LUND
UNIVERSITY

Probing the Structure of the $\mathbf{f_0(980)}$ via K^+K^- Coalescence in Proton–Proton Collisions at $\sqrt{s} = 5.02$ TeV

A PYTHIA-Based Coalescence Afterburner Study

Vishvam Patel

Bachelor's Thesis (15 hp), FYSK04

Duration: 6 months

Department of Physics, Lund University, Lund, Sweden

Division of Particle and Nuclear Physics

Supervisors: Peter Christiansen, Sushanta Tripathy

Bachelor of Science in Physics

Contents

Acknowledgements

I would like to begin by expressing my sincere gratitude to my supervisor, Professor Peter Christiansen, for welcoming me into this project and providing such an open and supportive research environment.

A special thank you goes to my co-supervisor, Dr. Sushanta Tripathy, whose guidance, insight, patience, and encouragement were invaluable at every stage of this work. This thesis would not have taken the shape it did without your involvement. I am particularly grateful for the clarity and care with which you introduced me to the foundational concepts of this problem, and for your tireless help with the computational and programming aspects of the project. Your support made a genuine difference.

I would also like to thank my parents and sister for their unwavering encouragement through every high and low of this journey.

Finally, I utilized ChatGPT and Claude for language editing to improve the clarity and flow of the English text in the thesis.

Abstract

The internal structure of the scalar meson $f_0(980)$ remains an open question, with competing interpretations as a conventional $q\bar{q}$ state, a tetraquark, or a loosely bound K^+K^- hadronic molecule. This thesis investigates the molecular hypothesis through a Monte Carlo coalescence calculation applied to proton–proton collisions at $\sqrt{s} = 5.02$ TeV. Kaon phase-space distributions are generated with PYTHIA 8 using the Monash 2013 tune, validated against ALICE measurements, and a p_{T}-dependent correction is applied to isolate discrepancies arising from the coalescence model itself. The formation probability is evaluated from the Wigner-function overlap of a Gaussian molecular wavefunction, with additional scans of sharp momentum and spatial cutoffs used to assess parameter sensitivity. The model successfully reproduces the overall shape and magnitude of the measured $f_0(980)$ spectrum at low p_{T}, supporting the presence of a significant K^+K^- molecular component in this region. At higher p_{T}, the predicted yield increasingly underestimates the data, a deficit that persists across all molecular sizes considered and follows naturally from the molecular picture. The results suggest a mixed internal structure, with the high-p_{T} excess pointing toward an additional contribution from a more compact configuration such as a $q\bar{q}$ state or tetraquark.

Chapter 1

Introduction and Motivation

One of the most fundamental goals of modern physics is to understand how matter is structured at the subatomic level. The strong nuclear force governs the internal structure of protons, neutrons, and all composite particles known as hadrons. Despite decades of experimental and theoretical progress, many basic questions remain open: how quarks combine to form bound states, what configurations are permitted by nature, and etc. An example of this is the scalar meson $f_0(980)$, whose internal structure remains unresolved and can be probed through its production at the Large Hadron Collider (LHC).

1.1 The Strong Interaction and QCD

The Universe is governed by four fundamental forces: gravity, electromagnetism, the weak nuclear force, and the strong nuclear force. At subatomic scales, the strong force is by far the most powerful of these, responsible for binding quarks and gluons into hadrons and for holding atomic nuclei together [2]. Its dynamics account for nearly all of the visible mass in the observable Universe [16, 18].

The theoretical framework describing the strong interaction is Quantum Chromodynamics (QCD), in which quarks interact by exchanging gluons, the force carriers of the strong force. The charge associated with the strong interaction is *color charge*, which quarks carry in one of three forms: red, green, or blue. Unlike the photon in electrodynamics, gluons themselves carry color charge and can interact with each other, which is responsible for different features of QCD [16]. A key requirement imposed by QCD is that all physical states must be color neutral, a property known as *color confinement*. This means that quarks and gluons cannot exist as free isolated particles but are always bound within hadrons. Many aspects of how this confinement gives rise to observed hadronic states are still not fully understood, making hadronic structure an active and important area of research [1, 2].

A complementary feature of QCD is *asymptotic freedom*, which is a direct consequence of the running of the strong coupling constant α_s with the momentum scale [3]. At low momentum transfers, α_s becomes large and binds quarks and gluons into hadrons, while at high momentum transfers it decreases and quarks behave as nearly free particles. This scale dependence is essential for treating QCD perturbatively in the high-energy regime probed at the LHC.

1.2 The Quark Model

In 1964, Gell-Mann and Zweig independently proposed the quark model, which provided a systematic way of organizing the known hadron spectrum. In this model, hadrons are composite objects built from quarks, which come in six flavours: up (u), down (d), strange (s), charm (c), bottom (b), and top (t). The first three are referred to as light quarks and the remaining three as heavy quarks. Quarks carry fractional electric charge, spin-$\frac{1}{2}$, baryon number $1/3$, and a color charge. In the simplest version of the model, hadrons are classified as either mesons, made of a quark–antiquark pair ($q\bar{q}$), or baryons, made of three quarks (qqq) [16]. Other quantum numbers, such as electric charge, isospin, and strangeness, follow from the contributions of

the constituent quarks.

This framework has been very successful in describing the known hadron spectrum and predicting the properties of many states [1]. However, a growing number of experimentally observed resonances do not fit naturally into the simple $q\bar{q}$ or qqq picture, motivating the study of hadronic states that go beyond the minimal quark model [2].

1.3 Exotic Hadrons

QCD itself does not restrict hadrons to only $q\bar{q}$ and qqq configurations. The only requirement is that physical states are color neutral, which in principle allows for a much wider range of configurations. These include tetraquarks ($qq\bar{q}\bar{q}$), pentaquarks ($qqqq\bar{q}$), and hadronic molecules, which are loosely bound states of two or more color-neutral hadrons held together by residual strong interactions. States beyond the minimal quark model are collectively referred to as *exotic hadrons* [1, 2].

The first experimental evidence for a state beyond the conventional quark model came in 2003, when the Belle collaboration discovered the $X(3872)$, a charmonium-like state whose properties are inconsistent with a simple $c\bar{c}$ assignment [14]. This was followed in 2015 by the LHCb observation of pentaquark-like states in the $J/\psi\, p$ invariant mass spectrum [15]. In the light quark sector, the scalar mesons, and the $f_0(980)$ in particular, have long been considered candidates for exotic configurations.

1.4 Heavy-Ion Collisions and the LHC

The LHC at CERN provides access to high-energy proton–proton (pp) and heavy-ion collisions, enabling detailed experimental studies of hadron production. A key observable in these experiments is the differential yield of produced particles as a function of transverse momentum (p_{T}), the component of momentum perpendicular to the beam axis. Since different internal structures correspond to different spatial sizes and formation mechanisms, they produce distinct p_{T} distributions, making the transverse momentum spectrum a useful observable for probing the internal structure of hadrons.

pp collisions at the LHC provide an environment for studying hadron production without the added complexity of a dense nuclear medium, and serve as a clean baseline for comparisons with heavy-ion results. In this work, Pythia Monte Carlo simulations are used to generate kaon phase-space distributions, and the results are compared with measurements from the ALICE (A Large Ion Collider Experiment) collaboration. ALICE is one of the main experiment at the LHC with strong particle identification capabilities across a wide range of transverse momenta, making it well suited for measuring resonances such as the $f_0(980)$ [1, 13].

1.5 Hadronization

In high-energy collisions, quarks and gluons are produced in a coloured state and must subsequently combine into colour-neutral bound states through a process known as hadronization. The simulations in this thesis rely on Pythia, which implements hadronization through the *Lund string fragmentation model*: colour flux tubes ("strings") stretched between separating partons break iteratively into smaller pieces, each piece corresponding to a produced hadron [20]. Once the hadrons are formed, their interactions gradually cease at *kinetic freeze-out*, after which their momentum distributions are no longer modified. In coalescence approaches such as the one used here, the formation of loosely bound states is assumed to occur at this kinetic freeze-out stage, where independently produced hadrons can combine if they are close enough in phase space [17].

Since hadronization is governed by nonperturbative QCD dynamics, it cannot be calculated directly from first principles [16]. Instead, phenomenological models are used to describe how quarks and gluons form bound states and how these states appear in measurable quantities such as particle yields and momentum spectra [6].

1.6 The $f_0(980)$ Scalar Meson

The $f_0(980)$ is a scalar meson with quantum numbers $J^{PC} = 0^{++}$, where J denotes the total angular momentum, P the intrinsic parity, and C the charge-conjugation parity. The notation 0^{++} therefore identifies the $f_0(980)$ as a spinless state with positive parity and positive C-parity. Its mass is quoted by the Particle Data Group as $m_{f_0} = 990 \pm 20$ MeV/c^2 [18]. The dominant decay mode is $f_0(980) \to \pi^+\pi^-$, with a secondary coupling to K^+K^- near threshold. Despite being known for over half a century, its internal structure remains unresolved, making it one of the key questions in the light flavour sector [2, 6].

1.7 Motivation and Goals

The unresolved structure of the $f_0(980)$ raises the question of whether exotic multiquark configurations already appear in the lightest part of the hadronic spectrum, and tests our understanding of how the confining dynamics of QCD give rise to composite bound states. Solving this question would be a significant advance in hadron physics [1, 2].

pp collisions at the LHC provide a useful environment for probing the internal structure of the $f_0(980)$ through its production mechanism. The central idea is that the production of a composite state depends on what it is made of. A compact $q\bar{q}$ or tetraquark state is formed during hadronization of quarks that are close together in phase space, whereas a K^+K^- molecule forms later, through the coalescence of two independently produced kaons. These different mechanisms predict different transverse momentum spectra, which can be measured with the ALICE detector [7, 8].

