

Lattice QCD at Finite Temperature with Wilson Fermions

Dissertation
zur Erlangung des Doktorgrades
der Naturwissenschaften

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

von
Christopher Pinke
aus Münster

Frankfurt am Main 2014
D30

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

Dekan: Prof. Dr. J. Stroth

Gutachter: Prof. Dr. O. Philipsen
Prof. Dr. H. Wittig (Universität Mainz)

Datum der Disputation:

CONTENTS

Abstract	v
Deutsche Zusammenfassung	vii
1 INTRODUCTION	1
2 QUANTUM CHROMODYNAMICS (QCD)	7
2.1 Continuum QCD at Finite Temperature and Density	7
2.2 Symmetries of QCD	10
2.2.1 Chiral Symmetry	10
2.2.2 Centre Symmetry	11
2.2.3 Roberge-Weiss Symmetry	12
2.3 The Twisted Mass Formulation of QCD	13
2.4 The Running Coupling of QCD	14
2.5 QCD on the Lattice	15
2.5.1 Wilson Fermions	17
2.5.2 Twisted Mass Wilson Fermions	19
2.5.3 Alternative Fermion Discretisations	20
2.5.4 Finite T and μ	21
2.5.5 Spectroscopy on the Lattice	22
2.6 The QCD Phase Diagram and Heavy Ion Collisions	24
3 NUMERICAL ASPECTS OF LATTICE QCD AND THE CL ² QCD SOFTWARE	29
3.1 Algorithms	29
3.1.1 Hybrid Monte Carlo	30
3.1.2 Heatbath Algorithm	32
3.1.3 Inversion of the Fermion Matrix	32
3.1.4 Reweighting	33
3.1.5 The Sign Problem	34
3.2 Data Analysis	35
3.3 Numerical Structure of LQCD	36
3.4 Graphics Processing Units (GPUs)	38
3.4.1 GPU Programming	39
3.5 Lattice QCD based on OpenCL: CL ² QCD	42
3.5.1 Implementation Details	43
3.5.2 Performance Results	47
3.6 Summary and Discussion	53
4 THE THERMAL TRANSITION IN TWISTED MASS LATTICE QCD	55
4.1 The Columbia Plot	55
4.2 Investigations of the Thermal Transition with $N_f = 2$ Twisted Mass LQCD . . .	58
4.3 Scale Setting	60

4.4	Simulation Setup	61
4.5	The Thermal Transition for the Z12 Ensemble	64
4.6	Strength of the Axial Anomaly	76
4.7	Summary and Discussion	82
5	THE ROBERGE-WEISS TRANSITION WITH WILSON FERMIONS	85
5.1	The QCD Phase Diagram at Imaginary μ	85
5.2	Simulation Setup & Observables	91
5.3	Roberge-Weiss Transition as a Function of the Fermion Mass	95
5.4	Summary and Discussion	102
6	CONCLUSIONS AND RESEARCH PERSPECTIVES	103
A	APPENDIX	107
A.1	Notations and Conventions	107
A.2	$SU(N)$ and Group Theory	108
A.3	Grassmann Calculus	109
A.4	Even-odd Preconditioning	110
A.5	Estimation of D^{-1}	110
A.6	CL ² QCD Details	112
	A.6.1 Compiling CL ² QCD	112
	A.6.2 Memory Requirements	112
	A.6.3 Even-odd Mapping	113
A.7	Clusters & Machines	114
A.8	Details of Simulations with Twisted Mass Fermions	116
A.9	Details of Simulations with Wilson Fermions	119
	Danksagung	137
	Lebenslauf	139
	Notizen	145

ABSTRACT

The subatomic world is governed by the strong interactions of quarks and gluons, described by *Quantum Chromodynamics (QCD)*. Quarks experience *confinement* into colour-less objects, i.e. they can not be observed as free particles. Under extreme conditions such as high temperature or high density, this constraint softens and a transition to a phase where quarks and gluons are quasi-free particles (*Quark-Gluon-Plasma*) can occur. This environment resembles the conditions prevailing during the early stages of the universe shortly after the Big Bang. The phase diagram of QCD is under investigation in current and future collider experiments, for example at the *Large Hadron Collider (LHC)* or at the *Facility for Antiproton and Ion Research (FAIR)*. Due to the strength of the strong interactions in the energy regime of interest, analytic methods can not be applied rigorously. The only tool to study QCD from first principles is given by simulations of its discretised version, *Lattice QCD (LQCD)*. These simulations are in the high-performance computing area, hence, the numerical aspects of LQCD are a vital part in this field of research.

In recent years, *Graphic Processing Units (GPUs)* have been incorporated in these simulations as they are a standard tool for general purpose calculations today. In the course of this thesis, the LQCD application CL^2QCD has been developed, which allows for simulations on GPUs as well as on traditional CPUs, as it is based on OpenCL. CL^2QCD constitutes the first application for *Wilson type fermions* in OpenCL. It provides excellent performance and has been applied in physics studies presented in this thesis.

The investigation of the QCD phase diagram is hampered by the notorious *sign-problem*, which restricts current simulation algorithms to small values of the chemical potential. Theoretically, studying unphysical parameter ranges allows for constraints on the phase diagram. Of utmost importance is the clarification of the order of the finite temperature transition in the $N_f = 2$ chiral limit at zero chemical potential. It is not known if it is of first or second order. To this end, simulations utilising Twisted Mass Wilson fermions aiming at the chiral limit are presented in this thesis. Another possibility is the investigation of QCD at purely imaginary chemical potential. In this region, QCD is known to possess a rich phase structure, which can be used to constrain the phase diagram of QCD at real chemical potential and to clarify the nature of the $N_f = 2$ chiral limit. This phase structure is studied within this thesis, in particular the nature of the Roberge-Weiss endpoint is mapped out using Wilson fermions.

DEUTSCHE ZUSAMMENFASSUNG

Die Bausteine hadronischer Materie sind die *Quarks*, von denen es sechs verschiedene Arten gibt. Sie wechselwirken mittels der *Gluonen* miteinander, beschrieben durch die *Quantenchromodynamik (QCD)*. Diese Quantenfeldtheorie ist Teil des *Standardmodells der Teilchenphysik*, welches erfolgreich alle heute bekannten Teilchen und ihre Wechselwirkungen beschreibt. Diese werden in aktuellen und zukünftigen Beschleunigerexperimenten untersucht, z.B. am *Large Hadron Collider (LHC)* am CERN oder an der *Facility for Antiproton and Ion Research (FAIR)* an der GSI in Darmstadt. Unlängst wurde am LHC der experimentelle Nachweis für den wichtigen *Higgssektor* gefunden, welcher dazu benötigt wird, den Elementarteilchen eine Masse zuzuschreiben ohne grundlegende Symmetrien zu verletzen. Quarks tragen Farbladung und physikalische Objekte müssen stets so zusammengesetzt sein, dass sie farbneutral erscheinen (*Confinement*). Aufgrund der nicht abelschen Natur der Eichgruppe der QCD wechselwirken die Gluonen auch untereinander, was enorme Auswirkungen auf die Struktur stark wechselwirkender Materie hat. Insbesondere besitzt QCD die Eigenschaft der *asymptotischen Freiheit*, was bedeutet das die Stärke der Wechselwirkung abnimmt je näher sich die Elementarteilchen kommen bzw. je mehr Energie im betrachteten Prozess enthalten ist. Andersherum nimmt die Anziehung bei sich vergrößerndem Abstand zu. Des Weiteren besitzt QCD eine intrinsische Energieskala, Λ_{QCD} , welche bei ungefähr 200 MeV liegt. Das entspricht einer Längenskala im Femtometerbereich. Bei Energien dieser Größenordnung ist die starke Wechselwirkung von der Ordnung Eins. Für die Energieskalen der subatomaren Welt bedeutet dies, dass, anders als in der Quantenelektrodynamik, man QCD hier nicht mittels Störungstheorie behandeln kann, sie ist explizit nicht-perturbativ.

QCD wird daher im Rahmen von Gittereichtheorien behandelt. In diesen wird die (kontinuierliche) Raumzeit auf einem Hyperkubus diskretisiert und benachbarte Punkte sind durch den Gitterabstand a voneinander entfernt. Hierdurch ist das Pfadintegral, also die Quantisierung der klassischen Theorie, wohldefiniert und lässt sich numerisch berechnen. Für die exakte Wahl der diskretisierten Form der QCD (Gitter-QCD) bestehen gewisse Freiheiten: Eine Gitterwirkung muss nur den korrekten Kontinuumslikes, $a \rightarrow 0$, besitzen. Sie unterscheiden sich daher bei endlichem a mitunter deutlich, sei es in ihren numerischen Eigenschaften oder ihrer Abhängigkeit vom Gitterabstand. Dies wird besonders bei den Fermionen deutlich. Wenn diese naiv diskretisiert werden beschreibt die Gitterwirkung nicht eins, sondern sechzehn Fermionen. Diese *Doppler* werden durch die Gitterregularisierung verursacht. Um sie loszuwerden gibt es verschiedene Methoden. In der von Wilson vorgeschlagenen Diskretisierung bekommen fünfzehn Doppler eine Masse welche antiproportional zum Gitterabstand ist, so dass sie von der Dynamik des Systems entkoppeln. Dies hat allerdings zur Folge, dass die *chirale Symmetrie* hierbei explizit gebrochen wird. Bei dieser sind im chiralen (masselosen) Limes Teilchen unterschiedlicher Chiralität unabhängig voneinander. Um die Abhängigkeit vom Gitterabstand von Wilsonfermionen zu verbessern werden weitere Terme zur Wirkung hinzugefügt. Zum Beispiel

kann man einen chiral verdrehten Massenterm hinzufügen, was die Kontinuumstheorie nicht ändert, aber die Diskretisierungsartefakte in führender Ordnung verschwinden lässt bei entsprechender Wahl der Parameter (*Twisted Mass Wilsonfermionen*). Ein anderer Ansatz wird bei *Staggeredfermionen* verfolgt. Hier wird durch geschickte Wahl der Diracmatrizen die Anzahl der Doppler effektiv auf vier verringert. Um diese weiter auf eins zu reduzieren wird das sogenannte *Rooting* angewendet, bei welchem nicht restlos geklärt ist, ob es zur korrekten Physik führt. *Staggeredfermionen* sind jedoch in numerischer Hinsicht um ein Vielfaches günstiger als *Wilsonfermionen* und sind daher oft verwendet worden. In jüngerer Zeit gibt es weitere Fermionformulierungen, welche jedoch erheblich kostspieliger in numerischer Hinsicht sind und daher noch selten eingesetzt werden. Mit dem Fortschreiten der Technik ist es heutzutage möglich, Gitter-QCD bei physikalischen Massen zu simulieren (auch wenn dies immer noch sehr teuer ist). Das Gitter-QCD in der Tat korrekte Ergebnisse liefert wurde in mehreren Studien belegt, welche mit Experimenten übereinstimmende Ergebnisse hervorbrachten.

Das Pfadintegral auf dem Gitter ist wohldefiniert. Die grassmannwertigen Fermionen können exakt ausintegriert werden und man erhält die *Fermiondeterminante*. Die Freiheitsgrade des Systems sind dann gegeben durch die Eichfelder. Das resultierende Integral kann allerdings aufgrund seiner hohen Dimensionalität praktisch nicht exakt ausgewertet werden. Daher werden Monte-Carlo Methoden verwendet, bei welchen nur solche Bereiche des Phasenraums abgetastet werden, die einen großen Beitrag zum Pfadintegral darstellen. Hierzu wird der Integrand als Wahrscheinlichkeit interpretiert. Die Fermiondeterminante läßt sich dabei als pseudofermionisches Integral schreiben, welches von der inversen Fermionmatrix abhängt. Die Berechnung dieser Größe wird mit sogenannten Krylov-Lösern durchgeführt, da die Matrix dünn besetzt ist. Dieser Teil der Simulationen ist sehr teuer, da es sich um einen iterativen Prozess handelt, im Zuge dessen die Wirkung der Fermionmatrix auf einen Vektor vielfach ausgerechnet werden muss. Vor allem der Ableitungsterm \not{D} ist numerisch aufwendig. Die benötigten Speicherobjekte stellen eine große Datenmenge dar im Vergleich zur Anzahl der Rechenoperationen, die auf ihnen auszuführen sind. Daher sind Gitter-QCD Anwendungen immer speicherbandbreitenlimitiert. Allerdings lassen sie sich auch auf natürliche Weise parallelisieren, d.h. auf mehreren Prozessoren werden zeitgleich verschiedene Teile des gesamten Gitters behandelt. Gitter-QCD ist eines der aufwändigsten Computerprobleme und Hochleistungsrechner werden hierfür verwendet.

In jüngerer Zeit ist die Benutzung von Grafikkarten (GPUs) interessant geworden. Diese stellen, vor allem im Vergleich zu traditionellen CPUs, enorme Rechenkapazitäten bereit, was sich durch die hohen Anforderungen in der Videospiegelindustrie erklärt. Insbesondere sind sie sehr gut für parallelisierte Anwendungen geeignet. Zusätzlich bieten GPUs ein attraktives Verhältnis von Rechenkapazität zu Kaufpreis, d.h. sie sind auch in der Anschaffung günstig. Daher sind sie integraler Bestandteil vieler moderner Hochleistungsrechner. So gibt es in Frankfurt den LOEWE-CSC, welcher eine hybride Architektur darstellt: Pro Rechenknoten gibt es zwei CPUs und eine GPU. Des Weiteren wurde SANAM an der Goethe-Universität Frankfurt entwickelt, welcher in höherem Maße auf GPUs als Recheneinheiten setzt: Hier gibt es vier GPUs pro Rechenknoten. Beide Rechner bestehen aus Hardware des Herstellers AMD.

Auch in Gitter-QCD sind GPUs in Gebrauch und können vor allem die Invertierung der Fermionmatrix beschleunigen. Allerdings sind die existierenden Anwendungen auf CUDA basierend. Diese Technik erlaubt es, Grafikkarten zu programmieren, ist allerdings nur auf Hardware des Herstellers NVIDIA anwendbar. Eine Alternative ist durch den offenen Standard OpenCL gegeben, welcher hardwareunabhängig ist. Zusätzlich ist es mit diesem möglich, CPUs und GPUs mit demselben Programm zu benutzen. Im Laufe der vorliegenden Arbeit wurde die neue Gitter-QCD Anwendung CL²QCD geschrieben.¹ Sie basiert auf einem C++-Hostprogramm, welches die algorithmische Logik und I/O-Aufgaben übernimmt. Zusätzlich steuert dieses die Ausführung der in OpenCL geschriebenen Gitter-QCD-Funktionen auf der Grafikkarte bzw. CPU. Mit CL²QCD ist es möglich, sowohl AMD als auch NVIDIA Hardware zu benutzen. Der Fokus der Entwicklung wurde auf Anwendungen mit Wilsonfermionen bei endlichen Temperaturen gelegt, aber auch andere Aspekte von Gitter-QCD, wie z.B. reine Eichtheorie, wurden verwirklicht. Es wurden alle im Umfang dieser Arbeit benötigten Algorithmen integriert. Die für diesen Anwendungsfall relevanten Funktionen zeigen eine sehr gute Performance auf der im LOEWE-CSC bzw. SANAM verbauten Hardware. Vor allem der oben erwähnte \mathcal{D} erzielt im Vergleich zu Literaturwerten bessere Resultate. Insbesondere können hier, für eine Vielzahl von Gittervolumina, 70-80% der theoretisch überhaupt möglichen Speicherbandbreite der GPU ausgenutzt werden, was zeigt das die Hardware optimal genutzt wird. Als Beispiel für die Rechenkraft von GPUs wurde die Anwendung von CL²QCD unter realistischen Bedingungen mit einem Referenzprogramm verglichen. Das Resultat ist eindrucksvoll: Die Leistung der LOEWE-CSC-GPU ist vergleichbar mit zwei ganzen LOEWE-CSC-CPU-Knoten (4 CPUs). Mit der moderneren SANAM-GPU erhöht sich die Beschleunigung um einen Faktor zwei. Diese Ergebnisse unterstreichen die Vorteile von GPUs für Gitter-QCD. CL²QCD ist darauf ausgelegt, leicht erweiterbar zu sein. Beispielsweise gibt es zurzeit ein Projekt, Staggeredfermionen zu integrieren. Von hohem Interesse ist die parallele Nutzung mehrerer GPUs, da der Hauptspeicher heutiger GPUs momentan der limitierende Faktor für physikalische Studien ist. Diese ist, z.B. auf SANAM, bereits möglich mit CL²QCD. Das Programm wurde bereits erfolgreich in physikalischen Studien verwendet, wie im Folgenden ausgeführt.

