

**Low-Mass Dielectron Production in
Pb–Pb Collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV and
Feasibility of a QGP-Temperature Measurement**

Dissertation

zur Erlangung des Doktorgrades
der Naturwissenschaften

vorgelegt beim Fachbereich Physik
der Johann Wolfgang Goethe-Universität
in Frankfurt am Main

von

Patrick Simon Reichelt
aus Frankfurt am Main

Frankfurt am Main 2021

(D 30)

Vom Fachbereich Physik der Johann Wolfgang Goethe-Universität
als Dissertation angenommen.

Dekan: Prof. Dr. Harald Appelshäuser

Gutachter: Prof. Dr. Harald Appelshäuser

Prof. Dr. Christoph Blume

Datum der Disputation:

Zusammenfassung

Der Forschungsschwerpunkt im Feld der Hochenergie-Schwerionenphysik liegt in der Studie des Quark–Gluon Plasmas (QGP). Dieses stellt einen extrem heißen Materiezustand dar, in dem sich die fundamentalen Bausteine stark wechselwirkender Materie quasi-frei bewegen können. Auf natürliche Weise waren die dafür nötigen Bedingungen lediglich in den ersten Mikrosekunden der Evolution des Universums erfüllt, bevor sich dieses aufgrund seiner Expansion abkühlte und die Materie einen Phasenübergang vom Quark–Gluon Plasma zum sogenannten Hadrongas vollzog. Dabei formten Quarks und Gluonen gebundene Zustände namens Hadronen, welche die Grundbausteine aller bekannten, im weiteren Verlauf der Evolution gebildeten Materie darstellen.

Zur experimentellen Untersuchung des Quark–Gluon Plasmas muss dieses im Labor erzeugt werden, indem schwere Ionen — zum Beispiel Blei-Kerne — auf ultrarelativistische Energien beschleunigt und dann zur Kollision gebracht werden. Aus der dabei freigesetzten Energie werden neue Teilchen und Antiteilchen erzeugt, die ein System hoher Temperatur und Dichte bilden, welches sich dann ähnlich dem frühen Universum entwickelt. Die Eigenschaften dieses Systems lassen sich aufgrund seiner Kurzlebigkeit von nur etwa 10^{-23} s ausschließlich indirekt anhand der finalen Reaktionsprodukte untersuchen. Dabei unterscheiden sich die vielfältigen verfügbaren Observablen insbesondere darin, ob sie auf stark wechselwirkenden Teilchen beruhen, die nach ihrer Erzeugung vom weiteren Kollisionsverlauf beeinflusst werden, oder auf sogenannten elektromagnetischen Sonden, die das heiße und dichte Medium ohne nennenswerte Interaktionen verlassen.

Thema der vorliegenden Arbeit ist die Messung von Elektron–Positron Paaren — kurz Dielektronen genannt —, welche zur zweitgenannten Kategorie gehören. Deren Analyse lässt somit direkte Rückschlüsse auf zentrale Eigenschaften des gebildeten Systems zu, sofern die experimentellen Daten eine ausreichende Präzision aufweisen. Insbesondere gilt die Messung der initialen Temperatur des Quark–Gluon Plasmas als ein “heiliger Gral” der Schwerionenphysik. Im Rahmen dieser Arbeit wurde deshalb zusätzlich zur Analyse bestehender Daten auch eine Machbarkeitsstudie durchgeführt, inwieweit dieses Ziel durch ein Upgrade des ALICE Experimentes am CERN erreichbar wird.

Schwerionenkollisionen lassen sich als Pfad im Phasendiagramm stark wechselwirkender Materie darstellen, indem der Verlauf der mittleren Temperatur und Dichte des produzierten Teilchensystems berechnet wird. Die höchsten Werte werden im Moment der initialen Nukleonenstöße erreicht, noch bevor die anschließende Thermalisierung und somit tatsächliche Bildung eines Quark–Gluon Plasmas einsetzt. Während der QGP-Phase erfolgt eine kollektive Expansion des Mediums, bis es auf die sogenannte kritische Temperatur abkühlt, bei der Quarks und Gluonen wieder zu Hadronen rekombinieren. Für hohe Kollisionsenergien liegt ein *Crossover*-Phasenübergang ohne abrupte Änderung der Systemeigenschaften vor, der auch nahezu mit dem chemischen *Freeze-out* koinzidiert, sodass die entstehenden Hadronen nur noch weitgehend elastisch interagieren, bevor sie frei zu den Teilchendetektoren der Experimente fliegen.

Dielektronen werden in allen Kollisionsstadien produziert und stellen eine vielversprechende Sonde dar, weil ihre invariante Masse die zum Zeitpunkt ihrer Erzeugung verfügbare Energie widerspiegelt. Somit werden die massereichsten Dielektronen in den initialen Stößen durch Quark–Antiquark Annihilation erzeugt. Auch aus den Zerfallsreihen früh produzierter Paare schwerer Quarks, die als *Heavy-Flavours* bezeichnet werden, können korrelierte Elektron–Positron Paare hervorgehen, die aufgrund der variablen Zerfallskinetik ihrer Mutterteilchen den gesamten Bereich niedriger bis hoher invarianter Masse abdecken. Da diese schweren Quarks das um sie herum entstandene Quark–Gluon Plasma durchqueren müssen, beeinflussen deren Streuung und Energieverlust im Medium diese Komponente des Dielektronenspektrums. Während das System sein QGP-Stadium durchläuft, emittiert es thermische Strahlung in Form von Photonen und Dileptonen¹, welche Informationen über die vorliegende Temperatur zum Beobachter transportieren. Nahe des Phasenübergangs wird auch im Hadrongas noch thermische Dileptonstrahlung produziert, die sensibel auf Modifikationen der elektromagnetischen Spektralfunktion aufgrund von Medium-Effekten ist. Im Endstadium der Kollision führen Zerfälle von *Light-Flavour* Mesonen², die bis zum *Freeze-out* produziert wurden, zu einem Cocktail zusätzlicher Beiträge zum Dielektronenspektrum.

Experimente bei einer Schwerpunktsenergie von $\sqrt{s_{NN}} = 17,3$ GeV konnten erstmals die Entstehung eines Quark–Gluon Plasmas nachweisen. Ihre Messungen des ρ -Mesons über seinen dileptonischen Zerfall belegen auch die *In-Medium* Modifikation der elek-

¹Dilepton ist ein Überbegriff für Elektron–Positron sowie Muon–Antimuon Paare, die gleichermaßen produziert werden können, sofern die verfügbare Energie hoch genug ist.

²Mesonen gehören zur Gruppe der Hadronen und bestehen aus einem Quark und einem Antiquark.

tromagnetischen Spektralfunktion. Nachfolgende Experimente bei $\sqrt{s_{\text{NN}}} = 200$ GeV charakterisierten das Quark–Gluon Plasma als nahezu perfekte Flüssigkeit, doch thermische Dileptonstrahlung aus dem QGP-Stadium konnte bisher nicht von den anderen Beiträgen separiert werden.

Die vorliegende Arbeit basiert auf frühen Daten des ALICE Experimentes, die von Blei–Blei Kollisionen bei einer Schwerpunktsenergie von $\sqrt{s_{\text{NN}}} = 2,76$ TeV aufgezeichnet wurden. ALICE ist das einzige dedizierte Schwerionenexperiment am *Large Hadron Collider* des CERN und ist daher darauf ausgelegt, eine Vielzahl an Observablen untersuchen zu können. Die verschiedenen Detektoren des Experimentes ermöglichen insbesondere die Identifizierung unterschiedlicher Spezies von Teilchen sowie eine präzise Rekonstruktion ihrer Trajektorien, wobei ein Fokus auf Teilchen mit niedrigen Impulsen besteht. Dies sind essentielle Voraussetzungen für die Analyse von Dielektronen niedriger invarianter Masse, da sie die Messung von Elektronen und Positronen mit minimalen Transversalimpulsen von nur $p_{\text{T}} \approx 0,1$ GeV/c erlauben.

Die in dieser Arbeit analysierten Daten beinhalten jeweils knapp 20 Millionen Ereignisse zentraler sowie semi-zentraler Blei–Blei Kollisionen. Da dies an der unteren Grenze für eine aussagekräftige Dielektronenmessung liegt, wurde während der Analyse, die in Kapitel 4 ausführlich beschrieben wird, ein Fokus auf das Erreichen hoher Effizienzen gelegt — besonders bei der Selektion von Elektron- und Positron-Kandidaten aus den detektierten Teilchenspuren. Dazu wurde eine spezielle Strategie zur Elektronidentifizierung entwickelt, die das Potential dreier Detektoren des Experiments, nämlich *Inner Tracking System*, *Time Projection Chamber* und *Time-Of-Flight*, ideal kombiniert. Zudem wurde die Spurselektion so angepasst, dass vorliegende Detektorprobleme weitgehend kompensiert und somit etwa 50 % mehr Teilchen ausgewählt werden konnten. In Summe wurde in dieser Analyse eine Paareffizienz von über 10 % erreicht, im Vergleich zu nur etwa 1 % in einer vorangegangenen Analyse des Datensatzes.

Das Dielektronenspektrum wurde auf statistischer Basis ausgewertet, indem pro Kollisionsereignis sämtliche Paarkombinationen zwischen Elektronen und Positronen gebildet wurden. Sie beinhalten das echte Signal sowie den kombinatorischen Untergrund, welcher in dieser Messung bis zu tausendfach größer ist. Zwei Kriterien zur Bestimmung der Signalqualität, nämlich das Signal-zu-Untergrund Verhältnis und die Signifikanz, wurden zur Optimierung einer Vorabfilterung von Paaren mit bestimmten Merkmalen verwendet, um den Anteil von Elektronen und Positronen aus unerwünschten Quellen und somit auch den kombinatorischen Untergrund zu reduzieren. Zudem wurde ein Artefakt des Spurrekonstruktionsalgorithmus ausgenutzt, um Paare aus Photonkonversionen im Detektormaterial zu verwerfen und um eine Korrektur des Dielektronenspektrums bei niedriger

invarianter Masse vorzunehmen. Zur Studie systematischer Unsicherheiten in den Daten wurde die Analyse für 20 stark variierende Wertekombinationen der Elektronselektionsparameter durchgeführt. Der große Bereich daraus resultierender Paareffizienzen zwischen 2 und 20 % und die damit verbundene Variation des Signal-zu-Untergrund Verhältnisses decken sowohl die Unsicherheiten auf Spurebene als auch diejenigen der Signalextraktion ab. Die durchschnittliche Abweichung der 20 Ergebnisse von ihrem Mittelwert wurde als Unsicherheit des Signals angenommen und beträgt etwa $\pm 20\%$ im Großteil des studierten Masse- und Impulsbereiches.

Die Ergebnisse der Datenanalyse werden in Kapitel 5 dieser Arbeit weiter diskutiert und mit verschiedenen Modellrechnungen verglichen. Die gemessene Dielektronenproduktion als Funktion der invarianten Masse und des Paar-Transversalimpulses kann bereits weitgehend durch die oben genannten *Light-Flavour* und *Heavy-Flavour* Beiträge, die den hadronischen Cocktail bilden, beschrieben werden. Nur bei niedriger Masse lässt sich eine geringe Erhöhung in den Daten erahnen, sowohl in den 0–10 % zentralsten als auch in 10–50 % semi-zentralen Kollisionen. Der initiale Vergleich zwischen Daten und Cocktail dient als Basis für die Quantifizierung der Modifikationen und der zusätzlichen Beiträge des Dielektronenspektrums, die in Blei–Blei Kollisionen im Vergleich zu Proton–Proton Kollisionen erwartet werden. In den vorgestellten Ergebnissen entsprechen die *Light-Flavour* Beiträge bereits dem Kenntnisstand der Bedingungen in Blei–Blei Kollisionen, wobei verbleibende Annahmen von den mit ihnen assoziierten systematischen Unsicherheiten abgedeckt werden. Der *Heavy-Flavour* Anteil hingegen entspricht zunächst einer einfachen Skalierung der in Proton–Proton Kollisionen gemessenen Ausbeute mit der Anzahl individueller Nukleonenstöße innerhalb einer Blei–Blei Kollision, wodurch aufgrund der oben genannten Medium-Effekte die tatsächliche Ausbeute überschätzt wird.

Im Rahmen dieser Arbeit werden zwei phenomenologische Ansätze zur Modellierung der *Heavy-Flavour* Unterdrückung durch Streuung und Energieverlust im Medium verfolgt. Zunächst wird die von ALICE gemessene nukleare Modifikation der Impulsverteilung einzelner Elektronen aus diesen Quellen in die Berechnung des *Heavy-Flavour* Beitrages zum Cocktail einbezogen, was zu einer Unterdrückung um etwa 50 % im Bereich mittlerer invarianten Masse führt, wo dieser Beitrag am relevantesten ist. Der zweite Ansatz ist unabhängig von externen Eingaben, indem lediglich die Winkelverteilungen der Elektronen auf diverse Art verschmiert werden, um die Streuung ihrer Mutterteilchen im Medium zu simulieren. Je nach Ausprägung werden in diesem Szenario 10–40 % Unterdrückung erreicht. Die in jedem dieser Fälle reduzierte und somit auch realistischere *Heavy-Flavour* Ausbeute schafft etwas mehr Raum für einen thermischen Beitrag zum

Dielektronenspektrum bei niedriger und mittlerer invarianter Masse.

Ein vollständiger Vergleich zwischen experimentellen Daten und theoretischer Erwartung ist nur durch die Einbeziehung von Modellrechnungen möglich. Dazu werden theoretische Modelle mit zwei unterschiedlichen Ansätzen vorgestellt und jeweils ein prominenter Vertreter seiner Kategorie herangezogen. Beide liefern Vorhersagen für thermische Dielektronenstrahlung aus dem Quark–Gluon Plasma und dem Hadrongas. Das *Thermal Fireball Model* basiert auf einer hydrodynamischen Beschreibung des Kollisionsverlaufes, aus der die Temperaturentwicklung abgeleitet und dann mithilfe einer modifizierten elektromagnetischen Spektralfunktion die thermische Strahlung berechnet wird. Im UrQMD Modell wird eine mikroskopische Transportsimulation auf Teilchenebene verwendet und dadurch auch ein eigener *Light-Flavour Cocktail* produziert. Zur Bestimmung der thermodynamischen Eigenschaften wird das Teilchensystem zunächst in Raumzeitzellen granularisiert, bevor die thermischen Beiträge ähnlich dem anderen Modell berechnet werden.

Zur direkten Vergleichbarkeit mit Messungen beziehen beide Modellvorhersagen die wichtigsten experimentellen Selektionskriterien ein, wozu insbesondere die Zentralität der Kollisionen und der messbare Transversalimpulsbereich der Elektronen zählen. Letzterer wird aus historischen Gründen typischerweise auf $p_T > 0,2 \text{ GeV}/c$ angesetzt, was für eine Messung mit dem ALICE Experiment aufgrund seiner Größe und Detektoranordnung eine Herausforderung darstellt. Die in dieser Arbeit entwickelte Strategie zur Elektronidentifizierung erlaubt jedoch das Senken dieser unteren Schwelle von sonst 0,4 auf 0,2 GeV/c und somit einen konsistenten Vergleich. Dieser ergibt für das *Thermal Fireball Model*, dass die Daten mit den zusätzlichen thermischen Beiträgen, die hauptsächlich im Massebereich zwischen Pion und ω -Meson relevant sind, teilweise gut beschrieben werden können. Manche Datenpunkte befinden sich jedoch unterhalb der Erwartung, wofür mehrere Faktoren wie statistische Fluktuationen, potentielle Überschubtraktion durch sehr hohen Untergrund, sowie die oben beschriebene Überschätzung der *Heavy-Flavour* Ausbeute verantwortlich sein können. Für den Vergleich mit dem zweiten Modell, UrQMD, wird der *Heavy-Flavour* Anteil des Cocktails von den Daten subtrahiert. Dabei ergibt sich eine bessere Übereinstimmung, wenn dessen nukleare Modifikation berücksichtigt, also eine geringere Ausbeute subtrahiert wird. Die Vergleiche mit beiden Modellen ergeben ein konsistentes Bild, doch für präzise weiterführende Analysen sind Daten mit höherer Statistik, sowie reduzierte Unsicherheiten besonders im *Heavy-Flavour* Cocktail notwendig. Beide Verbesserungen sind von Messungen nach dem aktuellen Upgrade des ALICE Detektors zu erwarten.

Im Frühstadium der vorliegenden Arbeit wurde eine Machbarkeitsstudie für eine Dielektronenmessung unter eben diesen experimentellen Bedingungen durchgeführt, die zusammen mit zwei weiteren Studien dazu diente, das angedachte Upgrade-Programm physikalisch zu motivieren. Als anschauliche und zugleich gut quantifizierbare Kenngröße wurde die Präzision gewählt, mit der die frühe Temperatur des Quark–Gluon Plasmas bestimmt werden kann.

Der Ablauf der Studie wird in Kapitel 6 beschrieben. Eine Vielzahl individueller Beiträge aus Simulationen, Tools, theoretischen Vorhersagen und Annahmen wurde verarbeitet und zu einer vollwertigen Dielektronenanalyse auf Basis künstlich erzeugter Kollisionsereignisse zusammengefasst. Daraus wurden zunächst das Signal-zu-Untergrund Verhältnis und die Signifikanz bestimmt. Ersteres spiegelt einen Teil der zu erwartenden systematischen Unsicherheiten wieder, während die Signifikanz in Kombination mit der geplanten Anzahl an Blei–Blei Kollisionen eine Stichprobenerzeugung anhand der Poissonverteilung ermöglicht, also eine realistische “Messung” mit statistischen Fluktuationen um das theoretisch erwartete Dielektronensignal generieren lässt. Dieses Referenzsignal wurde zusammengestellt aus einem hadronischen Cocktail inklusive *Heavy-Flavour* Anteil ohne nukleare Modifikation, sowie einer Vorhersage der thermischen Beiträge aus dem *Thermal Fireball Model*.

Da der *Heavy-Flavour* im mittleren Massebereich um $1,5 \text{ GeV}/c^2$ etwa eine Größenordnung über der zu messenden thermischen Strahlung des Quark–Gluon Plasmas liegt, wurde ein zusätzlicher Analyseschritt untersucht, der die verbesserte Spurrekonstruktion mit dem neuen *Inner Tracking System* ausnutzt, um dessen Beitrag zu senken. Mit einer strikten Selektion ist eine relative Unterdrückung um bis zu einen Faktor fünf zu erwarten, die allerdings auch die Paareffizienz auf ein zehntel reduziert. Letzteres kann jedoch durch weitere Aspekte des Upgrade-Programms gut kompensiert werden, die im Vergleich zu den Daten dieser Arbeit die Aufzeichnung einer hundertfachen Kollisionsrate zulassen. Der so reduzierte *Heavy-Flavour* sowie der restliche Cocktail wurden dann vom generierten Dielektronensignal subtrahiert, um das thermische Exzess-Spektrum zu extrahieren. Letzteres hat im relevanten Massebereich eine nahezu exponentiell abfallende Form, deren inverser Steigungsparameter direkt von der QGP-Temperatur abhängt. Bei der finalen Analyse dieses Spektrums und seiner propagierten Unsicherheiten zeigte sich, dass die inverse Steigung der Modellvorhersage im Rahmen statistischer und systematischer Variationen von etwa $\pm 10\%$ reproduziert werden konnte.

Die vielversprechenden Ergebnisse dieser Studie trugen einerseits zur Realisierung des ALICE Upgrades und zu einer Designentscheidung für das neue *Inner Tracking System* bei, und stellen zugleich spannende Vorhersagen für bevorstehende Messungen dar.

Contents

1. Introduction	3
2. Dileptons	9
2.1. Dilepton Production	9
2.1.1. Initial Scatterings	10
2.1.2. Thermal Radiation	13
2.1.3. Final-State Decays	22
2.2. Previous Measurements	23
2.2.1. CERES	23
2.2.2. NA60	25
2.2.3. PHENIX	26
2.2.4. STAR	28
3. Experiment	31
3.1. LHC	31
3.2. ALICE	33
3.2.1. Detector Performance	35
3.2.2. PID Calibration	36
4. Analysis	41
4.1. Events	41
4.1.1. Event Selection	42
4.1.2. Event Mixing	43
4.2. Tracks	44
4.2.1. Track Selection	45
4.2.2. Particle Identification	49
4.2.3. Resulting Electron Samples	52
4.3. Pairs	53
4.3.1. Signal Extraction	53

4.3.2.	Signal Quality	56
4.4.	Pair Prefiltering	59
4.4.1.	Prefilter Technique	60
4.4.2.	Prefilter Tuning	61
4.4.3.	Choice of Prefilter	64
4.5.	Final Pair Cuts	64
4.5.1.	ϕ_V Distributions	65
4.5.2.	ϕ_V Template Fit	69
4.6.	Efficiency Correction	71
4.6.1.	Generated vs. Measurable Quantities	72
4.6.2.	Factorization of Pair Efficiency	74
4.6.3.	Electron Efficiency	75
4.6.4.	Prefilter Efficiency	76
4.6.5.	Pair Efficiency	77
4.7.	Systematic Uncertainties	78
4.7.1.	Cut Variations	79
4.7.2.	Dielectron Signal Uncertainty	82
4.7.3.	Representative Signal	84
4.8.	Final Dielectron Yield	85
5.	Results	87
5.1.	Hadronic Cocktail	87
5.1.1.	Cocktail Production	90
5.1.2.	Cocktail Inputs	91
5.2.	Heavy-Flavour Modifications	96
5.2.1.	Nuclear Modification Factor of HFE	97
5.2.2.	Heuristic Charm Modifications	99
5.2.3.	Combination of Approaches	100
5.3.	Increasing Dielectron Acceptance	101
5.4.	Model Calculations	102
5.4.1.	Thermal Fireball	103
5.4.2.	Microscopic Transport	105
5.5.	Semi-Central Collisions	108
5.5.1.	Spectra and Signal Quality	108
5.5.2.	Track and Pair Efficiencies	109
5.5.3.	Systematic Uncertainty	111

5.5.4. Cocktail Comparison	112
6. Upgrade	115
6.1. Physics Motivation	115
6.2. Experimental Improvements	116
6.3. Physics Performance Study	118
6.3.1. Track Analysis	121
6.3.2. Pair Analysis	124
6.3.3. Dielectron Signal and Poisson Sampling	131
6.3.4. Spectra Analysis	137
7. Summary	143
A. ITS PID Calibration	149
B. Definition of ϕ_V	151
C. Upgrade Study Poster	153
D. Write Progression	155
Bibliography	157
Danksagung	169

Chapter 1.

Introduction

The scientific field of high-energy heavy-ion physics is a particularly exciting one, as it enables us to establish a connection between the smallest and largest scales in nature, from the behaviour of elementary particles to the evolution of the universe.

According to the standard model of cosmology, the universe began with a *hot big bang* and since then expands, cools and evolves in consistence with general relativity. A set of initial conditions, which are inferred from astronomical observations, is required for this description. These observations are, to date, based on electromagnetic radiation, which puts a natural limit on the earliest accessible evolutionary phase: The decoupling of energy and matter, due to the recombination of electrons with nuclei, provides the first source of photons still measurable today. This is the cosmic microwave background, which originates from around 400 k years after the big bang. For the evolution prior to this event, cosmology needs to rely on the description by models. The initial stages of the big bang were governed by exceptionally high temperatures, such that all energy and matter is expected to have existed in the form of its fundamental constituents. Their manifestations as known today are described by the standard model of particle physics, including their interactions under various conditions. These conditions can be differentiated into individual phases of matter with transitions between them. Since such transitions emerge when considering the properties of the early universe, the initial stages of the big bang are often described as the *era of phase transitions*. [CL02, CO07] The conditions during some of these stages can be recreated in the laboratory using heavy-ion collisions at ultra-relativistic energies. In contrast to the early universe, the observable final-states are mainly hadrons, providing insight into the properties of earlier phases. Additionally, the full time-evolution of temperature and baryon-number density can be inferred from simulations. This allows for charting the evolution of heavy-ion collisions into the phase diagram of strongly-interacting matter. [TI⁺03, FNT98, CBM05]

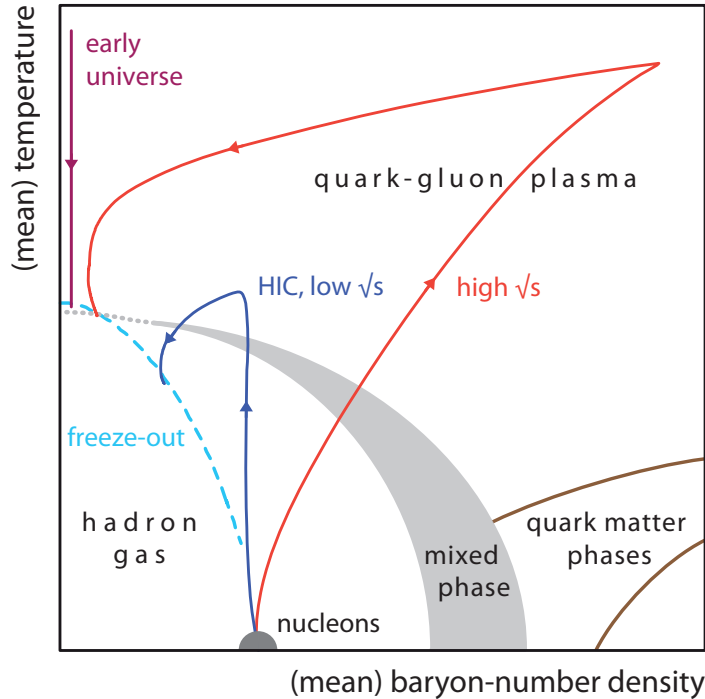


Figure 1.1.: Schematic view of the phase diagram of strongly-interacting matter, including evolutionary paths of heavy-ion collisions (HIC) and of the early universe. Based on [FH⁺11, RW00, BMRS03, FK04, TI⁺03, CBM05], see text for details.

A semi-quantitative view of our current understanding of this phase diagram is presented in Figure 1.1. It is drawn as a function of temperature and baryon-number density, stating their mean values to account for variations in the dynamics of heavy-ion collisions and their formation time for reaching local equilibrium. Nuclear matter, to which also the projectiles of heavy-ion collisions belong, is positioned at zero temperature and moderate baryon-number density. It is surrounded by the *hadron gas* phase, in which the constituents of strongly-interacting matter, quarks and gluons, are confined into composite particles. The *quark-gluon plasma* (QGP) phase covers the domain of high temperatures, at which quarks and gluons — jointly referred to as partons — become asymptotically free. At low baryon-number density, a smooth crossover transition between the hadronic and QGP phase exists, while at higher values a first-order phase transition is present, giving rise to a mixed phase of hadrons and partons. These two different phase boundaries are connected by a critical point, at which the phase transition is of second order and therefore expected to clearly show in various observables when the point is reached experimentally. At large baryon-number density and up to moderate temperatures, different *quark matter* phases exist, in which quarks form exotic condensates,

from color superconductors to color-flavor-locked states. Since the involved particles and their interactions are described by the theory of Quantum Chromodynamics (QCD), this diagram is also referred to as the QCD phase diagram. [FH⁺11, RW00, FK04]

In the laboratory, the highest achievable temperature and density depends on the collision energy \sqrt{s} and is reached as the projectiles intersect each other — at the kink of the trajectories shown in Figure 1.1. Quarks and gluons, which are liberated from the incoming nucleons or produced from the available energy, remain unconfined while the system expands. Even though the average system properties match those of the deconfined phase for a large fraction of the evolutionary path, the actual formation of a quark–gluon plasma only occurs if the system thermalizes during the expansion process, presumably only for high collision energies. As the system cools, it eventually reaches the condition of chemical freeze-out, where the relative abundances of hadron species become fixed. Statistical analysis of hadron compositions in various experiments leads to a unified freeze-out condition which reaches the *critical temperature* — marking the phase transition between QGP and hadron gas — for per-nucleon collision energies above $\sqrt{s_{\text{NN}}} \approx 20$ GeV. [TI⁺03, FH⁺11, BMRS03]

Heavy-ion collisions at LHC energy¹ consequently trace a path in the phase diagram, whose final cooling and freeze-out process essentially matches that of the early universe. The evolution of a heavy-ion collision proceeds within a few fm/c, i. e. of the order 10^{-23} s. In nature, the quark–hadron phase transition occurred $\approx 10 \mu\text{s}$ after the big bang. So by studying the properties of the quark–gluon plasma in high-energy heavy-ion collisions, even slightly earlier times of the universe’s evolution can be accessed experimentally. [CL02, FH⁺11]

The space–time diagram of a heavy-ion collision at high \sqrt{s} is shown in Figure 1.2. Evolving from bottom to top, the incoming nuclei intersect and experience initial hard scatterings at $t = 0$. After a short formation time $\tau \approx 0.1$ fm/c, during which multiple interactions and particle production occur, the QGP phase is reached as the majority of particles become thermalized. Partons from the initial processes are scattered in the hot medium and thermal radiation (real and virtual photons) is emitted from it. The phase transition from QGP to hadron gas is a crossover, without an intermediate mixed phase, which would be traversed only at lower \sqrt{s} . Here, the phase transition also coincides with the condition for chemical freeze-out. This means that essentially no inelastic collisions occur during the hadronic phase, such that the relative yields of the various hadron species are conserved throughout this phase — modulo their different lifetimes. Hadrons

¹The Large Hadron Collider at CERN delivers Pb–Pb collisions of up to $\sqrt{s_{\text{NN}}} \approx 5$ TeV.

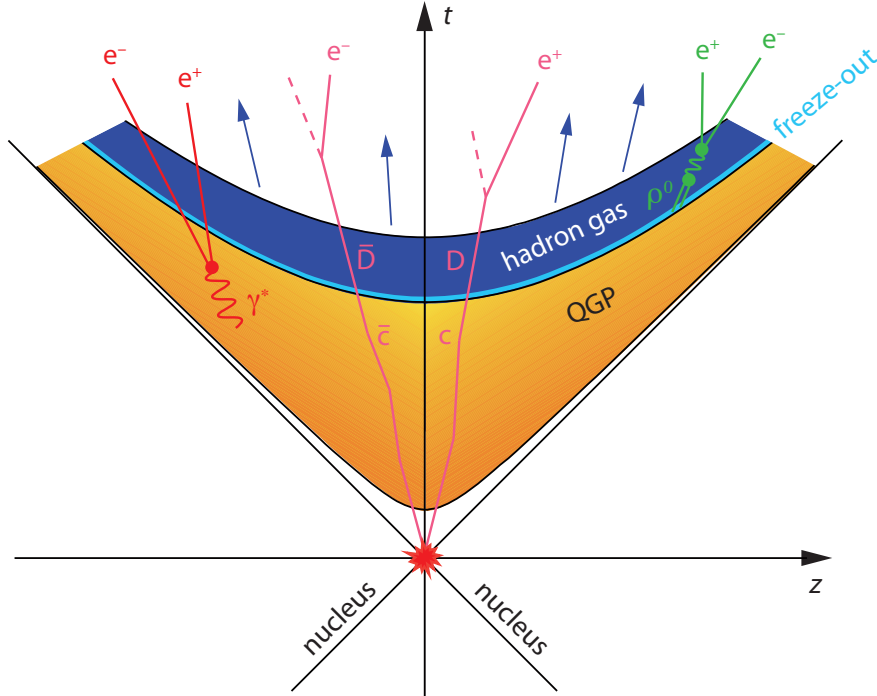


Figure 1.2.: Space–time evolution of a heavy-ion collision at high energy, which undergoes a crossover phase transition from QGP to hadron gas, rather than evolving through a mixed phase. Chemical freeze-out is reached during this phase transition. The bulk of free-streaming hadrons after kinetic freeze-out is indicated with arrows. Various sources of electron–positron pairs relevant to this work, originating from the different collision stages, are also included. Based on [SSS10].

merely perform elastic scatterings until also kinetic freeze-out is reached towards the top of the diagram. [SSS10, SS15]

With the aim to characterize the QGP phase of heavy-ion collisions, any measurement performed on hadrons needs further interpretation to infer from it the medium properties at earlier times. This is done by comparison to models, which often contain a hydrodynamic simulation of the QGP phase, a hadronization process and optionally a hadronic transport calculation. The early chemical freeze-out in case of high collision energies reduces the model dependence as it effectively places the hadron measurement close to the QGP phase.

Despite of their indirect nature, observables that are dominated by strongly-interacting probes provide valuable insight into various aspects of the collision evolution and properties of the created system [SS15, Kra12, Boo14]:

– The momentum distribution of charged particles and individual spectra of identified hadrons, as well as particle yields and multiplicity, characterize the final state of the collision. The observed mass-ordering in the slope of these spectra matches with the existence of collective motion in a thermalized system, and such a medium in thermal and chemical equilibrium also reproduces the measured hadron yields. This finding is used to establish the chemical freeze-out condition and locate it in the phase diagram, as discussed above.

– The pressure and viscosity of the medium as well as space–time quantities of the system are reflected in particle correlations. From the azimuthal anisotropy of measured particles the shape of the overlap region of the incoming nuclei can be reconstructed. The anisotropy is further influenced by the transport properties of the medium and reproduced best when describing the latter as an almost perfect fluid. Interferometric methods are used to evaluate the size of the homogeneity region as well as the decoupling time for hadrons.

– Hard probes, produced in the initial scatterings, carry information on initial-state effects in the colliding nuclei and interact with the medium as they traverse it, thereby probing its properties via their energy loss. High-momentum particles experience larger suppression with increasing collision energy and centrality, and no suppression in proton–nucleus collisions (both relative to proton–proton collisions), showing that it is a medium-effect of heavy-ion collisions. Jets of particles provide more differential measurements as they can be correlated with each other or with high-energetic photons to select individual production mechanisms, which allows for probing the energy loss of the scattered partons in the QGP phase.

– Heavy-flavour hadrons are associated to hard probes and provide similar insight, though particularly for the charm- and beauty-quark sector. Both their suppression and azimuthal anisotropy are comparable to those of light hadrons. Description of the former requires both radiative and collisional energy loss in the QGP phase. The latter shows that heavy quarks also participate in the collective expansion of the medium, again requiring considerable re-scattering. Quarkonia, i. e. bound states of a heavy-flavour quark–antiquark pair, experience additional effects that are exclusive to the deconfined phase. Predominantly, they melt because of color screening and regenerate in case of a sufficiently high abundance of the particular quark flavour in the medium. For charmonium states a combination of these effects leads to less suppression at the LHC compared to lower energy, while bottomonium states show sequential suppression depending on their binding energy.

Complementary insight to that from hadronic observables is gained from the study of electromagnetic probes, which do not experience the strong force. They therefore traverse the medium without significant interaction and carry direct information from all collision stages: Photons are produced early in hard processes such as initial scatterings and bremsstrahlung till late in particle decays after freeze-out. They are also emitted from the hot medium in the form of thermal radiation, which allows for a temperature measurement of the QGP phase, given that this contribution can be isolated. In general they may exist as real or virtual photons, leaving the system as they are or quickly decaying into lepton pairs, respectively. [A⁺04]

In the present thesis the latter kind are studied, by means of measuring electron–positron pairs, referred to as *dielectrons*. This approach results in additional sources to be considered, one kind are resonance decays of vector mesons and the other are semileptonic decays of correlated heavy-flavour hadrons. Two representative cases are included in Figure 1.2, besides a thermal virtual photon from the QGP phase: decays of D mesons originating from one charm-anticharm pair produced in a hard scattering and the decay of the ρ meson shortly after chemical freeze-out. Measurements of these three types of dielectron sources are pursued in the present work and discussed throughout the following chapters.

Chapter 2.

Dileptons

In this chapter the production of dileptons in heavy-ion collisions is described and put into context with the measurements in this thesis. A selection of previous dilepton measurements with related scopes, mainly covering the dielectron and dimuon continuum at low and intermediate mass, is presented subsequently.

2.1. Dilepton Production

The different dilepton production mechanisms vary with the amount of energy available in the respective phase of the collision and generally range from hard processes in the initial phase to soft reactions in the equilibrated late stages. This relation further extends to the main observable of dilepton measurements, namely the invariant mass m_{l+l-} of the lepton pair. In case of a two-body process, m_{l+l-} is equivalent to the mass of the particle that produced this pair, which on the lowest level is a virtual photon. To leading order, such a photon is produced in the annihilation of a quark–antiquark pair, as shown by the diagram in Figure 2.1. In this process the entire energy of the quarks is transferred to the lepton pair, such that earlier annihilation events produce dileptons with higher invariant mass. Related processes in the hadronic phase follow the same general trend. [Won94]

In the following, the dominant dilepton production mechanisms are discussed in the order of their occurrence throughout the collision evolution.




Figure 2.1.: Feynman diagram of quark–antiquark annihilation via the electromagnetic channel. A virtual photon is produced, which then decays into a lepton pair.

2.1.1. Initial Scatterings

The most energetic reactions occur in the pre-equilibrium phase of the collision amongst the partons of the opposing nuclei. Dileptons may be produced directly from two quarks as well as in the decay chain of strong interactions.

Drell–Yan Process

In Drell–Yan dilepton production, a quark from one projectile annihilates with its corresponding antiquark in the opposing projectile or target. The leading-order process therefore follows the diagram of Figure 2.1:

$$q + \bar{q} \rightarrow \gamma^* \rightarrow e^+ + e^-. \quad (2.1)$$

Due to the high center-of-mass energy available in this process, the mass distribution of produced lepton pairs reaches further than that of other contributions. [Won94, Vog07] The Drell–Yan cross section scales with invariant mass m and collision energy \sqrt{s} like

$$\frac{d\sigma_{\text{DY}}}{dm dy} \propto \frac{1}{m^3} \cdot \left(\frac{m}{\sqrt{s}}\right)^{a_1} \cdot e^{-m \cdot a_2 / \sqrt{s}}, \quad (2.2)$$

with parameters $a_1 \approx 0.4$ and $a_2 \approx 11.5\text{--}15$ [Won94]. The inverse slope of the exponential can be interpreted as an effective temperature $T_{\text{DY}} = \sqrt{s}/a_2$, which for high collision energies vastly exceeds the temperatures associated to other processes.

In collisions of nuclei with mass number A and impact parameter b , the number of produced Drell–Yan dileptons N_{DY} further scales with the nuclear overlap function $T_{\text{AA}}(b)$.




Figure 2.2.: Predictions for the dielectron yield in central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV coming from hard (Drell–Yan and correlated charm) and from thermal (QGP and hadron gas) processes. Not included are contributions from final-state hadron decays. The low-mass region (left panel) contains medium-modifications from the hadron gas, while the intermediate-mass region (right panel) is sensitive to QGP radiation. From [A⁺08], modified.

For the most central collisions, where b goes to zero, this results in

$$\frac{dN_{\text{ll}}}{dm dy} \propto A^{4/3} \cdot \frac{d\sigma_{\text{DY}}}{dm dy} \quad (2.3)$$

[Won94, Vog07]¹.

Calculations for the LHC design energy predict that the Drell–Yan contribution to the dilepton yield becomes relevant — meaning comparable to QGP radiation — only for $m_{ee} \gtrsim 5$ GeV/ c^2 in the inclusive mass spectrum, and possibly at very low mass for pair momenta of $p_{T,ee} > 8$ GeV/ c [Vog07, HW10]. Both of these regions are not accessible to the analysis in this work, mainly because of statistical limitations in the data.

In the intermediate-mass region from $1 < m_{ee} < 3$ GeV/ c^2 , which is studied in two parts of this thesis, Drell–Yan production is expected to be at least one order of magnitude smaller than QGP radiation. This can be seen in the right panel of Figure 2.2, which shows the predicted yield for these processes in central Pb–Pb collisions for $m_{ee} > 1.5$ GeV/ c^2 . Also the tail of in-medium radiation from the hadron gas and an estimate for dileptons from correlated charm hadrons is shown (see discussion below).

¹The books of Vogt and Wong provide comprehensive treatises on the geometry, kinematics, and cross sections involved in heavy-ion collisions.

Dileptons from Heavy Quarks

The production of heavy-quark pairs, i. e. $c + \bar{c}$ and $b + \bar{b}$, is limited to the initial parton scatterings since their masses are much higher than the temperature of the system once it thermalizes [HW10, SS15]. The high-energetic heavy quarks interact with the medium, but the QGP lifetime is too short for them to reach thermal equilibrium. While a few percent of the produced pairs may form quarkonia, most of them dissociate in the deconfined phase because of color screening. The solitary heavy quarks later hadronize by fragmentation or recombination with light quarks to form open-heavy-flavour hadrons, mainly D mesons, which in turn may decay semi-leptonically:

$$c + \bar{c} \xrightarrow{\text{hadr.}} D + \bar{D}, \quad (2.4)$$

$$D \rightarrow X + \ell^+ + \nu_\ell, \quad (2.5)$$

$$\bar{D} \rightarrow X + \ell^- + \bar{\nu}_\ell. \quad (2.6)$$

On the one hand, this provides an experimentally valuable final-state to study heavy-quark energy loss and diffusion, since the decay leptons preserve most of the medium-modifications imposed on the heavy-flavour particles. On the other hand, when both of the open-heavy-flavour hadrons from the same heavy-quark pair decay into leptons of the same species, these correlated electron or muon pairs cause the dominant background to measurements of thermal dileptons over a wide mass range [HW10]. This is visible in both panels of Figure 2.2, where an estimate of the charm contribution² is given up to $m_{ee} = 3 \text{ GeV}/c^2$. Not shown is the beauty contribution, which for $m_{ee} \gtrsim 2 \text{ GeV}/c^2$ amounts to $\approx 1/3$ of the charm yield according to computations within this thesis (see Figure 5.1) and thus would also exceed the QGP yield slightly in the intermediate-mass region.

The relative contribution of the open-heavy-flavour components to the dielectron yield increases with collision energy, as their cross sections $\sigma_{c\bar{c}}$ and $\sigma_{b\bar{b}}$ rise quicker than the inelastic nucleon-nucleon cross section σ_{inel} , which is relevant for the bulk of particle production. From RHIC to LHC energy, the latter rises by $\approx 50\%$ [TOTEM19], while the charm and beauty cross sections increase by about one order of magnitude — more pronounced for the heavier particle [ALICE12a, ALICE14a]³.

²The estimate is based on data from RHIC and scaled to a higher cross section, such that nuclear modification is included to the extent observed at $\sqrt{s_{\text{NN}}} = 200 \text{ GeV}$ [A⁺08]. Yield and shape match well with the charm contribution computed in this thesis for $\sqrt{s_{\text{NN}}} = 2.76 \text{ TeV}$ (cp. Figure 5.9), when considering the rapidity coverage and accounting for $\approx 50\%$ heavy-flavour suppression (see Section 5.2). At nominal LHC energy the relative charm yield further increases.

³See also the RHIC results summarized in Section 2.2 in comparison to the present results in Chapter 5.

