \frac{n_C^+(\vec{p}_0)}{(n_C^+(\vec{p}_0) + n_C^-(\vec{p}_0))},$$

the reciprocal of the needed weight. Justification and a detailed description of the procedure can be found in Ref. [153].




Figure 4.3: MagDown (top) and MagUp (bottom) weights distribution for 2017-2018 data. Taken from [2].

The weighting procedure is performed separately for MagUp and MagDown configurations. An example of weight distributions is reported in Figure 4.3.

4.7 Results and systematic effects

After weights are computed for each signal candidate it is possible to extract $\mathcal{A}^{CP}(K_S^0 K_S^0)$ fitting signal sample distributions. In the case of the PV-compatible (PV-incompatible) samples, the signal peak in the Δm distribution is well-described by a Johnson S_U (Gaussian) distribution [154]. The $m(K_S^0)$ peak is described by the sum of two Gaussians for both PV-compatible and PV-incompatible samples. An empirical threshold function is used to describe the background in the Δm distribution

$$b(x) \propto \{1 - \exp[-c(x - m_\pi)]\} + b(x - m_\pi).$$




Figure 4.4: The Δm distributions and fit projections for some representative candidate categories (2017–2018 data). See text for the definition of purities. Taken from [2].

The background in the $m(K_S^0)$ distribution is described with a first order Chebyshev polynomial $pol(x)$ for the long K_S^0 and with a second order Chebyshev polynomial for the downstream K_S^0 . The same modeling is exploited for the PV-compatible and PV-incompatible sample, except for the use of a simple Gaussian to describe the Δm peak instead of a Johnson S_U . The total background p.d.f. is given by the sum of 7 different components, comprising all possible combinations. An unbinned maximum-likelihood fit is exploited to fit the joint distribution of Δm and the two $m(K_S^0)$ observables. The D^{*+} and D^{*-} candidates distributions are fitted simultaneously, enabling a direct extraction of \mathcal{A}^{CP} . Fit is performed separately for different categories (LL, LD, DD), PV-compatible and PV-incompatible sub-samples, and data-taking period (2015-16, 2017-18), combining the results at the end. The fit is simultaneously performed to low- and high-purity bins in each identified category, as the distributions are expected to have the same shape.

The Δm distributions for some representative categories, together with relative fit projections, are reported in Figure 4.4.

The largest systematic uncertainties associated with the reported analysis are:

- the effect due to the choice of the models exploited to describe mass distributions in the fit, being these purely empirical. The $\mathcal{A}^{CP}(K_S^0 K_S^0)$ variation obtained by fitting exploited model to data simulated with alternative models is associated as a systematic effect;
- the effect on $\mathcal{A}^{CP}(K_S^0 K_S^0)$ due to the $D^0 \rightarrow K^+ K^-$ limited size, determining a statistical uncertainty on the computed weight. Its size has been determined bootstrapping the calibration sample. The $\mathcal{A}^{CP}(K_S^0 K_S^0)$ variation exploiting weights computed with different iterations of the bootstrap has been associated as systematic uncertainty;

Sample	2015 + 2016 (2 fb ⁻¹)				2017 + 2018 (4 fb ⁻¹)			
	Yield	\mathcal{A}^{CP} [%]			Yield	\mathcal{A}^{CP} [%]		
LL PV-comp.	1388 ± 41	0.3 ± 2.5	± 0.6		4056 ± 77	-4.3 ± 1.6	± 0.4	
LL PV-incomp.	178 ± 31	-11 ± 17	± 2		430 ± 41	-3.0 ± 7.9	± 1.1	
LD PV-comp.	411 ± 25	-7.2 ± 5.8	± 1.1		1145 ± 49	-2.9 ± 3.8	± 0.7	
LD PV-incomp.	58 ± 18	-10 ± 31	± 4		349 ± 64	-5 ± 17	± 2	
DD	—	—			87 ± 28	-35 ± 47	± 6	

Table 4.6: Yields and $\mathcal{A}^{CP}(D^0 \rightarrow K_S^0 K_S^0)$ results for the different analyzed samples. For asymmetry values, the first uncertainty is statistical and the second is systematic. Taken from [2].

- the effect due to the possible presence of a residual difference between signal and calibration channel secondary decays fraction. This is estimated to be lower than 0.2%;
- uncertainty associated to the limited knowledge on $\mathcal{A}^{CP}(D^0 \rightarrow K^+ K^-)$. This is an external parameter, given by the uncertainty of the most recent measurement of this parameter [150].

Results obtained from the fit, together with statistical and systematic uncertainties, are reported in Table 4.6. The measurement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$, utilizing the full Run 2 data sample and combining results across all sub-samples, is given by:

$$\mathcal{A}^{CP}(K_S^0 K_S^0) = (-3.1 \pm 1.2 \pm 0.4 \pm 0.2)\%.$$

Systematic uncertainties are treated as uncorrelated, with the exception of the uncertainty associated with the fit function shapes, which is considered fully correlated across the sub-samples. The measurement is compatible with zero at the 2.4σ level.

4.8 Considerations about the Run 2 measurement

The measurement briefly described in the previous section is the most precise $\mathcal{A}^{CP}(K_S^0 K_S^0)$ determination to date, significantly contributing to the world average. Achieving this level of precision was a remarkable accomplishment for the LHCb collaboration, particularly given the challenges associated with studying a fully neutral final state like $D^0 \rightarrow K_S^0 K_S^0$ in the complex environment of a hadron collider. In fact, this measurement is particularly challenging in an environment such as the one offered by the LHC, and it is usually better performed at b -factories experiments, like BaBar, Belle and Belle II, where the significantly lower event complexity and ease of triggering compensate for the much smaller cross-section.

It is crucial to notice how, despite the huge charm production cross section offered by the LHC and the exploitation of the full Run 2 dataset, the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement is still statistically limited. The cause for this resides in the selections exploited at trigger level to retain signal events, representing the major limiting factor for the collection of a larger statistics, and the achievement of a higher precision. In particular, the earlier trigger levels, L0 and HLT1, significantly limit the achievable efficiency.

The selections exploited to trigger $D^0 \rightarrow K_S^0 K_S^0$ candidates at the L0 level rely on the presence of a multi-GeV energy deposition on the hadronic calorimeter. This requirement has a limited efficiency for decays of heavy flavors, particularly for charmed hadrons, as can be seen from the plot reported in Figure 4.5. This was a well-known feature of the




Figure 4.5: The L0 hadron efficiency for different beauty and charm decays. Taken from [155].

Run 2 trigger system, but necessary to reduce the event rate to a level manageable by the HLT1 system.

Concerning HLT1, the selections exploited to trigger $D^0 \rightarrow K_S^0 K_S^0$ candidates in Run 2 were based on the **TrackMVA** and **TwoTrackMVA** lines, firing when one or two PV-detached long tracks are present. This trigger strategy imposes significant limitations on the efficiency of selecting $K_S^0 K_S^0$ decays due to two primary factors. First, the absence of downstream track reconstruction and triggering at this level prevents the direct selection of any K_S^0 decays occurring outside the VELO acceptance ($z > \sim 500$ mm) at the HLT1 stage, resulting in a very low efficiency for this fraction of decays. In fact, even if $\sim 2/3$ of K_S^0 decays occur between the VELO and the first T-station¹, downstream reconstructed K_S^0 make up only a small fraction of Run 2 $D^0 \rightarrow K_S^0 K_S^0$, as can be seen from Table 4.6 and Figure 4.4. Efforts to recover these tracks are crucial, as the reported analysis demonstrates their potential utility in measurements. The smaller contribution of these tracks to the final result is not primarily due to their slightly reduced p_T and impact parameter resolution but rather to their limited yield. A dedicated trigger strategy aimed at increasing the number of selected downstream tracks could significantly improve the final precision of the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement. Second, even for the fraction of $D^0 \rightarrow K_S^0 K_S^0$ decays where both K_S^0 candidates are reconstructed with long tracks, the efficiency of the current selections remains low. Pions making up the $K_S^0 K_S^0$ final state are characterized by larger impact parameters with respect to the primary vertex, while their momentum is, on average, lower. These features made these tracks less efficiently reconstructed at HLT1 level in Run 2, as the thresholds for the **TrackMVA** and **TwoTrackMVA** lines were tuned for

¹Considering the example of the $B^0 \rightarrow J/\psi K_S^0$ decays [65], the K_S^0 decays reconstructed with downstream tracks make up 70% of the sample.

selecting b -hadron decays, which limits their efficiency for K_S^0 decays, and for LLPs in general.

Starting with the next chapter, I will describe the work I performed to overcome all these limitations, and boost the LHCb sensitivity to this important measurement.

Chapter 5

Improving $D^0 \rightarrow K_S^0 K_S^0$ acceptance in Run 3

In this chapter I describe my efforts to improve the trigger efficiency for K_S^0 modes, culminating in the implementation of two novel trigger lines within the HLT1 level of LHCb. I describe the optimization of these selections, and analyze their performance in terms of both efficiency gains and background rejection.

As discussed in previous chapters, the LHCb apparatus has many features making it extremely well-suited for the collection, and analysis, of large samples of heavy-flavored hadrons. However, decays with K_S^0 particles in the final state are collected with lower efficiencies by the LHCb trigger. Amongst them, my decay of interest, $D^0 \rightarrow K_S^0 K_S^0$, is a notable example. Before the present work, this was considered to be an intrinsic limitation. There was no proof that by implementing specifically optimized K_S^0 -based selections at HLT1, it was actually possible to improve the LHCb efficiency in selecting such modes. Proving the feasibility here requires not just attaining a sufficient improvement of efficiency for relevant channels, but also showing that it can be obtained with an affordable rate cost. At this point, a trivial but important observation of general value is in order. The removal of the L0 trigger alone cannot produce any increase in efficiency, unless coupled with the introduction of new selection criteria in HLT1, yielding the same reduction in rate of L0 by means of smarter algorithms. This highlights the importance of studies like the one described in the present chapters.

Initially, only triggers based exclusively on *long* tracks will be considered; this means that it will not be necessary to assume the availability of computing power beyond what is already required by all other standard Run 3 triggers, except for the modest power necessary to perform the appropriate combinatorics. Having LLP trigger selections based on long tracks is a very solid basis for building a future version including the use of downstream tracks, that will unlock even greater efficiency gains when such tracks will be made available to HLT1. This will be discussed in detail in the following chapters.

5.1 Triggering on K_S^0 's in HLT1

5.1.1 Lessons learned from Run 2

Let's start by examining the performance of the Run 2 track lines in detail, to gain insights on how a different line could have better performance. Selections applied by the **TrackMVA** and **TwoTrackMVA** lines in Run 2 have been described in paragraph 4.3.2 and will not be repeated here.

In the **TrackMVA** line, while the requirements on χ^2/ndf and $\mathcal{P}_{\text{ghost}}$ help in rejecting fake tracks, the multivariate selection applied in the $\chi_{\text{IP}}^2-p_T$ is the one rejecting most of the background tracks. If we compare the $\chi_{\text{IP}}^2-p_T$ distribution for pions coming from the $D^0 \rightarrow K_S^0 K_S^0$, $K_S^0 \rightarrow \pi^+ \pi^-$ decay with the aforementioned selection, as done in Figure 5.1, we can note that in the majority of cases, such pions do not fire the **TrackMVA** line. The displayed distribution corresponds to a sample of $D^0 \rightarrow K_S^0 K_S^0$ decays simulated under Run 3 conditions, prior to the application of any selections beyond those imposed by acceptance and reconstruction criteria.




Figure 5.1: Distribution of tracks from simulated $D^0 \rightarrow K_S^0 K_S^0$ events on the p_T - χ_{IP}^2 plane. Four entries for each D^0 candidate are present. The selection applied by the Run 2 **TrackMVA** line cut is represented by a red curve.

As previously mentioned, using a two-track-based selection enables a looser selection on the individual track for the **TwoTrackMVA**, potentially achieving a higher efficiency for K_S^0 decays compared to **TrackMVA** line. This is however prevented by the $m_{\text{corr}} > 1 \text{ GeV}/c^2$ requirement. Since the mass of a K_S^0 is $m(K_S^0) \sim 498 \text{ MeV}/c^2$ [156], the applied selection is highly inefficient, as it can be seen in Figure 5.2, where the m_{corr} distribution, for pion pairs coming from the same simulated sample of $D^0 \rightarrow K_S^0 K_S^0$ decays, is shown. It can be noted that the majority of events have a value of m_{corr} around $500 \text{ MeV}/c^2$, with a small tail which goes also over the $1 \text{ GeV}/c^2$ threshold. It is therefore clear that the selections applied in Run 2 at HLT1 level were highly inefficient in selecting K_S^0 candidates, even when they were decaying inside the VELO acceptance.




Figure 5.2: Distribution of the corrected mass (top) and 2-dimensional distribution of the corrected mass as a function of $m(K_S^0)$ (bottom) for the simulated MC sample.

5.1.2 The vision for Run 3

With the removal of the L0, the HLT1 is charged to reduce the rate from an input of 30 MHz (on average) to 1 MHz. Although the rate reduction factor is similar to that of Run 2, it must be achieved at a 30 times higher throughput (30 MHz instead of 1.1 MHz) and with a greater track multiplicity due to the increased instantaneous luminosity. To cope with the higher input rate, inclusive lines may have tightened thresholds. This is evident in the case of the **TrackMVA** line. In the baseline configuration object of the studies presented in this chapter, the p_T threshold is increased from 1 GeV/ c to 2.3 GeV/ c (Figure 5.3). While this is a preliminary configuration to be fine-tuned later using real data, it underscores the need for tighter selection criteria to manage the higher input event rate, that unavoidably must result in an even lower efficiency for K_S^0 decays. The **TwoTrackMVA** selections is instead very similar to Run 2 - that is, still highly inefficient on K_S^0 mesons because of the presence of the $m_{corr} > 1 \text{ GeV}/c^2$ requirement.

It can be therefore seen very well that, without the addition of dedicated trigger lines, the efficiency in selecting K_S^0 mesons decaying inside the VELO acceptance (and consequently $D^0 \rightarrow K_S^0 K_S^0$) would have been lower in Run 3 than in Run 2.

5.2 Development of a K_S^0 line for HLT1

This line is designed to improve the efficiency in selecting K_S^0 mesons decaying inside the vertex detector during Run 3, aiming at the selection of single K_S^0 candidates, and




Figure 5.3: Distribution of Upgrade simulated $D^0 \rightarrow K_S^0 K_S^0$ events (without the application of any cut) on the p_T - χ_{IP}^2 plane. Four entries for each D^0 candidate are present. Run 2 and Run 3 version of the **TrackMVA** line cuts are represented.

will be indicated in the following as **Hlt1TwoTrackKs**. The use of a set of rectangular selections has been preferred to multivariate classifiers, such as *Bosted Decision Trees* (BDT) and *Neural Networks* (NN). While the latter might offer slightly better performance in terms of signal acceptance and background rejection, rectangular selections provide the advantage of being more easily interpretable and reproducible offline. Additionally, this method avoids unwanted correlations among selection variables, which introduced systematic effects in precision measurements in past measurements.

The $D^0 \rightarrow K_S^0 K_S^0$ decay has of course been used as a benchmark for the entirety of the studies here reported, both in the choice of the variables exploited in the selections and in the thresholds tuning, but this line clearly has efficiency for other modes as well. But the efficiency of selections on other decay channels of interest has also been studied, and will be reported separately in the next chapter.

The first step in the definition of the novel HLT1 line is the identification of variables to be exploited in the selections. These must be available at the HLT1 level, and provide some statistical separation between signal and background. In order to identify these variables, two samples are exploited as a benchmark for signal and background examples of K_S^0 decays. The same simulated $D^0 \rightarrow K_S^0 K_S^0$ sample mentioned earlier is utilized for this purpose. This sample is generated under Run 3 conditions, and includes all candidates passing the acceptance and reconstruction criteria, representing the signal candidates that HLT1 is designed to identify. The background sample is drawn from a set of generic pp collision events simulated under the same Run 3 conditions as the signal sample, referred to as *Minimum Bias* events. Specifically, any two-track candidate is constructed by matching tracks from this sample, applying no selection criteria other than requiring their invariant mass to lie within a $70 \text{ MeV}/c^2$ window around the PDG K_S^0 mass. This approach ensures a reliable representation of background candidates that the line is expected to reject during data taking. These samples are processed with offline-level reconstruction, which offers higher precision than the reconstruction performed at the

HLT1 level. The use of this approximation is unavoidable, as the HLT1 system does not output its internally reconstructed quantities, so their direct study has many practical complications. However, any differences are expected to have minimal impact for our present purpose.

5.2.1 Cuts excluded from optimization

I chose some of the selection cuts of the **TwoTrackKs** line without the help of any numerical optimization, as they are motivated by criteria other than optimal background rejection.

The first is the χ^2/ndf of the tracks for K_S^0 decay products. The distributions for signal and background candidates are shown in Figure 5.4, with both distributions normalized to have a unit integral. The distributions displayed are truncated at a value of 3, corresponding to the selection applied during offline reconstruction. A selection criterion of:

$$\chi^2/\text{ndf} < 2.5$$

has been chosen. While a tighter threshold could surely improve background rejection with minimal efficiency loss under nominal reconstruction conditions, this may not hold during the initial data-taking period, with preliminary alignments, where the χ^2 values of tracks could be higher. Applying a stricter cut in this case might therefore exclude a substantial fraction of signal candidates. Therefore, a more conservative selection was preferred.




Figure 5.4: Distributions of tracks χ^2/ndf for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events. Both distributions are normalized such that their integral is equal to one.

The invariant mass of K_S^0 candidates is another key observable considered in the **TwoTrackKs** line, offering substantial separation between signal and background events, as demonstrated by the distributions shown in Figure 5.5. A tight selection around the signal K_S^0 peak would effectively suppress background candidates; however, this approach would complicate the offline analysis. Specifically, a narrow mass window would hinder the extraction of the background component under the signal peak. Additionally, the same considerations discussed for the χ^2/ndf variable apply here. During the initial data-taking period, the invariant mass resolution might be worse than under design conditions,




Figure 5.5: Distributions of K_S^0 candidates invariant mass for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events. Both distributions are normalized such that their integral is equal to one.

potentially leading to a significant loss of signal efficiency with a narrow selection. To account for these effects, the following selection is applied:

$$|m(K_S^0) - 498 \text{ MeV}/c^2| < 45 \text{ MeV}/c^2$$

This mass window is broader than the $72 \text{ MeV}/c^2$ window used at the HLT2 level, ensuring robust signal efficiency while accommodating potential variations in resolution.

The χ^2 of the fit to the K_S^0 candidate vertex is another useful variable for the TwoTrackKs line. Real K_S^0 candidates are expected to exhibit lower χ^2 values on average, as confirmed by the distributions shown in Figure 5.6. Again, to accommodate




Figure 5.6: The $\chi_{vtx}^2(K_S^0)$ distributions for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events, on logarithmic scale. Both distributions are normalized such that their integral is equal to one.

imperfect alignments or other unexpected behaviors during the initial phase of data-taking, a conservative threshold is chosen:

$$\chi_{vtx}^2 < 20$$

This ensures robust signal efficiency while maintaining flexibility.

Once real data are available, these selections can be re-evaluated. If a higher level of background rejection is needed, thresholds can be tightened based on the observed distributions in real data.

The pseudorapidity (η) of the K_S^0 candidate is another variable considered for inclusion in the **TwoTrackKs** line selections. The distributions for signal and background samples are shown in Figure 5.7. As illustrated in the figure, this variable provides good discrimination




Figure 5.7: Distributions of K_S^0 candidates pseudorapidity for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events. Both distributions are normalized such that their integral is equal to one.

between signal and background candidates. In LHCb, reconstructed candidates are typically accepted within the pseudorapidity range $2 < \eta < 5$. However, a tighter selection has been implemented for this analysis:

$$2 < \eta(K_S^0) < 4.2$$

This refined cut significantly reduces the background fraction, as demonstrated by the reported distributions. The threshold is selected to exclude regions where candidates traverse a larger amount of material (exceeding 0.3 interaction lengths) [23], which are associated with increased detection asymmetries.

Two additional variables used in the selections are the angle between the K_S^0 daughters $\theta_{\pi\pi}$ and the K_S^0 θ_{DIRA} . Specifically, the cosine of these angles is employed in the selection criteria, with distributions shown in Figure 5.8. Thresholds for these variables are defined as:

$$\begin{aligned} \cos(\theta_{\pi\pi}) &> 0.99 \\ \cos(\theta_{DIRA}) &> 0.99 \end{aligned}$$

Since these variables provide only limited separation between signal and background candidates, the thresholds were chosen conservatively, to keep an efficiency close to 100%, while still contributing to background rejection.

All selections chosen without optimization are summarized in Table 5.1.

5.2.2 Cuts determined by numerical optimization

I started from considering the momentum (p) and transverse momentum (p_T) of the K_S^0 daughters. The distributions for signal and background candidates for these variables are




Figure 5.8: Distributions of $\cos(\theta_{\pi\pi})$ (top) and $\cos(\theta_{DIRA})$ (bottom, on logarithmic scale) for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events. All distributions are normalized such that their integral is equal to one.

Variable	Threshold
$\chi^2/\text{ndf}(\pi)$	< 2.5
$ m(K_S^0) - 498 \text{ MeV}/c^2 $	$< 45 \text{ MeV}/c^2$
$\chi_{vtx}^2(K_S^0)$	< 20
$\eta(K_S^0)$	$2 < \eta(K_S^0) < 4.2$
$\cos(\theta_{DIRA})$	> 0.99
$\cos(\theta_{\pi\pi})$	> 0.99

Table 5.1: Summary of selections applied by the **TwoTrackKs** line excluded by numerical optimization.

shown in Figure 5.9. Among these, p_T in particular demonstrates strong discriminating power between signal and background. Both observables are listed in the **TwoTrackKs** selection criteria to enhance separation between the two categories.

The next considered observable is the χ_{IP}^2 of the tracks. This provides substantial discriminating power, as tracks from genuine K_S^0 decays are expected to exhibit significantly larger detachment from the PV compared to background tracks, which predominantly




Figure 5.9: Distributions of $p(\pi)$ (top) and $p_T(\pi)$ (bottom) for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events. All distributions are normalized such that their integral is equal to one.

originate from the PV. This is corroborated by the signal and background distributions shown in Figure 5.10. Consequently, the χ_{IP}^2 of the tracks has been incorporated into the **TwoTrackKs** selection criteria to enhance signal-to-background separation.

The transverse momentum of K_S^0 candidates is also taken into account. While it is correlated with the $p_T(\pi)$ already considered, it provides additional discriminating power as it pertains to the entire K_S^0 candidate. The $p_T(K_S^0)$ distributions for signal and background, shown in Figure 5.11, clearly demonstrate the strong discriminating power of this variable. For this reason, $p_T(K_S^0)$ has been incorporated into the **TwoTrackKs** selection criteria.

The last observable considered is defined by the ratio between the product of daughters IP and K_S^0 candidate IP:

$$IP(\pi^+) \times IP(\pi^-) / IP(K_S^0)$$

The signal and background distribution for this variable are shown in Figure 5.12. It shows a strong discriminating power, and it is included in the selection criteria for the **TwoTrackKs** line.




Figure 5.10: Distributions of $\chi_{\text{IP}}^2(\pi)$ for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events, on a logarithmic scale. Both distributions are normalized such that their integral is equal to one.




Figure 5.11: Distributions of $p_T(K_S^0)$ for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events. Both distributions are normalized such that their integral is equal to one.

5.2.3 Observables purposely excluded from our trigger line

The IP and χ_{IP}^2 of heavy-flavored hadron decay products are often exploited in LHCb analyses to separate heavy flavor signals from background, both in online and offline selections, due to their measurable displacement from the primary vertex. A typical example of this is the displacement requirement applied by the **TrackMVA** and **TwoTrackMVA** lines at the HLT1 level to suppress background. I chose to do this in case of the **TwoTrackKs** for the pions coming from the K_S^0 decay, but not for the K_S^0 candidate itself. The primary reason is the limited discriminating power of these variables for K_S^0 candidates, at this stage of the data-taking process. For a long-lived neutral particle like the K_S^0 , determining a displacement on the scale of tens of microns is intrinsically challenging due to the large extrapolation. The low separation power offered by $IP(K_S^0)$ and $\chi_{\text{IP}}^2(K_S^0)$ is confirmed by the IP and χ_{IP}^2 distributions shown in Figure 5.13.

The long flight distance of a K_S^0 provides a strong discriminating power between a genuine K_S^0 decay and background candidates. This is confirmed by the distributions




Figure 5.12: Distributions of $IP(\pi^+) \times IP(\pi^-) / IP(K_S^0)$ for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events. Both distributions are normalized such that their integral is equal to one.

reported in Figure 5.14. This variable would be an obvious candidate to be included in the **TwoTrackKs** selection criteria. This is not implemented as a flight distance requirement would introduce a bias on the secondary decay fraction, making the application of the analysis method developed for the most recent measurement extremely difficult. It is crucial to note that this issue is particularly relevant for a cut applied at the HLT1 level. Such a cut is problematic to reproduce offline, as selections are applied to reconstructed objects that are unavailable in the offline data; HLT2 reconstruction, which is performed from scratch, is the one accessible offline.

5.2.4 Numerical optimization of thresholds

The tool used to optimize the rectangular selection criteria for the identified variables is the rectangular cut optimizer provided by the TMVA toolkit [157], included with the ROOT framework. This tool takes as input two sets of candidates representing the signal and background categories to be separated, along with the list of variables for which cuts will be applied. The output consists of a series of cut configurations, each achieving the maximum background rejection for a given signal efficiency. External constraints on selection thresholds can be specified before the optimization process. A lower limit of 425 MeV/ c^2 was imposed on the $p_T(\pi)$ threshold to ensure it does not fall below the 450 MeV/ c^2 reconstruction cut applied at the HLT1 level. A safety margin of 25 MeV/ c^2 was added to account for potential numerical issues. Similarly, an upper bound of 100 was set for the $\chi_{IP}^2(\pi)$ threshold. While this value is arbitrary, it is justified by the fact that typical selections for this variable remain well below this threshold, as demonstrated by the HLT2 selections in the $D^0 \rightarrow K_S^0 K_S^0$ HLT2 lines, where it is set to 36.

The optimal configuration for the new line's selections is designed to ensure that it occupies only a small fraction of the HLT1 output rate. This approach is driven by the fact that the HLT1 output bandwidth is already nearly fully allocated to 1-track and 2-track based selections. Consequently, the new line must achieve an enhancement




Figure 5.13: Distributions of $IP(K_S^0)$ and $\chi^2_{IP}(K_S^0)$ for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events. All distributions are normalized such that their integral is equal to one.

in efficiency while minimally impacting the existing system. To this end, the optimal configuration is defined as the one that maximizes efficiency while limiting its share of the total HLT1 output rate to approximately 5%, corresponding to an assumed 1 MHz output at this stage. This also fits the primary goal of the study at this level: to demonstrate the feasibility of achieving a significant efficiency increase with a limited rate cost. Further fine-tunings are planning to be conducted closer to the start of the data-taking for the overall HLT1 system.

5.2.5 Selection criteria for the TwoTrackKs HLT1 line

The selection criteria chosen for the TwoTrackKs HLT1 line are reported in Table 5.2. This set of selections is chosen as it satisfies the decided rate cost constraint, having an output rate of 60 kHz, 6% of the total budget of HLT1. The efficiency of the reported selections on $D^0 \rightarrow K_S^0 K_S^0$ decays that have at least one K_S^0 candidate decaying within the VELO acceptance is of 11%, considering candidates directly triggered by the new line. This is twice the efficiency of all other triggers for this mode, as we will discuss in greater detail in Section 5.6. Therefore, the optimized selections successfully achieve a sizeable efficiency on signal decays with a minimal rate cost.




Figure 5.14: Distributions of $FD(K_S^0)$ with respect to the PV for $D^0 \rightarrow K_S^0 K_S^0$ and Minimum Bias events. All distributions are normalized such that their integral is equal to one.

Variable	Threshold
$\chi^2/\text{ndf}(\pi)$	< 2.5
$ m(K_S^0) - 498 \text{ MeV}/c^2 $	$< 45 \text{ MeV}/c^2$
$\chi_{vtx}^2(K_S^0)$	< 20
$\eta(K_S^0)$	$2 < \eta(K_S^0) < 4.2$
$\cos(\theta_{DIRA})$	> 0.99
$\cos(\theta_{\pi\pi})$	> 0.99
$p_T(\pi)$	$> 470 \text{ MeV}/c$
$p(\pi)$	$> 5 \text{ GeV}/c$
$\chi_{IP}^2(\pi)$	> 50
$p_T(K_S^0)$	$> 2500 \text{ MeV}/c$
$IP(\pi^+) \times IP(\pi^-)/IP(K_S^0)$	$> 0.72 \text{ mm}$

Table 5.2: Selections applied by **TwoTrackKs** HLT1 line.

5.3 Extension to K_S^0 pairs selection

The work performed in formulating the **TwoTrackKs** line can be exploited to further extend the approach to create an even more exclusive HLT1 line dedicated to pairs of K_S^0 (**TwoKs**). This is an even stronger innovation of the trigger approach in LHCb, that assigns to HLT1 only inclusive selection tasks. However, it is clear that this innovation may be very fruitful, in a future-oriented vision in which ever increasing specificity of the triggers will be required to handle increasing luminosities.

The purpose of introducing this line in Run 3 is somewhat different from the **TwoTrackKs** line. In addition to experimenting with a highly exclusive selection at the first trigger level in the LHC environment, its primary objective is to ensure the availability of an HLT1 line with a good efficiency for $D^0 \rightarrow K_S^0 K_S^0$ decays and an extremely limited rate cost throughout the entire data-taking. This can be important at the beginning of the run with the upgraded trigger system, where detector and trigger

performance could deviate from expectations, potentially necessitating a tightening of the trigger selections. Such adjustments could lead to substantial efficiency losses, which can be mitigated by leveraging a line with minimal rate cost and sufficient efficiency—one that would not require further tightening. In addition, the potential of the single- K_S^0 line to accommodate a rather wide number of decay modes of interest, might lead in the future to its being tuned for wider purposes; the permanence of a dedicated line for K_S^0 pairs with stable criteria is clearly a useful additional tool. The requirement for the presence of two K_S^0 candidates within the same event substantially reduces acceptance for background candidates, enabling the application of looser selections on individual K_S^0 candidates while maintaining robust signal efficiency. It should be remembered that events are never duplicated in the HLT1 process, so the candidates firing both triggers are only acquired once; therefore, there is no issue of overcrowding of the output of HLT1 by adding lines with large overlaps.

The **TwoKs** line is naturally designed as a direct extension of the **TwoTrackKs** line, so its selections will be based on the already identified variables, applied identically to both K_S^0 candidates. For selections excluded from the numerical optimization in the **TwoTrackKs** line, the same thresholds will be retained, as the rationale behind their definition remains valid when selecting two K_S^0 candidates instead of one. A numerical optimization performed exploiting the same TMVA rectangular optimizer is performed to choose thresholds on remaining variables. No requirements are applied on the K_S^0 candidates mother particle in the baseline scenario. This configuration is preferred as it extends the utility of the **TwoKs** line to additional decays different from $D^0 \rightarrow K_S^0 K_S^0$, as the $B^0/B_s^0 \rightarrow K_S^0 K_S^0$, or decays with additional final state particles. In case rate cost will be too high in this configuration, background rejection can be increased by applying requirements on the invariant mass of the parent particle of the two K_S^0 candidates.

Following the numerical optimization, the thresholds are again determined as the configuration achieving the best efficiency at a given rate cost. In this case, the optimization targets a cost of approximately 1%. This is decidedly appropriate for such an exclusive selection at the very first trigger level. It is also expected and hoped that a non negligible fraction of this rate will be overlapping with the **TwoTrackKs** rate.

5.3.1 Selection criteria for the TwoKs HLT1 line

The selections determined for the **TwoKs** line are reported in Table 5.3. It is worth noting that the optimizer sets the $p_T(\pi)$ threshold to the lowest possible value within the established limits. This observation will become particularly relevant when analyzing the line’s efficiencies, as discussed in the following section. The reported selections match the required rate constraint, having a rate of 11 kHz. The efficiency achieved on $D^0 \rightarrow K_S^0 K_S^0$ decays is of 4%. The **TwoKs** line achieves an efficiency that is 36% of the **TwoTrackKs** line’s one, while having just 17% of its rate cost.

5.4 K_S^0 lines impact on HLT1 throughput

The impact of the two newly implemented K_S^0 -dedicated lines on HLT1 throughput has been measured to be negligible. This was verified by testing HLT1 throughput on the same A5000 GPU cards intended for Run 3 data-taking, using a simulated sample of

Variable	Threshold
$\chi_{trk}^2/dof(\pi)$	< 2.5
$ m(\pi^+\pi^-) - m(K_S^0) $	$< 45 \text{ MeV}/c^2$
$\chi_{vtx}^2(K_S^0)$	< 20
$\eta(K_S^0)$	$2 < \eta(K_S^0) < 4.2$
$\cos(\theta_{DIRA})$	> 0.99
$\cos(\theta_{\pi\pi})$	> 0.99
$p_T(\pi)$	$> 425 \text{ MeV}/c$
$p(\pi)$	$> 3 \text{ GeV}/c$
$\chi_{IP}^2(\pi)$	> 15
$p_T(K_S^0)$	$> 1150 \text{ MeV}/c$
$IP(\pi^+) \times IP(\pi^-)/IP(K_S^0)$	$> 0.23 \text{ mm}$

 Table 5.3: Selections applied by **TwoKs** HLT1 line.

Run 3 Minimum Bias events. The minimal impact of these lines is due to the fact that they do not require reconstructing additional tracks or vertices, as HLT1 already performs secondary vertex identification. The **TwoTrackKs** line introduces only a minor computational overhead, as it merely applies selection criteria to two-track candidates, incurring a negligible rate cost. Similarly, the **TwoKs** line applies selections to pairs of K_S^0 candidates without performing an additional vertex fit for the parent particle. The number of secondary vertex combinations is efficiently managed by applying minimal requirements to K_S^0 candidates, which remain looser than the final selection criteria.

5.5 Lines rate cost

A critical factor in the viability of the proposed HLT1 lines is their rate cost, which is stringently limited by the maximum allowable output rate of 1 MHz for HLT1. The estimates obtained from the minimization process must therefore be verified by a more careful analysis, accounting also for the context given by the presence of all other HLT1 lines. The output rate for any line, or set of lines is calculated as:

$$Rate = \frac{n_{pass}}{n_{tot}} \times R_{input}$$

where n_{pass} is the number of Minimum Bias events retained, n_{tot} is the total number of Minimum Bias events, and R_{input} represents the input rate at the given trigger stage. For HLT1, this input rate is 30 MHz.

The output rate for a given set of selections is estimated using the **MooreAnalysis** tool [158, 159], the official LHCb utility for calculating trigger line rates and efficiencies. This tool exploits the official LHCb codebase to provide rate and efficiency estimates for both HLT1 and HLT2 trigger stages. In this case, the line rates are evaluated using a sample of simulated Minimum Bias events under the nominal Run 3 conditions of LHCb.

The first parameter considered is the exclusive rate of the lines, calculated as if the respective line were the sole active component in the HLT1 system. This value is significant as it reflects the absolute acceptance of background events by the implemented selections,

and it is the one exploited in the choice of selections. The exclusive rates of the newly developed lines are measured and compared to those of the **TrackMVA** and **TwoTrackMVA** lines. The exclusive rate values for the HLT1 lines under consideration are summarized in Table 5.4. The selected configurations for the **TwoTrackKs** and **TwoKs** HLT1 lines

Line	Exclusive rate [kHz]
TrackMVA	351 ± 10
TwoTrackMVA	647 ± 14
TwoTrackKs	60 ± 4
TwoKs	11 ± 2

Table 5.4: Exclusive rates of considered HLT1 lines.

successfully achieve their goal of utilizing only a small fraction of the total HLT1 output rate budget. Specifically, the estimated exclusive rates for these lines correspond to approximately 6% and 1% of the total output budget of 1 MHz, respectively.

The impact of the new lines on the output rate of the overall HLT1 system can be lower than the simple sum of exclusive rates, because of the existing overlap between events selected by multiple lines. In this case, it has been decided to consider the impact of the addition of the novel lines to an HLT1 system only composed by the **TrackMVA** and **TwoTrackMVA** lines. This is motivated by the fact that these lines cover the majority of LHCb physics case, and in particular the selection of our processes of interest in the Run 2, and occupy $\sim 90\%$ of the HLT1 output rate, making them a good benchmark for the purposes considered in these studies. The rates of various combinations from the considered group of lines is reported in Table 5.5.

HLT1 selection	Rate [kHz]	Increase [kHz]
TrackMVA \vee TwoTrackMVA	857 ± 16	–
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs	896 ± 16	39
TrackMVA \vee TwoTrackMVA \vee TwoKs	864 ± 16	7
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs \vee TwoKs	901 ± 16	44
TwoTrackKs \vee TwoKs	69 ± 5	–

Table 5.5: Rates of the considered groups of HLT1 lines. Last column reports the difference between the considered HLT1 selection and **TrackMVA** \vee **TwoTrackMVA**.

The first parameter to assess is the output rate of events triggered by either the **TrackMVA** or **TwoTrackMVA** lines, as this serves as the baseline HLT1 output rate for the scenario under consideration. This baseline rate is 857 kHz, confirming that the two lines adequately represent the bulk of the HLT1 rate. When the **TwoTrackKs** line is added to

this system, the overall output rate increases by 39 kHz, corresponding to approximately 4% of the total HLT1 output budget. Notably, this represents a $\sim 33\%$ reduction compared to the exclusive rate of the **TwoTrackKs** line alone. Similarly, the addition of the **TwoKs** line increases the rate by 7 kHz. If both novel lines are included in the HLT1 system, the total output rate rises by 44 kHz, equivalent to 4.4% of the overall HLT1 budget.

From these numbers it can be seen that the novel lines' impact on HLT1 rate is at the level of a few percent of the total output rate budget. In particular, **TwoKs** line rate is almost negligible in the considered system, as its impact corresponds to less than 1% of the total HLT1 rate budget, thanks to the requirement for the presence of an additional K_S^0 candidate in the same event.

This result is particularly significant and deserves attention. The development of a set of selections capable of triggering on K_S^0 decays with a sizeable efficiency at the first trigger level in a hadron collider environment, while maintaining a finite output rate, is far from being an expected result. Given the enormous production rate of K_S^0 particles and the large track multiplicity in pp collisions at the LHC, these lines might in principle have required a substantially larger fraction of the HLT1 output rate. Such a scenario would have rendered their inclusion in the LHCb trigger unfeasible. The fact that this challenge has been successfully addressed highlights the feasibility of this general approach to triggering, and the quality of the new HLT1 system of LHCb Run 3.

5.6 Impact on $D^0 \rightarrow K_S^0 K_S^0$ collection efficiency

After having assessed the affordability of K_S^0 -based HLT1 selections, it is crucial to estimate their efficiency on selecting signal decays. In fact, the achievement of a sensible efficiency gain motivates the adoption of such selections during data-taking. The **MooreAnalysis** tool is utilized to estimate the efficiencies of HLT1 lines, using the same sample of simulated $D^0 \rightarrow K_S^0 K_S^0$ decays employed during the selection optimization process. To produce meaningful figures, these efficiencies are normalized to the number of simulated events in which all final-state particles fall within the LHCb acceptance and can be reconstructed as long tracks (referred to as 'reconstructible tracks'). The efficiency is then defined as:

$$\epsilon = \frac{n_{sel,match}}{n_{acc}}$$

where $n_{sel,match}$ is the number of events retained by the considered selection (or group of selections) in which the candidate triggering the decision is matched to a true simulated candidate. Meanwhile, n_{acc} is the number of events with all final-state tracks tagged as reconstructible. As with the rate analysis, a reference HLT1 system consisting solely of the **TrackMVA** and **TwoTrackMVA** lines is used as a baseline for comparison.

Efficiencies are computed considering both individual HLT1 lines and groups of them, and relative values are reported in Table 5.6.

Looking at the **TwoTrackKs** line performance, it can be seen that the chosen set of selections successfully reaches the target of achieving a high efficiency on $D^0 \rightarrow K_S^0 K_S^0$ decays, as its standalone efficiency is ~ 3.8 times higher than the **TrackMVA** line, and its addition to the system composed by the **TrackMVA** and **TwoTrackMVA** line produces an efficiency increase by a factor of ~ 2.6 . The efficiency of the **TwoKs** line demonstrates that the adopted set of selections successfully achieves its intended target. Despite occupying only approximately 1% of the HLT1 output rate, the line delivers an efficiency greater than

HLT1 selection	ϵ [%]
TrackMVA	2.9 ± 0.9
TwoTrackMVA	2.6 ± 0.8
TwoTrackKs	11.0 ± 1.6
TwoKs	3.9 ± 1.0
TrackMVA \vee TwoTrackMVA	5.5 ± 1.2
TwoTrackKs \vee TwoKs	12.3 ± 1.7
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs	14.4 ± 1.8
TrackMVA \vee TwoTrackMVA \vee TwoKs	9.1 ± 1.5
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs \vee TwoKs	15.7 ± 1.9

Table 5.6: Efficiency of considered HLT1 selections on $D^0 \rightarrow K_S^0 K_S^0$ decays.

that of the **TrackMVA** line. It is worth noting, however, that its efficiency is significantly lower than that of the **TwoTrackKs** line. This difference is primarily attributable to the reconstruction cuts applied at the HLT1 level, which must be satisfied by all four tracks in the final state for the **TwoKs** HLT1 line, ultimately limiting its efficiency.