Von besonderem Interesse, sowohl experimentell als auch theoretisch, ist die Erforschung des QCD Phasendiagramms, also die Erforschung der Phasenstruktur von stark wechselwirkender Materie bei Änderung der Temperatur oder der Dichte. Aufgrund der oben erwähnten asymptotischen Freiheit lässt sich für sehr hohe Temperaturbereiche eine Phase erwarten, in der Quark und Gluonen nicht länger dem Confinement genügen müssen, sondern sie quasi-freie Teilchen sind (sie sind *deconfined*). Dieser Zustand wird *Quark-Gluon-Plasma* genannt und ist Gegenstand aktueller Experimente. Phänomenologisch ist dies von Bedeutung da im frühen Universum, also kurz nach dem Big Bang, auch sehr hohe Temperaturen herrschten. Der Übergang vom Quark-Gluon-Plasma zur hadronischen Welt, der sogenannte *thermische Übergang der QCD*, kann also Aufschluss über die Entwicklung des Universums geben. Neben dem thermischen Übergang besitzt QCD noch weitere Phasen, beispielsweise bilden sich für hohe Dichten Cooperpaare und eine farbsupraleitende Phase entsteht. Im Vergleich zur Supraleitung der Elektrodynamik gibt es hier verschiedene Strukturen aufgrund der komplizierteren

¹ Siehe <http://code.compeng.uni-frankfurt.de/projects/clhmc> .

Struktur der Theorie. Der Bereich hoher Dichten ist in astrophysikalischen Zusammenhängen interessant, beispielsweise beim Aufbau von Neutronensternen.

Ungeklärt ist die Frage, ob der thermische Übergang der QCD in einem gewissen Temperatur- und Dichtebereich ein echter Phasenübergang ist. Für $\mu = 0$ ist aus Gitterstudien bekannt, dass es sich hier um einen schnellen, aber kontinuierlichen Übergang handelt der bei $T_c \approx 150 - 170$ MeV stattfindet. Bei einem Phasenübergang erster Ordnung hingegen würde eine Diskontinuität in den thermodynamischen Größen auftreten. Aufgrund von allgemeinen Überlegungen und Modellstudien wird so eine Art von Übergang für größere Werte von μ erwartet. Beide Regime müssten sich dann in einem kritischen Endpunkt zweiter Ordnung treffen. Allerdings ist dieser Bereich des Phasendiagramms (zurzeit) nicht mit Gitter-QCD, und damit in der vollen Theorie, erforschbar, da die vorhandenen Simulationsalgorithmen bei endlichem μ durch das *Vorzeichenproblem* unbrauchbar werden. In diesem Parameterbereich ist die Fermioneterminante komplex (insbesondere nicht mehr positiv (semi-)definit) und kann nicht mehr als Wahrscheinlichkeit interpretiert werden. Gitterstudien sind daher auf den Bereich $\mu = 0$ beschränkt. Hier lassen sich verschiedene Limites der Theorie untersuchen. Der thermische Übergang passiert bei Energien der Größenordnung 200 MeV. Hier sind nur die drei leichtesten Quarks relevante Freiheitsgrade, das charm-Quark ist bereits viel schwerer. Weiterhin lassen sich up- und down-Quark in sehr guter Näherung als gleich schwer behandeln.

Gegenstand aktueller Gitterstudien sind daher verschiedene Grenzfälle dieser Näherung, wobei die Massen der Quarks in den Simulationen unphysikalische Werte annehmen. Mit diesen ändert sich auch die Art des thermischen Übergangs. Generell lassen sich hier zwei Regime unterscheiden, der chirale Limes sehr leichter bzw. verschwindender Massen und der entgegengesetzten Limes sehr schwerer Massen. Bei letzterem spielen die (schweren) Quarks eine untergeordnete Rolle und die Dynamik wird hauptsächlich durch den gluonischen Teil der Wirkung bestimmt. Dieser zeigt einen Phasenübergang erster Ordnung von einer confined in eine deconfined Phase bei $T_c \approx 270$ MeV. Der Ordnungsparameter für diesen Übergang ist der Polyakovloop. Da die diesem Übergang zugrundeliegende Zentrumssymmetrie durch Quarks explizit gebrochen wird, schwächen diese den Phasenübergang ab, bis er schließlich, bei Verringerung der Quarkmasse, in einem kritischen Endpunkt zweiter Ordnung kontinuierlich wird. Im Regime leichter Massen ist der chirale Phasenübergang relevant. Die Wiederherstellung der chiralen Symmetrie wird durch das Schmelzen des chiralen Kondensates $\langle \bar{\psi}\psi \rangle$ angezeigt. Im chiralen Limes hängt die Ordnung dieses Übergangs von der Anzahl der Quarkflavours ab. Es kann allgemein gezeigt werden dass er für drei masselose Quarks ($N_f = 3$) erster Ordnung ist. Bei endlichen Quarkmassen wird der Übergang abgeschwächt, bis er schließlich kontinuierlich wird. Am Übergangspunkt beider Phasengebiete gibt es wiederum einen kritischen Endpunkt zweiter Ordnung, welcher in der $Z(2)$ Universalitätsklasse des Ising Modells ist. Für den Limes von zwei dynamischen Quarks, masselosen ($N_f = 2$) hingegen ist nicht bekannt wie der Phasenübergang aussieht. Hier gibt es zwei Möglichkeiten, nämlich dass er erster oder zweiter Ordnung ist. Letzteres Szenario wird allgemein als das Wahrscheinlichere gesehen, es gibt aber noch keinen rigorosen Beweis hiervon, z.B. aus Gittersimulationen. In diesem Fall wäre der Übergang in der $O(4)$ Universalitätsklasse. Beide Möglichkeiten haben jedoch enorme Implikationen für die Phasenstruktur der QCD bei $\mu = 0$. Im Fall eines Übergangs erster Ordnung würde sich,

ähnlich zu $N_f = 3$, eine Region solcher Übergänge bei endlichen Massen ergeben, welche dann in einem $Z(2)$ Endpunkt endet. Die beiden $N_f = 2$ und $N_f = 3$ Endpunkte sollten dann, bei entsprechender Änderung der Quarkmassen, durch eine Linie von $Z(2)$ Endpunkten verbunden sein. Im Fall eines Übergangs zweiter Ordnung im $N_f = 2$ chiralen Limes hingegen würde es einen trikritischen Punkt bei einer gewissen Masse des strange Quarks und verschwindender leichten Quarkmasse geben. An diesem trafen sich dann die $N_f = 2$ $O(4)$ Linie und die $Z(2)$ Linie, die den $N_f = 3$ ersten Ordnungsbereich einschließt. Bei endlichem μ werden diese Linien zu Oberflächen. Da bekannt ist dass der Übergang bei physikalischen Massen kontinuierlich ist (siehe oben), ist es nun entscheidend zu wissen wie sich diese Oberflächen bei Änderung von μ verhalten, um eine Aussage über einen kritischen Endpunkt im QCD Phasendiagramm treffen zu können. Hierfür muss aber auch die Art des Übergangs im $N_f = 2$ chiralen Limes geklärt sein, dieser Punkt ist daher von großer Bedeutung.

Die direkte Simulation im chiralen Limes ist aufgrund von algorithmischen Schwierigkeiten (noch) nicht möglich. In Gitterstudien wird daher bei endlichen Massen simuliert, signalisiert durch die Pionmasse m_π , die Masse des leichtesten Teilchens. Die tmft-Kollaboration hat hierzu eine Reihe von $N_f = 2$ Simulationen bei verschiedenen Pionmassen durchgeführt mit den oben erwähnten Twisted-Mass Wilsonfermionen, wobei die leichteste Masse bei $m_\pi \approx 316$ MeV lag. Für jedes m_π wurde bei mehreren Temperaturen simuliert und der chirale Übergang bestimmt. Aussagen über den chiralen Limes sind möglich indem man das Skalierungsverhalten des chiralen Kondensates $\langle \bar{\psi}\psi \rangle$ analysiert. Jedoch konnte nicht zweifelsfrei zwischen den beiden Szenarien unterschieden werden. Daher wurde im Zuge dieser Arbeit eine weitere Simulationsreihe bei $m_\pi \approx 270$ MeV hinzugefügt und die Analyse entsprechend erweitert. Diese Simulationen sind numerisch sehr aufwendig, da große Gittervolumen nötig sind um mögliche Effekte des endlichen Volumens auszuschließen. Sie wurden auf LOEWE-CSC und SANAM hauptsächlich unter Benutzung von CL²QCD durchgeführt. Hier konnte eine im Vergleich zu den vorigen Massenpunkten um ein Vielfaches höhere Datenmenge produziert werden, genauer gesagt besitzt die neue Simulationsreihe fast doppelt so viele Datenpunkte wie die vorherigen Reihen zusammengenommen. Die Analyse von $\langle \bar{\psi}\psi \rangle$ gestaltete sich jedoch schwieriger als bei den vorherigen Analysen. Dies ist mit der kleinen Masse zu begründen, bei welcher das Signal sehr verrauscht ist. Weiterhin ist die Skalensetzung in dem verwendeten Parameterbereich von zunehmend hohen Fehlern beeinträchtigt. Nichtsdestotrotz stimmt der extrahierte Übergangspunkt mit den Erwartungen aus den vorigen Simulationen überein. Ebenfalls analysiert wurde ein mögliches Verschwinden der axialen Anomalie am chiralen Übergang. Dieses würde für das $O(4)$ Szenario und gegen einen Übergang erster Ordnung im chiralen Limes sprechen. In der Simulation würde sich dies in gleichen Massen von chiralen Partnerteilchen manifestieren, was jedoch nicht gesehen wurde. Die Simulationen bestätigen somit das Bild der vorherigen Studien: Das $O(4)$ Szenario ist begünstigt, kann aber bei diesen Massen noch nicht verifiziert werden.

Ein weiter interessanter Aspekt von QCD ist dessen Phasenstruktur bei rein imaginären Werten des chemischen Potentials, μ_I . Zuallererst gibt es hier kein Vorzeichenproblem, d.h. die Fermiondeterminante ist reell und Simulationen können hier problemlos ausgeführt werden. Des Weiteren besitzt QCD hier die sogenannte *Roberge-Weiss (RW) Symmetrie*. Diese besagt,

das das Phasendiagramm in diesem Bereich periodisch ist. Weiterhin impliziert die Symmetrie das bei kritischen Werten μ_f^c das System einen Übergang zwischen verschiedenen Zentrumssektoren erfährt. Dieser ist erster Ordnung für hohe und ein kontinuierlicher Übergang für tiefe Temperaturen. Konsequenterweise gibt es einen Treffpunkt beider Regionen, den *RW Endpunkt*. Dessen Natur ist höchst nicht trivial und hängt von der Anzahl der Flavours und ihrer Masse ab. Dies ist darin begründet, dass die chiralen und deconfinement Übergänge bei $\mu = 0$ (siehe oben) sich in den Bereich rein imaginären μ fortsetzen. Genauer gesagt treffen sie bei μ_f^c mit dem RW Endpunkt zusammen. Die oben erläuterten Änderungen des Phasenübergangs bei $\mu = 0$ und sich ändernder Masse spiegeln sich nun im RW Endpunkt wider. Für schwere Massen ist der deconfinement Übergang erster Ordnung, so dass sich am RW Endpunkt drei erste Ordnungslinien treffen (Triplepunkt). Dasselbe ist der Fall für leichte Massen und den chiralen Phasenübergang. Für den mittleren Massenbereich ist der thermische Übergang kontinuierlich, das bedeutet für den RW Endpunkt das er ein kritischer Endpunkt zweiter Ordnung in der Ising Universalitätsklasse ist. Die drei Bereiche sind durch trikritische Punkte begrenzt. Diese Erkenntnisse wurde aus Simulationen mit Staggeredfermionen für $N_f = 2$ und $N_f = 3$ gewonnen.

In dieser Arbeit konnte dieses Bild in Simulationen mit $N_f = 2$ Wilsonfermionen exakt bestätigt werden. Hierzu wurde ein großer Massenbereich in den Simulationen abgedeckt. Bei jeder Masse wurden bis zu fünf räumliche Volumen, bei fixierter zeitlicher Ausdehnung, simuliert. Zusätzlich wurde der kritische Bereich mit einer sehr feinen Auflösung abgetastet. Die somit sehr große Anzahl von Simulationen wurde mit dem oben erwähnten CL^2 QCD auf LOEWECSC und SANAM durchgeführt. Auf Grundlage der erhobenen Daten wurden dann Finite-Size Analysen von mehreren Größen durchgeführt und die Ordnung des Phasenübergangs für jeden Massenpunkt extrahiert. Die Ergebnisse zeigen deutlich wie sich das System bei Änderung der Quarkmasse unterschiedlich verhält. Insbesondere konnten beide trikritische Punkte lokalisiert werden. Die Bestätigung der vorherigen Studien mit einer anderen Fermiondiskretisierung ist von großer Bedeutung, da sich aus dem Bereich der QCD bei imaginärem μ Einschränkungen für das physikalische Phasendiagramm ableiten lassen. Ähnlich zur Schilderung bei $\mu = 0$ werden bei $\mu = \mu_f^c$ und sich ändernden Massen die trikritischen Punkte zu trikritischen Linien. Für Werte von μ zwischen $\mu = 0$ und μ_f^c schließen diese dann die von $\mu = 0$ kommenden $Z(2)$ Oberflächen ab. Das Verfolgen der kritischen Oberfläche ist in Simulationen möglich. Die vorliegenden Ergebnisse sind somit ein wichtiger Schritt dorthin und sie haben des weiteren eine sehr gute Ausgangsposition geschaffen für weitere Studien mit Wilsonfermionen in diese Richtung. Insbesondere bietet dieser Ansatz die Möglichkeit, die Art des Übergangs im $N_f = 2$ chiralen Limes zu klären.

*If you hit a wrong note,
it's the next note that you play that determines
if its's good or bad.*

— Miles Davis

INTRODUCTION

It has been discovered many years ago that nuclei have a substructure and their binding can be explained by strongly interacting elementary particles, the *quarks* and *gluons*. There are six flavours of quarks with increasing mass: up (≈ 3 MeV), down (≈ 5 MeV), strange (≈ 95 MeV), charm (≈ 1.3 GeV), bottom (≈ 4.3 GeV) and top (≈ 173 GeV).¹ However, free quarks cannot be observed in experiment: Quarks experience *confinement* into baryons and mesons. This led to the postulation of a new quantum number, the so-called *colour*. Colour comes in three different charges, and elementary particles have to combine into colour-singlets to be physical. Introductions to this field of research can be found in the literature, for instance [Perkins, 1982; Povh et al., 1993] have been used for this presentation. A direct experimental verification of the number of colours N_c and the number of flavours N_f can be seen in electron-electron scattering experiments, where the cross-section is proportional to N_c and increases with energy each time a mass threshold for another quark flavour is reached. The eight massless gluons are the mediators of the strong force. In contrast to the photons in Quantum Electrodynamics, they carry colour charge and interact among themselves. This has severe implications for theoretical investigations of the theory as well as for the particle spectrum. In fact, particles made up solely of gluons (*glueballs*) have been postulated for long and are subject of current investigations (see [Eshraim et al., 2013] and references therein). The existence of gluons can be inferred from jet events observed in high-energy experiments.

From the theoretical point of view, the strong interactions are formulated as a quantum field theory with $SU(N_c = 3)$ gauge symmetry called *Quantum Chromodynamics (QCD)*. Together with the electroweak interactions QCD forms the *standard model of particle physics*, which successfully describes all known elementary particles. The part of the standard model necessary for the particles to have a (bare) mass, the famous *Higgs-sector* [Englert and Brout, 1964; Higgs, 1964] (Nobel prize 2013), has recently been verified experimentally with the discovery of a Higgs-like boson [Aad et al., 2012; Chatrchyan et al., 2012], confirming a theoretically long-predicted essential component.

An important property of the strong force discovered in deep inelastic scattering experiments is *asymptotic freedom*. It has been shown by Wilczek, Politzer and Gross [Gross and Wilczek, 1973a,b; Politzer, 1973] that non-abelian gauge theories have this property (Nobel prize 2004). The force between strongly interacting particles decreases with decreasing distance and gets

¹ For more information see reviews by the *Particle Data Group*, e.g. [Beringer et al., 2012] or <http://pdg.lbl.gov/>.

stronger if the particles are separated: The coupling α_s runs with the energy involved in the process, Q^2 (cf. (2.46)), and quarks are quasi-free particles if they are close together. In turn, the potential between two quarks grows linearly for larger distances. Even if the particles were largely separated, the flux tube between them would just break at some point and yield two additional, energetically favourable particles, but there would be still no free colour charges (*string breaking*). This may serve as a qualitative picture of confinement. The running of α_s can be investigated both in perturbation theory and in experiment and one finds very good agreement, for instance that at the mass scale of the Z-boson $\alpha_s \approx 0.1$. This is one proof of QCD being the correct theory of the strong interactions.