In heavy-ion collisions, two additional factors accelerate the aforementioned effect, since heavy-flavour yields scale with the number of binary nucleon-nucleon collisions given by

$$N_{\text{coll}}(\sqrt{s}, b) = \sigma_{\text{inel}}(\sqrt{s}) \cdot T_{\text{AA}}(b), \quad (2.7)$$

whereas the overall number of produced particles, which includes dilepton signals from soft processes, is found to scale well with the number of nucleons that participate in the collision, $N_{\text{part}}(b)$ [Vog07]. This number of participants is almost independent of the collision energy, and for the most central collisions peaks at the sum of all nucleons in the colliding ions, i. e. $N_{\text{part}}(0) = 2 \cdot A$, reaching ≈ 400 for central Pb–Pb collisions. In contrast, N_{coll} rises with collision energy owing to its dependence on σ_{inel} , plus the nuclear overlap function T_{AA} increases faster than N_{part} when heavier ions are chosen, since it is proportional to $A^{4/3}$ [Vog07]. In central Pb–Pb collisions at LHC energy, the combined effect leads to $N_{\text{coll}} \approx 1800$ [Vog07, ALICE13a].

One phenomenon that mildly counteracts this disproportionate rise of the open-heavy-flavour yield is the medium-modification of heavy quarks and hadrons due to interactions with the thermalized system created in nuclear collisions⁴. Since heavy quarks are produced early, they have to traverse the dense medium and thereby lose energy via collisional and radiative processes [HW10, SS15]. This shift to lower momenta leads to a suppression of heavy-flavour particles at high transverse momentum, when compared to the yield expected from N_{coll} -scaling of results obtained in proton–proton collisions. The scatterings of these particles also cause a diffusion of their trajectories and some degree of collective behaviour with the medium, which can be inferred from their azimuthal anisotropy in non-central collisions. For D mesons under LHC conditions, the suppression peaks at a factor of ≈ 4 and the anisotropy at $\approx 20\%$ [SS15]. This affects also the dielectron yield from correlated heavy-flavour hadrons, and different approaches are used in the present work to model these medium modifications (see Section 5.2).

2.1.2. Thermal Radiation

Once the majority of particles has been produced in the early phase of the collision, subsequent interactions lead to an equilibration of their energies, eventually creating a thermalized medium. This medium consists of quarks and gluons as long as its temperature is higher than the transition temperature T_c of the QGP, and mainly of pions once it reaches the hadronic phase. Processes amongst these particles create final states

⁴See [RvH10] for an extensive review on heavy-flavour physics in heavy-ion collisions.

which carry information on the thermodynamic properties of the medium in their energy distribution. [Won94, Vog07]

The leading-order process for dilepton production in the QGP is quark–antiquark annihilation, according to the diagram discussed above, and correspondingly $\pi^+\pi^-$ annihilation in the hadron gas. The reaction rate for either of these processes can be derived from kinetic theory⁵ and depends on the spacial density of the initial-state particles $n_a = \int \frac{d^3p}{(2\pi)^3} f_a(\vec{p})$ with occupation probability $f_a(\vec{p})$, on their relative velocity v_{rel} , and on the cross section $\sigma_{a+a^- \rightarrow l+l^-}(m) \equiv \sigma_a(m)$:

$$R_{a+a^- \rightarrow l+l^-} = \int \frac{d^3p_1 d^3p_2}{(2\pi)^6} f_a(\vec{p}_1) f_a(\vec{p}_2) \sigma_a(m) v_{\text{rel}} \quad (2.8)$$

[KK⁺86], where a denotes a quark or a pion.

Since this rate is based on a particle density, it is rather a rate per unit volume⁶ and can thus also be written as the number of produced dileptons per unit time and unit volume or per unit four-volume: $R = \frac{dN_{\text{ll}}}{dt d^3x} = \frac{dN_{\text{ll}}}{d^4x}$.

To discuss the spectral shape of dilepton radiation, the momentum vectors of the incident particles are expressed in terms of these particles' energies E_1 and E_2 and the invariant mass m of the process. This requires symmetry with respect to the orientation of their momenta, such that $f_a(\vec{p}) \rightarrow f_a(E)$, which is the case for a system in thermal equilibrium [Won94]. Further using $dm^2 = 2|\vec{p}_1||\vec{p}_2|d(\cos\theta_{\text{rel}})$ and $v_{\text{rel}} \cong \frac{m^2}{2E_1E_2} \cdot \sqrt{1 - \frac{4m_a^2}{m^2}}$ leads to

$$\frac{dN_{\text{ll}}}{d^4x} = \frac{1}{2(2\pi)^4} \int dE_1 dE_2 dm^2 f_a(E_1) f_a(E_2) \sigma_a(m) m^2 \sqrt{1 - \frac{4m_a^2}{m^2}} \quad (2.9)$$

[Won94].

In a thermalized system at high temperature T , the energy distribution of the incident particles has the form $f_a(E) = e^{-E/T}$ such that integration over E and m^2 gives a rate per unit volume and per invariant mass squared of

$$\frac{dN_{\text{ll}}}{dm^2 d^4x} = \frac{\sigma_a(m)}{2(2\pi)^4} \cdot m^3 \cdot T \cdot \sqrt{1 - \frac{4m_a^2}{m^2}} \cdot \sqrt{\frac{\pi T}{2m}} \cdot e^{-m/T} \quad (2.10)$$

[Won94, KK⁺86]⁷.

⁵A detailed derivation of dilepton production rates and cross sections is done by Wong ([Won94] chapter 14 and exercises).

⁶In their original paper on dilepton emission from the quark–gluon plasma, McLerran and Toimela did use R/V instead of R as the rate per unit spacial volume [MT85].

⁷Note that Wong derives this result for the specific case of QGP radiation, which leads to additional prefactors (see discussion below). Kajantie et al. present this general form, but express the last two

The produced lepton pairs have to escape from the hot system in order to be observed experimentally. Since leptons do not interact via the strong force, their interactions with the medium are only of the order α_{EM}^2 and thus their mean-free paths much larger than the system size. Therefore the production rate given above can readily be treated as a thermal emission rate. [Won94, KK⁺86]

The rate given in equation (2.10) features an exponential decrease with invariant mass and its slope depends on the present temperature $T(x)$ of the medium, already indicating the typical shape of predictions for QGP radiation. However, in order to obtain mass spectra for dilepton radiation from quark–gluon plasma and hadron gas as presented in Figure 2.2, the corresponding cross sections, $\sigma_q(m)$ and $\sigma_\pi(m)$ respectively, need to be considered and a space–time integration of equation (2.10) over the collision evolution be performed. Both steps are discussed in the following.

Cross Sections in QGP and Hadron Gas

In both phases, the dilepton production cross section can be expressed in terms of the elementary cross section $\sigma_{e^+e^- \rightarrow \mu^+\mu^-} \equiv \tilde{\sigma}$ given by

$$\tilde{\sigma}(m) = \frac{4\pi}{3} \frac{\alpha_{\text{EM}}^2}{m^2} \left(1 + \frac{2m_l^2}{m^2}\right) \sqrt{1 - \frac{4m_l^2}{m^2}} \quad (2.11)$$

[KK⁺86], where $\alpha_{\text{EM}} = e^2/4\pi$ is the electromagnetic coupling constant, m the invariant mass, and m_l the mass of the final-state leptons.

In the quark–gluon plasma, the cross section further depends on the number of quark flavours N_f that are part of the thermalized medium and their respective charges e_f , as well as the number of colors $N_c = 3$ and the number of spins $N_s = 2s + 1 = 2$:

$$\sigma_{q^+q^- \rightarrow l^+l^-}(m) = N_c N_s^2 \sum_f^{N_f} \left(\frac{e_f}{e}\right)^2 \cdot \tilde{\sigma}(m) \quad (2.12)$$

$$= F_q \cdot \tilde{\sigma}(m). \quad (2.13)$$

The coefficient F_q does not depend on the invariant mass of the process, if quarks are considered to be massless. It takes the value $F_q = 20/3$ for a system of up and down quarks, or $F_q = 24/3$ when strange quarks are included. [KK⁺86]

In the hadronic phase, $\pi^+\pi^-$ annihilation is the dominant dilepton production channel and according to the vector dominance model [Sak69] the process evolves via an intermediate

factors as the modified Bessel function $K_1(m/T)$ (and they appear to have missed the power of $\frac{1}{2}$ in the mass threshold term coming from v_{rel}).




Figure 2.3.: Coefficients for the cross sections of dilepton production in the partonic and hadronic phases. The numbers are based on two quark flavours, u and d, and correspondingly on a hadron gas made only of pions. From [KK⁺86].

ρ or ρ' vector meson, each having the same quantum numbers as a virtual photon:

$$\pi^+ + \pi^- \rightarrow \rho^{(\prime)} \rightarrow e^+ + e^-. \quad (2.14)$$

The coupling to these mesons is strongest when the center-of-mass energy of the process is equal to their mass, which is described by a form factor $F_\pi(m)$ that depends on their masses $m_{\rho^{(\prime)}}$ and widths $\Gamma_{\rho^{(\prime)}}$:

$$F_\pi(m) = \frac{m_\rho^4}{(m_\rho^2 - m^2)^2 + m_\rho^2 \Gamma_\rho^2} + \frac{1}{4} \frac{m_{\rho'}^4}{(m_{\rho'}^2 - m^2)^2 + m_{\rho'}^2 \Gamma_{\rho'}^2} \quad (2.15)$$

[KK⁺86]. The factor $\frac{1}{4}$ results from the different branching ratios to produce these mesons. In addition, the pion mass cannot be neglected and leads to a mass threshold of $m > 2m_\pi$ for this process, such that the cross section reads

$$\sigma_{\pi^+\pi^-\rightarrow l^+l^-}(m) = F_\pi(m) \cdot \tilde{\sigma}(m) \cdot \sqrt{1 - \frac{(2m_\pi)^2}{m^2}} \quad (2.16)$$

[KK⁺86, Won94].

For a quantitative comparison of the cross sections in both phases, the coefficients F_q and $F_\pi(m) \cdot \sqrt{1 - (2m_\pi/m)^2}$ are presented in Figure 2.3 for the case of two quark flavours and a pion gas. The resonance structure clearly determines the invariant-mass dependence of the hadronic cross section, which also vanishes below the two-pion threshold at $m < 280 \text{ MeV}/c^2$. In contrast to this, the partonic cross section is constant in mass and reaches down to essentially zero (not entirely shown).

Electromagnetic Spectral Function

The cross section for dilepton production can also be expressed in terms of the electromagnetic spectral function Π_{EM} , which provides a unified approach to describe dilepton emission from partonic and hadronic processes in the collision. It is defined via the correlation function of the electromagnetic current in the system, which depends on the respective electric-charge carriers, being either quarks or hadrons. The EM spectral function is therefore also referred to as the electromagnetic current-current correlator. [Rap01, Rap13]

In the vacuum, both hadronic and partonic degrees of freedom are involved in dilepton reactions, with changing relevance depending on invariant mass. This is seen in the time-reverse process to dielectron production, namely electron-positron annihilation into hadrons, which gives access to the EM spectral function in the vacuum $\Pi_{\text{EM}}^{\text{vac}}$ according to

$$\sigma_{e^+e^- \rightarrow \text{hadrons}}(m) = -\frac{12\pi}{m^2} \cdot \tilde{\sigma}(m) \cdot \text{Im} \Pi_{\text{EM}}^{\text{vac}}(m) \quad (2.17)$$

[RWvH09]. The term $\text{Im} \Pi_{\text{EM}}$ denotes the imaginary part of the spectral function, i. e. its expectation value or strength. Recall that $\tilde{\sigma}$ is the cross section to produce dimuons, so the ratio

$$R(m) \equiv \frac{\sigma_{e^+e^- \rightarrow \text{hadrons}}}{\sigma_{e^+e^- \rightarrow \mu^+\mu^-}} = -\frac{12\pi}{m^2} \cdot \text{Im} \Pi_{\text{EM}}^{\text{vac}}(m) \quad (2.18)$$

closely resembles the vacuum EM spectral function.

This quantity is presented in Figure 2.4, where \sqrt{s} is equivalent to the invariant mass m . The measurements show broad resonance structures from coupling to the ρ and ρ' mesons, qualitatively similar to the parameterization shown in Figure 2.3. In addition, they feature narrow peaks from the ω and ϕ mesons, which mainly decay into three pions or two kaons, respectively. These are not included above, since the reverse processes in a hadron gas are less likely to occur. For higher invariant masses, $m \gtrsim 1.5 \text{ GeV}/c^2$, the spectral strength can be explained by a coupling to free quark-antiquark states, as comparisons with the naive quark model and with a perturbative-QCD calculation demonstrate. These quarks then fragment into multiple hadrons that approximately follow a statistical distribution [RWvH09].

In the thermalized system produced by heavy-ion collisions, this spectral function is modified, since the inverse reactions in both the partonic and the hadronic phase are affected by the surrounding medium. Both contributions to the intermediate-mass




Figure 2.4.: Ratio of the e^+e^- annihilation cross section into hadrons over that into muons, which is closely related to the EM spectral function in the vacuum. Shown is the region of low to intermediate invariant mass, where only light quarks and the corresponding hadrons participate in the coupling. From [PDG18], modified.

region are however expected to remain structureless: multi-hadron annihilation such as $4\pi \rightarrow e^+e^-$ of thermalized particles leads to the same uniform spectral strength as quark–antiquark annihilation [RWvH09].

In the low-mass region, the hadronic part of the spectral function can still be described by the sum of the vector-meson states,

$$\text{Im } \Pi_{\text{EM}}^{\text{had}} = \sum_{V=\rho,\omega,\phi} \frac{m_V^4}{g_V^2} \cdot \text{Im } D_V, \quad (2.19)$$

except that their individual spectral functions $\text{Im } D_V$ are modified by the hot hadronic medium [RW99].

Figure 2.5 shows modifications of these mesons depending on the temperature and effective nucleon density of the medium, calculated from polarization of and resonant interactions with the surrounding mesons and baryons⁸. In comparison to their vacuum line shapes, the ρ and ω mesons experience a strong broadening near the expected phase boundary, while the strange-quark state ϕ appears to be more protected from medium effects [Rap01]. Despite of the spectral strengths given in the figure, the ρ meson is the dominant transition channel because of its large width and its isospin $I = 1$ [RWvH09].

For the evaluation of the QGP contribution to the low-mass part of the spectral function, multiple approaches exist, each suited for a particular region of the phase diagram: In

⁸See [Rap01] for a discussion of all self-energy corrections that are applied to the vector-meson propagators D_V .




Figure 2.5.: Spectral functions of the ρ (left panel), ω and ϕ mesons (right panel), which together form the low-mass part of the EM spectral function, in the vacuum and for a hadron gas of different temperatures and densities. From [Rap01].

the high-temperature limit, perturbative-QCD calculations are valid, so the spectral function should be structureless. As the temperature decreases towards the phase boundary, lattice-QCD computations allow for the inclusion of higher-order processes, which become significant as the strong coupling constant⁹ α_s approaches or exceeds one. For moderate temperatures and densities, many-body techniques are suitable and in principle also able to describe phase-transition dynamics. [RWvH09]

In the models considered throughout this thesis, the EM spectral function in the QGP is parameterized with a baseline strength given by leading-order perturbative-QCD, modified by finite-temperature corrections encoded in $\hat{f}_2(q, T)$, and a term describing the enhancement towards vanishing mass seen in the spectral function of lattice-QCD “data”:

$$\text{Im } \Pi_{\text{EM}}^{\text{QGP}} = -\frac{N_c}{12\pi} \sum_f \left(\frac{e_f}{e} \right)^2 \cdot m^2 \cdot \left(\hat{f}_2(q, T) + \frac{1}{3} Q_{\text{LAT}}^T(m, |\vec{q}|, T) \cdot \left(2 + \frac{m^2}{q_0^2} \right) \right) \quad (2.20)$$

[Rap13], where $q = (q_0, \vec{q})$ is the four-momentum transfer and Q_{LAT}^T denotes the transverse part of the lattice-QCD modification. For high temperatures, \hat{f}_2 approaches unity and Q_{LAT}^T vanishes, leading to the naive-quark-model baseline.

With the electromagnetic spectral function defined for both phases, the thermal dilepton production rate at any instant of the collision evolution can be computed from the

⁹Despite of its name, the value of α_s is not constant but rather a function of squared momentum transfer Q^2 and the number of quark flavours with $m_q < Q$ [Vog07].




Figure 2.6.: Dilepton rates in the low-mass region produced via the isovector (ρ) channel. While the vacuum and perturbative-QGP rates differ substantially, the in-medium hadronic and lattice-QGP rates approach each other near the phase-transition temperature at small baryochemical potential. From [Rap13], modified.

McLerran-Toimela formula [MT85]. In recent literature it is presented in the form

$$\frac{dN_{\text{II}}}{d^4x d^4q} = -\frac{\alpha_{\text{EM}}^2 L(m)}{\pi^3 m^2} \cdot f^{\text{B}}(q_0, T) \cdot \text{Im} \Pi_{\text{EM}}(m, |\vec{q}|, \mu_{\text{B}}, T) \quad (2.21)$$

[Rap13, EvHB16], where $L(m)$ describes the lepton phase-space and f^{B} is the thermal Bose distribution. The EM spectral function depends on the mass and momentum of the lepton pair as well as on the density and temperature of the medium, the former represented by its baryochemical potential μ_{B} .

The dilepton production rates obtained with this relation using EM spectral functions according to equations (2.19) and (2.20) are presented in Figure 2.6. They are only based on the dominant transition channel with isospin $I = 1$ for illustration purposes. All rates approach a common value and shape towards the intermediate-mass region, while in the low-mass region they deviate from one another. However — in contrast to the two extreme cases of a vacuum or a perturbative QGP — both realistic descriptions of the conditions near the phase boundary at LHC energy lead to rather similar rates, in particular towards low invariant mass: the in-medium hadronic and nonperturbative partonic rates clearly approach each other, as they should in case of parton–hadron duality, which also leads to the smooth crossover phase transition that is expected for LHC conditions [Rap13].

Integrated Dilepton Yields

Finally, in order to obtain observable dilepton spectra for the emissions of the thermalized medium, the rates need to be integrated over the space–time evolution of the collision from the moment of thermalization at the initial proper time τ_0 to chemical freeze-out at τ_{chem} . This requires knowledge of the systems’ equation of state (EoS), which in a hydrodynamic description relates its energy density to temperature and pressure, and thereby determines the system evolution [Won94]. In high-energy nucleus-nucleus collisions, the energy density is approximately constant over an extended rapidity range¹⁰, such that the system can be described by Bjorken hydrodynamics [Bjo83], which is Lorentz invariant under longitudinal translation. For an ideal gas of massless quarks and gluons, energy density and pressure then depend only on the proper time according to

$$\frac{\epsilon(\tau)}{\epsilon(\tau_0)} = \frac{p(\tau)}{p(\tau_0)} = \left(\frac{\tau_0}{\tau}\right)^{4/3}, \quad (2.22)$$

while the temperature decreases more slowly as

$$T(\tau) = T_0 \cdot \left(\frac{\tau_0}{\tau}\right)^{1/3} \quad (2.23)$$

[Won94]. Also the system size can be expressed as a function of proper time using

$$d^3x = d^2x_T \cdot \tau dy, \quad (2.24)$$

where x_T is the transverse dimension and y is the rapidity [Bjo83].

Integration of the dilepton rate given in equation (2.10) over all contributions from the QGP phase — i. e. from τ_0 to τ_c , at which the critical temperature T_c for the transition to a hadron gas is reached — leads to

$$\frac{dN_{\text{ll}}}{dm dy} \approx \frac{5}{3\pi^2} \sqrt{\frac{\pi}{2}} \cdot \alpha_{\text{EM}}^2 \tau_0^2 R_A^2 T_0^{2.5} \cdot \sqrt{m} \cdot e^{-m/T_0}, \quad (2.25)$$

using approximations that are valid in the limit $m \gg T_0 \gg T_c$ [Won94]¹¹. The radius of the nuclei R_A results from integration over the transverse system size d^2x_T . With $T_c \approx 170$ MeV and $T_0 \approx 560$ – 800 MeV at LHC energy [Rap13, EvHB16], the applicability

¹⁰Rapidity is a measure for the angle relative to the beam axis, with midrapidity $y = 0$ being perpendicular to the beam direction. For any given proper time, the rapidity of a particle is closely related to its longitudinal position [Won94].

¹¹In comparison to equation (14.11) of [Won94], further corrections of order $(m/T)^{-n}$ are skipped here, based on the premise of $m \gg T$.

of equation (2.25) is clearly limited to the intermediate-mass region and above.

For more precise descriptions covering the full invariant-mass range, models tend to implement the McLerran-Toimela formula in combination with electromagnetic spectral functions as discussed above and integrate over the space–time evolution using a realistic equation of state for both phases including a phase transition. Details on the models from which results are used in this thesis are given in Section 5.4.

The thermal emissions from each phase are computed separately, resulting in dilepton spectra of the kind presented in Figure 2.2. Two scenarios for the hadronic phase are shown, a free and an in-medium hadron gas, where the former only serves as a guideline using vacuum spectral functions. Up to the broad ρ resonance, the hadron gas yield is comparable to that from decays of correlated open-charm hadrons, before dropping off rapidly¹². The QGP yield is about a factor of five smaller than the charm yield over the full mass range. Experimental techniques to suppress the heavy-flavour contribution in order to measure the QGP yield in upcoming LHC data are discussed in Chapter 6.

2.1.3. Final-State Decays

For a full coverage of all dilepton emission sources, the last contribution comes from hadron decays after chemical freeze-out. Most relevant to the low-mass dilepton continuum are Dalitz decays of the π^0 and η mesons into a real and a virtual photon. Depending on the four-momentum fraction given to the virtual photons, they decay into lepton pairs of varying invariant mass between zero and the meson mass, thus creating a continuum in the dilepton yield [KW55]. In addition, two-body decays of the ρ , ω , ϕ and J/ψ mesons produce resonances in the dilepton mass spectrum. Since the ρ meson has a comparatively short lifetime of only $\tau \approx 1.3 \text{ fm}/c$ [ALICE19b], its decays create a broad distribution around its nominal mass, in contrast to the narrow peaks produced by the other mesons. This ρ contribution is distinct from its thermal production in the hadron gas — which is discussed above — and for differentiation called the “freeze-out ρ ”.

The summed dilepton yield from these meson decays and from the correlated heavy-flavours described in Section 2.1.1 is commonly referred to as the *hadronic cocktail*. It is usually constructed on an individual basis to match the requirements of the particular analysis and experimental setup. The approach followed in the present thesis is discussed

¹²A thermal contribution from the ϕ meson, seen at $m_{ee} = 1 \text{ GeV}/c^2$, does not emerge from newer model calculations (see Section 5.4).

in detail in Section 5.1. Other forms of the hadronic cocktail which are included in previous experimental results are shown in the following section¹³.

2.2. Previous Measurements

Experiments at two particle accelerators prior to the LHC have been able to measure nuclear matter at energy densities that are sufficient for the creation of a quark–gluon plasma [Bar08]. The Super Proton Synchrotron (SPS) at CERN provides heavy-ion beams with up to 158 GeV per nucleon to multiple fixed-target experiments¹⁴, which together achieved “compelling evidence [for] a new state of matter” [HJ00, and refs. therein]. And experiments at the Relativistic Heavy Ion Collider (RHIC) at BNL, measuring gold–gold collisions up to $\sqrt{s_{NN}} = 200$ GeV, established that the quark–gluon plasma behaves like a nearly perfect fluid [Shu04, Mül07, and refs. therein]. The main dilepton results of CERES and NA60 at the SPS and selected ones of the two large RHIC experiments PHENIX and STAR are summarized here.

2.2.1. CERES

The CERES/NA45 experiment is designed for the measurement of low-mass dielectrons, originally using two ring-imaging Cherenkov (RICH) counters for electron identification in combination with two silicon detectors near the target for charged-particle tracking and vertex reconstruction, which allowed for the observation of a dielectron excess in sulphur-gold collisions [CERES95]. First results in heavy-ion (Pb–Au) collisions using improved tracking detectors confirmed that this excess, which is most prominent in the mass region $0.3 < m_{ee} < 0.7$ GeV/ c^2 but reaches beyond 1 GeV/ c^2 , increases stronger-than-linear with the charged-particle multiplicity [CERES98] — consistent with a quadratic rise that is expected for radiation from thermal processes in a hot and dense medium [CERES95]. The CERES experiment was subsequently upgraded with a Time Projection Chamber for optimized tracking and additional electron identification, leading to a much improved invariant-mass resolution capable of separating the ρ/ω and ϕ resonances [CERES08]. Figure 2.7 shows the results obtained with this upgraded detector setup for the 7% most central collisions. In the left panel, the dielectron yield normalized by the number of

¹³Note that in measurements at the SPS, where the charm cross section is three orders of magnitude lower than at LHC energy [ALICE12a], heavy-flavour contributions were small and left to be treated or neglected by the theoretical models (see e. g. [vHR08]).

¹⁴For comparison to the energies reached at collider facilities, the center-of-mass energy of the collision can be computed from the beam energy and the target (nucleon) mass via $\sqrt{s} \cong \sqrt{2 \cdot m_N \cdot E_{\text{beam}}}$, leading to an energy per nucleon-nucleon collision of $\sqrt{s_{NN}} = 17.3$ GeV at the SPS [SSS10].

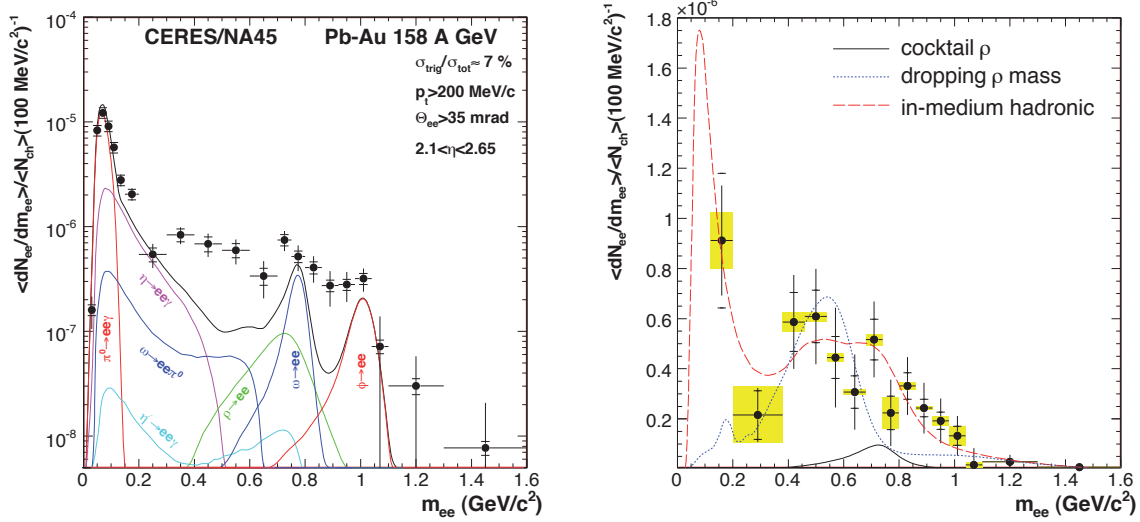


Figure 2.7.: Low-mass dielectron yield in the 7% most central Pb–Au collisions at $\sqrt{s_{\text{NN}}} = 17.3 \text{ GeV}$ measured by the CERES experiment. Left panel: total dielectron signal in comparison to the hadronic cocktail. Right panel: dielectron excess yield after cocktail (except ρ) subtraction, compared to two scenarios of additional thermal dilepton radiation (see text for details). From [CERES08].

charged particles ($\langle N_{\text{ch}} \rangle \approx 177$ in the acceptance) is compared to the hadronic cocktail, which contains the improved mass resolution. No heavy-flavour contribution is included. Data and cocktail are in very good agreement in the π^0 -Dalitz region; also the pair cut on the opening angle, leading to vanishing yield towards zero mass, is well-described by the cocktail. But starting from $m_{ee} \approx 0.2 \text{ GeV}/c^2$, the cocktail consistently underestimates the observed yield, most pronounced in the η - and ω -Dalitz region and between the ω and ϕ resonances.

The right panel presents on a linear scale the dielectron excess yield after subtraction of all hadronic sources except for the ρ contribution, which is shown as “cocktail ρ ”. This excess spans the whole low-mass region, with a total enhancement over the cocktail by a factor of ≈ 2.45 within $0.2 < m_{ee} < 1.1 \text{ GeV}/c^2$ [CERES08]. The measurement is compared to model calculations that employ a realistic space–time evolution of the collision starting at $T_0 = 190 \text{ MeV}$ [RW99] and two scenarios of ρ modification, one featuring a broadening of the spectral function as discussed above (“in-medium hadronic”) and the other a dropping of the ρ mass depending on the medium density as suggested by Brown-Rho scaling [BR91]. While the calculations do include dileptons from QGP radiation and the Drell–Yan process [CERES08], the observed excess is clearly dominated by the hadronic contribution at this collision energy. For $m_{ee} < 0.7 \text{ GeV}/c^2$ both scenarios can describe the measurement, but for higher masses only the in-medium spectral function




Figure 2.8.: Low-mass dimuon measurement by NA60 in In–In collisions at $\sqrt{s_{\text{NN}}} = 17.3$ GeV. Left panel: total dimuon signal (open circles) and hadronic cocktail (solid lines), as well as the excess signal obtained by subtraction (full triangles). Right panel: dimuon excess in semi-central collisions compared to two model predictions for thermal radiation and a scaled cocktail ρ (see text for details). Systematic uncertainties (not shown) are on the order of 25 %. From [NA6006].

is in agreement with the data, especially between the ω and ϕ mass, where systematic uncertainties are small.

2.2.2. NA60

The NA60 experiment follows an independent approach by measuring low-mass lepton pairs in the dimuon channel. It consists of a silicon tracker inside a dipole magnet near the target, followed by a hadron absorber and a muon spectrometer [NA6006]. Measurements in the dimuon channel benefit from a clean track sample due to the absorber. This in particular allows for triggering on (di)muons reaching the spectrometer, such that each recorded event actually contains a muon pair to analyze. One disadvantage compared to dielectrons is the reduced accessible phase-space due to the considerable muon mass.

Figure 2.8 shows the low-mass dimuon spectrum obtained by NA60 from indium-indium¹⁵ collisions at $\sqrt{s_{\text{NN}}} = 17.3$ GeV [NA6006]. The measured signal in minimum-bias collisions after subtraction of combinatorial background and fake matches is given in the left panel.

¹⁵With a mass of 115 u, indium also is a heavy ion and central In–In collisions lead to an average charged-particle multiplicity density of $\langle dN_{\text{ch}}/d\eta \rangle \approx 200$, while in Pb–Au collisions $\langle dN_{\text{ch}}/d\eta \rangle \approx 350$ are reached [NA6006, CERES98]. When the above Pb–Au data from CERES was taken, the NA60-predecessor experiment NA50 recorded Pb–Pb collisions mostly to study J/ψ suppression, but it lacked good acceptance for low-mass resonances [NA5005, NA5003].

From the fine mass resolution it is evident that the hadronic cocktail — irrespective of its absolute normalization — cannot saturate the data. Its contributions are therefore conservatively normalized and subtracted such that the remaining excess yield receives a smooth distribution and vanishes at the mass threshold.

The right panel shows the excess for semi-central collisions in comparison to model predictions. Full and open markers represent alternative η subtractions as a systematic variation. The theoretical spectra are normalized to the data by their integrals within $0.2 < m_{\mu\mu} < 0.9 \text{ GeV}/c^2$. As a vacuum baseline, the cocktail ρ (thin solid line) is upscaled (dashed line). The dashed-dotted line corresponds to the dropping-mass scenario and the solid line to the broadening scenario. Also included is the contribution from correlated charm decays, showing that it is almost negligible at SPS energy (thin dashed line).

In contrast to the CERES data, the NA60 dimuon excess exhibits a rather symmetrical broadening around the ρ mass¹⁶. This behaviour in combination with the high experimental precision rules out the dropping-mass scenario and challenges other models regarding the description of the upper part at $m_{\mu\mu} > 0.9 \text{ GeV}/c^2$ [NA6006]. Subsequently improved models are able to describe this part of the excess by including also small thermal contributions from the ω and ϕ mesons, four-pion annihilation, and some QGP radiation of similar size as the charm contribution, while remaining consistent with the CERES data [NA6007, vHR08].

2.2.3. PHENIX

PHENIX is one of the two large general-purpose experiments at RHIC. Both are designed rather symmetrical around their respective interaction point, in contrast to the fixed-target experiments at the SPS [Bar08]. The PHENIX detector is well suited both for dielectron and for dimuon measurements [PHENIX03]: The central arm spectrometers around midrapidity provide charged-particle tracking via drift and pad chambers, as well as electron identification using a RICH detector, an electromagnetic calorimeter (EMCal), and time-of-flight (TOF) information from the EMCal and from a dedicated TOF-detector. At forward and backward rapidity, two large muon spectrometers and detectors for triggering, timing and centrality determination are installed.

The main dielectron measurement by PHENIX with its techniques and results is described extensively in [PHENIX10]. Amongst other findings, the low-mass enhancement seen at the SPS was confirmed, but its increased strength could not be saturated by model

¹⁶Note however that the smaller excess on the low-mass side is caused by the falling acceptance of the muon spectrometer towards low pair- p_T . Model calculations are filtered by the respective experimental acceptance to account for this effect [NA6006].

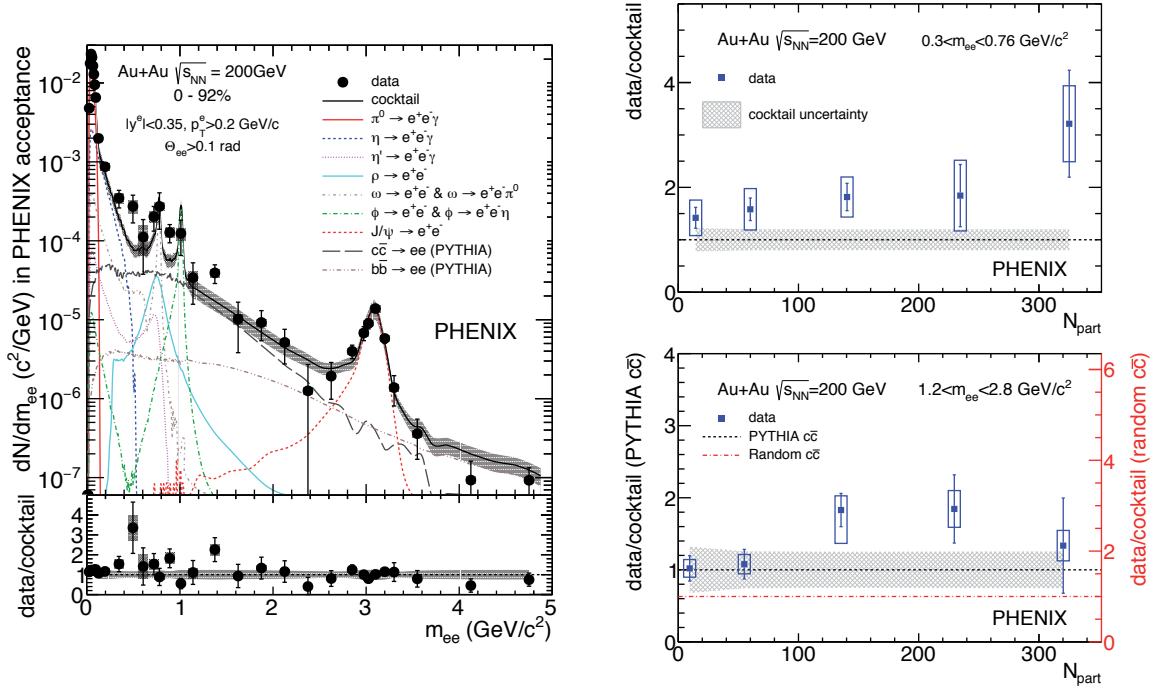


Figure 2.9.: Dielectron results from PHENIX in Au–Au collisions at $\sqrt{s_{NN}} = 200$ GeV. Left panel: total invariant-mass distribution in comparison to the hadronic cocktail. Right panels: data-to-cocktail ratio in the low-mass (upper) and intermediate-mass (lower) region as a function of centrality. Each region shows a dielectron excess, with their magnitudes sensitive to the $c\bar{c}$ generator (see text for details). From [PHENIX16].

calculations. The experiment was then upgraded with a hadron-blind detector (HBD), which specifically benefits the dielectron measurement, in order to address the open questions [PHENIX16]. The HBD adds another point of electron identification, placed close to the interaction point, which reduces hadron contamination from track mismatches and in particular allows for the exclusion of low-mass pairs from the analysis and thereby reduces the combinatorial background.

The left panel of Figure 2.9 shows the updated PHENIX dielectron measurement for minimum-bias Au–Au collisions at $\sqrt{s_{NN}} = 200$ GeV. In comparison to the SPS results, the mass range is increased up to 5 GeV/c^2 , thereby including the entire intermediate-mass region and the prominent J/ψ resonance. Also note the substantial contributions from open heavy-flavours denoted by $c\bar{c}$ and $b\bar{b}$.

In order to quantify the enhancement observed in the low-mass region and to investigate the intermediate-mass region, the analysis is divided into five centrality classes and the yields are compared to the respective cocktail expectation. This is presented in the right panels, where the data-to-cocktail ratio is given for two invariant-mass intervals as a

function of the number of participants. In the low-mass region, enhancement factors of ≈ 1.4 – 3.2 with increasing centrality are found, which can be described by thermal radiation from QGP and hadron gas using the ρ -broadening scenario [PHENIX16]¹⁷. The intermediate-mass region also shows an excess above the cocktail, which is again quite sensitive to the heavy-flavour contributions in the cocktail, primarily the one from correlated charm decays. The default baseline using PYTHIA corresponds to a proton–proton reference scaled by the number of binary collisions N_{coll} and the alternative y-axis is based on total kinematic decorrelation of the open-charm mesons. These two cases represent opposite extremes for the charm contribution, with the reality expected in between, leaving room for a significant additional contribution like thermal QGP radiation [PHENIX16].

2.2.4. STAR

The STAR experiment uses a mainly cylindrical detector geometry with full azimuthal and large rapidity coverage, designed to measure a wide range of observables and thereby achieve a broad understanding of high-energy heavy-ion collisions [STAR03]. Tracking of charged particles is performed by silicon detectors near the beam axis and a large Time Projection Chamber (TPC), which also provides particle identification (PID). Additional PID information is obtained by a RICH, TOF and EMCal, all originally with small solid-angle coverage [STAR03]. In the forward and backward direction, a radial-drift TPC extends the tracking capability and further small detectors provide trigger, time and vertex information. Centrality determination in Au–Au collisions is done via the measured charged-particle multiplicity density [STAR15b].

Following an upgrade of the TOF to cover the full TPC acceptance, the STAR detector is also well suited for low-mass dielectron measurements, since the combined PID information from TPC and TOF allows for effective separation of electrons from hadrons at low momenta [STAR15b]¹⁸. With this improved experimental setup, STAR participated in the RHIC beam energy scan program and collected Au–Au data at $\sqrt{s_{\text{NN}}} = 19.6, 27, 39, \text{ and } 62.4$ GeV [STAR15a, STAR18].

The dielectron invariant-mass spectra for minimum-bias Au–Au collisions¹⁹ at all avail-

¹⁷When a different generator (MC@NLO instead of PYTHIA) is used to compute the $c\bar{c}$ contribution, the low-mass enhancement factors are lower by $\approx 40\%$. In the intermediate-mass region, both generators produce charm yields consistent within 10% [PHENIX16].

¹⁸This PID strategy is also common in ALICE, where a similar set of detectors is available in the central barrel (see Section 3.2), and discussed later in this work (see Section 4.2.2 and Section ??).

¹⁹The minimum-bias trigger in STAR corresponds to the 0–80% most central events, while in PHENIX it covers 0–92% centrality [STAR15b].

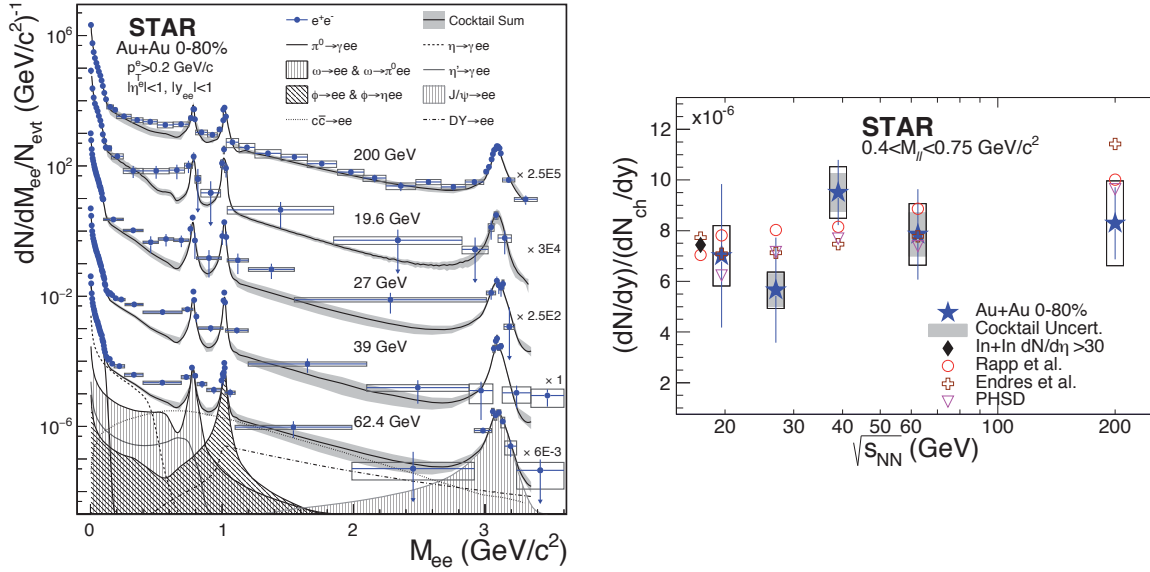


Figure 2.10.: Dielectron results from STAR in minimum-bias Au–Au collisions at multiple collision energies. Left panel: invariant-mass spectra for each energy in comparison to the cocktail expectation. Yields are scaled by the indicated factors for better visibility. Right panel: collision-energy dependence of the low-mass excess yield, normalized by dN_{ch}/dy , compared to different model calculations and to the NA60 measurement in In–In collisions. From [STAR18].

able energies, including the measurement at nominal RHIC energy, are presented in the left panel of Figure 2.10. The individual results are shifted as indicated for sufficient separation. Each measurement is compared to the respective hadronic cocktail, using N_{coll} -scaled PYTHIA simulations for the open-charm contribution [STAR18]. A low-mass enhancement is visible at all collision energies, whereas the remaining spectrum is well-reproduced by the cocktail expectation. The enhancement factor at $\sqrt{s_{\text{NN}}} = 200$ GeV amounts to ≈ 1.7 [STAR15b], which is in agreement with the (updated) PHENIX measurement, and the result at 19.6 GeV is consistent with the In–In data from NA60 at 17.3 GeV [STAR18].