The coalescence model provides the theoretical link between the internal structure of the $f_0(980)$ and its measured production properties. The relevant object is the Wigner function, the quantum-mechanical analogue of a classical phase-space distribution, which encodes the internal structure of a bound state in terms of the relative coordinate and relative momentum of its constituents. By integrating this Wigner function under different structural assumptions with the phase-space distributions of its constituents from Monte Carlo simulations, quantitative predictions can be generated and tested directly against data.

The aims of this project are therefore to:

(i) Implement a Monte Carlo-based coalescence calculation using phase-space distributions from Pythia-generated pp collision events as input;

(ii) Compute the transverse momentum spectrum of the $f_0(980)$ under the K^+K^- molecular hypothesis using a two-particle coalescence approach with a Gaussian wavefunction;

(iii) Compare the resulting simulated spectra with experimental data from the ALICE collaboration in pp collisions at $\sqrt{s} = 5.02$ TeV, to test if coalescence is consistent with the measured production properties.

Chapter 2

Background

2.1 Kinematics in High-Energy Collisions

The description of particle production in high-energy collisions relies on relativistic kinematics. Since the LHC probes particles at energies far exceeding their rest masses, a fully relativistic and Lorentz-invariant formulation is required throughout. Particles are described using four-momentum vectors,

$$p^\mu = (E, \mathbf{p}), \tag{2.1}$$

where E is the energy and $\mathbf{p}$ the three-momentum. Natural units are used throughout, so $c = 1$ in all formulae, with factors of c restored only when quoting masses in MeV/c^2 and momentum in GeV/c.

The invariant mass of the colliding system defines the centre-of-mass energy, $\sqrt{s}$. For two colliding protons with four-momenta p_1^μ and p_2^μ, the Lorentz-invariant squared centre-of-mass energy is

$$s = (p_1^\mu + p_2^\mu)^2 = (E_1 + E_2)^2 - |\mathbf{p}_1 + \mathbf{p}_2|^2, \tag{2.2}$$

so that $\sqrt{s}$ is the total energy available for particle production in the centre-of-mass frame. The data analysed in this thesis were collected at $\sqrt{s} = 5.02$ TeV [13].

It is important to consider how variables behave under Lorentz transformations to ensure that physical results do not depend on the choice of reference frame. A particularly important observable is the transverse momentum,

$$p_\mathrm{T} = \sqrt{p_x^2 + p_y^2}, \tag{2.3}$$

which is the component of momentum perpendicular to the beam direction. Since p_T is invariant under Lorentz transformations along beam axis (longitudinal boosts), see Section 2.2, it is a natural and convenient observable in collisions at the LHC [13]. Another important kinematic variable is the rapidity,

$$y = \frac{1}{2} \ln\left(\frac{E + p_z}{E - p_z}\right), \tag{2.4}$$

which has the property that differences in rapidity are invariant under longitudinal boosts. Particle production at colliders is approximately boost invariant near mid-rapidity ($y \approx 0$), which is the region least affected by beam remnants and provides a clean environment for studying production mechanisms [13].

2.1.1 The Invariant Yield

The fundamental Lorentz-invariant quantity in single-particle inclusive production is

$$E\frac{d^3N}{d^3p}, \tag{2.5}$$

which is invariant under arbitrary Lorentz transformations. Expressing the three-momentum element in cylindrical coordinates,

$$d^3p = p_T\, dp_T\, d\phi\, dp_z = E\, p_T\, dp_T\, d\phi\, dy = m_T \cosh y\; p_T\, dp_T\, d\phi\, dy, \tag{2.6}$$

where $m_T = \sqrt{m^2 + p_{\mathrm{T}}^2}$ is the transverse mass, and use has been made of $E = m_T \cosh y$. Substituting this gives

$$E\frac{d^3N}{d^3p} = \frac{1}{2\pi p_T}\frac{d^2N}{dy\, dp_T}, \tag{2.7}$$

where the factor of 2π comes from assuming azimuthal symmetry as is the case for minimum-bias collisions so that $\frac{dN}{d\phi} = \frac{N}{2\pi}$. Eq (2.7) is the standard form of the *invariant yield*. The main observable used in the present analysis is the invariant yield per inelastic event,

$$\frac{1}{N_{\mathrm{event}}}\frac{d^2N}{dy\, dp_T}, \tag{2.8}$$

where N_{event} is the total number of analysed events. Since both dy and dp_{T} are invariant under longitudinal boosts, the combination $d^2N/dy\, dp_T$ is invariant under boosts along the beam axis [4].

2.2 Lorentz Transformation

A fundamental requirement in high-energy physics is that physical observables must be independent of the reference frame in which they are measured. This principle is embodied in special relativity through *Lorentz invariance*, which ensures that the laws of physics remain unchanged under transformations between inertial frames. The Lorentz transformation is implemented numerically in the results where the relative momentum of each K^+K^- pair must be evaluated in the rest frame of the candidate $f_0(980)$.

2.2.1 Boost along the beam axis

For a Lorentz boost along the beam (z-) direction, characterised by $\beta = v/c$ and $\gamma = 1/\sqrt{1-\beta^2}$, the energy and longitudinal momentum transform as,

$$E' = \gamma(E - \beta p_z), \qquad p'_z = \gamma(p_z - \beta E), \qquad p'_{\mathrm{T}} = p_{\mathrm{T}}, \tag{2.9}$$

where the transverse momentum components are left unchanged. The p_{T} is unchanged under boosts along the beam direction, making it a useful observable for comparing particle production across events [4]. The rapidity defined in Eq (2.4) transforms under a longitudinal boost as

$$y' = y - \tanh^{-1}(\beta), \tag{2.10}$$

so that the *shape* of rapidity distributions is preserved. This property supports the invariance of the differential yield $d^2N/dy\, dp_{\mathrm{T}}$ and is widely used in experimental analyses to compare results across different collision systems and frames.

2.2.2 Boost in an arbitrary direction

The coalescence calculation requires boosting four-vectors into the rest frame of each K^+K^- pair, whose three-momentum is generally not aligned with the beam axis. For a general boost with velocity $\boldsymbol{\beta}$, magnitude $\beta = |\boldsymbol{\beta}|$, and Lorentz factor $\gamma = (1-\beta^2)^{-1/2}$, a four-vector $R^\mu = (t, \mathbf{r})$ transforms as

$$t' = \gamma(t - \boldsymbol{\beta}\cdot\mathbf{r})\,, \tag{2.11}$$

$$\mathbf{r}' = \mathbf{r} + \frac{\gamma-1}{\beta^2}(\boldsymbol{\beta}\cdot\mathbf{r})\,\boldsymbol{\beta} - \gamma\,\boldsymbol{\beta}\,t. \tag{2.12}$$

For $\boldsymbol{\beta}$ along $\hat{z}$, these reduce to Eq (2.9). They can be written compactly in matrix form as

$$\begin{pmatrix} t' \\ x' \\ y' \\ z' \end{pmatrix} = \begin{pmatrix} \gamma & -\gamma\beta_x & -\gamma\beta_y & -\gamma\beta_z \\ -\gamma\beta_x & 1+(\gamma-1)\dfrac{\beta_x^2}{\beta^2} & (\gamma-1)\dfrac{\beta_x\beta_y}{\beta^2} & (\gamma-1)\dfrac{\beta_x\beta_z}{\beta^2} \\ -\gamma\beta_y & (\gamma-1)\dfrac{\beta_y\beta_x}{\beta^2} & 1+(\gamma-1)\dfrac{\beta_y^2}{\beta^2} & (\gamma-1)\dfrac{\beta_y\beta_z}{\beta^2} \\ -\gamma\beta_z & (\gamma-1)\dfrac{\beta_z\beta_x}{\beta^2} & (\gamma-1)\dfrac{\beta_z\beta_y}{\beta^2} & 1+(\gamma-1)\dfrac{\beta_z^2}{\beta^2} \end{pmatrix} \begin{pmatrix} t \\ x \\ y \\ z \end{pmatrix}. \tag{2.13}$$

Given the boost velocity of a K^+K^- pair, this transformation allows one to move to the pair rest frame, where the total momentum vanishes. The four-momentum of each kaon is then expressed in this frame, making it straightforward to evaluate their relative momentum [4]. This quantity is essential in the coalescence model, as it governs the likelihood of forming a bound or resonant state.

2.3 Phase Space and Particle Distributions

A central concept in the description of particle production is *phase space*, the combined space spanned by the three position coordinates $\mathbf{x}$ and the three momentum components $\mathbf{p}$ of a particle. The state of a system can be described by a single-particle distribution function $f(\mathbf{x}, \mathbf{p}, t)$, which represents the phase-space density of particles at a given position, momentum, and time. Integrating f over the relevant variables gives observable quantities such as particle yields and momentum spectra [9].

Experiments typically measure momentum-space distributions of the form dN/d^3p, which give the number of particles produced per unit volume of momentum space [13]. Both $E\,d^3N/d^3p$ and $d^2N/dy\,dp_\mathrm{T}$ are experimentally accessible and are related to one another through the kinematic prefactor $1/(2\pi p_\mathrm{T})$, as shown in Section 2.1.1. The form $d^2N/dy\,dp_\mathrm{T}$ is preferred in this work because it is invariant under boosts along the beam axis. These distributions are directly connected to the underlying phase-space density and play a central role in coalescence models, where the probability of a bound state forming depends on the phase-space overlap between its constituents at freeze-out.

One should note here that, in quantum mechanics, position and momentum cannot both be known exactly at the same time. The phase-space description used in this thesis is therefore a semi-classical approximation, where particles are treated as having approximate positions and momenta at freeze-out [10]. The quantum nature of the bound state is then included through the Wigner function in the coalescence calculation as discussed in Section 3.1.2.