This represents a pivotal result. The simulation studies demonstrate the feasibility of implementing a K_S^0 -dedicated trigger selection at the earliest trigger level, achieving a substantial increase in efficiency for signal decays with a minimal rate cost. This achievement is unprecedented at a hadron collider and is made possible by the execution of an early reconstruction. This approach enables the application of decay-specific selections that leverage unique signatures to effectively suppress background acceptance while enhancing signal decay efficiency.

5.7 Run 3 HLT2 $D^0 \rightarrow K_S^0 K_S^0$ lines

5.7.1 Rate reduction of $D^0 \rightarrow K_S^0 K_S^0$ HLT2 lines

In addition to the development of the novel **TwoTrackKs** and **TwoKs** HLT1 lines, it is also necessary to work on the HLT2 side of the selections for collecting $D^0 \rightarrow K_S^0 K_S^0$ decays. The minimal goal was to adapt the HLT2 selections used during Run 2 to the updated Run 3 framework for the LL, LD, and DD decay channels. The responsibility for this effort was assigned to me, as one of the experts involved in past analyses for this specific decay channel. However, given the introduction of the novel HLT1 lines, and the significantly increased data flow due to the new DAQ system, coupled with a higher rate of signal events, a revision of the Run 2 selections, leading to reductions in the line rates, was a necessary step to take. While the $D^0 \rightarrow K_S^0 K_S^0$ HLT2 lines are not expected to consume a substantial portion of the output bandwidth, reducing the background event collection remains a valuable objective.

The proposed strategy to achieve a reduced line rate involves tightening certain existing selections. To avoid compromising the efficiency for signal decays, adjustments were limited to selections expected to maintain nearly 100% efficiency for signal events. Specifically, selections based on invariant mass distributions were targeted, in particular those for $m(\pi\pi)$, $m(K_S^0 K_S^0)$, and Δm . These mass windows were set in Run 2 much wider than the ranges occupied by the signal distributions, suggesting that modest tightening could safely reduce background acceptance without affecting signal efficiency. However, the width of these mass windows cannot be reduced indiscriminately. Regions populated by background candidates are crucial for accurately determining the background distribution shape. Therefore, dedicated studies were conducted to evaluate the extent to which the mass windows could be narrowed without compromising the precision of background shape estimation and extrapolation into the signal region.

To implement the proposed changes, three additional HLT2 lines labeled as “tight” have been introduced. These lines are designed to be nearly identical to the nominal lines, (re-labeled as “loose”), except for the application of tighter thresholds on the invariant mass selections. This strategy was adopted to provide the system with enhanced flexibility. This dual-line approach enables seamless switching between configurations. For example, the loose lines can be turned off if a rate reduction is required or reactivated if the system can accommodate a higher rate. This setup offers a straightforward and efficient mechanism to adapt the trigger configuration based on evolving operational needs, without the complication of variable thresholds that would unavoidably complicate the physics data analysis in a significant way.

The $D^0 \rightarrow K_S^0 K_S^0$ sample from the 2017–2018 dataset was used to tune the tighter selection criteria, with separate analyses conducted for the LL, LD, and DD decay channels. To reduce background contamination and replicate conditions similar to those in the physics analysis, specific pre-selections are applied. These include the L0 trigger requirements used in the analysis and invariant mass selections. When analyzing a particular mass distribution (e.g., $m(\pi\pi)$), selections are applied to the other two distributions ($m(K_S^0 K_S^0)$ and Δm) to improve sample purity.

The method for tuning tighter selections involves determining the minimum threshold that does not adversely affect the performance of the physics analysis. This threshold is then adopted for the tight lines. The procedure involves iteratively fitting the invariant mass distributions with models that include both signal and background components, progressively narrowing the mass window widths. For each iteration, the number of signal (N_{sig}) and background (N_{bkg}) events is estimated through the fits. Mass window reductions are applied symmetrically around the nominal $m(K_S^0)$ and $m(D^0)$ values in their respective distributions, while for the Δm quantity, only the upper threshold is adjusted.

It is important to note that when fitting the data with the same model, variations in N_{sig} are primarily attributable to statistical fluctuations, even though the estimates are correlated. To account for potential systematic effects, alternative background models that are equally valid are also used in the fits.

The trends in N_{sig} across different fits with varying models are then analyzed to set new thresholds. The new selections are established at the point where N_{sig} begins to show significant deviations, either due to the reduced fit range or the adoption of a different model. This ensures that the tighter selections maintain compatibility with the performance requirements of the physics analysis.

LL channel

The invariant mass distributions for LL candidates where studies are performed are reported in Figure 5.15. Considered candidates are the ones passing the same L0 selections




Figure 5.15: Δm (top), $m(K_S^0 K_S^0)$ (center) and $m(\pi\pi)$ (bottom) invariant mass distributions for LL events. The L0 and invariant mass selections applied are explicitly reported in the text.

exploited in Run 2 analysis and the following invariant mass selections:

- $|m(K_S^0 K_S^0) - 1865.6 \text{ MeV}/c^2| < 20 \text{ MeV}/c^2$ and $|m(\pi\pi) - 497.6 \text{ MeV}/c^2| < 10.5 \text{ MeV}/c^2$ selections are applied for candidates shown in the Δm distribution
- $|\Delta m - 145.5 \text{ MeV}/c^2| < 1.5$ and $|m(\pi\pi) - 497.6 \text{ MeV}/c^2| < 10.5 \text{ MeV}/c^2$ selections are applied for candidates shown in the $m(K_S^0 K_S^0)$ distribution
- $|m(K_S^0 K_S^0) - 1865.6 \text{ MeV}/c^2| < 20 \text{ MeV}/c^2$ and $|\Delta m - 145.5 \text{ MeV}/c^2| < 1.5 \text{ MeV}/c^2$ selections are applied for candidates shown in the $m(\pi\pi)$ distribution

A Gaussian model is used to describe the signal contribution in the Δm distribution. This choice is favored over alternative functions, such as the sum of two Gaussians or a JohnsonSU model, because, while it may not perfectly describe the data, it offers greater stability across the extensive number of fits required for the study. Additionally, the estimates provided by the Gaussian model are sufficiently accurate for the objectives of this analysis. Three different models are considered to describe the background distribution:

- $b(\Delta m) \propto \sqrt{\Delta m} \cdot (1 + a \cdot \Delta m + b \cdot \Delta m^2)$
- $b(\Delta m) \propto \sqrt{\Delta m} \cdot (1 + c \cdot \Delta m + d \cdot \Delta m^2 + e \cdot \Delta m^3)$
- $b(\Delta m) \propto 1 - e^{-f \cdot \Delta m} + g \cdot \Delta m$

where a, \dots, g are free parameters extracted in the fit.

The upper threshold for Δm is systematically varied from 170 MeV/c^2 to 145 MeV/c^2 in 10 equally spaced increments. For the sake of brevity, complete fit results are not included here. The corresponding N_{sig} estimates, obtained from fits performed across various ranges and models, are presented in Figure 5.16 as a function of the upper threshold applied to Δm . The reported values clearly indicate that the different fit models exhibit




Figure 5.16: The N_{sig} values extracted with fits to different Δm ranges and exploiting different background models for LL channel.

varying trends, underscoring the value of the adopted procedure. Based on these results, it is reasonable to set the Δm upper threshold for the LL HLT2 tight line at 160 MeV/c^2 , as this is the point where the estimates from different models begin to diverge, indicating a growing systematic uncertainty in the background subtraction.

A similar procedure is repeated for the $m(K_S^0 K_S^0)$ distribution. A Gaussian model is exploited to describe the signal contribution, and polynomials of grade zero, one, and two are adopted to describe the background component:

- $b(m(K_S^0 K_S^0)) \propto a$
- $b(m(K_S^0 K_S^0)) \propto b + c \cdot m(K_S^0 K_S^0)$
- $b(m(K_S^0 K_S^0)) \propto d + e \cdot m(K_S^0 K_S^0) + f \cdot m(K_S^0 K_S^0)^2$

where a, \dots, f are free parameters extracted in the fit.

The mass windows are kept symmetrical around the PDG $m(D^0)$ value, while their width is reduced from 180 MeV/c^2 to 90 MeV/c^2 in 10 equally spaced steps. The N_{sig} values extracted with different configurations are reported in Figure 5.17. Here a smaller




Figure 5.17: The N_{sig} values extracted with fits to different $m(K_S^0 K_S^0)$ ranges and exploiting different background models for LL channel.

spread is present between three adopted models, showing similar trends. However, it has been decided to set the tight line sideband mass window at the value of 150 MeV/c^2 . This is done to prioritize caution, even N_{sig} does not show a significant variation for thresholds below this value. Additionally, this decision is also supported by the exploitation of candidates falling in $m(K_S^0 K_S^0)$ sidebands in the training of the kNN algorithm in the most recent $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement, that should not be reduced too much.

Analyzing the $m(\pi\pi)$ distribution, a Gaussian model is again exploited to describe the signal component, while the same polynomial models exploited in the study of $m(K_S^0 K_S^0)$ are exploited to model the background. The mass windows are symmetrical around the PDG $m(K_S^0)$ value, and their width is reduced from 70 MeV/c^2 to 52 MeV/c^2 in 10 equally spaced steps. The N_{sig} values extracted in different configurations are reported in Figure 5.18. In this case, the considered estimates show a significant spread already with the unmodified sidebands. This is not alarming, considering that the difference is of $\sim 2\sigma$ and that exploited models are quite approximate. Only the increase of the spread between estimates has been considered for the purposes of this study, and the $m(\pi\pi)$ sideband range for the tight line has been set to 60 MeV/c^2 .

The Δm , $m(K_S^0 K_S^0)$, and $m(\pi\pi)$ thresholds for the LL loose and tight HLT2 lines are summarized in Table 5.7.




Figure 5.18: The N_{sig} values extracted with fits to different $m(\pi\pi)$ ranges and exploiting different background models for LL channel.

Variable	Loose threshold [MeV/c ²]	Tight threshold [MeV/c ²]
Δm	$\in [-75, 170]$	$\in [-75, 160]$
$m(D^0)$	$\in [1775, 1955]$	$\in [1790, 1940]$
$m(K_S^0)$	$\in [462.6, 532.6]$	$\in [467.6, 527.6]$

Table 5.7: Summary of the Δm , $m(K_S^0 K_S^0)$ and $m(\pi\pi)$ mass windows exploited for the LL tight and loose lines.

LD channel

A similar procedure is carried out for LD events using the previously mentioned 2017–2018 sample. In this case, the study focuses exclusively on the downstream-reconstructed K_S^0 , as no significant differences are expected compared to the LL sample for the long-reconstructed K_S^0 . Consequently, the thresholds identified for the LL sample are directly applied to the long-reconstructed K_S^0 in the LD configuration.

The same selections applied on the LL channel are exploited here to improve the sample's purity. The Δm , $m(K_S^0 K_S^0)$ and $m(\pi\pi)$ distributions after the application of these selections are reported in Figure 5.19. The signal and background models used to describe the Δm distribution for the LL channel are also employed here, along with the same incremental steps for mass window reduction. However, the exponential background model is excluded from this analysis due to its difficulty in achieving fit convergence, a consequence of the poorer S/B ratio in the LD sample. The N_{sig} values extracted from fits to the Δm distributions under various configurations are shown in Figure 5.20. The two considered models show a similar trend up to an upper value of ~ 155 MeV/c² for the accepted mass window. However, since this can be due to the low signal statistics, and the exploitation of just two models, I decided to apply a conservative approach and to set the upper limit for the tight lines threshold at the same value set for the LL line, hence 160 MeV/c².

For the $m(K_S^0 K_S^0)$ LD channel, the same signal and background models used for the LL channel are applied, along with identical mass window reduction steps. The number of signal events estimated from fits over varying ranges is presented in Figure 5.21. No




Figure 5.19: Δm (top), $m(K_S^0 K_S^0)$ (center) and $m(\pi\pi)$ (bottom) invariant mass distributions for LD events. Only the downstream-reconstructed K_S^0 candidates are reported in the $m(\pi\pi)$ distribution. The L0 and invariant mass selections applied are explicitly reported in the text.

significant differences are observed between the various estimates. However, due to the limited signal sample statistics and the use of $m(K_S^0 K_S^0)$ candidates from the sidebands in the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ analysis, a conservative approach has been adopted. The tight threshold is therefore set at 150 MeV/c^2 , consistent with the choice made for the LL sample.

In performing the study about the mass window adopted to select the K_S^0 -reconstructed




Figure 5.20: The N_{sig} values extracted with fits to different Δm ranges and exploiting different background models for LD channel.




Figure 5.21: The N_{sig} values extracted with fits to different $m(K_S^0 K_S^0)$ ranges and exploiting different background models for LD channel.

candidate in the LD channel a Gaussian model is again exploited to model the signal contribution. The models adopted to describe the background contribution are:

- $b(m(\pi\pi)) \propto a + b \cdot m(\pi\pi) + c \cdot m(\pi\pi)^2$
- Gaussian model

where a, b, c are parameters extracted in the fit. These are different from the models exploited in case of the long-reconstructed K_S^0 as a poorer S/B is present when downstream tracks are involved. The mass window is centered around the K_S^0 PDG mass value and it's varied from 128 MeV/c^2 to 110 MeV/c^2 , in 10 equally spaced steps. The N_{sig} values extracted from fits in different configurations are shown in Figure 5.22. Estimates obtained using different models reveal a discrepancy even at the nominal mass window width, with no additional spread observed as the range is reduced. Given the relatively small differences in N_{sig} estimates and the simplicity of the background models used, this is not expected to




Figure 5.22: The N_{sig} values extracted with fits to different $m(\pi\pi)$ ranges and exploiting different background models for the downstream-reconstructed K_S^0 of the LD channel.

create any significant issues. The tight line's sideband range has been set to 120 MeV/c^2 . This decision was not driven solely by reduced sensitivity to the number of signal events but rather by the observed background shape in the $m(\pi\pi)$ distribution. As shown in Figure 5.23, the background exhibits two distinct trends: one resembling the behavior observed in the background of long-reconstructed K_S^0 , and another with a narrower range that results in two "bumps" around approximately 460 MeV/c^2 and 540 MeV/c^2 . Due to this peculiar trend, the cause of which remains unclear, the mass window for the tight selection was not further reduced to ensure a reliable and accurate description of the background shape.

The Δm , $m(K_S^0 K_S^0)$, and $m(\pi\pi)$ thresholds for the LD loose and tight HLT2 lines are summarized in Table 5.8.

Variable	Loose threshold [MeV/c^2]	Tight threshold [MeV/c^2]
Δm	$\in [-75, 170]$	$\in [-75, 160]$
$m(D^0)$	$\in [1775, 1955]$	$\in [1790, 1940]$
$m(K_{SL}^0)$	$\in [462.6, 532.6]$	$\in [467.6, 527.6]$
$m(K_{SD}^0)$	$\in [433.6, 561.6]$	$\in [437.6, 557.6]$

Table 5.8: Summary of the Δm , $m(K_S^0 K_S^0)$ and $m(\pi\pi)$ mass windows exploited for the LL tight and loose lines.

DD channel

The Δm , $m(K_S^0 K_S^0)$, $m(\pi\pi)$ distributions obtained for the DD candidates after the application of the selections exploited for the LL and LD channels are reported in Figure 5.23. While the LL and LD candidates exhibit a clearly visible signal peak in all three distributions, enabling a fitting procedure, this is not the case for DD candidates. No discernible signal peak is observed in the Δm and $m(K_S^0 K_S^0)$ distributions, making it impossible to apply the same method used for the LL and LD samples. Nevertheless, since




Figure 5.23: Δm (top), $m(K_S^0 K_S^0)$ (center) and $m(\pi\pi)$ (bottom) invariant mass distributions for DD events. The L0 and invariant mass selections applied are explicitly reported in the text.

no significant differences are expected between the distributions of LD and DD events, the tighter selections identified for the LD channel can be extended to the DD channel for the Δm and $m(K_S^0 K_S^0)$ cuts. For the $m(\pi\pi)$ distribution, a signal peak is clearly present, and it would, in principle, be possible to perform a similar optimization exercise. However, given the expected consistency between LD and DD downstream-reconstructed K_S^0 distributions, the tighter selections determined for the LD sample can be confidently

extended to the DD sample as well.

The Δm , $m(K_S^0 K_S^0)$, and $m(\pi\pi)$ thresholds for the DD loose and tight HLT2 lines are summarized in Table 5.9. The rate reduction achieved through the use of the tight

Variable	Loose threshold [MeV/ c^2]	Tight threshold [MeV/ c^2]
Δm	$\in [-75, 170]$	$\in [-75, 160]$
$m(D^0)$	$\in [1775, 1955]$	$\in [1790, 1940]$
$m(K_S^0 K_S^0)$	$\in [433.6, 561.6]$	$\in [437.6, 557.6]$

Table 5.9: Summary of the Δm , $m(K_S^0 K_S^0)$ and $m(\pi\pi)$ mass windows exploited for the LL tight and loose lines.

HLT2 lines is estimated directly using the 2017–2018 sample employed in the studies, which serves as a representative dataset for candidates collected by the HLT2 lines with looser selections. The estimation is carried out by applying the tighter selection criteria to the sample that implements the baseline (loose) selection. The ratio of the size of the sample with tight selections to that with loose selections provides a direct measure of the rate reduction achieved by the tight HLT2 lines. This reduction is reported in Table 5.10 for LL, LD, and DD events. The thresholds identified for the tight HLT2 selections

HLT2 line	Rate(Loose)/Rate(Tight)
LL	2.7
LD	1.9
DD	2.2

Table 5.10: Ratio between rates of Loose and Tight HLT2 configurations for LL, LD and DD HLT2 lines. The value is extracted exploiting the 2017–18 $D^0 \rightarrow K_S^0 K_S^0$ sample.

successfully achieve a significant reduction in the line rates, fulfilling the purpose for which they were designed. It is important to highlight how this reduction does not compromise the efficiency for signal events, as the selection criteria remain much broader than the region occupied by the signal candidates. Additionally, the studies conducted confirm that the narrowed sidebands will not affect the ability to accurately model the background shape beneath the signal peak.

5.7.2 Tuning of HLT2 lines

The analysis methodology for the measurement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$ presented in Chapter 4 refrains from applying any bias to the fraction of secondary decays in the $D^0 \rightarrow K_S^0 K_S^0$ sample, also at the trigger level. However, selections were applied at the HLT2 level on $\tau(K_S^0)$, $\chi_{\text{IP}}^2(K_S^0)$ and χ_{vvd}^2 , as reported in Tables 4.2 and 4.4. Their impact on the secondaries fraction is expected to be small, and a systematic effect is assigned on the possible remaining difference in the fractions in the two samples.

However, since the implementation of the tight lines allows a significant rate reduction, it has been decided to release these thresholds in preparation for the start of Run 3. Even if this might lead to a lower sample purity, it can then be improved with offline selections.

The rate increase resulting from the removal of these selections could not be quantified due to the unavailability of a minimum bias simulated sample filtered with HLT1 selections. However, considering that the HLT2 lines responsible for selecting $D^0 \rightarrow K_S^0 K_S^0$ decays make up only a small fraction of the overall HLT2 bandwidth, and that the use of tight lines ensures a significant rate reduction, it has been decided to begin data collection in Run 3 with this configuration, keeping open the possibility of further tightening of the selections according to necessity.

This chapter described the development and preparation of the triggers responsible for selecting $D^0 \rightarrow K_S^0 K_S^0$ decays during Run 3, covering both the HLT1 and HLT2 levels. The newly introduced HLT1 selections turned out to be very successful, showing a significant efficiency enhancement. Given the large improvement and the potential for these advancements to benefit other K_S^0 final-state decays as well, I will devote the next chapter to the examination of some notable examples of these broader applications. These considerations have been important in the process of gathering enough motivation for obtaining the approval of the whole collaboration for the introduction of these modifications to the LHCb trigger.

Chapter 6

Spin-offs of K_S^0 triggers

The studies presented in the previous chapter demonstrated that the implementation of a K_S^0 -dedicated set of selections at the HLT1 level can significantly enhance the online selection efficiency for $D^0 \rightarrow K_S^0 K_S^0$ decays. These results can be extended both to other channels and to a larger acceptance, including also the possibility to trigger on downstream tracks. I start showing how the developed triggers, although specifically designed and optimized for $D^0 \rightarrow K_S^0 K_S^0$ decays, can yield substantial benefits for a variety of other decays involving K_S^0 final states, even with the current configuration. Then, I discuss the possibility of extending this trigger strategies to Λ decays. Finally, I detail the studies conducted to evaluate the efficiency gains and rate cost of extending the LLP-dedicated trigger strategies I developed and validated for long tracks to downstream tracks.

6.1 Impact on efficiency for K_S^0 final state decays

6.1.1 $B^0/B_s^0 \rightarrow K_S^0 K_S^0$

The decays $B^0/B_s^0 \rightarrow K_S^0 K_S^0$ present an interesting physics case, as outlined in Section 2.1, and stand to gain significantly from the introduction of the novel HLT1 lines, given their shared final state with $D^0 \rightarrow K_S^0 K_S^0$ decays. Employing the same methodology described in Section 5.6, the HLT1 selection efficiency for B^0 decays has been evaluated. The results are summarized in Table 6.1. For brevity, only the B^0 decay is considered in detail; however, no notable differences in performance are anticipated for B_s^0 decays. Due to the higher p_T of B^0 decay products compared to those from D^0 decays, the **TwoTrackKs** and **TwoKs** lines achieve efficiencies of approximately 60% and 25%, respectively—similar to the efficiencies of the **TrackMVA** and **TwoTrackMVA** lines. This is a remarkable result, highlighting the potential of the novel HLT1 lines for selecting decays of *beauty* hadrons. Despite the already substantial efficiency of the **TrackMVA** and **TwoTrackMVA** lines, the inclusion of the new lines enhances the overall efficiency by an additional 20%, resulting in a total efficiency nearing 80%.

6.1.2 $D^0 \rightarrow K_S^0 \pi^- \pi^+$

The $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays are selected at the HLT1 level based mostly on the detached π pair coming from the D^0 decay, by the **TrackMVA** and **TwoTrackMVA** lines. This is the

HLT1 selection	ϵ [%]
TrackMVA	57.6 ± 6.1
TwoTrackMVA	22.7 ± 5.2
TwoTrackKs	56.1 ± 6.1
TwoKs	24.2 ± 5.3
TrackMVA \vee TwoTrackMVA	63.6 ± 5.9
TwoTrackKs \vee TwoKs	57.6 ± 6.1
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs	77.3 ± 5.2
TrackMVA \vee TwoTrackMVA \vee TwoKs	68.2 ± 5.7
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs \vee TwoKs	77.3 ± 5.2

Table 6.1: Efficiency of considered HLT1 selections on $B^0 \rightarrow K_S^0 K_S^0$ decays.

selection providing most of the sample employed in the latest measurement of charm CP -violating parameters [160], made of 30M $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays. However, this trigger strategy results in a significant bias on the D^0 lifetime distribution, that is the source of an important systematic effect. In this analysis, the use of an alternative HLT1 selection solely based on the K_S^0 decay would allow the collection of a D^0 decay whose lifetime distribution is unbiased by online selection. The efficiencies of the new HLT1 lines on $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays are reported in Table 6.2 - only the **TwoTrackKs** line efficiency is reported, since only a single K_S^0 is present in the final state. The **TwoTrackKs** line shows a

HLT1 selection	ϵ [%]
TrackMVA	2.4 ± 0.5
TwoTrackMVA	5.2 ± 0.7
TwoTrackKs	2.7 ± 0.5
TrackMVA \vee TwoTrackMVA	6.3 ± 0.7
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs	8.2 ± 0.8

Table 6.2: Efficiency of considered HLT1 selections on $D^0 \rightarrow K_S^0 \pi^- \pi^+$ decays.

comparable efficiency to the traditional **TrackMVA** and **TwoTrackMVA** lines. Consequently, a dataset collected exclusively by the **TwoTrackKs** line would have approximately half the statistics of one collected using the traditional lines. However, the reduction in statistical precision would be compensated by the elimination of the systematic bias associated with the current selection strategy. This trade-off becomes particularly advantageous in scenarios with large integrated luminosities, such as those expected by the end of Run 3, when the sample size will not be the driver of the overall resolution anymore. This trigger strategy is also very appropriate for the purpose of providing a suitable calibration sample

for the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ analysis. The actual viability of this approach will be practically demonstrated in Chapter 8.

6.1.3 $D^+ \rightarrow K_S^0 \pi^+$ and $D_s^+ \rightarrow K_S^0 K^+$

It is interesting to consider the $D^+ \rightarrow K_S^0 \pi^+$ and $D_s^+ \rightarrow K_S^0 K^+$ decays as these have been used as calibration samples in the latest measurement of $\mathcal{A}^{CP}(K^+ K^-)$ performed by LHCb exploiting Run 2 data [161]. Efficiencies for the considered HLT1 lines are reported in Table 6.3 and 6.4 for $D^+ \rightarrow K_S^0 \pi^+$ and $D_s^+ \rightarrow K_S^0 K^+$ decays, respectively.

HLT1 selection	ϵ [%]
TrackMVA	5.1 ± 0.5
TwoTrackMVA	2.1 ± 0.3
TwoTrackKs	2.5 ± 0.3
TrackMVA \vee TwoTrackMVA	6.3 ± 0.5
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs	8.0 ± 0.6

Table 6.3: Efficiency of considered HLT1 selections on $D^+ \rightarrow K_S^0 \pi^+$ decays.

HLT1 selection	ϵ [%]
TrackMVA	3.5 ± 0.4
TwoTrackMVA	1.4 ± 0.2
TwoTrackKs	2.6 ± 0.3
TrackMVA \vee TwoTrackMVA	4.5 ± 0.4
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs	6.2 ± 0.4

Table 6.4: Efficiency of considered HLT1 selections on $D_s^+ \rightarrow K_S^0 K^+$ decays.

The efficiency of the **TwoTrackKs** line turns out comparable to that of the **TrackMVA** and **TwoTrackMVA** lines. Notably, when the **TwoTrackKs** line is included alongside **TrackMVA** and **TwoTrackMVA**, the overall efficiency increases by 27% for $D^+ \rightarrow K_S^0 \pi^+$ decays and 38% for $D_s^+ \rightarrow K_S^0 K^+$ decays. This efficiency boost could play a crucial role during Run 3 data-taking. Specifically, the statistics for $D^0 \rightarrow K^+ K^-$ decays may grow more rapidly than those of its calibration channels, primarily because the **TrackMVA** line requires stricter selection criteria compared to the **TwoTrackMVA** line with respect to the Run 2 configuration, as **TwoTrackMVA** benefits from a two-track requirement. This disparity could amplify the impact of limited statistics in the calibration sample on the measurement of $\mathcal{A}^{CP}(K^+ K^-)$, a challenge that would be mitigated by the increased efficiency provided by the **TwoTrackKs** line. Furthermore, the **TwoTrackKs** line remains strategically valuable if tighter **TrackMVA** selections become necessary due to data-taking constraints, as its low rate and reliance on large track detachments makes it well-suited to maintain strong efficiency under such conditions.

6.2 Studies on a Λ -dedicated HLT1 line

The approach to K_S^0 triggering presented up to now can be readily extended to the selection of Λ baryon decays. Like K_S^0 mesons, Λ baryons are neutral, long-lived particles with relatively low invariant mass. These shared characteristics result in Λ decays facing similar selection challenges as those encountered with K_S^0 mesons. This similarity provides a strong motivation for exploring the implementation of a Λ -dedicated HLT1 line, hereafter referred to as **TwoTrackLambda**.

6.2.1 TwoTrackLambda selections

As done with the **TwoTrackKs** and **TwoKs** lines, it has been decided to exploit a set of rectangular selections defining the line. Because of the similarities between K_S^0 and Λ decays, the same variables identified in the case of the K_S^0 -dedicated HLT1 lines will be exploited for the **TwoTrackLambda** line. Numerical optimization will be performed only for a part of the variable thresholds, as reported in Section 5.2.

Distributions for all the considered variables will not be reported here for brevity. However, it is necessary to discuss how the asymmetric final state of the Λ is handled during selections. While the K_S^0 decays into a pion pair, making the two final state particles interchangeable, this is not the case for the Λ , that is reconstructed in the $p\pi^-$ final state. Notably, in the case of $\Lambda \rightarrow p\pi^-$ decays, it is possible to distinguish between the proton and pion without relying on the standard LHCb PID information. This distinction is made possible by the significant mass difference between the proton and pion, which ensures that the proton consistently exhibits a higher p_T than the pion. This behavior is well illustrated in the distribution of the p_T difference between the proton and pion, shown in Figure 6.1, derived from a simulated sample of $\Lambda_b^0 \rightarrow \Lambda\mu^-\mu^+$ decays. This is a critical




Figure 6.1: Distribution of the difference between the p_T (p) and p_T (π^-) for Λ candidates present in a simulated sample of $\Lambda_b^0 \rightarrow \Lambda\mu^-\mu^+$ decays.

aspect of the trigger selections, as it ensures the correct mass hypothesis is assigned to the daughter particles when computing the Λ invariant mass. This, in turn, allows for a

significantly narrower mass window around the PDG value for $m(\Lambda)$, thereby reducing the acceptance of random track combinations. Specifically, a mass window of $\pm 35 \text{ MeV}/c^2$ around $m(\Lambda)$ has been chosen. The invariant mass distribution for Λ candidates from a simulated sample of $\Lambda_b^0 \rightarrow \Lambda \mu^- \mu^+$ decays is shown in Figure 6.2, with the applied selection indicated by the two vertical lines. Additionally, the ability to reliably distinguish between




Figure 6.2: Invariant mass distribution for Λ candidates from a simulated sample of $\Lambda_b^0 \rightarrow \Lambda \mu^- \mu^+$ decays. Vertical lines represent selections applied on the variable for the `TwoTrackLambda` HLT1 line, hence $|m(p\pi^-) - m(\Lambda)| < 35 \text{ MeV}/c^2$.

p and π^- enables the application of optimized selection criteria specifically targeting $p_T(p)$ and $p_T(\pi^-)$. This refinement has a significant impact, as the π^- is substantially softer compared to the p , as illustrated in the distributions shown in Figure 6.3.

After having identified this strategy for the distinction of final state particles, the selections are tuned. For the optimization procedure, a sample of Λ candidates originating from a simulated $J/\psi \rightarrow \Lambda \bar{\Lambda}$ decay sample was used as the signal. Only candidates decaying within the VELO are taken into account. The background sample consisted of two-track vertices reconstructed from a simulated minimum-bias dataset, with the sole requirement that their invariant mass fell within a $\pm 70 \text{ MeV}/c^2$ mass window around $m(\Lambda)$.

The variables whose selections have been set before the numerical optimization are: $\chi^2(p/\pi^-)$, $p/p_T(\pi^-)$, $m(p\pi^-)$, $\chi_{vtx}^2(\Lambda)$, $\eta(\Lambda)$, $\cos(\theta_{DIRA})$, $\cos(\theta_{p\pi^-})$. It is worth noting how $p(\pi^-)$ and $p_T(\pi^-)$ have been excluded from the numerical optimization procedure. Because of the low p and p_T of the final state pion, thresholds have been set to a value close to the HLT1 default reconstruction cuts, hence $p_T(\pi^-) > 425 \text{ MeV}/c$ and $p(\pi^-) > 2.5 \text{ GeV}/c$. Other selections have been tuned exploiting the same kNN rectangular cuts optimizer adopted in case of `TwoTrackKs` and `TwoKs` selections. The selections chosen for the `TwoTrackLambda` line are reported in Table 6.5.

6.2.2 HLT1 Λ line rate cost

The rate cost of the `TwoTrackLambda` line is estimated using the same procedure reported in Section 5.5, adding it to an HLT1 system composed of the two general-purpose `TrackMVA`

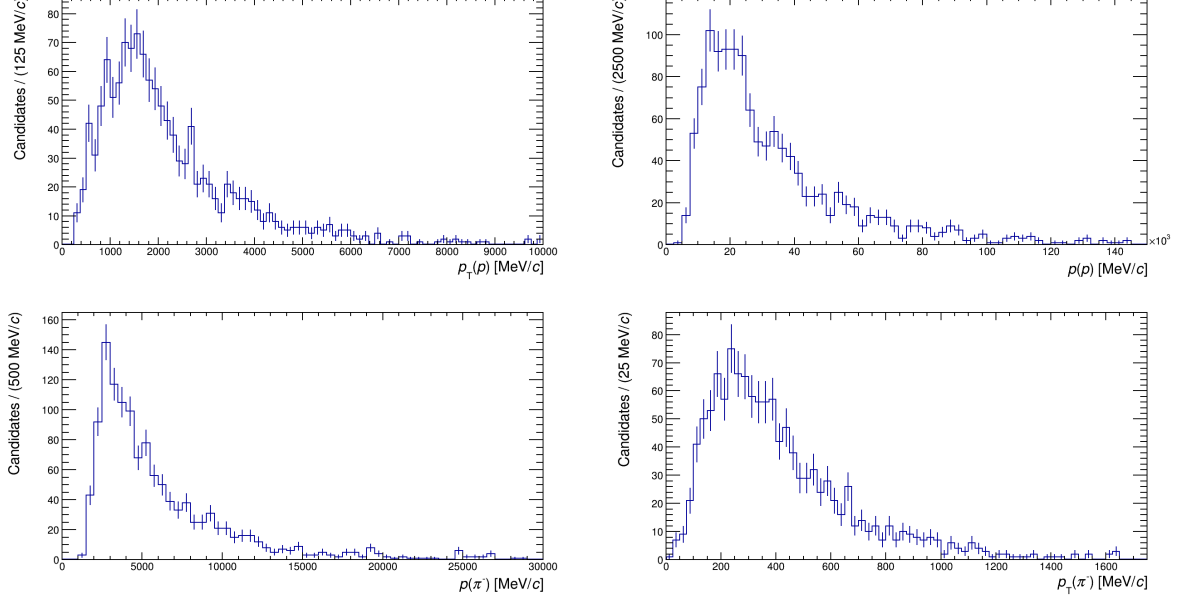


Figure 6.3: Distribution of p (left) and p_T (right) for the p (top) and π^- (bottom) candidates coming from the Λ decay for the simulated sample of $\Lambda_b^0 \rightarrow \Lambda \mu^- \mu^+$ decays.

Variable	Selection
$\chi^2(p, \pi^-)$	< 2.5
$ m(p\pi^-) - m(\Lambda) $	$< 35 \text{ MeV}/c^2$
$\chi_{vtx}^2(\Lambda)$	< 20
$\eta(\Lambda)$	$2 < \eta(\Lambda) < 5$
$\cos(\theta_{DIRA})$	> 0.99
$\cos(\theta_{p\pi^-})$	> 0.99
$p_T(\pi^-)$	$> 425 \text{ MeV}/c$
$p(\pi^-)$	$> 2.5 \text{ GeV}/c$
$\chi_{IP}^2(p, \pi^-)$	> 55
$p_T(p)$	$> 1.3 \text{ GeV}/c$
$p(p)$	$> 18 \text{ GeV}/c$
$p_T(\Lambda)$	$> 1150 \text{ MeV}/c$
$\frac{IP(p) \times IP(\pi^-)}{IP(\Lambda)}$	$> 1.3 \text{ mm}$

Table 6.5: Selections applied by **TwoTrackLambda** HLT1 line.

and **TwoTrackMVA** lines. The rate for the considered HLT1 selections are reported in Table 6.6. The identified selections have a limited impact on the existing HLT1 rate. The **TwoTrackLambda** has a negligible impact when added to the considered HLT1 system. Its exclusive rate is of 43 kHz and its addition to the considered HLT1 system causes a rate increase of 34 kHz. Its impact is therefore limited to 3 – 4% of the overall HLT1 output rate budget.

HLT1 selection	Exclusive rate [kHz]
TrackMVA	351 ± 10
TwoTrackMVA	647 ± 14
TwoTrackLambda	43 ± 4
TrackMVA \vee TwoTrackMVA	1033 ± 17
TrackMVA \vee TwoTrackMVA \vee TwoTrackLambda	1067 ± 18

Table 6.6: Rate of the considered selections when **TwoTrackLambda** line is added to the HLT1 system.

6.2.3 Impact of TwoTrackLambda line on selected measurements

The efficiency increase caused by the addition of the **TwoTrackLambda** line to the HLT1 system is estimated for a set of representative decays, using the same procedure reported in Section 5.6. In particular, the considered decays are: $J/\psi \rightarrow \Lambda \bar{\Lambda}$. All of the estimates are performed exploiting a sample of decays simulated under Run 3 nominal conditions. Efficiencies for the considered decays are reported in Table 6.7. Let's first consider the

HLT1 selection	$\epsilon(J/\psi \rightarrow \Lambda \bar{\Lambda})$ [%]	$\epsilon(\Lambda_b^0 \rightarrow \Lambda \mu^- \mu^+)$ [%]	$\epsilon(B^+ \rightarrow p \bar{\Lambda})$ [%]
TrackMVA	6.3 ± 1.9	44 ± 1	54 ± 2
TwoTrackMVA	3.8 ± 1.5	53 ± 1	20 ± 2
TwoTrackLambda	6.3 ± 1.9	7 ± 1	10 ± 1
TrackMVA \vee TwoTrackMVA	8.2 ± 2.2	62 ± 1	58 ± 2
TrackMVA \vee TwoTrackMVA \vee TwoTrackLambda	12.7 ± 2.6	64 ± 1	60 ± 2

Table 6.7: Efficiency of different HLT1 selections on the considered physics samples.

$J/\psi \rightarrow \Lambda \bar{\Lambda}$ decays. Its final state only composed of Λ decays, makes it extremely suited for the exploitation of a Λ -dedicated HLT1 line. The **TwoTrackLambda** line demonstrates an efficiency comparable to the **TrackMVA** line and higher than the **TwoTrackMVA** line, enhancing the overall system efficiency by $\sim 55\%$. This is a promising performance, given the **TwoTrackLambda** rate cost being ~ 20 times smaller than the two general-purpose HLT1 lines one. It is important to note that the impact of the **TwoTrackLambda** line is significantly smaller compared to the effect observed for **TwoTrackKs** on $D^0 \rightarrow K_S^0 K_S^0$ decays. This difference can be attributed to the pronounced momentum imbalance between the two final-state particles in a Λ decay. The high momentum carried by the proton enhances the efficiency of the **TrackMVA** line for Λ decays, whereas the pion, with its lower momentum, exhibits reduced efficiency already at the reconstruction stage. This combined effect decreases the relative impact of the **TwoTrackLambda** line on the HLT1 system efficiency. This condition is even worse in case of $\Lambda_b^0 \rightarrow \Lambda \mu^- \mu^+$ and $B^+ \rightarrow p \bar{\Lambda}$ decays. The efficiency for these decays presents a similar situation, where the **TwoTrackLambda** line has a smaller efficiency with respect to **TrackMVA** and **TwoTrackMVA** lines, and its impact on the overall system is limited to a few percent increase. In both cases decays

are efficiently triggered exploiting either the $\mu^-\mu^+$ pair or the p present in the final state, significantly reducing the impact of the addition of the **TwoTrackLambda** line.

6.2.4 Conclusions Λ -dedicated HLT1 line

The potential implementation of a Λ -dedicated HLT1 line was investigated as a natural extension of the **TwoTrackKs** HLT1 line, which has demonstrated a significant impact on the HLT1 efficiency for various decays. However, due to the pronounced momentum imbalance between the final-state particles of Λ decays and the stringent reconstruction-level cuts applied by the HLT1 tracking algorithm, the efficiency estimates for Λ decays indicate that the efficiency-to-rate cost ratio achieved by the **TwoTrackLambda** line is relatively modest. Despite the limited rate cost of the **TwoTrackLambda** line, accounting for only approximately 3% of the total HLT1 bandwidth, it was decided not to pursue its implementation at this stage. It is important to emphasize, however, that this conclusion is based on the constraints of the current HLT1 reconstruction algorithm. Should it become possible to relax these reconstruction cuts in the future, the implementation of a Λ -dedicated HLT1 line could prove to be beneficial.

This consideration becomes even more relevant in the context of the LHCb Upgrade II [162], where the experiment aims to increase the instantaneous luminosity to $1.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ - a factor of 7.5 higher than in Run 3 and Run 4. In this high-luminosity regime, tighter thresholds on one- and two-track triggers will be necessary to handle the increased event rate and track multiplicity. On the other hand, a low-rate selection strategy based on the identification of more complex candidates, such as the **TwoTrackLambda** might become extremely relevant, as it is expected to be more robust, maintaining good signal efficiency even under these demanding conditions. A clear example of this principle is the **TwoKs** HLT1 line. By requiring an additional K_S^0 candidate, it improves the efficiency-to-rate ratio compared to the **TwoTrackKs** line, as demonstrated in Section 5.6 through simulation studies. This advantage will also be confirmed using real data in Section 7.5.

6.3 Extension to downstream tracks

The usefulness of a trigger on long-lived states is obviously related to its acceptance as a function of the path length. In the current reconstruction, only long tracks are available at HLT1 level. This has limited the new trigger lines to the use of long tracks. This, of course, limits our acceptance to the particles decaying within the limited volume of VELO acceptance. This limitation is somewhat significant for K_S^0 , but becomes much more important for Λ^0 , and potentially interesting exotic states, of the kind mentioned in Chapter 2.