QCD inhabits an intrinsic mass scale $\Lambda_{\text{QCD}} \approx 200$ MeV. This is the relevant scale for the interactions of quarks and gluons since $1/\Lambda_{\text{QCD}} \approx \mathcal{O}(1\text{fm})$ is of the order of the length scale of the subatomic world. At these energies α_s is $\mathcal{O}(1)$ and perturbation theory is valid for $Q^2 \gg \Lambda_{\text{QCD}}$ only.² In fact, its non-perturbative nature is a feature of QCD and has many implications, such as the non-trivial topology of the QCD vacuum caused by instantons (see for example [Creutz, 2011; 't Hooft, 1999]), which plays a role in the strong CP problem and the axial anomaly. Furthermore, the bare masses of valence up- and down quarks in the nuclei make up only approximately one percent of their masses. This discrepancy can be explained by the sea of quarks and gluons, effectively giving mass to the nuclei. Thus, the Higgs-mechanism is responsible for a small part of the mass of nuclei only, its major part originates from QCD.

The non-perturbative regime can be studied from first principles using the lattice approach. Introductions to this field can be found in [Gattringer and Lang, 2010; Montvay and Münster, 1994]. Here, Euclidean space time is discretised and QCD is formulated on a hypercube as a lattice gauge theory. This yields a version of QCD afflicted with discretisation artifacts called *Lattice QCD (LQCD)*. LQCD can be evaluated by means of numerical methods, and constitutes one of the most compute-intensive problems, placing state-of-the-art LQCD simulations in the supercomputing area. Continuum QCD is recovered by taking the continuum limit. The validity of QCD has been successfully demonstrated, for example in the calculation of a variety of meson masses in perfect accordance with their experimental values [Durr et al., 2008].

Of great interest, both, experimentally and theoretically, is the phase diagram of QCD, which can be seen in Figure 1. For recent reviews see [Braun-Munzinger and Wambach, 2009; Fukushima and Hatsuda, 2011]. At low temperatures, as μ is increased, nuclear matter undergoes a first order liquid-gas phase transition. At zero T this happens at values of μ approximately the mass of the proton. At finite temperature, the transition bends to lower μ and finally ends in a second order *critical endpoint (CEP)* (*CEP of Nuclear Physics*). A colour-superconducting phase is expected at low T and high μ , where the quarks form Cooper pairs. This region has a rich phase structure as quarks can form various colour-flavour combinations and it is believed to be relevant for astrophysical objects, e.g. neutron stars. In the context of LQCD the *thermal transition* of QCD is also of particular interest. As T is increased, hadrons will eventually melt and QCD passes over to a deconfined phase, the *Quark-Gluon Plasma*. In this phase, quarks and gluons are quasi-free particles. It resembles the environment realised in the early universe shortly after the Big Bang and before hadronic matter is formed

² In fact, there are additional complications for perturbation theory, see e.g. [Linde, 1980].




Figure 1: QCD phase diagram as a function of temperature T and chemical potential μ . See the text for details.

as the universe cooled down. It can be recreated in collider experiments, in which particles are accelerated to nearly the speed of light and eventually collide, releasing a massive amount of energy. Studying this allows to draw conclusions about details of the formation of the universe and its present state.

In particular, the phase boundary of this transition is addressed in current and future high-energy experiments at the *Relativistic Heavy Ion Collider (RHIC)*, the *Large Hadron Collider (LHC)* or at the *Facility for Antiproton and Ion Research (FAIR)*.³ For zero μ , it is known from lattice calculations that the transition is actually a rapid crossover and happens at a critical temperature T_c of 150-170 MeV [Aoki et al., 2006]. At an expected QCD CEP, the thermal transition would become a first order line. It is not known how this line and the superconducting deconfinement transition are related, which is depicted by the yellow blob in the figure. Indeed, it is not known yet if the QCD CEP exists after all. It is predicted by various model calculations, but not known from first principles. In fact, direct explorations of the phase diagram in LQCD are prevented by the infamous *sign problem* (see Section 3.1), which spoils currently used simulation algorithms. Hence, the current knowledge about the phase diagram of QCD is restricted mainly to the $\mu = 0$ region. In addition, as the charm quark is much heavier than Λ_{QCD} , studies relevant to the thermal transition are carried out with the two or three lightest quarks only. To obtain information about the physical point, it can be advantageous to vary the quark masses to unphysical values as the phase structure here puts constraints on the physical case. In fact, QCD possesses several symmetries if the quark masses are set to the limiting cases of zero or infinity. For theoretical predictions about the QCD CEP from first principles it is of utmost importance to map out the different phases of QCD for different numbers of flavour and quark masses. Of particular importance is the *chiral symmetry* of chiral (massless) QCD. It is broken spontaneously by the QCD vacuum, yielding massless particle excitations, the *Goldstone bosons*. At physical quark masses, this symmetry is broken explicitly. Yet, for the light up and down quarks, it can be regarded as approximately

³ See <http://www.bnl.gov/rhic/>, <http://home.web.cern.ch/about/accelerators/large-hadron-collider> and <http://www.fair-center.de>.

realised and the pions can be identified with the Goldstone bosons. As temperature is increased, chiral symmetry is effectively restored. At the physical point, this happens around the same T_c as the deconfinement transition to the Quark-Gluon-Plasma.

Conclusions about the QCD CEP can be drawn once the nature of the chiral transition at $\mu = 0$ is known. For two massless flavours this is an unsettled issue. In this limit, it is possible that the transition is of first or second order. This depends crucially on the fate of the axial anomaly. Each of the scenarios has substantial implications for the overall phase diagram of QCD. In particular, the prediction of the CEP based on model calculations relies on the second order scenario. The settlement of this issue by means of LQCD simulations is the central topic of this thesis.

It is currently not possible for LQCD simulations to simulate directly in the chiral limit because of algorithmic problems and tremendously increasing computational costs when approaching this limit by decreasing the pion mass m_π . Furthermore, the discretisation of the fermionic part of QCD is not unique and various alternatives exist, which differ in terms of computational and algorithmic demands. Of particular importance for the continuum limit are the discretisation errors of the specific lattice action. In addition, symmetries are subject to change when discretising QCD. One particular choice of lattice fermions are the so-called *Wilson fermions* [Wilson, 1974], which are used in this thesis. They come with moderate computational costs and are widely used within the lattice community. A variant are the *Twisted Mass Wilson fermions* [Frezzotti and Rossi, 2004], which are obtained by adding a chirally twisted mass term. They have the great advantage that under a certain choice of parameters the leading discretisation artifacts vanish. Alternatives are given by *clover Wilson* [Sheikholeslami and Wohlert, 1985] or *staggered* [Kogut and Susskind, 1975] fermions. More recently, more sophisticated fermion formulations are put to use (*Domain-Wall* and *Overlap fermions* [Ginsparg and Wilson, 1982; Kaplan, 1992; Neuberger, 1998a,b; Shamir, 1993]), which preserve a lattice variant of chiral symmetry. The latter is broken for Wilson fermions. However, these new fermion types by default have very large computational costs.

To investigate the two flavour chiral limit, a series of simulation has been carried out by the *twisted mass finite temperature (tmft)* collaboration [Burger et al., 2013a] using Twisted Mass Wilson fermions at decreasing m_π , down to $m_\pi \approx 316$ MeV. The observed data is consistent with the second order scenario for the chiral limit, yet, it cannot be discriminated from the first order scenario. To this end, the study is extended within this thesis by adding a new gauge field ensemble at $m_\pi \approx 270$ MeV. This ensemble is analysed by means of the chiral transition itself and a possible restoration of the axial anomaly.

Another alternative approach is given by studying QCD at purely imaginary values of the chemical potential. In this region, QCD has a rich phase structure, in particular the so-called *Roberge-Weiss* transition takes place at critical values of μ [Roberge and Weiss, 1986]. This transition between adjacent centre sectors is of first order μ for high and a crossover for low temperatures. These two domains meet in the *Roberge-Weiss endpoint*. The nature of this point depends non-trivially on the number of flavours and the quark mass, i.e. it changes from a triple point to a second order endpoint and back as the mass is varied. These regions are connected by tricritical points. Going to lower values of μ , away from the critical value where

the Roberge-Weiss transition takes place, critical lines emerge from the tricritical points. It is possible to judge on the nature of the transition at $\mu = 0$ following the line emerging from the tricritical point at lighter mass towards the chiral limit [Bonati et al., 2013]. The current understanding of the Roberge-Weiss endpoint is based on studies with staggered fermions [Bonati et al., 2011c; de Forcrand and Philipsen, 2010]. This fermion type is subject to ongoing discussions about the *rooting procedure* [Creutz, 2007; Sharpe, 2006]. Hence, this topic is addressed utilising Wilson type fermions in this thesis in order to confirm these findings.

A central subject in lattice gauge theories is the numerical aspect and a lot of effort is spent on the efficient implementation of LQCD within the lattice community.⁴ In recent years, the simulations have been accelerated tremendously by using Graphic Processing Units (GPUs) for the calculations [Clark et al., 2010]. GPUs are particularly well-suited for LQCD purposes as they offer a high memory bandwidth, huge computational resources and are designed for parallelised problems. Most of modern computing clusters are equipped with GPUs besides traditional CPUs, for example the LOEWE-CSC [Bach et al., 2011] and SANAM [Kalcher et al., 2013] supercomputers developed at Frankfurt University. These are equipped with hardware from AMD. Virtually all GPU-enabled LQCD applications are based on CUDA by NVIDIA [NVIDIA, 2013], which allows to program GPUs, but it is vendor-specific. An alternative is given by the OpenCL standard [Khronos Working Group, 2013], which is vendor independent. In addition, it allows to use CPUs with the same code, too, i.e. it is also hardware independent. Within this thesis, a new OpenCL-based LQCD application called CL²QCD covering various aspects of lattice gauge theories has been developed.⁵

This thesis is organised as follows: In the next chapter, QCD is formally introduced as a quantum field theory with a focus on finite temperature and chemical potential. In addition, the lattice approach is introduced, in particular the (twisted mass) Wilson discretisation of the quarks. At the end of the chapter, the QCD phase diagram and heavy ion collision experiments are discussed in some detail. In the centre of Chapter 3 there are the numerical aspects of LQCD. Standard algorithms for LQCD's numerical realisation in software applications are given. The numerical structure of LQCD is presented, motivating the usage of GPUs. Then, the LQCD application CL²QCD is presented by means of implementation details and performance results. The Chapters 4 and 5 are devoted to the investigations of the two flavour chiral transition and the nature of the Roberge-Weiss endpoint mentioned above. At the beginning of each of these chapters, the current status in the respective fields is reviewed. The chapters are concluded with a short summary. In the end of the thesis, concluding remarks and possible research perspectives are given. Rather technical information is provided in the Appendix.

4 Manifest in proceedings from the *Algorithms & Machines* sessions at the annual Lattice conferences, e.g. <http://www.lattice2013.uni-mainz.de/> .

5 See <http://code.compeng.uni-frankfurt.de/projects/clhmc> .

 QUANTUM CHROMODYNAMICS

In this chapter *Quantum Chromodynamics* (QCD) is formally introduced as a quantum field theory and some of its properties important for the following are touched, in particular QCD at non-vanishing temperature and density (Section 2.1) and certain symmetries of QCD (Section 2.2). Sections 2.3 and 2.4 are dedicated to the Twisted Mass formulation and the running coupling of QCD, respectively. The lattice formulation of QCD will be motivated and introduced in Section 2.5. At the end of this chapter, the QCD phase diagram will be discussed together with current collider experiments. Aspects of statistical physics necessary for a coherent description of phase transitions are also introduced. Here and in the following, notation required for a clear understanding of the text is introduced as first encountered; rather technical properties and details of used quantities are collected and explained in the Appendix.

A variety of textbooks on the topics covered in this presentation can be found in literature. I mainly rely on the books by Peskin and Schröder [Peskin and Schroeder, 1995], Kapusta and Gale [Kapusta and Gale, 2006], Gattringer and Lang [Gattringer and Lang, 2010], deGrand and deTar [DeGrand and DeTar, 2006] and Montvay and Münster [Montvay and Münster, 1994].

2.1 CONTINUUM QCD AT FINITE TEMPERATURE AND DENSITY

QCD is formulated as an $SU(N_c)$ gauge theory describing the fermions as spinor fields $\psi = \psi_{a,\alpha,f}(x)$.¹ Their dynamics in Minkowski time are governed by

$$\mathcal{S}_f[\bar{\psi}, \psi, A_\mu] = \int d^4x \bar{\psi}(x) D[A_\mu] \psi(x), \quad (2.1)$$

and the Dirac operator D is commonly expressed as $i\mathcal{D}[A_\mu] - M$ with a flavour dependent mass matrix, $M = m_f \mathbb{1}$. \mathcal{S}_f is rendered gauge-invariant by the gauge field A_μ representing the gluons. This vector-valued field is an element of the algebra $\mathfrak{su}(N_c)$ and can be written in terms of the generators of $SU(N_c)$ T_a , $A_\mu = A_\mu^a T_a$, $a = 0, \dots, N_c^2 - 1$.² It mediates the interactions among fermions and couples to them via the covariant derivative

$$\mathcal{D}[A_\mu] = \gamma^\mu (\partial_\mu - igA_\mu). \quad (2.2)$$

¹ Dirac and colour indices will be denoted by Greek and Latin letters, respectively. f denotes flavour indices. Indices will be suppressed throughout when unambiguous. Summation over repeated indices is understood. Standard notation for contra- and covariant indices will be used.

² For a short overview over the group $SU(N_c)$ and the related algebra $\mathfrak{su}(N_c)$ see Appendix A.2.

g denotes the (bare) coupling constant of QCD and γ_μ are the Dirac matrices with the defining property

$$\{\gamma_\mu, \gamma_\nu\} \equiv \gamma_\mu \gamma_\nu + \gamma_\nu \gamma_\mu = 2g_{\mu\nu}. \quad (2.3)$$

$g_{\mu\nu}$ denotes the metric tensor.

Under local colour gauge transformations $\Omega(x) \in SU(N_c)$, ψ and A_μ transform according to

$$\psi(x) \rightarrow \Omega(x) \psi(x), \quad (2.4a)$$

$$A_\mu(x) \rightarrow \Omega(x) A_\mu(x) \Omega^\dagger(x) - \frac{i}{g} (\partial_\mu \Omega(x)) \Omega^\dagger(x). \quad (2.4b)$$

\not{D} transforms by construction in the adjoint representation, $\not{D}(x) \rightarrow \Omega(x) \not{D}(x) \Omega^\dagger(x)$. In general, $\bar{\psi}(z)$ and $\psi(y)$ are connected in a gauge-invariant way by the *Wilson line*

$$U_P(z, y) = \mathcal{P} \exp \left(-ig \int_P d^4x A_\mu(x) \right). \quad (2.5)$$

The integration is taken along some path P connecting z and y and \mathcal{P} denotes path ordering. The derivative \not{D} can be expressed in terms of the Wilson line,

$$\not{D} \psi(y) \propto [\psi(z) - U_P(z, y) \psi(y)] / |z - y|, \quad (2.6)$$

in the limit of $z \rightarrow y$. This will be useful on the lattice.

The dynamics of the gauge fields are described by

$$\mathcal{S}_g[A_\mu] = -\frac{1}{4} \int d^4x \mathcal{F}_{\mu\nu}^a(x) \mathcal{F}^{a,\mu\nu}(x), \quad (2.7)$$

with the field strength tensor

$$\mathcal{F}_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + gf_{abc} A_\mu^b A_\nu^c. \quad (2.8)$$

Here, f_{abc} denotes the structure constants of $\mathfrak{su}(N_c)$. The last term in $\mathcal{F}_{\mu\nu}^a$ shows the self-coupling of the A_μ : The gluons carry colour charge.³

For later studies of QCD it will be mandatory to go from Minkowski to Euclidean space.⁴ Formally, this can be done via a Wick rotation in time direction, $t \rightarrow i\tau$, which modifies the action of QCD in Minkowski space, $\mathcal{S}_{\text{QCD}} = \mathcal{S}_f + \mathcal{S}_g$, according to $\mathcal{S}_{\text{QCD}} \rightarrow i\mathcal{S}_{\text{QCD}}^E$, with

$$\mathcal{S}_{\text{QCD}}^E[\bar{\psi}, \psi, A_\mu] = \int d^3x \int d\tau \bar{\psi} \left(\not{D}^E[A_\mu] + M \right) \psi + \frac{1}{4} \mathcal{F}_{\mu\nu}^a \mathcal{F}_{\mu\nu}^a, \quad (2.9)$$

and $\not{D}^E = \gamma_\mu (\partial_\mu + igA_\mu)$.

Expectation values of an observable \mathcal{O} can conveniently be expressed in terms of the Feynman path integral:

$$\langle \mathcal{O} \rangle = \frac{1}{Z_{\text{QCD}}} \int \mathcal{D}[\bar{\psi}, \psi] \int \mathcal{D}[A_\mu] \mathcal{O}[\bar{\psi}, \psi, A_\mu] e^{-\mathcal{S}_{\text{QCD}}^E}. \quad (2.10)$$

3 A CP violating term $\sim \theta \tilde{\mathcal{F}}_{\mu\nu} \mathcal{F}^{\mu\nu}$ with dual field-strength tensor $\tilde{\mathcal{F}}_{\mu\nu} = \frac{1}{2} \epsilon_{\mu\nu\sigma\tau} \mathcal{F}^{\sigma\tau}$ will not be considered here.