2.4 Experimental Uncertainties

All measurements and simulations carry two main types of uncertainty: statistical and systematic. Statistical uncertainties arise from the finite number of events analysed and decrease with increasing sample size. Systematic uncertainties reflect limitations in the experimental setup or model assumptions and do not reduce with more data. Both contributions must be accounted for when comparing simulation results with

experimental data.

Statistical Uncertainties

Statistical uncertainties come from the finite number of observed events and follow Poisson statistics. For a count N, the uncertainty scales as $\sqrt{N}$, giving a relative uncertainty of $1/\sqrt{N}$. For yields normalised per event, the statistical uncertainty is

$$\sigma_{\text{stat}} = \frac{\sqrt{N}}{N_{\text{event}}}, \tag{2.14}$$

where N is the number of entries in a given bin and N_{event} is the total number of analysed events [13]. This expression applies to both experimental data and Monte Carlo simulations, including the coalescence-based calculations used in this work [7]. In practice, both the data sample and the simulated event sample used here are large enough that the statistical uncertainty is small compared to the systematic uncertainty discussed below.

Systematic Uncertainties

Systematic uncertainties arise from imperfect knowledge of the model and analysis inputs, and are estimated by varying key parameters within reasonable ranges. The dominant contributions in this work come from the value of the Gaussian width parameter σ_ρ in the Wigner function (corresponding to RMS radii of approximately 1.0–1.5 fm) [13]. These contributions are evaluated independently and combined to give the total systematic uncertainty. However, in the present thesis no dedicated evaluation of systematic uncertainties has been performed, and the comparison therefore includes only statistical uncertainties. A full treatment of systematic effects would require a more detailed study of the model assumptions, and fitting procedures.

Propagation of Uncertainties

Uncertainties on derived quantities are obtained using standard error propagation. For a function $f(x_1, x_2, \ldots, x_n)$, where x_i are measured input variables with variances $\sigma_{x_i}^2$, the total variance of f is

$$\sigma_f^2 = \sum_{i=1}^{n} \left(\frac{\partial f}{\partial x_i} \right)^2 \sigma_{x_i}^2 + 2 \sum_{i<j} \frac{\partial f}{\partial x_i} \frac{\partial f}{\partial x_j} \sigma_{x_i x_j}, \tag{2.15}$$

where $\partial f / \partial x_i$ is the partial derivative evaluated at the measured values and $\sigma_{x_i x_j}$ is the covariance between x_i and x_j. When the input uncertainties are small and uncorrelated, the covariance terms vanish, and the variances simply add in quadrature weighted by the partial derivatives. Two cases arise frequently. For $f = aA$, where a is a constant,

$$\sigma_f^2 = a^2 \, \sigma_A^2. \tag{2.16}$$

For $f = A + B$, the general result is

$$\sigma_f^2 = \sigma_A^2 + \sigma_B^2 + 2 \, \sigma_{AB}, \tag{2.17}$$

where $\sigma_{AB} = \rho_{AB} \, \sigma_A \, \sigma_B$ is the covariance between A and B, and ρ_{AB} is their correlation coefficient. When A and B are independent, $\sigma_{AB} = 0$ and the uncertainties add in quadrature. This formalism is widely used in particle physics. Statistical and systematic contributions are evaluated separately and then combined in quadrature [5].

2.5 $f_0(980)$ Structure Configurations and Experimental Evidence

The internal structure of a hadron directly influences how it is produced in high-energy collisions. The scalar meson $f_0(980)$ is a particularly interesting case because its internal configuration is not uniquely determined within the conventional quark model. Several competing interpretations exist, each with different spatial

scales, binding mechanisms, and production characteristics. These differences are especially relevant for coalescence models, where formation depends on the phase-space overlap between constituents.

Quark–Antiquark ($q\bar{q}$) Configuration

In the conventional picture, the $f_0(980)$ is described as a quark–antiquark meson, typically with a dominant $s\bar{s}$ component and quantum numbers $J^{PC} = 0^{++}$. The state is held together by color confinement, which gives rise to a compact spatial wavefunction. In recombination-type pictures, such states form at the partonic level when a q and $\bar{q}$ are sufficiently close in coordinate and momentum space [2]; in the Lund string-fragmentation framework used by Pythia, the same compact states emerge directly from the breaking of colour strings between separating partons rather than from explicit phase-space recombination.

Tetraquark ($qq\bar{q}\bar{q}$) Configuration

An alternative interpretation describes the $f_0(980)$ as a tetraquark state built from a diquark–antidiquark pair. Attractive color interactions bind quarks into correlated substructures, and this framework can naturally account for features of the scalar meson spectrum such as inverted mass ordering and suppressed electromagnetic couplings. Like the $q\bar{q}$ case, tetraquark states are expected to be compact and produced at the partonic level, with yields governed by quark-level phase-space overlap [2, 6].

Hadronic Molecule (K^+K^-) Configuration

A widely discussed alternative is that the $f_0(980)$ is a hadronic molecule, interpreted as a loosely bound K^+K^- state. Unlike compact quark configurations, the binding here arises from residual strong interactions between color-neutral mesons, resulting in a small binding energy and a larger spatial extent. The proximity of the $f_0(980)$ mass to the K^+K^- threshold supports this interpretation. In coalescence models, such states form at the hadronic level near kinetic freeze-out through the combination of independently produced kaons [1, 7].

Phenomenological Implications

The key distinction between these three configurations is their spatial size. Compact states ($q\bar{q}$ and tetraquarks) require strong localization in phase space at the partonic level, whereas molecular states are more spatially extended and can form from less correlated constituents. This leads to observable differences in production: extended states typically have reduced phase-space overlap at high transverse momentum, resulting in softer p_{T} spectra and suppressed yields at large p_{T} compared to compact configurations [1, 7, 8].

Experimental Status

The $f_0(980)$ has been studied experimentally through its $\pi^+\pi^-$ decay channel and its coupling to K^+K^-. However, different measurements have led to competing interpretations. CMS results based on elliptic flow scaling in p–Pb collisions are consistent with a compact $q\bar{q}$ structure, while analyses using Flatté parameterizations and compositeness arguments suggest a significant molecular component [1]. No consensus has yet been reached.

The approach taken in this work is to exploit these structural differences through coalescence model calculations, and to assess which configuration gives the most consistent description of the measured $f_0(980)$ transverse momentum spectrum [1, 2, 6, 7].

2.6 ALICE Data

The ALICE experiment at the LHC is designed to study particle production in high-energy collisions with excellent particle identification capabilities. The primary observable used in this analysis is the invariant yield introduced in Section 2.1.1, $\frac{1}{N_{\mathrm{event}}} d^2N/dy\, dp_{\mathrm{T}}$, which gives the normalised production rate as a function of

transverse momentum and rapidity. This quantity allows for a direct comparison between experimental data and theoretical predictions [11, 13].

In this work, the measured charged-kaon transverse momentum spectra are used as the benchmark for validating the PYTHIA event generator output, while the measured $f_0(980)$ spectra provide the reference for the final comparison. The shape of these spectra reflects the underlying production mechanisms: particles at low p_{T} are dominated by soft, thermal-like processes, while those at high p_{T} arise from hard parton scatterings [11].

Since coalescence depends on the phase-space distribution of the constituent particles, the resulting $f_0(980)$ spectrum is directly linked to the input kaon distributions. Comparing the simulated coalescence yield with the ALICE $f_0(980)$ measurement therefore provides a direct test of whether a K^+K^- molecular picture can reproduce the observed production behaviour [7, 11].

Chapter 3

Methodology

This chapter describes the theoretical framework and numerical methods used in this work to study the production of the $f_0(980)$ under the K^+K^- molecular hypothesis. First, the coalescence model and the Wigner-function formalism used to describe molecule formation are introduced. The event generation procedure using PYTHIA 8 is then presented, followed by the implementation of the coalescence afterburner and the Monte Carlo procedure used to determine whether a $+K^+K^-$ pair forms an $f_0(980)$ candidate.

3.1 Theoretical Framework

3.1.1 The Coalescence Model

The coalescence model is widely used to describe the production of loosely bound composite states in high-energy collisions [1, 22]. The main idea behind it is that if two particles are produced close together in phase space, meaning their relative coordinates and momenta are both small, then they can combine to form a composite bound state. The probability of this coalescence is determined by the overlap between the phase-space distributions of the two combining particles and the wavefunction describing their bound state.

The model is particularly natural for hadronic molecules such as the K^+K^- interpretation of the $f_0(980)$. Although the $f_0(980)$ is a resonant state rather than a stable bound state, its mass quoted by the Particle Data Group [18] lies close to the K^+K^- threshold. Taking the mass to lie just below threshold, the K^+K^- component of its wavefunction resembles that of a loosely bound state [1]. The $f_0(980)$ can then be modeled as forming via $K^+K^- \to f_0(980)$, similarly to the well-established coalescence production of the deuteron via $p + n \to d$ [22].

In contrast to fragmentation, which operates at the partonic level and is better suited for compact $q\bar{q}$ states, coalescence acts at the hadronic level. For this reason, and given the molecular hypothesis under investigation, the coalescence framework can provide the appropriate description of $f_0(980)$ production.