The right panel shows the low-mass excess yield above the hadronic cocktail as a function of the collision energy. Yields are normalized by the charged-particle multiplicity density to remove the system-size dependence, which also allows for the inclusion of the NA60 measurement at $\sqrt{s_{\text{NN}}} = 17.3$ GeV. The STAR data shows no significant energy dependence for the dilepton excess in minimum-bias Au–Au collisions, while different model calculations indicate a moderate increase of the excess with collision energy [STAR18]. Data and models are however in agreement within statistical uncertainties and precise measurements over an extended energy range will be useful to further study the nature

of the low-mass enhancement.

STAR also investigates the sensitivity of their data with respect to a decorrelation or suppression of electrons from open-heavy-flavour hadrons, which should lead to a modification of the charm contribution in nuclear collisions. Each tested scenario causes a softening relative to the original charm contribution, resulting in a better agreement between data and hadronic cocktail in central Au–Au collisions, where the nuclear modification is strongest [STAR15b].

Due to the dominance of the open-heavy-flavour contribution in the intermediate-mass region, a precise knowledge of its yield is a prerequisite for studies that aim to measure thermal radiation. A more recent upgrade of the STAR experiment, which includes a micro-vertex detector to resolve displaced vertices from heavy-flavour decays and a muon detector with 45% azimuthal coverage around midrapidity [STAR13], allows for a measurement of electron-muon correlations in a common rapidity window. This gives exclusive access to the open-heavy-flavour yield, offering the possibility to then subtract this contribution from a dielectron or dimuon measurement in order to extract a thermal excess yield [STAR14].

Chapter 3.

Experiment

The present work is set in the environment of ALICE¹, the dedicated heavy-ion experiment at CERN's Large Hadron Collider (LHC). In this chapter the experimental setup is described with a focus on the systems that are most relevant to this thesis. Some of their operational properties affecting the data used in the following analysis are also presented.

3.1. LHC

The main purpose of the LHC is to produce proton–proton (pp) collisions at the energy frontier achievable in a laboratory, in order to enable searches for undiscovered or yet unknown particles. In particular, the discovery of the Higgs boson was a set goal for the LHC physics program since its initiation in 1984 [SS15].

Besides this focus on high-energy particle physics, the LHC program also includes the field of heavy-ion physics to study the properties of strongly-interacting matter. To serve the latter, the LHC delivers one month of either lead–lead (Pb–Pb) collisions or proton–lead (p–Pb) collisions per year of operation.

Protons and ions are accelerated by a radio-frequency system based on superconducting cavities. It operates at a frequency of 400.8 MHz, which is constrained by the beam parameters — in particular the bunch length — at the injection energy and those required for beam storage during physics operation. The RF system increases the beam energy by ≈ 500 keV per revolution during ramp-up, while the energy loss due to synchrotron radiation amounts to only 7 keV per revolution at the highest energy. Most of the system's power is however required and reserved to control the beams and maintain stable beam conditions. [EB09]

¹A Large Ion Collider Experiment

The maximum beam energy of the LHC is more strictly determined by the bending power of its dipole magnets to keep the particles on a circular trajectory. It can therefore be estimated by equating the centripetal with the Lorentz force acting on the beams [Wie07], resulting in

$$E_{\text{beam}} \approx q \cdot B \cdot r \cdot c. \quad (3.1)$$

With the dipoles' nominal magnetic field $B_0 = 8.33 \text{ T}$ and their bending radius $r_b = 2804 \text{ m}$, a proton ($q = 1$) can reach $E_p = 7 \text{ TeV}$ [EB09]. The maximum energy of fully stripped $^{208}_{82}\text{Pb}$ ions is $E_{\text{Pb}} = 574 \text{ TeV}$ or, when given in multiples of the energy per nucleon, $E_{\text{Pb}} = (\frac{Z}{A} \cdot E_p) \cdot A = 2.76 \text{ ATeV}$. The corresponding center-of-mass energies are $\sqrt{s} = 2 \cdot E_{\text{beam}} = 14 \text{ TeV}$ in pp collisions and $\sqrt{s_{\text{NN}}} = 5.52 \text{ TeV}$ per nucleon-nucleon collision in Pb–Pb.

During LHC RUN 1 from 2009 to 2012, the maximum beam energy was limited to half of its nominal value for technical reasons², resulting in $\sqrt{s_{\text{NN}}} = 2.76 \text{ TeV}$ for the Pb–Pb data from 2011 [FL12] that is analyzed in the present thesis.

Both the searches for “new” particles in pp and the study of QGP signatures in Pb–Pb demand a large number of collisions to be provided by the LHC, from which those containing rare processes or producing the most particles can be selected and analyzed. Therefore not only the maximum beam energy of the accelerator is relevant, but also its luminosity \mathcal{L} , which is a measure for the intensity of its beams at their interaction point. Together with the cross section of the event category under study, σ_{event} , it determines the number of events per second according to

$$n_{\text{event}} (\text{s}^{-1}) = \mathcal{L} \cdot \sigma_{\text{event}}. \quad (3.2)$$

For two identical beams with N_b particles per bunch, n_b bunches per beam (i. e. bunch-bunch interactions per revolution), and the revolution frequency f_{rev} , the luminosity is given by

$$\mathcal{L} = \frac{N_b^2 n_b f_{\text{rev}}}{4\pi\sigma_x\sigma_y} \cdot F \cdot W \quad (3.3)$$

²A review of the cause is given in [SS15]. In 2012 the beam energy was slightly increased from 3.5 TeV to 4 TeV. During RUN 2 from 2015 to 2018, the LHC delivered $E_p = 6.5 \text{ TeV}$ and $E_{\text{Pb}} = 2.51 \text{ ATeV}$ [S⁺19, J⁺19], close to the design values.

[SS15, EB09]. This assumes Gaussian particle distributions in the bunches, with widths of σ_x and σ_y in the transverse direction. F and W are geometric luminosity reduction factors, which account for the necessary finite crossing angle of the beams and an optional transverse offset, respectively. The latter is used at the ALICE interaction point, in combination with increased beam widths, to reduce the luminosity in pp collisions from $\mathcal{L}_{\text{peak}} \approx 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ to $\mathcal{L}_{\text{ALICE}} \approx 10^{27} \text{ cm}^{-2}\text{s}^{-1}$ [SS15].

The ALICE design allows for the same luminosity also in Pb–Pb collisions, and in 2011 a peak luminosity $\mathcal{L} \approx 5 \cdot 10^{26} \text{ cm}^{-2}\text{s}^{-1}$ was delivered to the experiment. Considering the reduced beam energy, which leads to less adiabatic damping of the beam emittance, the achieved luminosity actually exceeded the LHC’s design performance [JAF⁺12, Hin08].

3.2. ALICE

Being the only LHC experiment dedicated to heavy-ion physics, ALICE is designed as a general-purpose detector to address a large variety of observables. It consists of 18 sub-detector systems to provide manifold information on the measured particles, in particular by employing essentially all available particle identification (PID) techniques. In combination with a solenoid magnet housing the central barrel of ALICE and a dipole in the forward direction, these detectors also provide momentum measurements ranging from below 100 MeV/ c to beyond 100 GeV/ c . [ALICE09]

Figure 3.1 shows the layout of the ALICE experimental setup, which is arranged around the interaction point of the LHC beams that is located in the left half of the figure. The innermost detector surrounding the interaction point is the Inner Tracking System (ITS), followed by the Time Projection Chamber (TPC). These detectors are specialized on tracking and particle identification³ and are used in the present work for measuring electrons at low momenta up to a few GeV/ c .

The ITS consists of six silicon detector layers at radial positions from 3.9 cm to 43 cm. Its two inner layers have the finest granularity to provide a precise extrapolation of tracks to the event vertex, reaching a DCAr⁴-resolution of $\sigma_{\text{DCAr}} = 60 \mu\text{m}$ for global tracks with $p_{\text{T}} = 1 \text{ GeV}/c$ in Pb–Pb collisions [ALICE14c]. The other four layers are crucial for ITS stand-alone tracking of soft particles and to match trajectories between ITS and TPC. These layers also provide PID information for charged particles, by analogous readout of the charge deposition due to their specific energy loss dE/dx in the silicon sensors.

³Detector properties mentioned in this chapter refer to the original ITS and TPC setup used during RUN 1. See Section 6.2 for a description of their upgraded versions for RUN 3.

⁴Distance of closest approach to the vertex in the radial direction.




Figure 3.1.: The ALICE detector setup, sliced open to view the interaction point and the individual subdetectors surrounding it. Detectors used in the present work are highlighted: Inner Tracking System (green), Time Projection Chamber (blue), and Time-Of-Flight (orange).

The TPC is the main tracking device of ALICE and provides primary PID information for charged particles. In contrast to all other sub-detectors, it consists of a single large active volume filled with a noble gas plus quencher mixture⁵, providing continuous three-dimensional tracking of charged particles between 85 cm and 247 cm radial distance from the event vertex. The primary ionization drifts towards the readout plane at the end caps of the cylindrical TPC volume, where it is amplified and measured by multi-wire proportional chambers. The average amount of deposited ionization depends on the particle species and momentum according to the Bethe-Bloch formula for the specific energy loss dE/dx . Favoured by the long trajectories through the active volume and correspondingly many dE/dx samples, a resolution of $\sigma_{dE/dx} = 6.5\%$ is reached in the most central Pb–Pb collisions, allowing the TPC to identify these particles (mostly electrons, pions, kaons and protons) on a track-by-track or statistical basis over most of the momentum range from $p \approx 0.2$ to 20 GeV/ c [ALICE14c].

Two more detector systems cover the full azimuthal angle and polar angles comparable to

⁵A mixture of Ne/CO₂ (90/10) was used during RUN 1, optimized for fast drift speed, low diffusion and small material budget. An addition of N₂ (5 parts) was used temporarily to improve quenching and stability [ALICE14c, ALICE13c].

ITS and TPC, namely the Transition Radiation Detector (TRD) and the Time-Of-Flight detector (TOF). While the former is a dedicated electron-identification detector, it is best suited for momenta above $1 - 2 \text{ GeV}/c$ and was not yet fully installed during RUN 1, such that it is not used here. The latter is primarily designed to distinguish between hadron species, but can serve to separate electrons from these hadrons in the low-momentum range, and thereby supplements or enables either of the two PID strategies in this work. In the periphery of the central barrel, a Cherenkov detector (HMPID) is installed to extend hadron identification to high momentum in a limited acceptance. Two electromagnetic calorimeters (PHOS and EMCal) allow for measurements of photons, neutral mesons, and jets. A scintillator array (ACORDE) on top of the solenoid is used to trigger on cosmic rays, mainly for alignment purposes. Located in the forward direction is the muon spectrometer, consisting of tracking and trigger stations, absorbers to block other particles, and the dipole magnet.

Several dedicated trigger detectors are positioned near the beampipe, providing the minimum-bias and multiplicity triggers (V0), the event start time (T0), and in case of Pb–Pb collisions the centrality trigger (ZDC and V0). Forward multiplicity detectors for charged particles (FMD) and photons (PMD) largely extend the rapidity coverage for measuring the event multiplicity. Together with the ZDC they also provide three complementary estimates for the reaction plane of a collision.

All details on the original detector setup of ALICE can be found in [ALICE09] and the Technical Design Reports referenced therein. The experimental performance achieved throughout RUN 1 is summarized in [ALICE14c]. Some aspects of the ITS and TPC performance during the heavy-ion period of 2011 are described in the following.

3.2.1. Detector Performance

Over an extended period of data taking, the operational parameters of a detector cannot be expected to remain constant. They are thus tracked while recording data and factored in when the events are being reconstructed, such that most variations can be hidden from the actual data analysis. In case of larger deviations from the expected performance, their influence on the analysis as well as corresponding reactions should be considered. While taking the 2011 Pb–Pb data, cooling issues in the ITS led to an overall reduction of its operational efficiency, as shown in the left panel of Figure 3.2. In particular, about 30 % of the Silicon Pixel Detector, which constitutes the two inner ITS layers, were inoperational during that period (white and dark blue ring segments). For comparison, the operational efficiency of the TPC is given in the right panel. Even though some of

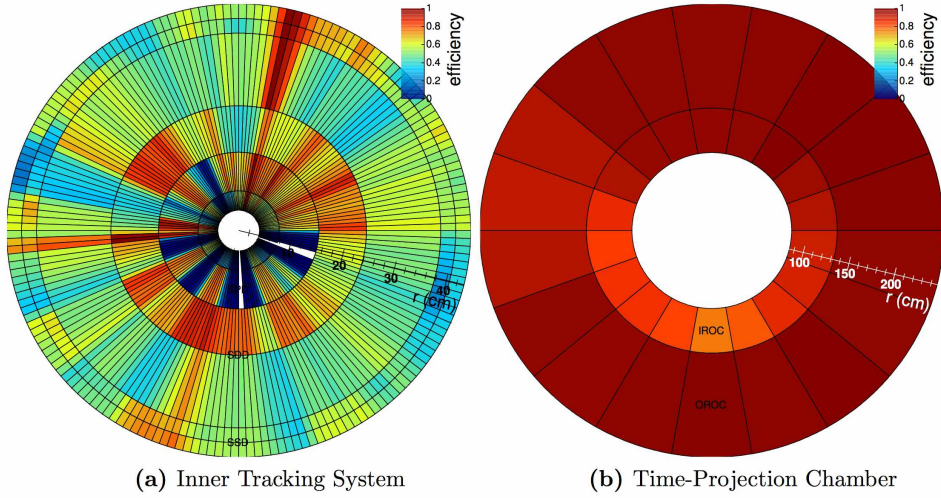


Figure 3.2.: ITS and TPC operation efficiencies during the Pb–Pb period of 2011 as a function of azimuth and radius. In case of the ITS, the inner radius of each ring gives the radial position of the corresponding detector layer. From [Boo14].

the inner and outer readout chambers had temporary high-voltage issues [Boo14], this only leads to a slight reduction of the overall efficiency.

The partial unavailability of the inner ITS layers has a significant impact on the following analysis and is therefore explicitly incorporated into the tracking strategy (Section 4.2.1), while the variation of the TPC performance is covered by the reconstruction process and reflected in the track efficiencies, for which the data will be corrected (Section 4.6).

3.2.2. PID Calibration

The PID measurements by the various detectors are centrally calibrated and provided in the form of their deviations from the expectation values of the most relevant particle species. They are normalized by the detector resolution $\sigma_{\text{det.}}$, such that each measured value is given in multiples of a standard deviation σ . For an electron measurement via specific energy loss in the ITS or the TPC, this becomes

$$n\sigma_e^{\text{det.}} = \frac{(dE/dx)_{\text{meas.}} - \langle dE/dx \rangle_e}{\sigma_{\text{det.}}} \quad (3.4)$$

This centralized calibration procedure thereby provides an important baseline for the following analysis, but it has some deficits in the description of the electron expectation. Besides its natural dependence on the electron momentum, the nominal value $\langle dE/dx \rangle_e$ as well as its resolution also depend on the pseudorapidity for geometrical reasons and on the event multiplicity because of increasing detector occupancy.

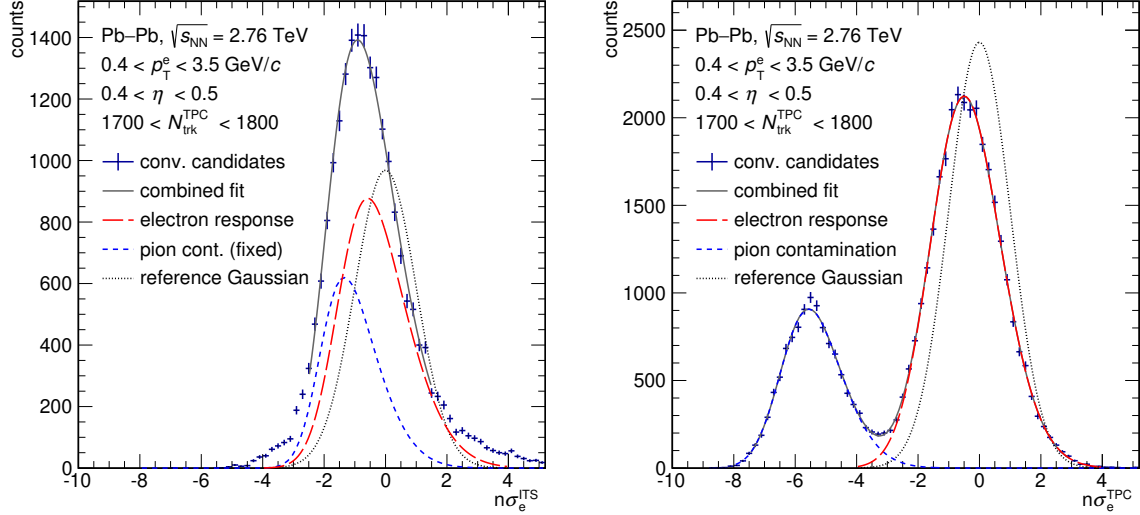


Figure 3.3.: Left panel: ITS PID response for conversion electron candidates in a small pseudorapidity and multiplicity interval. The residual pion contamination is extracted from a clean sample and fixed for the fit. The resulting electron response is slightly offset and tilted with respect to the reference Gaussian. Right panel: corresponding TPC PID response, where the electron signal is wider than the reference and shifted, but has negligible skewness.

A method to monitor these variations in the TPC was developed in [Boo14] and is adopted here: To obtain a relatively clean electron sample, tracks from photon conversions $\gamma + X \rightarrow e^+e^-$ are selected, if they fulfill similar quality criteria as the electron candidates of the analysis, except for hits in the ITS layers. The PID response for these particles is then evaluated differentially in the relevant variables and a correction map is constructed. This map is used during the final analysis for post-calibration of the PID response prior to the electron selection (Section 4.2.2). Since also ITS PID is used within this work, the post-calibration is performed for both detectors.

Figure 3.3 shows the ITS and TPC responses for the conversion track sample in the rapidity range $0.4 < \eta < 0.5$ and TPC multiplicity range $1700 < N_{\text{trk}}^{\text{TPC}} < 1800$, which corresponds to a centrality class of $\approx 8\text{--}10\%$. The track sample is integrated over the transverse-momentum range used in the analysis, presuming that the p_T dependence is correctly calibrated, which is approximately true. This integration leads to a notably wider distribution for the residual pion contamination in case of the TPC, due to the gradual relativistic rise of their energy loss in the detector gas, whereas electrons have reached the Fermi-plateau in this momentum range. The total distribution is well-described by a free fit with two generalized Gaussians, each including a skewness parameter to account for




Figure 3.4.: Mean (top) and width (bottom) of the TPC electron PID response as a function of pseudorapidity and multiplicity. Left panels: parameters obtained via Gaussian fits. Right panels: smoothed parameter distributions used for PID post-calibration during the analysis. See text for details.

the slight tilt of the bell shapes. The deviation of the resulting electron parameterization from its nominal shape and position is seen by comparison to the reference Gaussian. Since the relativistic rise is minor in a solid, the electron and pion distributions largely overlap in case of the ITS, as shown in the left panel of Figure 3.3. Therefore an additional clean sample of pions is constructed by pre-selecting them in the TPC with otherwise identical cuts and parameterized to fix their distribution in the ITS. The electron response is then extracted from a combined fit. This procedure is repeated for all needed pseudorapidity and multiplicity intervals and the parameters of the electron responses are collected.

Maps of the resulting means and widths of the electron response in the TPC are given in the top and bottom left panel of Figure 3.4, respectively. The mean varies from $+0.4\sigma$ at midrapidity in low-multiplicity events⁶ to -0.6σ at $\eta = \pm 0.5$ in central collisions. A similar behaviour is observed for the width of the electron distribution, which for any given pseudorapidity interval increases by 6–10% over the analyzed multiplicity range.

⁶The multiplicity range analyzed here corresponds to a centrality selection of 0–50%, covering all events of the following analysis.

In both cases the multiplicity dependence is linear, and the pseudorapidity could be approximated by a fifth-order polynomial. While the former is reasonable because of an increasing occupancy effect, there is no deeper argument for the latter to precisely match a polynomial. Therefore in the next step the multiplicity dependence is individually fitted for every pseudorapidity interval and then interpolated in η to create a smooth correction map for each parameter, as shown on the right side of Figure 3.4. The corresponding maps for the ITS can be found in Appendix A.

For both detectors, the residual difference between the corrected dE/dx response and its ideal values is below $\pm 0.03\sigma$ for both the mean and the width of the dE/dx distribution. This allows for a reliable electron selection and self-consistent efficiency correction⁷ in the following analysis.

⁷Electron responses in the used Monte-Carlo data are simulated such that they represent perfectly calibrated dE/dx distributions.

Chapter 4.

Analysis

In this chapter the dielectron analysis performed within this work is described. First the event-based analysis is discussed, mainly consisting of the electron selection and the pairing procedure. This leads to the statistical analysis of the obtained electron–positron pairs, including a discussion on measures for signal quality and their application to improve the analysis by using pair prefilters. After describing the efficiency correction and evaluation of systematic uncertainties, the chapter closes with a comparison between the result of this work and previously published data.

Particular techniques used in this analysis, such as the pair efficiency factorization and prefilter efficiency determination, as well as the unconventional PID strategy, have been initially developed and tuned on semi-central collisions. Within this chapter though, the 0–10% centrality class is used for presentation of the analysis process. Most effects and properties to consider, such as detector resolution and background size, are more pronounced in central collisions. This, along with significantly larger electron statistics¹ compared to the 10–50% centrality class, is beneficial for many of the figures and their discussion.

4.1. Events

The data used in this analysis was taken during LHC RUN 1, recorded over four weeks of heavy-ion beamtime in fall 2011. Since the conditions for physics operation (“runs”) and the resulting data properties vary within this period, only a subset of runs is suitable for a dielectron analysis. A diligent run selection for a similar analysis ($J/\psi \rightarrow e^+e^-$) was performed in [Boo14] and is adopted here. It allows for more variations in the detector

¹The term “statistics” is used in this work (and often in our community) to describe the size of a data sample, rather than referring to the respective field of mathematics.

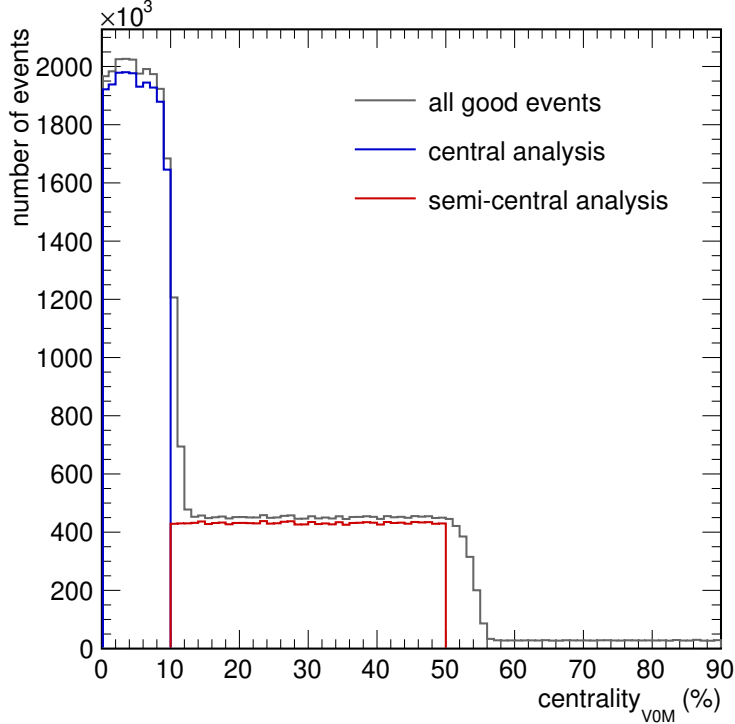


Figure 4.1.: Centrality distribution of all good events in the data set used for this analysis and the distributions of analyzed events in the two selected centrality intervals.

performance than the centrally provided run selection, yet ensuring that they do not compromise the measurement of electrons, which increases the number of events available to this analysis.

4.1.1. Event Selection

Figure 4.1 shows the centrality distribution of different event samples in the analyzed data set. All events with a reconstructed vertex within ± 10 cm around the nominal collision point in beam direction, including a pileup rejection, are presented as “good events”. Two subsets of these are selected for separate analysis of the 0–10 % and 10–50 % most central collisions, which we refer to as *central* and *semi-central* collisions. They correspond to the dedicated centrality triggers used during this period of data taking. The reduction compared to all good events is mainly caused by the success rate of the analysis computing jobs, which was $\approx 98\%$ in both cases. For central collisions, the distribution of events is not perfectly uniform in centrality, but this has no significant impact on the dielectron yield. 19.1 million central and 17.3 million semi-central events are contained in the final results.

4.1.2. Event Mixing

Mixed-event techniques are used to determine whether some relation between observables is caused by an immediate correlation between them or emerges from their properties also when probed individually. Event mixing removes any immediate correlation, while external influences may result in a residual correlation to persist in mixed-events.

In a dielectron analysis, the relation (or difference) between unlike-sign and like-sign pairs is caused by the different signal content primarily, but also by detector acceptance effects. Since the dielectron signal is a pure same-event occurrence, while acceptance effects also influence mixed-event pairs, event mixing is used here to study these effects. To ensure that mixed-event pairs reproduce the properties of (uncorrelated) same-event pairs, different event classes are defined based on quantities to which the pairing is sensitive. The used variables and their intervals are given in Table 4.1. A mixing depth of 15 is chosen, i. e. within each of the 120 event classes, each event is mixed with 14 other events to provide large statistics of mixed-event pairs.

event variable	intervals
z of primary vertex	$[-10, -5, 0, +5, +10]$ cm
centrality	$[0, 5, 10, 20, 30, 50]$ %
reaction plane (TPC)	$[-\frac{\pi}{2}, -\frac{\pi}{3}, -\frac{\pi}{6}, 0, +\frac{\pi}{6}, +\frac{\pi}{3}, +\frac{\pi}{2}]$

Table 4.1.: Variables and intervals defining the event-mixing classes.

Figure 4.2 shows the acceptance correction factor R for the nominal settings of the present analysis as a function of invariant mass. Details on its definition and the used variables are given within the pair analysis description in Section 4.3.1. Since R is sensitive to the magnetic field of the experiment, it is extracted separately for the two field configurations used during data taking and applied to the corresponding subset of data. The number of mixed-event unlike-sign and like-sign pairs shows a relative difference of up to 7% at very low invariant mass, which quickly falls off and vanishes for $m_{ee} > 0.5 \text{ GeV}/c^2$. To avoid propagation of statistical fluctuations in R to the final signal in kinematic regions where the event-mixing statistics become limited, R is set to unity for $m_{ee} > 1.5 \text{ GeV}/c^2$ in this analysis.




Figure 4.2.: Acceptance correction factor R for the two magnetic field configurations used during data taking.

4.2. Tracks

The track selection and particle identification procedures are described in this section. They both aim at accepting a very large fraction of all reconstructed electrons and positrons² for further analysis. This strategy is chosen to challenge the limited dielectron statistics contained in the available data set.

A track consists of a reconstructed particle trajectory and a collection of properties measured by any of the detectors, whose responses were associated to this trajectory. Two classes of track samples are defined in this analysis, which serve distinct purposes:

- *Electron candidates* is the term used for tracks which fulfill all criteria required to enter the dielectron signal extraction.
- *Prefilter electrons* refers to an extended track sample with relaxed requirements on kinematic, track quality and particle identification properties, which is used to identify and reject undesired dielectron sources.

²The explicit mentioning of positrons is situationally omitted, where the context does not require it.

4.2.1. Track Selection

Track Kinematics

The kinematic range in which electron candidates are selected is $0.4 < p_T^e < 3.5$ GeV/ c and $|\eta^e| < 0.8$ in full azimuth, while for prefilter electrons it is extended to $0.05 < p_T^e < 3.5$ GeV/ c and $|\eta^e| < 0.9$. The lower transverse-momentum limit for track reconstruction in the ALICE central barrel at its nominal magnetic field of $B = 0.5$ T is $p_T \approx 120$ MeV/ c for global tracks and $p_T \approx 70$ MeV/ c for ITS stand-alone tracks. Therefore the low- p_T cut on the prefilter sample adds no actual restriction, while ensuring the removal of potential artifacts. The high- p_T cut is chosen to restrict the analysis to a well-defined kinematic range with a relatively constant dielectron efficiency. At $p_T \approx 3.5$ GeV/ c , charged pions start to become indistinguishable from electrons with the utilized detectors, forcing either electron efficiency or purity to diminish.

Figure 4.3 shows the distributions of transverse momentum, pseudorapidity and azimuthal angle for electron candidates in data and Monte-Carlo for central collisions. The deviations in p_T and η have a common reason, namely that the particle composition and their momentum distributions in data are not well-reproduced by Monte-Carlo [Huh17]. This leads to different levels of hadron contamination in the electron candidate samples and causes the more significant kink in data around $p_T \approx 1.3$ GeV/ c . The additional entries near midrapidity in data are also caused by hadron contamination, but can be removed by a stricter PID selection. Details are presented in Section 4.2.2.

On the other hand, the φ distribution matches precisely between data and Monte-Carlo, showing that the detector performance itself is accurately modelled. Its distinct features, far from the expectation of being uniform, are caused by inoperative parts of the inner ITS layers (see Section 3.2.1). A recovery of tracks in the affected regions is discussed within this section.

In many analyses, a reweighting of the p_T -spectra of particles in Monte-Carlo is performed to reduce the occurring deviations. The present analysis removes this necessity by utilizing these inputs differentially, and in particular by decoupling the pair efficiency calculation from the actual input kinematics by factorization (see Section 4.6.2).

Track Quality

Track quality cuts mainly enhance the reliability of a track's reconstructed parameters, but some may also be utilized to select or reject particular physics processes.

In an electron analysis, a typical example for the latter is the requirement of tracks




Figure 4.3.: Comparison of the kinematic distributions of electron candidates between the actual data and a Monte-Carlo simulation.

having a hit in the innermost detector layer. This allows for the exclusion of electrons from photon conversions which occur in this layer or beyond, leaving only those in the track sample which originate from conversions in the beampipe. These remaining conversion electrons have to be dealt with by using other cuts and techniques, which will be discussed in Section 4.4 and Section 4.5. Consequently, a hit in the first ITS layer is required for the track samples presented in the following and a modification of this cut to recover some acceptance gaps is described afterwards.

Basic quantities to ensure proper track reconstruction are the number of measured trajectory points in each of the tracking detectors, referred to as *Number of clusters* ($N_{\text{cls}}^{\text{det}}$) and the spread of these points around the reconstructed path, χ^2 per cluster ($\chi^2/N_{\text{cls}}^{\text{det}}$).

In Figure 4.4, the distributions of these quantities for the ITS³ and the TPC are compared between data and Monte-Carlo. The number of clusters is well-reproduced in case of the ITS, while for the TPC on average 10% more clusters are assigned to each track. The minimum requirement of four ITS clusters ensures at least two of them with particle identification capability, and a cut at five is used for systematic checks. No cuts are applied to $N_{\text{cls}}^{\text{TPC}}$ to avoid an inconsistent efficiency correction. The $\chi^2/N_{\text{cls}}^{\text{TPC}}$ has a slightly higher mean but smaller tail in case of Monte-Carlo. By default, tracks with $\chi^2/N_{\text{cls}}^{\text{TPC}} > 4$ are rejected⁴, and an optionally stricter cut value is chosen at $\chi^2/N_{\text{cls}}^{\text{TPC}} = 3$.

³In the analyzed data set, the χ^2 per cluster is not stored for the ITS due to data compression. The value of this quantity has since been appreciated and it is kept in newer data sets.

⁴This is the threshold at which tracks are excluded from the standard track sample by the filter used during preparation of this data set.




Figure 4.4.: Comparison of the number of ITS and TPC clusters per track and the χ^2 per TPC cluster between data and Monte-Carlo.

Several geometrics-based quantities, for which a better modelling is achieved in Monte-Carlo productions, are provided as alternatives to $N_{\text{cls}}^{\text{TPC}}$ [Kal12]. Two of them are used in this analysis: one is the *Number of crossed rows* ($N_{\text{cross.rows}}^{\text{TPC}}$), which stores how many of the 159 TPC readout pad-rows in radial direction are crossed by the reconstructed trajectory. For tracks which exit the TPC through its end caps, this number is reduced accordingly. The second one is the ratio between $N_{\text{cross.rows}}^{\text{TPC}}$ and the number of readout pads on this trajectory which can potentially deliver a signal, and consequently called *Number of crossed rows over findable clusters* ($N_{\text{cross.rows}}^{\text{TPC}}/N_{\text{find.cls}}^{\text{TPC}}$). Tracks crossing TPC readout-chamber boundaries have a reduced number of findable clusters, while inactive readout pads are still considered to produce findable clusters.

Their distributions are presented in the two left panels of Figure 4.5 for electron candidates in data and Monte-Carlo. Since the acceptance cuts in η and p_{T} require all tracks to traverse the TPC entirely, $N_{\text{cross.rows}}^{\text{TPC}}$ is strongly peaked at its maximum value. This property is well-described by Monte-Carlo and applied cuts range from 80 to 130 minimum crossed rows. The narrow distribution of $N_{\text{cross.rows}}^{\text{TPC}}/N_{\text{find.cls}}^{\text{TPC}}$ around unity showcases the high reconstruction efficiency of charge depositions and therefore trajectories in the TPC in general. The lower cut is varied from 0.8 to 0.9 with the upper cut fixed at 1.1. The right panel of Figure 4.5 shows the *distance of closest approach* (DCA) between the reconstructed trajectory and the primary-vertex position in the transverse plane. It is well-described by Monte-Carlo. Only a very loose cut of $|\text{DCA}_{xy}| < 1.0$ cm is applied to guarantee an unbiased selection of electrons from open-heavy-flavour mesons, which have a wider DCA distribution than electrons from prompt sources.




Figure 4.5.: The number of crossed rows in the TPC, its ratio to the number of findable clusters, and the distance of closest approach between track and vertex in the transverse plane, compared between data and Monte–Carlo.

Track Recovery

Some distinct distributions of electron candidates with the typical track cuts explained previously, including a hit in the first ITS layer, are presented with the blue histograms in Figure 4.6. The large inhomogeneities in the azimuthal distribution are due to inoperative sectors of the Silicon Pixel Detector (SPD), which forms the two inner layers of the ITS. This causes down to zero acceptance in some regions, if their signal is required for a track.

As a means to recover tracks in these regions, an additional track sample is defined with the requirement of having a hit in the first layer of the Silicon Drift Detector (SDD), i. e. the third ITS layer, and no hit in the first two layers to avoid double-counting. In addition, the cut on $N_{\text{cls}}^{\text{ITS}}$ is relaxed to a minimum of three (four for systematic checks). This track sample is shown with the green histograms in Figure 4.6. It contributes an additional 46% to the total statistics of electron candidates in central (58% in semi-central) collisions and successfully fills the gaps in the φ distribution. Since information from the innermost detector is missing, the origin of these tracks cannot be reconstructed quite as precisely, resulting in a wider DCA distribution. Overall, the inclusion of this additional track sample has proven to be beneficial for the present analysis. It is therefore used for all final results, while the original track selection is included in the systematic studies.

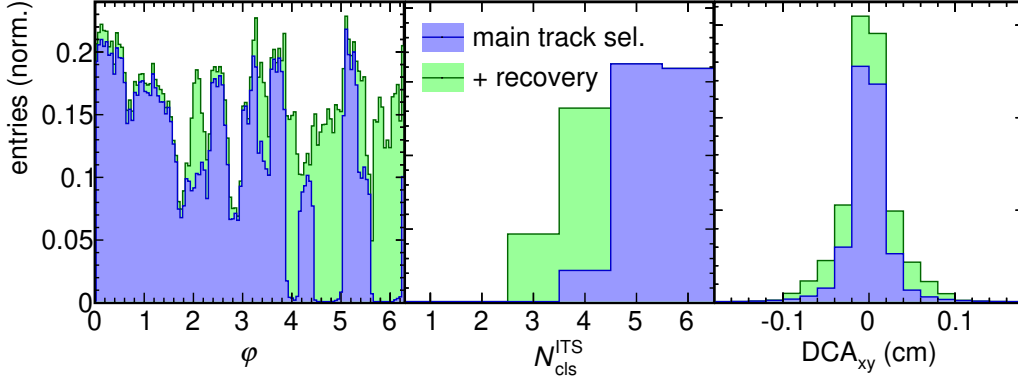


Figure 4.6.: Combination of the typical track selection and an additional track sample with hit in the first layer of the ITS-SDD to recover tracks in regions of inoperational ITS-SPD. The additional track sample has a relaxed requirement on the number of ITS clusters and a wider DCA distribution.

4.2.2. Particle Identification

The reference property known for any reasonably measured track in the central barrel of ALICE is its momentum, which is inferred from the trajectory's curvature in the magnetic field. Identification of electrons amongst other charged particles relies on the correlation of further properties to this reference. The two most relevant measurements for this purpose are those of the *specific energy loss* dE/dx , which is a function of $\beta\gamma$, and the particle velocity v . According to $p = \gamma mv = \beta\gamma mc$, both dE/dx and v are momentum-dependent. At momenta above 1–2 GeV/c, an energy loss mechanism directly depending on γ , namely *transition radiation*, becomes powerful to identify electrons [Rei11]. Since the present analysis is focussed on low momenta, it is not used here.

In the momentum range relevant to this analysis, the specific energy loss of electrons has already saturated in the Fermi-plateau of the Bethe-Bloch parameterization, while the one of light hadrons steeply falls towards minimum ionization, before entering the region of relativistic rise. This behaviour is best seen in detectors with gaseous active volume, because the extent of relativistic rise is limited by the (charge-)density of the medium. It amounts to 50–70 % for noble gases and only about 10 % for solids [GS08].

The left panel of Figure 4.7 shows the specific energy loss measured by the TPC for a subset of tracks from the analyzed data set, which have passed the track quality criteria described previously. dE/dx is given in units of the standard deviation σ around the




Figure 4.7.: Specific energy loss of particles in the TPC (left panel) and in the ITS (right panel). Electrons are located around zero, with significant separation from the minimum-ionizing region of hadrons in case of the TPC.

expected electron signal, which is thus situated at zero⁵. It is presented as a function of the track momentum measured by the TPC itself, which gives the most precise correlation, as opposed to using the momentum from globally re-evaluated trajectories. At lowest momenta, limited by a kinematic cut of $p_T > 150$ MeV/c, the rapidly dropping dE/dx of the dominant pion signal is visible, followed by the distinct bands of kaons and protons, crossing the electron signal around $p_{\text{TPC}} \approx 0.5$ GeV/c and $p_{\text{TPC}} \approx 1$ GeV/c, respectively. The analogous dE/dx -measurement from the ITS is presented in the right panel of Figure 4.7. Two major differences appear in comparison to the TPC measurement: since the relativistic rise in dE/dx is negligible in the silicon chips of the ITS, the electron signal does not significantly differ from that of minimum-ionizing particles (MIP), such that electrons already become indistinguishable from pions, kaons and protons as each of the latter approach their point of minimum ionization. Also, the hadron bands are much broader than in case of the TPC. This is mainly due to the fact that the ITS measures at most four PID clusters per track, while the TPC measures up to 159, therefore achieving a better dE/dx -resolution. Besides that, the momentum estimate from the TPC is used for this figure as well, which may differ from the particle momentum during ITS transition, thereby softening the correlation⁶.

⁵A slight downward shift is visible, since the data used for Figure 4.7 and Figure 4.8 does not contain the PID post-calibration described in Section 3.2.2.

⁶For the actual PID selection in the ITS, the global track momentum is used. It provides a good compromise by incorporating the ITS trajectory points and having a high momentum resolution.




Figure 4.8.: Left panel: time-of-flight measurement in the TOF. Right panel: specific energy loss in the TPC after electron selection by the ITS and by TOF (if available). See text for details.

The difference in dE/dx -behaviour between TPC and ITS is exploited in this analysis to exclude hadrons from the electron sample. In the momentum regions where any of the hadron signals intersects with the electron signal in the TPC, the ITS provides discrimination power by rejecting tracks with an energy loss above a certain threshold. Its value is to be tuned between high electron purity versus efficiency. The nominal electron selection is set to $-4 < \sigma_e^{\text{ITS}} < 1$, alternatively using 0.5 or 1.5 for the upper and -3 for the lower bound to check for systematic variations.

The typical PID strategy for electrons in the central barrel of ALICE foresees the Time-Of-Flight detector (TOF) for their preselection. It provides a much better separation to hadrons up to higher momenta than the ITS, theoretically reaching 100% electron purity in the TPC crossing regions when using a $|\sigma_e^{\text{TOF}}| < 3$ electron selection. Downsides of this strategy are the reduced track efficiency when a TOF hit is required and a threshold of $p_T \approx 0.4$ GeV/c, below which particles do not reach the TOF due to their deflection in the magnetic field.

These features can be seen in the left panel of Figure 4.8, where the time-of-flight measurement is presented in units of σ around the expected electron flight time. Since the velocity of electrons has already reached $v \approx c$ at the lowest accessible momenta, values below zero (and beyond resolution) are unphysical, being caused by event mismatches. The color scale is matched to the ITS and TPC measurements for an intuitive efficiency

comparison. To avoid this reduction in track statistics, the present analysis uses TOF PID — in the form of a $|\sigma_e^{\text{TOF}}| < 3$ electron selection — only if a TOF hit is associated to the track. This PID mode is referred to as *TOF if available* and applied here in combination with the approach of using ITS PID as discussed previously.

The result of this composite PID strategy is presented in the right panel of Figure 4.8, which shows the TPC signal after electron selection in the ITS and TOF. With the nominal track quality and PID cuts, $\approx 70\%$ of the original TPC electron statistics (Figure 4.7, left) are conserved. At the same time, the hadron bands have nearly vanished from the region of the electron signal. Only if tracks with $p_T < 0.4$ GeV/ c are included in the analysis, the residual pion contamination is considerable, depending on the final electron selection. The nominal range is $-1 < \sigma_e^{\text{TPC}} < 3$. As the lower limit is most sensitive to hadron contamination, it is varied between -1.5 and -0.5 . The upper limit is optionally reduced to 2.

It should be noted that even with a strict requirement of a TOF hit for every track, some hadron contamination would remain. This is caused by an effect called *TOF mismatching*, where the TOF signal of one particle is associated with the track of another. The relevant case here is that the TOF hit of a pion is accidentally assigned to the track of a kaon or proton, giving them time-of-flights similar to those of electrons. This is also the most frequent case, due to the high abundance of pions.

The consequences of a residual kaon and proton contamination in the electron sample have been studied in an extension [Kle14] of the physics performance study for the ALICE upgrade (Chapter 6): for an analysis using $p_T > 0.2$ GeV/ c , a contamination of 30% in the crossing regions contributes 20% of the combinatorial-background pairs. It thereby influences the dielectron signal quality, but does not contribute to the subtracted signal itself. In the present analysis, the overall background fraction from contamination is larger due to the smaller considered p_T range, yet potential pairs from correlated hadrons are not expected to have a significant influence on the signal. Nonetheless, by the variation of PID cuts as part of the systematic studies (Section 4.7), the effect of hadron contamination is incorporated into the resulting signal uncertainty.