However, this is not necessarily a limitation that cannot be overcome. To address the issue of the reconstruction of downstream tracks at the first level, a specific proposal has been advanced for a dedicated, real-time “downstream tracker” (DWT) [163], starting already back in 2016. This device, based on commercial FPGA cards, is intended to perform substantial portions of event reconstruction in the back-end of the DAQ system, before event building and HLT1 processing, enabling the inclusion of downstream tracks in the HLT1-level triggers starting from Run 4.

More recently, work has been ongoing to explore the possibility of introducing downstream reconstruction even in the current Run 3, by enlarging the GPU farm by 50%, and optimizing the algorithms for greater speed [164].

6.3.1 Benefits of trigger-level downstream tracking

It is therefore very interesting to extend the studies of previous sections to consider the advantages and feasibility of possible downstream-based trigger strategies. To this purpose, an important first step is to estimate the order of magnitude of the possible gain for several decay topologies.

Yield increase for decay modes with an LLP

The fraction of LLPs decaying in the downstream region has been estimated here for a selection of decay channels involving K_S^0 and Λ decays. A simple phase-space model is employed to describe the final-state particle distributions, while the underlying momentum distribution of promptly produced b - and c -hadrons is derived from measurements of b -hadron production at 8 TeV [165].

While this approach is an approximation, it is sufficient to meet the precision needs of a preliminary study. The results of this calculation, shown in Table 6.8, highlight the substantial fraction of LLP decays that could be recovered by incorporating downstream tracks into the trigger selection, underscoring the potential impact of this enhancement.

Decay Mode	(D+L)/L
$\Lambda_c^+ \rightarrow \Lambda \pi^+$	3.3
$D^0 \rightarrow K_S^0 \pi^+ \pi^-$	2.5
$B^0 \rightarrow J/\psi K_S^0$	2.1
$B_s^0 \rightarrow K_S^0 K^- \pi^+$	2.0
$B^+ \rightarrow D^0 K^+, D^0 \rightarrow K_S^0 \pi^+ \pi^-$	1.8
$\Lambda_b^0 \rightarrow J/\psi \Lambda$	2.5

Table 6.8: The lifetime acceptance for selected decay modes with K_S^0 and Λ in the final state is presented. The acceptance is expressed as the ratio (D+L)/L, which represents the number of decays that can be reconstructed using both downstream and long tracks compared to those reconstructed with long tracks alone.

Consider the specific example of the decay $D^0 \rightarrow K_S^0 \pi^+ \pi^-$, where the K_S^0 decays in the downstream region. In this scenario, the daughters of the K_S^0 are unavailable to the HLT1 long track-based trigger, meaning the $\pi^+ \pi^-$ system alone must satisfy the initial trigger requirements. This contrasts with cases where the K_S^0 decays within the VELO acceptance, allowing all four final-state pions to contribute to the HLT1 trigger conditions. (There may also be a limited number of TIS candidates, where tracks not belonging to the decay satisfy the trigger. However, these are often excluded from analyses.)

In the Run 1 mixing analysis of $D^0 \rightarrow K_S^0 \pi^+ \pi^-$, the downstream yield made up for approximately 50% of the long-track yield [166]. In the Run 2 update [167], candidates with a downstream K_S^0 in the final state were excluded from the final sample used for the measurement. Clearly, a downstream tracker capable of providing HLT1 trigger

information for downstream K_S^0 decays would significantly enhance the downstream yield in this channel.

A preliminary estimate of the potential improvement from incorporating downstream tracks in the HLT1 selection can be derived by examining physics channels where LHCb triggers efficiently without relying on the downstream portion of the decay. Table 6.9 presents the yields for two such decays. These modes rely on triggering via the $J/\psi \rightarrow \mu^+\mu^-$ portion of the final state, making the trigger independent of the decay location of the K_S^0 or Λ . In these examples, downstream decays of the K_S^0 and Λ account for 71% and 74% of the final sample, respectively.

Decay Mode	D/L yield in data	Reference
$B^0 \rightarrow J/\psi K_S^0$	2.5	[65]
$\Lambda_b^0 \rightarrow J/\psi \Lambda$	2.9	[168]

Table 6.9: The value of D/L ratio for candidate yields in data from LHCb analyses.

Modes with only LLPs in the final state

For decay modes involving only LLPs in the final state, such as $D^0 \rightarrow K_S^0 K_S^0$, the ability to trigger on downstream decays becomes even more crucial. Categorizing events into LL, LD, and DD categories, the lifetime acceptances are 12%, 48%, and 40%, respectively. This means that 40% of decays lack any long track, and a significant 88% of decays involve at least one downstream K_S^0 decay. The actual event yields will depend on the trigger and reconstruction efficiencies. Table 6.10 summarizes the event yields achieved in the two published LHCb analyses of this channel [44, 149, 169].

Category	Run 1		Run 2 2015 + 2016		Run 2 2017 + 2018	
	N_{sig}	% of total	N_{sig}	% of total	N_{sig}	% of total
LL	367	59%	1566	77%	4486	74%
LD	165	25%	469	23%	1494	25%
DD	95	15%	0	0%	87	1%

Table 6.10: Different K_S^0 categories yields for LHCb Run 1 and Run 2 $D^0 \rightarrow K_S^0 K_S^0$ analyses. The absence of DD candidates in 2015-16 sample is due to the requirement for at least one long track in the decay applied at the HLT2 level.

These numbers clearly highlight a significant deficit of downstream decays compared to the lifetime acceptance figures mentioned earlier. A similar observation applies to the channel $B_s^0 \rightarrow K_S^0 K_S^0$ [57], where the sample is evenly split between the LL and LD categories, while the DD category was excluded due to low trigger efficiency.

Improved trigger algorithms

The ability to trigger on downstream LLPs can also play a significant role in reducing biases in existing trigger selections. For instance, the current trigger for the channel

$D^0 \rightarrow K_S \pi^+ \pi^-$ imposes a requirement for large impact parameters on the final-state pions. This condition introduces a decay-time bias, which contributes to systematic uncertainties in the charm mixing analysis using this channel [166, 167], as previously discussed. By enabling early inclusion of downstream K_S candidates in the trigger, the impact parameter requirement could be relaxed, thereby reducing systematic uncertainties. It is important to notice that the feasibility of this trigger strategy has already been demonstrated in 2024 data-taking. A sample of $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ has been collected with a trigger without *any* impact parameter requirement on the pion pair coming from the D^0 , and it has been exploited as a calibration sample for the measurement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$, as it will be shown in Chapter 8. This sample only comprises long-reconstructed K_S^0 , but it can be trivially extended to downstream tracking, as selections do not rely on D^0 pointing variables.

More broadly, early triggering on LLP candidates would allow for more relaxed selection criteria on other aspects of the decay. For example, one could reduce p_T thresholds, ease particle identification requirements, or even adjust other kinematic constraints. Additionally, downstream tracks available at an early trigger stage could also be incorporated into muon identification, further enhancing trigger capabilities and reducing biases.

6.3.2 Efficiency gain on benchmark decay modes and rate cost

The efficiency gain due to the presence of downstream tracking at the first level of the trigger is discussed here. This is done by analyzing a set of benchmark decay modes, with one or two K_S^0 in the final state.

Strategy

Since the technical details of the tracking algorithm intended to run on the FPGA system are yet to be finalized, we assume that downstream tracks will have a quality comparable to HLT2-level tracks. To evaluate the efficiency, samples simulated under Run 3 conditions are utilized.

The efficiency is defined as:

$$\epsilon = \frac{\text{number of reconstructed candidates}}{\text{number of reconstructible candidates}},$$

where the definition of a reconstructible track is tailored to long and downstream tracks, considering their different detector acceptances.

For the rate estimation, a sample of Run 2 minimum bias data is used, where all possible two-track vertices have been reconstructed (opposite-charged tracks, without applying PID information). Real data is preferred over simulation because generators have significant uncertainties in accurately modeling minimum bias events.

The rate is then computed using the following formula:

$$\text{Rate} = 30000 \text{ kHz} \times a \times \frac{\text{PU}(\text{Run 3-4})}{\text{PU}(\text{Run 2})},$$

where 30000 kHz is the HLT1 input rate, a represents the acceptance of a given selection (defined as the number of events passing the selection divided by the total number of events in the sample), and $\text{PU}(\text{Run}^*)$ denotes the pile-up for a specific Run of data-taking.

The pile-up for Run 3 (and Run 4) is taken as 5.35 [170]. The Run 2 pile-up, extracted from the LHCb metadata, is 1.64. The pile-up ratio ensures that the acceptance calculated on Run 2 data is properly adjusted for the higher luminosity conditions anticipated in Run 3.

Efficiency gain

To estimate the potential increase in yields achievable with downstream tracking available at the first level of the trigger, the following set of trigger lines is considered:

1. The current cut-based Run 3 implementation of the **TrackMVA** and **TwoTrackMVA** lines, which require the presence of one or two long tracks. These are governed by the selections detailed in Section 4.3.2;
2. A prototype line specifically designed for selecting K_S^0 candidates, reconstructed from both long and downstream tracks. This line, derived directly from the **TwoTrackKs** line, is detailed in Table 6.11 and is the only trigger line capable of selecting downstream tracks. When considering the extension to downstream tracks, a wider mass window is used, to compensate for the different invariant mass resolution.

Next, we examine the impact of these trigger selections on a few representative decay channels.

Variable	Cut
$\chi_{trk}^2/ndof(h_{1/2})$	< 2.5
$p_T(h_{1/2})$	$> 470 \text{ MeV}/c$
$\chi_{IP}^2(h_{1/2})$	> 50
$\chi_{vtx}^2(h_1, h_2)$	< 20
$\eta(h_1, h_2)$	$> 2 \wedge < 4.2$
$ m(h_1, h_2) - m(K_S^0) $	$< 45 \text{ MeV}/c^2 \text{ (long)}$ $60 \text{ MeV}/c^2 \text{ (downstream)}$
$p_T(h_1, h_2)$	$> 2500 \text{ MeV}/c$
$\cos(\theta_{DIRA})(h_1, h_2)$	> 0.99
$\cos(\theta_{opening}(h_1, h_2))$	> 0.99
$IP(h_1) \cdot IP(h_2)/IP(h_1, h_2)$	> 0.72

Table 6.11: K_S^0 prototype-line's selections. As before, h_1, h_2 are the two particles in input to the trigger line. Two different mass windows are used for long and downstream tracks.

Using a simulated sample of benchmark decays simulated under Run 3 conditions, and considering the trigger selections previously described, the efficiency gain in selecting such decays is estimated when downstream tracks are reconstructed at the earliest stage of the trigger. In particular, the considered decays are $D^0 \rightarrow K_S^0 K_S^0$, $D^0 \rightarrow \pi^+ \pi^-$, $B^0 \rightarrow K_S^0 K_S^0$ and $B^0 \rightarrow \pi^+ \pi^-$. For the charmed hadron decays, the D^0 is required to come from a $D^{*+} \rightarrow D^0 \pi^+$ decay, to allow flavour tagging.

The efficiency gain obtained by extending HLT1 to downstream triggering is reported in Table 6.12. A significant improvement in efficiency is observed for decays with a fully LLP final state, as predicted based on acceptance considerations.

Sample	Efficiency long [%]	Efficiency long + downstream [%]	Efficiency gain
$D^0 \rightarrow K_S^0 K_S^0$	8.0%	26.3%	3.3
$D^0 \rightarrow K_S^0 \pi^+ \pi^-$	35.4%	40.4%	1.14
$B^0 \rightarrow K_S^0 K_S^0$	12.2%	68%	5.6
$B^0 \rightarrow K_S^0 \pi^+ \pi^-$	65.6%	71.7%	1.09

Table 6.12: Efficiency in case of long or long and downstream triggering for some benchmark K_S^0 final state samples. Samples simulated at the nominal Run 3 conditions have been exploited.

Rate cost

Introducing a trigger line that selects downstream tracks also increases the acceptance for background events, which in turn raises the overall HLT1 rate. To estimate this rate increase, we considered a trigger system consisting of the three lines currently under consideration: the **TrackMVA** line, the **TwoTrackMVA** line, and the K_S^0 prototype line, which is the only one capable of triggering downstream tracks. The results of this estimate are presented in Table 6.13. The presented numbers show that the exploited emulator is able

HLT1 selection	Rate [kHz]
1-track	583 ± 57
2-track	1036 ± 76
1-track \vee 2-track	1518 ± 92
K_S^0 line (long)	80 ± 20
1-track \vee 2-track \vee K_S^0 line (long tracks only)	1518 ± 92
K_S^0 -Line (long and downstream triggering)	343 ± 43
1-track \vee 2-track \vee K_S^0 line (long and downstream tracks)	1767 ± 99

Table 6.13: Rate estimates for the considered HLT1 configuration. The K_S^0 line ran both only on long tracks, and including downstream tracks.

to reliably simulate the behavior of the HLT1 system, even if at first order. In fact, the rate estimates obtained for the long-based lines approximately match those estimated using the official LHCb tool, reported in Table 7.4. The extension of the considered HLT1 K_S^0 -dedicated line to downstream tracking determines a rate increase for the line alone of approximately a factor of 4. This is an encouraging result, indicating that the rate remains well under control. The estimated factor-of-four increase aligns with expectations, given that the signal yield is predicted to double, while a slightly larger background contribution arises from the wider mass window used for downstream tracks. Regarding the overall

impact on the HLT1 system, the effect is modest, with the total output rate increasing by only about 15%.

These estimates strongly suggests the feasibility of the considered trigger strategy. In fact, by increasing the HLT1 rate by such a limited amount, the efficiency for the selected benchmark modes can be enhanced by factors ranging from 2 to nearly 7.

These promising results played a key role in motivating the installation of the considered DWT in Run 4. My contribution has been documented in the LHCb public note [171], and in the recently approved *Technical Design Report* of the project [172].

Chapter 7

New triggers commissioning with Run 3 data

In this chapter I detail the efforts undertaken to integrate the developed HLT1 lines in the LHCb trigger framework for use during Run 3 data-taking. The initial focus is on preparatory work to allow the required tunability of our K_S^0 -dedicated selections to accomodate different operational scenarios. Subsequently, the chapter describes the process of testing and commissioning the novel HLT1 lines with the earliest real data collected in Run 3.

7.1 Tunability of selections

The results of the simulation studies reported in Chapters 5 and 6 motivated the inclusion of the `TwoTrackKs` and `TwoKs` lines in the Run 3 LHCb trigger system. For actual operation, the triggers need to be integrated in the official LHCb framework, that requires all HLT1 triggers to be tunable through appropriate handles (tuning parameters). This means that, in addition to the optimal set of selections that was developed in previous Chapters, a number of other settings need to be pre-defined, allowing tuning of the line rate in the most effective way. The reason for this requirement is the possibility that the detector and trigger system performance might deviate from the design one, requiring tightening of selections to maintain an output rate below 1 MHz. This is likely to happen especially at the start of a new run with a completely renewed detector. In addition, the simultaneous presence of a long list of triggers, that can be changed or updated, with some of them having large rate sensitivities to small changes in data-taking conditions, and with unpredictable overlap to each other, makes it virtually impossible to predict with accuracy the trigger rate in real data taking. This situation is handled at LHCb by the so-called *bandwidth division*, an automated procedure for tuning the parameters of a set of trigger selections, aiming to ensure an optimized sharing of the available HLT1 output bandwidth among HLT1 lines, for an effective collection of data across the entire LHCb physics program. This has already been in existence in Run 1 and Run 2 [173] by LHCb, for the allocation of L0 line rates. In doing this integration, the sophistication of the new K_S^0 triggers presents a new challenge. Most of the other HLT1 trigger lines have a very small number of parameters that makes sense to tune, so they are readily incorporated in the system. In the K_S^0 case, changing a single parameter of the set the

was simultaneously optimized is likely to quickly affect the trigger effectiveness, unless care is taken to tune all the other parameters as well.

A first step was taking the decision, in agreement with the trigger managers, that additional selections configurations are only needed for the **TwoTrackKs** line. For the **TwoKs** line, the set of selections implemented for Run 3 data-taking is fixed to those reported in Table 5.3. This is motivated by the **TwoKs** line’s very small expected rate, making it a negligible perturbation of the HLT1 system. It should be noted that the line has in fact been designed to behave like this, and this decision perfectly matches its purposes of being able to provide a minimal efficiency for $D^0 \rightarrow K_S^0 K_S^0$ decays in any data-taking condition, as discussed in the previous chapters.

Let us go now into some more details of the bandwidth division. This procedure takes as input the list of HLT1 lines whose selections have to be optimized, and a series of physics channels for which efficiency is maximized. The output of the procedure is the thresholds configuration maximizing overall efficiency across all the considered physics channels, at a fixed output rate of 1 MHz. This is done by varying the tunable parameters provided for each line. For example, this is done for the **TrackMVA** line by tuning the χ_{IP}^2 - p_{T} selection reported in Figure 5.3. To overcome the problem of the lack of a single parameter that can be conveniently tuned, the **TwoTrackKs** line acceptance has been tuned by defining a set of discrete tuning steps, each including changed of multiple parameters simultaneously.

Additionally, these discrete steps are designed with monotonically increasing or decreasing thresholds, forming a series of “concentric” selections. That is, care has been taken to ensure that every event passing a tighter selection would be passing any other looser selection. This approach simplifies the offline analysis of data collected with this HLT1 line. By ensuring that each successive selection is a subset of the previous one, any data collected under different selection configurations during data-taking will naturally form nested sub-samples, facilitating consistent and straightforward analysis.

Each of these additional configurations is identified through a numerical optimization similar to the one already exploited in 5. The same TMVA rectangular cuts optimizer is adopted, together with the same sets of K_S^0 candidates from minimum bias and $D^0 \rightarrow K_S^0 K_S^0$ events as background and signal samples. The procedure followed to identify additional steps is:

- a set of selections is taken as starting point - the “nominal” ones reported in Table 5.2 are the first considered;
- the optimizer is constrained to search for selections that are entirely tighter/looser than the ones of the previous step;
- the numerical optimization is ran and the the tighter/looser cuts are identified;
- the procedure is repeated taking the newly identified configuration as new starting point.

In this procedure, an upper bound is applied during the optimization on the p_{T} (K_S^0) threshold, for the identification of selection configurations tighter than the nominal one. This avoids any additional tightening, as the nominal selection already applies a hard cut of $2.5 \text{ GeV}/c^2$.

Eight different configurations are identified exploiting this procedure. For brevity, just the efficiency and rate for different configurations will be reported here, see Figure 7.1. The rate for different configurations is computed following the same procedure and sample




Figure 7.1: Efficiency and rate of different **TwoTrackKs** line configurations for the bandwidth division procedure. The set of selections reported in Section 5.6 is the configuration identified with number 3.

reported in Section 5.5. The efficiency considered here considers solely the **TwoTrackKs** line and is computed using an offline emulator of the HLT1 selections, rather than the **MooreAnalysis** tool. This approach was chosen due to the significantly shorter processing time required by the offline emulator, without compromising the validity of the study's results. While small differences may exist compared to values calculated using the LHCb official code, these are not expected to affect the monotonicity of the observed trend.

The identified selections ensure the required flexibility for the **TwoTrackKs** line, as their rate spans from 75 kHz to 15 kHz, allowing an overall rate variation by a factor of 5. It is also important to notice how the tightest configuration has a rate close to that of the **TwoKs** line, which has been excluded by this procedure exactly because of its small rate. Therefore, the possibility to tune the **TwoTrackKs** line to have such a small rate should ensure its exploitation in the majority of data-taking conditions. The maximum efficiency variation is a factor of 4.6. It is important to note that the significantly lower efficiency of tighter configurations, compared to the nominal one, does not defeat their purpose. In fact, in case conditions require such stringent selections for the **TwoTrackKs** line, an even harder selections tightening would be needed for the high-rate lines as **TrackMVA**, that would end up having an almost negligible efficiency on K_S^0 final state decays. Therefore, even under such harsh conditions, the **TwoTrackKs** line would still enhance efficiency for $D^0 \rightarrow K_S^0 K_S^0$ decays, along with the **TwoKs** line, whose selections will remain unchanged.

Accompanied by this set of tunings, our K_S^0 lines have been integrated into the LHCb trigger framework, and adopted for physics data taking for the whole of Run 3.

7.2 Commissioning with early Run 3 data

After a three-year-long stop, the LHC resumed operations for Run 3 in July 2022, alongside LHCb, which began data-taking after the upgrade installation. During the initial months of Run 3, the commissioning of the detector and DAQ systems was carried out, enabling the first reconstruction of online tracks and vertices from pp collisions before the end of the year. The first data collected with the upgraded system were, of course, of critical importance to the collaboration, as their analysis allows for prompt identification and resolution of any discrepancies with expectations. This is particularly relevant for the novel K_S^0 -dedicated HLT1 lines, which implement a trigger strategy that had never been utilized before, carrying a higher risk of unexpected deviations from expected behavior. Therefore, after having designed, implemented and integrated those lines in the LHCb trigger, I focused on debugging and testing their correct behavior in the field.

The behavior of the K_S^0 lines was checked using early Run 3 data as soon as it became available. It is important to note that these data were necessarily collected with a new detector, whose behavior was initially scarcely understood. Therefore, the purpose of the studies described here was to identify, at the earliest possible stage, the presence of any bugs or significant discrepancies between the performance expected from simulations and that observed in real data.

The samples used during the early commissioning of the novel HLT1 lines were collected at an instantaneous luminosity below the nominal Run 3 level for LHCb. This allowed the entire system to be tested under less challenging conditions before transitioning to nominal operation. The instantaneous luminosity in these tests is characterized by the parameter μ , which represents the average number of visible pp interactions per bunch crossing. Nominal Run 2 and Run 3 conditions correspond to $\mu = 1.1$ and $\mu = 5.5$, respectively. The analyzed samples are:

- S_1 - 23 nb^{-1} of data collected at $\mu = 0.1$
- S_2 - 13 nb^{-1} of data collected $\mu = 1.1$
- S_3 - 69 nb^{-1} of data collected at $\mu = 3.1$

These samples were processed only by HLT1, which applied the nominal selections for the K_S^0 lines, as no tuning of the selections had been performed at such an early stage. A small fraction of these samples was also saved immediately after event building, without undergoing any processing. These unprocessed data are only used for rate computation.

7.2.1 Commissioning the TwoTrackKs line

The candidates collected by the K_S^0 lines have been accessed with an offline, CPU-compiled version of HLT1, dumping information for the selected candidates. This is a necessary step, as the online version of HLT1 does not store reconstruction for offline analysis.

The first crucial check is to verify the actual presence of real K_S^0 decays among the candidates selected by the **TwoTrackKs** line. This is a basic but crucial check in such early conditions, that might fail (as it actually often happens) because of several conditions, such as a preliminary alignment, a background component much higher than in simulation, and also simple bugs in the implementation of line's selections. This check is performed analyzing the invariant mass of candidates collected on sample S_1 , reported in Figure 7.2.

The sample S_1 has been preferred at the time of these studies because it was the first made available. The distribution clearly shows a peak in the expected invariant mass region,




Figure 7.2: Invariant mass of K_S^0 candidates selected by **TwoTrackKs** HLT1 line on sample S_1 . A fit to the distribution is superimposed. The σ_1 , σ_2 , μ parameters respectively refer to the standard deviation of the two Gaussians and the shared mean of the two. N_{sig} and N_{bkg} are the number of signal and background candidates extracted in the fit.

exhibiting a nearly Gaussian shape with good purity. This provides strong evidence that the novel **TwoTrackKs** HLT1 line successfully collected real K_S^0 candidates, marking a significant milestone in this line of work.

More quantitative information on the analyzed candidates is obtained by fitting the distribution. The signal component has been modeled with the sum of two Gaussian distributions having a shared mean, while the combinatorial background distribution has been modeled with a uniform distribution. To enable a meaningful comparison, the extracted parameters are compared with the results of a fit to a simulated $D^0 \rightarrow K_S^0 K_S^0$ sample of K_S^0 decays processed with the same CPU version of HLT1. An unfiltered Run 2 data sample is added to the comparison, as a reference for the performance on real data collected with a commissioned detector. The K_S^0 candidates for these samples are fitted using the same model used for Run 3 data. The invariant mass distributions for these are reported in Figure 7.3, with the fits overlapped. The location of the peak is the first parameter drawing attention. The values for the mean of Gaussian distributions extracted from fits to different samples are reported in Table 7.1. The peak on Run 3 data appears to be shifted with respect to the PDG value by ~ 6 MeV/ c toward lower masses, corresponding to approximately a -1% bias in the candidate's mass. This behavior is consistent with what is observed on simulated data processed in the same way, where the extracted value of μ is compatible with the one from real data, while Run 2 data peak position is closer to the expected value. These observations indicate the presence of a small but non-negligible bias in the reconstructed mass of K_S^0 candidates. However, this is not a cause for concern, as it turned out to be due to a small bug in the HLT1 reconstruction code, which was promptly fixed.

The mass resolution of the collected K_S^0 decays is another important parameter, as it

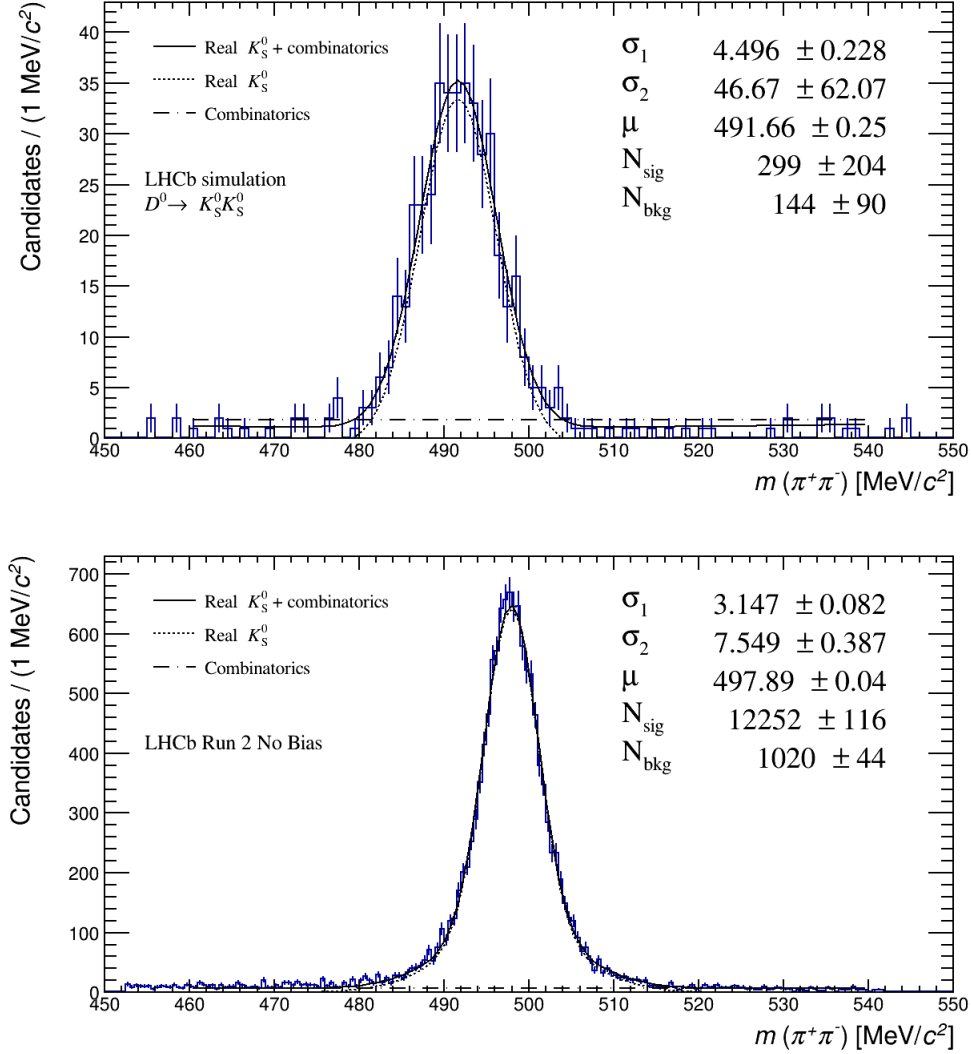


Figure 7.3: Invariant mass of K_S^0 candidates selected by **TwoTrackKs** HLT1 line on a $D^0 \rightarrow K_S^0 K_S^0$ Run 3 simulated sample (top) and an unfiltered Run 2 sample (bottom). The fit to the distributions are overlapped. The σ_1 , σ_2 , μ parameters respectively refer to the standard deviation of the two Gaussians and the shared mean of the two. N_{sig} and N_{bkg} are the number of signal and background candidates extracted in the fit.

determines the achievable signal-to-background (S/B) ratio for signal decays. Additionally, if the resolution deviates significantly from expectations, it may require loosening the mass window used during selection. To obtain a robust estimate of the peak width that is independent of the signal shape used in the fit, the Full Width at Half Maximum (FWHM) was chosen as the measure of peak width. The FWHM value for Run 3 data is compared to both Run 3 simulations and Run 2 data, with all results summarized in Table 7.2. A difference between Run 3 data and simulation is observed, with the FWHM for real data being approximately $\sim 3 \text{ MeV}/c^2$ larger, corresponding to a 30% increase over the expected width. This discrepancy can be attributed to the very preliminary detector alignment used during data-taking, which does not affect simulated data. However, this is not a cause for concern, as the impact on data-taking is minimal, and improvements

Sample	μ [MeV/ c^2]
$m(K_S^0)$ PDG	497.611
Run 3 data	491.80 ± 0.03
Run 3 simulation	491.66 ± 0.25
Run 2 data	497.98 ± 0.04

Table 7.1: The μ values extracted from fits to analyzed samples. The $m(K_S^0)$ value from PDG is reported as a reference.

Sample	FWHM [MeV/ c^2]
Run 3 data	13
Run 3 simulation	10
Run 2 data	7.5

Table 7.2: FWHM values for K_S^0 mass peaks in the reported samples.

are expected as operating conditions stabilize. A difference is also noted between Run 3 (both data and simulation) and Run 2 data. This can be explained by the accuracy of the offline-level reconstruction used in processing Run 2 data, which delivers better performance compared to the fast, HLT1-level reconstruction applied to the Run 3 data. Overall, the observed behavior indicates that the novel line is functioning as expected, and any discrepancies with predictions can be explained by the preliminary conditions of the detector and trigger system, and do not interfere with the efficient collection of genuine K_S^0 candidates. Therefore, no further modifications to the HLT1 line were deemed necessary.

After having verified that the novel K_S^0 line correctly selects genuine K_S^0 candidates, the quality of these is investigated. One interesting parameter that can be investigated to do this is the χ^2 of the fit to the vertex of K_S^0 candidate, reported in Figure 7.4. For the analyzed two-track vertex, the χ^2 is expected to be distributed as a χ^2 with one degree of freedom, and this function is superimposed with the plot in the figure. The distribution for K_S^0 candidates selected by the line shows a good agreement with the expected shape. This is a positive signal, suggesting that the exploited fit procedure is behaving as expected, with the uncertainties assigned to the vertex position being correctly assigned and considered in the vertex fit.

Another investigated variable is the DIRA of reconstructed K_S^0 candidates. This is interesting as the expected DIRA distribution for real K_S^0 decays is expected to be highly peaked towards 1, since flight direction and momentum of K_S^0 candidates are parallel, except for resolution effects. The distribution for candidates collected by the **TwoTrackKs** line is reported in Figure 7.5. The expected trend is visible on the reported distribution, confirming, together with the observed $m(K_S^0)$ distribution, that collected candidates are real genuine K_S^0 decays.

The tests performed on early Run 3 data strongly suggest that the novel **TwoTrackKs** is not showing any unexpected behavior, as it collected a sizeable sample of genuine K_S^0 candidates with a good purity in a few nb^{-1} of data. The observed discrepancies with Run 3 simulated data are not a matter of concern, as these can be attributed to the





Figure 7.4: Distribution of χ^2 for vertices of reconstructed K_S^0 candidates in logarithmic scale (top) and zoom at low values (bottom). Red line represents the shape of a χ^2 distribution for one degree of freedom.

preliminary conditions of data-taking, and not prevent the novel line from continuing to collect K_S^0 candidates.

7.2.2 Commissioning the TwoKs line

A similar analysis is repeated for the commissioning of the **TwoKs** line, studying the K_S^0 pair candidates collected by the line on Run 3 data. The S_1 sample is again exploited for these studies.

Being **TwoKs** line designed to collect K_S^0 pairs, the 2D distribution of K_S^0 pairs' invariant mass is considered, shown in Figure 7.6. First thing worth mentioning is the clear presence of a cluster of candidates close to the expected position for a pair of real K_S^0 candidates, suggesting that the **TwoKs** HLT1 line is collecting real K_S^0 pairs at the HLT1 level. This distribution is fitted to extract its parameters, the following components are adopted:


 Figure 7.5: Distribution of DIRA of reconstructed K_S^0 candidates in logarithmic scale.

Parameter	Value
N_{sig}	345 ± 27
μ	$490.6 \pm 0.2 \text{ MeV}/c^2$
σ_{sig}	$5.2 \pm 0.2 \text{ MeV}/c^2$
FWHM	$12 \text{ MeV}/c^2$

 Table 7.3: Results of the fit to the 2D distribution of K_S^0 candidate pairs collected by **TwoKs** HLT1 line on Run 3 data (sample S_1).

- signal contribution is modeled 2D Gaussian function modeling the contribution of events where both candidates are real K_S^0 decays;
- the background contributions where one real K_S^0 is paired with a random combination of tracks is modeled with two Gaussian \times uniform contributions;
- the background contribution given by random combination of four tracks is modeled with a uniform contribution in both dimensions.

The fit is superimposed on the bottom plot of Figure 7.6. Results of the fit are reported in Table 7.3. The extracted μ value indicates a mass shift similar to that observed in the single K_S^0 line. This shift is not a cause for concern, as the cluster of genuine K_S^0 candidates remains well-contained within the adopted mass window, and no corrective action will be taken. The FWHM of the 1D projection of the observed peak is $12 \text{ MeV}/c^2$, which is close to the value observed for candidates collected by the **TwoTrackKs** line, as expected. Additionally, the quantitative estimate of N_{sig} extracted from the fit confirms the significance of the observed peak, standing 10σ away from zero. Since the candidates collected by the **TwoKs** HLT1 line are pairs of those collected by the **TwoTrackKs** line, no additional checks are performed in this case, as no differences are expected compared to the results already presented. Before proceeding further, it is important to emphasize that the **TwoKs** HLT1 line successfully collected genuine K_S^0 pairs (not necessarily coming

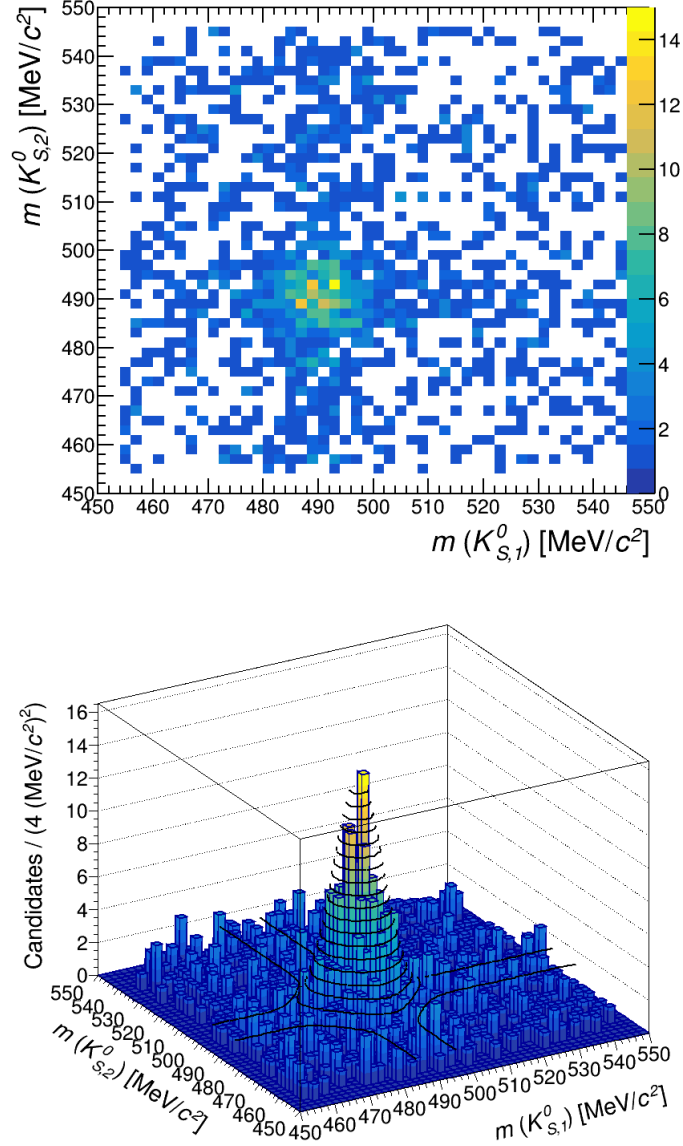


Figure 7.6: The 2D distribution of K_S^0 pairs invariant mass for candidates selected by TwoKs HLT1 line on Run 3 data (sample S_1). The bottom plot has the result of the 2D fit superimposed with the distribution.

from the decay of a common particle) at the first trigger level for the first time, without exhibiting any major unexpected behavior.

7.2.3 Rate measurements

Accurately reproducing generic pp collisions in simulation is very challenging, and the background acceptance of new selections can easily turn out to be significantly higher in real data than anticipated, potentially rendering the use of certain lines impractical. Consequently, verification of the rates of these lines on real data is a critical, anxiously awaited milestone. For this purpose, the unfiltered version of S_2 and S_3 samples was

analyzed, collected at different instantaneous luminosity levels. The rates computed at different μ values are summarized in Table 7.4.

μ	TwoTrackKs [kHz]	TwoKs [kHz]	TrackMVA [kHz]	TwoTrackMVA [kHz]
1.1	1.24 ± 0.08	0.05 ± 0.02	59.5 ± 0.6	90.6 ± 0.7
3.1	4.2 ± 0.5	2.5 ± 0.4	615 ± 6	1294 ± 9
5.32 (Simulation)	74 ± 7	13 ± 3	598 ± 18	817 ± 22

Table 7.4: Rate of HLT1 lines computed on 2022 commissioning data and simulation.

The results show that the rates are well-controlled, remaining comfortably within a small fraction of the available bandwidth, and below simulation predictions for both of the analyzed μ values. A rate increase is observed at higher μ , as expected, with the **TwoKs** line showing a sharper rate of increase with luminosity, consistent with expectations for a 4-track selection. The comparison with traditional trigger lines is also consistent with simulation expectations: the rates of **TrackMVA**, **TwoTrackMVA** are always at least two orders of magnitude larger with respect to the K_S^0 lines.

While these data do not reach up to the nominal LHCb Run 3 luminosity ($\mu = 5.5$) assumed in simulations and cannot therefore be directly compared, they are sufficient to confirm that the implemented selection rates remain within perfectly acceptable limits, and are qualitatively in agreement with expectations. This supports the continued use of the new lines in Run 3 physics data-taking.

However, the success of this new trigger strategy cannot be conclusively established until also the efficiency of the lines is measured on data, to check if the promising performances expected from simulation are confirmed by observation. This is discussed in the next section.

7.2.4 First attempt at evaluating efficiencies

The most appropriate parameter to evaluate this would be the comparison of yields per pb^{-1} for $D^0 \rightarrow K_S^0 K_S^0$ candidates with the old data-taking method. During Run 2, LHCb achieved a yield of approximately 1 $D^0 \rightarrow K_S^0 K_S^0$ candidate per pb^{-1} . However, the integrated luminosity of samples $S_{1,2,3}$ is too limited to expect the observation of even a single $D^0 \rightarrow K_S^0 K_S^0$ candidate. In principle, one might get a first idea of the efficiency of the new lines indirectly, by looking simply at the yields of K_S^0 or K_S^0 pairs. This is a physical parameter that is expected to remain constant, when normalized to the number of physical pp collisions, that can be estimated from the number of reconstructed PV in the sample. The estimate was attempted by performing a comparison with an ad-hoc real-data sample from Run 2, obtained by applying the **TwoTrackKs** and **TwoKs** selections to a sample of two-track pairs reconstructed in a special, unfiltered Run 2 subsample. Without going into the details of this study, now completely superseded by more recent and much better data, we will simply try to summarize why it failed to reach its goal of providing a credible estimate of trigger efficiency. The raw number of selected K_S^0 candidates (pairs) in the Run 3 data was found to be about a factor 10 (100) lower than what was achieved during Run 2. The relation between the two lines is as expected, as

it can be attributed to the requirement of a second candidate, which effectively reduces the overall efficiency by squaring it. However, it turned out to be impossible to translate those raw numbers into trigger efficiencies, as the (obviously huge) detector inefficiency during this early commissioning period could not be measured with any accuracy.

Too many detector parameters in these very early data were still to be calibrated. This included important effects like the time-alignment of many detectors reference to the LHC bunch timing and the voltage supplies (affecting the hit efficiencies), the detector alignment (affecting tracking efficiency and PV-finding efficiency); and the luminosity measurement itself was still quite uncertain. The product of all these inefficiency factors meant that the low observed overall efficiency was indeed far from being surprising, but its value could not be accurately estimated, and was indeed rapidly changing while data was being collected as all those calibrations were being performed by several teams of experts working in parallel on different things. So eventually we decided to give up and postpone a direct measurement of trigger efficiency to a later time, when larger samples of data with stabler and better understood detector conditions would have become available.