3.1.2 Wigner Function Coalescence

A more detailed version of the coalescence model describes the formation probability using the Wigner function, which gives a quantum-mechanical description of the particle's internal structure in phase space [1, 22]. Following Ref. [1], the production yield of the $f_0(980)$ can then be written as

$$\begin{aligned}\frac{dN}{d^3\mathbf{P}} = g \int p_1^\mu \, d^3\sigma_{1\mu} \, \frac{d^3\mathbf{p}_1}{E_1} \, p_2^\nu \, d^3\sigma_{2\nu} \, \frac{d^3\mathbf{p}_2}{E_2} \\ \times f_K^+(\mathbf{x}_1, \mathbf{p}_1, t_1) \, f_K^-(\mathbf{x}_2, \mathbf{p}_2, t_2) \\ \times \rho^W(\mathbf{x}_1^*, \mathbf{x}_2^*; \, \mathbf{p}_1^*, \mathbf{p}_2^*; \, t^*) \, \delta^3(\mathbf{P} - \mathbf{p}_1 - \mathbf{p}_2),\end{aligned} \tag{3.1}$$

where $dN/d^3\mathbf{P}$ is the number of $f_0(980)$ particles produced per unit momentum, and $\mathbf{P}$ is the three-momentum of the produced state. The terms $p_i^\mu \, d^3\sigma_{i\mu}$ and $d^3\mathbf{p}_i/E_i$ together count the kaon and antikaon contributions from the freeze-out surface, the point at which the particles stop interacting. The functions f_K^+ and f_{K^-} are the phase-space distributions of the kaon and antikaon at freeze-out, meaning they describe how many kaons and antikaons exist at a given position $\mathbf{x}_i$, with momentum $\mathbf{p}_i$, at time t_i. The Wigner function ρ^W describes the internal quantum-mechanical structure of the $f_0(980)$ and determines how likely the kaon and antikaon are to form a bound state given their positions and momenta. Finally, the delta function $\delta^3(\mathbf{P} - \mathbf{p}_1 - \mathbf{p}_2)$ simply ensures that momentum is conserved in the process. The statistical factor

$$g = \frac{2J+1}{\prod_{i=1}^{2}(2J_i+1)} = 1 \tag{3.2}$$

accounts for the spin states available to the $f_0(980)$ and its constituents. Since the $f_0(980)$ has spin $J = 0$ and is formed from a kaon and antikaon each with $J_i = 0$, there is only one possible spin configuration, so $g = 1$. The rest-frame coordinates $\mathbf{x}_i^*$ and momenta $\mathbf{p}_i^*$ entering the Wigner function are obtained by first boosting $\mathbf{x}_i$ and $\mathbf{p}_i$ to the rest frame of the $f_0(980)$, and then propagating the earlier freeze-out particle forward in time to t', when the later particle freezes out [1]. Assuming a Gaussian wavefunction for the relative motion of the K and K^-, the Wigner function takes the form [1, 22]

$$\rho^W(\boldsymbol{\rho},\, \mathbf{p}_\rho) = 8\exp\left(-\frac{\rho^2}{\sigma_\rho^2}\right)\exp\left(-p_\rho^2\,\sigma_\rho^2\right), \tag{3.3}$$

where $\boldsymbol{\rho}$ and $\mathbf{p}_\rho$ are the relative coordinate and relative momentum of the kaon and antikaon in the rest frame of the $f_0(980)$, defined as

$$\boldsymbol{\rho} = \frac{1}{\sqrt{2}}(\mathbf{x}_1^* - \mathbf{x}_2^*), \qquad \mathbf{p}_\rho = \frac{1}{\sqrt{2}}(\mathbf{p}_1^* - \mathbf{p}_2^*), \tag{3.4}$$

and σ_ρ is a width parameter that controls the spatial size of the bound state. The two exponentials in Eq (3.3) show that the coalescence probability falls off if the kaon and antikaon are either too far apart in position or too different in momentum. The Gaussian form is a natural choice as it corresponds to the ground-state wavefunction of a harmonic oscillator, and it has the convenient property that its Wigner transform remains Gaussian in both $\boldsymbol{\rho}$ and $\mathbf{p}_\rho$, making Eq (3.3) straightforward to evaluate numerically.

The coalescence probability for a given K^+K^- pair, obtained by evaluating Eq (3.3) at the corresponding relative coordinate and momentum in the $f_0(980)$ rest frame, determines whether the pair forms an $f_0(980)$ in the Monte Carlo implementation described in Section 3.3.

The Gaussian Width Parameter σ_ρ, the RMS Radius, and the Binding Energy

The Gaussian width parameter σ_ρ in Eq (3.3) is directly related to the physical size of the $f_0(980)$. Specifically, it is connected to the root-mean-square (RMS) radius $\sqrt{\langle r^2\rangle}$ by computing the expectation value $\langle r^2\rangle = \int r^2|\psi(r)|^2\, d^3r$ for a spherically symmetric Gaussian wavefunction,

$$\psi(\mathbf{r}) \propto \exp\left(-\frac{r^2}{2\sigma_r^2}\right), \tag{3.5}$$

which gives the relation [1]

$$\sigma_\rho = \frac{2}{\sqrt{3}}\sqrt{\langle r^2\rangle}. \tag{3.6}$$

This means σ_ρ directly sets the spatial size of the $f_0(980)$ under the molecular hypothesis, so estimating the RMS radius is equivalent to fixing σ_ρ.

The RMS radius can be estimated from the binding energy of the state. The binding energy E_B =

$m_1 + m_2 - M_{f_0}$ is the energy required to break the bound state apart into a free kaon and antikaon. Since the kaon masses sum to approximately 987 MeV/c^2 and the PDG mass of the $f_0(980)$ is $M_{f_0} = 990 \pm 20$ MeV/c^2 [18], the $f_0(980)$ sits very close to the K^+K^- threshold and is only very weakly bound. Following Ref. [1], a binding energy of $E_B \approx 10$–20 MeV is applied. From basic quantum mechanics, a weakly bound state has a wavefunction that extends far in space, and this is given by the relation [1, 2]

$$R \sim \frac{1}{\gamma_B}, \qquad \gamma_B = \sqrt{2\mu E_B}, \tag{3.7}$$

where $\mu = m_1 m_2/(m_1 + m_2)$ is the reduced mass of the K^+K^- system, with $m_1 = m_2 = m_{K^\pm} = 493.677 \pm 0.013$ MeV/c^2 [18]. Substituting the binding energy range gives $\sqrt{\langle r^2 \rangle} \approx 1.0$–$1.5$ fm, which is considerably larger than the typical size of a compact $q\bar{q}$ meson ($r_{\mathrm{rms}} \lesssim 0.5$ fm) [2], consistent with the expectation for a spatially extended molecular state. Using Eq (3.6), the corresponding width parameter is $\sigma_\rho \approx 1.15$–1.73 fm.

This constraint is explicitly built into the Wigner function in Eq (3.3), where the second Gaussian suppresses pairs with large relative momentum in the pair rest frame. Producing a high-p_{T} $f_0(980)$ therefore requires two kaons that are both energetic and have small relative momentum in their common rest frame. Although the Wigner suppression itself is Lorentz invariant and does not explicitly disfavour high-p_{T} particles, the underlying single-kaon p_{T} spectrum falls steeply, so finding two suitable kaons becomes increasingly unlikely as p_{T} grows.

It should be noted that the numerical RMS radius range of 1.0–1.5 fm quoted above is taken from Ref. [1], where it is obtained by applying Eq (3.7) to the binding energy interval $E_B \approx 10$–20 MeV. As an independent cross-check, the same calculation was repeated in the present work using the PDG values of the kaon and $f_0(980)$ masses, and the result obtained was instead $\sqrt{\langle r^2 \rangle} \approx 2.0$–$2.8$ fm, which is appreciably larger than the value quoted in the literature. The discrepancy may stem from a different choice of input masses, a different assumed binding energy, or additional model corrections beyond the simple weakly-bound estimate of Eq (3.7); the source of the difference could not be identified unambiguously. The detailed calculation is given in the Appendix 6 for completeness. In the analysis as follows, the literature value is adopted to allow direct comparison with previous work, but this small ambiguity in the input RMS radius should be kept in mind when interpreting the results.

It should be noted that the coalescence treatment described in this thesis is a semi-classical approximation. The model does not enforce the exact microscopic reaction $K^+K^- \to f_0(980)$ with full energy conservation. Instead, it assumes that the small binding energy can be neglected compared with the hadronic mass scale, so that the formation probability is governed mainly by the phase-space proximity of the two kaons. The afterburner, therefore, provides an effective phenomenological estimate of molecular $f_0(980)$ production rather than a complete dynamical calculation [21] .

3.2 Event Generation with PYTHIA

The kaon phase-space distributions $f_{K^+K^-}(x, p, t)$ that enter the coalescence formula, Equation (3.1), are obtained from minimum-bias pp collisions generated with PYTHIA 8 [19], using the hadronization framework described in Section 1.5. Collisions are simulated at $\sqrt{s} = 5.02$ TeV, matching the conditions of the ALICE data used for comparison. The Monash 2013 tune is adopted, which has been optimised specifically for minimum-bias and underlying-event observables at LHC energies and provides a reliable description of light hadron production across the relevant p_{T} range [20]. For each generated event, all stable K^+ and K^- particles are recorded along with their four-momenta, and the invariant yield is extracted in the form defined in Section 2.1.1 for comparison with ALICE measurements prior to the coalescence step.

3.3 Coalescence Afterburner

The coalescence calculation is implemented as an afterburner, meaning it is a post-processing step applied to the kaon four-momenta produced by PYTHIA after each event has been fully generated. The physical justification for this approach is that coalescence can be treated as an independent process. This is because the $f_0(980)$ is assumed to form from kaons that are produced independently after string fragmentation has completed. Separating the event generation from the coalescence step is also computationally convenient. In practice, each K^+K^- pair in an event is tested by the afterburner using the Wigner function overlap condition described in Section 3.1.2. Pairs that satisfy this condition are identified as $f_0(980)$ candidates, and their four-momenta are recorded.

The procedure for each generated event is as follows:

1. All K^+ and K^- particles produced within the rapidity window $|y| < 0.5$, consistent with the ALICE mid-rapidity acceptance, are collected.
2. All distinct K^+K^- pairs are formed.
3. For each pair, the Lorentz transformation of Section 2.2.2 is applied to boost both kaon four-momenta into the rest frame of the candidate $f_0(980)$.
4. The relative momentum $\mathbf{p}_\rho$ is computed in that frame using Eq (3.4), and the Wigner function is evaluated from Eq (3.3).
5. The resulting coalescence probability is used as the acceptance weight in the Monte Carlo procedure of Section 3.3.1.
6. If the pair is accepted, the $f_0(980)$ four-momentum is recorded as $P^\mu = p_1^\mu + p_2^\mu$ and its transverse momentum $p_T^{f_0} = \sqrt{(p_x)^2 + (p_y)^2}$ is added to the output histogram.