4.2.3. Resulting Electron Samples

With all track selection and particle identification criteria applied, the resulting number of electron candidates per event is presented in the left panel of Figure 4.9. Their distribution has a mean of 4.5 electron plus positron candidates. The right panel shows




Figure 4.9.: Number of selected electron candidates per event in central Pb–Pb collisions (left panel) and one choice of a corresponding prefilter electron sample (right panel).

the distribution of a prefilter electron sample, for which tracks are accepted in an extended kinematic range and with relaxed quality and PID criteria, resulting in a mean of 125 particles per event⁷. The application of the prefilter electron sample will be described in Section 4.4. Event fractions presented here are relative to the number of events in the given centrality class that have passed the event selection of this analysis. 92% of the 0–10% most central collisions contribute at least one pair (unlike-sign or like-sign) to the analysis. In comparison, the centrality class of 10–50% has a mean of 1.6 electron candidates per event, with only 44% of all selected events containing at least one pair.

4.3. Pairs

In this section the dielectron signal extraction and quantities to monitor the signal quality are described.

4.3.1. Signal Extraction

Since the size of the region in which prompt particles are created is orders of magnitude smaller than the pointing resolution of any detector, there is no means to check which two

⁷The samples for the final reference result have been modified concerning their track quality parameters. While not being differentially available, they show the same magnitude and relative sizes as the presented ones.

tracks belong to a true pair. Therefore the dielectron signal is extracted on a statistical basis: in each event, all n_- electron candidates are paired with all n_+ positron candidates, thus creating $N_{\text{ULS}} = N_{-+} = n_- \cdot n_+$ unlike-sign pairs. These contain the real signal (S) as well as two types of background, namely combinatorial and correlated non-signal pairs, which have to be subtracted:

$$N_{\text{S}} = N_{\text{ULS}} - N_{\text{B}}. \quad (4.1)$$

An elegant way of constructing the background (B) is to create like-sign pairs from the electron and positron candidates, since they contain on average the same combinatorics and correlations as the unlike-sign background pairs. This leads to $N_{--} = n_- \cdot (n_- - 1)/2$ and $N_{++} = n_+ \cdot (n_+ - 1)/2$ pairs, while the total number of like-sign pairs is computed as

$$N_{\text{LS}} = 2\sqrt{N_{--} \cdot N_{++}}. \quad (4.2)$$

Using twice the geometric mean rather than the sum (twice the arithmetic mean) ensures the correct behaviour in the presence of an average charge asymmetry, i. e. when $\langle n_- \rangle \neq \langle n_+ \rangle$.

However, the acceptance may be different for unlike- and like-sign pairs due to detector effects. This becomes evident especially at small opening angles, where two particles with same charge traverse adjacent detector regions which tend to be similarly functional, while oppositely charged particles will be measured in more unrelated regions. Therefore an acceptance correction factor R is introduced, such that the expected unlike-sign background can be described by

$$N_{\text{B}} = N_{\text{LS}} \cdot R. \quad (4.3)$$

This so-called *R-factor* is built via a mixed-event technique (see Section 4.1.2) as

$$R = \frac{N_{\text{ULS}}^{\text{mix}}}{N_{\text{LS}}^{\text{mix}}} = \frac{N_{-+}^{\text{mix}} + N_{+-}^{\text{mix}}}{2\sqrt{N_{--}^{\text{mix}} \cdot N_{++}^{\text{mix}}}}. \quad (4.4)$$

The superscript “mix” denotes that the paired particles are taken from two different events, which belong to the same event class. Further description and a result for the *R-factor* is given in Section 4.1.2.

For each pair, the invariant mass m_{ee} and transverse momentum $p_{\text{T},ee}$ are computed,

based on the reconstructed momenta \vec{p}_1 and \vec{p}_2 of the tracks assigned to this pair, called its *legs*, and the hypothesis that both legs have the electron mass m_e :

$$m_{ee} = \sqrt{2m_e^2 + 2(E_1E_2 - p_1p_2 \cos \theta_{ee})}, \quad (4.5)$$

$$p_{T,ee} = \sqrt{(p_{T,1}^2 + p_{T,2}^2 + 2p_{T,1}p_{T,2} \cos \theta_{ee})}, \quad (4.6)$$

with θ_{ee} being the angle between the momentum vectors, called *opening angle* of the pair. The signal extraction described above is then performed differentially⁸ in m_{ee} and $p_{T,ee}$:

$$N \rightarrow dN/dm_{ee}dp_{T,ee}. \quad (4.7)$$

For the 0–10% most central Pb–Pb collisions and the default electron selection using $0.4 < p_T^e < 3.5$ GeV/ c , Figure 4.10 shows the invariant-mass distribution of unlike-sign pairs and acceptance-corrected background, as well as their difference, the subtracted signal. Horizontal bars denote the bin width and vertical bars show the statistical uncertainty, given by one standard deviation σ :

$$\sigma_m = \sqrt{N_m} \quad \text{and} \quad (4.8)$$

$$\sigma_d = \sqrt{\sum_m \left(\frac{\partial N_d}{\partial N_m} \cdot \sigma_{N_m} \right)^2} \quad (4.9)$$

for measured ($m \in -+, --, ++$) and derived ($d \in S, LS, B$) quantities, respectively. The spectra are normalized per 1 GeV/ c^2 . They are not corrected for the electron selection efficiency and therefore denoted as *raw* spectra.

Except for very low masses, the unsubtracted spectra almost perfectly overlap, which apparently is not quite the case since the dielectron signal is a continuous function of mass. Therefore the background shape cannot be determined from the unlike-sign spectrum alone — in contrast to analyses of a sharp resonance flanked by pure background. It has to be actually measured, using a method as described above. There is no immediate verification for it and its robustness will be evaluated in Section 4.7.

⁸Due to the usage of the geometric mean in equations (4.2) and (4.4), the signal extraction is sensitive to the utilized granularity. To minimize effects due to statistical fluctuations, this granularity is identical to the $(m_{ee}, p_{T,ee})$ -bin sizes used for the figures. If in one bin $N_{--}^{(\text{mix})} = 0$ or $N_{++}^{(\text{mix})} = 0$, then $N_{LS}^{(\text{mix})} = N_{++}^{(\text{mix})}$ or $N_{LS}^{(\text{mix})} = N_{--}^{(\text{mix})}$, respectively (instead of zero), resulting in additional robustness.




Figure 4.10.: Raw invariant-mass distributions of unlike-sign pairs (open squares), background (crosses) and subtracted signal (full circles). A prefilter is used (Section 4.4) and final pair cuts are applied (Section 4.5).

4.3.2. Signal Quality

Given that the background is constructed correctly, there are two main quantities to judge the reliability of the signal extraction: the *signal-to-background ratio* and the *significance*. Both can be immediately derived from raw spectra and are therefore useful to optimize the analysis. High values are preferred for both quantities. Since they do not rely on an efficiency correction or even a comparison to a signal expectation, they are independent of simulated Monte-Carlo data and do not pose a risk of tuning the analysis towards a desired result.

Signal-to-Background Ratio

A straightforward way to judge the quality of a signal extracted via equation (4.1) is to examine the signal-to-background ratio N_S/N_B . It represents the fraction of true over fake pairs. For given electron selection criteria, its overall magnitude is antiproportional to the number of dielectron signal sources per event, because the background from fake




Figure 4.11.: Signal-to-background ratio constructed from the spectra of Figure 4.10.

pairs rises quadratically. It is therefore also approximately antiproportional to the event multiplicity. These relations require that only a small fraction of the electron candidate sample consists of tracks from non-dielectron sources or from pairs with one leg outside of the acceptance, since either of them only produce background.

To improve the signal-to-background ratio, one can in general reduce the phase space for accepting signals by restricting leg momenta or the opening angle. This is useful for measuring one particular source of dielectrons (e. g. the J/ψ), but when aiming to measure the whole low-mass continuum there is limited scope for this. One obvious source of pure background are pairs containing one hadron which passes the electron selection, so improving the electron purity increases N_S/N_B . This typically compromises the electron efficiency and thereby causes an anticorrelation between N_S/N_B and efficiency (see Section 4.7, Figure 4.27). While the signal-to-background ratio is in principle independent of the electron efficiency, this is violated if the change in efficiency effects the fraction of tracks from purely background-producing sources, as mentioned above.

In Figure 4.11 the signal-to-background ratio derived from the spectra of Figure 4.10 is presented. In the vicinity of $m_{ee} \approx 1 \text{ GeV}/c^2$, where the background peaks, it drops below one permille. Semi-central collisions feature a 2.5 times higher signal-to-background

ratio with very similar mass dependence. This difference precisely matches the change in charged-particle multiplicity density, increasing by $dN_{\text{ch}}/d\eta|_{0-10\%}/dN_{\text{ch}}/d\eta|_{10-50\%} = 2.5 \pm 0.2$ [ALICE11a] from semi-central to central collisions.

Significance

Significance quantifies the statistical precision of a signal. In this analysis it is defined as

$$\text{significance} := \frac{N_{\text{S}}}{\sqrt{N_{\text{S}} + 2N_{\text{B}}}} \quad (4.10)$$

and thereby essentially represents the inverse of the relative statistical uncertainty $\sigma_{\text{S}}/N_{\text{S}}$ of the signal, as we will see below.

At a constant signal-to-background ratio, $\frac{N_{\text{S}}}{N_{\text{B}}} \rightarrow \frac{c \cdot N_{\text{S}}}{c \cdot N_{\text{B}}}$, the significance is proportional to \sqrt{c} . Therefore a high significance is obtained if a large number of electron candidates enter the analysis: ideally by increasing the total amount of data or, for a given number of recorded events, by achieving a high electron selection efficiency, provided that any accompanying hadron contamination is under control (see Section 4.7, Figure 4.28).

To derive the significance we evaluate σ_{S} for $N_{\text{S}} = N_{\text{ULS}} - N_{\text{LS}} \cdot R$ based on equation (4.9). The uncertainty on R is neglected here, because pair statistics in event mixing is much larger than in single events. Also note that $\sigma_{\text{LS}} = \sqrt{N_{--} + N_{++}}$ is valid for both the geometric and arithmetic like-sign mean:

$$\begin{aligned} \sigma_{\text{S}} &= \sqrt{\left(\frac{\partial N_{\text{S}}}{\partial N_{\text{ULS}}} \cdot \sigma_{\text{ULS}}\right)^2 + \left(\frac{\partial N_{\text{S}}}{\partial N_{\text{LS}}} \cdot \sigma_{\text{LS}}\right)^2} \\ &= \sqrt{(1 \cdot \sigma_{\text{ULS}})^2 + (R \cdot \sigma_{\text{LS}})^2} \\ &= \sqrt{\sqrt{N_{-+}}^2 + R^2 \cdot \sqrt{N_{--} + N_{++}}^2}. \end{aligned}$$

Measured quantities are now substituted with their composition in terms of S and B, while distinguishing between the backgrounds in unlike-sign and like-sign, because they contribute statistics-independently: $N_{-+} = N_{\text{S}} + N_{\text{B}_{-+}}$ and $N_{--} + N_{++} = N_{\text{B}_{\text{LS}}}$. Additionally, in this analysis, R is near unity ($R^2 \approx 1$) and therefore also $N_{\text{B}_{\text{LS}}} \approx N_{\text{B}_{-+}}$. So we obtain

$$\begin{aligned} \sigma_{\text{S}} &= \sqrt{N_{\text{S}} + N_{\text{B}_{-+}} + R^2 \cdot N_{\text{B}_{\text{LS}}}} \\ &\approx \sqrt{N_{\text{S}} + 2 \cdot N_{\text{B}}}, \end{aligned}$$

which is the denominator of the significance as given above. Evidently, this definition depends on the signal extraction procedure and actual properties (e. g. size of R -factor) of the analysis.

Unnormalized and Normalized Significance

It is essential to realize that the significance defined by equation (4.10) is not a dimensionless quantity, but rather

$$[\text{significance}] = \sqrt{[N_s]}. \quad (4.11)$$

Therefore it depends on the normalizations performed on the spectra beforehand. Two versions are used within this work:

1. The significance is computed from signal and background spectra without any normalizations, so that they contain the actual number of pairs in each bin⁹. With this method each significance value gives the inverse of the signal's relative statistical uncertainty in that bin. This should be the default when referring to "significance", and is used throughout the present chapter. The unit in this case is 1, so one could also call this *unnormalized* significance.
2. Signal and background spectra are normalized per 1 GeV/ c^2 and per one event before computing the significance. The unit of this *normalized* significance consequently is (GeV/ c^2)^{-1/2} and the normalization factor is $\sqrt{\text{bin size} \cdot N_{\text{events}}^{-1}}$. This generic representation is useful to compare results across analyses, collision systems or experiments, also if different bin sizes are used. It can also serve as input for random sampling of hypothetical data points, when combined with an assumed number of recorded events and bin sizes. This method is utilized in Chapter 6, Section 6.3.3.

The results for both significance versions, from the data presented above, are shown in Figure 4.12. Besides the major difference in absolute numbers, also their shapes differ, due to the non-uniform bin sizes which enter in the normalization process.

4.4. Pair Prefiltering

Pair prefilters are intended to identify electron candidates from particular undesired sources and reject them before they enter the signal extraction procedure. The goal is to

⁹In case of non-uniform bin sizes those spectra do look distorted, but they are the correct input.




Figure 4.12.: Significance: based on signal and background spectra which are unnormalized (left panel) or normalized per event and $1 \text{ GeV}/c^2$ (right panel).

reduce the combinatorial background and consequently achieve a better signal quality with minimum impact on the signal contributions from other sources.

4.4.1. Prefilter Technique

The prefilter is essentially a copy of the pairing procedure which is used for the signal extraction and is executed beforehand. A prefilter cut defines the criteria which a pair must fulfill to be rejected by this prefilter, which means that both of its legs will be removed from the final analysis. Legs of rejected pairs are not immediately excluded from the procedure, such that one track may cause the removal of multiple others¹⁰.

The prefilter cuts may be applied to unlike-sign as well as like-sign pairs. For a typical prefilter cut, which aims at rejecting a specific (unlike-sign) signal contribution, the filtering on like-sign clearly does not contribute to achieving this goal. However, if this same prefilter cut also causes some rejection of combinatorial pairs, referred to as *random rejection*, then the corresponding filtering on like-sign ensures a consistent treatment of combinatorial like-sign pairs. This is vital for the signal extraction procedure, which uses like-sign pairs as approximation for the unlike-sign background, as described in

¹⁰Alternatively, considering that each track should belong to just one pair, the rejection could be limited to only the pair which has the highest probability of being from an undesired dielectron source. The relevant probability distributions as function of the cut variables, as well as the benefit of losing less signal at the cost of rejecting undesired sources less effectively, would have to be studied.

Section 4.3.1. Therefore the prefilter is applied to both unlike-sign and like-sign pairs throughout this analysis.

4.4.2. Prefilter Tuning

In a low-mass dielectron analysis the majority of electrons intrinsically comes from π^0 Dalitz decays — and from photon conversions, depending on detector and analysis properties, which will be discussed in Section 4.5. So besides dominating the signal below the pion mass, these electrons also produce the majority of the background, and their contribution spans the whole kinematic region of interest. The prefilter is therefore configured to reject a large portion of the π^0 signal via a mass cut with $m_{\text{cut}}^{\text{min}} = 0$ and $m_{\text{cut}}^{\text{max}} < m(\pi^0)$. The upper cut value needs to be tuned for a good compromise between rejection of signal pairs and combinatorial pairs, because each rejected combinatorial pair causes the loss of up to two dielectron signals at higher mass due to losing one of their leg electrons. A natural value for setting this cut is therefore the mass at which the signal-to-background ratio equals unity ($m_{ee} \approx 40 \text{ MeV}/c^2$, see Figure 4.11).

To further reduce random rejection, a cut on one of the following other pair variables is added to the prefilter, each having their individual benefits:

- The opening angle of the pair (θ_{ee}): π^0 signals prefer small opening angles while combinatorial pairs prefer large ones. Due to the correlation between θ_{ee} and m_{ee} , the opening-angle distribution also depends on the pair mass. The cut value is chosen near the most probable opening angle of π^0 Dalitz pairs whose mass is equal to the selected m_{ee} cut.
- The charge-ordered pair plane angle with respect to the magnetic field direction (ϕ_V)¹¹: prompt signals as well as combinatorial pairs are almost uniformly distributed in this variable at low invariant mass ($m_{ee} \lesssim 100 \text{ MeV}/c^2$)¹². Photon conversion pairs however populate at most half of its range and their ϕ_V distribution is strongly peaked towards the maximum value (π) as a result of the track reconstruction process. In this analysis a cut on $\phi_V > \frac{3}{4}\pi$ rejects more than 95 % of all reconstructed photon conversions, both in central and semi-central collisions, while approximately 75 % of the prompt signal survives in the mass region in which the cut is applied.

¹¹The exact definition of ϕ_V is given in Appendix B. Visualizations for prompt and conversion pairs are shown in Figure 4.16 below.

¹²At high invariant mass, phase-space effects due to limited detector acceptance lead to deviations.




Figure 4.13.: Comparison of the subtracted signal without prefilter and with two different prefilters, including efficiency corrections. No further pair cuts are applied. The signals are consistent within their statistical uncertainties, except for very low mass, where the prefilters cut into the signal.

A prefilter using θ_{ee} and m_{ee} cuts will cause vanishing yield for $m_{ee} \rightarrow 0$, while the one using ϕ_V and m_{ee} preserves a constant fraction of the yield, corresponding to the ϕ_V cut value. Figure 4.13 shows the subtracted signal for both prefilter configurations in comparison to the result without prefilter. The selected exclusion cuts are listed in Table 4.2. Corresponding efficiency corrections, to be discussed in Section 4.6, are applied to these spectra for a consistent comparison. Apart from the mentioned deviations at very low mass, the three signals are consistent within their statistical uncertainties. While this is a positive finding, it is also hard to judge which one has the best signal quality.

	mass exclusion	angular exclusion
prefilter #1	$m_{ee} < 40 \text{ MeV}/c^2$	$\theta_{ee} < 80 \text{ mrad}$
prefilter #2	$m_{ee} < 40 \text{ MeV}/c^2$	$\phi_V > \frac{3}{4}\pi$

Table 4.2.: Exclusion cuts of the two prefilter types. Pairs are rejected, if both requirements are fulfilled.

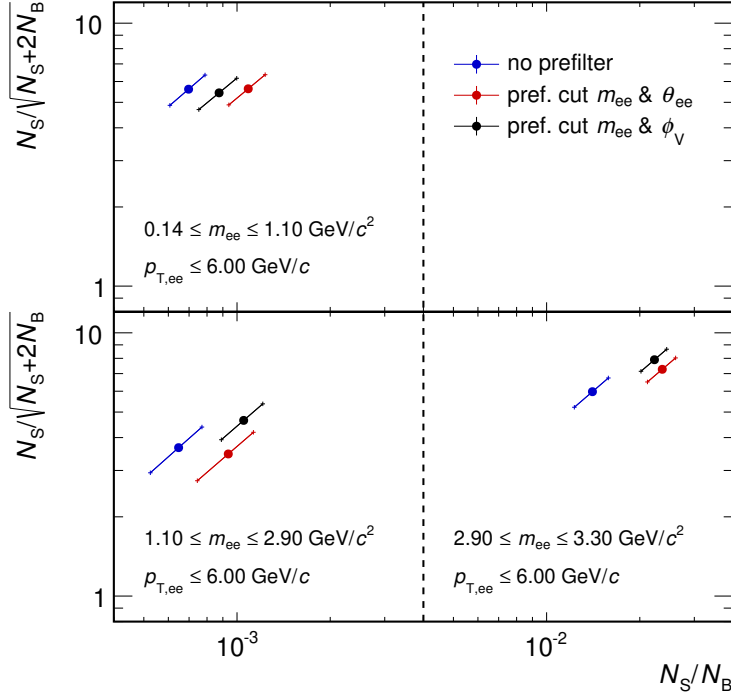


Figure 4.14.: Correlation between significance and signal-to-background ratio for the signals from Figure 4.13, integrated over the specified mass ranges.

To visualize the anticipated benefits of the prefilters, a different representation of the results is needed. Since both the signal-to-background ratio and the significance are needed to quantify the signal quality (Section 4.3.2), their correlation can provide further insight. It is created separately for the mass ranges dominated by the light-flavour resonances ($0.14 \leq m_{ee} \leq 1.1 \text{ GeV}/c^2$), the open-heavy-flavour mesons ($1.1 \leq m_{ee} \leq 2.9 \text{ GeV}/c^2$) and the J/ψ ($2.9 \leq m_{ee} \leq 3.3 \text{ GeV}/c^2$), to allow for a differential judgement.

The result is presented in Figure 4.14. The data points encode the significance along the ordinate and the signal-to-background ratio as abscissa for the three signals in the given mass ranges. The error bars are drawn to connect the points where both quantities commonly reach the upper or lower bound of their $\pm 1\sigma$ statistical uncertainties. These are both propagated from σ_S and σ_B and therefore correlated. Note that this correlation is not linear: the slope of the error bars would change, if they were computed for less or more than one standard deviation of N_S and N_B . A partial correlation of uncertainties across the three signals is also present due to their common data set. For these reasons, differences in the presented correlations are to be judged mainly from the pole values of

the data points.

When using one of the prefilters, the signal-to-background ratio rises by 20–60 % in the light-flavour mass range, 50–60 % in the intermediate-mass region, and 60–70 % for the J/ψ , compared to the unfiltered case. The only sizeable increase in significance occurs in the J/ψ region and amounts to 20–30 %.

4.4.3. Choice of Prefilter

Since both prefilter types perform similarly well in improving the signal quality, the final decision can be based on other concerns. An opening-angle cut causes a non-trivial change in the dielectron spectrum at low mass and high pair momentum. This has to be considered for comparisons to model calculations, into which the cut and corresponding experimental resolutions have to be implemented. In contrast, the dielectron spectrum can be corrected for a ϕ_V cut, thus preserving its natural shape down to zero mass. For these reasons prefilter #2 is chosen, such that pairs which fulfill both $m_{ee} < 40 \text{ MeV}/c^2$ and $\phi_V > \frac{3}{4}\pi$ are rejected.

The correction procedure, which includes evaluation and handling of any remaining yield from photon conversions, is described in the following section.

4.5. Final Pair Cuts

Pair cuts may in general be applied to define the phase-space of the results. In this analysis however, no restrictions are made beyond the acceptance cuts on single electrons. At very low mass, cuts to suppress conversion pairs are useful whether or not a prefilter is used. This is illustrated by Figure 4.15. Its left panel shows a zoomed-in mass distribution of the subtracted signal in data compared to its main contributions, the prompt signal¹³ and conversion pairs, obtained from Monte–Carlo. Open symbols show spectra without application of a prefilter. Pairs from photon conversions dominate the yield in the first two mass bins ($m_{ee} < 40 \text{ MeV}/c^2$). Beyond that, the prompt signal approximately describes the data up to the point, where pairs from heavy-flavour-meson decays enter the signal¹⁴. Closed symbols show these spectra after prefiltering with prefilter #2. Conversion pairs are reduced by 97 %, prompt pairs by ≈ 40 % in the first

¹³This term refers to all true electron–positron pairs from Dalitz and two-body decays of vector mesons.

¹⁴Some deviations occur, since kinematics and particle composition of the Monte–Carlo signal differs from data, due to the nature of its underlying generators. In this analysis no reweighting with proper inputs is performed, since none of the used correction procedures rely on them.

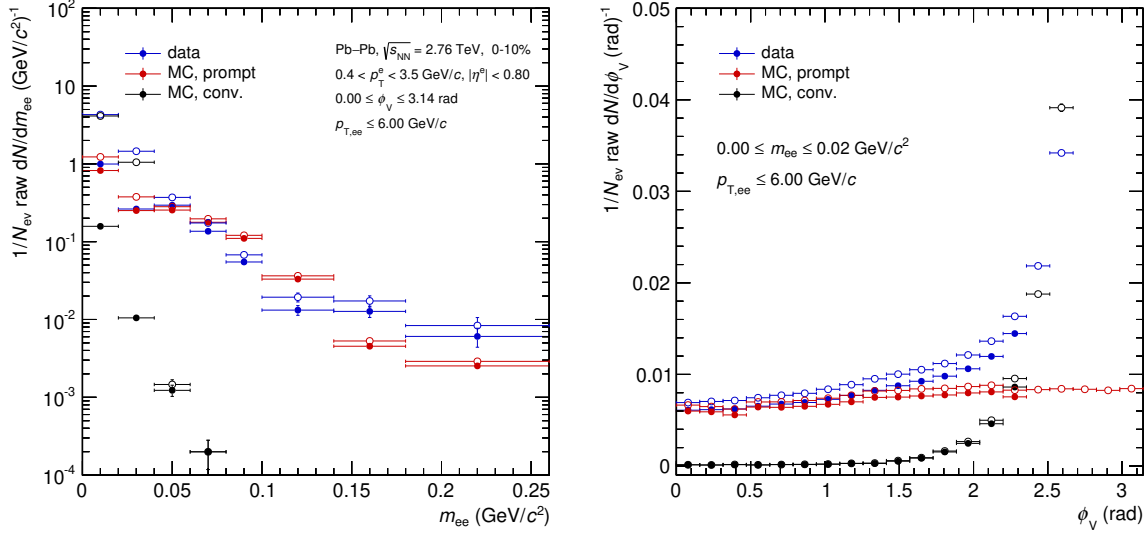


Figure 4.15.: Very-low-mass region of the subtracted signal compared between data and the prompt and conversion templates from Monte-Carlo (left panel) and the ϕ_V distribution of these signals in the first invariant-mass bin (right panel). Closed symbols show spectra with prefilter (m_{ee} & ϕ_V) applied, and open symbols represent unfiltered results.

two mass bins, where the cut is applied. The 10–15% reduction at higher mass is caused by random rejection of true signal, to be discussed in Section 4.6.4.

4.5.1. ϕ_V Distributions

A more differential view of the prefilter effects is given in the right panel of Figure 4.15, where the yields in the first mass bin are shown as a function of ϕ_V . Conversion pairs are strongly peaked towards $\phi_V = \pi$ as a consequence of the assumption during track reconstruction, that their electrons originate from the primary event vertex. This causes an artificial non-zero opening angle, with the pair plane always oriented perpendicular to the magnetic field direction and the trajectories of electron and positron intersecting each other (as opposed to diverging).

This effect is illustrated by the left graphic of Figure 4.16. Dashed lines indicate the true electron and positron paths with zero opening angle, originating from the conversion point in the beampipe. Solid lines show the reconstructed trajectories with finite opening angle between their initial directions. Their charge-ordered vector product \vec{v} is used in the computation of ϕ_V (see Appendix B). The value $\phi_V = \pi$ is reached when \vec{v} points in the negative z direction, i. e. opposite to the magnetic field vector \vec{B} .




Figure 4.16.: Left: photon conversion in the beampipe creating an electron–positron pair with zero opening angle (dashed lines). Track reconstruction constrains their trajectories to the event vertex, causing an artificial opening angle and overlap (solid lines). Right: prompt electron–positron pair with large opening angle. Arrows indicate the direction of reconstructed momentum vectors at the event vertex and span the shown pair planes. Their vector products \vec{v} relative to z determine the ϕ_V value.

An example for a prompt pair is given on the right, emphasizing that track curvature occurs in the x - y -plane rather than the pair plane. Prompt pairs should be uniformly distributed in ϕ_V , since their mother particles decay isotropically at the primary vertex. The deviation seen in Figure 4.15 is caused by different tracking resolutions in the transverse and longitudinal direction, such that for vanishing true opening angles at low mass, the reconstructed pair plane is more likely to be tilted into the conversion-like orientation.

In the first mass bin, a final pair cut at $\phi_V = \frac{1}{2}\pi$ removes the remaining conversion pairs. The data yield can then be corrected for the excluded phase-space either by an extrapolation using the shape of the prompt template, or — neglecting the resolution effect — by simply upscaling the yield by a factor of 2.

However, neither of these approaches works reliably for the subsequent mass bins, where the data still shows a peaked ϕ_V distribution even though the yield from conversion pairs becomes insignificant for $m_{ee} > 40 \text{ MeV}/c^2$. This is presented in Figure 4.17 for the mass bin of 40–60 MeV/c^2 . Closed symbols show results with prefilter and open ones




Figure 4.17.: ϕ_V distribution of data and individual signal components from Monte–Carlo in the third invariant-mass bin. Pairs from different photon conversions influence the signal shape and contribute with negative yield at low values of ϕ_V . Closed symbols show spectra with prefilter (m_{ee} & ϕ_V) applied, and open symbols represent unfiltered results.

without it, where the effects to be discussed are more pronounced than in prefiltered data. Electron–positron pairs belonging to the same conversion have zero yield, but combinatorial pairs from two different conversions have a significant impact on the signal shape and yield.

In particular, the yield from different conversions may become negative at low values of ϕ_V . As a consequence, the data yield in the region $\phi_V < \frac{1}{2}\pi$ is smaller than the prompt signal, such that both previously mentioned correction procedures will underestimate the expected yield in the affected mass bins.

In contrast to prompt pairs and true conversion pairs, the Monte–Carlo signal of different conversions has to be extracted with the approach $N_S = N_{ULS} - N_{LS}$ as in real data. The distortion of the leg trajectories due to the tracking constraint to the primary vertex does not significantly modify the opening angle of like-sign pairs from different conversions, since both of their momentum vectors are biased in the same direction. Their invariant




Figure 4.18.: Illustration of two adjacent photon conversions, showing only reconstructed trajectories (cp. Figure 4.16). The right one has a positive z-component, such that the connecting plane is rotated around y, and lower track momenta. The actual plane orientation varies for each pair of tracks. x-y-projections of trajectories from like-sign pairs may or may not intersect, allowing both orientations of their momentum-ordered vector product ($\vec{v}_{++}, \vec{v}_{--}$). All four unlike-sign pairs however show overlapping trajectories, causing the same ϕ_V for the true and the combinatorial pairs, which is not reproduced in the like-sign.

mass is thus approximately a function of the angle between the two conversion points. The combinatorial unlike-sign pairs however are shifted away from the like-sign mass, because their reconstructed opening angles deviate from the angle between the conversion points. A downward mass shift exceeding zero opening angle causes the vector product's direction to flip, such that: $\phi_V^{\text{ULS}} \rightarrow \phi_V^{\text{ULS}} + \frac{1}{2}\pi$.

Figure 4.18 shows a visualization of this effect based on two adjacent conversions with slightly different momenta. The like-sign pair with negative charges has overlapping x-y-projections (equal to viewing direction) of its trajectories, while the positive tracks do not intersect. Each case is true for half of the like-sign pairs, allowing the momentum-ordered vector product of their track momenta to point in either direction. All unlike-sign pairs however have overlapping track projections. In particular, the combinatorial pair

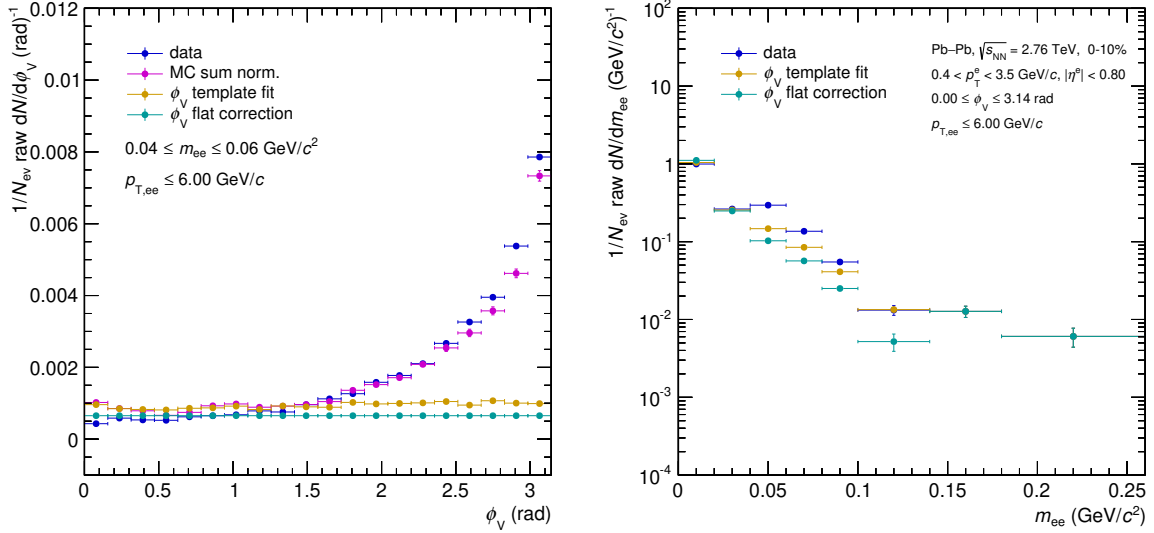


Figure 4.19.: Left panel: sum of the three Monte–Carlo components normalized to data in the third invariant-mass bin, with prefilter. The consistently scaled prompt template (“ ϕ_V template fit”) produces a higher yield than the flat correction method. Right panel: uncorrected prefiltered data in comparison to the results of the two correction methods.

whose true tracks (not drawn here) diverge also shows an overlap, which occurs close to the event vertex. This behaviour cannot be reproduced by the like-sign and causes the aforementioned oversubtraction at low ϕ_V values.

4.5.2. ϕ_V Template Fit

To extract the proper yield of the dielectron signal in the first mass bins, a Monte–Carlo template fit to data is performed. It is based on the ϕ_V distributions of true pairs from vector-meson decays, from photon conversions, and combinatorial pairs from different conversions. In each invariant-mass bin, the sum of these three Monte–Carlo contributions is normalized to the data over the full ϕ_V range, or up to $\phi_V = \frac{3}{4}\pi$ where the prefilter cut is applied.

The left panel of Figure 4.19 shows the normalization result for the mass bin of 40–60 MeV/ c^2 in prefiltered data. The Monte–Carlo sum matches the shape in data reasonably well, while some tension remains to be addressed. The normalization factor amounts to 0.58 in this bin (compared to the yields in Figure 4.17) and is applied to the prompt template. The result is shown here as “ ϕ_V template fit” and is compared to the simplistic approach of cutting the data at $\phi_V = \frac{1}{2}\pi$ and using the average yield per

remaining ϕ_V bin for a flat extrapolation.

In the right panel of Figure 4.19, the invariant-mass distribution of these two correction procedures is presented. Since the original data is reduced by the prefilter cut for $m_{ee} < 40 \text{ MeV}/c^2$, it coincidentally matches with the corrected results in these two bins. In the subsequent mass bins, the yield from the flat correction falls steeper than the one of the template fit. Both corrections are performed up to $m_{ee} = 140 \text{ MeV}/c^2$. The flat correction produces a yield in the last corrected mass bin, which is lower than the yields in the two shown untouched bins. This is also the case if instead the extrapolation according to the prompt template is used (not shown). This makes some kind of overcorrection apparent even without the ϕ_V -differential view. On the contrary, the result of the template fit approaches the original data as the signal contribution from different conversions falls off with increasing mass, and matches well in the last corrected bin.

Systematic Uncertainty

For the nominal results of the ϕ_V template fit it is assumed, that the relative sizes of the different signal components in Monte-Carlo correctly describe those fractions in data. This can be expected, since the low-mass prompt signal is dominated by π^0 decays, which also produce the majority of photons as source for conversion pairs. The material budget, being the second link in this correlation, is also precisely known and included in the Monte-Carlo production. However, the contribution of pairs from different conversions in the individual mass bins is expected to be particularly sensitive to resolution effects, due to their unusual geometry. Therefore a variation of their yield is performed to achieve an improved matching when normalizing the Monte-Carlo sum to data.

The result with a relative increase of the different-conversion yield by 50 % is shown in the left panel of Figure 4.20 for the mass bin of 40–60 MeV/c^2 . The shapes of Monte-Carlo sum and data are now almost identical and the resulting prompt component emerges closer to the flat correction than in the nominal case. An increased different-conversion yield is favoured by all corrected mass bins except the first. The accordingly reduced prompt yield is used as lower systematic uncertainty on the ϕ_V -corrected data, which is shown in the right panel of Figure 4.20. The upward uncertainty corresponds to a decreased different-conversion yield by 50 %. This is preferred by the ϕ_V distributions in the first mass bin, where the impact of this signal contribution is small though, but is less likely for the other bins. Based on the previous arguments, a redistribution of the different-conversion yield amongst the mass bins is more likely anyhow than an overall underestimate of its relative occurrence.

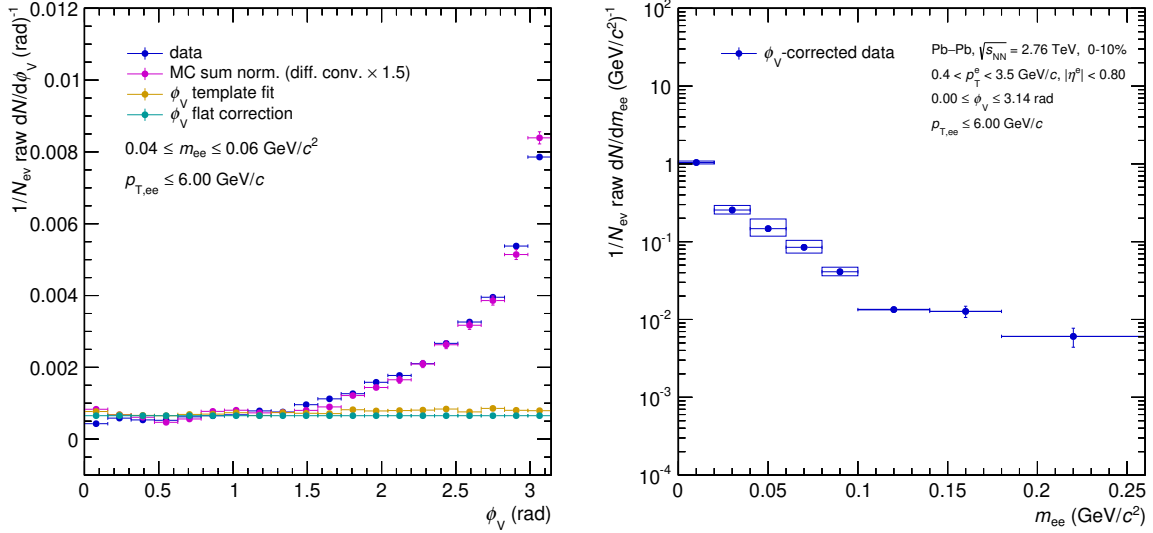


Figure 4.20.: Left panel: Monte–Carlo normalization result with 50 % increased yield from different conversions relative to prompt pairs. Right panel: ϕ_V -corrected data with corresponding systematic uncertainties, shown as boxes. Lower bounds are reached with 50 % increased and upper bounds with 50 % decreased different-conversion yield.

4.6. Efficiency Correction

Correcting the measured dielectron spectrum for the particle detection and identification efficiency of the experimental setup involves multiple steps. It is based on a production of Monte–Carlo data, which includes the propagation of simulated particles through an accurate model of the ALICE detector. This takes into account any interactions with the material as well as the responses of each subdetector.

In principle, the electron–positron pair efficiency can immediately be extracted from such a production by running the same analysis over it and dividing the number of successfully reconstructed pairs by all generated pairs in the acceptance:

$$\epsilon_{\text{pair}} = \frac{N_{\text{rec}}}{f_{\text{acc}}(N_{\text{gen}})}. \quad (4.12)$$

The function f_{acc} applies the acceptance cuts of the analysis on the generated pairs and their leg electrons, so that the remainder fulfills all kinematic criteria that the

reconstruction procedure imposes on measured pairs. This ensures that the efficiency correction takes into account only those regions of phase-space, where the measurement occurs¹⁵. The corrected yield is then obtained from the raw data via

$$N = \frac{1}{\epsilon_{\text{pair}}} \cdot N_{\text{data}}. \quad (4.13)$$

The implications of this general procedure and a modified approach, needed for the present analysis, are discussed in the following.

4.6.1. Generated vs. Measurable Quantities

In the process of obtaining N_{rec} for equation (4.12), the acceptance cuts are applied on kinematic variables, which were influenced by the simulated detector resolution and material interaction. These variables correspond to the ones available in the actual data analysis. We therefore refer to them as *measurable* quantities. In the denominator of equation (4.12), f_{acc} applies the same cuts on their generated counterparts. The pair efficiency is then constructed as a function of the generated pair variables: $\epsilon_{\text{pair}} = \epsilon_{\text{pair}}(m_{ee;\text{gen}}, p_{T,ee;\text{gen}})$. After all, no measured values exist for the portion of N_{gen} which was not reconstructed.

Consequently, to properly apply this pair efficiency to the data, all resolution effects would need to be unfolded from the raw spectra in advance to get $dN_{\text{data}}/dm_{ee;\text{gen}}dp_{T,ee;\text{gen}}$. This is not commonly attempted in low-mass dielectron measurements due to the continuous spectral shape. In return, ϵ_{pair} needs to be extracted as a function of the measurable quantities as well, in order to execute the efficiency correction according to

$$dN/dm_{ee;\text{meas}}dp_{T,ee;\text{meas}} = \frac{1}{\epsilon_{\text{pair}}} \cdot dN_{\text{data}}/dm_{ee;\text{meas}}dp_{T,ee;\text{meas}} \quad (4.14)$$

with

$$\epsilon_{\text{pair}} = \frac{N_{\text{rec}}(m_{ee;\text{meas}}, p_{T,ee;\text{meas}})}{f_{\text{acc}}(N_{\text{gen}}(m_{ee;\text{meas}}, p_{T,ee;\text{meas}}))}. \quad (4.15)$$

The differences between generated and measurable quantities are best seen in the variables of single tracks, as those are immediately influenced by the mentioned effects. They are

¹⁵An extrapolation into inaccessible phase-space is only allowed for dimensions in which the physics is invariant, such as the azimuthal angle, for collisions with unpolarized beams and observables which are averaged over events.




Figure 4.21.: Relations between measurable and generated values for the electron momentum, polar angle and azimuthal angle, each in correlation with the generated electron momentum.

extracted from Monte–Carlo data, using only true electrons and positrons, excluding secondaries, and presented in Figure 4.21. The left panel shows how the distribution of measurable over generated momentum (along the abscissa) correlates with the generated momentum. Since the main effect on this quantity is the energy loss of the particle, it is expected to see most of the smearing towards $p^{\text{meas}}/p^{\text{gen}} < 1$. The middle and right panels show the same momentum dependence for the polar and azimuthal angle respectively, except that the absolute differences are mapped.

Momentum and polar angle are shown for the sum of electrons and positrons, as these distributions are charge-symmetric. The azimuthal angle is presented for electrons only. Due to the oppositely curved trajectory of positrons, they experience detector inhomogeneities differently in the direction perpendicular to the magnetic field orientation. The resulting distribution is approximately mirrored along $\varphi^{\text{meas}} - \varphi^{\text{gen}} = 0$. In contrast to the momentum resolution, both angular resolutions continuously improve from low to high momenta.