Nevertheless, the outcome of these first tests performed on the novel HLT1 lines on real data was absolutely positive. No evidence was found of any obvious bug in the developed selections, that were collecting genuine K_S^0 decays and K_S^0 pairs, with a very good purity. The lines' rate was well under control and its ratio to the rate of other lines perfectly in line with expectations. We could afford to continue using them, while looking forward to collecting a large enough chunk of good data to complete their assessment.

7.3 K_S^0 triggers performance assessment on 2022 data

After the commissioning period, LHCb obtained a first data sample using the complete detector and data-processing chain at the end of 2022. These data were promptly made available to the whole collaboration for detailed analysis, as they could provide critical insights into the functionality of many crucial LHCb subsystems, including the trigger. This included our K_S^0 lines, that had kept running smoothly throughout the whole period.

The sample we analyzed corresponds to an integrated luminosity of 8.6 pb^{-1} . Only events selected by the $D^0 \rightarrow K_S^0 K_S^0$ LL HLT2 line were considered, as the UT subdetector was not yet included in data-taking at this stage, preventing downstream tracking to be performed, both at the HLT1 and HLT2 level. We however accepted decay candidates coming from any of the HLT1 lines in operation. This was motivated, not only to facilitate comparisons between lines, but also for including the largest possible sample in seeking a first charm signal. The bandwidth division procedure was not active, and the nominal configuration for K_S^0 HLT1 lines was available.

The first crucial check has been done to verify the correct functioning of HLT2 lines code, that although inspired by the previous Run 2 line, had been completely rewritten for Run 3. The first check is performed by looking at the $m(K_S^0)$, $m(D^0)$ and Δm distributions for candidates, before the application of any offline selection. Results are reported in Figure 7.7. Examining the 2D invariant mass distribution of K_S^0 candidates, a clear cluster is again observed around the expected mass for a real K_S^0 pair. This demonstrates that the HLT2 line is successfully collecting genuine K_S^0 pairs, albeit with inclusion of a significant background. However, the invariant mass distributions for D^0 and Δm do not show any hint of a signal peak, appearing perfectly consistent with a pure background distribution.

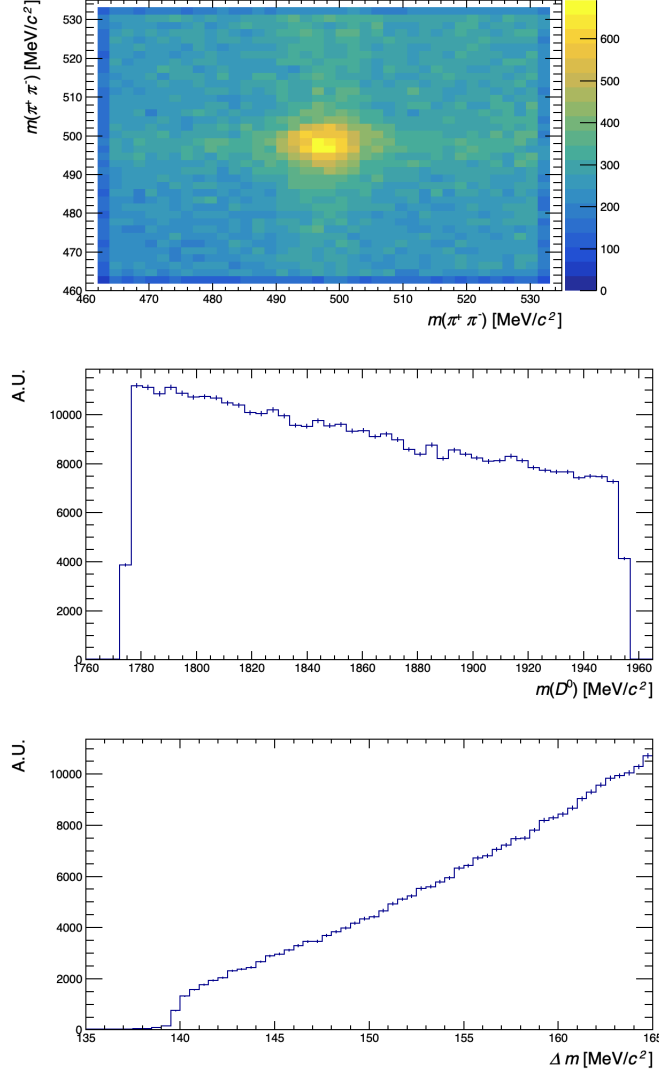


Figure 7.7: Invariant mass distribution for K_S^0 pairs (top), D^0 (center) and Δm (bottom) for candidates selected by the $D^0 \rightarrow K_S^0 K_S^0$ LL HLT2 line on 2022 physics sample. The full triggered sample is plotted, without any additional offline selections.

All that can be verified here is that the distributions fall within the expected ranges, confirming the correct implementation of the mass selection thresholds.

To check for the possible presence of any charm signals in this sample, a set of offline selections was applied, as detailed by Table 7.5. These selections were not formally optimized, but were formulated on the basis of physics intuition and empirical observations, as an initial attempt to reduce background in the sample. Notably, no requirements are applied at this stage on the displacement of the K_S^0 candidates. This choice is deliberate, as it allows $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays to contribute to the Δm peak, if present. Given that these decays are much more abundant than $D^0 \rightarrow K_S^0 K_S^0$, and are expected to be effectively selected by all trigger requirements, they serve as a valuable tool for checking the functionality of the HLT2 line, despite being effectively a background for our target signal at the analysis level. The D^0 and Δm invariant mass distributions after the application of the listed offline selections are reported in Figure 7.8. Although the $m(D^0)$ distribution

Variable	Selection
$\sqrt{(m(\pi^+\pi^-)_1 - m(K_S^0))^2 + (m(\pi^+\pi^-)_2 - m(K_S^0))^2}$	$< 15 \text{ MeV}/c^2$
$\cos(\theta_{\text{DIRA}}(D^0))$	> 0.99992
$p_T(K_S^0)$	$> 850 \text{ MeV}/c$
$\chi_{\text{vtx}}^2(K_S^0)$	< 4
$\chi_{\text{vtx}}^2(D^0)$	< 4

Table 7.5: Offline selections applied to the 2022 $D^0 \rightarrow K_S^0 K_S^0$ sample. These reject background contribution while keeping $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays.

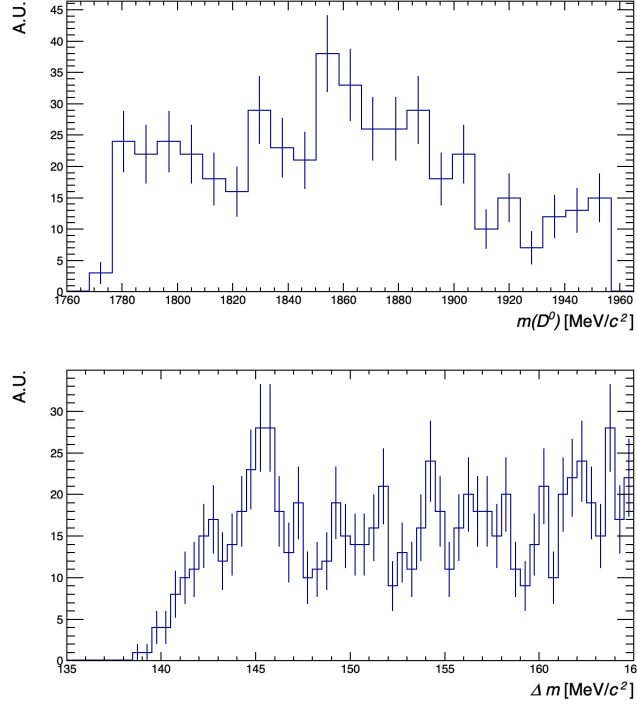


Figure 7.8: Δm (top) and D^0 (bottom) invariant mass distributions for candidates selected by the $D^0 \rightarrow K_S^0 K_S^0$ LL HLT2 line on 2022 physics sample after the application of selections reported in Table 7.5. On top of those a $143.5 < \Delta m < 147 \text{ MeV}/c^2$ ($|m(K_S^0 K_S^0) - m(D^0)| < 30 \text{ MeV}/c^2$) selection is applied for the D^0 (Δm) invariant mass.

shows only a weak hint of a broad peak, the Δm distribution reveals a clear indication of a peak emerging from the background at the expected location. The appearance of this first signal is a significant outcome, validating the proper functioning of the HLT2 line and more widely of the whole trigger setup.

After this check, it is of course interesting to step further forward, rejecting $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ candidates by applying a requirement on the K_S^0 displacement, to search for the presence of an actual $D^0 \rightarrow K_S^0 K_S^0$ signal. This is done by applying the selections reported in Table 7.6. 7.5. Because of the lower precision on the determination of the D^0 vertex position in the case of the $K_S^0 K_S^0$ final state, some selections have been slightly tuned with the aim of improving the efficiency for $D^0 \rightarrow K_S^0 K_S^0$ decays. The resulting Δm and D^0 invariant mass distributions are reported in Figure 7.9. This time, no clear signal

Variable	Selection
$\sqrt{(m(\pi^+\pi^-)_1 - m(K_S^0))^2 + (m(\pi^+\pi^-)_2 - m(K_S^0))^2}$	$< 15 \text{ MeV}/c^2$
$\cos(\theta_{\text{DIRA}}(D^0))$	> 0.9999
$p_T(K_S^0)$	$> 850 \text{ MeV}/c$
$\chi_{\text{vtx}}^2(K_S^0)$	< 3
$\chi_{\text{vtx}}^2(D^0)$	< 5
$\log \chi_{FD}^2(K_S^0)$	> 9

Table 7.6: Offline selections applied on the 2022 $D^0 \rightarrow K_S^0 K_S^0$ sample rejecting background contribution and $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays.




Figure 7.9: Δm (top) and D^0 (bottom) invariant mass distributions for candidates selected by the $D^0 \rightarrow K_S^0 K_S^0$ LL HLT2 line on 2022 physics sample after the application of selections reported in Table 7.6. On top of those a $143.5 < \Delta m < 147 \text{ MeV}/c^2$ ($|m(K_S^0 K_S^0) - m(D^0)| < 30 \text{ MeV}/c^2$) selection is applied for the D^0 (Δm) invariant mass.

is observed in either of the two distributions, although one might argue about a faint indication of a peak near the expected mass in the Δm distribution. A quick calculation shows that this is not a reason for concern at this stage. As previously mentioned, approximately $1 D^0 \rightarrow K_S^0 K_S^0 / \text{pb}^{-1}$ was collected in Run, and, even accounting for potential improvements, we would expect at most a few dozen genuine candidates in this sample before any further cuts. However, this yield is likely further reduced by the applied offline selections, which are not optimized and have an efficiency significantly below 100%, as well as by the preliminary detector conditions, which are likely impacting reconstruction and selection efficiency. Moreover, considering the $\mathcal{O}(100)$ observed $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ candidates and using simple branching fraction and acceptance estimates, the expected $K_S^0 K_S^0$ yield

should be smaller by a factor of several hundred. This confirms that the expected number of $K_S^0 K_S^0$ candidates is only a few, making them too scarce to be distinguishable from the background seen in Figure 7.9.

Nonetheless, the performed checks are critically important, as the observation of the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ signal, very similar to our target mode, confirms the proper functionality of the HLT2 line machinery.

7.4 Analysis of 2023 HLT2 $D^0 \rightarrow K_S^0 K_S^0$ data

Following the 2022 commissioning activities, LHCb resumed operations in 2023. In this year, the detector performances have been limited by an accident that caused deformation of the VELO RF-foil. This prevented LHCb from using the VELO in its closed configuration, which significantly limited the overall detector performance. The increased distance between the VELO and the interaction region primarily resulted in reduced precision in track position and pointing, as well as a roughly 50% reduction in detector acceptance. While the reduced acceptance affected all physics channels uniformly, the lack of precise vertex position and impact parameter determination had a smaller influence on K_S^0 final-state decays. This is because the resolution for these decays is already limited by the large flight distance of the K_S^0 particles. This motivated an early analysis of 2023 data that can anyway be exploited to measure the performance of the K_S^0 -dedicated HLT1 and HLT2 lines, despite the harsh data-taking conditions.

The analyzed data were collected during the summer, following an initial period dedicated to recommissioning the detector and ensuring the safe operation of the VELO during data-taking. These correspond to an integrated luminosity of 48 pb^{-1} . The considered candidates are collected by $D^0 \rightarrow K_S^0 K_S^0$ LL HLT2 line already exploited in the analysis of 2022 data. Candidates can be selected at the HLT1 level in two configurations: the D^0 candidates are TOS on `TwoTrackKs`, `TwoKs`, `TrackMVA` or `TwoTrackMVA`, or the D^{*+} is TIS on any HLT1 line. This selection is adopted as it corresponds to the one exploited during Run 2 analysis and it is expected to maximize S/B, at least at the first level. Also in this case, the bandwidth division procedure was not exploited, and the nominal configuration for K_S^0 HLT1 lines was used.

The inspection of $m(K_S^0)$, $m(D^0)$, and Δm distributions before the application of any offline selection is the first performed check. These are reported in Figure 7.10. A distinct peak near the expected position is observed in the 2D $m(K_S^0)$ invariant mass distribution, exhibiting a lower background level compared to the 2022 sample. Notably, peaks are also visible in the $m(D^0)$ and Δm distributions, this time even before applying any offline selections. Both $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ and $D^0 \rightarrow K_S^0 K_S^0$ decay contributions are included at this stage, but the 2023 dataset clearly demonstrates higher purity relative to the 2022 data. The presence of a signal peak confirms the correct functioning of the exploited HLT2 line in 2023.

Having checked this, a set of offline selections is applied to search for a $D^0 \rightarrow K_S^0 K_S^0$ signal evidence, rejecting combinatorial background and $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays. These are derived from the ones applied in the case of the 2022 sample, and are reported in Table 7.7. The Δm invariant mass distribution after the application of the selections is reported in Figure 7.11, before and after the exploitation of DTF variables computed constraining the D^{*+} and PV vertices to coincide. The reported distributions show evidence for the

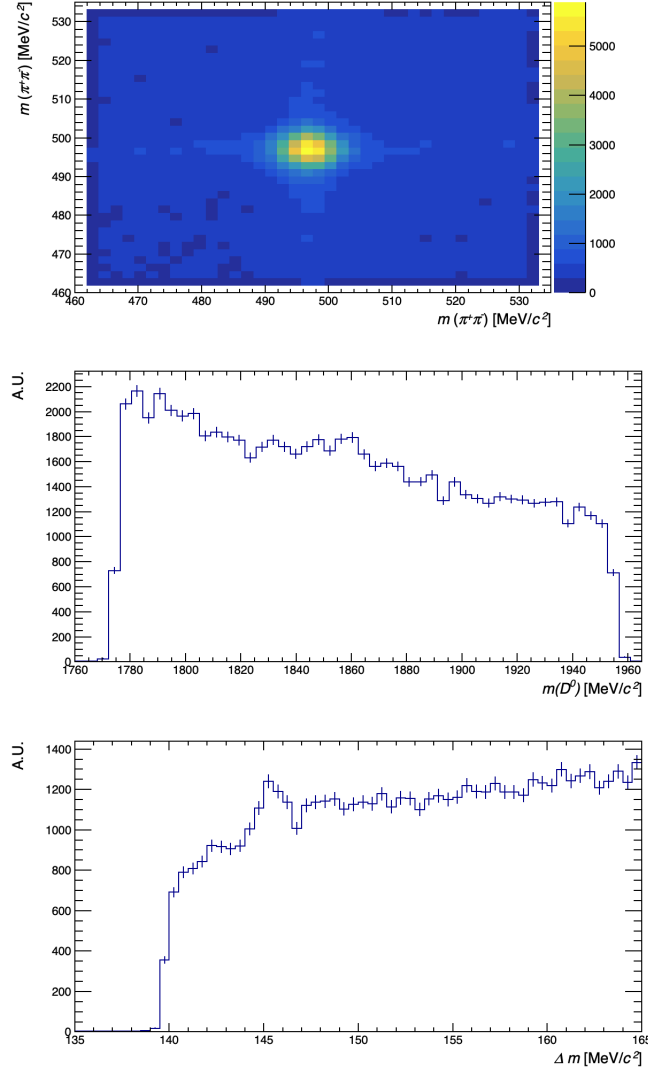


Figure 7.10: Invariant mass distribution for $K_S^0 K_S^0$ pairs (top), D^0 (center) and Δm (bottom) for candidates selected by the $D^0 \rightarrow K_S^0 K_S^0$ LL HLT2 line on 2023 physics sample. Only trigger selections are applied to reported candidates.

presence of a $D^0 \rightarrow K_S^0 K_S^0$ signal in the 2023 sample. The number of $D^0 \rightarrow K_S^0 K_S^0$ decays present in the sample is estimated by fitting the Δm distribution. A Gaussian and an empirical threshold function are used to model the signal and background components, respectively. The signal peak stands 3σ from zero using invariant masses computed without the exploitation of DTF variables. When this tool is exploited, constraining the D^{*+} and PV vertices to coincide, the peak reaches a 5σ significance. This is the first evidence for the presence of $D^0 \rightarrow K_S^0 K_S^0$ decays in Run 3 in a LHCb sample. The achieved yield/ pb^{-1} is approximately 1 $D^0 \rightarrow K_S^0 K_S^0/\text{pb}^{-1}$, similarly to the value for Run 2, with similar purity. This is highly promising, as it matches the performance achieved during Run 2, despite being obtained under the VELO open conditions of the 2023 sample.

With the presence of signal decays confirmed in the sample, it becomes possible, and interesting, to investigate which HLT1 line is responsible for their selection. This is analyzed by categorizing the sample based on the HLT1 line that selected the D^0 in

Variable	Selection
$\sqrt{[m(\pi\pi)_1 - m(K_S^0)]^2 + [m(\pi\pi)_2 - m(K_S^0)]^2}$	$< 12 \text{ MeV}/c^2$
$ m(K_S^0 K_S^0) - m(D^0) $	$< 30 \text{ MeV}/c^2$
$\sqrt{[\log \text{FD}_{\text{ORIVX}}(K_{S1}^0) - 6.2]^2 + [\log \text{FD}_{\text{ORIVX}}(K_{S2}^0) - 6.2]^2}$	< 3
$\chi_{\text{vtx}}^2(K_S^0)$	< 10
$\chi_{\text{vtx}}^2(D^0)$	< 10
$\chi_{\text{vtx}}^2(D^*)$	< 10
$\cos(\text{DIRA}(D^0))$	> 0.99995
$\chi_{\text{FD}}^2(D^0)$	> 7
$p_T(K_S^0)$	$> 800 \text{ MeV}/c$
$p_T(\pi_{\text{tag}})$	$> 225 \text{ MeV}/c$

Table 7.7: Selection criteria for the analysis of 2023 $D^0 \rightarrow K_S^0 K_S^0$ sample.

a TOS configuration. The Δm distributions for the resulting sub-samples are shown in Figure 7.12, both prior to and following the application of offline selections. These distributions are not fitted to extract signal yields, as in the majority of cases the statistics are quite low. However, some interesting conclusions can be anyway extracted. Both before and after the application of offline selections, the **TwoTrackKs** line is responsible for the selection of the majority of candidates. Notably, this is the case also for events present in the $D^0 \rightarrow K_S^0 K_S^0$ signal peak. In contrast, only a faint indication of a signal peak is observed for candidates selected by the **TwoKs** line. This is not concerning at this stage, as the **TwoKs** line is expected to have lower efficiency compared to the **TwoTrackKs** line, and this is still a small sample. The observed behavior of events selected by the **TrackMVA** and **TrackMVA** lines aligns with expectations. The single-track trigger recruits what is a fraction of the events triggered by the **TwoTrackKs** line, agreeing with simulation predictions. The two-track trigger is the line with the second largest contribution before the application of selections, being the line with the largest rate. However, because of the m_{corr} cut, it does not show any signal evidence after the selection of $D^0 \rightarrow K_S^0 K_S^0$ decays.

The tests performed on 2023 data finally allowed the observation of a $D^0 \rightarrow K_S^0 K_S^0$ signal on Run 3 data. The achieved yield/ pb^{-1} looks promising, as it matches the one observed on Run 2 data, despite the VELO open conditions. Even if the extraction of quantitative conclusions is complicated by the small available sample and the peculiar data-taking conditions, the analysis of HLT1 lines involved in the selection of these candidates finally confirms that the novel HLT1 K_S^0 lines are going to be the main source for the collection of these events in Run 3.

These observed promising performances, together with the lower sensitivity of the $K_S^0 K_S^0$ final state to the VELO open conditions, made the measurement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$ an excellent candidate for the physics exploitation of the data collected in 2023. The LHCb collaboration agreed to that, deciding to assign a larger weight to the $D^0 \rightarrow K_S^0 K_S^0$ in the bandwidth division procedure, so that the HLT1 selection would have been optimized to achieve a higher efficiency for these decays. The data-taking period target of this strategy was August 2023, where a statistics of approximately 1 fb^{-1} was planned to be collected. Unfortunately, an unexpected accident occurred to one of the LHC magnets, preventing LHCb from collecting these data, and no further collision data were produced until very close to the PbPb run planned at the end of the 2023 data taking. The plan for obtaining




Figure 7.11: Δm invariant mass distribution before (top) and after (bottom) the application of DTF with PV constraint. Reported candidates have been selected by the $D^0 \rightarrow K_S^0 K_S^0$ LL HLT2 line on 2023 physics sample and are required to pass the selections reported in Table 7.7.

a physics measurement exploiting the novel trigger was then postponed to 2024.

7.5 Analysis of 2024 HLT2 $D^0 \rightarrow K_S^0 K_S^0$ data

After the replacement of the RF-foil during the end of year stop of 2023, LHCb reached a stable data-taking condition, and started to consistently collect integrated luminosity in 2024.

The first dataset with physics-level performance became available in early May 2024. I promptly analyzed this data to assess its quality and evaluate the impact of the newly implemented trigger lines. Here the bandwidth division was exploited for the first time in Run 3. The nominal configuration for these lines has again been used to collect these data, as the bandwidth division optimization chose this configuration as the optimal one. The

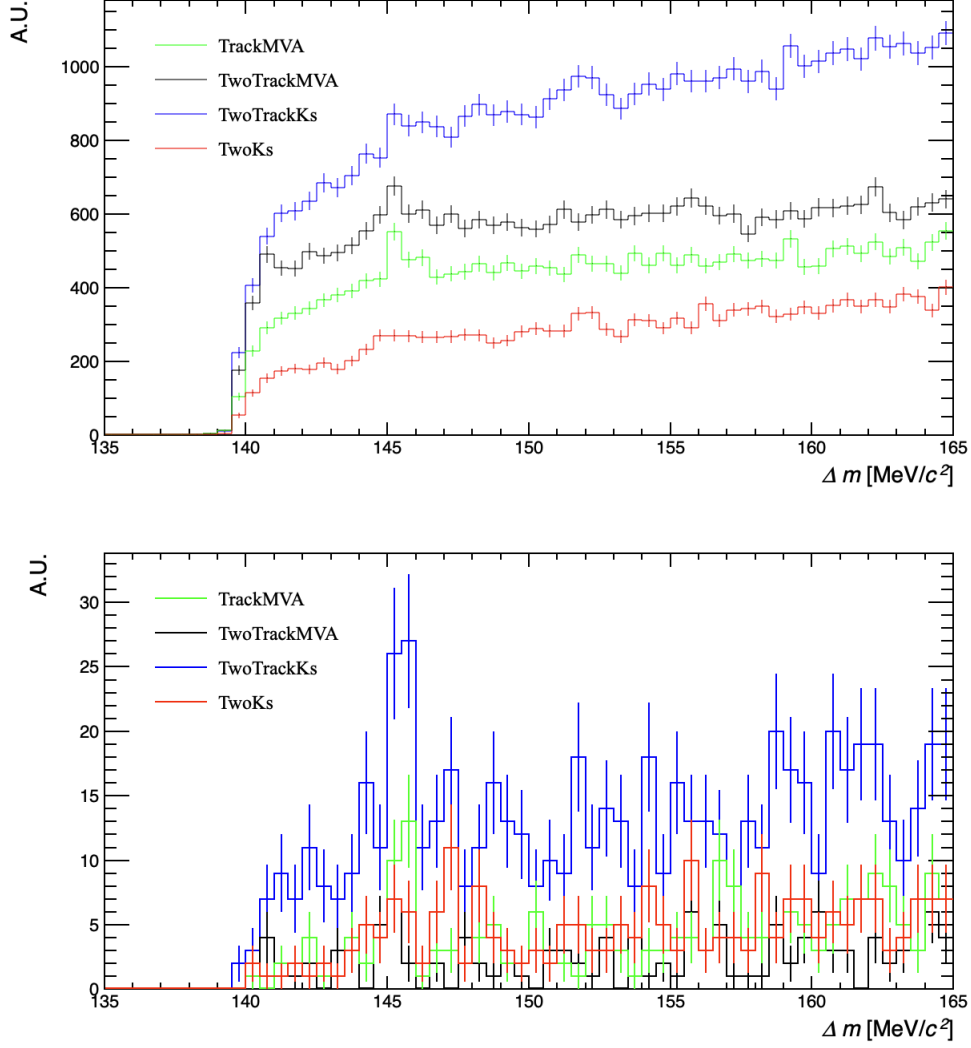


Figure 7.12: Δm invariant mass distribution for different HLT1 triggers before (top) and after (bottom) the application of offline selections reported in Table 7.7. The DTF constraint is applied after the application of selections. The D^0 candidate is required to be TOS on the considered HLT1 line in each sub-sample.

initial sample analyzed corresponds to an integrated luminosity of approximately 28 pb^{-1} . The candidates considered are those selected by the $D^0 \rightarrow K_S^0 K_S^0$ LL HLT2 line. The D^0 candidates are required to be TOS on the **TwoTrackKs**, **TwoKs**, or **TrackMVA** lines at the HLT1 level. Based on observations from the 2023 data, the **TwoTrackMVA** line has been excluded from this analysis. Additionally, candidates with D^{*+} TIS on any HLT1 line are removed to mitigate background contributions, as candidates selected on a voluntary basis are known to exhibit higher background contamination.

The same machinery tests performed for 2022 and 2023 data were repeated on 2024 data. The outcome of these was positive, confirming the correct functioning of HLT2 lines, and they are not reported here in detail, as they don't add new information to what has already been said.

The presence of $D^0 \rightarrow K_S^0 K_S^0$ decays is investigated by applying a set of offline selections,

reported in Table 7.8. These correspond to the preliminary selections applied during

Variable	Selection
$\rho_{xy}(D^0)$	$< 5 \text{ mm}$
$\sqrt{(m(\pi^+\pi^-)_1 - m(K_S^0))^2 + (m(\pi^+\pi^-)_2 - m(K_S^0))^2}$	< 12
$\log \chi_{\text{FD, ORIVX}}^2(K_S^0)$	> 4
$\chi_{\text{vtx}}^2(K_S^0)$	< 5
$\chi_{\text{vtx}}^2(D^0)$	< 5
$\chi_{\text{vtx}}^2(D^*)$	< 10

Table 7.8: Offline selections applied to the $D^0 \rightarrow K_S^0 K_S^0$ 2024 dataset.

Run 2 analysis for the measurement of $\mathcal{A}^{CP}(K_S^0 K_S^0)$. These have been exploited as they meet our purposes. Their application rejects the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays present in the sample with the $\chi_{\text{FD, ORIVX}}^2(K_S^0)$ selection, avoiding any bias in estimating the number of signal candidates. Furthermore, since these criteria were applied during the Run 2 analysis, they enable an easy, one-to-one quantitative determination of the improvement resulting from the implementation of the novel trigger system. Notice how the *preliminary*, and not the *full* Run 2 analysis selections (yielding the $1 D^0 \rightarrow K_S^0 K_S^0 / \text{pb}^{-1}$ quoted so far) are exploited here. This is done as the latter rely on the application of a selection based on the training of a kNN classifier, as reported in Chapter 4, whose effect is not easily reproducible in this sample. In fact, since it has been trained to separate a simulated (under Run 2 conditions) signal sample from Run 2 background data, the score produced for Run 3 signal candidates is likely to have a different efficiency with respect to the one achieved on Run 2 data. Therefore, it has been preferred to exploit just the cuts reported in Table 7.8, that achieved a yield/ pb^{-1} of $1.35 / \text{pb}^{-1}$.

The Δm distribution for candidates passing the listed offline selections, computed exploiting the DTF PV constraint, is reported in Figure 7.13. A signal peak is clearly evident in this sample. The same model is exploited to fit the distribution. The fit results show that the signal peak stands 5σ away from zero, containing $N_{\text{sig}} = 154 \pm 30$ candidates. It is now possible, for the first time in Run 3, to quantitatively estimate the gain due to the implementation of novel HLT1 lines on real data processed by the entire LHCb data-taking chain. The yield/ pb^{-1} achieved analyzing the 2024 sample is of 5.5, to be compared with the $1.35 / \text{pb}^{-1}$ achieved in Run 2. Let's take a moment to highlight the importance of this result. This test on early 2024 data shows that the inclusion of the K_S^0 -dedicated lines in the LHCb HLT1 determined an increase in the achieved yield/ pb^{-1} for $D^0 \rightarrow K_S^0 K_S^0$ decays by a factor of 4. This result clearly demonstrates the validity of the novel trigger strategy explored in this Thesis, now fully commissioned and validated on Run 3 data.

The test that has been performed is repeated with a larger sample, allowing a more precise determination of each HLT1 line contribution. Also the **TwoTrackMVA** line will be considered in this case. A sample of 1.8 fb^{-1} is analyzed. It is worth mentioning that this sample has approximately the same size as the one collected during the entire 2017 LHCb data-taking, but has been accumulated in less than two months, thanks to the higher instantaneous luminosity at which LHCb is able to take data in Run 3. The same selections reported in Table 7.8 are applied in this case. The obtained Δm distributions for different HLT1 configurations, together with fit results are reported in Figure 7.14. The




Figure 7.13: Δm invariant mass distribution for candidates present in the early 2024 sample after the application of selections reported in Table 7.8. The Δm is computed exploiting the DTF PV constraint.

HLT1 selection	N_{sig}	Yield/ pb $^{-1}$
TwoTrackKs	8218 ± 194	4.6
TwoKs	3393 ± 105	1.9
TwoTrackKs \vee TwoKs	9725 ± 216	5.4
TrackMVA \vee TwoTrackMVA	2562 ± 103	1.4
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs	9037 ± 213	5.0
TrackMVA \vee TwoTrackMVA \vee TwoTrackKs \vee TwoKs	10493 ± 226	5.8
TrackMVA (Run 2)	5283 ± 102	1.35

Table 7.9: N_{sig} extracted from fit to 2024 analyzed sample with different HLT1 configurations and relative yield/ pb $^{-1}$. The value for the 2017-18 LL sample is reported as a reference for Run 2.

considered HLT1 selections configuration, together with the extracted N_{sig} and the relative yield/ pb $^{-1}$ are reported in Table 7.9. 7.14. First thing worth noting is the number of signal candidates selected solely by the novel K_S^0 HLT1 lines, hence $N_{sig} = 9725 \pm 216$. This is larger by a factor of 1.62 with respect to the full LL Run 2 sample exploited in the latest $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement. This achievement becomes even more remarkable when considering the yield/ pb $^{-1}$, that for the 2024 sample corresponds to 5.4/ pb $^{-1}$, to be compared with the 1.35/ pb $^{-1}$ of Run 2, corresponding to an enhancement by a factor of 4 with respect to the Run 2 configuration.

It is interesting also to notice the yield/ pb $^{-1}$ obtained exploiting the **TrackMVA** and

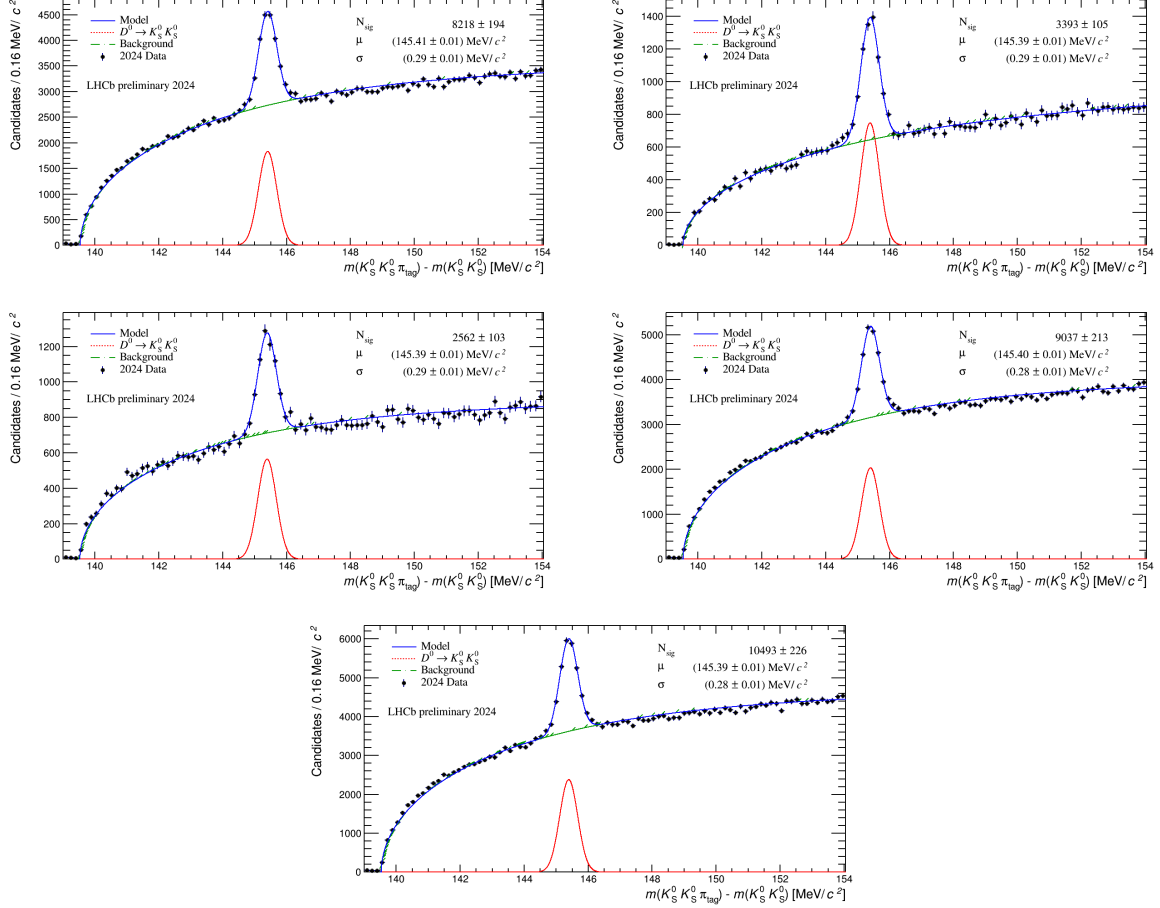


Figure 7.14: Δm invariant mass distributions for $D^0 \rightarrow K_S^0 K_S^0$ candidates present in the a sample of 1.8 fb^{-1} collected in 2024 with different HLT1 configurations. Candidates are required to be selected at the HLT1 level by: **TwoTrackKs** (top left), **TwoKs** (top right), **TrackMVA** \vee **TwoTrackMVA** (center left), **TrackMVA** \vee **TwoTrackMVA** \vee **TwoTrackKs** (center right), **TrackMVA** \vee **TwoTrackMVA** \vee **TwoTrackKs** \vee **TwoKs** (bottom). The fit is superimposed to each distribution.

TwoTrackMVA lines alone. This is essentially the same as the one achieved in Run 2. This is not unexpected in a channel like $D^0 \rightarrow K_S^0 K_S^0$, even with the exploitation of a software based trigger like HLT. In fact, the difficulties in triggering such a final state represent the fundamental motivations of this work. This should not be viewed as a limitation of the upgraded LHCb system. Rather, it highlights the need to adopt more sophisticated trigger strategies to fully leverage the early reconstruction capabilities provided by the upgraded HLT1. By implementing more complex selections that take full advantage of the decay topology, it becomes possible to significantly increase background rejection, while maintaining a high efficiency. This approach has been shown to be particularly effective for neutral, long-lived particle final states like $K_S^0 K_S^0$, but its principles are equally applicable to other decay channels.

The enhancement achieved by incorporating the **TwoTrackKs** line into the system alongside **TrackMVA** and **TwoTrackMVA** corresponds to a factor of 3.6, compared to the 2.6 predicted by simulation. This greater gain can be attributed to the fact that only LL decays are considered in this analysis, whereas the simulation study included all

$D^0 \rightarrow K_S^0 K_S^0$ decays occurring within the LHCb acceptance. Additionally, the **TwoTrackKs** line alone selects $\sim 80\%$ of the total $D^0 \rightarrow K_S^0 K_S^0$ decays. It is also extremely interesting to observe now that also the **TwoKs** line perfectly meets its purposes, being able to achieve alone $\sim 30\%$ of total efficiency, while occupying just a small fraction of the HLT1 output rate.

Concluding considerations

In summary, the tests performed on 2024 data showed on real data that the exploitation of the novel HLT1 lines enhances by a factor of approximately 4 the yields/ pb^{-1} for the $D^0 \rightarrow K_S^0 K_S^0$ decays, similarly to what had been predicted on simulation and validating the work that I performed and here reported. Furthermore, this achievement marks a significant milestone for the entire LHCb collaboration, as it provides real-world evidence of the upgrade's positive impact on physics data collection. Thanks to the increased instantaneous luminosity and the implementation of the new K_S^0 trigger lines, the experiment successfully collected thousands of $D^0 \rightarrow K_S^0 K_S^0$ decays within just a few weeks.

The impressive result of the novel HLT1 K_S^0 -dedicated lines motivated me, and the entire LHCb collaboration, to perform an early \mathcal{A}^{CP} ($K_S^0 K_S^0$) measurement with Run 3 data, that will be described in the next Chapter. This measurement fits extremely well the profile for an early Run 3 measurement, as it is one of the channels where LHCb is expected to achieve a result with a resolution close to the world average even with just a few fb^{-1} of data, thanks to the implementation of the novel HLT1 lines.

Chapter 8

Measurement of \mathcal{A}^{CP} ($D^0 \rightarrow K_S^0 K_S^0$) in Run 3

The LHCb 2024 data-taking continued smoothly throughout 2024, allowing the collaboration to collect a total integrated luminosity of 9.56 fb^{-1} , by the end of the year. This makes the 2024 sample larger than the ones collected during all past runs taken together, as can be seen in Figure 8.1.




Figure 8.1: Diagram showing the integrated luminosity collected by LHCb in different data-taking years.

During the summer of 2024, I started analyzing the part of this sample that was already available, to perform the measurement of \mathcal{A}^{CP} ($K_S^0 K_S^0$). This work has been performed within a small group, formed by me, with my advisor Prof. Giovanni Punzi, Dr. Giulia Tuci, and Erich Vito Butera, a master's student working on this for his Master's thesis work.

Before presenting this work, it is important to emphasize that the primary goal of this measurement is not to achieve the best possible performance with the available data, but rather to quickly produce a good result to take a step forward in the understanding of this important physics observable. At the same time, we aimed at validating the actual capabilities of the upgraded LHCb detector and trigger system to produce reliable physics, while the Run 3 is still underway. For this reason, we eschewed several optimizations we might otherwise have performed, in favor of a simpler and faster analysis. Obviously, it will make sense to perform a full optimization of all aspects of this measurement after the end of the Run 3, when the full sample will have been collected.

8.1 Analysis strategy

To expedite the process, the methodology adopted in this analysis follows whenever possible the procedures employed in the Run 2 measurement described in chapter 4.

The most important differences are:

1. exclusive exploitation of the novel HLT1 and HLT2 triggers;
2. the use of $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays as a calibration sample (instead of $D^0 \rightarrow K^+ K^-$);
3. use only K_S^0 reconstructed from long tracks.

All of those three changes are motivated by making the analysis simpler and faster. Point (1) simplifies the analysis by avoiding the complexities associated with multiple subsamples, taking advantage of the largest fraction of data being selected by a single line. Point (2), while superficially appearing to be a complication, allows to define the calibration sample exactly from the same trigger line of the signal. This is particularly important when considering that the $D^0 \rightarrow K^+ K^-$ sample comes from a completely different trigger in Run 3, and one that was repeatedly re-tuned and improved during 2024. Point (3) again avoids multiplying the number of samples included in the analysis, and sacrifices a modest amount of statistics, given that the K_S^0 trigger of point 1 does not make use of downstream tracks, but only of long tracks.

A sample of flavor-tagged $D^0 \rightarrow K_S^0 K_S^0$ decays is obtained by selecting D^{*+} candidates, with subsequent decay $D^{*+} \rightarrow D^0 \pi^+$. The sign of the pion in this decay gives the flavor of the accompanying D^0 . $D^0 - \bar{D}^0$ mixing is negligible to our level of precision. K_S^0 candidates are reconstructed in the $\pi^+ \pi^-$ decay channel. As previously mentioned, only K_S^0 candidates where the daughters pions are reconstructed as long tracks are used.

The applied offline selections consist of two main steps: the first rejects unwanted physics background, while the second is aimed at reducing combinatorial background to maximize the sensitivity to \mathcal{A}^{CP} . The $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays where the di-pion invariant mass falls within the K_S^0 mass window are an important source of physics background also in this case, and their contribution is reduced using K_S^0 decay-flight observables. The second step of the selection is again based on a cut on the output of a kNN classifier, basing its decision on kinematic, topological and PID quantities. As in the Run 2 analysis, secondary decays are included as part of the signal.