After all events have been processed, the histogram is normalised per event to give the yield $\frac{1}{N_{\text{event}}} \frac{dN_{f_0}}{dp_{\text{T}}}$, which is then compared directly to the ALICE $f_0(980)$ data.

3.3.1 Monte Carlo Accept-Reject Sampling

The Wigner function gives a coalescence probability for each K^+K^- pair, but rather than forming the $f_0(980)$ with a fractional weight, the decision is made probabilistically using the Monte Carlo accept-reject method. The idea is for each pair, a random number u is drawn uniformly from the interval $[0, 1]$, and the pair is accepted as an $f_0(980)$ if

$$u \leq P_{\text{coal}}, \tag{3.8}$$

where $P_{\text{coal}} \in [0, 1]$ is the coalescence probability computed from Eq (3.3), and rejected otherwise. This means pairs with a high coalescence probability are accepted most of the time, while pairs with a low probability are rarely accepted. Over a large number of events, the fraction of accepted pairs naturally converges to P_{coal}, correctly reproducing the Wigner function overlap as a formation probability [23].

The statistical uncertainty on the yield in any p_{T} bin is $\sqrt{N_{\text{acc}}}$, where N_{acc} is the number of accepted pairs in that bin. At high p_{T}, both the number of available K^+K^- pairs and the coalescence probability become small, so fewer pairs are accepted and the statistical uncertainty grows. This is why a large event sample is needed to get reliable results across the full p_{T} range.

3.3.2 Application of the Lorentz Boost to the $f_0(980)$ Rest Frame

The coalescence model must be evaluated in the rest frame of the $f_0(980)$ candidate, meaning both the relative momentum and the relative position of the kaon pair need to be computed after boosting into that frame. If these quantities were evaluated in the laboratory frame, the result would depend on the overall motion of the pair rather than its internal structure, making the coalescence probability unphysical [1, 22].

For each K^+K^- pair, the boost velocity is determined from the combined four-momentum of the two kaons,

$$\boldsymbol{\beta} = \frac{\mathbf{p}_1 + \mathbf{p}_2}{E_1 + E_2}, \tag{3.9}$$

where $E_i = \sqrt{|\mathbf{p}_i|^2 + m_K^2}$ is the energy of each kaon, computed using the PDG mass $m_{K^\pm} = 493.677 \pm 0.013$ MeV$/c^2$ [18]. This is simply the velocity the $f_0(980)$ would have in the laboratory frame if the pair coalesced. Both the momentum four-vectors p_i^μ and the space-time position four-vectors $r_i^\mu = (t_i, \mathbf{x}_i)$ of the two kaons are then boosted into this rest frame using the general Lorentz boost implemented in Eqs (2.11)–(2.13). Once in the rest frame, the momentum and position differences are computed as

$$\Delta\mathbf{p} = \mathbf{p}_1^* - \mathbf{p}_2^*, \qquad \Delta\mathbf{r} = \mathbf{x}_1^* - \mathbf{x}_2^*, \tag{3.10}$$

and the relative momentum and coordinate entering the Wigner function are obtained from these via Eq (3.4),

$$|\mathbf{p}_\rho| = \frac{|\Delta\mathbf{p}|}{\sqrt{2}}, \qquad \rho = \frac{|\Delta\mathbf{r}|}{\sqrt{2}}. \tag{3.11}$$

These two quantities are then substituted into the Wigner function, Eq (3.3), to evaluate the coalescence probability, after which the accept-reject step of Section 3.3.1 is applied.

Chapter 4

Results and Discussion

This chapter compares the coalescence calculation with the ALICE measurement of $f_0(980)$ production in pp collisions at $\sqrt{s}$ = 5.02 TeV. The kaon baseline from PYTHIA is checked first, since it provides the input to coalescence models that follows. A simple cut-based model is then discussed to see which features of the production mechanism the data are sensitive to, and the full Wigner-function calculation is presented afterwards, together with a discussion of what the results imply for the internal structure of the $f_0(980)$.

4.1 Kaon Baseline Validation

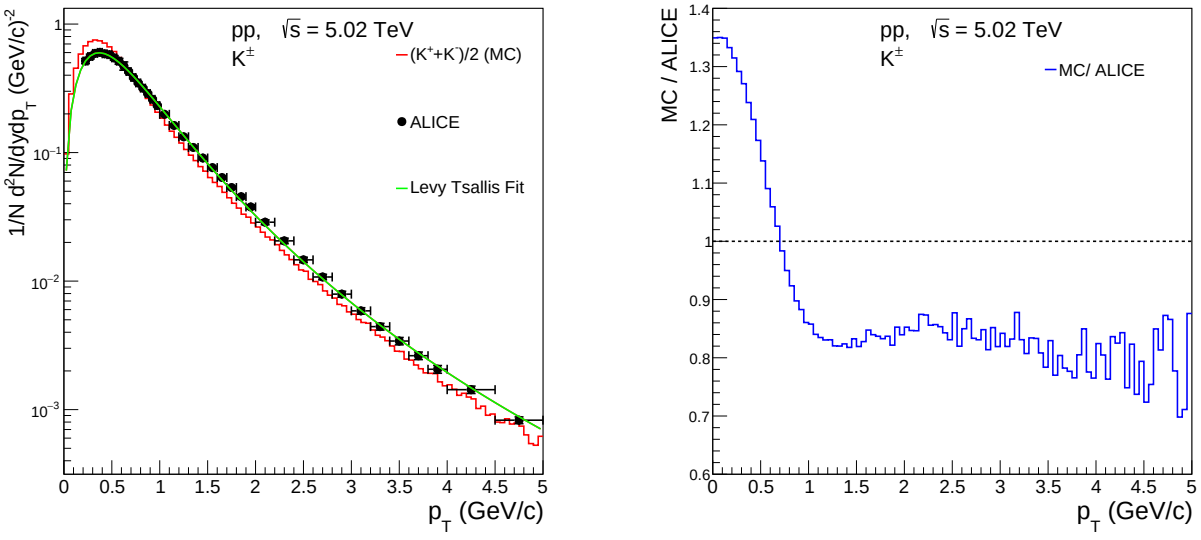

(a) (b)

Figure 4.1: (a) Charged-kaon invariant transverse momentum spectrum from PYTHIA (Monash 2013 tune) compared with ALICE measurements in pp collisions at $\sqrt{s}$ = 5.02 TeV [13]. The green curve is a Lévy–Tsallis fit to the data. (b) Ratio of the PYTHIA spectrum to the Lévy–Tsallis fit.

Before computing the $f_0(980)$ yield, it is important to verify that PYTHIA reproduces the charged-kaon transverse momentum spectrum reasonably well, since the kaon phase-space distributions are the direct input to the coalescence calculation. Figure 4.1(a) shows the charged-kaon invariant yield from PYTHIA compared to ALICE data at $\sqrt{s}$ = 5.02 TeV. The overall shape of the spectrum is well reproduced across the full p_{T} range shown. A Lévy–Tsallis function is fit to the ALICE data to provide a smooth reference curve for the comparison.

The ratio of the PYTHIA spectrum to the Lévy–Tsallis fit, shown in Figure 4.1(b), reveals a p_{T}-dependent deviation. At very low p_{T} (below ~ 0.5 GeV/c), PYTHIA overestimates the kaon yield by up to approximately

35%. At intermediate and high p_{T}, the ratio falls to approximately 0.8–0.9, meaning PYTHIA underestimates the yield by roughly 10–20% in this region. This level of agreement is typical for PYTHIA with the Monash 2013 tune.

To ensure that this kaon-level discrepancy does not propagate into the $f_0(980)$ coalescence result, the ratio shown in Figure 4.1(b) is used as a p_{T}-dependent weight function in the Wigner-function coalescence calculation. This means that any systematic difference between the simulated and measured kaon spectra is corrected for at the kaon level before the coalescence step is applied. As a result, any remaining discrepancy between the predicted $f_0(980)$ yield and the ALICE data would originate from the coalescence model and the molecular hypothesis being tested.

4.2 Simple Cut-Based Coalescence

A coalescence model is first tested in its simplest form, using sharp cutoffs on the relative momentum and spatial separation of each K^+K^- pair. This approach, referred to here as the cut-based coalescence, serves as a useful first test of whether the molecular hypothesis can reproduce the general features of the $f_0(980)$ spectrum before introducing the more physically motivated Gaussian suppression of the Wigner function. Three sets of results are presented: a comparison of the momentum-only and combined momentum-and-spatial criteria at fixed baseline parameters, followed by independent scans of the momentum cutoff p_0 and the source radius r_0 to assess the sensitivity of the prediction to each parameter.

4.2.1 Momentum and Spatial Criteria

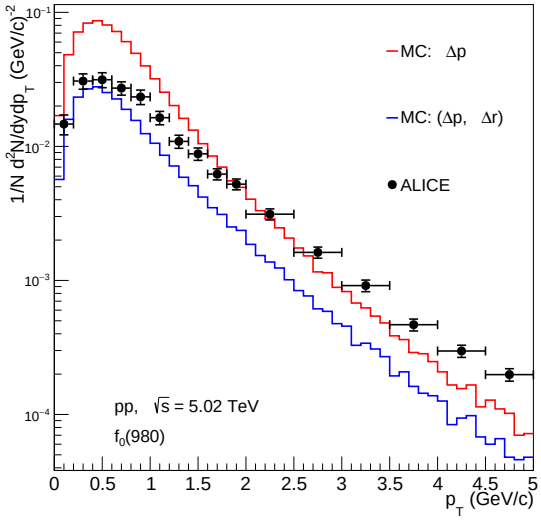

(a)

Figure 4.2: (a) $f_0(980)$ transverse momentum spectrum from the simple coalescence model with the momentum criterion only (red) and with both the momentum and spatial criteria (blue), compared with ALICE data [11]. Baseline values are $p_0 = 0.23$ GeV/c and $r_0 = 1.25$ fm.