With the matrices of Figure 4.21 it becomes possible to create the pair efficiency as a function of measurable pair variables, as needed for equation (4.14). For the non-reconstructed portion of N_{gen} , the expected leg properties are simulated by sampling measurable values according to the generated ones, thereby creating realistic pair kinematics. The other generated pairs inherit the kinematics of their associated reconstructed pair to avoid efficiency fluctuations. Before the pair efficiency is presented in Section 4.6.5, a second approach to its construction is discussed.

4.6.2. Factorization of Pair Efficiency

The number of simulated events in a Monte–Carlo production in ALICE usually amounts to one-tenth of the corresponding real data. For the analysis presented here, this provides insufficient statistics, especially in less populated regions of pair phase-space. Therefore a different procedure was developed, which is based on factorizing the pair efficiency into its single track components:

$$\epsilon_{\text{pair}} = \epsilon_{\text{track}}^{e^-} \cdot \epsilon_{\text{track}}^{e^+}. \quad (4.16)$$

This approach contains the assumption that, to first order, the detection efficiency of an electron with a particular trajectory through the detector does not depend on the pair it belongs to, i. e. is not directly correlated to the other leg’s efficiency¹⁶. The factorization enables an evaluation of the pair efficiency in any position of phase-space by folding the probabilities of all kinematically allowed two-body decays of a particle in that position with the products of their legs’ track efficiencies. It thereby completely overcomes statistical fluctuations in the pair efficiency, whereas sufficient statistics in the track efficiency remains mandatory to avoid the systematic propagation of fluctuations into the pair efficiency.

In each $(m_{ee}, p_{T,ee})$ -bin, particles with corresponding mass and transverse momentum are distributed uniformly in azimuth and pseudorapidity, then decayed into electron–positron pairs via a two-body decay without polarization. This method is also valid for dielectrons originating from Dalitz decays, since the relevant pair kinematics are determined by the virtual photon. The procedure can likewise represent pairs from decays of open-heavy-flavour mesons. This was verified by instead obtaining the pair and leg kinematics from semileptonic decays of correlated open-charm mesons using PYTHIA.

The resulting pair efficiency map does not carry any information on the actual phase-space distribution of dielectron sources, unless such distribution is explicitly used as weight for the number of particles per $(m_{ee}, p_{T,ee})$ -bin. However, as long as the map is created and applied to the data in fine granularity, no bias is introduced by an unphysical — e. g. uniform — distribution of particles per bin.

¹⁶The correlation due to both legs traversing similar detector regions is covered by using differential track efficiencies and realistic pair kinematics (see below). In contrast, the assumption would be violated by using a track cut on shared ITS clusters, because then most likely either both legs were measured or both lost, which is not reproduced by factorization.




Figure 4.22.: Illustrations of the three-dimensional electron efficiency map, as functions of the measurable p_T , η , and φ . The acceptance cuts in p_T and η are therefore undiffused. The non-flat φ distribution depicts the detector inhomogeneities.

4.6.3. Electron Efficiency

The electron reconstruction efficiency (also “track efficiency”) is obtained from a Monte-Carlo production using true electrons and positrons from the decays of light vector mesons, open-heavy-flavour mesons, and the J/ψ meson. The track quality, particle identification and acceptance cuts, described in Section 4.2, are applied to the reconstructed electrons. The kinematics of the generated ones are modified to achieve consistent, measurable quantities, in accordance to the procedure explained in Section 4.6.1: generated electrons with a reconstructed partner obtain its kinematics, while the properties of the others are modified by the resolution matrices. From this sample, a track efficiency map is created differentially in p_T , η , and φ :

$$\epsilon_{\text{track}} \rightarrow \epsilon_{\text{track}}(p_T, \eta, \varphi). \quad (4.17)$$

Figure 4.22 shows one-dimensional representations of this map, for which the numbers of reconstructed and generated electrons are integrated over the hidden dimensions before calculating the efficiency values. No distinct features are visible in the transverse-momentum and pseudorapidity distributions, except a slight p_T -dependent efficiency decline. The acceptance cuts of $0.4 < p_T^e < 3.5$ GeV/ c and $|\eta^e| < 0.8$ are visible as strict edges due to the usage of measurable quantities. The non-uniform azimuthal distribution is caused by detector performance issues, discussed in Section 4.2.1. The relatively coarse granularity presented here corresponds to the one of the three-dimensional efficiency

map, being a compromise between the realistic, differential modelling of the detector response and its statistical fluctuations.

4.6.4. Prefilter Efficiency

When using a pair prefilter to reduce the combinatorial background by rejecting electrons from certain sources, it is likely to also cause some *random rejection* of other electrons, namely those from desired signal contributions. This effect has to be quantified and included in the efficiency correction. It can be encoded in a track property that represents the fraction of accidentally rejected electrons over all measured ones, which we define as random rejection probability P_{rand} . More specifically, this property is based on the sample of electron candidates, which are meant to enter the signal extraction procedure, while the rejection of these candidates may be caused by any member of the potentially larger prefilter electron sample (see Section 4.2).

A data-driven procedure was developed to determine the rejection probability with high statistical precision:

In a given event, which contains a certain prefilter electron sample, the value of P_{rand} only depends on the kinematics of the electron candidate:

$$P_{\text{rand}} \rightarrow P_{\text{rand}}(p_{\text{T}}, \eta, \varphi). \quad (4.18)$$

It is determined by injecting a *testparticle* into the event and computing its rejection probability. The kinematic domain of testparticles is equal to the allowed acceptance for electron candidates. Since the original electron candidate is also part of the prefilter sample, the testparticle would experience one additional rejection partner. To avoid this slight multiplicity bias, one electron candidate is randomly removed from the prefilter sample¹⁷. This procedure is repeated n_{test} times per event with testparticles of varying kinematics to achieve large statistics for P_{rand} .

The resulting rejection probability naturally depends on the event multiplicity, to which the size of the prefilter electron sample is proportional. To achieve the correct weighting of events according to their multiplicity, the number of testparticles is set proportional to the number of electron candidates in each event: $n_{\text{test}} = i \cdot (n_{-} + n_{+})$. Then the map of mean random rejection probabilities in a given centrality class is obtained by the average values from all testparticles. Since this procedure is executed on the same data set as the analysis itself, there are no systematic uncertainties associated to it, while simultaneously

¹⁷Not any random electron is removed, but one which satisfies the final electron selection, because testparticles always match these stronger criteria.




Figure 4.23.: Illustrations of the three-dimensional random rejection map for electron candidates, as functions of the measurable p_T , η , and φ . It is extracted from testparticles which are injected into real data and cover the same acceptance as the electron candidates.

achieving high statistical precision.

The effect of a prefilter with cuts on $m_{ee} < 40 \text{ MeV}/c^2$ and $\phi_V > \frac{3}{4}\pi$ in central collisions is presented in Figure 4.23, where the random rejection probability is given as function of each considered track variable. While both angular distributions are relatively uniform, a strongly decreasing random rejection probability is observed as function of the transverse momentum, mainly because high- p_T tracks are less likely to form low-mass pairs.

4.6.5. Pair Efficiency

The pair reconstruction efficiency is determined by using the factorization approach discussed in Section 4.6.2. One necessary addition is to take the prefilter efficiency into account. It is incorporated into the factorization procedure by computing the *pair survival probability* of each simulated decay, based on the random rejection probabilities of its legs:

$$P_{\text{surv}}^{\text{pair}} = \left(1 - P_{\text{rand}}^{e^-}\right) \cdot \left(1 - P_{\text{rand}}^{e^+}\right). \quad (4.19)$$

The total pair efficiency is finally constructed via:

$$\epsilon_{\text{pair}} = \epsilon_{\text{track}}^{e^-} \cdot \epsilon_{\text{track}}^{e^+} \cdot P_{\text{surv}}^{\text{pair}}. \quad (4.20)$$




Figure 4.24.: Pair survival probability in the prefilter (left panel) and total pair efficiency (right panel) as function of invariant mass in narrow pair- p_T slices. Since electrons are selected only up to $p_T^e = 3.5$ GeV/c, there is a $p_{T,ee}$ -dependent mass cutoff, which is visible for the last momentum slice.

Both quantities are presented in Figure 4.24, where their mass dependence is given for four narrow slices in pair transverse momentum¹⁸. The survival probability has its minimum at low mass and momentum, which is expected from the strongly p_T -dependent random rejection of tracks (Figure 4.23). The opposite trend is observed for the total pair efficiency, showing that the random rejection of electrons is overcompensated by the higher track efficiencies at low momenta, with one exception in the vicinity of the acceptance hole.

4.7. Systematic Uncertainties

The systematic uncertainty of a measurement represents the expected degree of reproducibility of its result, if parameters and procedures of the analysis are changed within limits of indistinguishable plausibility. It is commonly quantified by the variation of the result, as track selection criteria, signal extraction methods, and correction procedures are modified. A typical approach is the individual modification of each property with respect to a reference setting, and the quadratic addition of all resulting deviations from the reference result. This method imposes the assumption that no correlations exist

¹⁸Integration over wide $p_{T,ee}$ -intervals requires a realistic phase-space distribution of pairs, which adds input uncertainties. Instead, pairs are sampled flat in m_{ee} and $p_{T,ee}$ and the correction is done differentially. Therefore the visualization is also done in narrow slices.




Figure 4.25.: Electron efficiencies for 20 different cut settings, used to study systematic uncertainties.

between the varied properties, which however cannot be guaranteed in general, and is in particular not true for sets of observables that relate to the same detector. Additionally, given the limited statistics available to the present analysis, no clear indications for an “optimal” reference setting can be found.

To overcome these issues, a different approach is followed here. For the selection of electron candidates, multiple sensible cut values are chosen for each track-quality and PID variable. Random combinations of these are then used to define 20 individual *cut settings*, for each of which the analysis is performed. Finally, all resulting dielectron signals are used to find a representative result and obtain its systematic uncertainty. These steps are described in the following.

4.7.1. Cut Variations

The individual cut values or ranges for all track-quality and PID variables are mentioned in Section 4.2 alongside the presentation of their distributions. While their combination into cut settings is done randomly, the 20 settings used for analysis are selected such that they cover a large variety of combinations and no cuts that supersede each other. The resulting diversity of electron selections is best seen in their track efficiencies, which are presented in Figure 4.25. The individual p_T and η distributions mainly differ by magnitude and show the wide spread of integrated efficiencies, lying between 15 % and 40 %. Seven of the cut settings use the track recovery in regions of inoperative ITS-SPD sectors, as discussed in Section 4.2.1. This is best seen in some of the rightmost bins of the φ distribution, where only these settings (red to yellow) have a non-zero efficiency. The overall efficiency spread is well-represented also by this subset of cut settings,




Figure 4.26.: Left panel: comparison of the pair survival probability for both prefilter types. It is identical for all cut settings, since they use the same prefilter electron sample. Right panel: pair efficiencies for the 20 cut settings, when prefilter #2 is used. One transverse-momentum slice is shown (cp. Figure 4.24).

ensuring the same coverage of systematic uncertainties on the other track selection criteria.

The corresponding pair efficiencies resulting from the factorization procedure are shown in the right panel of Figure 4.26 for $1.9 < p_{T,ee} < 2.1$ GeV/ c . They include the pair survival probability from prefilter #2, which cuts on $m_{ee} < 40$ MeV/ c^2 and $\phi_V > \frac{3}{4}\pi$. A comparison of $P_{\text{surv}}^{\text{pair}}$ for both prefilter types is given in the left panel of Figure 4.26. Prefilter #1, using $m_{ee} < 40$ MeV/ c^2 and $\theta_{ee} < 80$ mrad, causes between two and three times as much random rejection as prefilter #2, increasing towards large values of $p_{T,ee}$. Note that the pair survival probability is identical for all cut settings. While the electron candidate selection and therefore the number and distribution of testparticles varies, all cut settings use the same relaxed selection criteria for the prefilter electron sample. Therefore the probability to reject any given testparticle stays constant. The systematic check of this procedure is covered by comparing corrected results from both prefilters and the unfiltered result with each other, which are consistent as described and shown in Section 4.4.2.

It is worth elaborating on the spread of pair efficiencies considered for the systematic study. As a function of invariant mass, they cover a factor of 5 to 8 between the most extreme cut settings, ensuring an extensive coverage of uncertainties related to electron selection and the corresponding Monte–Carlo properties. The variation of hadron con-

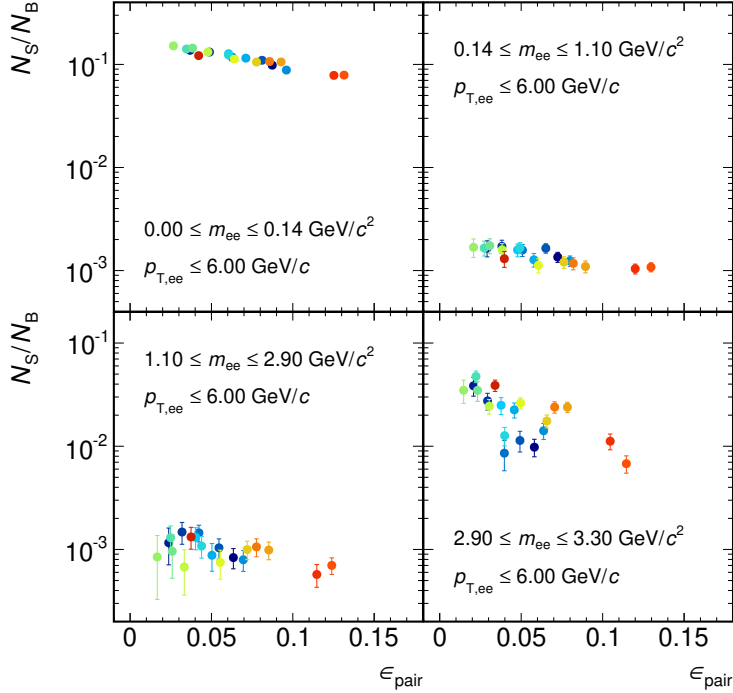


Figure 4.27.: Correlation between signal-to-background ratio and pair efficiency, integrated over the specified mass and pair- p_T ranges. See text for details.

tamination depending on the PID cuts is not reflected in the efficiency, because the latter is based on true electrons only. But it can be inferred from the correlation between the signal-to-background ratio and the pair efficiency, since S/B is independent of ϵ , as long as the fraction of electrons from findable dielectron sources in the electron candidate sample is unchanged. This is the case if the variation of electron efficiencies is constant as function of their kinematics. For p_T this is not quite fulfilled due to the pion exclusion cut which affects the lowest and highest momenta of some cut settings.

The correlation is presented in Figure 4.27, which shows a decreasing S/B as function of ϵ in all four invariant-mass regions. A reduction of S/B by 50% can be explained by hadron contamination if the latter causes no background in the cut settings at low ϵ_{pair} but half of the background at high ϵ_{pair} . While in practice a finite amount of hadrons is still present in the cut settings with strict PID, there may also be a higher background-contribution from single-electron sources in the relaxed cut settings. All things considered, a hadron contamination of 50% is a plausible value, also in relation to the estimates in [Kle14].

For comparison, the correlation between significance and pair efficiency is shown in

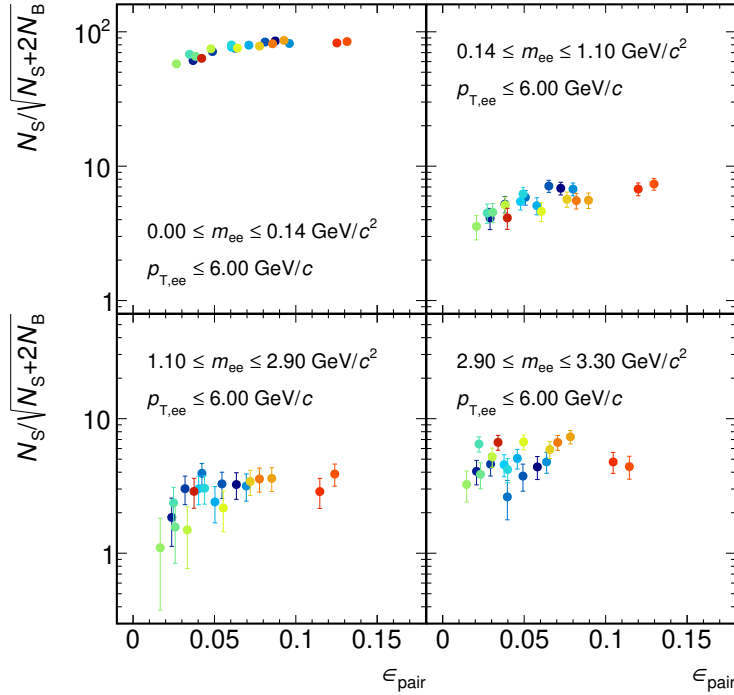


Figure 4.28.: Correlation between significance and pair efficiency.

Figure 4.28. The significance increases up to $\epsilon_{\text{pair}} \approx 8\%$, at which point the gain in pair statistics is compensated by the additional background.

4.7.2. Dielectron Signal Uncertainty

The majority of uncertainties on the dielectron signal is covered by the variation of results from the 20 cut settings. These cut variations primarily affect the electron selection, based on track quality criteria and particle identification measurements. But also the uncertainty on the background subtraction, which is the main contribution during dielectron signal extraction and includes effects from hadron contamination, is probed due to the spread in signal-to-background ratios.

According to [Bar02] — and common sense — only the significant part of a deviation should be considered as a systematic uncertainty. If the deviation is not significant, then the systematic check on the varied properties has passed and the deviation is not added to the systematic uncertainty of the final result. Since checks are usually performed on the same data, the results and fluctuations are strongly correlated, such that a significant deviation may be present even when their results “agree within their

statistical uncertainties”. Instead, the error σ_Δ on the deviation Δ has to be calculated and if the results differ by more than a reasonable number of σ_Δ , an incorporation into the systematic uncertainties has to be considered. For a check involving the total data T and a subset $U \subset T$, the error is given by $\sigma_\Delta^2 = \sigma_U^2 - \sigma_T^2$, while for a more complicated correlation only limits can be placed on it. [Bar02]

In the present analysis, the subtracted dielectron signal N_S of the cut setting with the highest pair efficiency unfortunately does not correspond to the total data T , meaning that not all signal pairs of U are also contained in T : $N_{S,U} \not\subset N_{S,T}$. This is caused by the change in combinatorial background, which is not limited to the kinematic region in which signal pairs are directly affected by tightening the cuts from T to U . This leads to U containing other and potentially more signal pairs than T . Therefore a different method was developed, which aims at minimizing statistical uncertainties on the subtracted signal before comparing results from the 20 cut settings. So initially, the dielectron yield is integrated over pair- p_T and divided into only five distinct invariant-mass regions.

The resulting distribution of signals and their statistical uncertainties is shown in Figure 4.29. At least for the more significant cut settings, the uncertainties are reduced to only 10–20% (with the exception of the last mass bin). They are thereby smaller than the spread of the individual signals, thus indicating that a systematic deviation is present. In the terminology of [Bar02], this procedure and its result correspond to a *check that fails*. Moderate deviations between cut settings are expected and reasonable however, so it becomes an *evaluation*, and the following method is chosen to evaluate the systematic uncertainty related to these cut variations: first, the average yield $\langle N_S(x) \rangle$ per mass bin x is computed, then the relative systematic uncertainty is defined as the RMS of the deviations between each signal and the average:

$$\frac{\Delta^{\text{syst.}}(x)}{N_S(x)} = \sqrt{\frac{1}{20} \sum_{i=1}^{20} (N_{S,i}(x) - \langle N_S(x) \rangle)^2}. \quad (4.21)$$

This avoids the inflation of uncertainties when more cut settings are added to the evaluation. It also removes the need for a reference cut setting and prevents double-counting of uncertainties from correlated quantities, as initially mentioned. The average yield is compared to the individual results in Figure 4.29 and the RMS from equation (4.21) is shown by its error bars.

The uncertainty from the ϕ_V template fit, which affects the first mass bins as described in Section 4.5.2, is quadratically added to the uncertainty from cut variations. The




Figure 4.29.: Dielectron signals of the 20 cut settings merged into wide invariant-mass bins to extract systematic uncertainties. The average yield is shown with square markers and their vertical error bars correspond to the RMS of all deviations from this average. The signal selected for the final result is shown in orange with dashed error bars.

combined systematic uncertainty is shown on the final result in Figure 4.30, which is further described below.

4.7.3. Representative Signal

Initially, all analysis settings are treated equal. Nevertheless, one of the cut settings has to be chosen for the final result¹⁹. A χ^2 -test on the two-dimensional yield is used to determine three signals which are most similar to the average computed previously. The one of them which features the highest visual similarity with the average is selected by hand as the representative signal. It is highlighted in Figure 4.29 with dashed lines.

¹⁹While the average of all cut variations could be considered for the result, its statistical errors cannot be properly defined, and it may contain further inconsistencies.

A good correspondence with the average is seen, except for a downward fluctuation by 18 % in the light-flavour resonance region²⁰.

4.8. Final Dielectron Yield

The final result of the analysis presented in this chapter is given in Figure 4.30 with blue circles. Shown is the subtracted dielectron yield from the selected cut setting as function of invariant mass, integrated over pair- p_T . Statistical uncertainties are given as error bars and the boxes show systematic uncertainties from cut variations and final pair cuts, added in quadrature. The invariant-mass distribution features a smooth transition from the low-mass to intermediate-mass region, due to the lack of statistical precision to resolve the ω and ϕ resonances in the vicinity of $\approx 1 \text{ GeV}/c^2$. A pronounced J/ψ peak is visible near its mass of $3.1 \text{ GeV}/c^2$. The expected yield from hadronic sources is further discussed in the following chapter.

For comparison to the present result, the measurement published by ALICE [ALICE19a] is also shown in Figure 4.30. It follows a largely different approach on the data, using more restrictive track quality and PID selection criteria, including a cut on the number of shared clusters in the ITS and the non-optional requirement of an electron-like signal in the TOF detector. This grants a higher purity ($\approx 90 \%$) but leads to lower pair efficiency ($\approx 1 \%$) and significance, which is reflected in the larger statistical uncertainties and fluctuations in the published data. The kinematic cuts for electrons are identical, except that the present analysis uses a fixed high- p_T cut of $p_T^e < 3.5 \text{ GeV}/c$ instead of relying on the pion exclusion cut in the TPC. The pair- p_T reach is therefore limited to $p_{T,ee}^{\text{max}} = 7 \text{ GeV}/c$ and for statistics reasons the kinematic range is further reduced to $p_{T,ee} < 6 \text{ GeV}/c$, while the published result does not contain an explicit restriction. Both analyses use a pair prefilter, the published one focussing primarily on conversion rejection, while the present one also reduces the π^0 Dalitz contribution. The methods to evaluate systematic uncertainties are similar, both using the RMS of the spread of corrected yields from cut variations to cover the main sources of uncertainties.

The degree of correlation between these results is hard to estimate, but based on the different analysis strategies and the factor 10 difference in pair efficiency, the correlation between the subtracted-pair statistics of both results is expected to be small. Systematic uncertainties are determined using similar approaches, and both analyses include the

²⁰The selection process was done on results without prefilter, while the figure shows the prefiltered data, from which systematic uncertainties are extracted. It was chosen not to change the selection in hindsight.




Figure 4.30.: Final dielectron signal of this analysis for the 0–10% most central collisions in comparison to the published result [ALICE19a]. See text for details.

study of a wide range of cut variations. While there may be some tension in a few data points, the overall agreement between these analyses gives confidence in the result of this challenging measurement.

Chapter 5.

Results

In this chapter the dielectron measurement is compared to different types of model calculations to investigate both the extent of agreement and the sensitivity towards expected deviations. The baseline for various further studies is the comparison of the data to all contributions from known hadronic sources of electron–positron pairs, called *hadronic cocktail*. Additional contributions from the QGP phase as well as further modifications due to the dense hadronic medium are expected in Pb–Pb collisions. Their influences on the dielectron yield are discussed subsequently.

5.1. Hadronic Cocktail

The hadronic cocktail includes light-flavour mesons, correlated open-heavy-flavour hadrons and heavy-flavour resonances. It is based on measurements of these particles in the same collision system and at the same center-of-mass energy, where available. Unavailable particle spectra are derived from others using first-principle approaches like scaling by transverse mass or by particle ratios measured in other collision systems. The heavy-flavour contributions are typically generated by pQCD-based models and scaled to the expected number of binary collisions.

Since the dielectron measurement is performed in a finite acceptance and also not corrected for the detector resolution, each effect has to be folded into the cocktail by first computing realistic hadron decays, then selecting and modifying the produced electrons correspondingly. Details on cocktail generation and inputs are presented in the following subsections.

By this construction, the hadronic cocktail is expected to agree with the dielectron measurement in proton–proton collisions, with one exception being the contribution from virtual direct photons [ALICE18a]. In Pb–Pb collisions it can still serve as a




Figure 5.1.: Comparison of the dielectron signal measured in 0–10% most central Pb–Pb collisions to the hadronic cocktail. Contributions from light-flavour mesons are based on a charged-pion measurement and m_T -scaling. The open heavy-flavour is generated with PYTHIA and scaled by the number of binary collisions in Pb–Pb. The J/ψ is constructed from measurements in pp and Pb–Pb, including its R_{AA} . Experimental resolution and acceptance are included in the cocktail.

baseline to study additional effects, but also to improve the understanding of detector and measurement properties.

Figure 5.1 shows the comparison between the measured dielectron signal and the hadronic cocktail in central Pb–Pb collisions. All particles contributing to the cocktail are drawn individually, while multiple decay channels per particle are merged. Overall, the invariant-mass distributions of data and cocktail are in reasonably good agreement, given the sizeable uncertainties on both sides. In the ratio a 2–3 σ upward shift of two datapoints is visible in the mass range dominated by the η meson and a 2 σ downward shift in the ω - ϕ region, both only considering statistical uncertainties. The latter is where the combinatorial background peaks, such that the data do not provide enough significance for a more differential view of the expected two-peak structure.

In Figure 5.2, the pair-transverse-momentum spectra for different invariant-mass regions are presented. The low-mass region with $m_{ee} < 0.14$ GeV/ c^2 is dominated by the π^0 -Dalitz

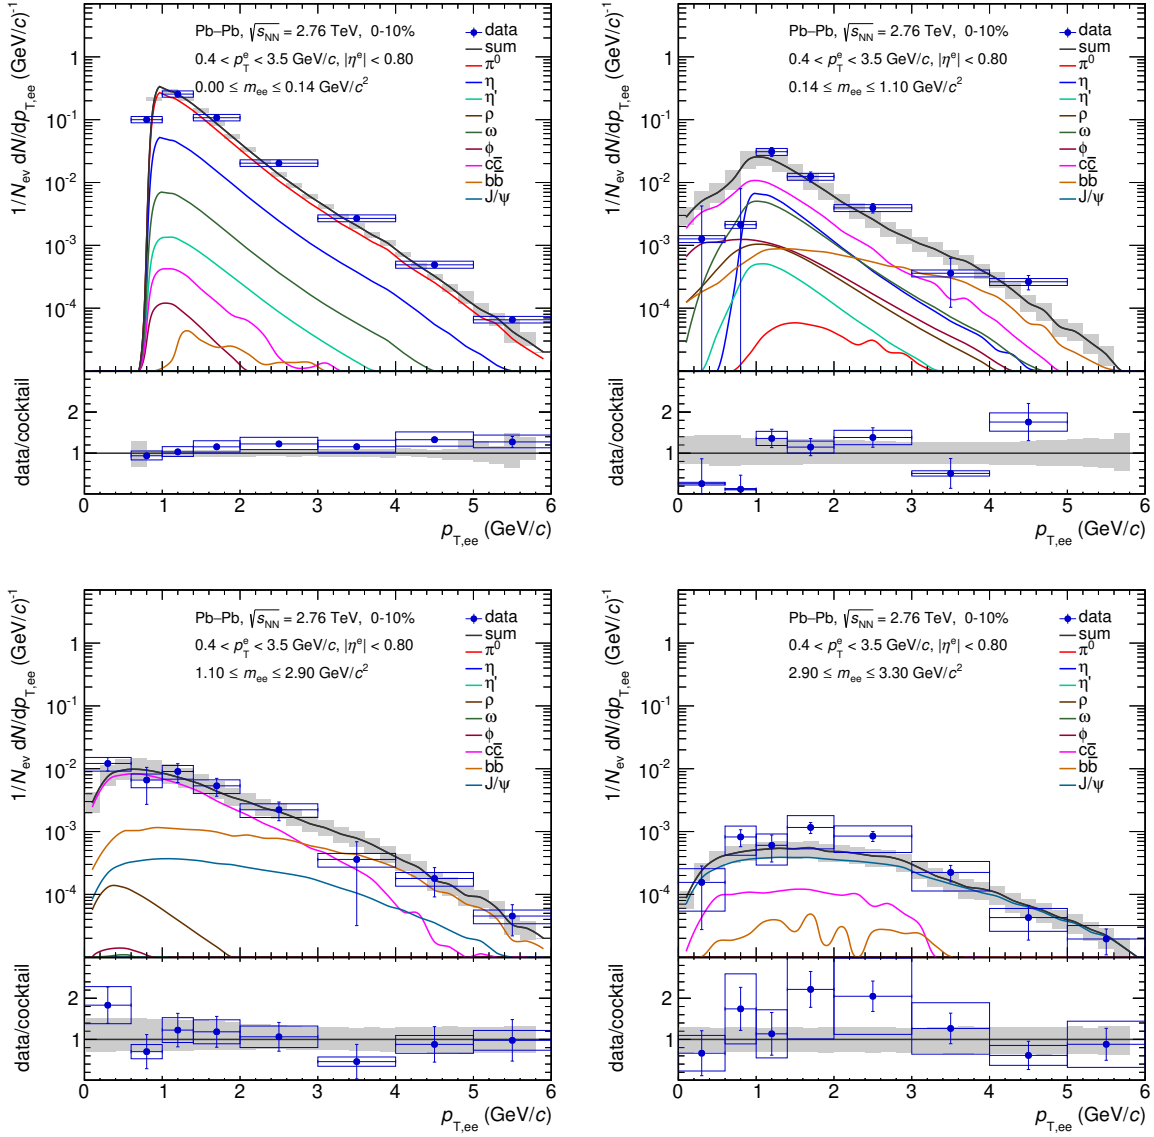


Figure 5.2.: Pair-transverse-momentum dependence of the comparison between data and cocktail from Figure 5.1, divided into four invariant-mass regions. See text for details.

yield, and the other mesons follow its spectral shape because of their construction via m_T -scaling (see Section 5.1.2). The slope of the cocktail is slightly steeper than that of the data, allowing for a contribution from virtual direct photons towards higher $p_{T,ee}$. In the resonance region up to $m_{ee} = 1.1$ GeV/ c^2 , large statistical fluctuations are visible in the data. They are caused by the small signal-to-background ratio, which reaches its minimum of $N_S/N_B \approx 10^{-4}$ at $m_{ee} \approx p_{T,ee} \approx 1$ GeV. The more significant data points are in reasonable agreement with the cocktail. Therefore the discrepancy between data

and cocktail seen in this mass range in Figure 5.1 can be attributed to the downward fluctuation at low $p_{T,ee}$.

The intermediate-mass region, being dominated by the open-charm and -beauty contributions, shows a good correspondence between measurement and cocktail. This however may be coincidental, since electrons from heavy-flavours are suppressed in nucleus-nucleus collisions (see Section 5.2.1), but additional thermal contributions are expected (see Section 5.4), which may compensate the suppression. The kinks at $p_{T,ee} \approx 3.9$ GeV/ c in the heavy-flavour contributions are caused by the upper transverse-momentum cut of $p_T^e < 3.5$ GeV/ c .

In the high-mass region, the measured $p_{T,ee}$ dependence of the J/ψ yield is not perfectly reproduced by the cocktail. The latter is based on a measurement in proton-proton collisions and does not account for the $p_{T,ee}$ dependence of the J/ψ suppression in Pb-Pb collisions (see Section 5.1.2).

5.1.1. Cocktail Production

The generation of each cocktail component consists of the following main steps:

1. A suitable input is selected to provide the kinematics of the respective dielectron source. For the light-flavour and J/ψ mesons, measured p_T -spectra are parameterized to obtain smooth distributions throughout the p_T -interval required in this analysis. The open-heavy-flavour contributions are modelled by PYTHIA 6 [SMS06] and scaled from pp to the number of binary collisions in (semi-)central Pb-Pb collisions. Details on the individual inputs are provided below.
2. Electrons and positrons are produced from this input using an appropriate decay mechanism. An EXODUS¹ implementation is used for the light-flavour and J/ψ mesons, ensuring the use of electromagnetic transition form factors and distribution functions to best knowledge [EO14]. PYTHIA is used for the semileptonic decays of open-heavy-flavours to provide a self-consistent estimate for these contributions.
3. Different influences on the measured electron kinematics are considered: effects that actually modify the electrons are multiple scattering and bremsstrahlung, while the detector resolution changes only the apparent kinematics. The resolution maps created for the efficiency correction (Figure 4.21 in Section 4.6) contain a superposition of mainly multiple scattering and detector resolution, and are applied

¹EXODUS is a phenomenological hadron decay generator developed by the PHENIX collaboration to simulate phase space distributions of electron sources and their decays [PHENIX10].

to the generated electrons at this point. Bremsstrahlung is taken into account already within EXODUS.

4. The kinematic track and pair cuts used during data analysis are applied². Thereby only dielectrons within the utilized detector acceptance enter the cocktail. The combination of this and the previous step, in the given order, ensures the consistent representation of data and cocktail based on *measurable* quantities, as discussed in Section 4.6.1.

5.1.2. Cocktail Inputs

Light Flavour

Ideally, the p_T -differential yields of all particle species with decay channels into dielectrons are known for the respective collision system and energy. The light-flavour cocktail then contains all medium-modifications on hadronic dielectron sources, attributing any deviation from the measurement to thermal radiation. For the present analysis however, a precise measurement down to low transverse momentum was only available from charged pions [ALICE14d], which is shown in Figure 5.3. The data and their upper and lower systematic variations are parameterized by a modified Hagedorn function, which transitions from an exponential at low p_T to a power law at high p_T [AB⁺17]. The fit range is chosen such that the low- p_T data, being most relevant to this analysis, are reproduced well. These parameterizations are used to generate the π^0 -Dalitz component and its uncertainty interval.

Lacking measurements of the other cocktail particles, their input spectra are constructed from the pion parameterization with a method called m_T -scaling, utilizing the transverse mass defined as $m_T = \sqrt{m_0^2 + p_T^2}$. It originates from the concept of statistical thermodynamical hadron production, according to which all hadron m_T -spectra should follow exponential distributions with a universal inverse-slope parameter, and has been tested for various measurements at different collision energies [AB⁺17, KSS11, KS⁺12]. According to m_T -scaling, the p_T -spectrum of a reference particle R can be transformed into a production spectrum for a different particle D by the substitution

$$p_{T,R} \rightarrow \sqrt{m_{0,D}^2 - m_{0,R}^2 + p_{T,D}^2}. \quad (5.1)$$

²One exception is the ϕ_V cut, since the data are corrected for it.




Figure 5.3.: Transverse-momentum spectrum of charged pions measured in central Pb–Pb collisions [ALICE14d]. Lines show parameterizations up to $p_T = 7$ GeV/c, created in this analysis as input for generating the light-flavour cocktail.

The yield of particle D relative to R is normalized using the ratio of their spectra at high p_T , which may be obtained from measurements at a lower collision energy.

It should be noted that m_T -scaling is not expected to be generally fulfilled in nucleus-nucleus collisions, owing to the presence of medium effects. In particular radial flow during the QGP phase causes spectral modifications that do not follow the mass hierarchy of the mesons, since it acts on the partonic level. This effect was recently studied to estimate the p_T -spectrum of the ω meson in Pb–Pb collisions [ALICE19b]. The result is shown in the left panel of Figure 5.4 relative to the charged-pion yield and compared to the ratio obtained via m_T -scaling. The ω yield is estimated to be enhanced at intermediate p_T due to a stronger boost from low- p_T .

In contrast to this, a recent measurement of the η/π -ratio in Pb–Pb collisions [ALICE18c] does not significantly indicate an enhancement with respect to pp collisions, as shown in the right panel of Figure 5.4. It is therefore reasonably well described via m_T -scaling, which tends to slightly overestimate the measured η/π -ratio in pp collisions at low p_T , when performed based on the charged-pion yield [AB⁺17].

An uncertainty of $\pm 20\%$ is assigned to all m_T -scaled contributions in addition to

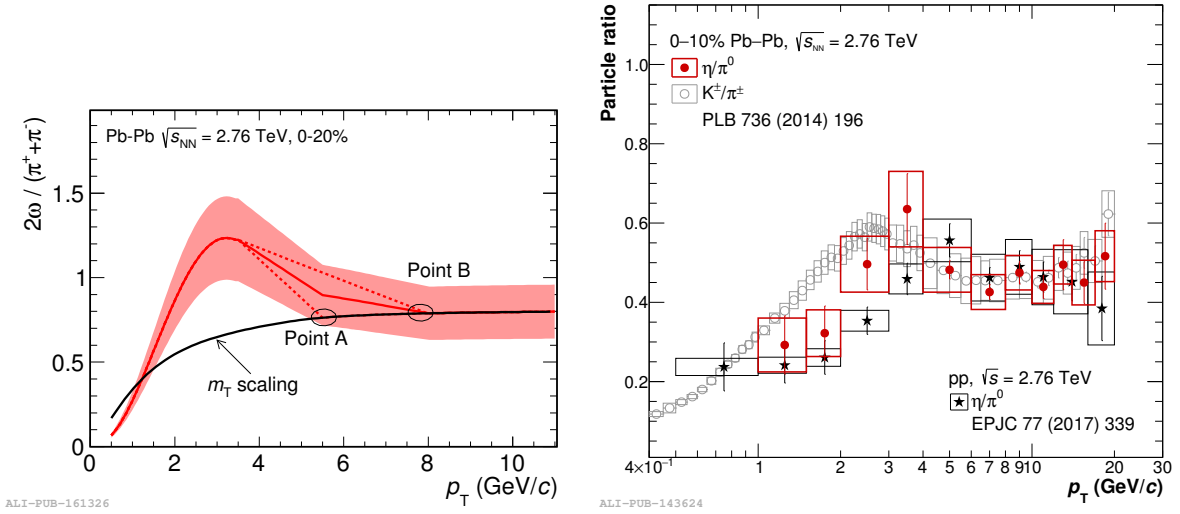


Figure 5.4.: Left panel: estimate of the ω/π -ratio in Pb–Pb-collisions based on Blast-Wave fits to evaluate the collective radial flow, in comparison to the expectation from m_T -scaling [ALICE19b]. Right panel: measurements of the η/π -ratio in Pb–Pb and pp collisions are consistent within current uncertainties [ALICE18c]. m_T -scaling is used in the present work for both contributions.

propagating the upper and lower pion parameterization. Uncertainties on the decay branching ratios are much smaller but also included.

Heavy Flavour

For the contributions from correlated open-heavy-flavour mesons a different approach is needed, since neither the p_T -spectra of the individual mesons nor those of their leptonic decay products carry information on the resulting pair kinematics of the latter. Their construction therefore relies on the PYTHIA 6 event generator [SMS06] using the Perugia 2011 tune, which includes insights from early LHC data [Ska10]. It is configured to generate at least one $c\bar{c}$ or $b\bar{b}$ pair per event, then hadronize it and decay its products. Dielectrons from individual heavy-flavour quark pairs consist not only of unlike-sign but also of like-sign signal pairs, owing to oscillations of B^0 and more rarely of D^0 mesons, and because the $B\bar{B}$ decay chain may produce more than two charged leptons. Both effects are reproduced by PYTHIA and in addition the generator may create more than one heavy-flavour quark pair in an event if multiple hard interactions occur. Therefore a separate pairing for unlike-sign and like-sign is performed on the output and the like-sign yield is subtracted.

The per-event dielectron yield from $c\bar{c}$ is thus obtained via

$$\frac{1}{N_{\text{ev}}} N_{\text{ee}\leftarrow c\bar{c}} = \frac{N_{\text{ee}\leftarrow c\bar{c}}^{\text{ULS}} - N_{\text{ee}\leftarrow c\bar{c}}^{\text{LS}}}{N_{c\bar{c},\text{gen}}} \cdot \frac{\sigma_{c\bar{c}}^{\text{tot}}}{\sigma_{\text{INEL}}} \cdot N_{\text{coll}}, \quad (5.2)$$

and for $b\bar{b}$ accordingly.

The normalization per generated heavy-flavour quark pair, $N_{c\bar{c},\text{gen}}$ resp. $N_{b\bar{b},\text{gen}}$, allows for the use of cross section measurements from ALICE to scale the results to the expectations for Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. The corresponding cross sections are listed in Table 5.1. They contain quite sizeable uncertainties, dominated by the extrapolation into full phase space in case of charm and by the uncertainty on the background from charm-hadron decays in case of beauty, which are propagated into the total cocktail uncertainty.

cross section (2.76 TeV)	value	uncertainty		ref.
		statistical	systematic	
$\sigma_{c\bar{c}}^{\text{tot}}$	4.8 mb	$\pm 17\%$	$^{+58}_{-29}\%$	[ALICE12a]
$\sigma_{b\bar{b}}^{\text{tot}}$	0.13 mb	$\pm 12\%$	$^{+33}_{-39}\%$	[ALICE14a]
σ_{INEL}	62.8 mb		$^{+4}_{-7}\%$	[ALICE13b]

Table 5.1.: Total charm and beauty cross sections as well as inelastic proton–proton cross section at 2.76 TeV, used to normalize the heavy-flavour cocktail contributions.

For the number of binary collisions, results from Glauber Monte–Carlo simulations are applied: $N_{\text{coll}}(0\text{--}10\%) = 1500.5$ and $N_{\text{coll}}(10\text{--}50\%) = 492.2$ [ALICE13a]. Their uncertainties are on the order of 1–10% and not taken into account here. Instead, alternatives to a pure N_{coll} -scaling are discussed in Section 5.2.

Charmonium

The J/ψ meson is the single quarkonium state considered for this cocktail. Its transverse-momentum and rapidity distributions measured in pp collisions at $\sqrt{s} = 7$ TeV [ALICE11b] are used in combination with the EXODUS decayer to obtain realistic dielectron kinematics [Gun13]. Figure 5.5 shows the J/ψ p_{T} -spectrum and its parameterization with an exponential function serving as generator input. Acceptance, resolution, and bremsstrahlung effects are applied to the decay electrons as for the other cocktail contributions.

For the normalization to Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV, the result is scaled by




Figure 5.5.: Transverse-momentum spectrum of the J/ψ meson measured in pp collisions [ALICE11b, Gun13]. Lines show parameterizations used in this analysis for generating the J/ψ cocktail contribution, in combination with scalings for collision energy and R_{AA} .

the ratio of J/ψ cross sections and the nuclear modification factor R_{AA} :

$$N_{ee \leftarrow J/\psi}^{\text{Pb-Pb}}(2.76 \text{ TeV}) = N_{ee \leftarrow J/\psi}^{\text{pp}}(7 \text{ TeV}) \cdot \frac{d\sigma_{J/\psi}^{\text{pp}}/dy(2.76 \text{ TeV})}{d\sigma_{J/\psi}^{\text{pp}}/dy(7 \text{ TeV})} \cdot R_{AA}^{J/\psi}. \quad (5.3)$$

The corresponding differential cross sections are listed in Table 5.2. An interpolation between results from multiple experiments was used for the value at 2.76 TeV, leading to a small combined uncertainty [ALICE14b].