4 See Appendix A.1 for implications on the metric.

$\mathcal{D}[\bar{\psi}, \psi]$ and $\mathcal{D}[A_\mu]$ denote integration over all possible field configurations and the normalisation is chosen to be

$$Z_{\text{QCD}} = \int \mathcal{D}[\bar{\psi}, \psi] \int \mathcal{D}[A_\mu] e^{-S_{\text{QCD}}^E} . \quad (2.11)$$

The Euclidean index will be suppressed in the following.

In quantum statistics, the canonical partition function Z of a system described by a Hamiltonian H at temperature T is known to be the quantum-mechanical trace over the Boltzmann weights. The system is fully described by the partition function, all thermodynamical quantities like the pressure or the free energy can be extracted from it. Z can also be expressed in terms of the Feynman path integral for some bosonic field ϕ with conjugate momenta π [Kapusta and Gale, 2006]

$$\begin{aligned} Z(T) &= \text{Tr} \left[e^{-\beta H} \right] \\ &= \int \mathcal{D}[\pi] \int_{\text{periodic}} \mathcal{D}[\phi] \exp \left(- \int_0^\beta \left(H(\phi, \pi) - \int d^3x \, i\pi \frac{\partial \phi}{\partial \tau} \right) \right) , \end{aligned} \quad (2.12)$$

with the inverse temperature $\beta = 1/T$. The periodic integration over ϕ means that $\phi(\vec{x}, 0) = \phi(\vec{x}, \beta)$, i.e. one has periodic boundary conditions (BC) in the compactified Euclidean time direction. Similar considerations in QCD show that Z_{QCD} (2.11) with compact time direction of extent β can be identified with a statistical system of temperature T . The fields have to fulfil (anti-)periodic BC in time direction,

$$A_\mu(\vec{x}, \tau) = A_\mu(\vec{x}, \tau + 1/T) , \quad (2.13a)$$

$$\psi(\vec{x}, \tau) = -\psi(\vec{x}, \tau + 1/T) , \quad (2.13b)$$

to ensure Bose-Einstein and Fermi-Dirac statistics, respectively. The compactness of the time direction has implications for the theory, since the energy levels ω_n ($n \in \mathbb{Z}$) of fermions and gluons can only take on integer values (*Matsubara frequencies*), reflecting the BC:

$$\begin{aligned} \omega_n[A_\mu] &= 2\pi n T , \\ \omega_n[\psi] &= (2n + 1)\pi T . \end{aligned} \quad (2.14)$$

A finite (fermionic) chemical potential μ enters the partition function as the Lagrange multiplier of the quark number operator Q ,

$$Z(T, \mu) = \text{Tr} \left[e^{-\beta(H - \mu Q)} \right] , \quad (2.15)$$

which is given by the temporal component of the conserved vector current:

$$Q = \bar{\psi} \gamma_0 \psi . \quad (2.16)$$

Thus, μ can be readily introduced into the action of QCD:

$$S_f[\bar{\psi}, \psi, A_\mu] \rightarrow S_f[\bar{\psi}, \psi, A_\mu, \mu] = \int d^3x \int d\tau \bar{\psi} (\not{D}[A_\mu] + M - \mu \gamma_0) \psi . \quad (2.17)$$

One might also interpret μ as a (constant) temporal $U(1)$ background field $\tilde{A}_0 = i\mu/\tilde{g}$ coupling to the fermion current, cf. (2.16). \tilde{A}_0 can then also be introduced in terms of the Wilson line (2.5) as

$$\tilde{U} = \exp(-i\epsilon\tilde{g}\tilde{A}_0) = \exp(-\epsilon\mu) \quad (2.18)$$

multiplied to all A_0 in \mathcal{D} , in the limit of $\epsilon \rightarrow 0$. With this procedure, μ appears as a phase $\exp(-\epsilon\mu)$ of the A_0 only. At finite temperature, it is then also possible to move μ completely into the boundary conditions of the fermions via a $U(1)$ gauge transformation:

$$\psi(\vec{x}, \tau) = -\exp(-\mu/T) \psi(\vec{x}, \tau + 1/T) . \quad (2.19)$$

2.2 SYMMETRIES OF QCD

A theory under consideration, i.e. its Hamiltonian H or its partition function Z , can be invariant under certain symmetry transformations. On the one hand, this symmetry can be broken spontaneously if the ground state of the system is not invariant. But on the other hand, it is also possible that the symmetry is broken explicitly, e.g. if parameters exceed a certain threshold.

QCD has numerous symmetries, in particular it is manifestly Lorentz and CPT invariant and by construction invariant under $SU(N_c)$ symmetry transformations. In the following, those symmetries important for upcoming considerations will be introduced.

2.2.1 Chiral Symmetry

The left and right components of ψ are defined via the projectors $P_{R,L} = \frac{1}{2}(1 \pm \gamma_5)$ as $\psi_{R,L} = P_{R,L}\psi$. In the massless case (*chiral limit*), $M \equiv 0$, ψ_R and ψ_L can be rotated independently in flavour space. The symmetry group then reads

$$\begin{aligned} G(N_f) &= U_L(N_f) \times U_R(N_f) \\ &= SU_L(N_f) \times SU_R(N_f) \times U_L(1) \times U_R(1) \\ &\simeq (SU_L(N_f) \times SU_R(N_f))/SU_V(N_f) \times SU_V(N_f) \times U_A(1) \times U_V(1) . \end{aligned} \quad (2.20)$$

It can be summarised in the fact that the Dirac operator anti-commutes with γ_5 ,

$$\{D, \gamma_5\} = 0 . \quad (2.21)$$

These chiral transformations are typically transformed into vector and axial vector transformations which treat the chiral components equally (V) or oppositely (A):⁵

$$\psi(x) \rightarrow \exp(i\theta^a T_a) \psi(x) \quad (SU_V(N_f)) , \quad (2.22a)$$

$$\psi(x) \rightarrow \exp(i\gamma_5 \theta^a T_a) \psi(x) \quad ((SU_L(N_f) \times SU_R(N_f))/SU_V(N_f)) , \quad (2.22b)$$

$$\psi(x) \rightarrow \exp(i\theta) \psi(x) \quad (U_V(1)) , \quad (2.22c)$$

$$\psi(x) \rightarrow \exp(i\gamma_5 \theta) \psi(x) \quad (U_A(1)) . \quad (2.22d)$$

⁵ Note that (2.22b) does not form a group, hence the notation $(SU_L(N_f) \times SU_R(N_f))/SU_V(N_f)$.

However, only the vector symmetries would be manifest in nature since both axial symmetries are broken. The $U_A(1)$ is broken during quantisation since the measure in the partition function is not invariant under this symmetry. This is related to the non-trivial topology of the QCD vacuum [Creutz, 2011] and is referred to as *axial anomaly* or *gauge anomaly*. It can be shown that the axial current connected to this symmetry, $J_\mu^A = \bar{\psi}\gamma_\mu\gamma_5\psi$, is not conserved (*Adler-Bell-Jackiw anomaly*) [Hooft, 1976]:

$$\partial_\mu J_\mu^A = \frac{g^2 N_f}{16\pi^2} \text{Tr} \tilde{\mathcal{F}}_{\mu\nu} \mathcal{F}_{\mu\nu} . \quad (2.23)$$

In two limits this anomaly vanishes, $N_c = \infty$ and $N_f = 0$. Nevertheless, even in the broken scenario there is a remnant $\mathbb{Z}(N_f)_A$ symmetry [Pisarski and Wilczek, 1984].

$SU_L(N_f) \times SU_R(N_f)$ is broken spontaneously by the QCD vacuum, manifest in a non-vanishing chiral condensate $\langle \bar{\psi}\psi \rangle$ (there is no anomaly since $\partial_\mu J_{\mu,a}^A \sim \text{Tr} T_a = 0$). The symmetry breaking pattern reads

$$SU_L(N_f) \times SU_R(N_f) \rightarrow SU_V(N_f) .$$

The spontaneous breaking of chiral symmetry implies massless bosonic field excitations for each broken generator T_a (*Goldstone theorem*). Its order parameter is the *chiral condensate*

$$\langle \bar{\psi}\psi \rangle = \frac{T}{V} \frac{\partial \ln Z}{\partial m} . \quad (2.24)$$

A mass term $m\bar{\psi}\psi$ explicitly breaks chiral symmetry. Nonetheless, it is approximately realised for $N_f = 2$ since $m_l \equiv m_u \approx m_d \ll m_s$. The lightest observed particles, the pions, can be identified as the would-be Goldstone bosons. One can treat the masses as perturbations to the chiral limit in the so-called chiral perturbation theory (χ PT). Here, one builds an effective theory around fields $\Sigma \in SU(2)$ related to the (would-be) Goldstone particles via

$$\Sigma(x) = \exp(2i\pi_a(x)\tau_a/f) . \quad (2.25)$$

τ_a are the Pauli matrices, related to the generators of $SU(2)$. The constant f can be identified with the pion decay constant. With this field at hand, one can write down an effective Lagrangian that has all symmetries one wishes to have realised in order to mirror “real” QCD. A famous result in leading order is the *Gell-Mann-Oakes-Renner* relation [Gell-Mann et al., 1968]

$$f^2 m_\pi^2 = -\langle \bar{\psi}\psi \rangle M_r + \mathcal{O}(m^2) . \quad (2.26)$$

This relates the renormalised quark mass M_r and the chiral condensate $\langle \bar{\psi}\psi \rangle$ to the physical quantities f and m_π which take on approximate values of 93 MeV and 135 MeV, respectively. χ PT is applicable up to around 300 MeV. For a review see [Bernard and Meissner, 2007].

2.2.2 Centre Symmetry

In the limit of $m_f \rightarrow \infty$, quarks become static and *Pure Gauge Theory* can be studied. At finite temperature, gauge transformations Ω have to fulfil periodic boundary conditions, cf.

(2.13). Furthermore, the transformation of the gauge field (2.4b) does not change if Ω picks up a constant phase z ,

$$\Omega(\vec{x}, \tau) = z \Omega(\vec{x}, \tau + 1/T) . \quad (2.27)$$

z is an element of the centre of $SU(N_c)$, $\mathcal{Z}(N_c) = \left\{ z \in SU(N_c) \mid z = \exp(i\frac{2\pi k}{N_c}) , k \in \mathbb{N} \right\}$, and the theory exhibits centre symmetry. As it is infinitely heavy, a quark will propagate in time only, which is described by the trace of a closed Wilson Loop in temporal direction, the *Polyakov Loop*:

$$L(\vec{x}) = \text{Tr} \mathcal{P} \exp \left(-ig \int_0^{1/T} d\tau A_0(\vec{x}) \right) . \quad (2.28)$$

As the loop winds around the temporal direction, it picks up a phase from $\mathcal{Z}(N_c)$. A particular sector can be identified from the phase of L : $L = |L|e^{i\phi}$. If centre symmetry is realised, L will cycle through the different $\mathcal{Z}(N_c)$ sectors and the phases average to zero. Thus, also its expectation value $\langle L \rangle$ vanishes. Then again, if the symmetry is spontaneously broken $\langle L \rangle \neq 0$.

Physically, $\langle L \rangle$ corresponds to the free energy difference of a system with and without a static quark Δ_F :

$$\langle L \rangle = \frac{1}{Z} \int \mathcal{D}[A_\mu] L e^{-S_g} = \exp(-\Delta_F/T) . \quad (2.29)$$

The quark carries colour charge, so if the theory is confining it costs an infinite amount of energy to put it into the system, $\Delta_F = \infty$, and $\langle L \rangle = 0$. Correspondingly, $\langle L \rangle \neq 0$ signals deconfinement and this can be expected as $T \rightarrow \infty$. As will be elucidated below, this *deconfinement transition* actually takes place at some finite temperature T_c . Centre symmetry is broken explicitly by quarks which pick up a factor of z under (2.27). In fact, the fermion determinant selects the real centre sector. Hence, there is no true phase transition in this case. Nevertheless, $\langle L \rangle$ can still serve as an indicator for deconfinement.

2.2.3 Roberge-Weiss Symmetry

The grand canonical partition function (2.15) can be related to the canonical one via the fugacity expansion:

$$Z(T, \mu) = \sum_n \left(e^{\beta\mu} \right)^n Z_{\text{can}}(T) , \quad (2.30)$$

where the quark number n runs from positive to negative integer values. This implies

$$Z(T, \mu) = Z(T, -\mu) . \quad (2.31)$$

In case of purely imaginary chemical potential, $\mu = i\theta/\beta = i\theta T$, Z has a periodicity of 2π . Centre symmetry is broken explicitly by ψ at $\mu = 0$ as was mentioned above. However, for a transformation as in (2.27), the phase picked up by ψ can be accounted for by an appropriate shift in μ : The transformations

$$A_\mu \longrightarrow \Omega A_\mu \Omega^\dagger - \frac{i}{g} (\partial_\mu \Omega) \Omega^\dagger \quad (2.32a)$$

$$\psi \longrightarrow \Omega \psi \quad (2.32b)$$

$$\theta \longrightarrow \theta + \frac{2\pi k}{N_c} , k \in \mathbb{N}, \quad (2.32c)$$

leave the partition function invariant and unveil a periodicity of $2\pi/N_c$:

$$Z(\theta) = Z(\theta + 2\pi k/N_c), k \in \mathbb{N}. \quad (2.33)$$

This can easily be understood if μ is considered as a boundary condition, cf. (2.19), and is called extended centre symmetry or simply Roberge-Weiss (RW) symmetry [Roberge and Weiss, 1986]. An order parameter can be defined by introducing the *modified Polyakov loop*

$$\hat{L} = Le^{i\theta} = |\hat{L}|e^{i\varphi}, \quad (2.34)$$

which is invariant under (2.32). Similar to the Polyakov Loop (2.28) its phase φ again indicates the $\mathbb{Z}(N_c)$ sector of the system. Below T_c , the phase averages to zero, whereas above T_c it takes on the values

$$\langle \varphi \rangle = n(2\pi/N_c), n = 0, 1, \dots, N_c - 1. \quad (2.35)$$

2.3 THE TWISTED MASS FORMULATION OF QCD

The following considerations will be carried out for two flavours of mass-degenerated fermions, i.e. ψ is a doublet (u, d) .

The mass term $m\bar{\psi}\psi$ commonly used for the continuum action of QCD is not of the most general form. In fact, a chirally twisted mass term $i\hat{\mu}\bar{\psi}\gamma_5\tau_3\psi$ may be added [Frezzotti and Rossi, 2004], with $\hat{\mu}$ the *twisted mass*.⁶ Note that this mass term breaks parity and flavour symmetry. The fermionic part of the action of QCD, (2.9), then reads:

$$\mathcal{S}_{\text{QCD}}[\bar{\chi}, \chi, A_\mu] = \int d^3x \int d\tau \bar{\chi} (\not{D}[A_\mu] + m + i\hat{\mu}\gamma_5\tau_3) \chi. \quad (2.36)$$

The fields in this *twisted basis*, χ , and those in the physical basis, ψ , are related by axial transformations

$$\chi \rightarrow \psi = e^{-i\omega\gamma_5\tau_3/2}\chi, \quad (2.37)$$

$$\bar{\chi} \rightarrow \bar{\psi} = \bar{\chi} e^{-i\omega\gamma_5\tau_3/2}. \quad (2.38)$$

If one identifies the *twist angle* with

$$\omega = \arctan(\hat{\mu}/m), \quad (2.39)$$

the mass term again takes the standard form $M\bar{\psi}\psi$ with the polar mass $M = \sqrt{\hat{\mu}^2 + m^2}$ and one has

$$m = M \cos(\omega) \text{ and } \hat{\mu} = M \sin(\omega). \quad (2.40)$$

Thus, (2.37) relates both formulations of QCD in a unique way, the mass term can always be brought into the standard form. Therefore, both share the same symmetries and describe the same physics. In fact, the parity and flavour symmetry breaking induced by the twisted mass

⁶ One may also add a term $\propto \bar{\psi}\tau_3\psi$ resulting in a difference in the masses of the quarks or a CP violating term $\propto \bar{\psi}\gamma_5\psi$, see [Creutz, 2011] for a general discussion.

to the mass term is irrelevant for the remaining terms in the action as renormalisability ensures that they are not affected (*accidental symmetry*) [Shindler, 2008].