As a first step, the coalescence model is applied using simple sharp cutoffs on the relative momentum and spatial separation of each K^+K^- pair, rather than the full Wigner-function prescription. Two cases are compared: one in which only a relative-momentum criterion is applied ($\Delta p < p_0$), and one in which both a momentum and a spatial criterion are applied simultaneously ($\Delta p < p_0$ and $\Delta r < r_0$). The baseline values are $p_0 = 0.23$ GeV/c and $r_0 = 1.25$ fm. The results are shown in Figure 4.2. Both predictions reproduce the overall shape of the $f_0(980)$ spectrum, and both peak at approximately the same transverse momentum of

$p_{\mathrm{T}} \approx 0.4$ GeV/c, in agreement with the ALICE data. This result shows that the K^+K^- coalescence model, even in its simplest version, is able to reproduce the main features of the measured spectrum.

The momentum-only criteria (red curve) produce a significantly higher yield than the combined criteria (blue curve). This is expected: without a spatial condition, any two kaons that happen to have similar momenta can be accepted as an $f_0(980)$ candidate, regardless of how far apart they are produced. Adding the spatial criterion substantially reduces the yield, because it additionally requires the two kaons to be close together in space. This behaviour makes physical sense for a hadronic molecule: the $f_0(980)$, if it is a K^+K^- bound state, must form from kaons that are produced in close spatial proximity. Most kaon pairs in a typical pp collisions are simply too far apart to coalesce into a bound state, and the spatial criterion reflects this localised nature of molecular formation.

At low p_{T}, the ALICE data points lie between the two criteria, sitting below the momentum-only curve and above the combined-criterion curve. The momentum-only cut tends to overestimate the yield because it allows kaon pairs that are spatially too far apart to realistically form a bound state. Introducing the spatial condition reduces this excess, but the sharp cutoff at $r_0 = 1.25$ fm may be slightly too strict, which could explain why the combined prediction undershoots the data at low p_{T}. Both models sit on either side of the data in this region, so the real formation probability is likely somewhere in between. At higher p_{T} (above about 2 GeV/c), both models begin to underestimate the data, with the predicted yields falling below the ALICE measurements. This behaviour appears in both the momentum-only and combined cases, indicating that the deficit is not driven by the spatial cut, but is instead a more general feature of the coalescence framework used here.

This observation motivates the use of the Wigner-function approach, which replaces the sharp cutoffs with a smooth Gaussian wave function that provides a more realistic description of the formation probability.

4.2.2 Relative-Momentum Scan

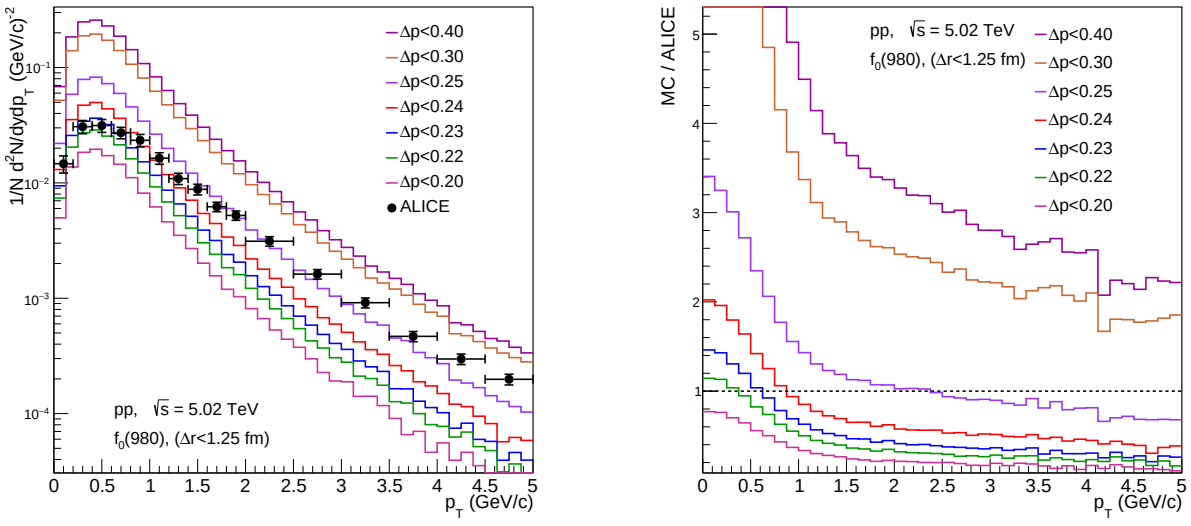

(a) (b)

Figure 4.3: (a) Scan of the relative-momentum cutoff p_0, with the spatial cutoff fixed at $\Delta r < 1.25$ fm. Larger values of p_0 admit pairs with larger relative momentum and so raise the predicted yield. (b) Ratio of each MC prediction to the Lévy–Tsallis fit of the ALICE data [11]. The dashed line at unity indicates perfect agreement.

To study the sensitivity of the prediction to the choice of momentum cutoff, p_0 is varied from 0.20 to 0.40 GeV/c while the spatial cutoff is held fixed at $\Delta r < 1.25$ fm. The results are shown in Figure 4.3.

As p_0 is increased, the yield grows and the spectrum extends to higher p_T. A larger momentum cutoff admits kaon pairs with larger relative momentum, which tend to be produced at larger pair p_T, so the high-p_T end of the spectrum becomes more populated. At low p_T, the spectra for different values of p_0 are already well separated, while at high p_T the differences become even more emphasised.

A brief note on the physics of the cutoff is useful here. The relative momentum p_0 between the two kaons in the pair rest frame is related to the binding energy of the molecular state. A loosely bound state, like the $f_0(980)$ under the molecular hypothesis, has a small binding energy ($E_B \approx 10$–20 MeV), meaning the kaons inside it can only have a small relative kinetic energy. This sets a natural physical scale for p_0: accepting pairs with very large relative momentum is unphysical, because such pairs may not have the internal kinematics consistent with forming a loosely bound molecule.

The ratio panel in Figure 4.3(b) shows how well each value of p_0 reproduces the ALICE data. At low p_T up to approximately 1.5 GeV/c, the value $p_0 = 0.23$ GeV/c gives the closest agreement, with the ratio staying within approximately 10–20% of unity. Above $p_T \approx 1.5$ GeV/c, the ratio for $p_0 = 0.23$ falls below unity and the slightly larger value $p_0 = 0.25$ GeV/c sits closer to the data. The best overall agreement across the full range is therefore achieved for p_0 in the range 0.23–0.25 GeV/c. Despite this, all values of p_0 increasingly underestimate the data at high p_T, with ratios dropping to 50% or below beyond $p_T \approx 2$ GeV/c, and no single cutoff value is able to simultaneously describe both the low- and high-p_T regions well.

This limitation is a direct consequence of using a sharp cutoff. A real bound state does not have a hard boundary in relative momentum: its quantum-mechanical wavefunction smoothly suppresses contributions from fast-moving pairs rather than excluding them entirely. The Wigner-function approach, described in Section 4.3, addresses this by replacing the sharp cut with a smooth Gaussian wave function, which is more physically motivated.

4.2.3 Source Radius Scan

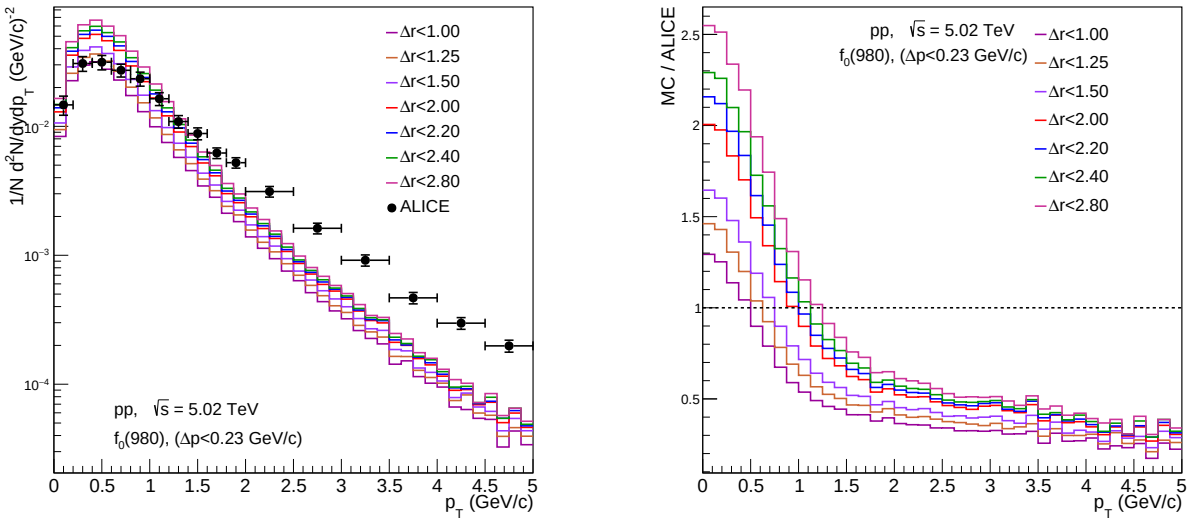

(a) (b)

Figure 4.4: (a) Scan of the source radius r_0, with the relative-momentum cutoff fixed at $\Delta p < 0.23$ GeV/c. Larger values of r_0 admit pairs that are further apart in space and so raise the predicted yield, with the effect most visible at low p_T. (b) Ratio of each MC prediction to the Lévy–Tsallis fit of the ALICE data [11]. The dashed line at unity indicates perfect agreement.