\sqrt{s}	$d\sigma_{J/\psi}^{\text{pp}}/dy _{ y <0.9}$	total uncertainty	ref.
2.76 TeV (interp.)	$4.25 \mu\text{b}$	$\pm 12 \%$	[ALICE14b]
7 TeV	$6.90 \mu\text{b}$	$^{+26}_{-33} \%$	[ALICE11b]

Table 5.2.: Differential J/ψ cross sections around midrapidity in pp collisions, used for scaling the J/ψ cocktail contribution to the Pb–Pb collision energy.

Measurements of the J/ψ R_{AA} in three centrality intervals are summarized in Table 5.3 [Boo14]. The semi-central event class is defined differently but used here for the cocktail of the 10–50 % centrality class. This approximation can be justified with the similar R_{AA} observed in peripheral collisions. Two types of systematic uncertainties are quoted. The ones correlated between centrality intervals are dominated by the pp reference uncertainty, which is connected to the pp cross sections whose uncertainties also have to be accounted for according to equation (5.3). Therefore only the latter and the uncorrelated R_{AA} uncertainty are added in quadrature to obtain the total uncertainty on the J/ψ cocktail contribution.

centrality	J/ψ R_{AA}	systematic uncertainty	
		correlated	uncorrelated
0–10 %	0.65		± 12.8 %
10–40 %	0.74	± 13.2 %	± 12.5 %
40–90 %	0.79		± 9.4 %

Table 5.3.: Nuclear modification factor of the J/ψ meson in 2.76 TeV Pb–Pb collisions of different centralities, measured at midrapidity [Boo14].

Also a p_T dependence of the J/ψ R_{AA} is observed, featuring little suppression at low momentum as expected from charm quark recombination in the medium and a strong suppression towards high p_T due to energy loss [ALICE14b, Boo14]. Despite integrating over wider centrality ranges, statistics are however insufficient for a reliable parameterization of this dependence, in particular for the midrapidity measurement, which should be used here. Since the present analysis is divided into centrality classes and not primarily focussed on pair- p_T spectra, only the centrality dependence of the J/ψ R_{AA} is taken into account.

The dielectron yield from the Drell–Yan process is not included in the cocktail, since its contribution is subdominant in the kinematic region relevant to this analysis (see Section 2.1.1).

5.2. Heavy-Flavour Modifications

Starting from $m_{ee} \approx 350$ MeV/ c^2 , the dielectron signal is dominated by semileptonic decays of correlated open-heavy-flavour mesons, with the only exception being the J/ψ

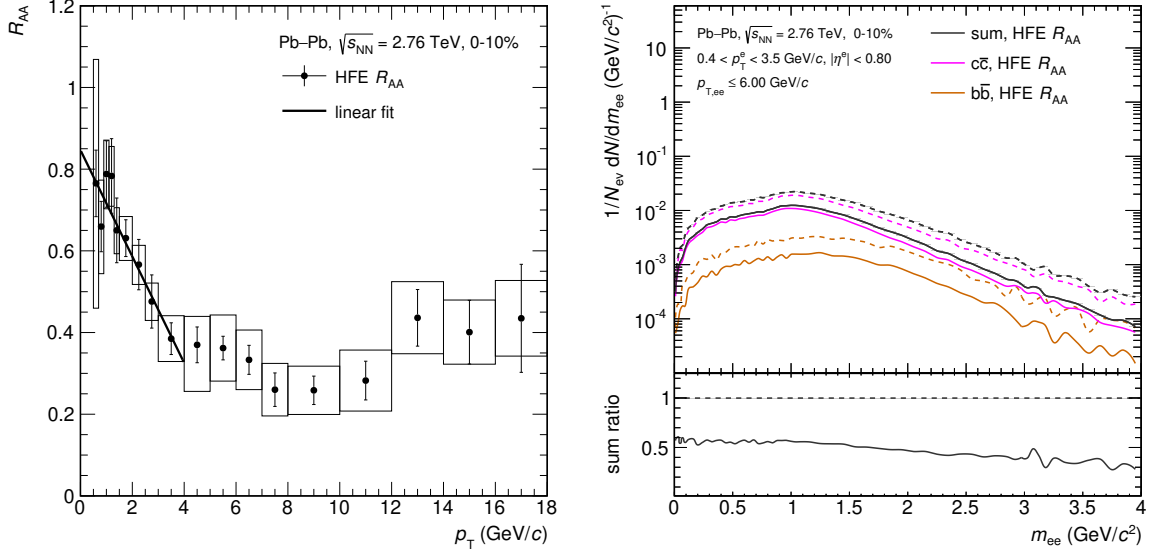


Figure 5.6.: Left panel: heavy-flavour electron R_{AA} measurement for 0–10% centrality [ALICE18b] and parameterization at low p_T used for the alternative heavy-flavour-cocktail production. Right panel: mass dependence of both heavy-flavour contributions with the R_{AA} applied to their electrons (solid lines) in comparison to the original case (dashed lines). The ratio of the heavy-flavour sum with R_{AA} to the original sum is given at the bottom.

peak. Further studies are thus performed to evaluate the sensitivity of the data on modifications of these contributions.

5.2.1. Nuclear Modification Factor of HFE

For the initial cocktail generation, a binary-collision scaling with respect to pp collisions is applied to both open-heavy-flavour contributions. This preserves the original mass and momentum distributions and thereby assumes no interactions with the medium. In contrast to this, the nuclear modification factor R_{AA} of electrons from heavy-flavour-hadron decays (“heavy-flavour electrons”, HFE) has been measured to be below unity [ALICE18b]. The p_T dependence of the HFE R_{AA} for central Pb–Pb collisions is shown in the left panel of Figure 5.6. The low- p_T part, which is relevant to the present analysis, is parameterized with a linear function that can be applied as a weight to the heavy-flavour electrons during cocktail production. This leads to a more data-driven approximation of the heavy-flavour contributions, more consistent to the other cocktail components. Yet the electron kinematics are still based on the PYTHIA generator and correlated suppression or non-suppression of HFE due to in-medium effects on quarkonia are not included.




Figure 5.7.: Data compared to hadronic cocktail with R_{AA} from heavy-flavour electrons applied to the charm and beauty contributions. The suppressed heavy-flavour gives more room for a low-mass enhancement, but underestimates the data in the intermediate mass range and beyond the J/ψ mass.

The result of this procedure is presented in the right panel of Figure 5.6, where the modified mass distributions of charm and beauty are compared to their original versions. Below that, the ratio of the modified to the original heavy-flavour sum is drawn. Following this approach, a sizeable suppression of the open-heavy-flavour contributions to the dielectron spectrum is expected, ranging from 40% to 70% with increasing invariant mass.

The modified hadronic cocktail is compared to the data in Figure 5.7, where the light-flavour components are merged into a single contribution. While the suppressed heavy-flavour contributions lead to a better match for $m_{ee} \lesssim 1 \text{ GeV}/c^2$ and give slightly more room for a low-mass enhancement, they do cause the cocktail to significantly underestimate the data in the intermediate-mass region and beyond the J/ψ mass. Some tension is also visible in the J/ψ peak, where the data exceed the cocktail by 30–40%. The method presented here only applies a weight — the R_{AA} — to the heavy-flavour electrons, but it preserves the kinematics of each individual track and pair. Given the dense medium in a Pb–Pb collision, with known effects like jet quenching or flow, this

certainly is a simplification. The influence of modified kinematics on the invariant-mass distribution of open heavy-flavours is therefore discussed in the following.

5.2.2. Heuristic Charm Modifications

To estimate medium-like effects, which may influence the kinematics of open-heavy-flavour pairs as an addition or alternative to the HFE R_{AA} , three types of kinematic modifications are applied to the electrons of the original open-charm contribution individually. Each modification should be considered as or actually is the upper limit of its type:

- Randomization of the azimuthal angle φ^e . This allows for all opening angles, while preserving the rapidity of the pair.
- Sampling of the pseudorapidity η^e from the original distribution. Initially, η^e of the two electrons is correlated, with the $\Delta\eta^e$ having an RMS of 1.9 over the whole pair-rapidity range. When this sampling is applied, the correlation is lost and the RMS increases to 3.7 for pairs around midrapidity.
- Smearing of φ^e and η^e around the original value according to a Gaussian distribution. This causes only a partial decorrelation in both dimensions. A width of $\sigma_\varphi = \sigma_\eta = 1$ is chosen for the Gaussians. In this case the $\Delta\eta^e$ RMS only moderately increases to 2.4.

For a comparison of these modified charm contributions to data, all other components of the hadronic cocktail are subtracted from the measured yield. The result is shown in Figure 5.8. It includes the original charm as a reference, which also serves as denominator for the ratio. When comparing to data, uncertainties of the subtracted cocktail and the charm have to be considered in addition to the data uncertainties presented here.

Of the three open-charm modifications, the randomized azimuthal angle has the least effect on the invariant-mass distribution, suggesting a moderate redistribution of pairs from high to low invariant mass. The smeared φ^e and η^e lead to an almost mass-independent suppression of 30–40 % and the sampled η^e causes the strongest suppression, reaching 50 % at low and ≈ 30 % at high invariant mass, which is opposite to the effect from applying the heavy-flavour-electron R_{AA} .

No electrons are explicitly rejected or down-weighted within these scenarios. Therefore the net suppression in the two latter cases is caused by the widening of the $\Delta\eta^e$ distributions as stated above, thus shifting legs from more pairs out of the detector acceptance.




Figure 5.8.: Various modified charm contributions and cocktail-subtracted (except $c\bar{c}$) data in comparison to the nominal N_{coll} -scaled charm. Randomization of φ^e slightly softens the mass distribution and preserves all pairs. Sampling from η^e causes a net increase in m_{ee} and a suppression due to more pairs missing the acceptance. Gaussian smearing of φ^e and η^e provides a mix of these effects, leading to moderate, mass-independent suppression.

5.2.3. Combination of Approaches

Given that the open-heavy-flavour modification using the HFE R_{AA} only applies weights on electrons while the heuristic scenarios only modify their kinematics, combinations of the two can be sensible. This is clearly true for the φ^e -randomization, which by itself causes no suppression. Both other heuristic scenarios cause a pair suppression, but behave differently on the track level. In the smearing scenario, more electrons are shifted out of the acceptance than back in, due to the bell-shaped η^e distribution. It therefore features its own electron suppression and should be considered as an alternative to the modification caused by HFE- R_{AA} scaling. On the other hand, in the η^e -sampling scenario the distributions and the yield of single electrons remain unchanged, which corresponds to an HFE R_{AA} still at unity. The pair suppression is solely caused by the decorrelation of the legs, suggesting that the measured R_{AA} values may be applied in addition. This

would lead to an enormous suppression. It can be argued however, that the rapidity change a particle can undergo during the expansion of the medium is limited by the volume of causally connected particles, which may be as small as one unit in rapidity. This makes the free sampling of η^e appear to be a too strong assumption.

In a similar study by the PHENIX collaboration, the $c\bar{c}$ contribution is created by pairing electrons sampled from HFE p_T -spectra using random opening angles [PHENIX10]. It thereby is essentially a combination of the φ^e -randomization and the HFE- R_{AA} scaling presented here. PHENIX finds a substantial softening of the mass distribution, attributing it to less pairs from back-to-back particles in comparison to PYTHIA.

Indeed the randomized azimuthal angle does cause a slight softening of the mass distribution, as seen by the long-dashed line in Figure 5.8. However, the present study suggests that instead the p_T dependence of the HFE R_{AA} , which is similar in central collisions at $\sqrt{s_{NN}} = 200$ GeV and 2.76 TeV [PHENIX07, ALICE18b], is the dominant reason for the observed mass dependence (see ratio in the right panel of Figure 5.6). Also the integrated suppression of the dielectron yield in the intermediate-mass region by a factor of ≈ 2 is similar between PHENIX and the result from HFE- R_{AA} scaling shown in Figure 5.7.

In conclusion, as also suggested by PHENIX, the results from PYTHIA should be considered as the upper and the ones with HFE R_{AA} (and randomized angle) as the lower limit on the heavy-flavour contributions to the cocktail.

5.3. Increasing Dielectron Acceptance

For studies of the intermediate-mass region, an electron- p_T range of $p_T^e > 0.4$ GeV/ c seems ideal, since the selected tracks essentially saturate the dielectron phase-space for invariant masses of $m_{ee} \gtrsim 1.5$ GeV/ c^2 . Below this mass, softer electrons start to significantly contribute to the dielectron yield and increase the pair acceptance to lower $p_{T,ee}$.

In the present analysis, the low- p_T cut can be seamlessly reduced to $p_T^e = 0.2$ GeV/ c , facilitated by the specialized PID strategy, which does not require a time-of-flight signal (see Section 4.2.2). This lower limit on the electron- p_T is well-established from low-mass dielectron measurements by the CERES, PHENIX, and STAR collaborations, as summarized in Section 2.2. Achieving a corresponding result is therefore essential for a comparison to model calculations that aim for a description of the low-mass region and have adopted the experimentally accessible phase-space for their post- and predictions on dielectron yields.

Figure 5.9 shows the measured signal in comparison to the unmodified hadronic cocktail,




Figure 5.9.: Comparison of the dielectron signal measured in 0–10% most central Pb–Pb collisions to the unmodified hadronic cocktail. Same as Figure 5.1, except that electrons are selected in an increased transverse-momentum range of $0.2 < p_T^e < 3.5$ GeV/ c to match with existing measurements and model calculations.

using the reduced electron- p_T cut. The result is qualitatively similar to the original one, with the main difference being that the yield for $m_{ee} \lesssim 0.6$ GeV/ c^2 increases by a factor of five.

5.4. Model Calculations

In addition to electron–positron pairs from hadronic final-state decays, i. e. the hadronic cocktail, thermal contributions from the different stages of a heavy-ion collision are present in the dielectron signal. Their characterization lies beyond the realm of an experimental analysis, such that theoretical descriptions are consulted.

Various approaches exist on the modelling of heavy-ion collisions, which can be grouped into two types. The created medium is either treated as a collective system and described by its effective properties, or as individual particles with microscopic propagation and interactions between them. Results from one of either model types are compared to the experimental data in this section.

5.4.1. Thermal Fireball

Since heavy-ion collisions at LHC energy predominantly evolve through the QGP phase, in which the medium can be characterized as a perfect fluid, the natural theory to describe this evolution is relativistic hydrodynamics. The thermal fireball model [Rap01] provides a simplification to this approach, which is more suitable for calculations of the thermal dielectron yield, since the system evolution can be parameterized. This is done by assuming a cylindrical initial volume, determined by the impact parameter and a formation time after the nuclei intersect, and expanding it using a constant longitudinal velocity and transverse acceleration. The latter is tuned to reproduce a realistic radial flow and freeze-out time [Rap13].

The time evolution of the fireball temperature is determined from an equation of state that is based on lattice-QCD calculations using a crossover phase transition to a hadron resonance gas around $T = 160$ MeV [HFR12]. This temperature dependence influences the emission of thermal dileptons in both phases, which is reflected in the corresponding electromagnetic spectral function. The lattice-QCD results show that, in case of a thermal rather than perturbative QGP, the annihilation of quark–antiquark pairs no longer dominates the production of dileptons, except for high invariant masses. A wide resonance contribution is therefore added at low energy to the otherwise mass-independent spectral function of the QGP [Rap13]. In the hadronic phase, the spectral function is dominated by the short-lived ρ meson, which has a quite temperature-sensitive lineshape [RW99]. Using these two contributions, the partonic and hadronic rates of thermal dilepton emission are calculated via the McLerran-Toimela formula [MT85] (see Section 2.1.2, equation (2.21)). The rates are separately integrated over the collision evolution, leading to dielectron yields from both phases.

For 0–10% central Pb–Pb collisions at 2.76 TeV, an initial fireball temperature $T_{\text{init}} = 565$ MeV and a charged-particle multiplicity density $\langle dN_{\text{ch}}/dy \rangle = 1600$ are chosen [Rap16]. The resulting thermal yields from QGP and hadron gas (HG) are presented in Figure 5.10 with blue and green lines, respectively³. The thermal fireball model does not include a simulation of the hadronic cocktail, with the exception of the ρ meson. In its case, the model can provide a more consistent and accurate description of the final-state decays, influenced by medium effects. Its contribution is added to the hadron gas component in the figure.

³Model yields are provided up to $m_{ee} = 1.5$ GeV/ c^2 and extrapolated here using the functional form $N \propto m^{3/2} \cdot \exp(-m/p_1)$. Due to the low electron- p_T cut, the saturation of the dilepton phase-space starts below $m_{ee} = 1.5$ GeV/ c^2 , allowing for robust fits. These fits also constrain the extrapolation of alternative results with $p_T^e > 0.4$ GeV/ c shown below.




Figure 5.10.: Data of Figure 5.9 compared to hadronic cocktail plus thermal contributions from the hadronic and QGP phases based on a thermal fireball model [Rap13, Rap16]. The vacuum- ρ is excluded from the cocktail since the model provides its final-state decays including medium effects. The low-mass enhancement predicted by the model is only partially matched by the data. See text for details.

For comparison to the measurement, all remaining hadronic contributions are taken from the cocktail described in Section 5.1 — without heavy-flavour modification. They are summarized into the red line, which also serves as the reference for the ratio. The additional thermal contributions from the model predict a significant enhancement of the dielectron yield in the low-mass region between the π^0 and the ω mass (black line), ranging from $\approx 40\%$ to 80% above the cocktail. The data seem to confirm this trend for the lower part of this mass range, while they appear to fluctuate around the pure cocktail in the upper part.

The model also provides predictions for the higher electron- p_T cut of $p_T^e > 0.4$ GeV/ c , which are compared to the corresponding data in Figure 5.11. In this case the expected low-mass enhancement is slightly increased to $\approx 50\text{--}100\%$. The data may again favour its lower part, but they barely match the cocktail in the upper three mass bins. Various causes could be blamed for the latter, including a likely cocktail overestimate due to the N_{coll} -scaled heavy-flavour contributions as discussed above, the peaking combinatorial




Figure 5.11.: Data compared to hadronic cocktail plus thermal contributions from the hadronic and QGP phases based on a thermal fireball model. Same as Figure 5.10, but electrons are selected in a transverse-momentum range of $0.4 < p_T^e < 3.5$ GeV/ c , which is the default analysis setting in this work. The individual cocktail contributions are shown in Figure 5.1.

background in this region which may cause oversubtraction, and the below-average yield of the final analysis setting in comparison to its systematic variations (see Section 4.7.3). Yet even without these arguments, the model prediction is just within reach of the data, considering the illustrated statistical and systematic uncertainties.

5.4.2. Microscopic Transport

An alternative approach to describe the evolution of heavy-ion collisions is realized in transport models. By simulating the particle interactions on a microscopic level, they provide one formalism that is consistently used during all collision stages. This is their main advantage over hydrodynamic models, which require external inputs for the initial state and an additional mechanism for final-state particle production, if desired. Transport models do however have their limitations, particularly in the hot and dense collision phases, where different QCD processes are interfering with one another.

One such transport model is UrQMD [BB⁺98, BZ⁺99], which describes the collision evolution on the level of hadrons, propagating them on classical trajectories with a stochastic treatment of their interactions. These interactions include binary elastic and inelastic scatterings and resonance decays, as well as color string formation for masses of $m \gtrsim 2 \text{ GeV}/c^2$. The UrQMD approach thereby corresponds to a Monte–Carlo solution of Boltzmann equations for the evolution of the various particle distributions [BZ⁺99].

By simulating all relevant hadron–hadron interactions and the decays of their final-states, the model provides its own version of the light-flavour cocktail. However, no thermal contributions emerge from the microscopic transport itself. Newer versions of UrQMD therefore apply a “coarse-graining” technique, where the particle distributions are averaged over multiple events and then sliced into small space–time cells [EvHB16]. Particles in these cells are in approximate equilibrium, such that thermodynamic quantities can be extracted by applying a corresponding equation of state. Similarly to the thermal fireball model described above, a hadron gas EoS is used for cells up to $T = 170 \text{ MeV}$ and a lattice-QCD EoS⁴ for higher energy densities. Finally, the thermal dilepton rates from both phases are calculated via the McLerran-Toimela formula, using vector-meson dominance for the low-mass region and multi-pion states for the intermediate-mass region of the hadronic phase, and applying a lattice-QCD spectral function for the QGP phase [EvHB16] — again similarly to the thermal fireball model.

Figure 5.12 shows the combined UrQMD plus coarse-graining results for the low-mass dielectron yield in Pb–Pb collisions with 0–10% centrality. The UrQMD part corresponds to the light-flavour fraction of the hadronic cocktail, including medium effects and freeze-out contributions from the ω and ρ mesons. Coarse-graining provides three thermal contributions from the hadron gas: a dominating in-medium ρ , a smaller in-medium ω , and a tail from multi-pion states reaching into the intermediate-mass region. For a comparison to the hadronic yield from the thermal fireball model, the freeze-out ρ should be added here, leading to $\approx 60\%$ more thermal HG yield in the ρ peak and clearly less broadening towards lower mass. The QGP yield predicted by coarse-graining dominates over the remaining components throughout most of the mass region, and is larger than the one predicted by the fireball model by approximately a factor of two at low mass⁵.

While the coarse-grained microscopic transport approach presented here provides a more

⁴The same lattice equation of state is applied by both models, except that the chosen transition temperature differs by 10 MeV.

⁵This discrepancy increases in the intermediate-mass region, but there the fireball model yield is extrapolated (see above), thus to be taken with caution. However, the use of a partonic lattice EoS in UrQMD introduces some inconsistency, given the underlying hadronic degrees of freedom. If a hadron gas EoS is also applied for the partonic phase, the QGP yield above $m_{ee} \approx 1 \text{ GeV}/c^2$ softens [EvHB16].




Figure 5.12.: Low-mass dielectron yield in central Pb–Pb collisions, excluding the heavy-flavour and J/ψ contributions. Predictions of a coarse-grained transport model [EvHB16], using $p_T^e > 0.2$ GeV/ c and $|y_{ee}| < 0.85$, are compared to two sets of data. Contributions not included in the model are subtracted from the data, using the original N_{coll} -scaled (HFE R_{AA} -modified) heavy-flavour components for the blue (red) data points.

holistic description of the dielectron continuum than the thermal fireball model, one essential contribution is missing for a full prediction of the low-mass and intermediate-mass dielectron yield, which is the semileptonic decay of correlated open-charm as well as open-beauty mesons. In the given invariant-mass region, the also omitted J/ψ resonance is of secondary importance, but surpasses the QGP yield for $m_{ee} > 2.5$ GeV/ c^2 .

For a comparison between the model and the present results, these contributions (cp. Figure 5.9) are subtracted from the data in Figure 5.12. To account for the conceptual uncertainties related to the in-medium modifications of the open-heavy-flavour contributions, either of the extreme cases concluded from the discussion in Section 5.2 is used for subtraction. While a good overall agreement is seen, clearly more statistical precision is needed in the data for a conclusive comparison. Additional efforts are also needed to reduce the large uncertainties on the heavy-flavour cross sections (not propagated here) to achieve a reliable subtraction.




Figure 5.13.: Raw invariant-mass distributions of unlike-sign pairs (open squares), background (crosses) and subtracted signal (full circles) in 10–50% semi-central Pb–Pb collisions. A pair prefilter is used and final pair cuts are applied.

5.5. Semi-Central Collisions

In this section, selected results from the analysis of Pb–Pb collisions with 10–50% centrality are presented. Details on the analysis are described in Chapter 4, which focusses on the 0–10% centrality class, but includes some additional notes on semi-central collisions.

5.5.1. Spectra and Signal Quality

Figure 5.13 shows the invariant-mass distribution of unlike-sign pairs and acceptance-corrected background, as well as the raw signal after background subtraction. The only prominent structure in the subtracted signal is, as in central collisions, the peak of the J/ψ meson near $m_{ee} = 3$ GeV/ c^2 .

From these spectra, the signal-to-background ratio and the raw significance are computed as measures for the signal quality. They are presented in Figure 5.14. In comparison to central collisions, $N_{\text{S}}/N_{\text{B}}$ is larger by a factor of ≈ 2.5 , while the significance is similar on average.




Figure 5.14.: Signal quality in 10–50% semi-central Pb–Pb collisions. Left panel: signal-to-background ratio. Right panel: significance based on unnormalized signal and background spectra.

5.5.2. Track and Pair Efficiencies

The track selection efficiency for electron candidates is shown in Figure 5.15 for the 20 cut combinations used in this analysis. The transverse-momentum and pseudorapidity projections show no distinct features except for the quite large overall efficiency spread. The latter leads to a wide coverage of systematic uncertainties, not only from track cuts but also from signal extraction due to the resulting variation in pair statistics, especially in background size. The projection on the azimuthal angle is far from uniform, due to malfunctioning segments of the ITS, which however can be partially recovered by reducing certain track quality restrictions, as discussed in Section 4.2.1.

Pair efficiencies based on a factorization approach to obtain large statistics are presented in Figure 5.16. They are $\approx 20\%$ larger than in central collisions. This is caused by a slightly higher electron efficiency and because random rejection by the pair prefilter is less by a factor of two.




Figure 5.15.: Electron efficiencies for 20 different cut settings in semi-central collisions. Shown are projections of a 3-dimensional map, which is used to generate corresponding pair efficiencies.




Figure 5.16.: Pair efficiency in semi-central collisions. Left panel: pair efficiencies for the 20 cut settings in one transverse-momentum slice. Right panel: pair efficiency of the selected cut setting in different transverse-momentum slices. The prefilter using $m_{ee} < 40 \text{ MeV}/c^2$ & $\phi_V > \frac{3}{4}\pi$ is applied.




Figure 5.17.: Dielectron signals of the 20 cut settings in semi-central collisions merged into wide invariant-mass bins to extract systematic uncertainties. The average yield is shown with square markers and their vertical error bars correspond to the RMS of all deviations from this average. The signal selected for the final result is shown in orange with dashed error bars. At high mass, the signal of most cut settings is oversubtracted.

5.5.3. Systematic Uncertainty

The spread of efficiency-corrected dielectron spectra from the 20 cut settings is shown in Figure 5.17. For the estimate of the systematic uncertainty, wide invariant-mass bins are chosen to reduce statistical fluctuations in the spectra, which however remain sizeable for some cut settings.

The signal which is used to represent the final result corresponds to the same cut setting as in central collisions and matches well with the average yield of all 20 spectra. The RMS of all deviations from the average is used as the main systematic uncertainty on the result.




Figure 5.18.: Comparison of the dielectron signal measured in 10–50% semi-central Pb–Pb collisions to the hadronic cocktail. Contributions from light-flavour mesons are based on a charged-pion measurement and m_T -scaling. The open heavy-flavour is generated with PYTHIA and scaled by the number of binary collisions in Pb–Pb. The J/ψ is constructed from measurements in pp and Pb–Pb, including its R_{AA} . Experimental resolution and acceptance are included in the cocktail.

5.5.4. Cocktail Comparison

A comparison between the measured dielectron signal and the hadronic cocktail in semi-central Pb–Pb collisions is given in Figure 5.18. In the η -Dalitz region, the cocktail may leave some room for additional contributions, similarly to but not as pronounced as in central collisions (see Figure 5.1). At intermediate mass, the heavy-flavour contributions, scaled by $N_{\text{coll}}(10\text{--}50\%) = 492.2$ [ALICE13a], tend to overestimate the data. For $m_{ee} < 100$ MeV/ c^2 , a constant (flat) ϕ_V correction is applied to these data, since statistics in semi-central collisions and the corresponding Monte–Carlo production are insufficient for a $p_{T,ee}$ -differential correction using a ϕ_V template fit. This procedure is expected to underestimate the data for masses of $m_{ee} \approx 50$ MeV/ c^2 , as discussed in Section 4.5.2.

The pair-transverse-momentum spectra for individual invariant-mass regions are presented in Figure 5.19. Overall, they feature similar trends as in central collisions. In the low-mass

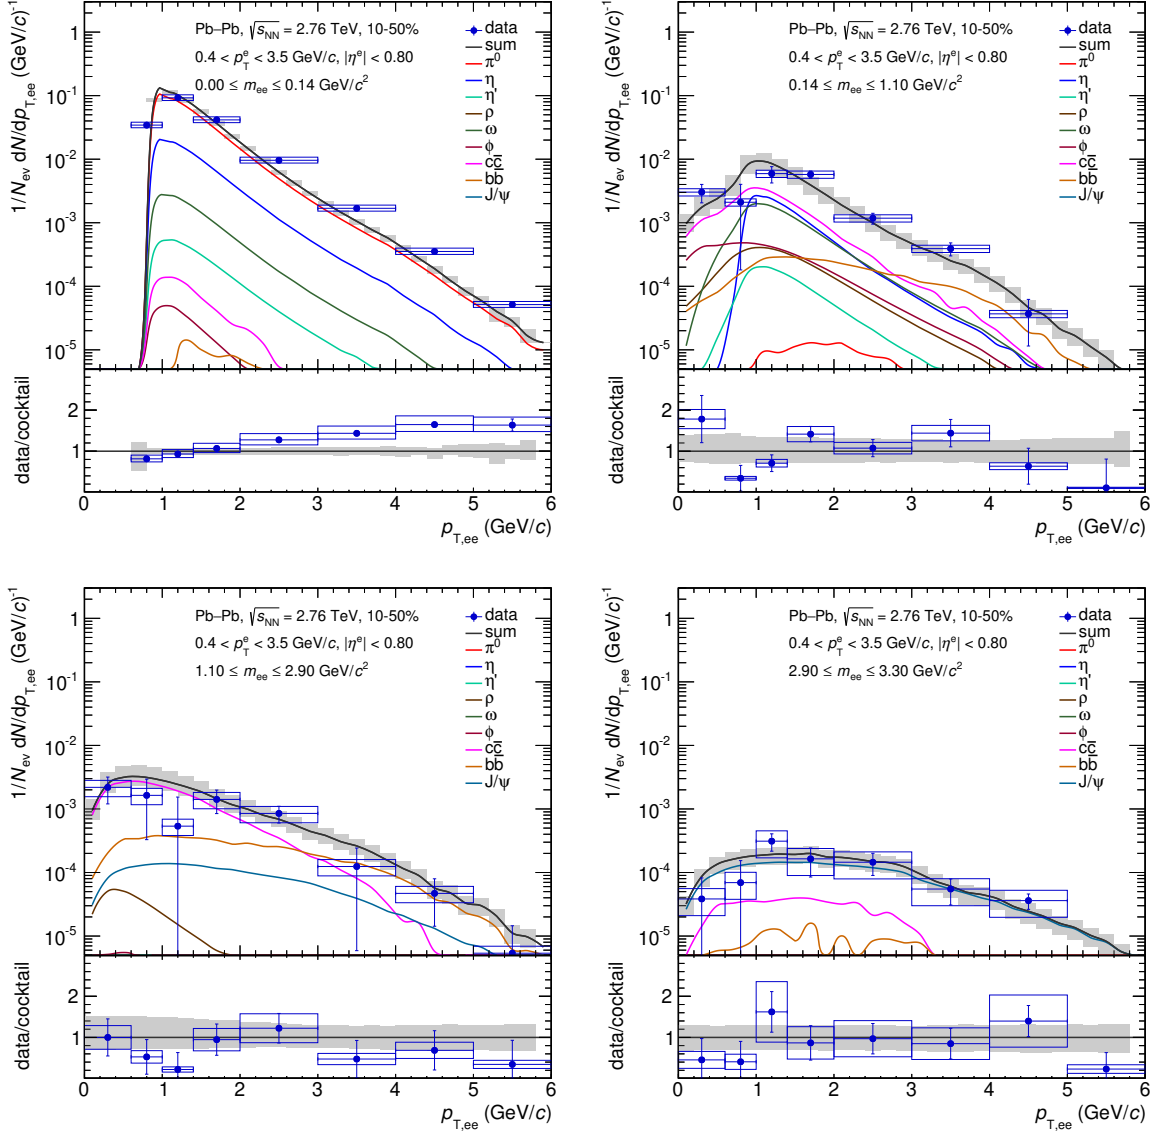


Figure 5.19.: Pair-transverse-momentum dependence of the comparison between data and cocktail from Figure 5.18, divided into four invariant-mass regions. See text for details.

region, the difference in slope between cocktail and data is more prominent, and not consistent within the uncertainty from cut variations and the ϕ_V correction. However, the ϕ_V distributions are less accurately described by Monte-Carlo in semi-central collisions, such that stricter conversion rejection cuts should be applied when focussing the analysis on very low mass.

The resonance region shows slightly less fluctuations than seen in central collisions, and a decent agreement between data and cocktail. The minimum signal-to-background ratio

is $N_S/N_B \approx 10^{-3}$ at $m_{ee} \approx p_{T,ee} \approx 1$ GeV.

In the intermediate-mass region, N_S/N_B is approximately identical for both centralities, resulting in less significance for the result shown here, due to smaller pair statistics. The data prefer a reduction of the heavy-flavour components with respect to N_{coll} -scaling, as already visible in the invariant-mass spectrum. This is expected considering the heavy-flavour electron R_{AA} , which is measured to be below unity also in semi-central collisions [ALICE18b].

The high-mass region features a better agreement between data and cocktail than for 0–10% centrality. However, the $p_{T,ee}$ dependence of the J/ψ -suppression in Pb–Pb collisions is not included in the cocktail, as discussed in Section 5.1.2.

No predictions by the models described in Section 5.4 are available for the 10–50% centrality class. Given the larger measurement uncertainty relative to central collisions, also no additional insights would be expected from a comparison. This situation will change with near-future dielectron measurements using the upgraded ALICE detector. Their expected potential is discussed in the following chapter.

Chapter 6.

Upgrade

The ALICE collaboration developed a substantial upgrade program to prepare the experiment for the high-luminosity data-taking during LHC RUN 3 [ALICE13d, ALICE13c, ALICE14e, ALICE15b, ALICE15a]. As a prerequisite to this, three physics performance studies were conducted to define the particular requirements for the upgraded detector setup as well as to verify its anticipated benefits for the benchmark measurements of the ALICE upgrade program: heavy-flavour production, quarkonia, and low-mass dileptons [ALICE12b].

The study on the last item was covered in the early phase of this work and is presented in this chapter¹.

6.1. Physics Motivation

The purpose of the upgrade is to improve and extend the capabilities of the ALICE experiment for its set goal, namely the study of the quark–gluon plasma. Owing to the confinement of strongly-interacting particles, many of the QGP properties are hidden from experimental observation, but amongst the accessible ones are its temperature and its transport coefficients [ALICE14e]. The upgrade aims at revealing these properties by facilitating a set of specific measurements with largely increased statistics and precision. The largest category are new differential measurements of heavy-flavour production in the charm and beauty sectors, which enable an investigation of the mass dependence of parton energy loss and thermalization in the QGP — together giving access to the heavy-flavour transport coefficients [ALICE12b]. At nominal LHC energy, the temperature in the early QGP phase may even be high enough to thermally produce $c\bar{c}$ pairs. A precise

¹The initial study is published in [ALICE12b] and was first presented in [ALICE12c], see Appendix C. Results of the first re-analysis are included in [ALICE14e].

measurement of the total charm production yield, which requires D meson reconstruction down to zero transverse momentum, will address this question and help to constrain the QGP temperature.

The measurement of the initial QGP temperature is indeed considered to be “one of the ‘holy grails’ of heavy-ion physics”, and the only unambiguous approach may be the measurement of thermal radiation in the dilepton channel at invariant masses above $\approx 1 \text{ GeV}/c^2$ [SS15]. The invariant-mass degree of freedom allows for the exclusive selection of radiation from the early, hot collision phase and it is unaffected by blue shifts due to radial expansion of the system, thereby minimizing the model dependence in the interpretation of experimental results.

The anticipated dielectron measurement also contains a unique probe of hot QCD at low invariant mass, namely the electromagnetic spectral function of light vector mesons, which is modified in the medium (details on thermal dilepton radiation at low and intermediate mass are described in Section 2.1.2). The behaviour of this spectral function may represent the only experimental observable related to the question of chiral symmetry restoration. This symmetry of QCD is fulfilled in the QGP but broken in cold nuclear matter, and it should become restored near the phase transition between QGP and hadron gas [ALICE12b]².

The following study is designed around the quest of measuring the initial QGP temperature, since it is the more tangible topic of the two, with less difficulty concerning the requirement of its quantification. The actual experimental goal covers both topics and can be summarized as to achieve a high-precision dielectron measurement, with results comparable to the NA60 dimuon excess spectra, in heavy-ion collisions at the present energy frontier. This will provide invaluable new insight to constrain and refine the theoretical models, which are all in reasonable agreement with the currently available experimental data, as described in Section 2.2, and improve our understanding of the fundamental properties of matter.

6.2. Experimental Improvements

The ALICE upgrade program includes major changes for the two detector systems that are most crucial to the low-mass dielectron analysis: a completely redesigned Inner Tracking System is installed around an also renewed beampipe section with reduced diameter [ALICE14e], and the Time Projection Chamber receives new readout chambers using Gas Electron Multiplier (GEM) technology that allows for continuous readout

²A review on chiral symmetry and the connection to dilepton observables can be found in [RWvH09].

[ALICE13c]. These changes to the detectors are complemented by upgrades of the Readout & Trigger System³ [ALICE13d] as well as the Online-Offline Computing System [ALICE15b]. All of these improvements are mandatory in order to record and process the foreseen event rate of 50 kHz for minimum-bias Pb–Pb interactions during LHC RUN 3. The main purpose of the TPC upgrade is to overcome the rate limitations of the original readout design by replacing the Multi-Wire Proportional Chambers (MWPC) with a GEM system, while preserving the TPC’s particle identification capabilities and the combined momentum resolution with other central-barrel detectors [ALICE13c]. GEM chambers amplify a primary charge deposition by accelerating the electrons through small holes in stacked GEM foils — which are oppositely charged on either surface — causing secondary ionization on their path to the anode pad readout. A large fraction of the produced ions are collected on the outer surface of each foil and thus do not reach the drift volume of the TPC, where they would cause a space-charge. Stacks of four GEM foils are used in the new readout chambers to meet the requirements of less than 1 % ion backflow at a gas gain of 2000 [ALICE13c]. This removes the need for a gating grid, which must be closed after each event readout in case of MWPCs to collect the ions. In combination with new front-end electronics, the GEM system therefore allows for continuous, untriggered readout. Since both the precision of global tracking and the PID capability shall remain unchanged after the upgrade, only the increased data-taking rate is explicitly included in the following performance study.

The new ITS consists of seven cylindrical layers of silicon pixel detectors compared to the six layers of the original⁴ ITS. The additional layer is part of its inner barrel, which thus provides three trajectory points close to the interaction point. Also the radii of the inner detector layers are substantially reduced, starting from 22 mm instead of 39 mm. These changes improve the pointing resolution, which represents the precision of extrapolating a trajectory to the event vertex, by a factor of ≈ 3 [ALICE12b]. This improvement is equally crucial for the low-mass dielectron, heavy-flavour and quarkonia analyses, as it allows for better separation of particles from prompt, charm and beauty sources. To reduce the amount of multiple scattering that affects momentum resolution and the background from photon conversions, the material budget of the new ITS is targeted at only 0.3 % of a radiation length X_0 per layer of the inner barrel and 0.8 % X_0 per outer layer, compared to 1.1 % X_0 per layer of the current ITS [ALICE14e].

³This includes a new readout for the TOF detector and a new combined interaction trigger and centrality detector, both needed for the dielectron analysis and implicitly assumed in the performance study presented below.

⁴In the remainder of this chapter, the term “current ITS” is used to refer to the original ITS, which was used during LHC RUN 1 and RUN 2.

Initially, two different designs were considered for the four outer layers of the new ITS, either with or without particle identification capabilities. Without PID, low- p_T tracking is slightly better by utilizing pixel sensors instead of strips and the omission of analogous readout reduces design complexity. An intermediate solution using pixels with analogous readout preserves the optimal tracking but leads to a reduced PID capability compared to the current ITS [ALICE12b]. Since the low-mass dielectron analysis is amongst the two most sensitive ones with respect to the availability of low-momentum-PID information⁵, the following study and a subsequent iteration, see [Kle14], also addressed this topic and helped to converge on the ITS design without PID capabilities.

Running Conditions

Besides the above improvements to detectors, their readout, and software, the low-mass dielectron measurement in RUN 3 also requires a reduction of the magnetic field in the central barrel from the nominal $B = 0.5$ T down to 0.2 T. The corresponding reduction of track curvatures provides two major benefits to the analysis: Primarily, it enables tracks with $p_T \gtrsim 150$ MeV/ c to reach the TOF detector, allowing for the use of the conventional TPC+TOF electron-PID strategy over the entire kinematic range down to the typically chosen limit at $p_T = 200$ MeV/ c , below which charged pions become indistinguishable from electrons in both detectors⁶. In addition, more of the very soft particles, with transverse momenta down to $p_T \approx 60$ MeV/ c , can be tracked by the ITS and used to reject photon conversions and π^0 Dalitz decays in order to improve the signal-to-background ratio of the measurement. In the meantime, several periods of low-field data-taking during proton–proton beamtimes were realized and a diligent low-mass dielectron analysis performed [Jun19], together successfully establishing some of the operational and analytical foundations for this upcoming measurement.

6.3. Physics Performance Study

The main goal of this study is to prove the feasibility and quantify the precision of measuring the QGP temperature of the early phase of heavy-ion collisions with low-mass dielectrons in the upgraded ALICE experiment. This requires in essence:

⁵This refers to the foreseen collection of RUN 3 analyses, where the other one profiting (slightly) from low- p_T PID is the $\Lambda_c \rightarrow pK\pi$ reconstruction as part of the heavy-flavour measurements.

⁶In fact, a recent dielectron analysis on low-field proton–proton data shows that the TPC is capable of identifying electrons again below the pion-crossing region, thus extending the low- p_T reach down to $p_T = 75$ MeV/ c [Jun19]. If a similar PID performance can be achieved in Pb–Pb collisions under RUN 3 conditions, this may allow for additional studies of the low-pair- p_T excess seen in that analysis.

- A full dielectron analysis based on simulated Pb–Pb data for the experimental and physical conditions of LHC RUN 3.
- Predictions for the hadronic cocktail with light- and heavy-flavour contributions and for the thermal radiation components which shall be measured.
- Estimates on uncertainties for all signal contributions that need to be subtracted to reveal the thermal radiation.

Since no such data set exists, it is successively constructed from individual inputs, governing much of this studies methodology. All ingredients and the workflow of the study are summarized in the diagram of Figure 6.1. The diagram is structured into three columns, where the central one gives an overview of the study process itself, while the side columns list additional inputs. Vertically the diagram can be divided into four sections, corresponding to the three main steps of a dielectron analysis plus a procedure to model realistic spectra based on the given ingredients. These components of the study are individually covered in the following:

Section 6.3.1 addresses the creation of electron candidate and prefilter electron samples as they would be measured by the upgraded detector setup. In order to optimally benefit from the improved performance, a reduced magnetic field of the ALICE solenoid is foreseen, which extends the acceptance and PID capabilities down to lower track momenta. A full Monte–Carlo production is used to have access to primary as well as secondary electrons and hadron contamination. Since the detector implementation corresponds to the RUN 1 setup of ALICE, particle samples are modified to match the expected tracking efficiencies and conversion probabilities of the RUN 3 setup.