A 3D maximum likelihood fit to Δm and the two $m(K_S^0)$ distributions is used to extract the CP asymmetry, weighting each event to correct for production and detection charge asymmetries. Weights are based on the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ calibration sample. Similarly

to what was done in Run 2 for the $D^0 \rightarrow K^+ K^-$, this sample has been selected with dedicated lifetime-unbiased selections, both at the HLT1 and HLT2 level, avoiding any bias on the fraction of secondary decays. Since in our offline selection we mainly select the CF decay $D^0 \rightarrow K^{*-} \pi^+$, $K^{*-} \rightarrow K_S^0 \pi^-$, we assume $\mathcal{A}^{CP}(K_S^0 \pi^+ \pi^-) = 0$ in this analysis. The central value of our measurement is kept blind to avoid any experimenter bias.

8.2 Data sample

A subset of the pp -collision data collected in 2024 has been included in the analysis, corresponding to an integrated luminosity of about $\int \mathcal{L} dt \simeq 6 \text{ fb}^{-1}$. These have been collected at an average instantaneous luminosity of approximately $1.6 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ ($\mu \simeq 4.4$). These data have no overlap with the sample of section 7.5, that cannot be used due to the HLT2 line for the calibration sample not being yet in operation.

An important feature of the 2024 sample is that it was collected during a period of rapidly evolving data-taking conditions. Frequent adjustments were made to improve data quality and handle increasing luminosities. To minimize internal disuniformities in the sample, it has been split into distinct blocks, each corresponding to a certain set of data-taking conditions. The analysis is performed separately in each block, and then the results are combined to produce the final measurement. This approach helps in the proper cancellation of nuisance asymmetries, as different blocks may be influenced by varying detector-induced asymmetries. The definitions of these blocks and their respective integrated luminosities are listed in Table 8.1.

Block	$\int \mathcal{L} dt [\text{fb}^{-1}]$	Polarity
4.1	0.7	MagDown
4.2	0.4	MagDown
3	0.9	MagUp
2	0.6	MagUp
1	1.2	MagUp
5	1.2	MagUp
6	0.9	MagDown

Table 8.1: Data blocks definition. Ordering is chronological top to bottom, with numbering following a different criterion.

An accurate simulation model of the LHCb detector and trigger conditions during each of those periods is not available at this time. For this reason, the present measurement is essentially data-driven, and only relies on simulations for general guidance or optimization of sensitivity. A sample of $\sim 7\text{k}$ (3.5k for each polarity) simulated $D^0 \rightarrow K_S^0 K_S^0$ decays was used for the purpose of training the kNN classifier that helped in separating signal from combinatorial background. This includes both prompt and secondary-produced charm particles. Also, a sample of $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ containing $\sim 7\text{k}$ events with both prompt and secondary decays is used to tune offline selections on calibration sample.

8.3 Trigger selection

8.3.1 HLT1

Based on the considerations outlined in the previous chapter, the HLT1 lines most relevant for this analysis are **TwoTrackKs** and **TwoKs**. However, since the **TwoTrackKs** HLT1 line is responsible for the selection of the majority of $D^0 \rightarrow K_S^0 K_S^0$ candidates, it has been decided to limit the analysis to just the candidates selected by the **TwoTrackKs** HLT1 line, both for the signal and calibration channels. This simplifies all analysis procedures, while having a minimal impact on the overall sensitivity, and ensures a more robust result given the varying data-taking conditions. This choice provides two samples collected with identical trigger selections, both based on a neutral candidate, minimizing any possible HLT1-induced selection asymmetry.

Two different configurations for the **TwoTrackKs** HLT1 line have been used in collecting the present sample. They are detailed in Table 8.2. They have only minimal differences

Variable	Configuration #1	Configuration #2
χ^2/ndf	< 2.5	< 2.5
$\chi_{\text{IP}}^2(\pi)$	> 50	> 50
$p_{\text{T}}(\pi)$	$> 445 \text{ MeV}/c$	$> 460 \text{ MeV}/c$
$p(\pi)$	$> 5 \text{ GeV}/c$	$> 5 \text{ GeV}/c$
$ m(\pi^+\pi^-) - m(K_S^0) $	$< 45 \text{ MeV}/c^2$	$< 45 \text{ MeV}/c^2$
$\chi_{\text{vtx}}^2(K_S^0)$	< 20	< 20
$\eta(K_S^0)$	$2 < \eta(K_S^0) < 4.2$	$2 < \eta(K_S^0) < 4.2$
$p_{\text{T}}(K_S^0)$	$> 2440 \text{ MeV}/c$	$> 2474 \text{ MeV}/c$
$\cos(\theta_{\text{DIRA}})$	> 0.99	> 0.99
$\cos(\theta_{\pi\pi})$	> 0.99	> 0.99
$\frac{IP(\pi^+) \times IP(\pi^-)}{IP(K_S^0)}$	$> 0.72 \text{ mm}$	$> 0.72 \text{ mm}$

Table 8.2: Selections applied by **TwoTrackKs** HLT1 line in the analyzed 2024 sample.

between each other, and with the nominal configuration.

8.3.2 HLT2

$$D^0 \rightarrow K_S^0 K_S^0$$

The analysis relies on the $D^0 \rightarrow K_S^0 K_S^0$ LL HLT2 Turbo line. The K_S^0 and D^0/D^{*+} candidates are selected according to selections reported in Table 8.3 and 8.4, respectively.

The invariant mass distributions, obtained directly from the HLT2 line all with no additional selections, are displayed in Figure 8.2.

$$D^0 \rightarrow K_S^0 \pi^+ \pi^-$$

At HLT2 level, the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays are selected by a dedicated Turbo line, whose selections do not rely on IP/χ^2 for pion pairs coming from the D^0 , avoiding any bias on the secondaries fraction of the sample. The selections applied by this line are summarized




Figure 8.2: $m(K_S^0)$ (top left), $m(D^0)$ (top right), and Δm (bottom) distributions for $D^0 \rightarrow K_S^0 K_S^0$ candidates selected by the HLT2 line. To produce the plots, 1% of data collected in block 4 have been used. No further selections are applied.

Variable	$K_{SL}^0 \rightarrow \pi_L^+ \pi_L^-$
$\chi^2/\text{ndf}(\pi)$	< 4
$\mathcal{P}_{\text{ghost}}(\pi)$	< 0.7
$\chi_{\text{IP}}^2(\pi)$	> 36
$ m(\pi^+ \pi^-) - m(K_S^0) $	$< 35 \text{ MeV}/c^2$
$p_T(K_S^0)$	$> 500 \text{ MeV}/c$
$z(K_S^0)$	$\in [-100, 500] \text{ mm}$

Table 8.3: Summary of K_S^0 selection in HLT2: these candidates are combined to form D^0 candidates.

Variable	$D^0 \rightarrow K_S^0 K_S^0$
$\sum_{K_S^0} p_T$	$> 1500 \text{ MeV}/c$
$m(K_S^0 K_S^0)$	$\in [1775, 1955] \text{ MeV}/c^2$
$\chi_{\text{vtx}}^2/\text{ndf}(D^0)$	< 10
$\cos(\theta_{\text{DIRA}}(D^0))$	> 0.9994
$m(D^0 \pi_{\text{tag}}) - m(K_S^0 K_S^0)$	$\in [0., 170] \text{ MeV}/c^2$
$p_T(\pi_{\text{tag}})$	$> 200 \text{ MeV}/c$
$p(\pi_{\text{tag}})$	$> 1 \text{ GeV}/c$
$\mathcal{P}_{\text{ghost}}(\pi_{\text{tag}})$	< 0.7
$\chi^2/\text{ndf}(\pi_{\text{tag}})$	< 4
$\chi_{\text{vtx}}^2/\text{ndf}(D^*)$	< 25

Table 8.4: Summary of HLT2 trigger selection on D^0 and D^* candidates.

in Table 8.5. The invariant mass distributions for candidates selected by the HLT2 line are shown in Figure 8.3.

8.4 Offline selections

A set of offline selections is applied on the $D^0 \rightarrow K_S^0 K_S^0$ sample. The selection optimisation relies on an estimate of the statistical sensitivity of the measurement, *i.e.*, the statistical uncertainty on \mathcal{A}^{CP} under the hypothesis that $\mathcal{A}^{CP} = 0$. This is evaluated as the relative error on the signal yield $\frac{\sigma(N_S)}{N_S}$.

8.4.1 Rejection of peaking physical backgrounds

At this stage, the $D^0 \rightarrow K_S^0 K_S^0$ sample is contaminated by other decays, as $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ and $D^0 \rightarrow \pi^+ \pi^- \pi^+ \pi^-$. The same method used in Run 2 analysis to deal with these will be adopted here, relying on the 3D fit for the disentanglement of this background from the $D^0 \rightarrow K_S^0 K_S^0$ component, but reducing its contribution with a selection on the K_S^0 flight distance significance, to improve the S/B of the sample. Since no significant difference is expected with the Run 2 configuration, the same selection is exploited to

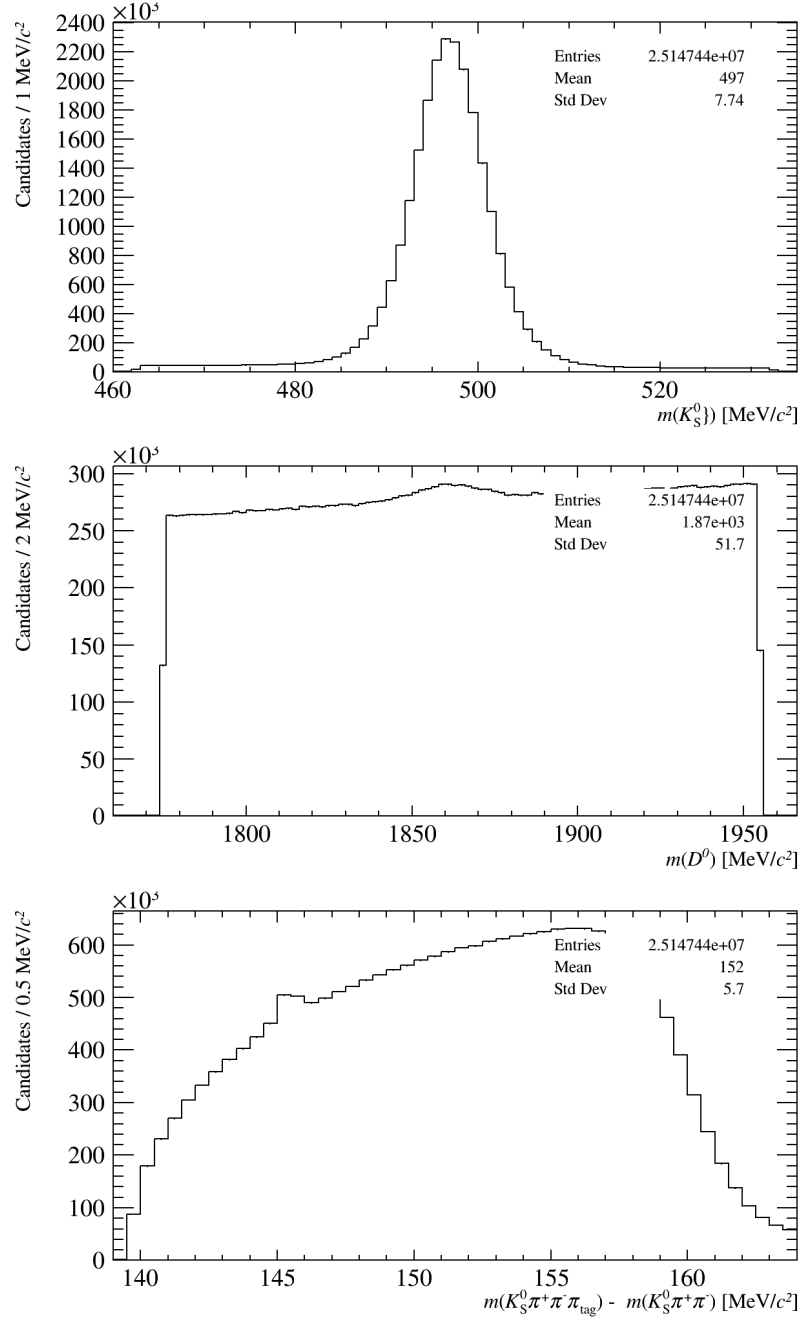


Figure 8.3: $m(K_S^0)$ (top left), $m(D^0)$ (top right), and Δm (bottom) distributions for $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ candidates selected by the HLT2 line. To produce the plots, 1% of data collected in block 4 has been used. No further selections are applied.

Variable	
$\chi^2/\text{ndf}(\pi^\pm)$	< 4
$\mathcal{P}_{\text{ghost}}(\pi^\pm)$	< 0.7
$p_T(\pi^\pm)$	$> 400 \text{ MeV}/c$
$p(\pi^\pm)$	$> 1500 \text{ MeV}/c$
$p_T(\pi^\pm_{\text{from } K_S^0})$	$> 400 \text{ MeV}/c$
$p(\pi^\pm_{\text{from } K_S^0})$	$> 3000 \text{ MeV}/c$
$\chi_{\text{IP}}^2(\pi^\pm_{\text{from } K_S^0})$	> 36
$ m(\pi^+\pi^-)_{\text{from } K_S^0} - m(K_S^0) $	$< 35 \text{ MeV}/c^2$
$\chi_{\text{vtx}}^2/\text{ndf}(K_S^0)$	< 9
$p_T(K_S^0)$	$> 1150 \text{ MeV}/c$
$\text{DOCA}(\pi^+\pi^-)_{\text{from } K_S^0}$	$< 0.15 \text{ mm}$
$\sum_{K_S^0\pi^+\pi^-} p_T$	$> 1500 \text{ MeV}/c$
$\text{DOCA}(\pi^+\pi^-)$	$< 0.11 \text{ mm}$
$\text{DOCA}(K_S^0\pi^-)$	$< 0.11 \text{ mm}$
$\text{DOCA}(K_S^0\pi^+)$	$< 0.11 \text{ mm}$
$\cos(\theta_{\pi^+\pi^-})$	> 0.992
$\cos(\theta_{K_S^0\pi^-})$	> 0.992
$\cos(\theta_{K_S^0\pi^+})$	> 0.992
$\cos(\theta_{\text{DIRA}})(D^0)$	> 0.99
$\chi_{\text{vtx}}^2/\text{ndf}(D^0)$	< 10
$m(K_S^0\pi^+\pi^-)$	$\in [1775, 1955] \text{ MeV}/c^2$
$m(K_S^0\pi^-)$	$< 1851 \text{ MeV}/c^2$
$m(D^0\pi_{\text{tag}}) - m(K_S^0\pi^+\pi^-)$	$\in [0, 160] \text{ MeV}/c^2$
$p_T(\pi_{\text{tag}})$	$> 200 \text{ MeV}/c$
$p(\pi_{\text{tag}})$	$> 1 \text{ GeV}/c$
$\chi^2/\text{ndf}(\pi_{\text{tag}})$	< 4
$\mathcal{P}_{\text{ghost}}(\pi_{\text{tag}})$	< 0.7
$\chi_{\text{vtx}}^2/\text{ndf}(D^*)$	< 25

Table 8.5: Summary of HLT2 selection for $D^0 \rightarrow K_S^0\pi^+\pi^-$ candidates.

reject $D^0 \rightarrow K_S^0\pi^+\pi^-$ and $D^0 \rightarrow \pi^+\pi^-\pi^+\pi^-$ decays, that is $\log \chi_{FD}^2(K_S^0) > 4^1$, applied on both K_S^0 . The effect of this selection can be better understood by looking at the 2D distribution of $\log(\chi_{FD}^2(K_S^0))$ for the two K_S^0 candidates reported in Figure 8.4. A clear distinction is observed among the different decay modes: $K_S^0K_S^0$ decays, where both K_S^0 candidates exhibit a large flight distance; $K_S^0\pi^+\pi^-$ decays, where only one K_S^0 displays a significant flight distance; and 4π decays (along with combinatorial background), where all four tracks originate close to the primary vertex.

¹Definition in Section 4.2.




Figure 8.4: The 2-dimensional distribution of $\log \chi_{FD}^2(K_S^0)$ for K_S^0 candidates of $D^0 \rightarrow K_S^0 K_S^0$ sample. A minimal selection on $m(D^0)$ and Δm is applied to produce the plots: $1845 \text{ MeV}/c^2 < m(D^0) < 1885 \text{ MeV}/c^2$, $143.5 \text{ MeV}/c^2 < \Delta m < 147 \text{ MeV}/c^2$.

8.4.2 D^0 candidate invariant mass selection

A selection is applied on the invariant mass of reconstructed D^0 candidates, as done in Run 2 analysis. This rejects some physical background (such as $D_s^+ \rightarrow K_S^0 K_S^0 \pi^+$) and significantly reduces combinatorial background, as can be seen in the distribution in Figure 8.2.

In the analysis of the 2024 data, special attention must be given to fine-tuning the position and width of the applied invariant mass cut. Due to the varying alignment conditions, the position of the mass peaks, including the D^0 peak, can shift by several MeV/c^2 between different blocks. This effect is compensated by centering the cut on the observed peak position $\mu(D^0)$, instead of the PDG value for $m(D^0)$. The $\mu(D^0)$ values for different blocks are reported in Table 8.6.

Block	$\mu(D^0)$ [MeV/c^2]
4.1	1860.6
4.2	1865.6
3	1866.4
2	1865.9
1	1866.6
5	1866.6
6	1866.6

Table 8.6: Central value of $m(D^0)$ selection for different blocks.

A shift has also been observed when analyzing the invariant mass distributions of D^0 and \bar{D}^0 candidates separately. This behavior is further clarified by examining the fit results reported in Figure 8.5. Because of this effect, a charge-integrated cut on $m(D^0)$




Figure 8.5: $m(D^0)$ peak positions obtained fitting charge-integrated ("tot") and charge-separated distributions in different blocks.

can have a different efficiency for D^0 and \bar{D}^0 , if tight enough, potentially introducing a bias in the CP asymmetry measurement.

In order to avoid this, it has been decided to apply a $|m(D^0) - \mu(D^0)| < 26 \text{ MeV}/c^2$ cut (looser with respect to the $\pm 20 \text{ MeV}/c^2$ window exploited in Run 2 analysis). This value was chosen because it marks the point beyond which the observed raw asymmetry becomes independent of the selection window, ensuring the avoidance of any bias.

8.4.3 Secondary decays

To ensure a proper cancellation of nuisance asymmetries, the fraction f of secondary decays needs to be the same in the signal and calibration sample. This is achieved by refraining from the application of any lifetime-biasing selection having a different resolution for signal and calibration decays, both at the trigger and offline level. To leverage the significant improvement in mass resolution provided by the DTF PV constraint, events are categorized into PV-compatible and PV-incompatible groups. This classification is based on the $\Delta\chi_{DTF}^2$ variable, defined in Equation 4.5. The $\Delta\chi_{DTF}^2$ distribution for simulated $D^0 \rightarrow K_S^0 K_S^0$ decays is shown in Figure 8.6. Because of a minor bug in the Run 3 code implementing DTF, the $\Delta\chi_{DTF}^2$ variable appears to be inverted relative to zero. After verifying that this is the only deviation from the expected distribution, and the one observed in Run 2, the threshold separating PV-compatible and PV-incompatible samples has been adjusted to -15 (as opposed to 15 used in Run 2). The same selection is also applied to the calibration sample.

The Δm distributions for PV-compatible and PV-incompatible samples after the application of all exposed selections are reported in Figure 8.7. It can be seen that the PV-incompatible sample shows no evidence for a signal peak. It has therefore been decided to drop it from further analysis, leaving only one subsample.




Figure 8.6: Distribution of $\Delta\chi^2$ for simulated $D^0 \rightarrow K_S^0 K_S^0$ decays. Primaries and secondaries distributions are normalized to have a unitary integral.




Figure 8.7: Δm distribution for PV-compatible (top) and PV-incompatible (bottom) $D^0 \rightarrow K_S^0 K_S^0$ candidates, after the application of the previously mentioned selections. To produce the plots, data of blocks 1-4 has been used.

8.4.4 Combinatorial background

A cut on the output of a kNN multivariate classifier is used for this purpose. The signal sample provided in input to the classifier consists of the available simulated $D^0 \rightarrow K_S^0 K_S^0$ sample, where all the selections discussed up to this point are applied. As a background sample, real data candidates are coming from $m(D^0)$ sidebands ($|m(D^0) - 1865| > 30$ MeV/ c^2). The value of k and the variables used in the training have been fixed to the values used in Run 2. In particular, $k = 15$ and the following variables are used in the training:

- $p(K_{S1}^0)$: the momentum of the first K_S^0 ;
- $p_T(K_{S2}^0)$: the transverse momentum of the second K_S^0 ;
- $p(D^0)$: the momentum of the D^0 candidate;
- $p_T(D^0)$: the transverse momentum of the D^0 candidate;
- $\max_{\pi, K_{S1}^0} \mathcal{P}_{ghost}$: the maximum value of the ghost probability \mathcal{P}_{ghost} among the two pions originating from the decay of the first K_S^0 ;
- $\max_{\pi, K_{S2}^0} \mathcal{P}_{ghost}$: the maximum value of the ghost probability \mathcal{P}_{ghost} among the two pions originating from the decay of the second K_S^0 ;
- $\max_{\pi, K_{S1}^0} DLL_{K\pi}$: the maximum value of $DLL_{K\pi}$ among the two pions originating from the decay of the first K_S^0 ;
- $\max_{\pi, K_{S2}^0} DLL_{K\pi}$: the maximum value of $DLL_{K\pi}$ among the two pions originating from the decay of the second K_S^0 ;
- $\max_{\pi, K_{S1}^0} DLL_{p\pi}$: the maximum value of $DLL_{p\pi}$ among the two pions originating from the decay of the first K_S^0 ;
- $\min(\log(\chi_{IP}^2)_{K_{S1}^0})$: the minimum value of the logarithm of χ_{IP}^2 among the two pions originating from the decay of the first K_S^0 ;
- $\min(\log(\chi_{IP}^2)_{K_{S2}^0})$: the minimum value of the logarithm of χ_{IP}^2 among the two pions originating from the decay of the second K_S^0 ;
- $\chi_{vtx}^2(D^0)$: the χ^2 of the vertex fit for the D^0 candidate;
- $\chi_{vtx}^2(D^*)$: the χ^2 of the vertex fit for the D^* candidate;
- $\chi_{vtx}^2(K_{S2}^0)$: the χ^2 of the vertex fit for the second K_S^0 candidate;
- $\theta_H(K_{S1}^0)$: the helicity angle for the K_{S1}^0 candidate;
- $\theta_H(K_{S2}^0)$: the helicity angle for the second K_S^0 candidate;
- $\theta_H(D^0)$: the helicity angle for the D^0 candidate.




Figure 8.8: Performance of kNN classifier on training and test samples.

As a cross-check, the performance of the training sample was compared against results obtained from an independent, labeled dataset of identical size and composition (referred to as “test sample” in the following). The results of this comparison are presented in Figure 8.8.

Variables that could bias the fraction of secondary decays, such as $\chi_{IP}^2(K_S^0)$, were intentionally excluded from the kNN training. However, $\chi_{IP}^2(K_S^0)$ is a highly effective variable for reducing background, particularly from real K_S^0 candidates combining to form a fake D^0 , as demonstrated in Figure 8.9. To leverage its discriminating power, a fixed cut




Figure 8.9: Comparison between signal and background distributions for the $\chi_{IP}^2(K_S^0)$ variable in $D^0 \rightarrow K_S^0 K_S^0$ decays. Simulated candidates (MC) are used as a proxy for the signal, while data taken from $m(D^0)$ sidebands are used as a proxy for the background. The distributions are normalized to have a unitary integral.

is applied to this variable, specifically targeting the K_S^0 candidate that triggered HLT1. If both K_S^0 candidates triggered HLT1, one is randomly chosen. This approach allows the same selection to be applied to the calibration channel, ensuring a consistent impact on the fraction of secondaries f , as both samples were collected using the same trigger. The effectiveness of this strategy is supported by the $\chi_{IP}^2(K_S^0)$ distributions shown in Figure 8.10, which reveals the excellent agreement between signal and calibration samples.

Therefore, to suppress the combinatorial background component, selections are applied to

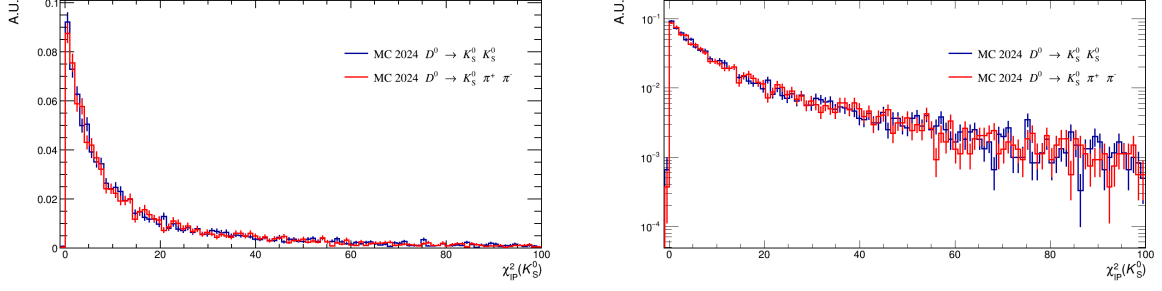


Figure 8.10: Comparison between $\chi_{\text{IP}}^2(K_S^0)$ distribution for $D^0 \rightarrow K_S^0 K_S^0$ and $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ simulated samples in linear (left) and logarithmic (right) scale.

both the kNN output and the $\chi_{\text{IP}}^2(K_S^0)$ of the K_S^0 candidate that triggered HLT1.

Following the approach of the Run 2 analysis, to maximize the use of the available statistics a direct cut on the classifier output is avoided. Instead, three bins are defined based on the kNN output: the lowest-purity bin is rejected, while the analysis is performed separately on the remaining two higher-purity bins. The same $\chi_{\text{IP}}^2(K_S^0)$ selection is applied for both kNN output bins. The thresholds for $\chi_{\text{IP}}^2(K_S^0)$ and the kNN output are optimized simultaneously by identifying the configuration that minimizes the combined value of σ_S/N_S across the two bins. The optimization chooses as optimal point the one given by the following thresholds:

- bin 1 (low-purity): $0.1 < \text{kNN out} \leq 0.9$;
- bin 2 (high-purity): $\text{kNN out} > 0.9$;
- $\chi_{\text{IP}}^2(K_S^0) > 11$.

8.4.5 Equalization of kinematic region with calibration channel

In this analysis, regions with significant detection asymmetries are not excluded using fiducial cuts, as the procedure outlined in Section 8.7 for subtracting nuisance asymmetries from the signal sample remains effective even in the presence of large asymmetries, provided they do not reach exactly 100%. Notably, fiducial requirements were also omitted in the Run 2 measurement. However, to ensure accurate local correction of detector-induced asymmetries, it is essential that the yield ratio between calibration and signal candidates does not deviate significantly from the selected scaling factor, $k = 30$. This condition has been validated for the signal and calibration samples following the application of the offline selections described here and in Section 8.6. As a result, a $p_z(\pi_{\text{tag}}) > 3 \text{ GeV}/c$ selection was applied to exclude a kinematic region where the calibration sample lacked sufficient statistics.

8.4.6 Removal of multiple candidates

After applying the described offline selections, multiple candidates — defined as $D^0 \rightarrow K_S^0 K_S^0$ decays reconstructed within the same bunch crossing — are rejected. These

make up approximately 13.5% of the sample after the offline selections (24% before the application of those). When multiple candidates are identified, a single, randomly selected candidate is retained.

8.5 Results on charge-integrated sample

The projections of the 3D fit to the charge-integrated Δm distribution for the two bins, after applying the optimized selections and the rejection of multiple candidates, are shown in Figure 8.11, with the corresponding fit results summarized in Table 8.7. Combining the results from both bins yields a $\sigma(N)/N$ value of 1.46%.

sample	parameter	result 3-d fit
$\text{kNN out} \geq 0.9 \wedge \chi_{\text{IP}}^2(K_S^0) > 11$	n_{sig}	4311 ± 73
	n_{bkg}	3204 ± 66
$0.1 < \text{kNN out} < 0.9 \wedge \chi_{\text{IP}}^2(K_S^0) > 11$	n_{sig}	10729 ± 163
	n_{bkg}	52570 ± 259
	$\mu(\Delta m)$	$(145.42 \pm 0.01) \text{ MeV}/c^2$
	$\sigma(\Delta m)$	$(0.49 \pm 0.02) \text{ MeV}/c^2$
	$\alpha(\Delta m)$	-0.03 ± 0.04
	$\nu(\Delta m)$	1.73 ± 0.08

Table 8.7: Results of the 3D charge-integrated fit after the application of the reported offline selections. The parameters of the Δm distribution are fitted simultaneously in the two bins.

8.6 The calibration sample

As mentioned before, a calibration sample of $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays is used to cancel out the spurious effect of production and detection asymmetries in the signal sample. This calibration sample must be carefully prepared for the cancellation to be accurate.

The sample was obtained from exactly the same HLT1 line of the signal, **TwoTrackKs**. At HLT2, a dedicated, vertex-unbiased line was used, to avoid introducing any possible bias on the fraction of secondaries. This line is different and separate from the "standard" $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ line, used for obtaining signal samples for other measurements. Offline selections are then applied to replicate the requirements applied to the signal sample, to ensure an accurate cancellation of detector effects. In particular, this includes the $\chi_{\text{IP}}^2(K_S^0) > 11$ and $\Delta\chi_{DTF}^2 > -15$ requirements.

However, this is not sufficient to provide a calibration sample of the needed purity to ensure an accurate cancellation. In fact, the lack of the customary $\text{IP}/\chi_{\text{IP}}^2$ trigger selections on the D^0 daughters leads to a pretty low purity of the collected sample, as seen in Figure 8.3. Therefore, further offline selections are necessary to improve S/B .

The further selections applied are listed in Table 8.8. The $m(D^0)$ cut is centered around $\mu(D^0)$, as for the signal sample (see section 8.4.2). The cuts on $\cos\theta(\pi^+ \pi^-)$, $\text{DOCA}(\pi^+ \pi^-)$ are chosen by the rectangular cuts optimization method available in TMVA. Multiple candidates are removed with the same procedure applied to the signal sample. From




Figure 8.11: Projections of the 3D charge-integrated fit to the signal sample after the application of the reported offline selections. On the left (right) column plots for candidates satisfying $kNN_{out} > 0.9$ ($0.1 < kNN_{out} \leq 0.9$) are shown. A $\chi^2_{IP}(K_S^0) > 11$ cut is applied to both samples. The top row shows the Δm distributions. The invariant mass distributions of the two K_S^0 candidates are shown in the other two rows.

Variable	
$\chi^2_{IP}(K_S^0)$	> 11
$\Delta\chi^2_{DTF}$	> -15
$ m(K_S^0) - 497.6 \text{ MeV}/c^2 $	$< 10 \text{ MeV}/c^2$
$ m(D^0) - \mu(D^0) $	$< 30 \text{ MeV}/c^2$
$ m(K_S^0 \pi_{OS}) - 892 \text{ MeV}/c^2 $	$< 50 \text{ MeV}/c^2$
$\chi^2_{vtx}/ndf (D^0, K_S^0, D^*)$	< 5
$\cos \theta(\pi^+ \pi^-)$	> 0.9977
$DOCA(\pi^+ \pi^-)$	< 0.022
$p_z(\pi_{tag})$	$> 3 \text{ GeV}/c$
Removal of multiple candidates	

Table 8.8: Summary of selections applied to the calibration sample.

the Dalitz plot of the candidates, reported in Figure 8.12, it clearly appears that the decay proceeds mainly through the $K^*(892)$ subresonance. For this reason, the decision was made to specifically select the $K^*(892)$ resonance region in the D^0 Dalitz plot, by

requiring $|m(K_S^0 \pi_{OS}) - 892 \text{ MeV}/c^2| < 50 \text{ MeV}/c^2$, where π_{OS} represents the D^0 daughter pion that has an opposite sign charge with respect to the tagging pion. This allows to reject a large chunk of the combinatorial background, that is spread over the whole Dalitz space, while keeping a substantial fraction of real decays, that are concentrated in this resonance region.




Figure 8.12: Top plots shows the distribution of Δm for $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ candidates after application of all offline selections. Bottom plot shows the Dalitz plot distribution for D^0 candidates, after application of all offline selections with the exception of $|m(K_S^0 \pi_{OS}) - 892 \text{ MeV}/c^2| < 50 \text{ MeV}/c^2$, which is reported in red. π_{OS} and π_{SS} are the D^0 daughter pions with opposite sign charge and same sign charge with respect to the tagging pion respectively. Dataset blocks 1-4.2 are exploited to produce these plots.

The combined effect of all of those offline selections leads to the achievement of a $S/B \sim 23$ in the Δm distribution (Figure 8.12). To additionally improve the purity of the sample used for the nuisance asymmetries cancellation, only candidates satisfying $|\Delta m - 145.45 \text{ MeV}/c^2| < 1.5 \text{ MeV}/c^2$ are used.

8.7 Evaluation of event weights

To correct the signal sample for the spurious asymmetry effects, we use the same procedure of the Run 2 analysis, of weighting each individual event with a correction factor, derived

from the calibration sample described in the previous section. The CP asymmetry of $D^0 \rightarrow K_S^0 K_S^0$ is then extracted through a 3D fit to the Δm and the two $m(K_S^0)$ distributions in the weighted sample, based on the fit model described in the next section.

However, a key difference with Run 2 is the use of a calibration mode whose final state is not charge-symmetric, unlike the $D^0 \rightarrow K^+ K^-$ mode. It is worth noting that this asymmetry is made worse by the application of the Dalitz cut mentioned in the previous section, but would still be present even in the absence of any such cut. To eliminate this asymmetry, that does not match a corresponding asymmetry in the signal sample needing cancellation, a pre-processing was applied to the calibration sample before the step of weight determination. This process, described in detail in the next section, ensures that the calibration sample behaves symmetrically with respect to the D^0 final state. After this preliminary step, the weighting procedure applied to the signal sample then follows exactly the same procedure of Run 2: each $D^0(\bar{D}^0) \rightarrow K_S^0 K_S^0$ event is weighted by the factor:

$$w^\pm(\vec{p}_0) = \frac{n_C^+(\vec{p}_0) + n_C^-(\vec{p}_0)}{2n_C^\pm(\vec{p}_0)} [1 \pm \mathcal{A}^{CP}(K_S^0 \pi^+ \pi^-)],$$

where \vec{p}_0 is the D^0 tri-momentum, and n_C^\pm is the numerical density of $D^{*\pm}$ decays in the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ calibration channel, lying in the region of the D^0 parameter space where the $K_S^0 K_S^0$ event lies. The density functions n_C , and in particular the ratio $n_C^\pm(\vec{p}_0)/(n_C^+(\vec{p}_0) + n_C^-(\vec{p}_0))$, are determined with the help of a kNN classifier, by splitting the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ sample according to π_{tag} charge. For each $D^0 \rightarrow K_S^0 K_S^0$ event, the value of the corresponding kNN output is evaluated, and it is used to calculate the weight for the fit. Since we cannot properly calibrate asymmetries in regions with 100% asymmetry, candidates with kNN output larger than 0.9 or smaller than 0.1 are removed from the final sample.

A difference between the procedure here reported and the one used in Run 2 analysis is the adoption of π_{tag} observables, in place of D^0 ones, in the weights calculation. While, in principle, using D^0 or π_{tag} variables should yield equivalent results in describing the \vec{p}_0 phase-space, in practice, their resolutions differ, as D^0 is a composite object while π_{tag} is a single track. Since detection asymmetry arises from where the π_{tag} intersects the subdetector modules, we prefer to use its direct kinematic variables. For this purpose, three observables were chosen: $|p_x(\pi_{tag})|$, $p_y(\pi_{tag})$ and $p_z(\pi_{tag})$ to map the \vec{p}_0 phase-space. $|p_x(\pi_{tag})|$ is preferred to $p_x(\pi_{tag})$, because its exploitation automatically reduces the regions with high asymmetries. The main motivation for this choice is the reduction of the systematic uncertainty related to the finite size of the calibration sample. A value of $k = 30$ is chosen for the kNN algorithm. This choice reflects the yield ratio between $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ and $D^0 \rightarrow K_S^0 K_S^0$, which is approximately 40. Selecting $k > 40$ risks inaccuracies in the description of local asymmetries, as it includes candidates from slightly different kinematic regions. Conversely, choosing a much smaller k leads to weights with larger statistical uncertainties. With this weighting procedure, it is not necessary for the kinematic variables in the signal and calibration samples to have identical distributions. The only requirement is that regions with sufficient statistics in the calibration sample are utilized. This requirement was verified, and led to the addition of a further cut on $p_z(\pi_{tag})$ (see Section 8.4.5).

The distributions of the weights for different blocks are shown in Figure 8.13.

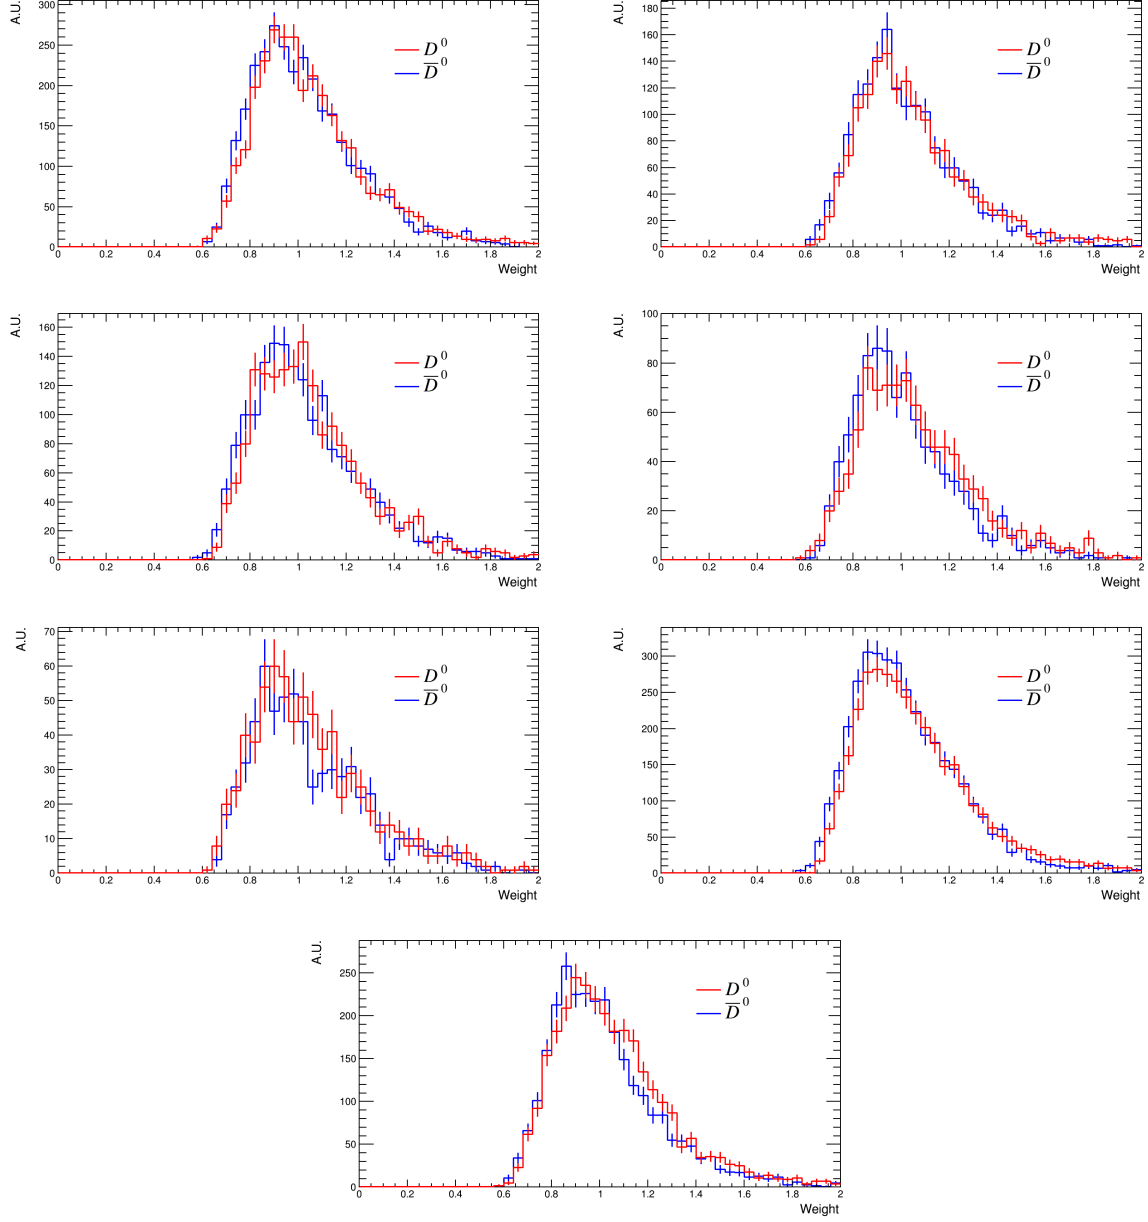


Figure 8.13: Distributions of weights applied to $K_S^0 K_S^0$ sample. Reading from top left to bottom right distribution for blocks 1, 2, 3, 4.1, 4.2, 5 and 6 (bottom) are shown.