Yet, the symmetries have different transformation rules in the twisted basis. For example, the chiral transformations (2.22) take on the form [Shindler, 2008]:

$$\chi(x) \rightarrow e^{-i\omega\gamma_5\tau_3/2} \exp(i\theta^a \frac{\tau_a}{2}) e^{i\omega\gamma_5\tau_3/2} \chi(x) \quad (SU_V(N_f)) , \quad (2.41a)$$

$$\chi(x) \rightarrow e^{-i\omega\gamma_5\tau_3/2} \exp(i\gamma_5\theta^a \frac{\tau_a}{2}) e^{i\omega\gamma_5\tau_3/2} \chi(x) \quad (SU_L(N_f) \times SU_R(N_f)) / SU_V(N_f) , \quad (2.41b)$$

$$\chi(x) \rightarrow \exp(i\theta) \chi(x) \quad (U_V(1)) , \quad (2.41c)$$

$$\chi(x) \rightarrow \exp(i\gamma_5\theta) \chi(x) \quad (U_A(1)) . \quad (2.41d)$$

While the (pseudo-)scalar transformations, (2.41d) and (2.41c), and those vector-like proportional to τ_3 do not change⁷, the vector-like transformations orthogonal to τ_3 will in general depend on the twist angle.

A special case is given when the mass is solely given by the twisted mass, $m = 0$. Then, $\omega = \pi/2$ and the system is said to be tuned to *Maximal Twist* [Frezzotti and Rossi, 2004]. In this case, the non-trivial chiral transformations (2.41) reveal a particular interesting property [Shindler, 2008]:

$$\chi(x) \rightarrow \exp(\pm i\theta^{1,2} \tau_{2,1} \gamma_5/2) \chi(x) \quad (SU_V(N_f)) , \quad (2.42a)$$

$$\chi(x) \rightarrow \exp(\pm i\theta^{1,2} \tau_{2,1}/2) \chi(x) \quad ((SU_L(N_f) \times SU_R(N_f)) / SU_V(N_f)) . \quad (2.42b)$$

That is, their role is reversed with respect to the standard form of the transformation, caused by the different basis chosen.

2.4 THE RUNNING COUPLING OF QCD

Beyond tree level, observables require renormalisation to remove divergences. This can be achieved by introducing some momentum cutoff C . The regularised theory is then expressed in terms of physical quantities and the cutoff is finally removed. See e.g. [Creutz, 2011] for an introduction on renormalisation group techniques. The bare parameters of the theory now depend on C and have to be varied as C is varied, keeping physical observables fixed. QCD has only few bare parameters: The fermion masses m_f and the bare coupling g . Their “running” is governed by *Renormalisation group (RG) equations* and commonly expressed via the functions:

$$\beta(g) \equiv C \frac{dg}{dC}; \quad \gamma_f(m_f) \equiv C \frac{dm_f}{dC} . \quad (2.43)$$

To leading order in QCD, $\beta(g)$ is known to be [Gross and Wilczek, 1973a,b; Politzer, 1973]

$$\beta(g) = -\frac{g^3}{(4\pi)^2} \left[\frac{11}{3} N_c - \frac{2}{3} N_f \right] \equiv -\frac{g^3}{(4\pi)^2} \beta_0 . \quad (2.44)$$

⁷ For $a = 3$ the components of the vector-like transformations simply commute.




Figure 2: Sketch of lattice discretisation at lattice spacing a . k, n, m, l denote lattice sites, ψ a fermion field and U a gauge link, respectively. The plaquette $P_{\mu\nu}(m)$ and the rectangle product of links $R_{\mu\nu}(l)$ are also shown. See text for details.

For $N_f < \frac{11}{2}N_c$ (which is valid in physical QCD), this means that g^2 goes to zero as the momentum Q^2 involved in the process goes to infinity (*asymptotic freedom*):

$$g^2(Q^2) = \frac{g_C^2}{1 + \frac{g_C^2}{(4\pi)^2} \beta_0 \log(Q^2/C)} \xrightarrow{Q^2 \rightarrow \infty} 0. \quad (2.45)$$

Since the renormalisation point $g_C^2 = g^2(C)$ is somewhat arbitrary, it is conveniently replaced by a mass scale Λ_{QCD} that satisfies $1 = g^2(\beta_0/8\pi^2) \log(C/\Lambda_{\text{QCD}})$. This explicitly represents a dimensioned quantity introduced in the theory via the RG, even though there is only a dimensionless quantity g^2 to start with (for $m_f = 0$). This goes by the name of *dimensional transmutation*; conformal symmetry is anomalously broken. The strong coupling constant $\alpha_s = \frac{g^2}{4\pi}$ can be expressed as

$$\alpha_s(Q^2) = \frac{2\pi}{\beta_0 \log(Q^2/\Lambda_{\text{QCD}})}. \quad (2.46)$$

Similar relations hold for γ , with Λ_{QCD} replaced by a constant M_f denoting the renormalised quark mass. Λ_{QCD} can be estimated to be $\Lambda_{\text{QCD}} \approx 200$ MeV, and (2.46) explicitly states that Λ_{QCD} is the scale at which α_s is large. In this regime QCD is non-perturbative.

2.5 QCD ON THE LATTICE

At the energy scales of hadronic matter, traditional perturbation theory fails in QCD because the running coupling is strong: $\alpha_s(\Lambda_{\text{QCD}}) \approx \mathcal{O}(1)$. The only method to study QCD in this region from first principles is the lattice approach, which will be introduced in the following.

The lattice is a discretised version of Euclidean space time on a four dimensional hypercube, with temporal and spatial extents N_τ and N_σ , lattice spacing a and total volume $a^4 V_{\text{tot}} = a^4 N_\sigma^3 N_\tau$ (see fig. 2). Each space time point x_μ is characterised by four integer coordinates, which allows a unique identification via a super index n , $n = 0, \dots, V_{\text{tot}} - 1$. The discrete step size a acts as a momentum cutoff, all components of momenta are restricted to lie within the Brillouin zone $(-\pi/a, \pi/a]$.

Furthermore, the finite lattice renders the path integral well defined, as the integration measures are given as

$$\mathcal{D}[\bar{\psi}, \psi] = \prod_{n,a,f,\alpha} d\bar{\psi}_{a,\alpha,f}(n) d\psi_{a,\alpha,f}(n) , \quad (2.47)$$

$$\mathcal{D}[U] = \prod_{n,a} d[U_\mu](n) . \quad (2.48)$$

$d[U_\mu]$ denotes the *Haar measure* of $SU(N_c)$.

The action of QCD will be modified to give a version of QCD afflicted with discretisation artifacts, *Lattice QCD* (LQCD):

$$\mathcal{S}_{\text{LQCD}} = \mathcal{S}_{\text{QCD}} + a\mathcal{S}_1 + a^2\mathcal{S}_2 + \dots . \quad (2.49)$$

There is some freedom in the actual choice, as it has to coincide with \mathcal{S}_{QCD} in the continuum limit, $a \rightarrow 0$, only. This limit is governed by the RG equations (2.43) with cutoff $C \sim 1/a$. To leading order in the coupling a goes like

$$a(g^2) = \frac{1}{\Lambda_{\text{QCD}}} e^{-(4\pi)^2/2\beta_0 g^2} \xrightarrow{g^2 \rightarrow 0} 0 \quad (2.50)$$

and similar for $a(m)$. The continuum limit thus corresponds to tuning the bare coupling g^2 to zero, keeping some physical quantity $\mathcal{O}_{\text{phys}} = 1/a \mathcal{O}_{\text{lat}}$ fixed. $\mathcal{S}_{\text{LQCD}}$ may also be viewed as a spin system which is tuned to criticality, since some correlation length ξ is diverging: $\xi \equiv \mathcal{O}_{\text{lat}}^{-1} \rightarrow \infty$. To keep $\mathcal{O}_{\text{phys}}$ fixed, the *thermodynamic limit*, $V_{\text{tot}} \rightarrow \infty$, has to be taken, too.

On the lattice, all quantities are expressed in units of a and thus are dimensionless. To relate a simulation result to the corresponding quantity in physical units, the value of a has to be determined. For example, this may be done by measuring the mass of pion and rho meson, $am_\pi(m_i)$ and $am_\rho(m_i)$, for several values m_i of the bare quark mass. Then, extrapolating the mass ratio to its physical value yields a bare quark mass \tilde{m} where the physical situation is realised in the simulation. Now, the extrapolated $am_\rho(\tilde{m})$ equals am_ρ^{exp} and a is given in physical units by $am_\rho(\tilde{m})/am_\rho^{\text{exp}}$. Note that these simulations are performed at fixed β as $a = a(\beta)$. A similar procedure has to be applied to each dimensioned parameter in the theory. Afterwards, other quantities are related to the physical world in a clear way. In practice, most LQCD simulations are performed at $m_\pi \neq m_\pi^{\text{phys}}$ because of the increase of numerical costs when going towards the physical point.

In order to render \mathcal{S}_{QCD} gauge invariant, gauge fields A_μ were introduced within the covariant derivative (2.2). Since on the lattice, n and its neighbour in μ direction ($n + \hat{\mu}$) can no longer be connected infinitesimally, gauge invariance is imposed via the lattice analogue of Wilson lines, called *links*:⁸

$$U_\mu(n) = e^{-iaA_\mu(n)} . \quad (2.51)$$

Whereas the A_μ are elements of the algebra of the gauge group, U_μ is an element of the group itself.

⁸ To account for the explicit factor of $\frac{1}{g^2}$ in (2.55), A_μ has been rescaled $A_\mu \rightarrow \frac{1}{g}A_\mu$.

Links transform in the adjoint representation,

$$U_\mu(n) \rightarrow \Omega(n) U_\mu(n) \Omega^\dagger(n + \hat{\mu}) , \quad (2.52)$$

and the simplest gauge-invariant object is the trace of the so-called *plaquette* (see fig. 2)

$$P_{\mu\nu}(n) = U_\mu(n) U_\nu(n + \hat{\mu}) U_\mu^\dagger(n + \hat{\nu}) U_\nu^\dagger(n) . \quad (2.53)$$

As a is taken to zero, $P_{\mu\nu}$ can be expressed in terms of the field strength tensor:

$$P_{\mu\nu}(n) = \exp\left(ia^2 \mathcal{F}_{\mu\nu}(n) + \mathcal{O}(a^3)\right) . \quad (2.54)$$

Thus, the Wilson gauge action [Wilson, 1974],

$$\mathcal{S}_{\text{gauge}} = \frac{2N_c}{g^2} \sum_n \sum_{\mu, \nu > \mu} \{1 - \text{Re Tr}_c(P_{\mu\nu}(n))\} , \quad (2.55)$$

approximates \mathcal{S}_g up to $\mathcal{O}(a^2)$ effects.⁹ A lattice coupling $\beta = \frac{2N_c}{g^2}$ is conveniently introduced.

Via the Symanzik improvement program, one can reduce the discretisation errors of $\mathcal{S}_{\text{gauge}}$ to be of $\mathcal{O}(a^4)$ and $\mathcal{O}(a^2 g^2)$, yielding the improved gauge action [Weisz, 1983; Weisz and Wohlert, 1984]:

$$\mathcal{S}_{\text{tIsym}} = \beta \sum_n \left(c_0 \sum_{\mu, \nu > \mu} \{1 - \text{Re Tr}_c(P_{\mu\nu}(n))\} + c_1 \sum_{\mu, \nu} \{1 - \text{Re Tr}_c(R_{\mu\nu}(n))\} \right) . \quad (2.56)$$

Here, $R_{\mu\nu}(n)$ denotes the path-ordered rectangle product of links (see fig. 2) and the parameters are $c_0 = 1 - 8c_1$ and $c_1 = 1/12$. The unimproved gauge action is regained setting c_1 to zero.

2.5.1 Wilson Fermions

For the discretisation of the fermions, one flavour will be considered for the moment. On the lattice, \mathcal{S}_f takes the general form

$$\mathcal{S}_f[\bar{\psi}, \psi, U] = a^4 \sum_{n, m} \bar{\psi}_{a, \alpha}(n) D(n, m)_{a, b, \alpha, \beta}[U] \psi_{b, \beta}(m) , \quad (2.57)$$

where $D(n, m)_{a, b, \alpha, \beta}[U]$ is the discretised Dirac operator (*fermion matrix*).

The \mathcal{D} can be naturally expressed in terms of the Wilson line, as was mentioned above. However, discretising it in a straight-forward manner runs into problems already in the free case ($U \equiv \mathbb{1}$). In momentum space, D then reads:

$$D_{\text{naive}} = m\mathbb{1} + \frac{i}{a} \sum_{\mu} \gamma_{\mu} \sin(p_{\mu} a) . \quad (2.58)$$

On the lattice, momenta are replaced by sinusoidal functions, $p_{\mu} \rightarrow \sin(p_{\mu} a)/a$, the propagator has unphysical poles at the edges of the Brillouin zone, called *doublers*. In order to avoid them,

⁹ The trace over colour indices Tr_c is normalised by $1/N_c$.

Wilson proposed to add a term of the form $-a\partial_\mu\partial_\mu$ to the Dirac operator. This term vanishes in the continuum limit, but produces a momentum dependent mass $\propto \frac{2}{a}\cos(p_\mu a)$ at nonzero a , rendering only the doublers heavy. It is convenient to introduce a new mass parameter κ ,

$$\kappa = (2(am + 4))^{-1}, \quad (2.59)$$

and rescale $\psi \rightarrow \sqrt{(2a\kappa)^{-1}} \psi$ to obtain the Wilson fermion matrix¹⁰

$$\begin{aligned} D_{\text{Wilson}}(n, m)_{\alpha, \beta, a, b} &= \delta_{nm} \delta_{\alpha\beta} \delta_{ab} - \kappa \sum_{\pm\mu} (1 - \gamma_\mu)_{\alpha\beta} U_\mu(n)_{ab} \delta_{n+\hat{\mu}, m} \\ &\equiv M_{\text{diag}} + \mathcal{D}. \end{aligned} \quad (2.60)$$

Shorthand notation $\gamma_{-\mu} = -\gamma_\mu$ and $U_{-\mu}(n) = U_\mu^\dagger(n - \vec{\mu})$ has been used.

A useful property of D_{Wilson} is that it is hermitian up to factors of γ_5 (γ_5 -hermiticity):

$$D_{\text{Wilson}}^\dagger(n, m) = \gamma_5 D_{\text{Wilson}}(m, n) \gamma_5. \quad (2.61)$$

A significant drawback of Wilson's approach is that it breaks chiral symmetry explicitly. This is in accordance with the *no-go theorem* by Nielsen and Ninomiya [Nielsen and Ninomiya, 1981a,b,c], which states that it is not possible to construct a chirally symmetric, local fermion action on the lattice that has no doublers and the correct continuum limit.

As it is no longer protected by chiral symmetry, the mass m experiences additive renormalisation. The renormalised quark mass M_f vanishes at some value $\kappa_c(\beta)$ which coincides with the naive chiral value $\kappa = 1/8$ only for $\beta \rightarrow \infty$. In the strong coupling region at small β and at non-zero a , there is an unphysical phase where the parity-flavour symmetry is spontaneously broken (*Aoki phase*). This phase meets the chiral critical line in a cusp point β_{cusp} . As κ is increased beyond $\kappa_c(\beta)$, a region with doubler physics is eventually encountered. For a general discussion of the phase structure of Wilson fermions see [Creutz, 2011]. At finite temperature (see below), increasing β means increasing T , i.e. at fixed κ , one sees a deconfining transition/crossover at some value of β . As κ is varied, these points lie on the thermal line κ_t , which meets the β -axis at the value of the quenched limit β_{qu} [Ilgenfritz et al., 2009]. For a visualisation, see Figure 3 at $\hat{\mu} = 0$.

For the evaluation of many observables a crucial ingredient is the fermionic two-point function. It is related to the inverse fermion matrix, the quark propagator D^{-1} , via

$$\langle \bar{\psi}(n)_{a\alpha} \psi(m)_{b\beta} \rangle = a^{-4} D^{-1}(n, m)_{ab\alpha\beta}. \quad (2.62)$$

This follows from *Wick's Theorem* (A.14). For example, in a theory with N_f mass-degenerated flavours the chiral condensate $\langle \bar{\psi}\psi \rangle$ defined in (2.24) can be determined as

$$\langle \bar{\psi}\psi \rangle \stackrel{(2.62)}{=} N_f \text{Tr} D^{-1}. \quad (2.63)$$

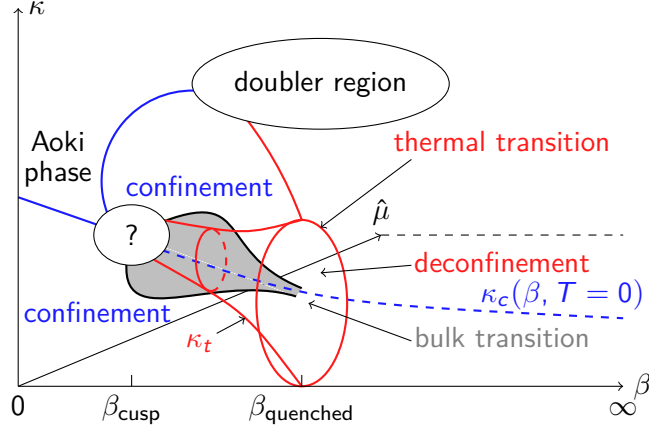


Figure 3: Phase diagram at finite temperature for twisted mass Wilson fermions following [Ilgenfritz et al., 2009]. The phase diagram for pure Wilson fermions can be read off at $\hat{\mu} = 0$.