Section 6.3.2 describes the pair analysis. Initially, the predicted charged-particle multiplicity density of 5.5 TeV Pb–Pb collisions is reproduced by superimposing events of the available production. The subsequent pairing steps correspond to a real analysis, which delivers the signal quality estimators S/B and significance.

In Section 6.3.3 complete dielectron spectra are created, based on model predictions and the expected number of events to be recorded. The input of the previous step is used to define statistical and systematic uncertainties of the data points.

Section 6.3.4 contains the final analysis on the spectra to evaluate at which precision the thermal radiation input can be re-extracted from the modelled data.

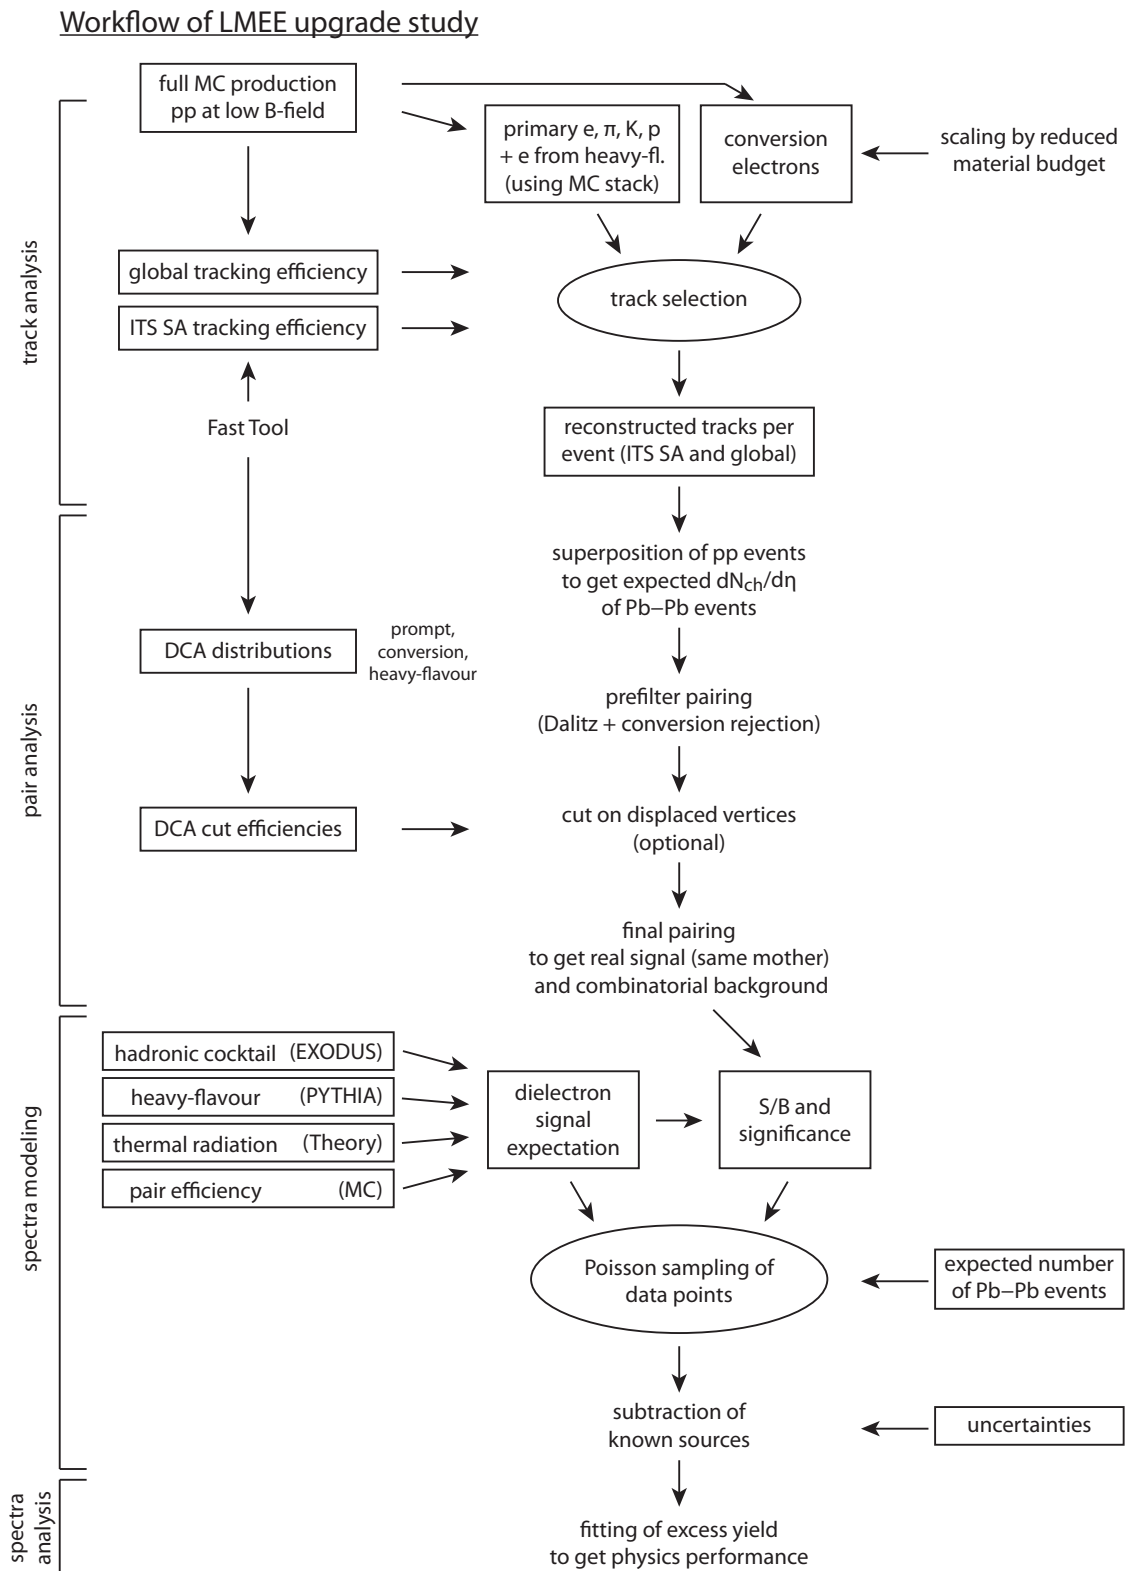


Figure 6.1.: Workflow diagram of the physics performance study. See text for details.

6.3.1. Track Analysis

The purpose of this first analysis step is to create realistic track samples of electron candidates and prefilter electrons as input for the pair analysis. This sample mainly consists of primary electrons from hadron decays (the light- and heavy-flavour cocktail⁷), secondary electrons from photon conversions, as well as contamination from misidentified hadrons. In real data it also contains electrons from the “unknown” sources, which we aim to measure, but these are a minor contribution and therefore neglected at this stage. The sample is constructed from several inputs, based on different tools and analyses, to cover all particular aspects of the anticipated running conditions. In short: based on tracking efficiency estimates, random number generation and further constraints, each relevant simulated particle gets a chance of entering one of the track samples.

Efficiencies

The ITS stand-alone (SA) tracking efficiency, which is listed in Figure 6.1 as the third input to this analysis step, is provided by a custom application called “Fast Tool” [Mus12, Sha12]. It contains the present and planned geometries of ITS and beampipe and propagates a given particle up to the TPC. It then performs a realistic reconstruction in the environment of a Pb–Pb event, using a Kalman filter. The resulting ITS SA efficiencies for the reduced magnetic field of $B = 0.2$ T are presented in the left panel of Figure 6.2. Both variants of the new ITS maintain a reasonable efficiency down to $p_T \approx 70$ MeV/ c . This value is chosen as the low- p_T cut for prefilter electrons in this analysis, since towards lower momenta the predictions of the Fast Tool are expected to become less accurate.

The reconstruction efficiency of electron candidates, which we refer to as *global tracking efficiency*, is known from [Hin12] for the RUN 1 detector setup at nominal and reduced magnetic field. Since the global tracking efficiency combines the ITS, TPC and TOF detectors, it needs to be modified in the present study to account for the improved performance of the upgraded ITS. In a simplified approach, the global efficiency is divided by the current ITS efficiency, then multiplied by the new ITS efficiency. All global tracking efficiencies are collected in the right panel of Figure 6.2. Most of the visible acceptance gain is achieved by the magnetic field reduction, while the new ITS version provides an efficiency increase of up to 50% near $p_T = 200$ MeV/ c , which is the low- p_T cut for electron candidates in this analysis. The p_T dependence of the global

⁷In the discussion below, electrons from heavy-flavours are distinguished from the remaining primary electrons by the term “charm electrons”.




Figure 6.2.: Left panel: ITS stand-alone tracking efficiency at $B = 0.2$ T for the current and two variants of the upgraded ITS. Right panel: combined reconstruction and PID efficiency of global electrons for different ITS versions and magnetic field strengths. Electron PID uses TPC and TOF information.

efficiency is governed by the TOF turn-on curve at low- p_T and the pion exclusion cut in the TPC at high- p_T .

Particle Identification

For electron identification, the classical PID scheme involving the TPC and TOF detectors is used. This choice is based on the two initial design options of the new ITS, which foresee either no or a potentially reduced PID capability, such that TOF is the only available hadron veto in the relevant p_T range. In addition, the requirement of operating at the reduced magnetic field allows for the use of TOF information down to $p_T \approx 150$ MeV/c, providing unified PID for all electron candidates.

Figure 6.3 shows the TPC signal for the Monte-Carlo track sample remaining after a PID cut of $|\sigma_e^{\text{TOF}}| < 3$ in the TOF. Electrons are selected if they fulfill $|\sigma_e^{\text{TPC}}| < 3$, plus exceed a parameterized dE/dx -threshold to exclude potential hadron contamination towards higher momenta, as indicated by the lines. Given this nearly perfect electron purity, hadrons are neglected for the global track sample to simplify subsequent analysis steps. No lower p_T -threshold is applied to the tracks in this figure, while the final sample of global tracks is restricted to $p_T^e > 200$ GeV/c.




Figure 6.3.: Specific energy loss in the TPC after electron selection by TOF, based on a subset of the low-field Monte–Carlo production. The lines indicate the final electron selection in the TPC.

Primary Electrons and Hadron Contamination

Besides providing the global tracking efficiency, the full Monte–Carlo production created in [Hin12] is also chosen as the source of particles for this study. Some assumptions have to be made, because it is a simulation of pp-events at $\sqrt{s} = 7$ TeV and uses detector representations corresponding to RUN 1 data taking, but it provides in particular a realistic sample of conversion electrons for the low-field conditions (see below).

Since the upgraded ITS provides stand-alone tracks far below the p_T -threshold of the current ITS, the reconstructed ITS SA tracks of this production are not sufficient as a prefilter electron sample. Therefore the generated particles of the Monte–Carlo stack are used instead, and declared as measured or lost on a random basis according to the ITS SA efficiency. Only true electrons are selected in this process. This is a strong assumption, since in practice all charged particles will make it into the prefilter electron sample, if no ITS PID is available⁸.

For global tracks, i.e. the electron candidate sample, a reweighting of the existing reconstructed tracks would be possible to match the improved efficiency, since their kinematic range does not increase from the current to the new ITS. However, the good

⁸The simplification was necessary because of time constraints in the initial study. Subsequent re-analyses including hadrons (this work and [Kle14], for [ALICE14e]) show that the increased prefilter track sample can be compensated with stricter prefilter pair cuts (see Section 6.3.2).

resolution of global track reconstruction and high electron purity of the chosen PID scheme also allow for using generated electrons instead, equivalent to the process for prefilter electrons. The consistency between both approaches is ensured by adjusting the selection of stack particles, such that after efficiency application they match with the reconstructed sample. If a particle does not enter the global track sample owing to its random number pull, it still gets a chance to enter the prefilter electron sample.

Conversion Electrons

Conversion electrons contribute to the combinatorial background of this analysis, if they pass the global track quality cuts and cannot be rejected otherwise. The first is true if the photon conversion occurs in the beampipe or in the innermost ITS layer⁹. To match the reduced material budget after the upgrade, the total number of photon conversions is scaled by the fraction

$$\frac{1 - e^{-7/9 \cdot (X/X_0)_{\text{upgrade}}}}{1 - e^{-7/9 \cdot (X/X_0)_{\text{current}}}} = 0.54, \quad (6.1)$$

where $(X/X_0)_i$ is the combined number of radiation lengths of beampipe and first ITS layer for the current or upgraded experimental setup [Kwe12]. The corresponding electron-positron pair is then either consistently excluded from the track sample, or based on the tracking efficiencies it is decided for each leg if it enters one of the track samples.

6.3.2. Pair Analysis

The pair analysis performed here is very similar to the actual data analysis presented in Section 4.3. In this case the goal is to simulate a realistic combinatorial background in order to extract the signal-to-background ratio and significance for the RUN 3 conditions. This requires event-by-event track samples of the right magnitude, which is achieved by superimposing n pp-events of the previous step to reach the expected multiplicity. n is computed via

$$n = \frac{dN_{\text{ch}}/d\eta|_{\text{Pb-Pb}}}{dN_{\text{ch}}/d\eta|_{\text{pp,INEL}}}, \quad (6.2)$$

⁹Technically, conversions are selected from the Monte-Carlo stack up to the radial distance from the vertex, at which the corresponding amount of conversion electrons equals the one in the reconstructed event. The radial conversion positions are kept also for simulating the upgraded scenario. To account for the new geometry, modified DCA distributions are used instead (see Section 6.3.2).

where $dN_{\text{ch}}/d\eta|_{\text{pp,INEL}} = 4.91$ is used for the given 7 TeV pp production. The values estimated for Pb–Pb at 5.5 TeV are summarized in the following table, in comparison to measured values for $\sqrt{s_{\text{NN}}} = 2.76$ TeV. This leads to $n = 356$ and $n = 51$ to obtain central and semi-peripheral (40–60 %) Pb–Pb collisions, respectively.

Centrality	$dN_{\text{ch}}/d\eta$	
	2.76 TeV	5.5 TeV
0–10 %	1448	1750
40–60 %	205	248

Table 6.1.: Charged-particle multiplicity density $dN_{\text{ch}}/d\eta$ of Pb–Pb collisions for two centrality classes and two center-of-mass energies. Values for $\sqrt{s_{\text{NN}}} = 2.76$ TeV are from [ALICE11a], the ones for $\sqrt{s_{\text{NN}}} = 5.5$ TeV are our estimates for RUN 3.

Pair Prefiltering

Upon obtaining a realistic track density, the prefilter pairing is performed on the combined track sample of prefilter electrons and electron candidates. A combined invariant-mass and opening-angle cut is applied to all unlike-sign and like-sign pairs. The cut values are presented in Table 6.2. As noted before, in the initial study only true electrons are considered as rejection partners, allowing for slightly looser cuts than in the RUN 1 data analysis. With these cuts the re-analysis assuming no ITS PID — therefore adding all low- p_{T} hadrons ($p_{\text{T}} < 200$ MeV/ c) to the prefilter electron sample — shows too strong *random rejection*, such that the signal is lost completely. However, the initially predicted performance can largely be restored by tightening the prefilter cut as specified.

	mass exclusion	angular exclusion
prefilter (initial study)	$m_{ee} < 50$ MeV/ c^2	$\theta_{ee} < 100$ mrad
prefilter (re-analysis)	$m_{ee} < 10$ MeV/ c^2	$\theta_{ee} < 50$ mrad

Table 6.2.: Prefilter exclusion cuts of the initial and secondary study (for [ALICE12b] and [ALICE14e], respectively). Pairs are rejected, if both requirements are fulfilled.

Figure 6.4 visualizes the effect of the prefilter for both cases. It shows the ratio of accepted pairs after the prefilter over all reconstructed pairs as function of their invariant mass. For the true Monte–Carlo signal only light-flavour resonances are available, thus limiting the mass range. On average 20 % of these pairs are rejected by the prefilter of




Figure 6.4.: Ratio of pairs accepted by the prefilter over all reconstructed pairs, for Monte-Carlo signal and background pairs. Left panel: initial study using only true electrons. Right panel: re-analysis including low- p_T hadrons. The double-ratios (not shown) are consistent within uncertainties, meaning that S/B can be conserved when hadron contamination is included.

the initial study, while combinatorial-background pairs are suppressed by $\approx 70\%$ in the corresponding mass range. The hadron contamination in the re-analysis without ITS PID information leads to 60% of random signal rejection, despite the refined prefilter cuts. At low mass, the background is also suppressed more strongly, dropping from 30% to 15%. Thereby the double-ratio, which is proportional to S/B , remains constant. Consequently, only the significance of the measurement is reduced according to the additional signal loss when removing the assumption of “perfect” ITS PID in favour of no ITS PID¹⁰. A change in mass dependence is clearly visible in the background, but is present in the signal as well. This is caused by the smaller prefilter cut window in the re-analysis, which affects less of the pairs at higher mass.

DCA Distributions and Displaced-Vertex Cut

The improved pointing resolution of the new ITS allows for further depletion of pairs from undesired sources in the dielectron sample by constraining the *distance of closest approach* (DCA) between the reconstructed particle trajectory and the event vertex.

¹⁰The difference between the two proposed variants of the new ITS, either with (realistic) PID or without PID but better tracking, is studied in [Kle14]. With PID, both signal-to-background ratio and significance are 10 – 20% higher than without PID.




Figure 6.5.: Distributions of track DCA in radial direction for electrons from prompt sources (left), photon conversions (center), and open-charm mesons (right). The new ITS provides an improved separation power between prompt and other sources (see text for details).

This observable is briefly called *track DCA* and only its radial (or xy -) component DCA_r is considered in this analysis. Measuring a value larger than the pointing resolution of the detector suggests, that the particle originates from a displaced secondary vertex, which is the case for photon conversions and heavy-flavour decays.

The DCA_r distributions of primary, conversion, and charm electrons are presented in Figure 6.5. They are generated by the Fast Tool, using as input track samples from Section 6.3.1 for self-consistent track kinematics [Sha12]. DCA_r distributions are compared between the current and new ITS and therefore normalized to the corresponding detector resolution¹¹. The new ITS features less outliers beyond $\pm 3\sigma$ for primary and conversion electrons, i. e. a more robust track reconstruction. The distribution of conversion electrons is shifted because of the fixed radial position of the beampipe¹². This offset has the same magnitude for both ITS versions by coincidence: the smaller radius of the new beampipe is just compensated by the improved pointing resolution of the new ITS. The distribution of charm electrons is considerably wider owing to the natural spread of the heavy-flavour-meson decay positions.

To enhance the relative amount of electrons from prompt sources, a $|DCA_r| < n\sigma$ inclusion cut can be applied, which rejects a larger fraction of the other sources. The DCA cut efficiency, which is the fraction of electrons surviving the cut, is given in the left

¹¹In hindsight, the normalization within the Fast Tool seems to be off, since the gaussian part of the primary-electron distributions should be identical by definition, and consequently the charm distribution wider for the new ITS. This inconsistency is removed by constructing relative quantities below.

¹²Also, only electrons and no positrons are used here, which would otherwise add a mirrored distribution, causing an artificial widening for the combined track sample.




Figure 6.6.: Left panel: DCA cut efficiency for electrons from different sources as a function of the fraction of primary electrons surviving the cut. Both ITS versions allow for mild suppression of conversion electrons, while electrons from open-charm mesons can be strongly suppressed, especially with the new ITS. Right panel: ratio of squared cut efficiencies, representing the obtainable enhancement of prompt dielectron sources relative to the charm contribution.

panel of Figure 6.6 as a function of the percentage of remaining primary electrons. Primary electrons themselves define the diagonal. Conversion electrons are slightly suppressed because of their offset DCA_r distribution, but no significant benefit is achieved by the new ITS because of the reduced beampipe radius. Charm electrons show a strong relative suppression for any given primary-electron cut efficiency and a clear improvement with the new ITS. A quite aggressive cut at $\epsilon_{\text{DCA,prim}} = 32\%$ is chosen, leading to $\epsilon_{\text{DCA,conv}} = 25\%$ and $\epsilon_{\text{DCA,charm}} = 14\%$. Since the track DCA cut acts quadratically on the final pair statistics, this sacrifices 90% of the prompt signal, while the charm contribution is suppressed down to 2%.

The achievable relative pair suppression as a function of the cut value, given by the ratio of the squared DCA cut efficiencies of primary and charm particles, is presented in the right panel of Figure 6.6. The ratio increases when tightening the cut and saturates near a value of three for the current ITS and ≈ 5.5 for the new ITS, modulo some fluctuations. With the cut at 32%, the charm contribution can be suppressed by a factor of five. This allows for a corresponding enhancement of the relative yield of thermal radiation in the intermediate-mass region, which will be presented at the end of Section 6.3.3.




Figure 6.7.: Result of the final pairing, when no DCA cut is used. Only unlike-sign pairs are considered and separated into signal and background based on Monte–Carlo–truth information. The total dielectron signal expected for LHC RUN 3 conditions (see Section 6.3.3), multiplied by its pair efficiency to represent the raw yield, is given as dashed line.

Final Pairing

After prefiltering and an optional DCA cut, the final pairing is performed on the remaining global tracks. Since Monte–Carlo–truth information exists on which tracks belong to a pair, this pairing step is done only on unlike-sign pairs. This avoids the need for an acceptance correction between unlike-sign and like-sign, which is done in real data via event mixing (see Section 4.1.2).

The result for 0–10% centrality is given in Figure 6.7. It shows the combinatorial background as histogram and the Monte–Carlo signal from light-flavour resonances as data points, each for the cases of using the current or new ITS version. With the new ITS, the background is reduced by a factor of two over the full invariant-mass range, while approximately 10–30% of the signal are lost because of a higher probability for random rejection in the prefilter.




Figure 6.8.: Signal-to-background ratio (left) and normalized significance (right) for the current and the new ITS. No DCA cut is used here.

Signal-to-Background Ratio and Normalized Significance

From the final pair statistics, the signal quality quantifiers S/B and significance can be computed. However, since the signal extracted from the Monte–Carlo production only contains light-flavour resonances, these pairs cannot serve as the sole signal input for the calculation. Instead, the expected dielectron signal for LHC RUN 3 conditions is used, whose composition is described in Section 6.3.3. For comparison, its sum is added in Figure 6.7 as dashed line. Its (light-flavour) cocktail component is largely consistent to the Monte–Carlo signal pairs. From this modelled signal and the background of each ITS version, S/B and significance are computed.

Both quantities are presented in Figure 6.8. The *normalized* version of significance, as defined in Section 4.3.2, is given here. It is favourable in this case, since the number of Pb–Pb events and the binning of the invariant-mass spectrum is not fixed at this point. The signal-to-background ratio is improved by a factor of two with the new ITS, as expected from the background reduction. The significance correspondingly rises by $\approx 70\%$.

S/B is used in the following to quantify the systematic uncertainty on the final dielectron spectrum due to background subtraction. The normalized significance serves as the main input for generating a realistic set of data points around the nominal signal expectation on a statistical basis, which we refer to as *sampling*. This process is described in the next section.

6.3.3. Dielectron Signal and Poisson Sampling

Signal Composition

The dielectron signal considered for this study contains the following elements:

- A light-flavour hadronic cocktail, similar to that of the RUN 1 data analysis, generated via EXODUS [Ave12]. It is based on a π^0 measurement in Pb–Pb at $\sqrt{s_{\text{NN}}} = 2.76$ TeV and m_{T} -scaling for the other resonances, then scaled to the expected $dN_{\text{ch}}/d\eta = 1750$ (248) in central (semi-peripheral) collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV. The ρ meson contribution is excluded, as it is provided by the model prediction below. The cocktail is folded by the estimated momentum resolution of the ALICE TPC at a reduced magnetic field of $B = 0.2$ T [Iva12].
- A heavy-flavour cocktail, for which only open charm is considered¹³. It is produced from PYTHIA and scaled by $N_{\text{coll}} = 1625$ (140) for central (semi-peripheral) collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV [Toi12, App12]. This binary-collision scaling is a conservative (or pessimistic) assumption for the purpose of this study, given that the measured heavy-flavour R_{AA} is significantly below unity, as discussed in Section 5.2.1.
- Predictions for thermal dielectron yields in Pb–Pb collisions of both centralities. They are based on an expanding-thermal-fireball model [Rap01, vHR08], which uses the in-medium spectral function of the ρ meson as main contributor to radiation from the hadronic phase, and lattice-QCD results as input for QGP radiation. Further aspects of the model are described in Section 5.4.1. A presentation of this particular model version including a brief discussion on the results of this study can be found in [Rap13].

Figure 6.9 shows the invariant-mass distribution of this signal composition for central Pb–Pb collisions. Acceptance cuts of $|y_e| < 0.84$ and $p_{\text{T}}^e > 0.2$ GeV/ c are applied on individual electrons and a cut of $p_{\text{T},ee} < 3$ GeV/ c on the pairs.

While it is hoped that this dielectron signal is a reasonable estimate, it does not aim to be a complete and precise prediction for the actual yield in future Pb–Pb collisions. Diligence is mainly attributed to the relative sizes of the contributions, in particular between the model and the cocktails. This is verified by comparing the ρ meson yield between cocktail and model. In the latter it is estimated by computing one generation

¹³Beauty is the subdominant component in the intermediate-mass region and determining its contribution to the dielectron cocktail is conceptionally and computationally extensive because of the various decay channels and B meson oscillation. It is therefore neglected here.




Figure 6.9.: Estimated dielectron signal in central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV, composed of light- and heavy-flavour cocktails from generators and a model prediction for thermal radiation (see text for details). No DCA cut is used, the intermediate-mass region is dominated by open-charm-meson decays.

of radiation around the temperature of chemical freeze-out, $T_{\text{chem}} = 160$ MeV [Rap12]. Consistency is achieved by upscaling the ρ meson yield of the model by 35%, and consequently its in-medium spectral function and the QGP radiation as well.

Sampling of Signal Spectra

The next step in the workflow is the creation of hypothetical data points with realistic fluctuations. This is the key element of the upgrade study as it combines all previously established inputs.

To simulate a data point for a given invariant-mass interval $\Delta m_{ee;i} =: x$, the number of hypothetically measured signal pairs $N_{\text{meas}}(x)$ is sampled from a Poisson distribution¹⁴. The distribution has a mean equal to the expected number of signal pairs in this mass

¹⁴We therefore refer to this procedure as *Poisson sampling*, while the more accurate term for it is “executing a Poisson point process”.

interval, $N_{\text{exp}}(x)$, and the probability to measure N pairs is given by:

$$P(N_{\text{meas}}(x) = N) = \frac{e^{-N_{\text{exp}}(x)} \cdot (N_{\text{exp}}(x))^N}{N!}. \quad (6.3)$$

The value of $N_{\text{exp}}(x)$ is determined by using the planned number of recorded Pb–Pb events in combination with the normalized significance extracted from the pair analysis. As shown in Section 4.3.2, significance is equal to the inverse relative statistical uncertainty of the signal¹⁵. The standard deviation of a sufficiently large signal in the absence of background is $\sigma_S = \sqrt{N_S}$, so the significance simplifies to:

$$\begin{aligned} \text{significance} &= \frac{N_S}{\sigma_S} \\ &= \sqrt{N_S} \\ &= \sqrt{N_{\text{exp}}}. \end{aligned}$$

Therefore it is generally true — and may serve as a descriptive relation — that significance corresponds to the square root of the number of background-free counts.

This relation is used here to find the mean of the Poisson distribution, N_{exp} , for the sampling process. Since this requires raw counts, the *normalized* significance provided by the previous analysis step (Figure 6.8, right panel), is supplemented by additional inputs: a bin width of $\Delta m_{ee} = 25 \text{ MeV}/c^2$ is chosen for the sampled spectra and the number of Pb–Pb events is set to $N_{\text{events}} = 2.5 \cdot 10^7$ per 10% centrality interval¹⁶. The Poisson mean is calculated as

$$N_{\text{exp}}(x) = (\text{significance}(x))^2 \cdot N_{\text{events}} \cdot \Delta m_{ee}. \quad (6.4)$$

Values for the normalized significance found in this study are on the order of $10^{-2} (\text{GeV}/c^2)^{-1/2}$ in the relevant invariant-mass range. This means that in a large mass bin of $\Delta m_{ee} = 1 \text{ GeV}/c^2$, on average one signal pair (after background subtraction) is measured per 10^4 events. Consequently, for the bin width chosen here, $4 \cdot 10^5$ events are required for one signal pair per bin.

For each mass bin, a number $N_{\text{meas}}(x)$ is drawn from the Poisson distribution and its relative statistical uncertainty set to $\sqrt{N_{\text{meas}}(x)}/N_{\text{meas}}(x)$. The data points are then

¹⁵Modulo acceptance differences between unlike-sign and like-sign pairs, which are not relevant in this study since only unlike-sign pairs are used, in combination with Monte–Carlo information.

¹⁶The latter is similar to the RUN 1 statistics, as it corresponds to four weeks of heavy-ion beamtime, given the rate limitations of the current TPC readout.




Figure 6.10.: Expected precision for a measurement of the dielectron signal in central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV, based on $2.5 \cdot 10^7$ events. The signal is multiplied by the pair efficiency to represent only the measured raw yield. No DCA cut is used.

scaled to match the normalization of the final signal yield (“sum”) via

$$\frac{dN_{\text{meas}}(x)}{dm_{ee}dy} = \frac{N_{\text{meas}}(x)}{N_{\text{exp}}(x)} \cdot \frac{dN_{\text{sum}}(x)}{dm_{ee}dy}. \quad (6.5)$$

Figure 6.10 shows the resulting sampled spectrum in comparison to the dielectron signal. All signal components are multiplied by the pair efficiency, which is obtained from Monte–Carlo [Ave12], to resemble only the measured fraction of the yield. Since a constant bin width is chosen, the fluctuation of the data points around the signal sum is entirely governed by the mass dependence of the significance.

Systematic Uncertainties

Estimates on different systematic uncertainties and their propagation are relevant in the further steps of this study.

Already present in the signal spectrum (Figure 6.10) is the one due to subtraction of

the combinatorial background. The background uncertainty is set to $\Delta_{\text{bkg}} = \pm 0.25\%$ ¹⁷. This number is put into perspective by considering the R -factor, which quantifies the systematic difference between the typically used background estimate from like-sign pairs and the true background in unlike-sign pairs. Any uncertainty on the R -factor thereby contributes to the background uncertainty. In the RUN 1 data analysis, two independent measurements of the R -factor are done, since the data set is separated according to different settings of the magnetic fields. The two results (see Figure 4.2 in Section 4.1.2) are identical within $\pm 0.1\%$ for $m_{ee} > 250 \text{ MeV}/c^2$, except for the bins near the J/ψ peak, which fluctuate because of limited statistics. Note that some deviation, especially at low invariant mass, is expected from the different field settings. Therefore it can only be understood as an upper limit on the R -factor's systematic uncertainty. The uncertainty is propagated to the signal by weighting it with the inverse of the signal-to-background ratio (Figure 6.8, left panel) and shown as green band in Figure 6.10.

Systematic uncertainties on the light- and heavy-flavour cocktails become relevant in this study starting from the point where these contributions are subtracted from the signal. For the light-flavour cocktail, an overall systematic uncertainty of $\Delta_{\text{cocktail}} = \pm 10\%$ is chosen. This precision may not be reached for all particles, but the η meson, which is most crucial to study a low-mass enhancement, has now been measured in Pb–Pb down to $p_{\text{T}} = 1 \text{ GeV}$ [ALICE18c]. The employed photon conversion method may benefit from the low-field setting in RUN 3 to even further reduce this threshold for improved yield estimates. Regarding the open-charm contribution, an uncertainty of $\Delta_{c\bar{c}} = \pm 20\%$ is assumed to be feasible. This is a challenging goal, given the large current uncertainty on the order of $\Delta_{c\bar{c}} = \begin{smallmatrix} +60\% \\ -30\% \end{smallmatrix}$. But the estimate is supported by the fact that precise heavy-flavour measurements exploiting the improved DCA resolution of the new ITS are also amongst the main objectives of the upgrade program.

Excess Spectra

To evaluate the precision at which the thermal radiation components can be extracted from the dielectron signal, all known sources are subtracted from the sampled yield. The remaining thermal excess yield is presented in Figure 6.11. Systematic uncertainties on the charm and cocktail contributions are propagated and shown together as magenta error bars. This is the dominant uncertainty for $m_{ee} > 0.9 \text{ GeV}/c^2$, even though statistical fluctuations also become sizeable in this mass range.

¹⁷The value was originally stated by the PHENIX collaboration for a similar analysis [PHENIX10] and adopted in the interest of time. Given the non-uniform geometry of the PHENIX detector, leading to a very large R -factor, this value is a quite conservative estimate for an ALICE measurement.




Figure 6.11.: Expected precision for the extraction of the dielectron excess in central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV, based on $2.5 \cdot 10^7$ events. All known sources are subtracted, no DCA cut is used. Systematic uncertainties in the intermediate-mass region are dominated by the uncertainty on the open-charm contribution.

Alternative Configurations

As discussed in Section 6.3.2, an inclusion cut on small electron DCAs can suppress the charm contribution by a factor of 5 relative to the prompt dielectron yield. In return, this cut also reduces the remaining prompt signal by one order of magnitude. This can be seen in the raw yields and it is considered during the sampling process. The combined effect is visible in Figure 6.12. The left panel shows the modified signal composition along with strongly fluctuating and partly missing data points due to the statistics loss. In the right panel, the corresponding excess yield is presented. While systematic uncertainties are under control in this case, the remaining statistics is insufficient.

In practice, the DCA cut will be tuned such that uncertainties of all types become similar-sized. This is approximately achieved when the second major improvement of the ALICE setup for RUN 3 is taken into account. The high-rate upgrade of the TPC readout using GEM detectors is designed to allow for a minimum-bias Pb–Pb event rate of 50 kHz, surpassing the current readout speed by a factor of 100 [ALICE13c]. A

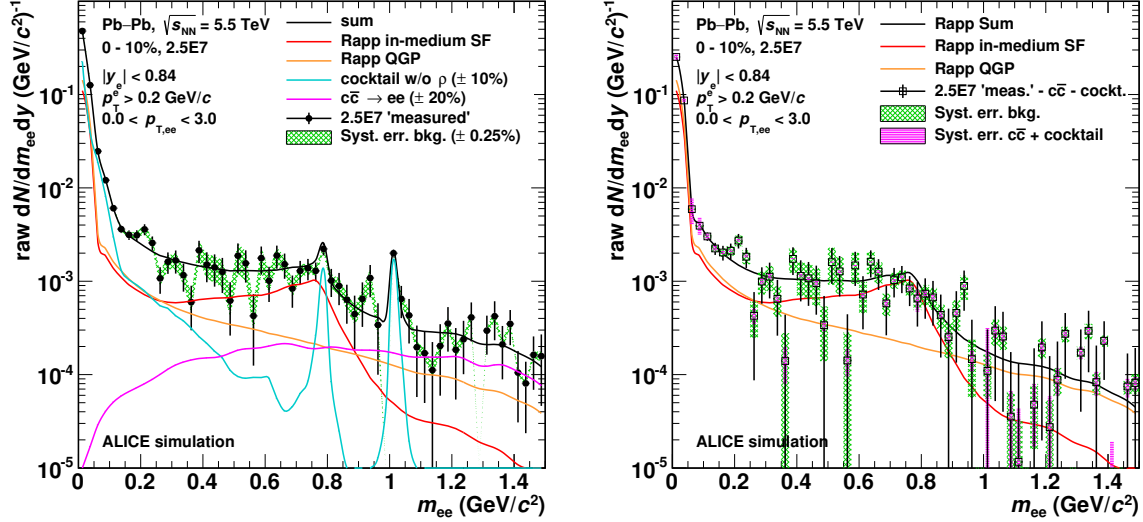


Figure 6.12.: Expected dielectron signal (left) and excess yield (right) in central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV, when a tight DCA cut is applied on the individual electrons to suppress the charm contribution. The sampled yield corresponds to $2.5 \cdot 10^7$ events.

dielectron signal sampled according to the increased event statistics of $2.5 \cdot 10^9$ events and its excess spectrum are presented in Figure 6.13. The same DCA cut is used. Systematic uncertainties from background and charm plus cocktail are comparable in the intermediate-mass region of the excess yield. The result suggests that intriguing statistical precision can be maintained in this case.

6.3.4. Spectra Analysis

In this last analysis step, the excess spectra are evaluated for their extractable physics content. With the experimental goal of measuring the QGP temperature, the key observable for this study is the inverse-slope parameter T of the thermal radiation. No transformation from the extracted (“measured”) inverse-slope parameter T_{meas} to the actual model input temperature is performed to avoid additional complexity and uncertainties, as well as an explicit model dependence. Instead, T_{meas} is compared directly to the inverse-slope parameter which is expected from the shape of the summed model contribution, T_{exp} . The result is thereby independent of the detailed composition of the dielectron excess.

In the onset of the intermediate-mass region, the combined radiation from QGP and




Figure 6.13.: Same as Figure 6.12, but sampled for $2.5 \cdot 10^9$ events. A tight DCA cut is applied to suppress the charm contribution. Systematic as well as statistical uncertainties remain under control in the excess yield.

hadron gas is well-described by the exponential function

$$N \propto \exp\left(-\frac{m_{ee}}{T}\right). \quad (6.6)$$

It is used for a fit in the invariant-mass range $0.9 < m_{ee} < 1.5 \text{ GeV}/c^2$ to make maximum use of the available model input¹⁸.

The fit to the sampled excess yield for $2.5 \cdot 10^9$ events is shown in the top-left panel of Figure 6.14 on a linear scale. The bottom-left panel shows the reference fit to the model input, providing T_{exp} . In the central panels, the data points are jointly shifted to the upper or lower limits of their systematic uncertainty from background subtraction and the fit is repeated for each case. To avoid any correlation with the statistical uncertainty, the data points are explicitly fixed to these limits instead of fluctuating around them. The right panels show the corresponding fits to the uncertainty limits from subtraction of charm and cocktail.

Since the fit result from the sampled data changes with each repetition of the sampling process — especially when investigating lower event rates — a more robust method is needed to determine a reliable result. For this purpose the sampling and fitting is repeated 100 times per configuration and the mean of the resulting T distribution is taken as T_{meas} . Its RMS is consequently used as the statistical uncertainty of the fit

¹⁸The mass range of the model calculations was extended for subsequent iterations of this study.




Figure 6.14.: Fit procedure to extract the “measured” inverse-slope parameter and its uncertainties. Top left panel: exponential fit to the excess yield sampled for $2.5 \cdot 10^9$ events with applied DCA cut. Bottom left panel: corresponding fit to the model input to extract the nominal inverse-slope parameter. Middle (right) panels: fits to the upper and lower bounds of the systematic uncertainties from background (charm plus cocktail) subtraction.

result. In summary, the fits lead to a relative inverse-slope parameter of

$$\frac{T_{\text{meas}}}{T_{\text{exp}}} = 1.02 \pm 0.05 \text{ (stat.) } \begin{matrix} +0.30 \\ -0.10 \end{matrix} \text{ (background) } \begin{matrix} +0.02 \\ -0.01 \end{matrix} \text{ (} \bar{c}\bar{c} \text{ + cocktail) } . \quad (6.7)$$

This result for $T_{\text{meas}}/T_{\text{exp}}$ in the 0–10% centrality class using the upgraded ITS and a high interaction rate is compared to five further experimental configurations in Figure 6.15. They all use a tight cut on the electron DCA in the ITS to reduce the uncertainty from charm. This in turn increases statistical fluctuations in the sampled yield, as discussed in the context of Figure 6.6.

Each pair of markers represents a measurement of T in central (0–10%) as well as semi-peripheral (40–60%) collisions with a certain experimental setup. The left pair corresponds to the current ITS and the RUN 1 interaction rate of 500 Hz. While both fits do reproduce the expected result on average, the meaning of the statistical error bar




Figure 6.15.: Compilation of fit results on the inverse-slope parameter relative to its expected value for the three reasonable combinations of ITS version and interaction rate. Closed (open) markers show the most probable value measured in central (semi-peripheral) collisions, including its standard deviation (see text for details). Green and magenta boxes represent the systematic uncertainties from background and charm plus cocktail subtraction, respectively.

is, that the analysis of data from one heavy-ion beamtime has a 68% chance (1σ) of measuring an inverse-slope parameter that lies within $\approx \pm 100\%$ around the true value. A meaningful measurement is therefore unfeasible in this case.

Using the upgraded ITS at low rate, shown in the center, provides a similar picture, except that systematic uncertainties are significantly reduced. In this configuration the DCA cut may be relaxed until uncertainties from charm and statistics become similar-sized, but a large discrepancy between measurement and true value would remain possible.

The right pair of markers shows results with the added benefit of increasing the interaction rate to 50 kHz, leading to $2.5 \cdot 10^9$ events in central and $5 \cdot 10^9$ in semi-peripheral collisions. All uncertainty types are reduced to within $\pm 10\%$ except for the upper background value in central collisions, as stated in equation (6.7). To recall, the background uncertainty is fortunately based on a rather conservative estimate in comparison to the other contributions.

In summary, the combined improvement from the upgraded ITS and the new TPC readout is expected to make a measurement of the early medium temperature feasible in LHC RUN 3.

Chapter 7.

Summary

Heavy-ion collisions at ultra-relativistic energies enable us to recreate in the laboratory the conditions that existed during the first microseconds after the big bang. They provide the only experimental approach to study strongly-interacting matter in its deconfined form, where the fundamental particles known today become asymptotically free. This state of matter is called the quark–gluon plasma (QGP), and it is the main focus of research in contemporary heavy-ion physics. Its theoretical description is given by Quantum Chromodynamics (QCD), as part of the standard model of particle physics.

A large variety of observables can be studied in order to improve our understanding of the QGP phase, which in heavy-ion collisions is embedded in the evolution of the hot and dense medium of newly created particles. Many of these observables are based on strongly-interacting particles and thus are influenced by the whole system evolution, such that they need to be complemented with theoretical model calculations in order to infer from them the properties of the QGP phase exclusively. In contrast to this, electromagnetic probes such as photons and electrons traverse the medium without significant interaction and thus carry unmodified information from all collision stages. This grants direct access to some of the QGP properties, most prominently its temperature, given that contributions from early and late stages can be differentiated.

Dilepton observables, i. e. electron–positron or muon–antimuon pairs, are of particular interest in this context, since their invariant mass reflects the amount of energy available at the time of their formation. The most energetic reactions occur in the initial scatterings of the colliding nuclei and produce dileptons directly from quark–antiquark annihilation and in the decay chains of correlated heavy-flavour particles. As the system subsequently evolves through the QGP phase, it emits thermal radiation in the form of photons and dileptons, which carry the temperature information to the observer. Thermal

dilepton radiation is also produced in the hadron gas until chemical freeze-out is reached, which for high collision energies occurs near the quark–hadron phase transition. In this phase, dilepton emission is sensitive to in-medium modifications of the electromagnetic spectral function, providing unique experimental access to study the chiral properties of hot QCD. Finally, electromagnetic decays of various light-flavour mesons lead to a cocktail of additional signals in the dilepton spectrum.