8.7.1 Correction of the intrinsic asymmetry of the calibration sample

The $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ mode primarily proceeds via the $D^0 \rightarrow K^{*-} \pi^+$ intermediate decay. This is a so-called "self-tagging" decay, meaning that observation of the final state allows to determine the flavor of the parent D^0 particle. After the decay of the intermediate state as $K^{*-} \rightarrow K_S^0 \pi^-$, part of this determination is lost in the possibility of confusing the two final pions; but not completely, as the pion originating directly from the D^0 has a different kinematics from the pion produced by the resonance decay. This residual difference can easily couple to disuniformities in the detector efficiency charge asymmetry to produce as a final result a difference in efficiency between D^0 and \bar{D}^0 candidates. This spurious asymmetry has no match in the $K_S^0 K_S^0$ mode, and must therefore be eliminated before proceeding to the cancellation. This effect has been canceled by a pre-processing of the sample, consisting in weighting each $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ candidate in a way to symmetrize the kinematic distributions of the two pions, that is to ensure that the joint kinematics probability density $p(\pi^+, \pi^-) = p(\pi^-, \pi^+)$ is symmetrical for pion exchange, for both flavors of the initial charm. This has the effect of zeroing any D^0 detection asymmetries related to the asymmetry of the decay. Therefore, after the application of this pre-processing, it is possible to exploit the weighted sample to correct the signal sample just as it was done with the symmetrical decay $D^0 \rightarrow K^+ K^-$.

In this procedure, the two pions coming from the D^0 are labeled as π_1 and π_2 , where π_1 (π_2) is defined as the pion with higher (lower) momentum. The labeling choice is arbitrary and does not impact on the result. The sample is then split in two halves depending on the π_1 charge. Then, the weights that need to be applied to the calibration sample to cancel D^0 detection asymmetry are defined for the D^0 candidates as:

$$w_+(D^0) = \left(\frac{n_{D^0}^+(p, q, P_0)}{n_{D^0}^-(p, q, P_0)} \right)^{-1/2}; \quad w_-(D^0) = \left(\frac{n_{D^0}^+(p, q, P_0)}{n_{D^0}^-(p, q, P_0)} \right)^{1/2}$$

where $w_+(D^0)$ ($w_-(D^0)$) is the weight applied to candidates in which the π_1 is positive (negative). Similarly, weights for \bar{D}^0 candidates are defined as:

$$w_+(\bar{D}^0) = \left(\frac{n_{\bar{D}^0}^+(p, q, P_0)}{n_{\bar{D}^0}^-(p, q, P_0)} \right)^{-1/2}; \quad w_-(\bar{D}^0) = \left(\frac{n_{\bar{D}^0}^+(p, q, P_0)}{n_{\bar{D}^0}^-(p, q, P_0)} \right)^{1/2}$$

where $w_+(\bar{D}^0)$ ($w_-(\bar{D}^0)$) is the weight applied to candidates in which the π_1 is positive (negative).

The weights are calculated using a kNN classifier, as the ratio in the weight definition can be derived directly from the kNN output. Candidates with kNN output values of out kNN = 0 or out kNN = 1 are excluded from the analysis since these correspond to regions of complete asymmetry, resulting in weights of 0 or ∞ . While this exclusion eliminates only a small fraction of $K_S^0 K_S^0$ candidates in the signal sample, it reduces the statistics of the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ calibration sample by approximately 50%. This significant reduction is due to the larger asymmetry affecting the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ final state compared to the asymmetry associated with the π_{tag} in $D^0 \rightarrow K_S^0 K_S^0$ decays.

The weights are computed as a function of the kinematic space of the D^0 decay by training the kNN algorithm using the variables that describe the kinematics of the two

pions ($p(\pi_1), \eta(\pi_1), \phi(\pi_1), p(\pi_2), \eta(\pi_2), \phi(\pi_2)$). In principle, the kinematic variables of the D^0 itself should also be included in the training to fully describe $n^\pm(p, q, \vec{p}_0)$. However, it has been verified that, in this specific case, excluding these variables from the training does not produce a significant impact on the final results. The kNN training and weighting is performed separately for each data block and particle species (D^0 and \bar{D}^0).

The positive outcome of this procedure can be checked by looking at Figure 8.14 and 8.15, displaying the kinematic distributions of π^+ and π^- particles, before and after weighting, respectively. The shape of opposite-sign particles matches after the application of the procedure, as expected.

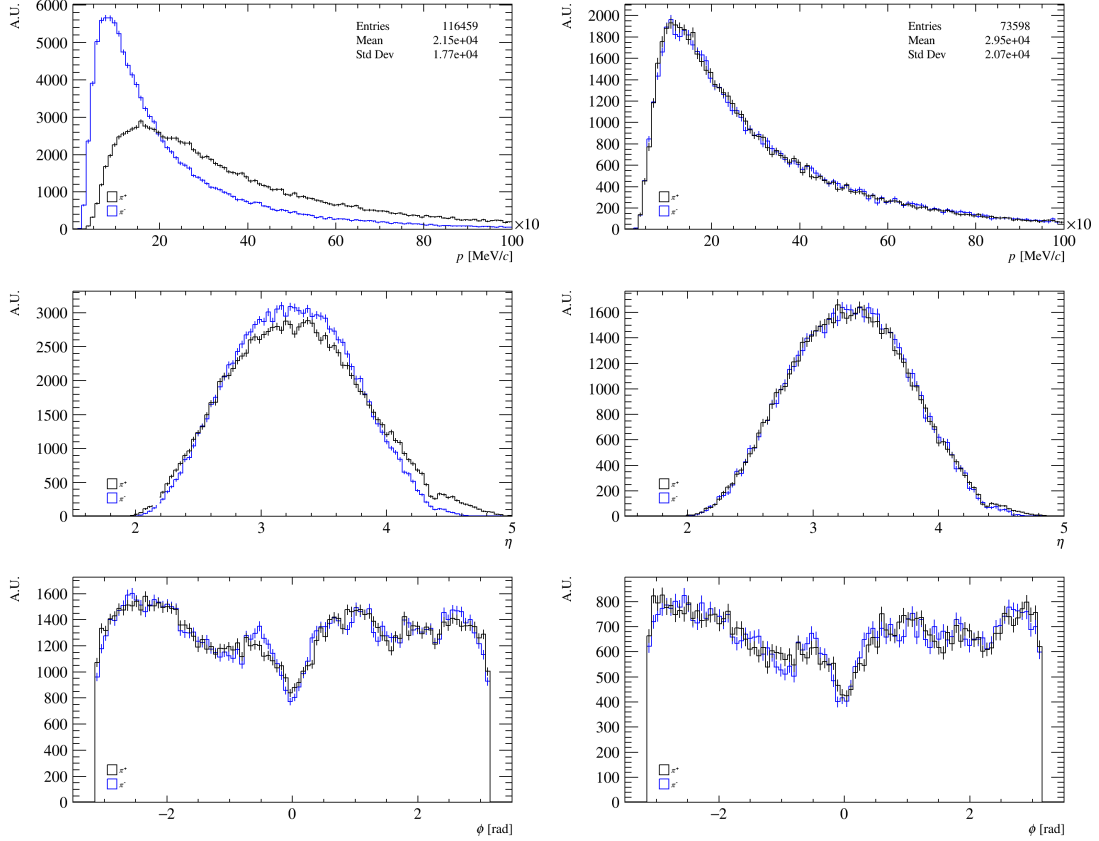


Figure 8.14: Distributions of p (top row), η (central row), ϕ (bottom row) for π^+ and π^- for the block 1 D^0 candidates. The distributions before (after) the weighting are plotted on the left (right) column.

The values of \mathcal{A}_{raw} for each block for every step of the procedure (without any correction, and after the weighting for the correction of the D^0 detection asymmetry) are summarized in Table 8.9. The \mathcal{A}_{raw} value for MagUp (blocks 1-2-3-5) and MagDown (blocks 4.1-4.2-6) before the weighting procedure is $\mathcal{A}_{raw} = (-1.47 \pm 0.12)\%$, $\mathcal{A}_{raw} = (-0.71 \pm 0.17)\%$ respectively.

For much more details about the procedure and the numerous cross-checks performed to validate it, the reader is referenced to Ref. [174], a Master's thesis devoted to this specific topic, that has been carried out with my co-supervision.

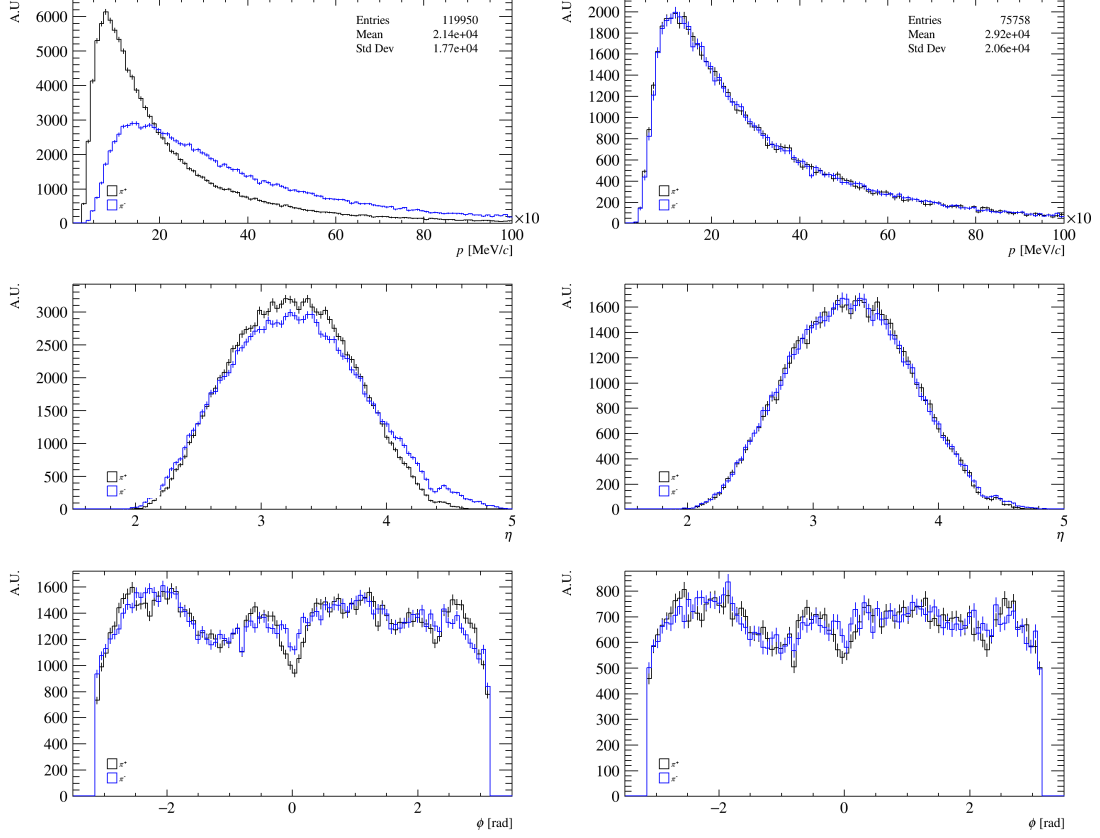


Figure 8.15: Distributions of p (top row), η (central row), ϕ (bottom row) for π^+ and π^- for the block 1 \bar{D}^0 candidates. The distributions before (after) the weighting are plotted on the left (right) column.

Block	$\mathcal{A}_{raw} [\%]$	\mathcal{A}_{raw} after (D^0) correction $[\%]$
1	-1.3 ± 0.2	-1.1 ± 0.3
2	-1.1 ± 0.3	-0.8 ± 0.4
3	-1.4 ± 0.3	-1.0 ± 0.4
4.1	-0.9 ± 0.3	-1.9 ± 0.5
4.2	-0.8 ± 0.4	-1.2 ± 0.6
5	-1.9 ± 0.2	-1.8 ± 0.3
6	-0.6 ± 0.2	-1.0 ± 0.3

Table 8.9: \mathcal{A}_{raw} values before and after D^0 detection asymmetry correction, for each block.

8.8 Fit procedure

After applying the offline selections and calculating the weights to cancel nuisance asymmetries in the $D^0 \rightarrow K_S^0 K_S^0$ sample, $\mathcal{A}^{CP}(K_S^0 K_S^0)$ is extracted through a binned, weighted maximum likelihood fit to the Δm and the two $m(K_S^0)$ distributions. The fit is performed simultaneously on the D^{*+} and D^{*-} candidates, where the signal asymmetry, \mathcal{A}^{CP} , is treated as a free parameter and is kept blind. To avoid an involuntary unblinding, in the

report of fit results only residuals of the fit projections on the three dimensions will be reported, and not the distributions of the relative fits. The background asymmetry, A_{bkg} , is also allowed to vary freely. Within each data block, the two kNN output bins are fitted simultaneously.

The total probability density function *pdf* is modeled as the sum of eight components: one for the signal, which peaks in all three observables, and seven for different background sources. These background components include a residual $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ background, which peak in Δm and one $m(K_S^0)$ distribution but not the other, as well as combinations of unrelated particles. These include cases such as two real K_S^0 forming a fake D^* , one real K_S^0 combined with a fake one, and random pion combinations resulting in fake K_S^0 and fake D^* candidates. The parameters defining the signal and background PDFs are partially shared between D^{*+} and D^{*-} mass distributions.

The signal component in the Δm distribution is described with a Johnson S_U distribution [154] for all the blocks:

$$j(x) \propto \left[1 + \left(\frac{x - \mu}{\sigma} \right)^2 \right]^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \left[-\alpha + (\nu \sinh)^{-1} \left(\frac{x - \mu}{\sigma} \right) \right]^2 \right\}. \quad (8.1)$$

The $m(K_S^0)$ peak is modeled as the sum of two Gaussian functions with a shared mean ($G(x)$). The standard deviations (σ) of the two Gaussian components are related by the equation $\sigma_2 = k_\sigma \sigma_1$, where σ_1 and k_σ are free parameters determined by the fit. This parameterization is mathematically equivalent to directly fitting σ_1 and σ_2 ; however, it has been chosen as it reduces correlations between the parameters and improves the fit's convergence. Both $m(K_S^0)$ distributions are described by the same probability density function (*pdf*). The signal *pdf* is therefore described as the product of the three reported *pdfs*

$$S(\Delta m, m(K_{S1}^0), m(K_{S2}^0)) = j(\Delta m) \cdot G(m(K_{S1}^0)) \cdot G(m(K_{S2}^0)). \quad (8.2)$$

An empirical threshold function is used to describe the non-peaking component in the Δm distribution of bin 2, that is:

$$b(x) \propto \{1 - \exp[-c(x - m_\pi)]\} + b(x - m_\pi). \quad (8.3)$$

Because of a lower background contribution, a first-order Chebyshev polynomial is used to describe the same component in bin 1. The non-peaking component in the $m(K_S^0)$ distribution is described with a first-order Chebyshev polynomial $pol(x)$. The total background *pdf* is the sum of seven different components:

$$\begin{aligned} B(\Delta m, m(K_{S1}^0), m(K_{S2}^0)) = & f_1 \cdot j(\Delta m) \cdot G(m(K_{S1}^0)) \cdot pol(m(K_{S2}^0)) \\ & + f_2 \cdot j(\Delta m) \cdot pol(m(K_{S1}^0)) \cdot pol(m(K_{S2}^0)) \\ & + f_3 \cdot b(\Delta m) \cdot G(m(K_{S1}^0)) \cdot G(m(K_{S2}^0)) \\ & + f_4 \cdot b(\Delta m) \cdot G(m(K_{S1}^0)) \cdot pol(m(K_{S2}^0)) \\ & + f_5 \cdot b(\Delta m) \cdot pol(m(K_{S1}^0)) \cdot G(m(K_{S2}^0)) \\ & + f_6 \cdot j(\Delta m) \cdot pol(m(K_{S1}^0)) \cdot G(m(K_{S2}^0)) \\ & + f_7 \cdot b(\Delta m) \cdot pol(m(K_{S1}^0)) \cdot pol(m(K_{S2}^0)). \end{aligned} \quad (8.4)$$

The first (and sixth) component peaks in the Δm and $m(K_{S1}^0)$ distributions but does not peak in the $m(K_{S2}^0)$ variable. This corresponds to $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ decays. The second

component, which peaks only in the Δm distribution, corresponds to $D^0 \rightarrow \pi^+\pi^-\pi^+\pi^-$ decays. The third component, peaking only in the two $m(K_S^0)$ distributions, corresponds to random combinations of true K_S^0 mesons that result in a fake D^{*+} . In principle, $D_s^+ \rightarrow K_S^0 K_S^0 \pi^+$ decays could contribute to this component, but they are expected to be fully removed by the applied cut on $m(D^0)$. The fourth and fifth components (which are identical by construction) arise from events where a real K_S^0 is combined with random pions to create a fake D^{*+} . Finally, the seventh component accounts for random combinations of pions that result in two fake K_S^0 mesons and a fake D^{*+} .

The full fit, performed without constraining any parameter to zero, was tested on data from Block 1, which contains the highest statistics. The parameters f_1 (equal to f_6 by construction) and f_2 were found to be very small and statistically consistent with zero. Therefore, to improve the fit convergence, these parameters were fixed to zero in all subsequent fits.

8.9 Results

The CP asymmetry is extracted both with the event weights discussed in Section 8.7, and with unweighted events as a check. This led to extracting both a raw asymmetry \mathcal{A}^{raw} , and the corrected asymmetry \mathcal{A}^{CP} , representing the measurement of the actual physics parameter. For both of these, only the *blind* values are reported, where an unknown random number has been added. This number is however the same for all the reported results, allowing for easy comparisons amongst them. This is necessary because at the time of this writing, this measurement, although already officially approved by the competent physics working group, has not yet received the final collaboration-wide approval. This final step, preceding public dissemination of the result, is expected to happen in the first months of 2025. The full result will hopefully be available in time for inclusion in the final version of the present document.

The complete fit results for all analyzed data blocks, together with the pull distribution of fit projections for unweighted and unweighted fit to all blocks can be found in Appendix A. The blind \mathcal{A}^{CP} results obtained in different blocks are summarized in Table 8.10, together with the results of unweighted fits and the yield of each bin.

Block	Yield	$\mathcal{A}^{raw}(K_S^0 K_S^0)$ (blind) [%]	$\mathcal{A}^{CP}(K_S^0 K_S^0)$ (blind) [%]
1	3413 ± 94	52.00 ± 2.19	52.83 ± 2.26
2	1703 ± 67	51.74 ± 3.11	53.27 ± 3.20
3	1933 ± 67	53.33 ± 2.97	54.57 ± 3.07
5	3796 ± 114	51.50 ± 2.13	54.42 ± 2.18
4.1	1067 ± 50	57.51 ± 4.05	59.54 ± 4.19
4.2	772 ± 45	60.14 ± 4.78	59.58 ± 4.68
6	3132 ± 100	59.47 ± 2.32	60.40 ± 2.27
Tot	15816 ± 213	54.27 ± 1.03	55.83 ± 1.05

Table 8.10: Values for yield, blind $\mathcal{A}^{raw}(K_S^0 K_S^0)$ (unweighted fits) and blind $\mathcal{A}^{CP}(D^0 \rightarrow K_S^0 K_S^0)$ (weighted fits) obtained from fits in different blocks.

The final (blind) result given by the combination of all different blocks is:

$$\mathcal{A}^{CP}(K_S^0 K_S^0) = (55.83 \pm 1.05)\%.$$

The achieved resolution of 1.05% makes this measurement the most precise determination of this quantity to date.

For comparison, the result of unweighted fits combining all blocks is $\mathcal{A}^{raw}(K_S^0 K_S^0) = (54.27 \pm 1.03)\%$. The total correction applied through weighting amounts then to 1.56%, representing the cumulative effect of all nuisance asymmetries affecting the $D^0 \rightarrow K_S^0 K_S^0$ sample prior to the correction. Their magnitude is comparable to that observed in Run 2 charm analyses.

8.10 Cross-checks and systematics

Several cross-checks have been performed to ensure that everything is under control in the measurement. Only the most important ones will be reported here.

8.10.1 Check of consistency of \mathcal{A}^{CP} between data-taking periods

An important cross-check is to ensure consistency between the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurements obtained in each block/bin. Results are plotted in Figure 8.16, and are all fully statistically consistent with each other. The results split by magnet polarity (Table 8.11) are consistent at the 2.8σ level, considering only statistical uncertainties. It is important to note that this difference stems from a difference in the low-statistics signal sample, while the much larger calibration sample shows no significant asymmetry difference between the two samples, with a much better resolution (this can be seen from the same Table 8.11). This leads to conclude that this difference (that is anyway averaged out in the final result) must be of purely statistical origin. This is further reinforced by the results of the check in the next subsection.




Figure 8.16: \mathcal{A}^{CP} values in different blocks (1, 2, 3, 4.1, 4.2, 5, 6) and in the different bins (LP = low purity, *i.e.* bin 2; HP = high purity, *i.e.* bin 1). The red line corresponds to the weighted average. Left: results before applying the correction for nuisance asymmetries. Right: results after applying the correction for nuisance asymmetries.

	$\mathcal{A}^{CP}(K_S^0 K_S^0)$ [%]	
	before weighting	after weighting
Mag Up	52.02 ± 1.24	53.76 ± 1.28
Mag Down	59.16 ± 1.84	60.11 ± 1.84
Tot	54.27 ± 1.03	55.83 ± 1.05

Table 8.11: $\mathcal{A}^{CP}(D^0 \rightarrow K_S^0 K_S^0)$ [%] obtained for the two magnet polarities, before and after the weighting.

8.10.2 Closure test of cancellation of spurious asymmetries

To validate the procedure for the cancellation of detector and production asymmetries, it has been applied to the $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ calibration sample itself, verifying that the resulting asymmetry is compatible with zero. In this procedure, half of the sample is used to evaluate the weights, which are then applied to the other half of the sample before measuring its asymmetry. The results, presented in Table 8.12, show values compatible with zero, as expected if the calibration procedure is functioning correctly. This is true both for the single blocks and for the integrated value.

Datablock	$\mathcal{A}(K_S^0 \pi^+ \pi^-)$ [%]	
	before weighting	after weighting
1	-1.5 ± 0.3	0.0 ± 0.4
2	-1.2 ± 0.3	-0.2 ± 0.6
3	-1.5 ± 0.3	0.2 ± 0.6
4.1	-0.9 ± 0.5	-0.5 ± 0.7
4.2	0.1 ± 0.6	0.6 ± 0.9
5	-1.7 ± 0.3	0.3 ± 0.4
6	-0.8 ± 0.3	0.1 ± 0.4
Tot	-1.25 ± 0.13	0.08 ± 0.19

Table 8.12: $\mathcal{A}(D^0 \rightarrow K_S^0 \pi^+ \pi^-)$ [%] before and after the nuisance asymmetries correction, for each block. To note that the MagUp vs MagDown comparison is already included in the splitting of the blocks. For the values after the weighting, half of the sample is used to evaluate the weights and half of the sample is used to measure the asymmetry. The value $k = 30$ is used for the kNN training.

Systematic effects

8.10.3 K_S^0 decays and regeneration effects

The CP violation, mixing, and regeneration effects due to the neutral kaons present in the final state of signal and calibration samples could in principle alter the results of this measurement. For the $K_S^0 K_S^0$ final state, where two K_S^0 are present, these effects cancel out in the \mathcal{A}^{CP} definition. However, they may still be relevant for the calibration sample.

These effects have been extensively studied during Run 2 and, for K_S^0 reconstructed from long tracks, are found to be $\leq 10^{-3}$ [161].

However, because of the different detector, material distribution is different in Run 3, and the impact of this effect has been evaluated. The variation of \mathcal{A}^{CP} (*i.e.* after the application of all the weights) has been studied as a function of the K_S^0 flight distance in the calibration sample (Figure 8.17). The distribution is compatible with being flat, confirming that any effect of this kind is negligible. A systematic uncertainty is assigned, from fitting the distribution with a non-zero slope (Figure 8.18). The value for the angular coefficient extracted from the fit is 0.0007 ± 0.0009 . If we consider the value of 0.0016 (hence a 1σ deviation of the angular coefficient) and an average flight distance of about 150 mm, we obtain an average variation for the asymmetry of 0.24%, which is taken as a systematic uncertainty on the measurement.




Figure 8.17: Left: flight distance(K_S^0) distributions for D^0 and \bar{D}^0 candidates. Right: $A_{CP}(K_S^0 \pi^+ \pi^-)$ as a function of the K_S^0 flight distance. Candidates from all blocks are included and have been selected with the application of all offline cuts. The p-value of the constant fit is 0.12.




Figure 8.18: $\mathcal{A}^{CP}(K_S^0 \pi^+ \pi^-)$ as a function of the K_S^0 flight distance. Candidates from all blocks are included and have been selected with the application of all offline cuts. A linear fit is performed. The slope is found to be 0.0007 ± 0.0009 , and the p-value is 0.09.

8.10.4 Systematics on the choice of the k parameter

The weighting procedure used to cancel nuisance asymmetries may be influenced by the choice of the number of nearest neighbors, k , as the calculated weights are subject to smearing due to the finite size of the multidimensional sphere defined by the k nearest neighbors. In regions where the functions vary rapidly, this smearing can introduce significant distortions. To evaluate the impact of k on the procedure, weights were computed for the $D^0 \rightarrow K_S^0 K_S^0$ sample using various values of k . Specifically, the values $k = 5, 10, 30$ (nominal), 50, 100, 200, and 500 were tested. The $\mathcal{A}^{CP}(K_S^0 K_S^0)$ values corresponding to each choice of k were extracted using a one-dimensional fit to the Δm distribution, integrating over all blocks. The larger statistical uncertainty affecting values obtained with $k = 5, 10$ is expected. Because of the lower k , a larger fraction of candidates fall in regions where the kNN output is 0 or 1 and are rejected, reducing the sample statistics. The difference observed between the value obtained with $k = 30$ and $k = 5, 10$ is compatible with a statistical fluctuation and is not considered in the systematic effect. For $k = 30$ and larger, the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ values show variations that are consistently smaller than 0.1%. This variation is taken as the systematic uncertainty associated with the choice of k .

k	$\mathcal{A}^{CP}(K_S^0 K_S^0)$ [%] (blind)
5	55.30 ± 1.22
10	54.58 ± 1.10
30	55.06 ± 0.99
50	55.00 ± 1.00
100	55.03 ± 0.99
200	55.03 ± 0.99
500	55.10 ± 0.99

Table 8.13: Blind $\mathcal{A}^{CP}(K_S^0 K_S^0)$ values computed for different values of k . Asymmetries have been extracted exploiting a 1D fit to the Δm distribution integrating all the blocks.

8.10.5 \mathcal{A}^{CP} in $D^0 \rightarrow K_S^0 \pi^+ \pi^-$

Events in the region of the $D^0 \rightarrow K^{*-}(892)\pi^+$ decay, which is a CF process, have been selected for the analysis. For these decays $\mathcal{A}^{CP}(K_S^0 \pi^+ \pi^-)$ is expected to be negligible on the scale of our resolution, and no systematic uncertainty is attributed to this assumption.

8.10.6 Fit model uncertainty

The largest systematic effect in the Run 2 measurement arose from a potential dependence of the measured CP asymmetry on the choice of the PDF parameterization used in the fit. The models employed to describe the signal and background are empirical, and alternative parameterizations also provide a good description of the data. To estimate the systematic uncertainty associated with model selection, the same approach as in Run 2 will be used. This involves generating pseudo-experiments based on an alternative model and comparing the \mathcal{A}^{CP} values obtained when fitting with the nominal and alternative models. For each model and each \mathcal{A}^{CP} value, approximately 200 pseudo-experiments will be generated. The systematic uncertainty will then be determined as the quadratic sum of the mean and standard deviation of the Gaussian that best describes the resulting distribution of differences.

It is worth noting that this systematic effect is expected to diminish with the larger statistics available in Run 3, meaning its impact will likely be smaller than in the Run 2 analysis.

8.10.7 Finite statistics of the calibration sample and other effects

The cancellation of production and detection asymmetries might vary because the weights are affected by statistical fluctuations of the control sample. The magnitude of this effect is being estimated by a *bootstrap* procedure on the original $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ dataset, generating for each event a random number from a Poisson distribution with mean value equal to 1, and accepting the event a number of times equal to the generated value. This will be repeated many times, re-evaluating the weights for the $D^0 \rightarrow K_S^0 K_S^0$ sample using the alternative generated calibration samples. The width of the obtained distribution is then taken as a systematic uncertainty. This computation, that is conceptually simple but highly CPU-time-consuming, together with a number of other smaller systematic effects, is still ongoing at the time of this writing. All these effects mirror corresponding

effects of the Run 2 analysis (see Section 4.7), so their results are quite predictable and are not anticipated to have a significant impact on the overall precision of the result. They are planned to be added as the final touch to the measurement in the very last stage of approval.

8.11 Conclusions

In this chapter I described the $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurement, performed analyzing a fraction of the data collected by LHCb in 2024, corresponding to an integrated luminosity of 6 fb^{-1} . With the sole exploitation of the novel K_S^0 -dedicated HLT1 line, it has been possible to collect statistics at the *analysis level* larger by a factor of 2.8 with respect to the one we had in Run 2, for equivalent integrated luminosity. The final (blind) result obtained analyzing this dataset is:

$$\mathcal{A}^{CP}(K_S^0 K_S^0) = (55.83 \pm 1.05)\%.$$

As discussed in previous section, the final value of the systematic uncertainty is still being evaluated and it is not quoted here, but the total uncertainty is anyway expected to be dominated by the statistical uncertainty, making this the most precise determination of this quantity to date.

This result represents the culmination of my PhD work, providing clear confirmation of the enhanced physics reach of LHCb made possible by the trigger strategy I developed and meticulously commissioned at the start of Run 3 data-taking. In addition, this is a significant milestone for the entire LHCb collaboration. The ability to collect such a substantial $D^0 \rightarrow K_S^0 K_S^0$ sample in just a few months highlights the success and potential of the upgraded system, which enables data collection at a rate five times higher than in Run 2, while achieving a signal efficiency improvement of over 2.5 times. This is particularly remarkable for a challenging channel like $D^0 \rightarrow K_S^0 K_S^0$, and foreshadows the future success of a wider set of measurements, that can benefit from—or are directly enabled by—the new dedicated K_S^0 triggers.

Chapter 9

Concluding remarks and future prospects

The studies presented in this thesis demonstrated that the implementation of a dedicated trigger strategy for K_S^0 at first level significantly enhances the collection of decay samples involving LLPs, thereby expanding the physics potential of LHCb. This work culminated in an actual physics measurement of the quantity that motivated it, coronating the whole process with its final confirmation.

The resolution of this measurement, obtained from just a few months of Run 3 data and not optimized for resolution, is already the most precise in existence. This bodes extremely well for the potential of the results that will be obtained from the full Run 3 data sample, that will be more than doubled, and with full optimization. It will then be fully established whether or not the $D^0 \rightarrow K_S^0 K_S^0$ mode really has the large CPV effect that several calculations predict.

In addition, as shown with specific examples in Chapter 6, our first-level K_S^0 trigger line is bound to produce precious samples leading to several other important physics measurements on that same timescale; and likely some that we have not even imagined.

However, the potential of this line of development does not end here. All that was achieved was based solely on triggering on the relatively minor fraction of K_S^0 decaying within the acceptance of the vertex detector. This is strictly due to the limitations of the current HLT1, that has not enough computing capacity to accommodate the heavier job of reconstructing tracks originating outside of the VELO acceptance. If this can be solved, there is a clear possibility for a further expansion of the LHCb potential in the LLP field in the future. I have explored also this other possibility within my simulation studies, quantitatively showing that a much more efficient first-level trigger based on displaced downstream tracks is definitely possible within the current allowed DAQ bandwidth (see Section 6.3.2 and Ref. [175]). The successful verification on data attained by my other rate estimates lends some credibility to this prediction. My study has been an important motivation in the formulation of a dedicated device [171] for real-time reconstruction of downstream tracks in the next Run 4, and is reported in the recently approved Technical Design Report [172] of the project. This perspective has further motivated an effort at inserting downstream-reconstruction capability even in the ongoing run [164], by adding further GPU processors in HLT1 and reoptimizing the tracking code to push its capabilities to their maximum for this specific purpose.

From a wider perspective, the `TwoTrackKs` and `TwoKs` lines are a first example of

30 MHz triggering on K_S^0 at the LHC. The excellent efficiency/rate ratio achieved was anything but trivial before shown studies, because of the huge production rate of light particles as K_S^0 in multi-TeV pp collisions. This, in principle, could have easily made rate and throughput cost of such selections unmanageable. Indeed, this is an encouragement to further pursue the idea of triggering on semi-exclusive, or even exclusive, signatures at the very first trigger level; signatures that in the past have been restricted to the higher levels of triggers.

This can be a particularly important path to pursue on a longer timescale, looking at higher and higher luminosities, as planned for LHCb Upgrade II and possibly even further. As we move towards larger luminosities, this kind of selections may become more and more important to the physics program. The more exclusive **TwoKs** and **TwoTrackLambda** lines, currently overshadowed by simpler lines, may well become crucial at higher luminosities, due to their very small bandwidth cost. It may turn out to be an effective strategy to exploit more lines of this kind, selecting other convenient intermediate states at first level (ϕ , ψ 's ...), or even some specific charm and beauty modes. The consequential large increase of computing load on the first-level trigger system might well be compensated by the rise of the new generation of custom real-time reconstruction devices like the DWT, providing ready-made reconstruction primitives to the HLT1 systems of the future, empowering them to perform much more refined selections than possible today.

Appendix A

Complete 2024 fit results

A.1 Unweighted fit results

The results obtained with the unweighted fit are summarized in Tables A.1 and A.2 for MagUp and MagDown blocks, respectively. The fit projection residuals on the Δm and the two $m(K_S^0)$ observables are shown in Figures A.1,A.2,A.3,A.4,A.5,A.6,A.7.




Figure A.1: Distributions of residuals for the block 1 sample computed from non-weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

Parameter	Block 1	Block 2	Block 3	Block 5
bin 1				
$ACP_{sig}^{CP} [\%]$ (blind)	52.68 ± 3.51	46.90 ± 4.80	53.38 ± 4.78	52.87 ± 3.43
$A_{bkg} [\%]$	0.65 ± 4.34	3.16 ± 5.72	-4.82 ± 5.67	-2.74 ± 3.80
$N_{tot}(sig)$	954 ± 36	508 ± 26	519 ± 25	1018 ± 38
$N_{tot}(bkg)$	666 ± 31	377 ± 23	387 ± 22	855 ± 35
$q(\Delta m)$	0.28 ± 0.07	0.19 ± 0.09	0.25 ± 0.09	0.27 ± 0.06
$\sigma(K_S^0) [\text{MeV}/c^2]$	7.9 ± 0.8	4.8 ± 0.4	6 ± 1	8 ± 1
$k_\sigma(K_S^0)$	0.47 ± 0.04	0.54 ± 0.07	0.63 ± 0.06	0.45 ± 0.04
$n(K_S^0)$	0.79 ± 0.06	0.29 ± 0.15	0.67 ± 0.2	0.8 ± 0.05
$m(K_S^0)$	-0.02 ± 0.23	0.04 ± 0.19	-0.04 ± 0.2	-0.1 ± 0.16
f_3^+	0.89 ± 0.04	0.75 ± 0.05	0.78 ± 0.05	0.84 ± 0.04
f_3^-	0.87 ± 0.04	0.71 ± 0.05	0.77 ± 0.05	0.78 ± 0.05
f_4^+	0.05 ± 0.02	0.11 ± 0.03	0.12 ± 0.03	0.07 ± 0.02
f_4^-	0.06 ± 0.02	0.13 ± 0.03	0.12 ± 0.03	0.11 ± 0.02
bin 2				
$ACP_{sig}^{CP} [\%]$ (blind)	51.56 ± 2.81	55.23 ± 4.08	53.30 ± 3.80	50.64 ± 2.71
$A_{bkg} [\%]$	0.71 ± 1.03	0.88 ± 1.41	0.94 ± 1.37	-0.61 ± 0.93
$N_{tot}(sig)$	2289 ± 82	1107 ± 56	1312 ± 54	2533 ± 93
$N_{tot}(bkg)$	11013 ± 124	5822 ± 89	6313 ± 89	13358 ± 140
$c(\Delta m)$	0.6 ± 0.2	0.5 ± 0.2	0.5 ± 0.2	0.5 ± 0.2
$b(\Delta m)$	0.03 ± 0.02	0.01 ± 0.02	0.02 ± 0.03	0.03 ± 0.03
$\sigma(K_S^0) [\text{MeV}/c^2]$	7.3 ± 0.4	6.2 ± 0.3	6.6 ± 0.3	7.1 ± 0.3
$k_\sigma(K_S^0)$	0.48 ± 0.02	0.49 ± 0.01	0.48 ± 0.01	0.48 ± 0.01
$n(K_S^0)$	0.67 ± 0.03	0.51 ± 0.06	0.56 ± 0.04	0.62 ± 0.03
$m(K_S^0)$	-0.18 ± 0.03	-0.22 ± 0.04	-0.17 ± 0.04	-0.16 ± 0.03
f_3^+	0.70 ± 0.01	0.68 ± 0.02	0.70 ± 0.01	0.71 ± 0.01
f_3^-	0.69 ± 0.01	0.69 ± 0.02	0.70 ± 0.02	0.71 ± 0.01
f_4^+	0.14 ± 0.01	0.144 ± 0.01	0.14 ± 0.01	0.13 ± 0.01
f_4^-	0.14 ± 0.01	0.14 ± 0.01	0.14 ± 0.01	0.13 ± 0.01
bin 1 and bin 2				
$\mu^+(\Delta m) [\text{MeV}/c^2]$	145.47 ± 0.06	145.32 ± 0.08	145.69 ± 0.19	145.39 ± 0.04
$\mu^-(\Delta m) [\text{MeV}/c^2]$	145.50 ± 0.06	145.29 ± 0.08	145.78 ± 0.22	145.40 ± 0.04
$\sigma^+(\Delta m) [\text{MeV}/c^2]$	0.59 ± 0.12	0.4 ± 0.1	1.14 ± 0.39	0.5 ± 0.1
$\sigma^-(\Delta m) [\text{MeV}/c^2]$	—	—	0.99 ± 0.34	—
$\alpha(\Delta m)$	0.16 ± 0.20	-0.49 ± 0.30	1.08 ± 0.75	-0.11 ± 0.16
$\nu(\Delta m)$	2.03 ± 0.42	1.7 ± 0.4	3.85 ± 1.28	1.9 ± 0.4
$\mu(K_S^0) [\text{MeV}/c^2]$	497.78 ± 0.03	497.72 ± 0.04	497.68 ± 0.04	497.84 ± 0.03

Table A.1: Results of the non-weighted simultaneous fit for blocks 1, 2, 3 and 5. Bin 1 and bin 2 refers to the definitions in Section 8.4.4, and they are fitted simultaneously. The parameters which describe the Δm peak and the mean value of the K_S^0 mass are shared between the two bins. The parameter $q(\Delta m)$ is used in the first order Chebyshev polynomial describing the Δm background for bin 1. The parameters n and k_σ describe the ratio between the two Gaussian distributions describing the K_S^0 peak and the ratio between their σ , respectively. The parameter m is instead used in the first order Chebyshev polynomial to describe the K_S^0 background. According to Eq. 8.4, background fractions sum to 1 and $f_7 = 1 - \sum_i f_i$.