2.5.2 Twisted Mass Wilson Fermions

On the lattice, after performing a discretisation of the twisted mass formulation of QCD (2.36) similar to Wilson fermions, the fermion matrix of the so-called *twisted mass* Wilson fermions reads:

$$D_{\text{tm}}^{\pm} = (1 \pm 2ia\kappa\hat{\mu}\gamma_5) \delta_{nm}\delta_{\alpha\beta}\delta_{ab} + \not{D} = M_{\text{diag}}^{\pm} + \not{D}.$$

The sign in the diagonal mass matrix M_{diag}^{\pm} corresponds to up and down flavour. For an extensive review see [Shindler, 2008].

As before, the Wilson term explicitly breaks chiral symmetry and the untwisted bare quark mass experiences additive and multiplicative renormalisation. The twisted mass term is renormalised only multiplicatively because it is protected by parity. Thus, the renormalised mass can be expressed as

$$M_r = \sqrt{(Z_m(m - m_c))^2 + (Z_{\hat{\mu}}\hat{\mu})^2}, \quad (2.64)$$

with renormalisation constants Z_m and $Z_{\hat{\mu}}$. m_c is the value where the renormalised untwisted mass vanishes (converted to κ_c in the Wilson formulation, (2.59)). Thus, tuning to maximal twist on the lattice means to tune to κ_c . Here, twisted mass fermions have the enormous advantage that at maximal twist, $\mathcal{O}(a)$ discretisation errors in (2.49) vanish, too [Frezzotti and Rossi, 2004]. Maximal twist can be achieved by tuning a parity-odd operator to zero. Commonly, the *partially conserved axial current* mass m_{PCAC} is chosen. It is defined as [Shindler, 2008]

$$m_{\text{PCAC}} = \frac{\sum_{\mathbf{n}} \langle \partial_0 A_0^a(\mathbf{n}, t) P^a(0) \rangle}{2 \sum_{\mathbf{n}} \langle P^a(\mathbf{n}, t) P^a(0) \rangle}, \quad a = 1, 2. \quad (2.65)$$

In here the A_{μ}^a and P^a are the axial vector current and the pseudo scalar density, respectively. They are defined as

$$A_{\mu}^a(n) = \bar{\chi}(n)\gamma_{\mu}\gamma_5\frac{\tau_a}{2}\chi(n) \quad P^a(n) = \chi(n)\gamma_5\frac{\tau_a}{2}\chi(n). \quad (2.66)$$

10 The Wilson parameter r will be set to 1 throughout the text.

It was shown above that the twisted mass formulation is equivalent to the standard form of QCD in the continuum. On the lattice, this is no longer true because the Wilson term breaks the symmetry of the action under the axial transformations (2.37). Both formulations are equal in the continuum limit only. This means that at finite lattice spacing, parity and flavour symmetry are broken for twisted mass Wilson fermions. However, parity can for example be recovered as a combined symmetry together with a flavour flip [Shindler, 2008]:

$$\psi(x_0, \mathbf{n}) \rightarrow \gamma_0 \psi(x_0, -\mathbf{n}) \Rightarrow \chi(x_0, \mathbf{n}) \rightarrow i\gamma_0 \tau_{1,2} \chi(x_0, -\mathbf{n}) . \quad (2.67)$$

This can have implications for the results obtained from simulations. For example, a scalar particle can generally decay to a pseudo-scalar one as this is no longer forbidden by a symmetry [Shindler, 2008].

The change of basis (2.37) causes various changes in the form of symmetries and observables, for example for the chiral condensate $\langle \bar{\psi} \psi \rangle$, which is given in the physical basis in (2.63). In the twisted basis, the expression changes:

$$\langle \bar{\psi} \psi \rangle \xrightarrow{(2.37)} \langle \bar{\chi} (\cos \omega + i\gamma_5 \tau_3 \sin \omega) \chi \rangle \quad (2.68)$$

and at maximal twist one has:

$$\langle \bar{\chi} i\gamma_5 \chi \rangle = \langle \bar{u} i\gamma_5 u \rangle - \langle \bar{d} i\gamma_5 d \rangle = i \text{Tr}(\gamma_5 (D_u^{-1} - D_d^{-1})) = -2 \text{Im} \text{Tr}(\gamma_5 D^{-1}) . \quad (2.69)$$

In the last step, up and down flavours are related by (2.61). It is numerically convenient to employ the *one-end trick* [Boucaud et al., 2008]. This relies on the identity:

$$D_u - D_d = -4i\kappa\hat{\mu}\gamma_5 \Leftrightarrow D_u^{-1} - D_d^{-1} = -4i\kappa\hat{\mu}\gamma_5 (D_u^{-1})^\dagger D_u^{-1} , \quad (2.70)$$

which yields:

$$\langle \bar{\psi} \psi \rangle = -4\kappa\hat{\mu} |D_u^{-1}|^2 \quad (2.71)$$

The phase diagram of twisted mass fermions (cf. [Creutz, 2011; Ilgenfritz et al., 2009] and references therein) includes that of pure Wilson fermions as the latter can be reobtained at zero twist, $D_{\text{Wilson}} \equiv D_{\text{tm}}(\hat{\mu} = 0)$. At nonzero $\hat{\mu}$ there is an additional unphysical phase, the *Sharpe-Singleton* phase that includes the chiral line. It extends into the $\hat{\mu}$ -plane and is perpendicular to the Aoki-phase plane. This restricts simulations at maximal twist from going to arbitrary small $\hat{\mu}$ without encountering this phase. At finite temperature β_c is constant on lines of constant M_r . Thus, the thermal line κ_t of pure Wilson fermions becomes an ellipse wrapping around κ_c . As β_c or T_c is decreased, the ellipses shrink since also M_r decreases with critical temperature, giving a conical shape to the transition surface. See Figure 3 for a sketch of the phase diagram.

2.5.3 Alternative Fermion Discretisations

Apart from (twisted mass) Wilson fermions, there are various different kinds of discretisations, which will be briefly commented on now. For more details we refer to the lattice-related books mentioned at the beginning of this chapter.

The $\mathcal{O}(a)$ effects of Wilson fermions can also be removed by the Symanzik improvement program, which adds a plaquette-like term (the *clover* term) to the action [Sheikholeslami and Wohlert, 1985]. Its coefficient is best tuned non-perturbatively.

In addition, there are different ways to discretise \mathcal{S}_f . Staggered fermions [Kogut and Susskind, 1975], which are obtained by diagonalising the Dirac matrices, are numerically cheaper than Wilson fermions. With regard to the naive fermion discretisation, in this formulation the lattice spacing is effectively doubled or, equivalently, the momentum region is halved removing the origins of the unphysical poles in the propagator. This *staggering* reduces the number of doublers to four (which are called *tastes*) and also preserves a remnant of chiral symmetry. Nevertheless, this discretisation is subject of ongoing debate because of the rooting-procedure which is used to remove the remaining tastes [Creutz, 2007; Sharpe, 2006].

Another approach to chiral fermions on the lattice is undertaken via a lattice version of chiral symmetry (2.21), the *Ginsparg-Wilson relation*:

$$\{D, \gamma_5\} = aD\gamma_5D, \quad (2.72)$$

which states that the Dirac operator anti commutes with γ_5 only up to lattice effects. Overlap and domain wall fermions fulfil this relation and thus principally allow for chiral studies [Ginsparg and Wilson, 1982; Kaplan, 1992; Neuberger, 1998a,b; Shamir, 1993]. Nonetheless, apart from algorithmic problems, they come with an increase of numerical cost which is currently at least one order of magnitude and thus, they are hardly used.

2.5.4 Finite T and μ

On the hypercube, the system's extents are finite. Regarding temperature, the same considerations as in Section 2.1 hold and T can again be identified with the inverse temporal extent:

$$T = 1 / (a(\beta)N_\tau). \quad (2.73)$$

At nonzero μ , a naive discretisation of the quark number operator as in (2.17) on the lattice turns out to give a divergent continuum limit, even for free quarks. The reason for this is that this discretisation violates gauge invariance [Hasenfratz and Karsch, 1983]. The correct way is to introduce a function $f(a\mu) = \exp(a\mu)$ and weight every temporal link in positive direction by it (and every temporal link in negative direction by the inverse, $f(-a\mu)$). This is in analogy of interpreting μ as a constant gauge field, cf. (2.18).

For Wilson fermions, the action including chemical potential reads:

$$\begin{aligned} D_{\text{Wilson}}(n, m) = & M_{\text{diag}} - \kappa \sum_{i=1}^3 \left\{ (1 - \gamma_i) U_{\pm i}(n) \delta_{n+\hat{i}, m} \right\} \\ & - \kappa \left\{ (1 - \gamma_0) e^{+a\mu} U_0(n) \delta_{m, n+\hat{0}} + (1 + \gamma_0) e^{-a\mu} U_0^\dagger(m) \delta_{m, n-\hat{0}} \right\}. \end{aligned} \quad (2.74)$$

At finite μ , D_{Wilson} is no longer γ_5 -hermitian, but fulfils:

$$D_{\text{Wilson}}^\dagger(\mu) = \gamma_5 D_{\text{Wilson}}(-\mu^*) \gamma_5. \quad (2.75)$$

In the spatial directions, periodic BC are imposed everywhere. In temporal direction, the gauge fields have periodic BC, while the fermions have antiperiodic ones (cf. (2.13)). The BC are conveniently shifted from the boundary of the lattice to each site via a gauge transformation [Jansen and Urbach, 2009].

The finiteness of the system may have severe effects on the simulation: Due to the BC, the system sees infinitely copies of itself. If the de-Broglie wavelength of the pion is shorter than the system extent, so-called *finite-size effects* (FEV) will appear. To avoid this in finite temperature studies, as a rule of thumb $aN_\sigma m_\pi \approx 3 - 4$ and $N_\sigma/N_\tau \approx 3$ should be realised. Opposed to that, zero temperature simulations usually have $N_\tau \approx 2N_\sigma$, yielding negligible temperature effects.

2.5.5 Spectroscopy on the Lattice

In this section a short introduction on the spectroscopic methods frequently used on the lattice is given. For more detailed overviews see for example the dedicated sections in [DeGrand and DeTar, 2006; Gattringer and Lang, 2010; Montvay and Münster, 1994].

In a system with a Euclidean metric two-point functions (*correlators*) between operators \mathcal{O}_1 and \mathcal{O}_2 at Euclidean times t_1 and t_2 have the general form

$$C \equiv \langle \mathcal{O}_1(t_1) \mathcal{O}_2(t_2) \rangle = \sum_n \langle \Omega | \mathcal{O}_2(0) | n \rangle \langle n | \mathcal{O}_1(0) | \Omega \rangle e^{-tE_n} . \quad (2.76)$$

In here, $|\Omega\rangle$ is the vacuum, $t = t_2 - t_1$ is the time difference and it is assumed that $t_2 > t_1$. Additionally, a complete set of eigenstates of the Hamiltonian H , $\{|n\rangle\}$, has been inserted and E_n is the energy of the n -th eigenstate. For large t the ground state will dominate the signal and can be extracted from the correlator. It is easy to show that summing over the spatial indices projects onto zero momentum, thus yielding the mass of the state under consideration.

In the context of this thesis, flavour-multiplet mesonic correlators for two mass-degenerated flavours will be of interest. Generally, such an operator at space-time point n can be written as

$$\mathcal{O}_m(n) = \bar{\psi}(n) \Gamma \tau^\pm \psi(n) , \quad (2.77)$$

with ladder operator $\tau^\pm = \tau_1 \pm i\tau_2$ acting in flavour space. Γ is an expression of γ matrices encoding certain quantum numbers. The correlator then reads

$$\begin{aligned} C_\Gamma \equiv \langle \bar{\mathcal{O}}(m) \mathcal{O}(n) \rangle &= \langle \bar{u}(m) \Gamma d(m) \bar{d}(n) \Gamma u(n) \rangle \\ &= -\text{Tr} \left[\Gamma (D^{-1})^{(u)}(n, m) \Gamma (D^{-1})^{(d)}(m, n) \right] . \end{aligned} \quad (2.78)$$

γ_5 -hermiticity (2.61) can be used to relate the propagators for up- and down flavour, yielding

$$C_\Gamma = -\text{Tr} \left[\Gamma (D^{-1})(n, m) \Gamma \gamma_5 ((D^{-1})(n, m))^\dagger \gamma_5 \right] . \quad (2.79)$$

Note that for flavour singlets, which are built from $\bar{\psi}(n) \Gamma \tau_3 \psi(n)$, the Wick contraction would yield terms which are proportional to $\text{Tr} \left[\Gamma (D^{-1})^{(u)}(n, n) \right]$. These *disconnected* parts are computationally demanding as they represent a particle propagating over the lattice, returning

state	J^{PC}	Γ	Particles
Scalar	0^{++}	$\mathbb{1}, \gamma_0$	$f_0, a_0, K_0^* \dots$
Pseudo-scalar	0^{-+}	$\gamma_5, \gamma_0\gamma_5$	$\pi^\pm, \pi^0, \eta, K^\pm, K^0, \dots$
Vector	1^{--}	$\gamma_i, \gamma_0\gamma_i$	$\rho^\pm, \rho^0, \omega, K^*, \phi, \dots$
Axial vector	1^{+-}	$\gamma_i\gamma_5, \gamma_i\gamma_0\gamma_5$	a_1, f_1, \dots
Tensor	1^{+-}	$\gamma_i\gamma_j$	h_1, b_1, \dots

Table 1: Overview of Γ for various channels following [Gattringer and Lang, 2010]. The γ matrices are used in physical basis and C classification is with respect to flavour neutral states. Examples for particles having the given quantum numbers are also given.

to its origin. This requires a large number of inversions of the fermion matrix with different sources in order to sample all possible paths sufficiently (see also Section 3.1.3). An overview of the possible Γ matrices for two flavours is given in Table 1.

The form of the correlator for large times will be dominated by the ground state, $C \sim e^{-tE_0}$. An *effective mass* at time t can be defined as

$$m_{\text{eff}}(t) = \ln \frac{C(t)}{C(t+1)}. \quad (2.80)$$

This expression will become constant when the ground state dominates. For mesons, one uses back- and forward propagation, i.e. that $C(t) \sim C(N_\tau - t)$. Then, the correlator will have a cosh-like behaviour and one can extract m_{eff} from

$$\frac{C(t)}{C(t+1)} = \frac{\cosh(m_{\text{eff}}(t - N_\tau/2))}{\cosh(m_{\text{eff}}(t+1 - N_\tau/2))}. \quad (2.81)$$

This kind of analysis is typically carried out in zero temperature studies, e.g. to extract meson masses. Generally, the same reasoning may be applied to correlators in one of the spatial directions. However, the exponential damping is no longer governed by the mass of a particle, but the exponential fall-off is rather interpreted as the screening of the particle in the medium, and the extracted quantity M_{scr} is called *screening mass*:

$$C_\Gamma(z) \sim e^{-M_{\text{scr}}z}. \quad (2.82)$$

This is the situation in finite T studies.

As was mentioned above, in the twisted mass formulation of QCD observables may be subject to change due to the change of basis, see (2.37). At maximal twist, this means for the general operator (2.77):

$$\begin{aligned} \bar{\psi}(n)\Gamma\tau^\pm\psi(n) &\rightarrow \frac{1}{2} \bar{\chi}(1+i\gamma_5\tau_3)\Gamma\tau^\pm(1+i\gamma_5\tau_3)\chi \\ &= \frac{1}{2} \bar{\chi} (\gamma_5\{\gamma_5, \Gamma\} - i[\gamma_5, \Gamma]\tau_3) \tau^\pm \chi. \end{aligned} \quad (2.83)$$




Figure 4: Illustrative examples of effective potentials V for some field Φ showing first order and second order transitions, respectively, as the temperature is varied passing a critical value T_c .

Here, the formula has been expressed in terms of the commutator $[\gamma_5, \Gamma]$ and the anticommutator $\{\gamma_5, \Gamma\}$ of γ_5 with general Γ . This means that if Γ is made up of an even number of γ matrices (different from γ_5), $\{\gamma_5, \Gamma\} = 2\gamma_5\Gamma$ and the general operator is of the same form as in the untwisted formulation of QCD. Yet, particles described by a Γ matrix which is made up of an odd number of γ matrices will appear quite different at maximal twist: The operators are of the form $\pm i\bar{\chi} \Gamma \tau^\pm \gamma_5 \chi$.