The sensitivity of the prediction to the spatial cutoff r_0 is investigated by varying it from 1.0 to 2.8 fm while keeping the momentum cutoff fixed at $\Delta p < 0.23$ GeV/c. The results are shown in Figure 4.4.

The values $r_0 = 1.0$–1.5 fm correspond to the literature estimate of the K^+K^- RMS radius for the $f_0(980)$ [1], derived from the binding energy via Eq (3.7). The larger values, $r_0 = 2.0$–2.8 fm, cover the alternative range obtained from the independent cross-check calculation described in Section 3.1.2. Using both ranges therefore covers the full range of possible spatial sizes suggested by the two estimates.

As r_0 increases, the predicted yield grows, most noticeably at low p_T. A larger spatial cutoff admits kaon pairs that are further apart, so more pairs are accepted and the total yield increases. The effect is most visible at low p_T because soft kaons, which make up the bulk of the spectrum, are more spread out in space on average. Looking at the ratio panel in Figure 4.4(b), the spread between the smallest and largest r_0 values at low p_T amounts to a factor of roughly 2–2.5, while at high p_T all curves converge toward each other and all still fall significantly below the data, with ratios dropping to around 20–50% of the measured yield beyond $p_T \approx 2$ GeV/c.

This result shows two main things. First, the p_T spectrum alone does not give a precise way to determine the source radius, since the predicted spectra have very similar shapes for all reasonable values of r_0. Second, the underestimation at high p_T is not due to a specific choice of r_0, since it appears for all tested values, making it a stable feature of the model.

Taken together, the three scans show that the simple cut-based coalescence model is able to reproduce the low- to intermediate-p_T yield of the $f_0(980)$ for physically reasonable parameter choices. The model consistently underestimates the high-p_T region, and this behaviour does not change when varying the momentum or spatial cutoffs within the tested ranges. These observations motivate the use of the more physically motivated Wigner-function approach, which provides a smoother and more realistic description of how kaon pairs form a bound state in phase space.

4.3 $f_0(980)$ Spectrum from Wigner-Function Coalescence

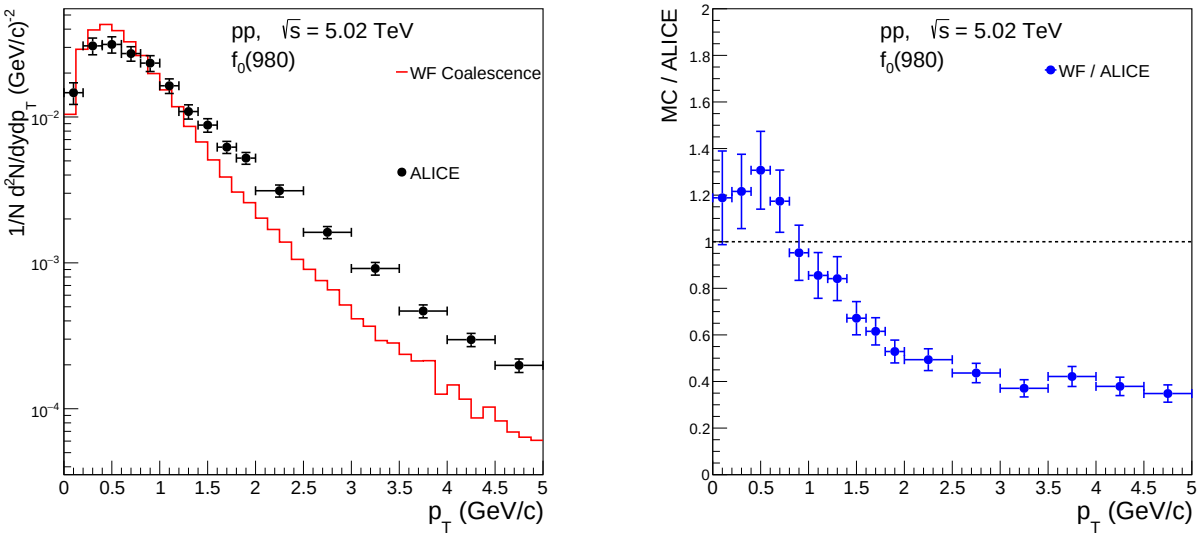

(a) (b)

Figure 4.5: (a) $f_0(980)$ transverse momentum spectrum from Wigner-function coalescence (red) compared with ALICE data [11]. (b) Ratio of the simulation to a Lévy parameterisation of the data.

The Wigner-function coalescence replaces the sharp momentum and spatial cutoffs of the previous section with a smooth Gaussian wave function, as described by Eq (3.3). The width parameter is fixed to $\sigma_\rho = 1.44$ fm, corresponding to an RMS radius of approximately 1.25 fm via Eq (3.6). The results are shown in Figure 4.5.

The Wigner-function prediction describes the ALICE data well at low p_T, roughly up to $p_T \approx 1$ GeV/c.

The ratio panel in Figure 4.5(b) shows values of approximately 1.2–1.3 in the lowest p_{T} bins, meaning the model slightly overestimates the yield there by about 20–30%, though the data uncertainties are also largest in this region. Around $p_{\mathrm{T}} \approx 0.5$–1 GeV/$c$ the ratio approaches unity, indicating good agreement with the data at the level of $\sim 10\%$ or better. This is a meaningful result: a coalescence model assuming a K^+K^- molecular state, with a physically motivated Gaussian wavefunction, is able to reproduce the measured $f_0(980)$ production rate at low transverse momentum. The overall spectral shape rising to a peak near $p_{\mathrm{T}} \approx 0.4$ GeV/c and falling steeply thereafter is well captured. Fig. 2 of Ref. [11] shows a comparison of the AMPT model as input with a K^+K^- molecule treatment [12] against the experimental data. The coalescence model therein also uses a Gaussian wavefunction and is found to be two orders of magnitude below the data. The coalescence afterburner developed in this thesis, with input from Pythia, is found to be much closer to the data, especially at low p_{T}. At higher p_{T}, the deviation observed in the AMPT-based coalescence is similar to that of the coalescence afterburner presented in this thesis.

The improved performance at low p_{T} compared to the sharp-cut model can be understood physically. The Wigner function suppresses kaon pairs with relative momentum much larger than $\hbar/\sigma_\rho$. For $\sigma_\rho = 1.44$ fm, the momentum scale $\hbar/\sigma_\rho$ is approximately 140 MeV/c. This is smaller than the sharp cutoff $p_0 = 230$ MeV/c used in the simple cut-based model, and is more consistent with the expectation for a spatially extended, weakly bound K^+K^- configuration. The corresponding relative kinetic energy is of order a few tens of MeV, comparable to the expected molecular binding-energy scale for $f_0(980)$ formation. The Wigner function therefore provides a smoother and more physically motivated selection of kaon pairs than an abrupt hard cutoff.

Above $p_{\mathrm{T}} \approx 1$ GeV/c, the prediction increasingly falls below the data. By $p_{\mathrm{T}} \sim 2$ GeV/c the ratio has dropped to approximately $0.6 - 0.5$, and by $p_{\mathrm{T}} \sim 3$ GeV/c it falls to around 0.4, meaning the model underestimates the yield by a factor of roughly 2.5 at this point. This is not a failure of the calculation itself, but may be a natural consequence of the molecular picture. Producing an $f_0(980)$ at high p_{T} requires two kaons that are both highly energetic and moving in nearly the same direction, so that their relative momentum in the pair rest frame remains small enough to satisfy the Wigner-function condition. In pp collisions, kaons are produced from many independent string fragmentations, and finding two such well-matched kaons becomes increasingly unlikely as p_{T} grows. The steeply falling single-kaon p_{T} spectrum makes this ever more rare. The suppression at high p_{T} is therefore a direct prediction of the molecular picture: a loosely bound, spatially extended state cannot easily be produced at high transverse momentum.

The overall conclusion from Figure 4.5 is that the Wigner-function coalescence model gives a good description of the $f_0(980)$ spectrum at low p_{T}, which is consistent with the K^+K^- molecular hypothesis. The model was designed to test whether kaon coalescence can reproduce the observed magnitude and spectral shape, and the answer is yes for the magnitude and for the low-p_{T} region where the molecular formation mechanism is most relevant.

4.4 Implications for the $f_0(980)$ Internal Structure

The results presented in the previous sections show that the K^+K^- coalescence model gives a good description of the $f_0(980)$ transverse momentum spectrum at low p_{T}. It demonstrates that kaon coalescence, under the molecular hypothesis, is consistent with the measured production properties in this kinematic region. The overall spectral shape is reproduced, and the model parameters used in particular the RMS radius of approximately 1.25 fm are physically motivated by the binding energy of the state.

At high p_{T}, the prediction falls below the data. As discussed in Section 4.3, this may be expected within the molecular picture due to the decreasing probability of finding two kaons with sufficiently similar momenta at high p_{T}. However, the ALICE data extend to higher p_{T} with a flatter slope than the coalescence prediction, suggesting that there may be an additional production component contributing at high p_{T} that is not captured by the molecular model alone.

A natural candidate for this additional component is a more compact configuration of the $f_0(980)$, such as a $q\bar{q}$ or tetraquark state. Compact states have spatial sizes well below one femtometre and are produced through partonic fragmentation rather than hadronic coalescence [2, 6]. Because they are formed at the partonic level through string fragmentation and/or partonic coalescence, they are not subject to the same phase-space suppression at high p_{T} as extended molecular states. Their harder p_{T} spectrum could therefore account for the excess seen in the data beyond $p_{\mathrm{T}} \approx 1$–2 GeV/c.