Parameter	Block 4.1	Block 4.2	Block 6
bin 1			
$A_{sig}^{CP}[\%]$ (blind)	54.27 ± 6.84	59.73 ± 7.62	56.93 ± 3.78
$A_{bkg}[\%]$	-2.27 ± 5.03	12.27 ± 9.62	3.26 ± 4.63
$N_{tot}(sig)$	259 ± 19	200 ± 16	845 ± 36
$N_{tot}(bkg)$	204 ± 17	133 ± 14	600 ± 32
$q(\Delta m)$	0.02 ± 0.13	-0.06 ± 0.16	0.19 ± 0.08
$\sigma(K_S^0)$ [MeV/ c^2]	7 ± 1	7 ± 2	7 ± 1
$k_\sigma(K_S^0)$ [MeV/ c^2]	0.46 ± 0.05	0.6 ± 0.1	0.57 ± 0.06
$n(K_S^0)$	0.7 ± 0.1	0.7 ± 0.3	0.8 ± 0.1
$m(K_S^0)$	-0.4 ± 0.5	0.0 ± 0.4	-0.2 ± 0.2
f_3^+	0.83 ± 0.09	0.76 ± 0.09	0.73 ± 0.04
f_3^-	0.89 ± 0.07	0.85 ± 0.08	0.71 ± 0.05
f_4^+	0.08 ± 0.05	0.11 ± 0.05	0.13 ± 0.02
f_4^-	0.05 ± 0.04	0.12 ± 0.02	0.21 ± 0.04
bin 2			
$A_{sig}^{CP}[\%]$ (blind)	59.26 ± 5.03	60.40 ± 6.13	61.01 ± 2.94
$A_{bkg}[\%]$	-2.92 ± 2.01	-1.27 ± 2.29	-0.09 ± 1.05
$N_{tot}(sig)$	767 ± 45	531 ± 38	2178 ± 91
$N_{tot}(bkg)$	3072 ± 66	2347 ± 57	10637 ± 130
$c(\Delta m)$	0.24 ± 0.15	0.7 ± 0.3	0.9 ± 0.4
$b(\Delta m)$	-0.03 ± 0.04	0.03 ± 0.03	0.07 ± 0.02
$\sigma(K_S^0)$ [MeV/ c^2]	7.6 ± 0.6	9.2 ± 0.7	6.8 ± 0.3
$k_\sigma(K_S^0)$ [MeV/ c^2]	0.47 ± 0.03	0.39 ± 0.03	0.48 ± 0.01
$n(K_S^0)$	0.70 ± 0.05	0.71 ± 0.03	0.61 ± 0.04
$m(K_S^0)$	-0.17 ± 0.07	-0.34 ± 0.08	-0.17 ± 0.03
f_3^+	0.83 ± 0.09	0.73 ± 0.03	0.68 ± 0.01
f_3^-	0.75 ± 0.02	0.76 ± 0.03	0.70 ± 0.01
f_4^+	0.12 ± 0.01	0.13 ± 0.02	0.14 ± 0.01
f_4^-	0.11 ± 0.01	0.12 ± 0.02	0.14 ± 0.01
bin 1 and bin 2			
$\mu^+(\Delta m)$ [MeV/ c^2]	145.27 ± 0.14	145.2 ± 0.3	145.47 ± 0.03
$\mu^-(\Delta m)$ [MeV/ c^2]	145.33 ± 0.12	145.2 ± 0.3	145.48 ± 0.03
$\sigma^+(\Delta m)$ [MeV/ c^2]	0.63 ± 0.27	0.8 ± 0.6	0.37 ± 0.06
$\sigma^-(\Delta m)$ [MeV/ c^2]	0.52 ± 0.23	0.8 ± 0.6	0.37 ± 0.06
$\alpha(\Delta m)$	-0.4 ± 0.5	-0.7 ± 1.0	0.16 ± 0.11
$\nu(\Delta m)$	2.1 ± 0.9	2.7 ± 2.0	1.3 ± 0.2
$\mu(K_S^0)$ [MeV/ c^2]	497.78 ± 0.06	497.76 ± 0.07	497.70 ± 0.03

Table A.2: Results of the non-weighted simultaneous fit on blocks 4.1, 4.2 and 6. Bin 1 and bin 2 refers to the definitions in Section 8.4.4, and they are fitted simultaneously. The parameters which describe the Δm peak and the mean value of the K_S^0 mass are shared between the two bins. The parameter $q(\Delta m)$ is used in the first order Chebyshev polynomial describing the Δm background for bin 1. The parameters n and k_σ describe the ratio between the two Gaussian distributions describing the K_S^0 peak and the ratio between their σ , respectively. The parameter m is instead used in the first order Chebyshev polynomial to describe the K_S^0 background. According to Eq. 8.4, background fractions sum to 1 and $f_7 = 1 - \sum_i f_i$.

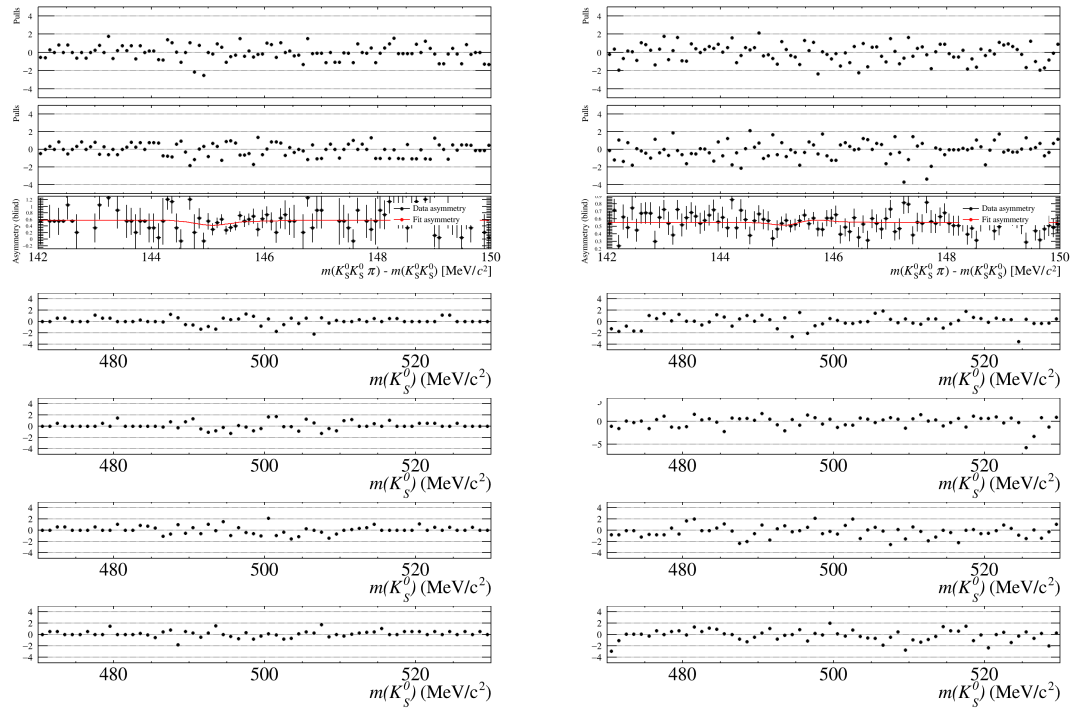


Figure A.2: Distributions of residuals for the block 2 sample computed from non-weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

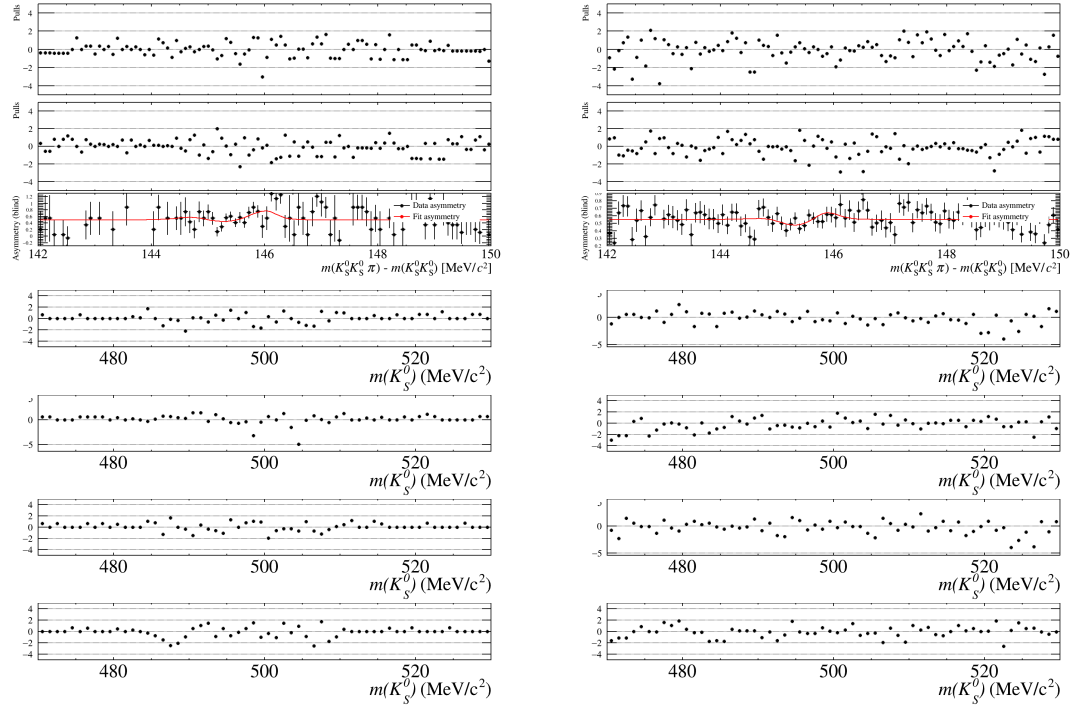


Figure A.3: Distributions of residuals for the block 3 sample computed from non-weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.




Figure A.4: Distributions of residuals for the block 5 sample computed from non-weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

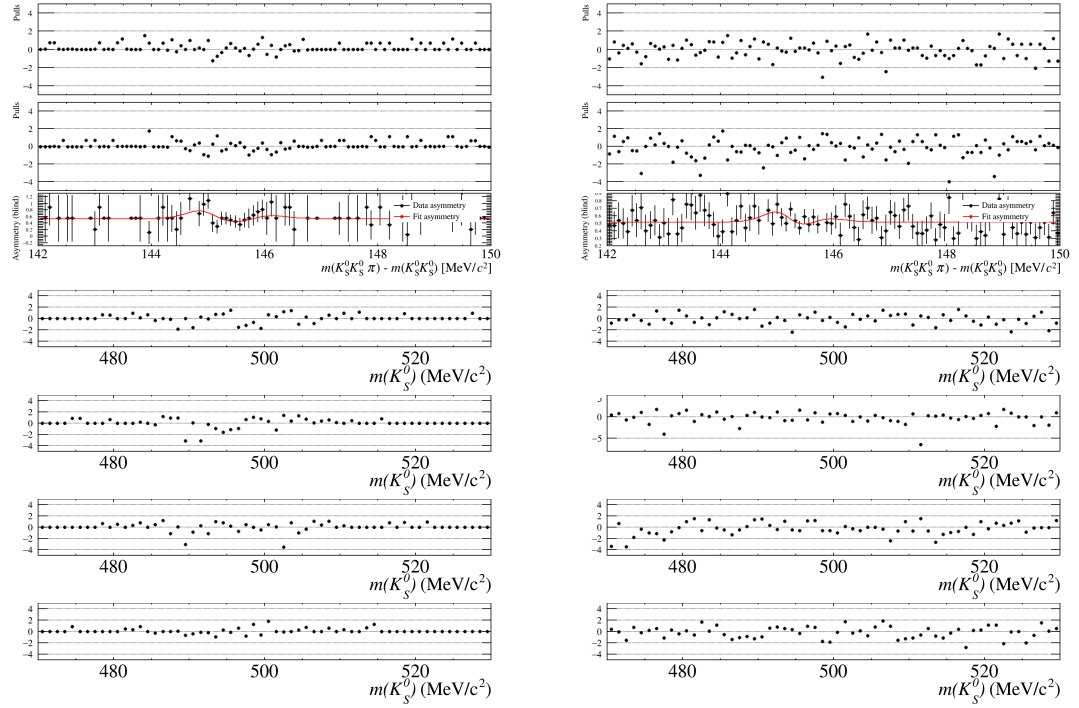


Figure A.5: Distributions of residuals for the block 4.1 sample computed from non-weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

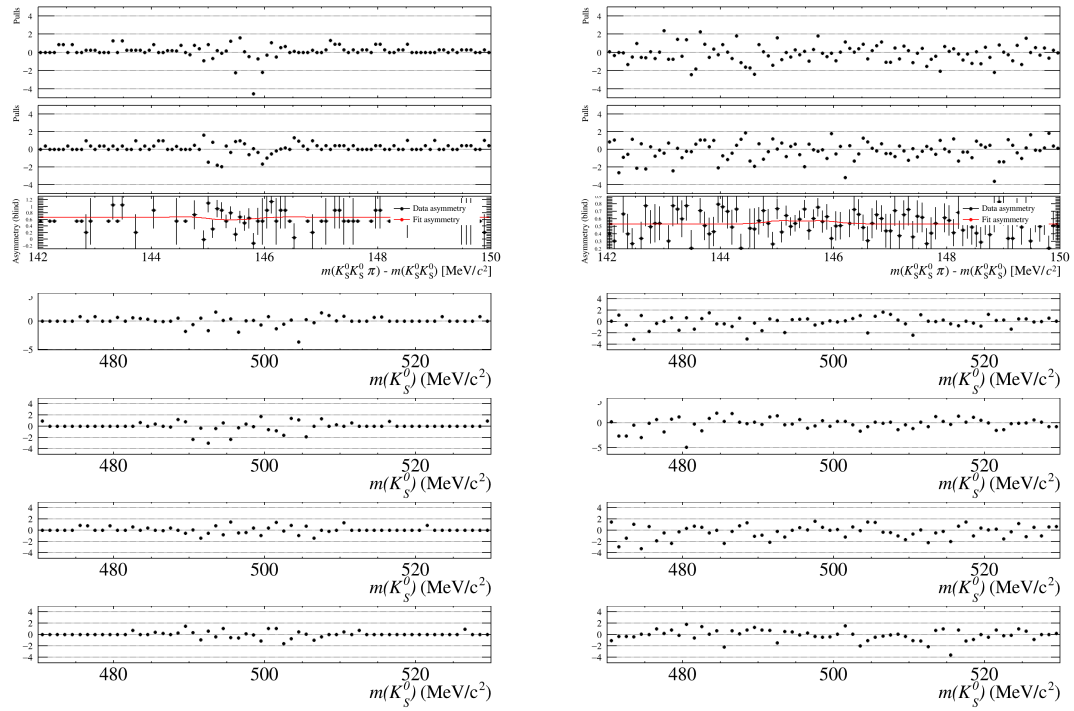


Figure A.6: Distributions of residuals for the block 4.2 sample computed from non-weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

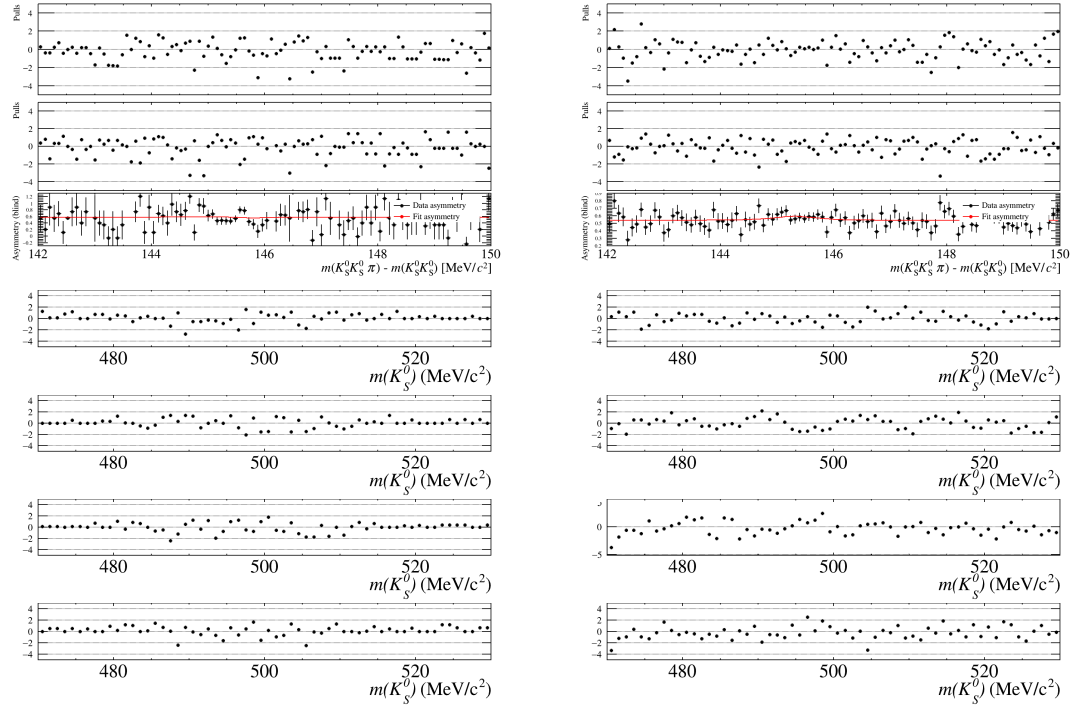


Figure A.7: Distributions of residuals for the block 6 sample computed from non-weighted fit. On the left, the distributions of D^{*+} decays are plotted, while on the right the distributions of the D^{*-} decays are plotted. Top: residuals for Δm distribution in the bin 1. Second line: residuals for Δm distribution in the bin 2. Third line: residuals for $m(K_{S1}^0)$ distribution in the bin 1. Fourth line: residuals for $m(K_{S1}^0)$ distribution in the bin 2. Fifth line: residuals for $m(K_{S2}^0)$ distribution in the bin 1. Bottom line: residuals for $m(K_{S2}^0)$ distribution in the bin 2.

A.2 Weighted fit results

The results obtained with the weighted fit are summarized in Tables A.3, A.4. The fit projection residuals on the Δm and the two $m(K_S^0)$ observables are shown in Figures A.8, A.9, A.10, A.12, A.13, A.11, A.14.

Parameter	Block 1	Block 2	Block 3	Block 5
bin 1				
$A_{sig}^{CP}[\%]$ (blind)	53.87 ± 3.62	49.46 ± 4.96	52.62 ± 4.95	56.89 ± 3.52
$A_{bkg}[\%]$	3.25 ± 4.47	4.07 ± 5.88	-3.34 ± 5.85	-1.48 ± 3.94
$N_{tot}(sig)$	1014 ± 39	536 ± 28	541 ± 27	1079 ± 41
$N_{tot}(bkg)$	705 ± 34	397 ± 25	407 ± 24	881 ± 38
$q(\Delta m)$	0.29 ± 0.07	0.2 ± 0.1	0.29 ± 0.09	0.31 ± 0.06
$\sigma(K_S^0)$ [MeV/ c^2]	7.7 ± 1.0	5.1 ± 0.9	5.9 ± 0.6	8.6 ± 1.2
$k_\sigma(K_S^0)$	0.475 ± 0.044	0.52 ± 0.08	0.64 ± 0.05	0.44 ± 0.05
$n(K_S^0)$	0.78 ± 0.08	0.4 ± 0.3	0.65 ± 0.15	0.83 ± 0.05
$m(K_S^0)$	-0.07 ± 0.23	0.16 ± 0.21	-0.09 ± 0.22	-0.03 ± 0.16
f_3^+	0.88 ± 0.04	0.77 ± 0.05	0.79 ± 0.05	0.85 ± 0.04
f_3^-	0.88 ± 0.04	0.72 ± 0.06	0.80 ± 0.05	0.78 ± 0.04
f_4^+	0.05 ± 0.02	0.10 ± 0.03	0.12 ± 0.03	0.07 ± 0.02
f_4^-	0.06 ± 0.02	0.13 ± 0.03	0.09 ± 0.03	0.12 ± 0.01
bin 2				
$A_{sig}^{CP}[\%]$ (blind)	52.17 ± 2.90	55.98 ± 4.18	55.79 ± 3.91	52.89 ± 2.77
$A_{bkg}[\%]$	1.54 ± 1.06	2.73 ± 1.45	1.98 ± 1.41	0.51 ± 0.96
$N_{tot}(sig)$	2399 ± 86	1167 ± 61	1392 ± 61	2717 ± 106
$N_{tot}(bkg)$	11610 ± 133	6091 ± 96	6631 ± 98	13971 ± 154
$c(\Delta m)$	0.59 ± 0.12	0.52 ± 0.2	0.51 ± 0.21	0.6 ± 0.2
$b(\Delta m)$	0.028 ± 0.016	0.01 ± 0.03	0.02 ± 0.03	0.05 ± 0.03
$\sigma(K_S^0)$ [MeV/ c^2]	7.4 ± 0.4	6.2 ± 0.3	6.7 ± 0.3	7.3 ± 0.3
$k_\sigma(K_S^0)$	0.47 ± 0.02	0.49 ± 0.02	0.48 ± 0.01	0.47 ± 0.01
$n(K_S^0)$	0.68 ± 0.04	0.51 ± 0.01	0.57 ± 0.04	0.64 ± 0.03
$m(K_S^0)$	-0.18 ± 0.03	-0.21 ± 0.04	-0.17 ± 0.04	-0.17 ± 0.03
f_3^+	0.70 ± 0.01	0.68 ± 0.02	0.70 ± 0.02	0.72 ± 0.01
f_3^-	0.69 ± 0.01	0.68 ± 0.02	0.70 ± 0.02	0.72 ± 0.01
f_4^+	0.138 ± 0.007	0.144 ± 0.008	0.138 ± 0.008	0.131 ± 0.006
f_4^-	0.137 ± 0.007	0.147 ± 0.008	0.138 ± 0.008	0.126 ± 0.006
bin 1 and bin 2				
$\mu^+(\Delta m)$ [MeV/ c^2]	145.51 ± 0.07	145.32 ± 0.07	145.80 ± 0.44	145.39 ± 0.04
$\mu^-(\Delta m)$ [MeV/ c^2]	145.48 ± 0.07	145.30 ± 0.07	145.71 ± 0.38	145.38 ± 0.04
$\sigma^+(\Delta m)$ [MeV/ c^2]	0.62 ± 0.16	0.43 ± 0.09	1.1 ± 0.5	0.5 ± 0.1
$\sigma^-(\Delta m)$ [MeV/ c^2]	—	—	0.9 ± 0.5	—
$\alpha(\Delta m)$	0.2 ± 0.2	-0.5 ± 0.3	1.2 ± 1.5	-0.13 ± 0.15
$\nu(\Delta m)$	2.1 ± 0.6	1.6 ± 0.4	3.7 ± 2.0	1.7 ± 0.3
$\mu(K_S^0)$ [MeV/ c^2]	497.78 ± 0.03	497.72 ± 0.042	497.67 ± 0.04	497.83 ± 0.03

Table A.3: Results of the weighted simultaneous fit on blocks 1, 2, 3 and 5 (MagUp). Bin 1 and bin 2 refers to the definitions in Section 8.4.4, and they are fitted simultaneously. The parameters which describe the Δm peak and the mean value of the K_S^0 mass are shared between the two bins. The parameter $q(\Delta m)$ is used in the first order Chebyshev polynomial describing the Δm background for bin 1. The parameters n and k_σ describe the ratio between the two Gaussian distributions describing the K_S^0 peak and the ratio between their σ , respectively. The parameter m is instead used in the first order Chebyshev polynomial to describe the K_S^0 background. According to Eq. 8.4, background fractions sum to 1 and $f_7 = 1 - \sum_i f_i$.

Parameter	Block 4.1	Block 4.2	Block 6
bin 1			
$A_{sig}^{CP}[\%]$ (blind)	56.09 ± 7.03	63.23 ± 7.61	57.03 ± 3.73
$A_{bkg}[\%]$	2.09 ± 8.11	14.62 ± 9.25	5.99 ± 4.46
$N_{tot}(sig)$	270 ± 20	204 ± 17	866 ± 36
$N_{tot}(bkg)$	211 ± 18	146 ± 15	639 ± 33
$q(\Delta m)$	0.03 ± 0.13	-0.01 ± 0.16	0.17 ± 0.07
$\sigma(K_S^0)$ [MeV/ c^2]	7.2 ± 1.4	6.7 ± 1.6	6.8 ± 0.9
$k_\sigma(K_S^0)$ [MeV/ c^2]	0.46 ± 0.06	0.59 ± 0.09	0.57 ± 0.06
$n(K_S^0)$	0.67 ± 0.13	0.7 ± 0.3	0.8 ± 0.1
$m(K_S^0)$	-0.43 ± 0.56	-0.01 ± 0.39	-0.2 ± 0.2
f_3^+	0.84 ± 0.09	0.77 ± 0.09	0.74 ± 0.04
f_3^-	0.90 ± 0.07	0.87 ± 0.07	0.71 ± 0.05
f_4^+	0.09 ± 0.06	0.11 ± 0.05	0.12 ± 0.02
f_4^-	0.05 ± 0.04	0.09 ± 0.06	0.21 ± 0.04
bin 2			
$A_{sig}^{CP}[\%]$ (blind)	61.44 ± 5.22	57.35 ± 5.94	62.40 ± 2.87
$A_{bkg}[\%]$	-1.33 ± 2.07	0.05 ± 2.23	1.54 ± 1.03
$N_{tot}(sig)$	797 ± 46	568 ± 42	2266 ± 93
$N_{tot}(bkg)$	3248 ± 70	2470 ± 61	11185 ± 132
$c(\Delta m)$	0.23 ± 0.11	0.6 ± 0.3	0.9 ± 0.4
$b(\Delta m)$	-0.03 ± 0.02	0.03 ± 0.04	0.07 ± 0.02
$\sigma(K_S^0)$ [MeV/ c^2]	7.7 ± 0.8	9.5 ± 0.8	6.9 ± 0.3
$k_\sigma(K_S^0)$ [MeV/ c^2]	0.47 ± 0.03	0.39 ± 0.03	0.48 ± 0.01
$n(K_S^0)$	0.71 ± 0.06	0.73 ± 0.03	0.61 ± 0.04
$m(K_S^0)$	-0.17 ± 0.07	-0.35 ± 0.08	-0.15 ± 0.03
f_3^+	0.73 ± 0.03	0.74 ± 0.03	0.68 ± 0.01
f_3^-	0.75 ± 0.02	0.76 ± 0.03	0.70 ± 0.01
f_4^+	0.12 ± 0.01	0.12 ± 0.01	0.14 ± 0.01
f_4^-	0.11 ± 0.01	0.12 ± 0.02	0.13 ± 0.01
bin 1 and bin 2			
$\mu^+(\Delta m)$ [MeV/ c^2]	145.32 ± 0.11	145.16 ± 0.45	145.48 ± 0.03
$\mu^-(\Delta m)$ [MeV/ c^2]	145.28 ± 0.13	145.15 ± 0.45	145.48 ± 0.03
$\sigma^+(\Delta m)$ [MeV/ c^2]	0.63 ± 0.28	0.83 ± 0.82	0.39 ± 0.06
$\sigma^-(\Delta m)$ [MeV/ c^2]	0.54 ± 0.26	0.81 ± 0.78	0.39 ± 0.06
$\alpha(\Delta m)$	-0.36 ± 0.46	-0.92 ± 1.61	0.2 ± 0.1
$\nu(\Delta m)$	2 ± 1	3 ± 3	2.6 ± 0.7
$\mu(K_S^0)$ [MeV/ c^2]	496.77 ± 0.06	497.77 ± 0.07	497.70 ± 0.03

Table A.4: Results of the weighted simultaneous fit on blocks 4.1, 4.2 and 6 (MagDown). Bin 1 and bin 2 refers to the definitions in Section 8.4.4, and they are fitted simultaneously. The parameters which describe the Δm peak and the mean value of the K_S^0 mass are shared between the two bins. The parameter $q(\Delta m)$ is used in the first order Chebyshev polynomial describing the Δm background for bin 1. The parameters n and k_σ describe the ratio between the two Gaussian distributions describing the K_S^0 peak and the ratio between their σ , respectively. The parameter m is instead used in the first order Chebyshev polynomial to describe the K_S^0 background. According to Eq. 8.4, background fractions sum to 1 and $f_7 = 1 - \sum_i f_i$.




Figure A.8: Distributions of residuals for the block 1 sample computed from weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

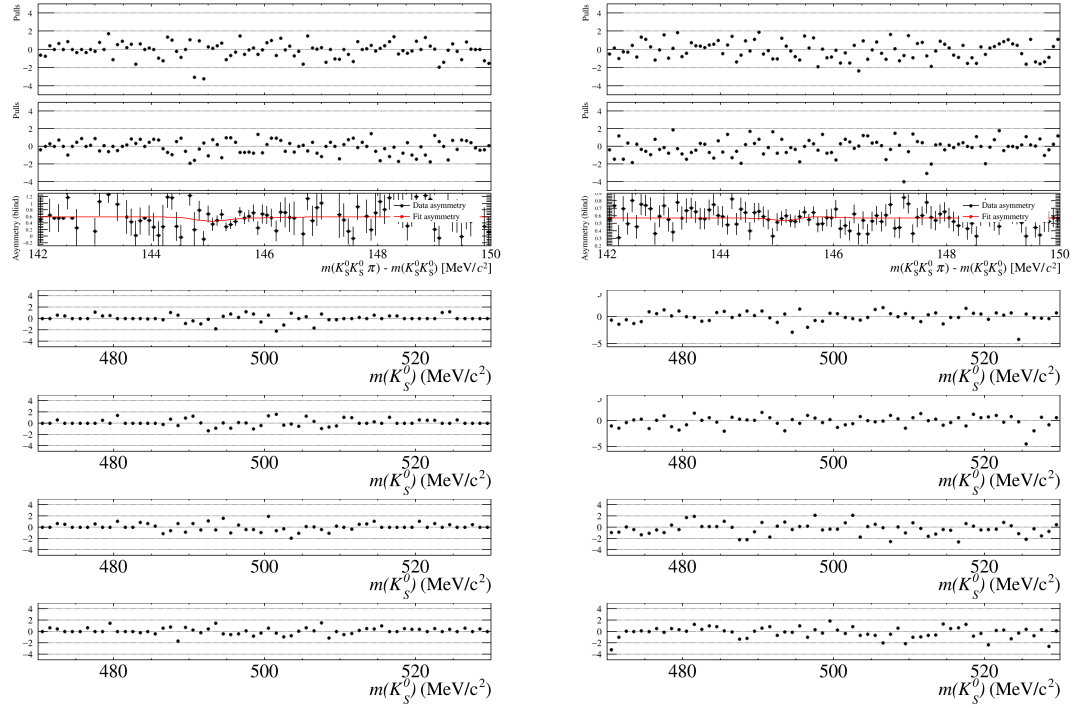


Figure A.9: Distributions of residuals for the block 2 sample computed from weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.




Figure A.10: Distributions of residuals for the block 3 sample computed from weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.




Figure A.11: Distributions of residuals for the block 5 sample computed from weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

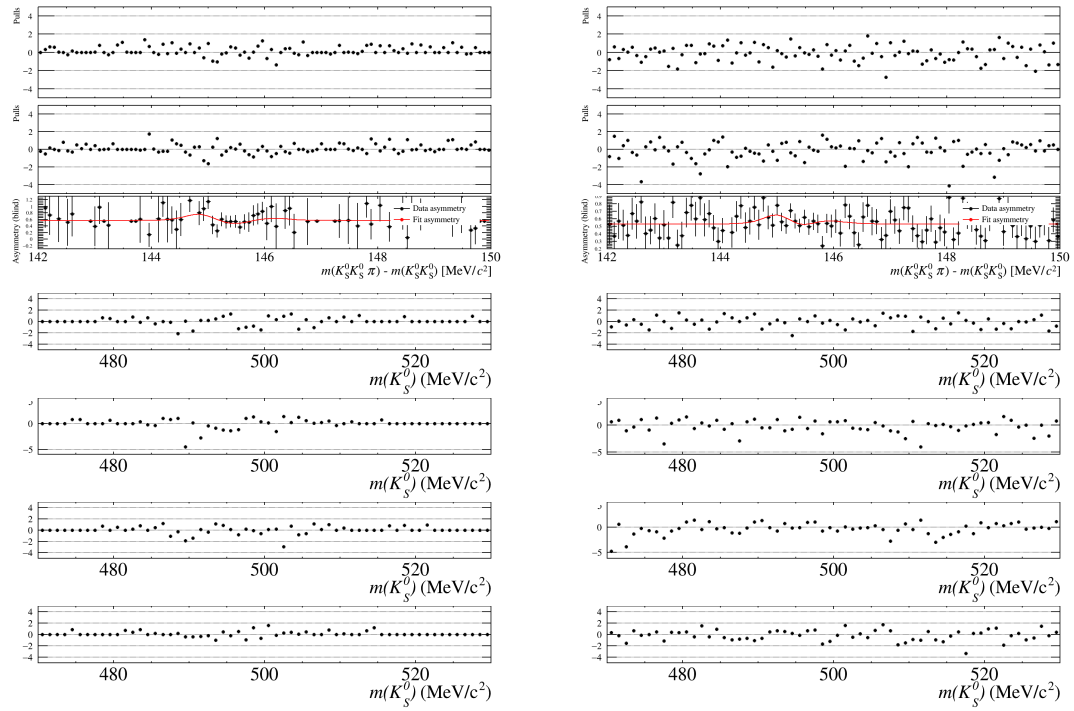


Figure A.12: Distributions of residuals for the block 4.1 sample computed from weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

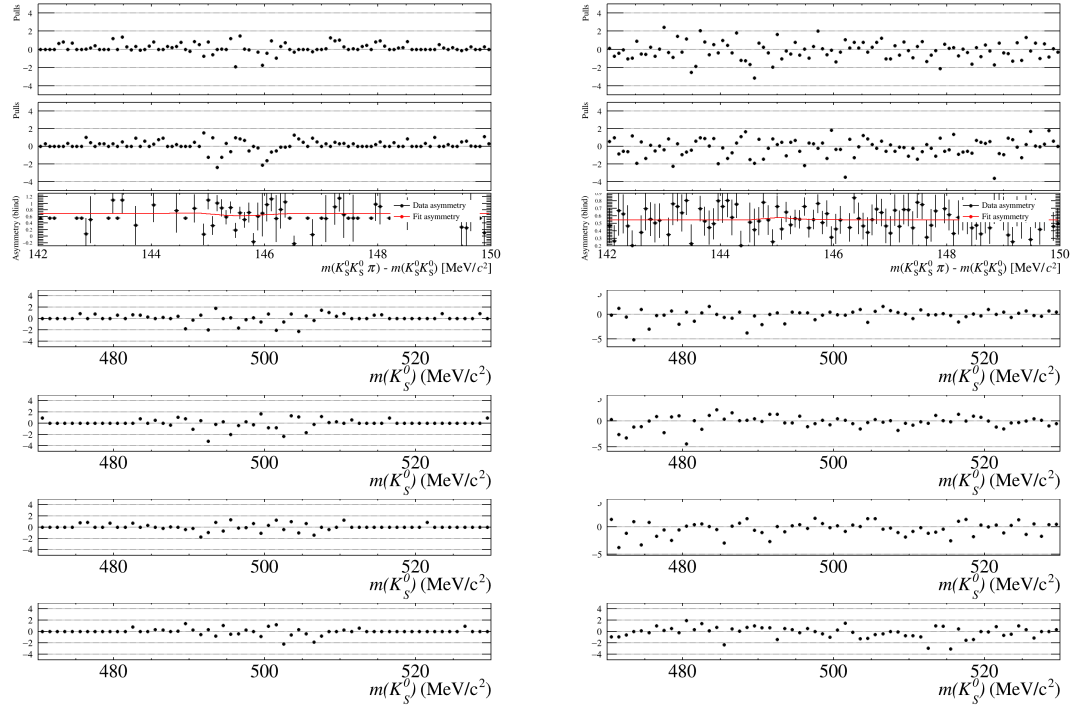


Figure A.13: Distributions of residuals for the block 4.2 sample computed from weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

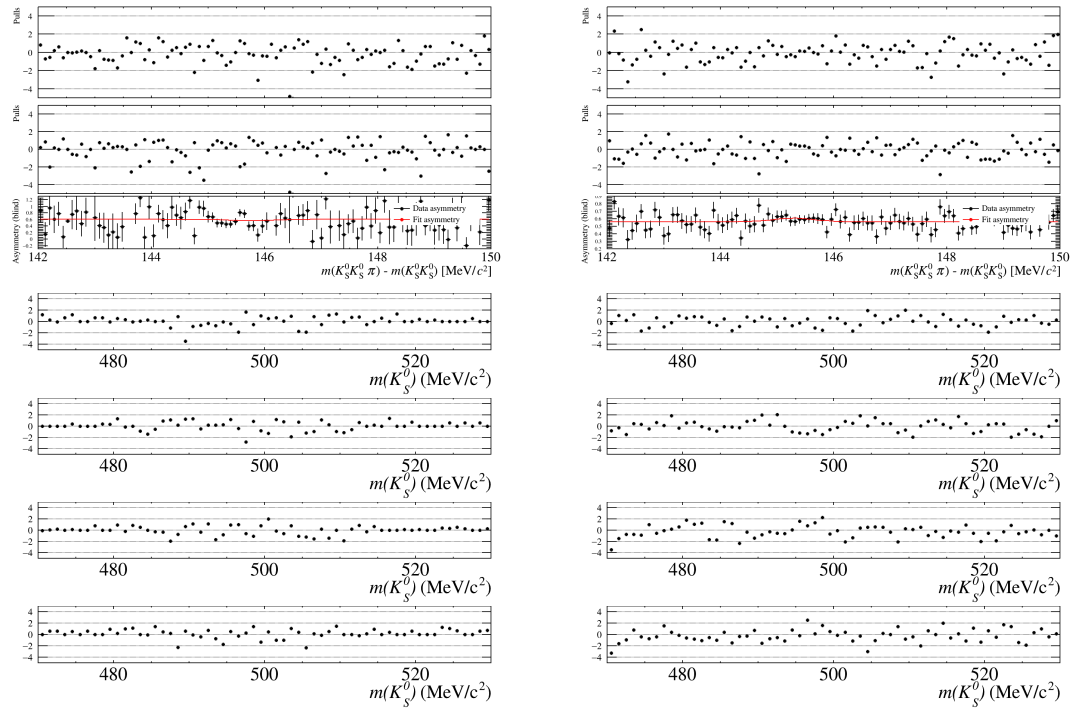


Figure A.14: Distributions of residuals for the block 6 sample computed from weighted fit. On the left, the distributions for bin 1 are plotted, while on the right the distributions for the bin 2 are plotted. Top line: residuals for D^{*+} , D^{*-} and asymmetry as a function of Δm (top to bottom). Second line: residuals for $m(K_{S1}^0)$ distribution for D^{*+} decays. Third line: residuals for $m(K_{S1}^0)$ distribution for D^{*-} decays. Fourth line: residuals for $m(K_{S2}^0)$ distribution for D^{*+} decays. Bottom line: residuals for $m(K_{S2}^0)$ distribution for D^{*-} decays.

Bibliography

- [1] LHCb collaboration, R. Aaij *et al.*, *Observation of CP violation in charm decays*, Phys. Rev. Lett. **122** (2019) 211803, [arXiv:1903.08726](#).
- [2] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP asymmetry in $D^0 \rightarrow K_S^0 K_S^0$ decays*, Phys. Rev. **D104** (2021) L031102, [arXiv:2105.01565](#).
- [3] CMS collaboration, A. Hayrapetyan *et al.*, *Search for CP violation in $D^0 \rightarrow K_S^0 K_S^0$ decays in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **84** (2024) 1264, [arXiv:2405.11606](#).
- [4] Belle-II, Belle collaboration, I. Adachi *et al.*, *Measurement of the time-integrated CP asymmetry in $D^0 \rightarrow K_S^0 K_S^0$ decays using Belle and Belle II data*, [arXiv:2411.00306](#).
- [5] AMS collaboration, M. Aguilar *et al.*, *The Alpha Magnetic Spectrometer (AMS) on the international space station: Part II — Results from the first seven years*, Phys. Rept. **894** (2021) 1.
- [6] C. S. Wu *et al.*, *Experimental Test of Parity Conservation in β Decay*, Phys. Rev. **105** (1957) 1413.
- [7] L. D. Landau, *On the conservation laws for weak interactions*, Nucl. Phys. **3** (1957) 127.
- [8] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, *Evidence for the 2π Decay of the K_2^0 Meson*, Phys. Rev. Lett. **13** (1964) 138.
- [9] A. D. Sakharov, *Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe*, Pisma Zh. Eksp. Teor. Fiz. **5** (1967) 32.
- [10] N. Cabibbo, *Unitary symmetry and leptonic decays*, Phys. Rev. Lett. **10** (1963) 531.
- [11] M. Kobayashi and T. Maskawa, *CP-violation in the renormalizable theory of weak interaction*, Prog. Theor. Phys. **49** (1973) 652.
- [12] L. Wolfenstein, *Parametrization of the Kobayashi-Maskawa Matrix*, Phys. Rev. Lett. **51** (1983) 1945.
- [13] Particle Data Group, P. A. Zyla *et al.*, *Review of Particle Physics*, PTEP **2020** (2020) 083C01.
- [14] S. Luo and Z.-z. Xing, *Radiative Corrections to the CKM Unitarity Triangles*, J. Phys. G **37** (2010) 075018, [arXiv:0912.4593](#).

- [15] C. Jarlskog, *Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Nonconservation*, Phys. Rev. Lett. **55** (1985) 1039.
- [16] Belle collaboration, K. Abe *et al.*, *Observation of large CP violation in the neutral B meson system*, Phys. Rev. Lett. **87** (2001) 091802, [arXiv:hep-ex/0107061](#).
- [17] BaBar collaboration, B. Aubert *et al.*, *Observation of CP violation in the B^0 meson system*, Phys. Rev. Lett. **87** (2001) 091801, [arXiv:hep-ex/0107013](#).
- [18] F. Buccella *et al.*, *Nonleptonic weak decays of charmed mesons*, Phys. Rev. D **51** (1995) 3478, [arXiv:hep-ph/9411286](#).
- [19] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP asymmetry in $D^0 \rightarrow K^- K^+$ and $D^0 \rightarrow \pi^- \pi^+$ decays*, JHEP **07** (2014) 041, [arXiv:1405.2797](#).
- [20] LHCb collaboration, R. Aaij *et al.*, *Measurement of the difference of time-integrated CP asymmetries in $D^0 \rightarrow K^- K^+$ and $D^0 \rightarrow \pi^- \pi^+$ decays*, Phys. Rev. Lett. **116** (2016) 191601, [arXiv:1602.03160](#).
- [21] HFLAV collaboration, Y. S. Amhis *et al.*, *Averages of b-hadron, c-hadron, and τ -lepton properties as of 2018*, Eur. Phys. J. C **81** (2021) 226, [arXiv:1909.12524](#).
- [22] LHCb collaboration, R. Aaij *et al.*, *Measurement of the CP violation parameter A_F in $D^0 \rightarrow K^+ K^-$ and $D^0 \rightarrow \pi^+ \pi^-$ decays*, Phys. Rev. Lett. **118** (2017) 261803, [arXiv:1702.06490](#).
- [23] LHCb collaboration, R. Aaij *et al.*, *Search for time-dependent CP violation in $D^0 \rightarrow K^+ K^-$ and $D^0 \rightarrow \pi^+ \pi^-$ decays*, Phys. Rev. **D104** (2021) 072010, [arXiv:2105.09889](#).
- [24] LHCb collaboration, R. Aaij *et al.*, *Updated measurement of decay-time-dependent CP asymmetries in $D^0 \rightarrow K^+ K^-$ and $D^0 \rightarrow \pi^+ \pi^-$ decays*, Phys. Rev. **D101** (2020) 012005, [arXiv:1911.01114](#).
- [25] A. L. Kagan and M. D. Sokoloff, *On Indirect CP Violation and Implications for D^0 - anti- D^0 and $B(s)$ - anti- $B(s)$ mixing*, Phys. Rev. D **80** (2009) 076008, [arXiv:0907.3917](#).
- [26] Y. Grossman, A. L. Kagan, and Y. Nir, *New physics and CP violation in singly Cabibbo suppressed D decays*, Phys. Rev. D **75** (2007) 036008, [arXiv:hep-ph/0609178](#).
- [27] A. Khodjamirian and A. A. Petrov, *Direct CP asymmetry in $D \rightarrow \pi^- \pi^+$ and $D \rightarrow K^- K^+$ in QCD-based approach*, Phys. Lett. B **774** (2017) 235, [arXiv:1706.07780](#).
- [28] M. Chala, A. Lenz, A. V. Rusov, and J. Scholtz, *ΔA_{CP} within the Standard Model and beyond*, JHEP **07** (2019) 161, [arXiv:1903.10490](#).
- [29] Y. Grossman and S. Schacht, *The emergence of the $\Delta U = 0$ rule in charm physics*, JHEP **07** (2019) 020, [arXiv:1903.10952](#).