2.6 THE QCD PHASE DIAGRAM AND HEAVY ION COLLISIONS

In the following, terminology of phase transitions and statistical physics will be introduced, for a more elaborated approach see e.g. [Goldenfeld, 1992; Kapusta and Gale, 2006]. The phase diagram of (physical) QCD as shown in Figure 1 is discussed, with an emphasis on the thermal transition. This is based on recent reviews [Braun-Munzinger and Wambach, 2009; Fukushima and Hatsuda, 2011], for details on the other regions see details in these. The discussions of the current status of LQCD studies in the $\mu = 0$ region and in the region of purely imaginary μ are postponed to the beginnings of Chapters 4 and 5.

The partition function Z depends on external parameters like the temperature T or the chemical potential μ . As these are varied, the symmetry properties of the system may change, in particular, symmetries can be broken or restored. Then, the system is said to undergo a *phase transition*. Figure 4 illustrates the two types of phase transitions in terms of an effective potential. It is energetically favourable for the system to reside in a minimum of such a potential. In a *first order* phase transition there is an order parameter X whose expectation value $\langle X \rangle$ shows a discontinuity at the transition point. This can be seen in the effective potential, whose minimum abruptly changes at T_c . The system is said to tunnel from one minimum to the other. In turn, in a *second order* transition, the position of the minimum changes continuously

and there is no discontinuity in the order parameter. Here, the susceptibility $\chi(X)$ exhibits a divergence at the transition point and the system shows an infinite correlation length ξ . If no quantity shows a discontinuity, there is no true phase transition but the system will continuously change from one phase to the other, which is denoted as a *crossover* transition. Consequently, there is no true order parameter and the transition point will depend on the observable by which it is determined.

In the phase diagram spanned by the external parameters, the transitions mark the boundaries between regions of certain phases. As mentioned earlier, this boundary can not be clearly determined for crossover transitions. This indicates some further structure: A line of (true) phase transitions must end before a crossover region is encountered. As a first order transition can not just end, it will weaken and eventually end up in a second order *critical end point* (CEP). Furthermore, lines of first order transitions may meet in a *triple point*, where three phases coexist (e.g. ordinary water has this property). It is also possible that a line of triple points extends into the phase diagram and ends in a *tricritical point*.

A ferromagnet is a well-known example of a system showing phase transitions, and it can also be nicely related to QCD. At high T the system will reside in an *unordered phase* where the individual components fluctuate and the net magnetisation M averages to zero, $\langle M \rangle = 0$. At low T , in turn, neighbouring components will align in a specific, but random direction. On a macroscopic level, this will result in a net magnetisation, $\langle M \rangle \neq 0$, and the system is in an *ordered phase*.¹¹ Hence, at some T_c the system will undergo a second order phase transition between the two phases. An external field h applied to the system will predetermine the direction of the ordered phase (explicit symmetry breaking). As h changes from positive to negative values (reversing the direction), M changes its sign, too. At low T , one then has a first order phase transition, visible in the abrupt change of M . The difference ΔM at the transition point is called *latent heat*. As T is increased, the latter will shrink until the non-analyticity vanishes at T_c : The first order line terminates in a second order endpoint. For high temperatures, the system will show crossover behaviour as the external field is varied.

An example of a tricritical point can be seen in a metamagnet, which can show ferromagnetic as well as anti-ferromagnetic behaviour. The more complicated phase structure of the system with respect to the ferromagnet discussed above comes about from the interplay of two external fields applied to the system, h and h^\dagger . Of these fields, only h breaks a symmetry, h^\dagger , does not. In the phase diagram spanned by h , h^\dagger and T , the tricritical point is located at some finite temperature and finite h^\dagger in the plane of zero h . In particular, there are second order lines emerging from this point into positive and negative h direction, respectively. See for example [Bonati et al., 2011b] for a sketch of the phase diagram.

In a region around a second order point, the behaviour of the system is solely determined by its *universality* class and few critical exponents characteristic to this class governing the scaling of observables. The universality depends on the dimension of the system and the symmetries of H . It is convenient to introduce the *reduced temperature* $t = (T - T_c)/T_c$. In this region, one can describe the system by an effective Hamiltonian (*Ginzburg-Landau theory*) and analyse

¹¹ This is an example of spontaneous symmetry breaking.

the power-law behaviour of its describing quantities. For example, the behaviour of M and its susceptibility $\chi = \partial M / \partial h$ is dictated by:

$$M \sim |t|^{\tilde{\beta}}, \quad (2.84)$$

$$M \sim |h|^{1/\delta}, \quad (2.85)$$

$$\chi \sim |t|^\gamma. \quad (2.86)$$

Here, $\tilde{\beta}$, δ and γ are (universality-class specific) exponents.¹² The scaling of M is encoded in a universal scaling function $f(z)$ (see [Parisen Toldin et al., 2003; Pelissetto and Vicari, 2002])

$$M = h^{1/\delta} f(z), \quad (2.87)$$

$$\chi = h^{1/\delta-1} \left(\frac{1}{\delta} f(z) - \frac{1}{\delta \tilde{\beta}} z f'(z) \right) = g(z). \quad (2.88)$$

The scaling variable z is defined as $z = t/h^{1/(\delta\tilde{\beta})}$. (2.87) is called *Magnetic Equation of State*.

These considerations strictly hold only in the thermodynamic limit where the spatial volume goes to infinity, $V \rightarrow \infty$. In a finite system there is never a true phase transition but one has always a crossover, since Z , (2.11), as a finite sum is always analytic and no divergences can occur. To infer on the situation in the thermodynamic limit, one therefore has to do a *finite size scaling analysis* taking into account multiple system extents.

The QCD phase diagram is subject to intensive theoretical and experimental investigations. A sketch is given in Figure 1. At low temperatures and densities hadronic matter is in the confined phase. This phase can be well described by the *Hadron Resonance Gas Model (HRG)*, in which hadrons are treated as a free gas. At a certain mass threshold, all particles start to overlap and are treated as a Boltzmann distributed gas. Then, as the energy of the system is increased, the partition function will diverge at a certain temperature (*Hagedorn temperature* [Hagedorn, 1965]) of around 160-180 MeV. In QCD, confinement of quarks and gluons eventually has to soften as asymptotically large values for T or μ are reached, due to the running of the coupling. Then, quarks and gluons can be described as a non-interacting gas of free particles. The degrees of freedom of the system change from hadronic to quarks and gluons. Between these two phases there must be a transition, called the *thermal transition*. At $\mu = 0$ the critical temperature T_c has been determined on the lattice to be around 150-170 MeV [Aoki et al., 2006], depending on the observable. This result has been obtained after settling a long on-going debate of different groups, known as the *T_c crisis* (see e.g. [Borsanyi et al., 2010]). Note that T_c agrees with the Hagedorn temperature within errors. Indeed, the lattice data can be well described by the HRG and a gas of free particles in the respective domains $T < T_c$ and $T \gg T_c$. For the theoretical investigation of Heavy Ion collisions by means of transport models [Petersen et al., 2008] the QCD equation of state is an important ingredient. It can be extracted from LQCD simulations, see [Philipsen, 2013] for a recent review on this topic. On the lattice it is found that the transition is not a true phase transition, but a rapid crossover. Up to now, explorations at $\mu \neq 0$ are hampered to the sign-problem which is

¹² To distinguish it from the lattice coupling β , the critical exponent is denoted as $\tilde{\beta}$.

explained in the next chapter. The expectations in this region are based on models of QCD that resemble certain aspects of the full theory. In the case of chiral symmetry QCD can be described in terms of magnetic systems. For instance, the chiral condensate $\langle\bar{\psi}\psi\rangle$ signals the breaking of chiral symmetry. Thus, one can readily identify M with $\langle\bar{\psi}\psi\rangle$. The role of the external field h breaking the symmetry is adopted by the mass m . This analogy is deepened further below when discussing the phase structure obtained from LQCD in the next chapters. The model calculations predict a first order transition for higher values of μ . Consequently, an CEP is expected where the first order line ends.

The phase structure of QCD is important to understand what happened shortly after the Big Bang. In this environment of extremely high temperature the *Quark-Gluon-Plasma* was formed, where quarks and gluons were quasi-free particles. As the universe expanded and cooled down, matter underwent a transition to the hadronic world. This situation is investigated experimentally at current Heavy Ion collision experiments, where nuclei are shot at each other relativistically. In the emerging fireball, a Quark-Gluon-Plasma is created after a quick equilibration. It behaves like a liquid, which can be inferred from the experimental data in terms of a Fourier analysis (in particular the Fourier coefficient v_2 can be described by ideal hydrodynamics). The fireball then starts to expand, the details are highly influenced by the anisotropies of the system prevailing at the impact. Strongly interacting particles of high momentum (*jets*) can not leave the fireball without energy loss in the medium (*jet quenching*). This can be made quantitatively in Au-Au collisions, where the ratio of number of particles at a given transverse momentum p_T to that in proton-proton collisions, R_{AA} , is measured. In the data, R_{AA} is one for photons, indicating no energy loss, but strongly-interacting hadrons are strongly suppressed, $R_{AA}(p_T) < 1$. As the fireball expands, the temperature sinks and hadrons start to form. The hadrons interact inelastically until the energy drops below a certain threshold (*chemical freeze out*). Elastic collisions cease at even lower temperatures (*thermal freeze out*) and the hadrons can be measured in detectors surrounding the impact point. The regions of the phase diagram which can be explored are specific to the experiment. At the LHC and at RHIC very high temperatures are reached at moderate densities, similar to the situation in the early universe, where the thermal transition occurred at very small μ . On the other hand, the future FAIR facility will be able to go to higher densities at moderate temperatures.¹³

13 See <http://www.bnl.gov/rhic/>, <http://home.web.cern.ch/about/accelerators/large-hadron-collider> and <http://www.fair-center.de>.

NUMERICAL ASPECTS OF LATTICE QCD AND THE CL²QCD SOFTWARE

The lattice formulation of QCD provides a mathematically solid formulation of the path integral and thus allows for its numerical evaluation. As it constitutes one of the most compute intensive problems, to have an efficient numerical implementation of Lattice QCD (LQCD) at hand is mandatory for investigations in this field. In recent years, the usage of *Graphics Processing Units* (GPUs) has become an integral part of state-of-the-art LQCD application, as they provide a well-suited architecture. During the development of this thesis, a new LQCD application called CL²QCD ¹ has been developed which was reported on in [Bach et al., 2013a]. Various aspects of LQCD are covered in CL²QCD and at the time of writing, it consists of $\mathcal{O}(100k)$ lines of code.² The presentation of its features is the topic of this chapter. In particular, CL²QCD is based on OpenCL, which is a certain way of utilising GPUs, and CL²QCD constitutes the first OpenCL implementation for Wilson fermions. Furthermore, CL²QCD is able to utilise GPUs and traditional CPUs in the same fashion and shows excellent performance compared to a reference code.

This chapter is organised as follows: Algorithmic concepts used for the efficient numerical realisation of LQCD are presented in Section 3.1. The analysis of data gained in LQCD simulations is addressed in Section 3.2. In Section 3.3 the numerical structure of LQCD is discussed. The use of GPUs for LQCD simulations is motivated and introduced in Section 3.4. Afterwards, the CL²QCD software, which was developed within this thesis, is presented. The chapter is concluded with a summary and a discussion, Section 3.6.

3.1 ALGORITHMS

Currently, the integration over the fermion fields cannot be performed efficiently because of the Grassmann nature of the fermions (see Appendix A.3).³ Hence, it is more convenient to integrate the fermions out using an identity for Gaussian Grassmann integrals (A.13):

$$\int \mathcal{D}[\bar{\psi}, \psi] \exp \left(\sum_{i,j=1}^N \bar{\psi}_i M_{ij} \psi_j \right) = \det M \quad (3.1)$$

¹ See <http://code.compeng.uni-frankfurt.de/projects/clhmc> .

² This was measured with the *CLOC* software, <http://cloc.sourceforge.net/> .

³ There are attempts to do the integration numerically, for example [Creutz, 1998].

for an $N \times N$ matrix M . After the integration, the determinant depends only on the gauge field. In QCD, however, the fermion matrix D is a huge square matrix of order $N = N_c N_{\text{Dirac}} V_{\text{tot}}$, which renders the direct evaluation of $\det D[U]$ unfeasible. One can (formally) express $\det D$ as $\exp \text{Tr} \ln D$. This yields a formulation of Z_{QCD} that is numerically accessible:

$$Z_{\text{QCD}}[U] = \int \mathcal{D}[U] e^{-S_{\text{eff}}}, \quad (3.2)$$

with an effective action

$$S_{\text{eff}}[U] = S_{\text{gauge}}[U] - \text{Tr} \ln D. \quad (3.3)$$

If D is positive definite, one can express the determinant in terms of bosonic degrees of freedom ϕ (*pseudo fermions*),

$$\det D[U] \propto \int \mathcal{D}[\phi^\dagger] \mathcal{D}[\phi] \exp \left\{ -\phi^\dagger D^{-1}[U] \phi \right\}. \quad (3.4)$$

The fact that this introduces the inverse of the fermion matrix, D^{-1} , is of great importance for the numerical evaluation of (3.4) and will be commented on further below. In particular, this means that the effective action has a non-local component. The real fields $\phi^\dagger(n)$ and $\phi(n)$ are conveniently treated as one complex valued field which will also be denoted as ϕ .

In general, the eigenvalues λ_i of the fermion matrix D are complex. However, at zero μ the former have to come in complex pairs or be real because of γ_5 -hermiticity (2.61). Hence, the fermion determinant $\det D = \prod_i \lambda_i$ is real. For Wilson fermions, this means that the determinant at $\mu = 0$ is positive (semi-)definite if an even number of mass-degenerate flavours are considered, $\det^2 D = \det D^\dagger D \geq 0$. It is conveniently written as the Hermitian matrix $D^\dagger D$ using (2.61). As twisted mass fermions always come in even numbers, their determinant is naturally positive definite and even strictly non-zero for $\hat{\mu} \neq 0$:

$$\det D_{\text{tm}}^+ \det D_{\text{tm}}^- = \det \left(|D_{\text{Wilson}}|^2 + 4\kappa^2 \hat{\mu}^2 \right). \quad (3.5)$$

In the following, the $N_f = 2$ flavour case is considered only. With pseudo fermions, the effective action then reads:

$$S_{\text{eff}}[U, \phi^\dagger, \phi] = S_{\text{gauge}}[U] + \phi^\dagger (D^\dagger D)^{-1}[U] \phi \quad (3.6)$$

and $Z_{\text{QCD}}[U] \rightarrow Z_{\text{QCD}}[U, \phi]$. The degrees of freedom are thus ϕ , which is a $(N_{\text{Dirac}} \times N_c)$ -dimensional complex vector per lattice site, and the gauge links $U_\mu(n)$, which are complex valued $N_c \times N_c$ matrices.

3.1.1 Hybrid Monte Carlo

Z_{QCD} (3.2) is the integral over all possible gauge field configurations (the whole grand-canonical ensemble) and as such it cannot be evaluated fully. Thus, the standard numerical approach to QCD employs importance sampling methods. A numerical evaluation of Z_{QCD} then consists of the stochastic generation of a number N of gauge configurations $\{U_n\}$ (these will be called *ensemble*, too, which is strictly true for $N \rightarrow \infty$ only). They form a so-called *Markov chain*.⁴

⁴ For details on the properties of Markov chains see standard textbooks, e.g. [Gattringer and Lang, 2010; Montvay and Münster, 1994].

If the $\{U_n\}$ are distributed according to the Boltzmann-weight $p = \exp\{-\mathcal{S}_{\text{eff}}\}$, the true ensemble average $\langle \mathcal{O} \rangle$ can be approximated by

$$\langle \mathcal{O} \rangle \approx \frac{1}{N} \sum_n \mathcal{O}[U_n] . \quad (3.7)$$

The time needed to reach the correct distribution in the simulation is called *thermalisation*. To be a well-defined probability measure, p must be positive and non-zero, i.e. there must be no sign-problem (see below). In the *Hybrid Monte-Carlo (HMC)* algorithm [Duane et al., 1987] the effective action is embedded in a fictitious classical molecular dynamics system.

This is governed by the Hamiltonian $H = \frac{1}{2}P^2 + \mathcal{S}_{\text{eff}}[U, \phi]$. The momenta P , conjugate to U , are elements of $\mathfrak{su}(N_c)$ and $P = P_\mu(n)$ is conveniently expressed as a real-valued, $(N_c \times N_c - 1)$ -dimensional vector. At the beginning of each evolution ϕ and P are generated according to the Gaussian distributions $\exp\{-(D^{-1}\phi)^\dagger D^{-1}\phi\}$ and $\exp\{-\frac{1}{2}P^2\}$, respectively. Given a configuration (U, P) , the system is evolved over a time τ to a new configuration (U', P') according to the Hamiltonian equations of motion

$$\dot{P} = -\partial \mathcal{S}_{\text{eff}} / \partial U \equiv F , \quad (3.8)$$

$$\dot{U} = P . \quad (3.9)$$

F is the force term and is of the same dimensionality as P .⁵ The evolution is accomplished by some suitable numerical integration scheme, for example the leapfrog- or the second order minimal (2MN) scheme (see [Takaishi and de Forcrand, 2006] and references therein). Their accuracy is controlled by the number of integration steps l , which define some minimal step size $\Delta\tau = \tau/l$. Each evolution step is completed with a Metropolis accept-reject step [Metropolis et al., 1953]. This means (U', P') is accepted only with a probability

$$P_{\text{acc}} = \min(1, \exp(H[P', U'] - H[P, U])) . \quad (3.10)$$

The Metropolis step is mandatory to ensure detailed balance. Such an evolution is called a *trajectory*.