The picture that emerges is therefore one in which the $f_0(980)$ may not be a pure state of a single configuration. At low p_{T}, the molecular K^+K^- component dominates and is well described by the coalescence model. At high p_{T}, a compact component may contribute additional yield. This is consistent with the broader discussion in the literature, where the $f_0(980)$ is often described as having a mixed internal structure [1, 6]. The results of this work support that interpretation: the coalescence model successfully reproduces the low-p_{T} part of the spectrum, which can be an evidence for a molecular component, while the high-p_{T} excess points to the possible presence of a more compact contribution.

It should be emphasised that the primary goal of this project was to test whether a K^+K^- coalescence model can reproduce the magnitude and overall spectral shape of the $f_0(980)$, and the answer is yes. The model gives a reasonable and physically motivated description of the data across a wide p_{T} range, and the deviations observed at high p_{T} could be due to limitation of molecular hypothesis.

Chapter 5

Conclusion and Outlook

This thesis investigated whether the $f_0(980)$ scalar meson can be described as a K^+K^- hadronic molecule by implementing a coalescence model and comparing its predictions with ALICE measurements in pp collisions at $\sqrt{s} = 5.02$ TeV. A p_T-dependent correction derived from the kaon baseline comparison was applied to ensure that any remaining discrepancy in the $f_0(980)$ prediction originates from the coalescence model itself rather than from the event generator. Both the cut-based and Wigner-function coalescence models reproduced the overall shape and approximate magnitude of the measured spectrum at low p_T, with the peak near $p_T \approx 0.4$ GeV/c correctly captured in both cases. The Wigner-function calculation with $\sigma_\rho = 1.44$ fm provided the most realistic description, with the ratio to the ALICE data approaching unity around $p_T \approx 0.5$–1 GeV/c, supporting the idea that the $f_0(980)$ contains a significant K^+K^- molecular component in this region.

At higher p_T, both models increasingly underestimated the data, with the Wigner-function ratio falling to approximately 0.4 by $p_T \sim 3$ GeV/c. It is worth noting, however, that the coalescence of K^+K^- pairs still reproduces the data within a similar order of magnitude across the full p_T range, which is itself a non-trivial result. The growing discrepancy at high p_T may nevertheless be a natural consequence of the molecular picture, since a loosely bound spatially extended state can only form from kaons with very similar momenta, which becomes increasingly unlikely at high p_T. The source-radius scan also showed that the p_T spectrum alone cannot clearly distinguish between source sizes in the range 1.0–2.8 fm. Overall, the results are consistent with a mixed internal structure: the low-p_T region is well described by the K^+K^- molecular component, while the excess at high p_T may point towards a more compact contribution such as a $q\bar{q}$ state or tetraquark, produced through partonic fragmentation rather than hadronic coalescence.

As an outlook, a natural next step would be to test other internal structures, a $q\bar{q}$ meson or tetraquark within the same coalescence framework, and compare the predicted p_T spectra with each other and with data.

Extending the analysis to Pb–Pb collisions would also be valuable. The nuclear modification factor R_{AA} could help distinguish a loosely bound molecular state, which would be strongly suppressed in the quark–gluon plasma, from a more compact structure. Elliptic flow measurements of v_2 could further constrain the formation mechanism: a molecular state formed after freeze-out should inherit kaon collective flow, while a compact state formed at the partonic level would not.

Finally, studying $f_0(980)$ production as a function of event multiplicity in pp collisions would provide an independent test. A large molecular component would lead to enhanced production in high-multiplicity events, where more kaons are produced in a smaller volume, making coalescence more likely.

Bibliography

[1] Y. Wang *et al.*, "Probing the structure of $f_0(980)$ from the elliptic flow in p–Pb collisions at the LHC," arXiv:2511.03638 (2025).

[2] F. K. Guo *et al.*, "Hadronic molecules," *Rev. Mod. Phys.* **90**, 015004 (2018); see also S. L. Olsen, T. Skwarnicki and D. Schlatter, "Nonstandard heavy mesons and baryons: Experimental evidence," *Rev. Mod. Phys.* **90**, 015003 (2018).

[3] H. D. Politzer, "Reliable perturbative results for strong interactions?," *Phys. Rev. Lett.* **30**, 1346 (1973).

[4] J. D. Jackson, *Classical Electrodynamics*, 3rd ed. (John Wiley & Sons, New York, 1999), Chaps. 11–12.

[5] J. R. Taylor, *An Introduction to Error Analysis: The Study of Uncertainties in Physical Measurements*, 2nd ed. (University Science Books, Sausalito, 1997), Chap. 3.

[6] Z. Wang *et al.*, "Theoretical study of $f_0(980)$, $a_0(980)$ and $\Xi(1/2^-)$," arXiv (2023).

[7] A. Gu and F. Wang, "Transverse momentum spectra of $f_0(980)$ from coalescence model," *Phys. Lett. B* **846**, 138399 (2023), arXiv:2306.08584.

[8] A. Gu, T. Edmonds, J. Zhao and F. Wang, "Elliptic flow coalescence to identify the $f_0(980)$ content," *Phys. Rev. C* **101**, 024908 (2020), arXiv:1902.07152.

[9] F. Reif, *Fundamentals of Statistical and Thermal Physics* (McGraw-Hill, New York, 1965), Chap. 12.

[10] A. Bellini *et al.*, "A realistic coalescence model for deuteron production," *Eur. Phys. J. C* **83**, 879 (2023), arXiv:2302.12696.

[11] ALICE Collaboration, "$f_0(980)$ production in inelastic pp collisions at $\sqrt{s} = 5.02$ TeV," *Phys. Lett. B* **835**, 137644 (2022), arXiv:2206.06216.

[12] A. Gu, T. Edmonds, J. Zhao and F. Wang, "Elliptical flow coalescence to identify the $f_0(980)$ content," Phys. Rev. C **101** (2020) no.2, 024908

[13] S. Acharya *et al.* (ALICE Collaboration), "Production of charged pions, kaons and (anti-)protons in Pb–Pb and inelastic pp collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV," *Phys. Rev. C* **101**, 044907 (2020), arXiv:1910.07678.

[14] S.-K. Choi *et al.* (Belle Collaboration), "Observation of a narrow charmonium-like state in exclusive $B^\pm \to K^\pm \pi^+ \pi^- J/\psi$ decays," *Phys. Rev. Lett.* **91**, 262001 (2003).

[15] R. Aaij *et al.* (LHCb Collaboration), "Observation of $J/\psi\, p$ resonances consistent with pentaquark states in $\Lambda_b^0 \to J/\psi\, K^- p$ decays," *Phys. Rev. Lett.* **115**, 072001 (2015).

[16] D. H. Rischke, "QCD and strongly coupled gauge theories," *Prog. Part. Nucl. Phys.* (2015), `https://pmc.ncbi.nlm.nih.gov/articles/PMC4413533/`.

[17] J. A. Oller and E. Oset, "The nature of the $f_0(980)$ and $a_0(980)$ resonances from meson-meson scattering," arXiv:0709.3366 [hep-ph] (2007), `https://arxiv.org/abs/0709.3366`.

[18] R. L. Workman *et al.* (Particle Data Group), "Review of Particle Physics," *Prog. Theor. Exp. Phys.* **2022**, 083C01 (2022).

[19] T. Sjöstrand *et al.*, "An introduction to PYTHIA 8.2," *Comput. Phys. Commun.* **191**, 159 (2015), arXiv:1410.3012.

[20] P. Skands, S. Carrazza and J. Rojo, "Tuning PYTHIA 8.1: the Monash 2013 tune," *Eur. Phys. J. C* **74**, 3024 (2014), arXiv:1404.5630.

[21] M. Mahlein, L. Barioglio, F. Bellini, L. Fabbietti, C. Pinto, B. Singh and S. Tripathy, "A realistic coalescence model for deuteron production," Eur. Phys. J. C **83** (2023) 804

[22] R. Mattiello, H. Sorge, H. Stöcker and W. Greiner, "Nuclear clusters as a probe for expansion flow in heavy-ion collisions at 10–15 A GeV," *Phys. Rev. C* **55**, 1443 (1997).

[23] C. P. Robert and G. Casella, "Monte Carlo Statistical Methods," Springer, New York (2004).

Chapter 6

Appendix

This appendix gives the calculation referred to in Section 3.1.2.

Modelling the $f_0(980)$ as a non-relativistic two-body bound state of a kaon and antikaon, the relative motion is described by

$$-\frac{\hbar^2}{2\mu}\nabla^2\psi(\mathbf{r}) = -E_B\,\psi(\mathbf{r}), \tag{6.1}$$

with reduced mass $\mu = m_K/2$ and binding energy $E_B > 0$. The asymptotic s-wave solution is $\psi(r) \propto e^{-\gamma r}$ with $\gamma = \sqrt{2\mu E_B}$, giving an RMS radius

$$R_{\mathrm{RMS}} = \sqrt{\langle r^2\rangle} \sim \frac{1}{\gamma} = \frac{1}{\sqrt{2\mu E_B}}. \tag{6.2}$$

Using $m_K \approx 495$ MeV, $\mu \approx 247.5$ MeV, and $E_B \approx 10$–20 MeV,

$$\gamma_{(10\,\mathrm{MeV})} \approx 70.4\ \mathrm{MeV}, \qquad \gamma_{(20\,\mathrm{MeV})} \approx 99.5\ \mathrm{MeV}. \tag{6.3}$$

Converting with $\hbar c = 197$ MeV $\cdot$ fm,

$$R_{(10\,\mathrm{MeV})} \approx 2.8\ \mathrm{fm}, \qquad R_{(20\,\mathrm{MeV})} \approx 2.0\ \mathrm{fm}, \tag{6.4}$$

giving $R_{\mathrm{RMS}} \approx 2.0$–$2.8$ fm, which is larger than the 1.0–1.5 fm range quoted in Ref. [1].