-
- [30] H.-Y. Cheng and C.-W. Chiang, *Revisiting CP violation in $D \rightarrow PP$ and VP decays*, Phys. Rev. D **100** (2019) 093002, [arXiv:1909.03063](#).
 - [31] H.-N. Li, C.-D. Lü, and F.-S. Yu, *Implications on the first observation of charm CPV at LHCb*, [arXiv:1903.10638](#).
 - [32] U. Nierste, *Charm decays*, PoS **Beauty2019** (2020) 048, [arXiv:2002.06686](#).
 - [33] S. Schacht and A. Soni, *Enhancement of charm CP violation due to nearby resonances*, Phys. Lett. B **825** (2022) 136855, [arXiv:2110.07619](#).
 - [34] I. Bediaga, T. Frederico, and P. C. Magalhães, *Enhanced Charm CP Asymmetries from Final State Interactions*, Phys. Rev. Lett. **131** (2023) 051802, [arXiv:2203.04056](#).
 - [35] H.-Y. Cheng and C.-W. Chiang, *Direct CP violation in two-body hadronic charmed meson decays*, Phys. Rev. D **85** (2012) 034036, [arXiv:1201.0785](#), [Erratum: Phys.Rev.D 85, 079903 (2012)].
 - [36] J. Brod, A. L. Kagan, and J. Zupan, *Size of direct CP violation in singly Cabibbo-suppressed D decays*, Phys. Rev. D **86** (2012) 014023, [arXiv:1111.5000](#).
 - [37] U. Nierste and S. Schacht, *CP Violation in $D^0 \rightarrow K_S K_S$* , Phys. Rev. D **92** (2015) 054036, [arXiv:1508.00074](#).
 - [38] U. Nierste and S. Schacht, *CP asymmetries in D decays to two pseudoscalars*, PoS **ICHEP2016** (2016) 527, [arXiv:1704.04523](#).
 - [39] S. Müller, U. Nierste, and S. Schacht, *Topological amplitudes in D decays to two pseudoscalars: A global analysis with linear $SU(3)_F$ breaking*, Phys. Rev. D **92** (2015) 014004, [arXiv:1503.06759](#).
 - [40] LHCb collaboration, R. Aaij *et al.*, *Measurement of the time-integrated CP asymmetry in $D^0 \rightarrow K_S^0 K_S^0$ decays*, JHEP **11** (2018) 048, [arXiv:1806.01642](#).
 - [41] H.-n. Li, C.-D. Lu, and F.-S. Yu, *Branching ratios and direct CP asymmetries in $D \rightarrow PP$ decays*, Phys. Rev. D **86** (2012) 036012, [arXiv:1203.3120](#).
 - [42] F. Buccella, A. Paul, and P. Santorelli, *$SU(3)_F$ breaking through final state interactions and CP asymmetries in $D \rightarrow PP$ decays*, Phys. Rev. D **99** (2019) 113001, [arXiv:1902.05564](#).
 - [43] CLEO collaboration, G. Bonvicini *et al.*, *Search for CP violation in $D^0 \rightarrow K_S^0 \pi^0$ and $D^0 \rightarrow \pi^0 \pi^0$ and $D^0 \rightarrow K_S^0 K_S^0$ decays*, Phys. Rev. D **63** (2001) 071101, [arXiv:hep-ex/0012054](#).
 - [44] LHCb collaboration, R. Aaij *et al.*, *Measurement of the time-integrated CP asymmetry in $D^0 \rightarrow K_S^0 K_S^0$ decays*, JHEP **10** (2015) 055, [arXiv:1508.06087](#).
 - [45] Belle collaboration, A. Abashian *et al.*, *The Belle Detector*, Nucl. Instrum. Meth. A **479** (2002) 117.

- [46] LHCb collaboration, R. Aaij *et al.*, *Measurements of prompt charm production cross-sections in pp collisions at $\sqrt{s}=13$ TeV*, JHEP **03** (2016) 159, Erratum *ibid.* **09** (2016) 013, Erratum *ibid.* **05** (2017) 074, [arXiv:1510.01707](#).
- [47] Belle collaboration, R. Seuster *et al.*, *Charm hadrons from fragmentation and B decays in $e^+ e^-$ annihilation at $s^{*(1/2)} = 10.6$ -GeV*, Phys. Rev. D **73** (2006) 032002, [arXiv:hep-ex/0506068](#).
- [48] ALEPH collaboration, R. Barate *et al.*, *Search for pair production of longlived heavy charged particles in $e^+ e^-$ annihilation*, Phys. Lett. B **405** (1997) 379, [arXiv:hep-ex/9706013](#).
- [49] DELPHI collaboration, J. Abdallah *et al.*, *Final results from DELPHI on the searches for SM and MSSM neutral Higgs bosons*, Eur. Phys. J. C **32** (2004) 145, [arXiv:hep-ex/0303013](#).
- [50] J. P. Chou, D. Curtin, and H. J. Lubatti, *New Detectors to Explore the Lifetime Frontier*, Phys. Lett. B **767** (2017) 29, [arXiv:1606.06298](#).
- [51] FASER collaboration, A. Ariga *et al.*, *FASER: ForwArd Search ExpeRiment at the LHC*, [arXiv:1901.04468](#).
- [52] A. Blondel *et al.*, *Searches for long-lived particles at the future FCC-ee*, Front. in Phys. **10** (2022) 967881, [arXiv:2203.05502](#).
- [53] J. Alimena *et al.*, *Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider*, J. Phys. G **47** (2020) 090501, [arXiv:1903.04497](#).
- [54] M. Ciuchini, M. Pierini, and L. Silvestrini, *$B(s)$ to $K(^*)0$ anti- $K(^*)0$ decays: The Golden channels for new physics searches*, Phys. Rev. Lett. **100** (2008) 031802, [arXiv:hep-ph/0703137](#).
- [55] S. Descotes-Genon, J. Matias, and J. Virto, *An analysis of $B_{d,s}$ mixing angles in presence of New Physics and an update of $B_s \rightarrow \bar{K}^{0*} \text{anti} - K^{0*}$* , Phys. Rev. D **85** (2012) 034010, [arXiv:1111.4882](#).
- [56] Y. Amhis, Y. Grossman, and Y. Nir, *The branching fraction of $B_s^0 \rightarrow K^0 \bar{K}^0$: three puzzles*, JHEP **02** (2023) 113, [arXiv:2212.03874](#).
- [57] LHCb collaboration, R. Aaij *et al.*, *Measurement of the branching fraction of the decay $B_s^0 \rightarrow K_S^0 K_S^0$* , Phys. Rev. D **102** (2020) 012011, [arXiv:2002.08229](#).
- [58] Belle collaboration, B. Pal *et al.*, *Observation of the decay $B_s^0 \rightarrow K^0 \bar{K}^0$* , Phys. Rev. Lett. **116** (2016) 161801, [arXiv:1512.02145](#).
- [59] A. B. Carter and A. I. Sanda, *CP Violation in B Meson Decays*, Phys. Rev. D **23** (1981) 1567.
- [60] I. I. Y. Bigi and A. I. Sanda, *Notes on the Observability of CP Violations in B Decays*, Nucl. Phys. B **193** (1981) 85.

-
- [61] BaBar collaboration, B. Aubert *et al.*, *Measurement of Time-Dependent CP Asymmetry in $B^0 \rightarrow c \text{ anti-}c K^{(*)}0$ Decays*, Phys. Rev. D **79** (2009) 072009, [arXiv:0902.1708](#).
 - [62] Belle collaboration, I. Adachi *et al.*, *Precise measurement of the CP violation parameter $\sin 2\phi_1$ in $B^0 \rightarrow (c\bar{c})K^0$ decays*, Phys. Rev. Lett. **108** (2012) 171802, [arXiv:1201.4643](#).
 - [63] Belle collaboration, H. Sahoo *et al.*, *Measurements of time-dependent CP violation in $B^0 \rightarrow \psi(2S)K_S^0$ decays*, Phys. Rev. D **77** (2008) 091103, [arXiv:0708.2604](#).
 - [64] Belle-II collaboration, I. Adachi *et al.*, *Measurement of decay-time-dependent CP violation in $B^0 \rightarrow J/\psi K_S^0$ decays using 2019-2021 Belle II data*, [arXiv:2302.12898](#).
 - [65] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP violation in $B^0 \rightarrow J/\psi K_S^0$ decays*, Phys. Rev. Lett. **115** (2015) 031601, [arXiv:1503.07089](#).
 - [66] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP violation in $B^0 \rightarrow J/\psi K_S^0$ and $B^0 \rightarrow \psi(2S)K_S^0$ decays*, JHEP **11** (2017) 170, [arXiv:1709.03944](#).
 - [67] LHCb collaboration, R. Aaij *et al.*, *TBD*, LHCb-PAPER-2023-013, in preparation.
 - [68] G. D'Ambrosio and T. Kitahara, *Direct CP Violation in $K \rightarrow \mu^+\mu^-$* , Phys. Rev. Lett. **119** (2017) 201802, [arXiv:1707.06999](#).
 - [69] LHCb collaboration, R. Aaij *et al.*, *Search for the rare decay $K_S^0 \rightarrow \mu^+\mu^-$* , JHEP **01** (2013) 090, [arXiv:1209.4029](#).
 - [70] LHCb collaboration, R. Aaij *et al.*, *Improved limit on the branching fraction of the rare decay $K_S^0 \rightarrow \mu\mu$* , Eur. Phys. J. **C77** (2017) 678, [arXiv:1706.00758](#).
 - [71] LHCb collaboration, R. Aaij *et al.*, *Constraints on the $K_S^0 \rightarrow \mu^+\mu^-$ branching fraction*, Phys. Rev. Lett. **125** (2020) 1231801, [arXiv:2001.10354](#).
 - [72] M. Claudson, S. L. Glashow, and M. B. Wise, *Isospin violation in $\frac{J}{\psi} \rightarrow \text{baryon} + \text{antibaryon}$* , Phys. Rev. D **25** (1982) 1345.
 - [73] C. Carimalo, *Quark Mass Effects in $\psi \rightarrow B\bar{B}$ Decays*, Int. J. Mod. Phys. A **2** (1987) 249.
 - [74] F. Murgia and M. Melis, *Mass corrections in $j/\psi \rightarrow b\bar{b}$ decay and the role of distribution amplitudes*, Phys. Rev. D **51** (1995) 3487.
 - [75] M. W. Eaton *et al.*, *Decays of the $\psi(3097)$ to baryon-antibaryon final states*, Phys. Rev. D **29** (1984) 804.
 - [76] D. Pallin *et al.*, *Baryon pair production in j/ψ decays*, Nuclear Physics B **292** (1987) 653.
 - [77] J. Z. Bai *et al.*, *Decays of the j/ψ to $\lambda\lambda, \lambda\lambda\gamma$ and $\lambda\lambda\pi^0$ final states.*, Physics Letters B **424** (1998) 213.

- [78] BES collaboration, J. Z. Bai *et al.*, *Measurement of $\psi(2S)$ decays to baryon pairs*, Phys. Rev. D **63** (2001) 032002, [arXiv:hep-ex/0008031](#).
- [79] BES collaboration, M. Ablikim *et al.*, *Study of J/ψ decays to Λ anti- Λ and Σ^0 anti- Σ^0* , Phys. Lett. B **632** (2006) 181, [arXiv:hep-ex/0506020](#).
- [80] M. Ablikim *et al.*, *Measurements of $\psi(2s)$ decays to octet baryon–antibaryon pairs*, Physics Letters B **648** (2007) 149.
- [81] BaBar collaboration, B. Aubert *et al.*, *Study of $e^+e^- \rightarrow \Lambda\bar{\Lambda}$, $\Lambda\bar{\Sigma}^0$, $\Sigma^0\bar{\Sigma}^0$ using initial state radiation with BABAR*, Phys. Rev. D **76** (2007) 092006, [arXiv:0709.1988](#).
- [82] Particle Data Group, R. L. Workman *et al.*, *Review of particle physics*, Prog. Theor. Exp. Phys. **2022** (2022) 083C01.
- [83] H.-Y. Cheng, *Charmed baryons circa 2015*, Front. Phys. (Beijing) **10** (2015) 101406.
- [84] L. L. Chau and H. Y. Cheng, *Quark Diagram Analysis of Two-body Charm Decays*, Phys. Rev. Lett. **56** (1986) 1655.
- [85] Y. Kohara, *Quark diagram analysis of charmed baryon decays*, Phys. Rev. D **44** (1991) 2799.
- [86] BESIII collaboration, M. Ablikim *et al.*, *Observation of the Singly Cabibbo Suppressed Decay $\Lambda_c^+ \rightarrow n\pi^+$* , Phys. Rev. Lett. **128** (2022) 142001, [arXiv:2201.02056](#).
- [87] P. Fayet, *Supersymmetry and Weak, Electromagnetic and Strong Interactions*, Phys. Lett. B **64** (1976) 159.
- [88] P. Fayet, *Spontaneously Broken Supersymmetric Theories of Weak, Electromagnetic and Strong Interactions*, Phys. Lett. B **69** (1977) 489.
- [89] G. R. Farrar and P. Fayet, *Phenomenology of the Production, Decay, and Detection of New Hadronic States Associated with Supersymmetry*, Phys. Lett. B **76** (1978) 575.
- [90] P. Fayet, *Relations Between the Masses of the Superpartners of Leptons and Quarks, the Goldstino Couplings and the Neutral Currents*, Phys. Lett. B **84** (1979) 416.
- [91] ATLAS collaboration, *Search for long-lived, massive particles in events with displaced vertices and multiple jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, .
- [92] ATLAS collaboration, G. Aad *et al.*, *Search for Displaced Leptons in $\sqrt{s} = 13$ TeV pp Collisions with the ATLAS Detector*, Phys. Rev. Lett. **127** (2021) 051802, [arXiv:2011.07812](#).
- [93] ATLAS collaboration, G. Aad *et al.*, *A search for the decays of stopped long-lived particles at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **07** (2021) 173, [arXiv:2104.03050](#).

-
- [94] ATLAS collaboration, G. Aad *et al.*, *Search for long-lived charginos based on a disappearing-track signature using 136 fb^{-1} of pp collisions at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, Eur. Phys. J. C **82** (2022) 606, [arXiv:2201.02472](#).
 - [95] R. Davis, Jr. D. S. Harmer, and K. C. Hoffman, *Search for neutrinos from the sun*, Phys. Rev. Lett. **20** (1968) 1205.
 - [96] R. N. Mohapatra, *Massive neutrinos as probe of higher unification*, in *Puri Winter School in Physics: Particle Physics and Cosmology at the Interface*, 214–224, 1993, [arXiv:hep-ph/9306283](#).
 - [97] GALLEX collaboration, W. Hampel *et al.*, *GALLEX solar neutrino observations: Results for GALLEX IV*, Phys. Lett. B **447** (1999) 127.
 - [98] Kamiokande collaboration, Y. Fukuda *et al.*, *Solar neutrino data covering solar cycle 22*, Phys. Rev. Lett. **77** (1996) 1683.
 - [99] Super-Kamiokande collaboration, Y. Fukuda *et al.*, *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett. **81** (1998) 1562, [arXiv:hep-ex/9807003](#).
 - [100] SNO collaboration, Q. R. Ahmad *et al.*, *Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **89** (2002) 011301, [arXiv:nucl-ex/0204008](#).
 - [101] P. Minkowski, $\mu \rightarrow e\gamma$ at a Rate of One Out of 10^9 Muon Decays?, Phys. Lett. B **67** (1977) 421.
 - [102] M. Gell-Mann, P. Ramond, and R. Slansky, *Complex Spinors and Unified Theories*, Conf. Proc. C **790927** (1979) 315, [arXiv:1306.4669](#).
 - [103] R. N. Mohapatra and G. Senjanovic, *Neutrino Mass and Spontaneous Parity Non-conservation*, Phys. Rev. Lett. **44** (1980) 912.
 - [104] T. Yanagida, *Horizontal gauge symmetry and masses of neutrinos*, Conf. Proc. C **7902131** (1979) 95.
 - [105] J. Schechter and J. W. F. Valle, *Neutrino Masses in $SU(2) \times U(1)$ Theories*, Phys. Rev. D **22** (1980) 2227.
 - [106] CMS collaboration, A. Tumasyan *et al.*, *Search for long-lived heavy neutral leptons with displaced vertices in proton-proton collisions at $\sqrt{s}=13\text{ TeV}$* , JHEP **07** (2022) 081, [arXiv:2201.05578](#).
 - [107] A. M. Abdullahi *et al.*, *The present and future status of heavy neutral leptons*, J. Phys. G **50** (2023) 020501, [arXiv:2203.08039](#).
 - [108] ATLAS, CMS collaboration, G. Aad *et al.*, *Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV* , JHEP **08** (2016) 045, [arXiv:1606.02266](#).

- [109] LHCb collaboration, R. Aaij *et al.*, *Updated search for long-lived particles decaying to jet pairs*, Eur. Phys. J. C **77** (2017) 812, [arXiv:1705.07332](#).
- [110] D0 collaboration, V. M. Abazov *et al.*, *Search for Resonant Pair Production of long-lived particles decaying to b anti- b in p anti- p collisions at $s^{*}(1/2) = 1.96$ -TeV*, Phys. Rev. Lett. **103** (2009) 071801, [arXiv:0906.1787](#).
- [111] CDF collaboration, T. Aaltonen *et al.*, *Search for heavy metastable particles decaying to jet pairs in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, Phys. Rev. D **85** (2012) 012007, [arXiv:1109.3136](#).
- [112] ATLAS collaboration, G. Aad *et al.*, *Search for a light Higgs boson decaying to long-lived weakly-interacting particles in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, Phys. Rev. Lett. **108** (2012) 251801, [arXiv:1203.1303](#).
- [113] ATLAS collaboration, G. Aad *et al.*, *Search for long-lived, weakly interacting particles that decay to displaced hadronic jets in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Phys. Rev. D **92** (2015) 012010, [arXiv:1504.03634](#).
- [114] ATLAS collaboration, G. Aad *et al.*, *Search for pair-produced long-lived neutral particles decaying in the ATLAS hadronic calorimeter in pp collisions at $\sqrt{s} = 8$ TeV*, Phys. Lett. B **743** (2015) 15, [arXiv:1501.04020](#).
- [115] CMS collaboration, V. Khachatryan *et al.*, *Search for Long-Lived Neutral Particles Decaying to Quark-Antiquark Pairs in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV*, Phys. Rev. D **91** (2015) 012007, [arXiv:1411.6530](#).
- [116] LHCb collaboration, R. Aaij *et al.*, *Search for Dark Photons Produced in 13 TeV pp Collisions*, Phys. Rev. Lett. **120** (2018) 061801, [arXiv:1710.02867](#).
- [117] CMS collaboration, V. Khachatryan *et al.*, *A search for pair production of new light bosons decaying into muons*, Phys. Lett. B **752** (2016) 146, [arXiv:1506.00424](#).
- [118] CMS collaboration, A. M. Sirunyan *et al.*, *A search for pair production of new light bosons decaying into muons in proton-proton collisions at 13 TeV*, Phys. Lett. B **796** (2019) 131, [arXiv:1812.00380](#).
- [119] ATLAS collaboration, G. Aad *et al.*, *Search for long-lived neutral particles decaying into lepton jets in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, JHEP **11** (2014) 088, [arXiv:1409.0746](#).
- [120] ATLAS collaboration, G. Aad *et al.*, *Measurement of light-by-light scattering and search for axion-like particles with 2.2 nb^{-1} of Pb+Pb data with the ATLAS detector*, JHEP **03** (2021) 243, [arXiv:2008.05355](#), [Erratum: JHEP **11**, 050 (2021)].
- [121] CMS collaboration, A. M. Sirunyan *et al.*, *Evidence for light-by-light scattering and searches for axion-like particles in ultraperipheral PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, Phys. Lett. B **797** (2019) 134826, [arXiv:1810.04602](#).
- [122] *LHC Machine*, JINST **3** (2008) S08001.

-
- [123] M. Benedikt *et al.*, *Future Circular Collider - European Strategy Update Documents*, CERN, Geneva, 2019.
- [124] LHCb collaboration, A. A. Alves, Jr. *et al.*, *The LHCb Detector at the LHC*, JINST **3** (2008) S08005.
- [125] LHCb collaboration, C. Elsasser, *$\bar{b}b$ production angle plots*, .
- [126] LHCb collaboration, *LHCb reoptimized detector design and performance: Technical Design Report*, CERN-LHCC-2003-030, 2003.
- [127] LHCb collaboration, *LHCb VELO (VERtex LOcator): Technical Design Report*, CERN-LHCC-2001-011, 2001.
- [128] LHCb collaboration, *LHCb technical design report: Reoptimized detector design and performance*, .
- [129] LHCb collaboration, *LHCb VELO Upgrade Technical Design Report*, CERN-LHCC-2013-021, 2013.
- [130] LHCb collaboration, A. A. Alves Jr. *et al.*, *The LHCb detector at the LHC*, JINST **3** (2008) S08005.
- [131] S. T. U. Group, *Material for publications*, .
- [132] LHCb collaboration, R. Aaij *et al.*, *LHCb detector performance*, Int. J. Mod. Phys. **A30** (2015) 1530022, [arXiv:1412.6352](#).
- [133] LHCb collaboration, *LHCb outer tracker: Technical Design Report*, CERN-LHCC-2001-024, 2001.
- [134] R. Arink *et al.*, *Performance of the LHCb Outer Tracker*, JINST **9** (2014) P01002, [arXiv:1311.3893](#).
- [135] LHCb collaboration, *LHCb RICH: Technical Design Report*, CERN-LHCC-2000-037, 2000.
- [136] M. Adinolfi *et al.*, *Performance of the LHCb RICH detector at the LHC*, Eur. Phys. J. **C73** (2013) 2431, [arXiv:1211.6759](#).
- [137] I. Machikhiliyan and (for the LHCb calorimeter group), *Current status and performance of the lhc calorimeters*, Journal of Physics: Conference Series **293** (2011) 012052.
- [138] T. Head, *The LHCb trigger system*, JINST **9** (2014) C09015.
- [139] R. Aaij *et al.*, *The LHCb Trigger and its Performance in 2011*, JINST **8** (2013) P04022, [arXiv:1211.3055](#).
- [140] R. Aaij *et al.*, *Tesla : an application for real-time data analysis in High Energy Physics*, Comput. Phys. Commun. **208** (2016) 35, [arXiv:1604.05596](#).

- [141] LHCb collaboration, R. Aaij *et al.*, *The LHCb Upgrade I*, JINST **19** (2024) P05065, [arXiv:2305.10515](#).
- [142] LHCb collaboration, *Framework TDR for the LHCb Upgrade: Technical Design Report*, CERN-LHCC-2012-007, 2012.
- [143] LHCb collaboration, *LHCb Tracker Upgrade Technical Design Report*, CERN-LHCC-2014-001, 2014.
- [144] LHCb collaboration, R. Aaij *et al.*, *The LHCb Upgrade I*, [arXiv:2305.10515](#), to appear in JINST.
- [145] LHCb collaboration, *RTA and DPA dataflow diagrams for Run 1, Run 2, and the upgraded LHCb detector*, .
- [146] LHCb collaboration, *LHCb Upgrade GPU High Level Trigger Technical Design Report*, CERN-LHCC-2020-006, 2020.
- [147] R. Aaij *et al.*, *Allen: A high level trigger on GPUs for LHCb*, Comput. Softw. Big Sci. **4** (2020) 7, [arXiv:1912.09161](#).
- [148] D. H. Campora Perez, *Optimization of high-throughput real-time processes in physics reconstruction*, 2019. Presented 29 Nov 2019.
- [149] LHCb collaboration, R. Aaij *et al.*, *Measurement of the time-integrated CP asymmetry in $D^0 \rightarrow K_S^0 K_S^0$ decays*, JHEP **11** (2018) 048, [arXiv:1806.01642](#).
- [150] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP asymmetry in $D^0 \rightarrow K^- K^+$ decays*, Phys. Lett. **B767** (2017) 177, [arXiv:1610.09476](#).
- [151] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning*, Springer Series in Statistics, 2001.
- [152] W. D. Hulsbergen, *Decay chain fitting with a Kalman filter*, Nucl. Instrum. Meth. **A552** (2005) 566, [arXiv:physics/0503191](#).
- [153] G. Tuci, *Searching for confirmation of charm CP violation in K_S^0 final states at LHCb*, 2021, PhD thesis, University of Pisa, Pisa, CERN-THESIS-2020-325.
- [154] N. L. Johnson, *Systems of frequency curves generated by methods of translation*, Biometrika **36** (1949) 149.
- [155] J. Albrecht, V. V. Gligorov, G. Raven, and S. Tolk, *Performance of the LHCb High Level Trigger in 2012*, J. Phys. : Conf. Ser. **513** (2014) 012001, [arXiv:1310.8544](#), [arXiv:1310.8544](#).
- [156] Particle Data Group, P. A. Zyla *et al.*, *Review of particle physics*, Prog. Theor. Exp. Phys. **2020** (2020) 083C01.
- [157] H. Voss, A. Hoecker, J. Stelzer, and F. Tegenfeldt, *TMVA - Toolkit for Multivariate Data Analysis with ROOT*, PoS **ACAT** (2007) 040.

-
- [158] Hunter, Ross *et al.*, *An automated tool to facilitate consistent test-driven development of trigger selections for lhcb 's run 3*, EPJ Web Conf. **251** (2021) 04024.
 - [159] R. J. Hunter, *Precise measurements of the W boson's mass and lepton flavour universality, and trigger development with the LHCb experiment at CERN*, 2022. Presented 25 Nov 2022.
 - [160] LHCb collaboration, R. Aaij *et al.*, *Observation of the mass difference between neutral charm-meson eigenstates*, Phys. Rev. Lett. **127** (2021) 111801, [arXiv:2106.03744](#).
 - [161] LHCb collaboration, R. Aaij *et al.*, *Measurement of the time-integrated CP asymmetry in $D^0 \rightarrow K^- K^+$ decays*, [arXiv:2209.03179](#), to appear in Phys. Rev. Lett.
 - [162] LHCb collaboration, *Framework TDR for the LHCb Upgrade II: Opportunities in flavour physics, and beyond, in the HL-LHC era*, CERN, Geneva, 2021.
 - [163] A. Abba *et al.*, *An 'artificial retina' processor for track reconstruction at the full LHC crossing rate*, Nucl. Instrum. Meth. A **824** (2016) 260.
 - [164] V. Gorkavenko *et al.*, *LHCb potential to discover long-lived new physics particles with lifetimes above 100 ps*, Eur. Phys. J. C **84** (2024) 608, [arXiv:2312.14016](#).
 - [165] LHCb collaboration, R. Aaij *et al.*, *Measurement of B^0 , B_s^0 , B^+ and Λ_b^0 production asymmetries in 7 and 8 TeV proton-proton collisions*, Phys. Lett. **B774** (2017) 139, [arXiv:1703.08464](#).
 - [166] LHCb collaboration, R. Aaij *et al.*, *Measurement of the mass difference between neutral charm-meson eigenstates*, Phys. Rev. Lett. **122** (2019) 231802, [arXiv:1903.03074](#).
 - [167] LHCb collaboration, R. Aaij *et al.*, *Observation of the Mass Difference Between Neutral Charm-Meson Eigenstates*, Phys. Rev. Lett. **127** (2021) 111801, [arXiv:2106.03744](#).
 - [168] LHCb collaboration, R. Aaij *et al.*, *Measurement of the ratio of branching fractions of the decays $\Lambda_b^0 \rightarrow \psi(2S)\Lambda$ and $\Lambda_b^0 \rightarrow J/\psi\Lambda$* , JHEP **03** (2019) 126, [arXiv:1902.02092](#).
 - [169] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP asymmetry in $D^0 \rightarrow K_S^0 K_S^0$ decays*, Phys. Rev. D **104** (2021) L031102, [arXiv:2105.01565](#).
 - [170] *Nu, Mu and Pile-Up. The LHCb definitions of what we see and what we don't see.*, <https://twiki.cern.ch/twiki/bin/view/LHCb/NuMuPileUp>.
 - [171] M. Morello *et al.*, *Proposal for FPGA-based tracking in the LHCb downstream region*, CERN, Geneva, 2024.
 - [172] L. collaboration, *LHCb Data Acquisition Enhancement TDR*, CERN-LHCC-2024-001, 2024.
 - [173] R. Aaij *et al.*, *Performance of the LHCb trigger and full real-time reconstruction in Run 2 of the LHC*, JINST **14** (2019) P04013, [arXiv:1812.10790](#).

- [174] E. V. Butera. 2025, Master’s thesis, University of Pisa, Pisa, to appear on ETD (<https://etd.adm.unipi.it/ETD-db/ETD-search/search>).
- [175] L. Pica, *Selecting long-lived particles in the first trigger level at the LHC*, PoS **CORFU2023** (2024) 015.
- [176] Belle and Belle II collaborations, I. Adachi *et al.*, *Measurement of the time-integrated CP asymmetry in $D^0 \rightarrow K_S^0 K_S^0$ decays using Belle and Belle II data*, Phys. Rev. **D111** (2025) 012015, [arXiv:2411.00306](#).
- [177] Belle and Belle II collaborations, I. Adachi *et al.*, *Measurement of the time-integrated CP asymmetry in $D^0 \rightarrow K_S^0 K_S^0$ decays using opposite-side flavor tagging at Belle and Belle II*, [arXiv:2504.15881](#).
- [178] G. Punzi on behalf of the LHCb Collaboration, *Mixing and CP violation in charm decays at LHCb*, EPS-HEP 2025, 2025.
- [179] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP asymmetry in $D^0 \rightarrow K_S^0 K_S^0$ decays with the LHCb Upgrade I detector*, [arXiv:2510.14732](#).

Acknowledgements

Questo viaggio non sarebbe stato lo stesso senza tutte le persone che mi hanno accompagnato lungo il cammino, tra le tante soddisfazioni, i momenti felici e quelli un po' meno.

Sono entrato in questo percorso come uno studente appena laureato e, se oggi ne esco somigliando un po' di più ad uno scienziato, il merito è in gran parte del mio relatore Giovanni Punzi. A lui va la mia più profonda riconoscenza per le innumerevoli lezioni, di fisica e non solo, per avermi guidato con successo nel superare i miei limiti e per avermi insegnato che, ad ogni problema, esiste *sempre* una soluzione.

Una parte importante del merito va anche a tutti i miei colleghi, a Pisa e nel resto del mondo, che ringrazio per essere stati fonte continua di idee, confronto e intuizioni, insieme a tanta amicizia, compagnia e conforto.

Un pensiero particolarmente sentito va a tutte le persone che ho incontrato e con cui ho condiviso l'anno da simillfellow al CERN. Grazie per avermi fatto sentire pienamente parte di una comunità così straordinaria e per avermi insegnato a vivere *ogni* momento con naturalezza e spensieratezza – fra pranzi, crêpe e birre. Mi avete tanto riempito il cuore che ora è più grande.

Un grazie mai abbastanza profondo va alla mia compagna di tutto, Laura, per avermi accompagnato mano nella mano anche in questa avventura, esultando con me ad ogni successo e stringendomi forte in ogni difficoltà, nonostante tutto.

Un pensiero speciale va ai miei genitori e a mio fratello: grazie per essere sempre stati pronti ad accogliermi a braccia aperte, a qualunque ora del giorno e della notte, da qualunque parte del mondo io fossi di ritorno. Spero di aver ripagato almeno in parte tutto ciò che avete fatto per me.

Un immenso grazie anche a tutti i miei amici di Pisa, per le mille avventure, le cene, le birre sulle spallette dell'Arno e chi più ne ha più ne metta. Senza di voi la vita di ogni giorno sarebbe stata molto più vuota e noiosa. Un ringraziamento speciale va a Claudia e Donato, compagni sempre presenti, soccorritori nei momenti più bui e confidenti senza mai giudicare.

Ed infine, non credo ci sia bisogno di ringraziare me stesso: dovrei sapere già tutto.

That's all, folks!

Post Scriptum

This post scriptum has the purpose of informing the reader about the final results this thesis has led to, that could not be included in the original submission (occurred on 31/1/2025) due to the pending approval of the Collaboration. After the thesis submission, I kept working according to the requests of the internal LHCb review of the measurement, leading to the finalization and approval of the analysis for public diffusion, and preparation of the paper. The final, unblinded result is:

$$\mathcal{A}^{CP}(K_S^0 K_S^0) = (1.86 \pm 1.04 \pm 0.41)\%,$$

where the first uncertainty is statistical, and the second is systematic. This is the most precise determination of this quantity by a single experiment to date. It is compatible with CP symmetry and with the average of previous determinations [2, 3, 43, 44, 176, 177]. The combined LHCb result, obtained by combining the present measurement with the previous LHCb determinations [2, 44], is: $\mathcal{A}^{CP}(K_S^0 K_S^0) = (-0.37 \pm 0.78 \pm 0.29)\%$. This is graphically shown in Figure PS.1, together with other $\mathcal{A}^{CP}(K_S^0 K_S^0)$ measurements.

The result has been publicly presented for the first time in July 2025 at the EPS-HEP international conference [178]. The paper describing the measurement and the result is now available as a preprint [179] and it's going to be submitted to the Journal of High Energy Physics.

During the review finalization, which confirmed the correctness of the analysis in all its aspects, no changes have been made to the analysis strategy described in this thesis, but an additional piece of data has been added that had become available in the meantime, and some minor fine-tunings of the selection have been applied. The main difference with the reported version is a small revision of $K_S^0 K_S^0$ and $K_S^0 \pi^+ \pi^-$ sample offline selections. These changes aimed at improving the kinematic overlap between the signal and calibration samples.

The revised calibration sample offline selections allowed to increase the statistics by a factor of 10 while keeping an $S/B \sim 10$ ($S/B \sim 20$ in the previous configuration). This background level can still be considered negligible, as proven in the Run 2 analysis, where the calibration sample exhibited a similar S/B ratio [2]. The Δm distribution for the calibration sample after the application of the updated offline selections is reported in Figure PS.2.

The signal offline selections revision just led to an additional cut, applied on top of what has been described in Chapter 8, that has otherwise been left unchanged. This aims at rejecting candidates in regions where the calibration-to-signal density ratio falls below 30, as this is the number of calibration candidates used in the asymmetry correction (Section 8.7). The resulting statistics loss is compensated by the increased inclusion of two additional data blocks, referred to as blocks 7 and 8. These increase the signal




Figure PS.1: Visual representation of $\mathcal{A}^{CP}(D^0 \rightarrow K_S^0 K_S^0)$ measurements [2, 3, 43, 44, 176, 177]. The average of all of the LHCb measurements and the updated world average are also reported.

sample statistics by approximately 20%, offsetting the impact of the kinematic-overlap selection. The invariant mass distributions for the signal sample after the application of offline selections and integrating all blocks are reported in Figure PS.3, PS.4. After the application of updated selections on signal and calibration samples, $\mathcal{A}^{CP}(K_S^0 K_S^0)$ is extracted block by block in the same way reported in Chapter 8. The results for each data block are reported in Figure PS.5, showing good compatibility. This plot has been added to the supplemental material published with this result, due to its value in demonstrating that the innovative and sophisticated correction procedure developed in this work, actually achieves its intended purpose of compensating the variations of detector response in different data-taking conditions.

During the analysis finalization, the systematic uncertainties associated with the measurement were determined, as summarized in Table PS.1. The lack of any effect due to K_S^0 decays and regeneration effects is given by the increase of the calibration sample statistics. With the updated selections, the associated uncertainty decreases to 0.05%, that is considered negligible with respect to the statistical one.




Figure PS.2: The Δm distribution for calibration sample candidates after application of revised offline selections.

Systematic source	Uncertainty [%]
Calibration sample statistics	0.24
Choice of k in weighting	0.20
Fit model choice	0.27
Total	0.41

Table PS.1: Summary of systematic uncertainties on $\mathcal{A}^{CP}(K_S^0 K_S^0)$ in percent.




Figure PS.3: Invariant mass distributions and fit for the high-purity bin sample, combining all data blocks. The histograms for the D^0 (\bar{D}^0) candidates are reported in the left (right) column. The top row reports Δm distributions. The middle and bottom rows report histograms for the masses of the two K_S^0 candidates. The red area indicates the $D^0 \rightarrow K_S^0 K_S^0$ component. The combinatorial component refers to random combination of pions for the $m(K_S^0)$ distributions.

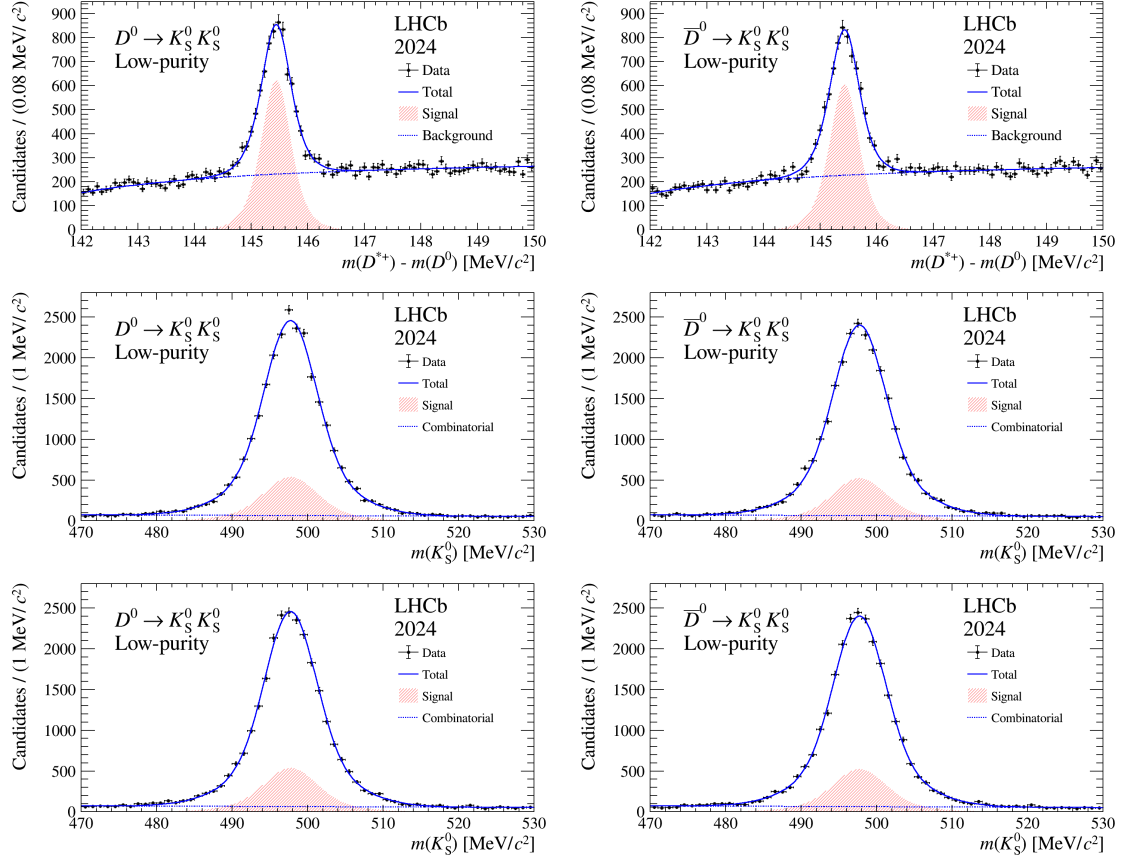


Figure PS.4: Invariant mass distributions and fit for the low-purity bin sample, combining all data blocks. The histograms for the D^0 (\bar{D}^0) candidates are reported in the left (right) column. The top row reports Δm distributions. The middle and bottom rows report histograms for the masses of the two K_S^0 candidates. The red area indicates the $D^0 \rightarrow K_S^0 K_S^0$ component. The combinatorial component refers to random combination of pions for the $m(K_S^0)$ distributions.




Figure PS.5: Measurements of $\mathcal{A}^{CP}(D^0 \rightarrow K_S^0 K_S^0)$ in different data blocks, rearranged in chronological order. The quoted uncertainties are statistical only. The band represents the statistical uncertainty associated with the combined result.