Previous experiments have found evidence for the in-medium modification of the ρ meson, which dominates the electromagnetic spectral function, but could not yet isolate thermal QGP radiation from the other sources.

In the work presented here, a new dielectron measurement has been performed in the context of the ALICE experiment at CERN’s Large Hadron Collider, using data from Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV recorded in 2011 during LHC RUN 1.

ALICE is the dedicated heavy-ion experiment at the LHC and is designed to address a wide range of observables with its multiple subsystems. In particular its particle identification and low-momentum tracking capabilities are unique amongst the LHC experiments. They are essential for a low-mass dielectron analysis by enabling the measurement of electrons and positrons with transverse momenta down to $p_{\text{T}} \approx 0.1$ GeV/ c .

In the presented analysis, a composite PID strategy has been developed in order to best combine the capabilities of the Inner Tracking System, the Time Projection Chamber and the Time-Of-Flight detector for electron identification. The track selection procedure has been customized to compensate for detector issues, allowing for the recovery of an additional $\approx 50\%$ of otherwise lost electron tracks. The dielectron signal needs to be extracted on a statistical basis, dealing with combinatorial background that is up to a factor of 1000 larger than the signal. Two signal quality criteria, namely the signal-to-background ratio and the significance, have been discussed and employed to tune a pair prefilter for the removal of electrons from unwanted sources at minimum signal loss. An artifact of the tracking algorithm has been exploited to suppress electrons from photon conversions and to correct the raw dielectron yield for a contribution from different-conversion pairs.

With all efforts combined, a dielectron reconstruction efficiency above 10% has been achieved — compared to only $\approx 1\%$ in a previously published result — which is crucial for this statistics-hungry measurement. A factorization approach has been developed to precisely determine this efficiency from the limited amount of simulated Monte–Carlo data, while taking into account detector resolution effects and the pair survival probability in the prefilter. Systematic uncertainties in the data have been studied by repeating

the analysis for 20 strong and combined variations of multiple electron selection criteria, resulting in dielectron efficiencies of $\approx 2\text{--}20\%$. This not only covers track-based uncertainties, but also those related to background subtraction and signal extraction, since the signal-to-background ratio also varies by at least a factor of 2. The RMS deviation between the 20 results and their average is assigned as the signal uncertainty, which amounts to $\approx \pm 20\%$ in most of the studied kinematic range.

Results have been presented for two data samples, separated into the 0–10% most central and 10–50% semi-central Pb–Pb collisions, each containing close to 20 million events. The dielectron signal has been extracted differentially in its two main kinematic dimensions, namely the invariant mass and transverse momentum of the lepton pairs, and discussed as a function of these variables.

Overall, the measured dielectron yield can be reasonably well described by a cocktail of known hadronic sources, which has been constructed based on ALICE measurements and event generators. As a function of invariant mass and integrated up to $p_{T,ee} = 6\text{ GeV}/c$, a hint of an enhancement in the data is visible in the low-mass region, where the hadronic cocktail is dominated by the η meson and where an additional contribution from thermal radiation is expected. While this trend is present in both studied centrality classes, it is more significant in central collisions. Transverse-momentum spectra have been extracted for four separate intervals of invariant mass in order to analyze the relevant mass regions individually. Data and cocktail are in good overall agreement, despite the changing spectral shapes from low to intermediate mass. This $p_{T,ee}$ -differential comparison also helps to understand the tension between data and cocktail in the resonance region, by revealing some downward fluctuations below and around $p_{T,ee} = 1\text{ GeV}/c$ that are caused by large combinatorial background.

The measurement presented here is also consistent with results published previously by ALICE, while achieving a much smoother mass distribution up to $m_{ee} = 4\text{ GeV}/c^2$ due to higher significance. This has allowed for the extension and shift of the analysis focus towards the intermediate-mass region.

In-medium effects in Pb–Pb collisions lead to a suppression of the heavy-flavour cocktail contributions in comparison to N_{coll} -scaled proton–proton results. They have been accounted for in the J/ψ component of the cocktail, but studied explicitly for the open-heavy-flavour contributions. Two approaches have been discussed, either incorporating a measurement of the nuclear modification factor of electrons from heavy-flavour hadrons into the event generator, or applying heuristic modifications to the generated electron kinematics. The former causes on average 50% suppression of the open-charm yield in

the intermediate-mass region, where this contribution is most relevant, while different scenarios of the latter lead to 10–40 % suppression. This reduced, more realistic heavy-flavour yield creates some additional space at low and intermediate mass for a thermal contribution to the dielectron signal.

The special PID strategy used in the presented work has allowed for the reduction of the low- p_T electron selection threshold from 0.4 to 0.2 GeV/ c , thereby extending the dielectron kinematic range such that consistent comparisons with two theoretical predictions have become possible. One of them is provided by the thermal fireball model, which is based on a hydrodynamic description of the collision evolution from which the temperature profile is inferred. Expected yields of thermal dielectron radiation from the quark–gluon plasma and the hadron gas are separately computed within the model using individual electromagnetic spectral functions. The most notable change in the combined dielectron yield from hadronic cocktail plus model appears between the pion and ω meson mass, mainly due to the broad contribution from thermal ρ mesons. The presented measurement tends to agree with this enhancement, even though in some mass bins the data is clearly below the expectation for various reasons. The other prediction comes from the UrQMD model, which contains a microscopic transport simulation of the system evolution and thereby provides its own light-flavour cocktail. Thermal contributions are computed similarly to the previous model, with the additional step of coarse-graining the system into space–time cells to extract its thermodynamic quantities. For a comparison with this model, the heavy-flavour cocktail contributions have been subtracted from the data. The agreement is better when the nuclear modification factor of heavy-flavour electrons is taken into account instead of using the N_{coll} -scaled charm and beauty yields. A conclusive comparison will need data with high statistical precision as well as reduced uncertainties in particular for the heavy-flavour cocktail. Both improvements are anticipated with the upgraded ALICE detector.

One of three physics performance studies to motivate and justify the ALICE upgrade program has been realized within this thesis by studying the feasibility and scientific potential of a low-mass dielectron measurement with the upgraded experimental setup. A holy grail of heavy-ion physics is the measurement of the initial QGP temperature, and dileptons may be the only experimental probe to provide unbiased access to this property. Therefore the main goal of the presented study has been to determine the achievable precision of such a measurement after the upgrade.

A multitude of inputs based on simulations, tools, theoretical predictions, and some assumptions have been processed and merged into a complete dielectron analysis. Signal-

to-background ratio and significance have been extracted from the customized event sample and used to generate a realistic dielectron “measurement” including statistical and systematic uncertainties via a Poisson sampling technique. Dielectrons from correlated open-charm hadrons are the dominant physics background for a QGP-temperature measurement in the intermediate-mass region. A procedure to reduce their contribution by a topological cut on the electron kinematics has been investigated and is expected to suppress their relative yield by up to a factor of five. The open-charm and light-flavour cocktail contributions have then been subtracted from the sampled data to isolate the thermal excess yield provided by the theoretical model. Based on the propagated signal, background, and cocktail uncertainties, it has been evaluated that the inverse slope of the excess — which is directly related to the QGP temperature — can be re-extracted from the data with statistical and systematic uncertainties on the order of $\pm 10\%$.

In conclusion, the presented low-mass dielectron feasibility study and its promising results have helped to enable the ALICE upgrade program and to converge on an ITS design with precise tracking and no PID capabilities. It will be exciting to compare the predictions of this study to actual measurements with the upgraded ALICE detector in the near future.

Appendix A.

ITS PID Calibration

As described in Section 3.2.2, the PID responses for electrons in the ITS and the TPC are post-calibrated for a precise electron selection in this analysis. Figure A.1 shows the input values (left) and the smoothed correction parameters (right) for the ITS.




Figure A.1.: Mean (top) and width (bottom) of the ITS electron PID response as a function of pseudorapidity and multiplicity. Left panels: parameters obtained via generalized double Gaussian fits to a sample of conversion electrons. Right panels: smoothed parameter distributions from linear fits to the $N_{\text{trk}}^{\text{TPC}}$ dependence and interpolation in η , used for PID post-calibration during the analysis.

Appendix B.

Definition of ϕ_V

The angle ϕ_V is introduced to measure the orientation of the e^+e^- pair plane with respect to the magnetic field direction [Dah08]. It is used to identify photon conversions and exclude them from the analysis, which helps to reduce the combinatorial background and clean up the π^0 Dalitz yield, as discussed in Section 4.4 and Section 4.5, respectively. The magnetic field is aligned with the z axis of the ALICE coordinate system. For a pair with track momentum vectors \vec{p}_1 and \vec{p}_2 , ϕ_V is defined by the following construction:

$$\vec{u} = \frac{\vec{p}_1 + \vec{p}_2}{|\vec{p}_1 + \vec{p}_2|} \quad (\text{B.1})$$

$$\vec{v} = \vec{p}_1 \times \vec{p}_2 \quad (\text{B.2})$$

$$\vec{w} = \vec{u} \times \vec{v} \quad (\text{B.3})$$

$$\vec{u}_a = \frac{\vec{u} \times \hat{z}}{|\vec{u} \times \hat{z}|} \quad (\text{B.4})$$

$$\phi_V = \arccos\left(\frac{\vec{w} \cdot \vec{u}_a}{|\vec{w}| |\vec{u}_a|}\right). \quad (\text{B.5})$$

For unlike-sign pairs, \vec{p}_1 is assigned to the positron and \vec{p}_2 to the electron, so that for conversion pairs the vector product \vec{v} points into the negative z direction. For like-sign pairs, where charge-ordering is ambiguous, \vec{p}_1 and \vec{p}_2 are ordered by momentum such that pairs with overlapping track geometry reproduce the ϕ_V of corresponding unlike-sign pairs. This is crucial for the combinatorial description of the signals' ϕ_V distribution by the like-sign background. The combinatorial yield from different conversions requires further efforts, as discussed in Section 4.5.1.

Prospects of Low-Mass Dielectron Measurement in ALICE with an upgraded Central Barrel Detector

P. Reichelt¹ for the ALICE Collaboration

Motivation

Physics Case:

The measurement of electron-positron pairs in the low invariant mass region allows to study the vacuum and in-medium properties of light vector mesons. Dielectrons also probe the production of thermal photons in heavy-ion collisions. ALICE is well-suited to perform this measurement due to its excellent tracking and particle identification capabilities at very low momenta.

Experimental Challenge:

- Dalitz decays and photon conversions lead to a high combinatorial background.
- Coincident semi-leptonic decays of charm and anti-charm hadrons produce a continuum signal, which dominates over a thermal dielectron signal.

Strategy:

- **Upgrade of Inner Tracking System (ITS) [1]**
 - will improve the tracking efficiency at very low transverse momentum
 - will provide excellent detection capabilities for electrons from secondary vertices
- **Reduction of central barrel magnetic field to $B = 0.2$ T**
 - will extend global tracking efficiency and electron PID to lower momenta
- **High rate upgrade of Time Projection Chamber (TPC) [1]**
 - will increase the data taking rate by a factor of 100




Fig. 1: Comparison of ITS tracking from current to new ITS and global tracking from nominal to reduced magnetic field.

Settings of this Study

► Dielectron signal contributions

- Hadronic cocktail (Exodus)
- Correlated open charm decay
- Thermal signal expectations (R. Rapp) [R. Rapp and J. Wambach, Eur. Phys. J. A6 (1999) 415]

► Background study

- Electron sample:
- Pythia pp events (superimposed)
 - Geant photon conversions
- Kinematic cuts:
- $|p_{el}| < 0.84$
 - $p_{T,e} > 0.2$ GeV/c (global tracking)
 - $p_{T,e} > 0.06$ GeV/c (ITS standalone)
- Conversion & Dalitz rejection:
- $M_{ee} < 50$ MeV/c² & angle < 100 mrad
- Optional cut on significance of Impact Parameter (DCA) measured with ITS:
- at primary e efficiency of 50% (current)
 - at primary e efficiency of 32% (new ITS)




Fig. 2: Invariant mass vs. opening angle for real pairs (left) and fake pairs. Pair cuts are indicated.




Fig. 3: DCA significance distributions (left) and efficiencies after cut of electrons from different sources, for current and new ITS.

Predictions for future e^+e^- Measurements

Simulation of Pb-Pb collisions at $\sqrt{s_{NN}}=5.5$ TeV:

- 1. Compute total signal**
 - for given $dN_{ch}/d\eta$ and N_{coll}
 - correct for pair efficiency
- 2. For each ITS version**
 - generate combinatorial background
 - compute S/B and significance
- 3. Sample signal spectrum**
 - according to number of events and significance
- 4. Extract excess spectrum**
 - subtract known contributions with assumed errors (hadronic cocktail $\pm 10\%$, open charm $\pm 20\%$)
 - propagate statistical and systematic errors

► Number of events per 10% centrality

- $2.5 \cdot 10^7$ at present readout rate
- $2.5 \cdot 10^9$ at high TPC rate (50 kHz)

► Multiplicity at $\sqrt{s_{NN}}=5.5$ TeV

- $dN_{ch}/d\eta = 1750$ (0-10%) / 250 (40-60%)
- $N_{coll} = 1625 / 140$ (for charm)

Expectations for 40 - 60% semi-central collisions:




Fig. 5: Signal (upper) and excess spectra as function of M_{ee} (left) and $p_{T,ee}$ for new ITS with DCA cut and high rate.

Expectations for 0 - 10% central collisions:




Fig. 4: e^+e^- Backgrounds and Signal, S/B, significance for current and new ITS, without DCA cuts.




Fig. 6: Expected absolute statistical uncertainty on dilepton elliptical flow, calculated from the excess spectra of Fig. 5.




Fig. 7: Average values of thermal fits to excess mass spectra from central and semi-central collisions.




Fig. 8: Mass-integrated pair- p_T spectrum for central collisions. Outlook: thermal fit after acceptance correction. High rate statistics also allow for mass-differential spectra (e.g. in the p region).

Further Implications

► Elliptic flow

- statistical uncertainties in excess spectra allow to estimate flow significance: $\sigma_{v_2} \approx 0.7 / (N/\sigma_N)$ (from MC)
- relevant for semi-central collisions




Fig. 6: Expected absolute statistical uncertainty on dilepton elliptical flow, calculated from the excess spectra of Fig. 5.

► QGP temperature

- exponential fit to excess mass spectrum gives effective temperature of thermal radiation: $dN/dM_{ee} \sim \exp(-M_{ee}/T)$ for $M_{ee} > 0.9$ GeV/c²
- mass region dominated by charm \rightarrow DCA cut




Fig. 7: Average values of thermal fits to excess mass spectra from central and semi-central collisions.

From current to new ITS:

- reduced charm & cocktail error in both centralities

New ITS with high rate:

- T -measurement possible
- background error in central to be optimized

Summary & Outlook

► Summary:

- Presentation of feasibility study for dilepton measurement at $\sqrt{s_{NN}}=5.5$ TeV featuring ALICE upgrade plans:
 - New Inner Tracking System & high rate upgrade of TPC
- This will allow - for the first time - a multi-differential low-mass dielectron measurement in heavy-ion collisions

► Expectations for other observables:

- Dilepton v_2 : absolute uncertainty of the order of $\sigma_{v_2} \approx 1\%$
- Thermal radiation T : uncertainty of approximately 10%

► Outlook:

- Perform QGP fits also on excess pair- p_T spectra (see Fig. 8)
- Repeat this study for current ITS at nominal B-field to compare to analysis of available Pb-Pb data
- Try to improve analysis strategy based on experience gained from this study




Fig. 8: Mass-integrated pair- p_T spectrum for central collisions. Outlook: thermal fit after acceptance correction. High rate statistics also allow for mass-differential spectra (e.g. in the p region).

Appendix C.

Upgrade Study Poster

The low-mass dielectron feasibility study has first been publicly presented by means of the conference poster shown on the left. In addition to the results discussed in Chapter 6, the dielectron excess expectations for the 40–60 % centrality class have been analyzed concerning their sensitivity to elliptic flow.

The asymmetric interaction region in non-central heavy-ion collisions — in combination with collective effects in the expanding medium — leads to an azimuthal anisotropy of particles, which is quantified by the elliptic flow coefficient v_2 [SSS10]. Therefore a flow measurement of thermal radiation from the quark–gluon plasma would give direct access to its collective properties.

The achievable precision of the elliptic-flow measurement depends on the statistical uncertainty σ_N of the corresponding integrated yield N . Its absolute statistical uncertainty σ_{v_2} can be approximated by the relation

$$\sigma_{v_2} \approx 0.7/(N/\sigma_N), \quad (\text{C.1})$$

which is inferred from Monte–Carlo [App12].

In order to obtain slightly differential results, N and σ_N are evaluated from the dielectron excess spectra (poster Fig. 5) after rebinning them into four (three) wide invariant-mass (pair- p_T) bins. This represents the ability to distinguish between thermal radiation components from the QGP and the hadronic phase by making use of the dielectron invariant mass. Values of $\sigma_{v_2} \approx 0.01$ are expected to be achieved in most of the studied kinematic region (Fig. 6). No model prediction is provided to be reproduced, such that the absolute uncertainties are stated here — in contrast to the thermal fits (Fig. 7). If for example an elliptic flow of $v_2^{\text{exp}} = 5\% = 0.05$ were present, the measurement should obtain $v_2^{\text{meas}} \approx 0.05 \pm 0.01$, i. e. achieve a relative statistical uncertainty of 20 %.

Appendix D.

Write Progression

Figure D.1 shows the progress over time for writing this thesis after completing the analysis¹. The pace of writing is estimated at $\approx \frac{2}{3}$ pages per whole day spent on the document. This partially includes work on figures and reading literature. Sometimes the parallel job required full-time attention, as seen by the plateaus in the enveloping graph.

¹Tracking this has been motivated by a lecture on scientific writing and the result may happily serve as an input for a potential recurrence.




Figure D.1.: Number of written pages per chapter as a function of time.

Bibliography

- [A⁺04] F. Arleo et al., *Hard probes in heavy-ion collisions at the LHC: Photon physics in heavy ion collisions at the LHC*, (2004), hep-ph/0311131.
- [A⁺08] S. Abreu et al., *Proceedings, Workshop on Heavy Ion Collisions at the LHC: Last Call for Predictions*, J. Phys. **G35**, 054001 (2008), 0711.0974.
- [AB⁺17] L. Altenkämper et al., *Applicability of transverse mass scaling in hadronic collisions at energies available at the CERN Large Hadron Collider*, Phys. Rev. **C96**(6), 064907 (2017), 1710.01933.
- [ALICE09] K. Aamodt et al. (ALICE Collaboration), *The ALICE experiment at the CERN LHC (2008 JINST 3 S08002)*, in *The CERN Large Hadron Collider: Accelerator and Experiments - Volume 1*, edited by A. Breskin and R. Voss (ALICE Collaboration), CERN, 2009.
- [ALICE11a] K. Aamodt et al. (ALICE Collaboration), *Centrality Dependence of the Charged-Particle Multiplicity Density at Midrapidity in Pb-Pb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Rev. Lett. **106**, 032301 (Jan 2011).
- [ALICE11b] K. Aamodt et al. (ALICE Collaboration), *Rapidity and transverse momentum dependence of inclusive J/ψ production in pp collisions at $\sqrt{s} = 7$ TeV*, Phys. Lett. **B704**, 442–455 (2011), 1105.0380, [Erratum: Phys. Lett.B718,692(2012)].
- [ALICE12a] B. Abelev et al. (ALICE Collaboration), *Measurement of charm production at central rapidity in proton-proton collisions at $\sqrt{s} = 2.76$ TeV*, JHEP **07**, 191 (2012), 1205.4007.
- [ALICE12b] B. Abelev et al. (ALICE Collaboration), *Upgrade of the ALICE Experiment: Letter of Intent*, Technical Report CERN-LHCC-2012-012. LHCC-I-022. ALICE-UG-002, Aug 2012.

- [ALICE12c] P. Reichelt (for the ALICE Collaboration), Prospects of Low-Mass Dielectron Measurement in ALICE with an upgraded ITS, in *Quark Matter Conference, Washington D.C., USA, August 13-18, 2012*, <https://indico.cern.ch/event/181055/contributions/308754/>, 2012.
- [ALICE13a] B. Abelev et al. (ALICE Collaboration), *Centrality determination of Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV with ALICE*, Phys. Rev. **C88**(4), 044909 (2013), 1301.4361.
- [ALICE13b] B. Abelev et al. (ALICE Collaboration), *Measurement of inelastic, single- and double-diffraction cross sections in proton-proton collisions at the LHC with ALICE*, Eur. Phys. J. **C73**(6), 2456 (2013), 1208.4968.
- [ALICE13c] B. Abelev et al. (ALICE Collaboration), Upgrade of the ALICE Time Projection Chamber, Technical Report CERN-LHCC-2013-020. ALICE-TDR-016, Oct 2013.
- [ALICE13d] P. Antonioli, A. Kluge and W. Riegler (ALICE Collaboration), Upgrade of the ALICE Readout & Trigger System, Technical Report CERN-LHCC-2013-019. ALICE-TDR-015, Sep 2013.
- [ALICE14a] B. Abelev et al. (ALICE Collaboration), *Beauty production in pp collisions at $\sqrt{s} = 2.76$ TeV measured via semi-electronic decays*, Phys. Lett. **B738**, 97–108 (2014), 1405.4144.
- [ALICE14b] B. Abelev et al. (ALICE Collaboration), *Centrality, rapidity and transverse momentum dependence of J/ψ suppression in Pb-Pb collisions at $\sqrt{s_{NN}}=2.76$ TeV*, Phys. Lett. **B734**, 314–327 (2014), 1311.0214.
- [ALICE14c] B. Abelev et al. (ALICE Collaboration), *Performance of the ALICE Experiment at the CERN LHC*, Int. J. Mod. Phys. **A29**, 1430044 (2014), 1402.4476.
- [ALICE14d] B. Abelev et al. (ALICE Collaboration), *Production of charged pions, kaons and protons at large transverse momenta in pp and Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Lett. **B736**, 196–207 (2014), 1401.1250.
- [ALICE14e] B. Abelev et al. (ALICE Collaboration), *Technical Design Report for the Upgrade of the ALICE Inner Tracking System*, Journal of Physics G: Nuclear and Particle Physics **41**(8), 087002 (2014).

- [ALICE15a] J. Adam et al. (ALICE Collaboration), Technical Design Report for the Muon Forward Tracker, Technical Report CERN-LHCC-2015-001. ALICE-TDR-018, Jan 2015.
- [ALICE15b] P. Buncic, M. Krzewicki and P. Vande Vyvre (ALICE Collaboration), Technical Design Report for the Upgrade of the Online-Offline Computing System, Technical Report CERN-LHCC-2015-006. ALICE-TDR-019, Apr 2015.
- [ALICE18a] S. Acharya et al. (ALICE Collaboration), *Dielectron production in proton-proton collisions at $\sqrt{s} = 7$ TeV*, JHEP **09**, 064 (2018), 1805.04391.
- [ALICE18b] S. Acharya et al. (ALICE Collaboration), *Measurements of low- p_T electrons from semileptonic heavy-flavour hadron decays at mid-rapidity in pp and Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, JHEP **10**, 061 (2018), 1805.04379.
- [ALICE18c] S. Acharya et al. (ALICE Collaboration), *Neutral pion and η meson production at mid-rapidity in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Rev. **C98**(4), 044901 (2018), 1803.05490.
- [ALICE19a] S. Acharya et al. (ALICE Collaboration), *Measurement of dielectron production in central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Rev. **C99**(2), 024002 (2019), 1807.00923.
- [ALICE19b] S. Acharya et al. (ALICE Collaboration), *Production of the $\rho(770)^0$ meson in pp and Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Rev. **C99**(6), 064901 (2019), 1805.04365.
- [App12] H. Appelshäuser, priv. communication, 2012.
- [Ave12] R. Averbeck, priv. communication, 2012.
- [Bar02] R. Barlow, Systematic errors: Facts and fictions, in *Advanced Statistical Techniques in Particle Physics. Proceedings, Conference, Durham, UK, March 18-22, 2002*, pages 134–144, 2002.
- [Bar08] J. Bartke, *Introduction to Relativistic Heavy Ion Physics*, World Scientific Publishing, 2008.
- [BB+98] S. Bass et al., *Microscopic models for ultrarelativistic heavy ion collisions*, Progress in Particle and Nuclear Physics **41**, 255 – 369 (1998).

- [Bjo83] J. D. Bjorken, *Highly relativistic nucleus-nucleus collisions: The central rapidity region*, Phys. Rev. D **27**, 140–151 (Jan 1983).
- [BMRS03] P. Braun-Munzinger, K. Redlich and J. Stachel, Particle production in heavy ion collisions, in *Quark-Gluon Plasma 3*, edited by R. C. Hwa and X.-N. Wang, pages 491–599, World Scientific Publishing, 2003.
- [Boo14] J. Book, *J/ψ Production in Pb-Pb Collisions with ALICE at the LHC*, PhD thesis, Johann Wolfgang Goethe-Universität Frankfurt am Main, 2014.
- [BR91] G. Brown and M. Rho, *Scaling effective Lagrangians in a dense medium*, Phys. Rev. Lett. **66**, 2720–2723 (1991).
- [BZ⁺99] M. Bleicher et al., *Relativistic hadron hadron collisions in the ultrarelativistic quantum molecular dynamics model*, J. Phys. **G25**, 1859–1896 (1999), hep-ph/9909407.
- [CBM05] D. Adamova et al. (CBM Collaboration), *Compressed Baryonic Matter Experiment - Technical Status Report*, Technical report, Jan 2005.
- [CERES95] G. Agakichiev et al. (CERES Collaboration), *Enhanced Production of Low-Mass Electron Pairs in 200 GeV/Nucleon S-Au Collisions at the CERN Super Proton Synchrotron*, Phys. Rev. Lett. **75**, 1272–1275 (Aug 1995).
- [CERES98] G. Agakichiev et al. (CERES Collaboration), *Low-mass e^+e^- pair production in 158 A GeV Pb-Au collisions at the CERN SPS, its dependence on multiplicity and transverse momentum*, Phys. Lett. B **422**, 405–412 (1998), nucl-ex/9712008.
- [CERES08] D. Adamova et al. (CERES Collaboration), *Modification of the ρ meson detected by low-mass electron-positron pairs in central Pb-Au collisions at 158 A GeV/c*, Phys. Lett. **B666**, 425–429 (2008), nucl-ex/0611022.
- [CL02] P. Coles and F. Lucchin, *Cosmology: The Origin and evolution of cosmic structure*, Wiley & Sons, New York, 2nd edition, 2002.
- [CO07] B. W. Carroll and D. A. Ostlie, *An Introduction to Modern Astrophysics*, Addison-Wesley, San Francisco: Pearson, 2nd (international) edition, 2007.
- [Dah08] T. Dahms, *Dilepton spectra in p+p and Au+Au collisions at RHIC*, PhD thesis, SUNY, Stony Brook, 2008.

- [EB09] L. Evans and P. Bryant, LHC Machine (2008 *JINST* **3** S08001), in *The CERN Large Hadron Collider: Accelerator and Experiments - Volume 1*, edited by A. Breskin and R. Voss, CERN, 2009.
- [EO14] I. Erdemir-Özdemir, Dielectron cocktail simulation in pp, p-Pb and Pb-Pb collisions at LHC energies, Master's thesis, Johann Wolfgang Goethe-Universität Frankfurt am Main, 2014.
- [EvHB16] S. Endres, H. van Hees and M. Bleicher, *Energy, centrality, and momentum dependence of dielectron production at collider energies in a coarse-grained transport approach*, *Phys. Rev. C* **94**, 024912 (Aug 2016).
- [FH⁺11] B. Friman et al., *The CBM physics book: Compressed baryonic matter in laboratory experiments*, *Lect. Notes Phys.* **814**, pp.1–980 (2011).
- [FK04] Z. Fodor and S. D. Katz, *Critical point of QCD at finite T and mu, lattice results for physical quark masses*, *JHEP* **04**, 050 (2004), hep-lat/0402006.
- [FL12] M. Ferro-Luzzi, Review of 2011 LHC run from the experiments perspective, in *Proceedings, Chamonix 2012 Workshop on LHC Performance, Chamonix, France, 6 - 10 Feb 2012*, pages 38–53, 2012.
- [FNT98] B. Friman, W. Nörenberg and V. D. Toneev, *The Quark condensate in relativistic nucleus-nucleus collisions*, *Eur. Phys. J.* **A3**, 165–170 (1998), nucl-th/9711065.
- [GS08] C. Grupen and B. Shwartz, *Particle Detectors, Second Edition*, Cambridge University Press, 2008.
- [Gun13] T. Gunji, priv. communication, 2013.
- [HFR12] M. He, R. J. Fries and R. Rapp, *Ideal hydrodynamics for bulk and multistrange hadrons in $\sqrt{s_{NN}} = 200A$ GeV Au-Au collisions*, *Phys. Rev. C* **85**, 044911 (Apr 2012).
- [Hin08] F. Hinterberger, *Physik der Teilchenbeschleuniger und Ionoptik*, Springer, 2nd edition, 2008.
- [Hin12] O. Hinrichs, Simulationen zur Messung von Elektron-Positron-Paaren im ALICE-Experiment, Master's thesis, Johann Wolfgang Goethe-Universität Frankfurt am Main, 2012.

- [HJ00] U. W. Heinz and M. Jacob, *Evidence for a new state of matter: An Assessment of the results from the CERN lead beam program*, (1 2000), nucl-th/0002042.
- [Huh17] P. Huhn, Datenbasierte Effizienzkorrektur zum Nachweis geladener Teilchen im ALICE Experiment, Master's thesis, Johann Wolfgang Goethe-Universität Frankfurt am Main, 2017.
- [HW10] R. C. Hwa and X.-N. Wang, *Quark-Gluon Plasma 4*, World Scientific Publishing, 2010.
- [Iva12] M. Ivanov, priv. communication, 2012.
- [J⁺19] J. Jowett et al., The 2018 heavy-ion run of the LHC, in *Proceedings, 10th International Particle Accelerator Conference (IPAC2019): Melbourne, Australia, May 19-24, 2019*, 2019.
- [JAF⁺12] J. M. Jowett et al., Heavy Ions in 2012 and the Programme up to 2022, in *Proceedings, Chamonix 2012 Workshop on LHC Performance, Chamonix, France, 6 - 10 Feb 2012*, pages 200–211, 2012.
- [Jun19] J. Jung, Dielectron production in pp collisions at $\sqrt{s} = 13$ TeV measured in a dedicated low magnetic-field setting with ALICE, Master's thesis, Johann Wolfgang Goethe-Universität Frankfurt am Main, 2019.
- [Kal12] A. Kalweit, *Production of light flavor hadrons and anti-nuclei at the LHC*, PhD thesis, Technische Universität Darmstadt, 2012.
- [KK⁺86] K. Kajantie et al., *Dilepton emission and the QCD phase transition in ultrarelativistic nuclear collisions*, Phys. Rev. D **34**, 2746–2754 (Nov 1986).
- [Kle14] C. Klein, Studie zum Einfluss der Teilchenidentifizierung mit dem ALICE-ITS auf die Messung von Dileptonen kleiner invarianter Masse, Master's thesis, Johann Wolfgang Goethe-Universität Frankfurt am Main, 2014.
- [Kra12] F. Kramer, *J/ψ Production in $\sqrt{s} = 7$ TeV pp Collisions*, PhD thesis, Johann Wolfgang Goethe-Universität Frankfurt am Main, 2012.
- [KS⁺12] P. K. Khandai et al., *Transverse Mass Spectra and Scaling of Hadrons at RHIC and LHC Energies*, (2012), hep-ph/1205.0648.
- [KSS11] P. K. Khandai, P. Shukla and V. Singh, *Meson spectra and m_T scaling in*

- p + p, d + Au, and Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. C **84**, 054904 (Nov 2011).
- [KW55] N. M. Kroll and W. Wada, *Internal pair production associated with the emission of high-energy gamma rays*, Phys. Rev. **98**, 1355–1359 (1955).
- [Kwe12] M. Kweon, priv. communication, 2012.
- [MT85] L. D. McLerran and T. Toimela, *Photon and dilepton emission from the quark-gluon plasma: Some general considerations*, Phys. Rev. D **31**, 545–563 (Feb 1985).
- [Mül07] B. Müller, *From Quark-Gluon Plasma to the Perfect Liquid*, (2007), nucl-th/0710.3366.
- [Mus12] L. Musa, Conceptual Design Report for the Upgrade of the ALICE ITS, Technical Report CERN-LHCC-2012-005. LHCC-G-159, CERN, Geneva, Mar 2012.
- [NA5003] B. Alessandro et al. (NA50 Collaboration), *ϕ production in Pb-Pb collisions at 158 GeV/c per nucleon incident momentum*, Phys. Lett. B **555**, 147–155 (2003), [Erratum: Phys.Lett.B 561, 294–294 (2003)].
- [NA5005] B. Alessandro et al. (NA50 Collaboration), *A new measurement of J/ψ suppression in Pb-Pb collisions at 158 GeV per nucleon*, Eur. Phys. J. C **39**, 335–345 (2005), hep-ex/0412036.
- [NA6006] R. Arnaldi et al. (NA60 Collaboration), *First Measurement of the ρ Spectral Function in High-Energy Nuclear Collisions*, Phys. Rev. Lett. **96**, 162302 (Apr 2006).
- [NA6007] S. Damjanovic et al. (NA60 Collaboration), *NA60 results on the ρ spectral function in In-In collisions*, Nucl. Phys. **A783**, 327–334 (2007), nucl-ex/0701015.
- [PDG18] M. Tanabashi et al. (PDG Collaboration), *Review of Particle Physics*, Phys. Rev. D **98**, 030001 (2018).
- [PHENIX03] K. Adcox et al. (PHENIX Collaboration), *PHENIX detector overview*, Nucl. Instrum. Meth. A **499**, 469–479 (2003).
- [PHENIX07] A. Adare et al. (PHENIX Collaboration), *Energy Loss and Flow of Heavy*

- Quarks in Au+Au Collisions at $s(NN)^{1/2} = 200\text{-GeV}$* , Phys. Rev. Lett. **98**, 172301 (2007), nucl-ex/0611018.
- [PHENIX10] A. Adare et al. (PHENIX Collaboration), *Detailed measurement of the e^+e^- pair continuum in $p + p$ and Au+Au collisions at $\sqrt{s_{NN}} = 200\text{ GeV}$ and implications for direct photon production*, Phys. Rev. **C81**, 034911 (2010), 0912.0244.
- [PHENIX16] A. Adare et al. (PHENIX Collaboration), *Dielectron production in Au+Au collisions at $\sqrt{s_{NN}}=200\text{ GeV}$* , Phys. Rev. **C93**(1), 014904 (2016), 1509.04667.
- [Rap01] R. Rapp, *Signatures of thermal dilepton radiation at ultrarelativistic energies*, Phys. Rev. C **63**, 054907 (Apr 2001).
- [Rap12] R. Rapp, priv. communication, 2012.
- [Rap13] R. Rapp, *Dilepton Spectroscopy of QCD Matter at Collider Energies*, Advances in High Energy Physics **2013**, 1–17 (2013).
- [Rap16] R. Rapp, priv. communication, 2016.
- [Rei11] P. Reichelt, *Simulationsstudien zur Entwicklung des Übergangsstrahlungszählers für das CBM-Experiment*, Master's thesis, Johann Wolfgang Goethe-Universität Frankfurt am Main, 2011.
- [RvH10] R. Rapp and H. van Hees, *Heavy Quarks in the Quark-Gluon Plasma*, in *Quark-Gluon Plasma 4*, edited by R. C. Hwa and X.-N. Wang, pages 111–206, World Scientific Publishing, 2010.
- [RW99] R. Rapp and J. Wambach, *Low mass dileptons at the CERN SPS: Evidence for chiral restoration?*, Eur. Phys. J. **A6**, 415–420 (1999), hep-ph/9907502.
- [RW00] K. Rajagopal and F. Wilczek, *The Condensed matter physics of QCD*, in *At the frontier of particle physics. Handbook of QCD. Vol. 1-3*, edited by M. Shifman and B. Ioffe, pages 2061–2151, World Scientific Publishing, 2000.
- [RWvH09] R. Rapp, J. Wambach and H. van Hees, *The Chiral Restoration Transition of QCD and Low Mass Dileptons*, (2009), hep-ph/0901.3289.
- [S⁺19] R. Steerenberg et al., *Operation and performance of the CERN Large*

- Hadron Collider during proton Run 2, in *Proceedings, 10th International Particle Accelerator Conference (IPAC2019): Melbourne, Australia, May 19-24, 2019*, 2019.
- [Sak69] J. J. Sakurai, *Currents and Mesons*, University of Chicago Press, Chicago, 1969.
- [Sha12] R. Shahoyan, priv. communication, 2012.
- [Shu04] E. Shuryak, *Why does the quark gluon plasma at RHIC behave as a nearly ideal fluid?*, Prog. Part. Nucl. Phys. **53**, 273–303 (2004), hep-ph/0312227.
- [Ska10] P. Z. Skands, *Tuning Monte Carlo Generators: The Perugia Tunes*, Phys. Rev. **D82**, 074018 (2010), 1005.3457.
- [SMS06] T. Sjostrand, S. Mrenna and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*, JHEP **05**, 026 (2006), hep-ph/0603175.
- [SS15] T. Schörner-Sadenius, editor, *The Large Hadron Collider : Harvest of Run 1*, Springer International Publishing, Cham, 2015.
- [SSS10] S. Sarkar, H. Satz and B. Sinha, *The physics of the quark-gluon plasma*, Lect. Notes Phys. **785**, pp.1–369 (2010).
- [STAR03] K. Ackermann et al. (STAR Collaboration), *STAR detector overview*, Nucl. Instrum. Meth. A **499**, 624–632 (2003).
- [STAR13] F. Videbaek (STAR Collaboration), *Heavy flavor physics in STAR*, J. Phys. Conf. Ser. **420**, 012024 (2013).
- [STAR14] Q. Hao (STAR Collaboration), *Status and Physics Opportunities of the STAR Heavy Flavor Tracker and the Muon Telescope Detector Upgrades*, J. Phys. Conf. Ser. **509**, 012026 (2014).
- [STAR15a] L. Adamczyk et al. (STAR Collaboration), *Energy dependence of acceptance-corrected dielectron excess mass spectrum at mid-rapidity in Au+Au collisions at $\sqrt{s_{NN}} = 19.6$ and 200 GeV*, Phys. Lett. B **750**, 64–71 (2015), 1501.05341.
- [STAR15b] L. Adamczyk et al. (STAR Collaboration), *Measurements of Dielectron Production in Au+Au Collisions at $\sqrt{s_{NN}} = 200$ GeV from the STAR Experiment*, Phys. Rev. C **92**(2), 024912 (2015), 1504.01317.

- [STAR18] J. Adam et al. (STAR Collaboration), *Measurements of Dielectron Production in Au+Au Collisions at $\sqrt{s_{NN}} = 27, 39, \text{ and } 62.4 \text{ GeV}$ from the STAR Experiment*, (10 2018), 1810.10159.
- [TI⁺03] V. D. Toneev et al., Three fluid simulations of relativistic heavy ion collisions, in *Proceedings, 12th International Conference on Selected Problems of Modern Physics, dedicated to the 95th anniversary of the birth of D.I. Blokhintsev (1908-1979): Dubna, Russia, June 8-11, 2003*, 2003.
- [Toi12] A. Toia, priv. communication, 2012.
- [TOTEM19] G. Antchev et al. (TOTEM Collaboration), *First measurement of elastic, inelastic and total cross-section at $\sqrt{s} = 13 \text{ TeV}$ by TOTEM and overview of cross-section data at LHC energies*, Eur. Phys. J. **C79**(2), 103 (2019), 1712.06153.
- [vHR08] H. van Hees and R. Rapp, *Dilepton Radiation at the CERN Super Proton Synchrotron*, Nucl. Phys. **A806**, 339–387 (2008), hep-ph/0711.3444.
- [Vog07] R. Vogt, *Ultrarelativistic Heavy-Ion Collisions*, Elsevier Science B.V., 2007.
- [Wie07] H. Wiedemann, *Particle Accelerator Physics*, Springer, 3rd edition, 2007.
- [Won94] C.-Y. Wong, *Introduction to High-Energy Heavy-Ion Collisions*, World Scientific Publishing, 1994.

Danksagung

An dieser Stelle bedanke ich mich ganz besonders bei meinem Doktorvater Prof. Dr. Harald Appelshäuser, der durch seine Leidenschaft schon früh mein Interesse an diesem Forschungsfeld geweckt und mir diese wissenschaftliche Laufbahn ermöglicht hat. Ich bin ihm auch sehr dankbar dafür, dass er mir mit den Dielektronen eines seiner Herzenthemen als Doktorarbeit anvertraut hat. Vielen Dank für die intensive Zusammenarbeit während der spannenden Studie zur Mitgestaltung der Dielektronen-Zukunft in ALICE, die meine Arbeit gleichermaßen bereichert wie verlängert hat, und für die tolle Betreuung im weiteren Verlauf meiner Doktorarbeit.

Hiermit bedanke ich mich ebenfalls bei Prof. Dr. Christoph Blume für die Zweitbegutachtung dieser Arbeit und für die vielmalige Verlängerung meines CERN-Vertrages. Prof. Dr. Marcus Bleicher danke ich für das Einbringen der theoretischen Perspektive zu meinem Thema während der PhD-Komitee-Meetings.

In Bezug auf die Upgrade-Studie bedanke ich mich auch bei allen externen Beteiligten und speziell bei Prof. Dr. Ralf Rapp für seine Modellrechnungen, damit verbundene Diskussionen und die Übernahme meines Farbschemas in seine Publikation.

Ein großer Dank gilt auch den Verantwortlichen der Helmholtz-Graduiertenschule HGS-HIRe for FAIR für die Vergabe eines Stipendiums und für die Aufnahme ins Graduiertenprogramm H-QM mit den damit verbundenen intensiven Vorlesungswochen, Softskill-Seminaren und weiteren Veranstaltungen. Insbesondere danke ich dabei Prof. Dr. Henner Büsching für sein großartiges Engagement zu deren Realisierung, sowie für viele freundliche Gespräche.

Außerdem bedanke ich mich bei allen Kolleginnen und Kollegen am IKF für ein sehr freundschaftliches Miteinander, für gesellige Zeiten abseits des Schreibtisches, sowie besonders bei Theo Bröker, Dr. Carsten Klein, Sebastian Scheid und Jerome Jung für die gute Zusammenarbeit und die gemeinsame Entwicklung unseres Analyse-Frameworks.

Ganz herzlich danke ich auch allen mir während dieser langen Zeit nahestehenden für ihre Unterstützung sowie ihr Interesse an meiner Arbeit und ihrem Fortschritt.