Within the scope of this work, there can be three contributions to the force:

$$F = F_{\text{plaquette}}^{\text{gauge}} + F_{\text{rectangles}}^{\text{gauge}} + F^{\text{fermion}} , \quad (3.11)$$

where the gauge part of the action has been divided into plaquette and rectangle parts. These are proportional to sums of link products (*staples*), which themselves are no longer elements of $SU(N_c)$ for $N_c \geq 2$, projected back to the algebra. The plaquette contribution $(F_{\text{plaquette}}^{\text{gauge}})_\mu(n)$, for example, is proportional to $\sum_{\mu \neq \nu} \tilde{P}_{\mu\nu}(n)$, see also Appendix A.2 for an explicit calculation. The staple $\tilde{P}_{\mu\nu}(n)$ is like the plaquette $P_{\mu\nu}(n)$ product of links in Figure 2, but without $U_\mu(n)$. The fermion force can be evaluated by means of the identity $\partial_U D^{-1} = D^{-1}(\partial_U D)D^{-1}$ and receives contributions from the \not{D} part of the fermion matrix only since the diagonal part does not depend on U . Hence, the force has the same expression for pure Wilson and twisted mass fermions.

⁵ For a definition of the derivative with respect to U see Appendix A.2.

There are techniques to refine the numerical integration. It turns out to be beneficial to choose the number of integration steps of the different contributions to the force according to their numerical costs (*multiple timescales*) [Urbach et al., 2006], i.e. choose $\Delta\tau$ smaller for numerically cheaper parts. The inverter can be preconditioned by introducing additional pseudo fermions via $\det(D) = \det(A) \frac{\det(D)}{\det(A)}$, where A denotes a fermion matrix with a different mass than D (*mass preconditioning* or *Hasenbusch trick* [Hasenbusch, 2001]). For details and further readings see [Jansen and Urbach, 2009].

3.1.2 Heatbath Algorithm

The HMC algorithm is perfectly suited to simulate pure gauge theory, where the fermion determinant simply drops out of the equations, $\det D \equiv 1$. However, without the inverse fermion matrix the action is not non-local anymore. Hence, there are more direct ways to evolve the system. The *heatbath* algorithm [Cabibbo and Marinari, 1982; Creutz, 1980; Kennedy and Pendleton, 1985] is based on an exact algorithm for $SU(2)$, which is extended to general $SU(N_c)$ case by reducing the $SU(N_c)$ links systematically to $SU(2)$ subgroups (for $N_c = 3$ these are usually also 3 groups). A sketch of the algorithm is shown in Figure 5. In here, `project` and `extend` denote reduction and extension to and from the i th subgroup, whereas W , U , V , L denote $SU(N_c)$ links and w , v $SU(2)$ links, respectively. The `update` routine performs either the exact $SU(2)$ update mentioned above or an overrelaxation step [Petronzio and Vicari, 1990]. The latter is applied to cover the whole configuration space more quickly. Random numbers are needed in the $SU(2)$ update only.

```
L(0) = U
for i < m = # subgroups:
  W = L(i) * Staple
  w = project(W, i)
  v = update(w)
  V = extend(v, i)
  L(i) = L(i-1) * V
U = L(m)
```

Figure 5: Heatbath algorithm

3.1.3 Inversion of the Fermion Matrix

It was argued in Section 2.5.5 that the inverse fermion matrix is a crucial ingredient for mesonic correlation functions. For more involved observables, more occurrences of D^{-1} will appear. In addition, it was argued that D^{-1} is an essential ingredient to the effective action (3.3). To calculate D^{-1} is thus important, although it can not be determined analytically. In addition, it is sufficient to know its action on a vector and not all elements of the full matrix.

As D is a sparse square matrix, Krylov space based methods are the best choice for the inversion [Meister, 2008]. Given a source \mathbf{b} these solve the equation $D\mathbf{x} - \mathbf{b} = 0$ iteratively and yield the solution

$$D\mathbf{x} = \mathbf{b} \Rightarrow x_a = (D^{-1})_{ab} b_b . \quad (3.12)$$

Each entry of \mathbf{x} represents the probability of a fermion propagation. For more detailed information about the matrix inversion and the choice of sources see also Appendix A.5.

```

r=b-D*x
p=r
rsold=r'*r
for i < # steps:
  Dp=D*p
  alpha=rsold/(p'*Dp)
  x=x+alpha*p
  r=r-alpha*Dp
  rsnew=r'*r
  if sqrt(rsnew)<eps
    break
  beta = rsnew/rsold
  p=r+beta*p
  rsold=rsnew

```

Figure 6: CG algorithm

In LQCD, the generic example is the *Conjugate Gradient* (CG) solver, see Figure (6). Apart from the matrix-vector product Dx , it needs a couple of linear algebra operations in each iteration. The number of iterations it takes to converge to some precision ϵ depends heavily on m_f and V_{tot} , but will be at most equal to N . The total costs can be approximated as [Ukawa, 2002]

$$\text{cost} \sim V_{\text{tot}}^{5/4} a^{-7} m_{\pi}^{-6}. \quad (3.13)$$

However, the CG may be applied to positive definite Hermitian matrices only. A commonly used variant of the CG is the *Bi-conjugate Gradient Stabilised* solver (BiCGStab), which may be applied to any matrix, yet, without ensured convergence. To decrease the number of iterations, various techniques may be applied. The \mathcal{D} term connects ψ at different sites and is the only non-local term needed in the CG.

However, regarding space time, it connects only nearest neighbours. Since the parity of a site n , $(-1)^{x_0+x_1+x_2+x_3}$, and its neighbour $n + \hat{\mu}$ are always different, a decomposition of \mathcal{D} into an even (parity 1) and odd (parity -1) part is possible (*even-odd preconditioning*) [Degrand and Rossi, 1990], see also Figure 63b and Appendix A.4. This leads to a reformulation of the problem where a matrix must be inverted by the solver that acts on sites of one parity only. This has several advantages: The matrix is only half the size of M , thus memory objects are half the size, too. In addition, since the smallest eigenvalue is approximately doubled in the reformulated problem, faster convergence is possible. (3.12) can then be trivially reconstructed.

The inversion may be ill-defined for an *exceptional configuration* U , where the eigenvalues of D are very small. Twisted mass fermions prevent this from happening and act as an infrared regulator (see (3.5)).

3.1.4 Reweighting

An observable X measured at one set of parameters, p , is exactly related to another set of parameters, q , via

$$\langle X \rangle_q = \frac{1}{Z_q} \int \mathcal{D}[U] e^{-S_{\text{eff}}^q} X = \frac{1}{Z_q} \frac{Z_p}{Z_p} \int \mathcal{D}[U] e^{-S_{\text{eff}}^q} \frac{e^{-S_{\text{eff}}^p}}{e^{-S_{\text{eff}}^p}} X \equiv \frac{Z_p}{Z_q} \langle \frac{e^{-S_{\text{eff}}^q}}{e^{-S_{\text{eff}}^p}} X \rangle_p. \quad (3.14)$$

Thus, by means of the *reweighting factor* $e^{-S_{\text{eff}}^p}/e^{-S_{\text{eff}}^q}$, it is possible to calculate $\langle X \rangle_q$ from an ensemble generated at parameters p .

However, a significant drawback is that this ensemble has been generated according to a probability distribution specific to p , i.e. $\mathcal{D}[U]$ has not been fully evaluated. At q , this weight can be quite different and its maxima need not have great overlap with those from p . Therefore, it is possible to sample rather unimportant sections of phase space, this is called the *Overlap problem*.

It is possible to combine the information of multiple ensembles in the reweighting (*multi-histogram technique*) [Ferrenberg and Swendsen, 1989]. In particular, for β -scans this method is called *Ferrenberg-Swendsen reweighting*.

3.1.5 The Sign Problem

At $\mu \neq 0$ one has $\det^* D(\mu) = \det D(-\mu^*)$, cf. (2.75), which is real for purely imaginary values of μ only.⁶ Thus, at real μ $\det D$ is really complex and in particular not positive definite, also for an even number of flavours. This is called the *sign problem*. Consequently, one cannot express the fermion determinant with pseudo fermions. In addition, it is not even possible to rewrite $\det D = \exp \ln \det D$ unambiguously because of the cut of the logarithm on the imaginary axis.

One way to access finite (real) values of μ is by means of reweighting, either of $\mu = 0$ data or of simulations in the *phase quenched approximation*. Here, the determinant is replaced by its absolute value, $|\det^* D(\mu) \det D(\mu)| = \det D(-\mu) \det D(\mu)$. This corresponds to up- and down-flavour having a relative sign on μ (*isospin chemical potential*). Apart from the overlap problem during reweighting, the sign-problem manifests itself in huge fluctuations and cancellations due to the complexity of the determinant. This is shown in the *average sign*, $\langle e^{i\theta} \rangle$. As this quantity approaches zero, the reweighted signal is lost in noise and the sign problem is said to be severe. The average sign is actually proportional to the reweighting factor itself, i.e. it is a ratio of partition functions. Their difference in free energy is an extensive quantity and grows with volume:

$$e^{-S_{\text{eff}}^q} / e^{-S_{\text{eff}}^p} = e^{-\Delta F/T} \sim e^{-\text{const } V} . \quad (3.15)$$

This means that the sign problem gets worse at larger volumes.

Alternatively to reweighting, one can expand the partition function in a Taylor series, e.g. for the pressure:

$$p(\mu) = T^4 \sum_n \frac{\partial^{2n} \ln Z(\mu)}{\partial \mu^{2n}} \Big|_{\mu=0} (\mu/T)^{2n} , \quad (3.16)$$

which has only even terms because of (2.31). This approach can be trusted until a non-analyticity of Z is encountered and a number of coefficients sufficient for convergence of the series (3.16) are known.

As the determinant is complex for real μ only, imaginary values can be simulated without additional algorithmic problems. QCD has interesting properties in this regime, which are addressed in Section 2.2.3 and Chapter 5. Measurements can be *analytically continued* to the real μ region [de Forcrand and Philipsen, 2002], however, the continuation is limited by the periodicity of the partition function at imaginary μ to $\mu/T \lesssim \pi/N_c$.

For a detailed overview about the sign-problem and methods to circumvent it, see for example [Levkova, 2011].

⁶ The situation is different for unphysical values of colours, e.g. $N_c = 2$ [Alles et al., 2007].

3.2 DATA ANALYSIS

In the following, basic terms of the analysis of LQCD simulation data will be introduced. For more details see e.g. [Gattringer and Lang, 2010; Montvay and Münster, 1994].

The N gauge configurations generated in the Markov process display only a subset or a sample of the whole ensemble. Consequently, the true ensemble mean $\langle \mathcal{O} \rangle$ can only be approximated by the sample mean $\langle \tilde{\mathcal{O}} \rangle$, see (3.7). Similarly, the true *variance*,

$$\sigma^2 = \langle (\mathcal{O} - \langle \mathcal{O} \rangle)^2 \rangle, \quad (3.17)$$

is estimated by the sample variance $\tilde{\sigma}^2$. For uncorrelated measurements one can show that

$$\tilde{\sigma}^2 = \frac{1}{N} \sigma^2. \quad (3.18)$$

This means that the statistical error decreases with \sqrt{N} .

As in the Markov chain the configurations are generated subsequently, in general there will be *autocorrelations*, which can be quantified in terms of the *autocorrelation function*

$$\Gamma_{\mathcal{O}}(\tau) \equiv \langle (\mathcal{O}_i - \langle \mathcal{O} \rangle)(\mathcal{O}_{i+\tau} - \langle \mathcal{O} \rangle) \rangle. \quad (3.19)$$

Here, \mathcal{O}_i denotes $\mathcal{O}[U_i]$ and τ is the separation in simulation time (not to be confused with the HMC trajectory length of the same symbol). As the Markov chain is generated in equilibrium, Γ is a function of τ only. In the correlated case, the variance of the sample is now

$$\tilde{\sigma}^2 = \sum_{i=1}^N \sum_{\tau=-N}^N \frac{N-|\tau|}{N^2} \Gamma_{\mathcal{O}}(\tau) \xrightarrow{N \rightarrow \infty} \sigma^2 \frac{2\tau_{\text{int}}}{N}, \quad (3.20)$$

with *integrated autocorrelation time*

$$\tau_{\text{int}} = \frac{1}{2} \sum_{\tau=-\infty}^{\infty} \frac{\langle \Gamma(\tau) \rangle}{\langle \Gamma(0) \rangle}. \quad (3.21)$$

For an estimate of the error on τ_{int} see e.g. [Wolff, 2004]. Comparing (3.18) and (3.20) one sees that the autocorrelation effectively decreases the number of measurements N .

To obtain uncorrelated data *binning* or *blocking* can be applied. For instance, this can be done by dividing the data into M subsamples of subsequent measurements. If the subsamples are large enough, their means can be considered as a new set of independent data points. In general, the blocksize M has to be tuned. This can be done by varying M and identifying a plateau-like behaviour of σ^2 .

Deleting the m -th subsample, a “new” dataset \mathcal{O}^m can be obtained with average

$$\mathcal{O}^m = \frac{1}{N-M} \sum_i \mathcal{O}_i^m. \quad (3.22)$$

This method is called *jackknife* and will be used for error estimates throughout. The variance

$$\sigma_{\text{jack}}^2 = \frac{M-1}{M} \sum_{m=1}^M (\mathcal{O}^m - \langle \mathcal{O} \rangle)^2, \quad (3.23)$$

on this data set is an estimator for σ^2 . The autocorrelation of the original data can then be estimated by (3.20).

3.3 NUMERICAL STRUCTURE OF LQCD

In the previous sections, various ingredients of LQCD simulations have been introduced. As an example of how these appear in the actual implementation, the \mathcal{D} will be analysed more closely now. The fermion matrix itself is not needed explicitly but only its action on a vector, i.e. one has to numerically calculate

$$\tilde{\psi}(n) = \mathcal{D}\psi(n) \sim (1 - \gamma_\mu)U_\mu(n)\psi(n + \hat{\mu}). \quad (3.24)$$

The (constant) Dirac part $(1 - \gamma_\mu)$ can be reconstructed analytically from the colour part (for details see [Jansen and Urbach, 2009]), so what is needed for the actual computation is the $SU(N_c)$ matrix $U_\mu(n)$ and the spinor $\psi(n + \hat{\mu})$ in each (positive and negative) direction μ . In the numerical application, these objects need to be transferred from computer-memory to the processor. The processor stores this data in *registers*. Afterwards, the result has to be transferred back to memory for storage. All in all, for $N_c = 3$ these are 2880 Bytes per site.⁷ The number of floating point operations (FLOP) the processor has to perform on this data is 1632 FLOP.⁸

An algorithmic density ρ_{alg} may be defined as the number of FLOP performed per Byte. For the \mathcal{D} this would be of $\mathcal{O}(0.5)$. A typical Central Processing Unit (CPU) nowadays has a peak performance of $\mathcal{O}(150)$ GFLOP per second (GFLOPS) and a peak memory bandwidth (BW) of $\mathcal{O}(50)$ GByte per second (GByte/s), see also Table 3. These two numbers are, of course, to be taken with caution as they correspond to the maximum theoretically possible performance on a given hardware. Nevertheless, for optimal hardware utilisation a function should have $\rho_{\text{alg}} \sim \mathcal{O}(3)$, which is clearly not the case for \mathcal{D} . The \mathcal{D} performance is dominated solely by memory transfer. In addition it is important to note that \mathcal{D} is an ultra-local function and can be easily parallelised as such, i.e. an update of one element of ψ does not depend on the other updated elements.

This is a generic pattern for LQCD: Its functions are always bandwidth limited and ultra-local.⁹ Table 2 shows an example for a couple of HMC and Heatbath related functions. For all HMC related kernels ρ_{alg} is always $\lesssim 1$. The heatbath algorithm might actually be expected to be less bandwidth bound, but it is not FLOP dominated either. There is also a clear hierarchy in the share of the various functions in overall performance time of a typical simulation. Clearly, of those the matrix inversion is most expensive, as it is an iterative process that in addition depends on the simulation parameters. The \mathcal{D} application is much more expensive than the linear algebra functions used in the solver. For a well performing LQCD application a well tuned \mathcal{D} is thus mandatory, and this process will heavily depend on the used hardware. Note, however, that Amdahl's Law [Amdahl, 1967] applies: By having a well-tuned \mathcal{D} , the other parts of the simulation become more important for overall performance. Generally, the tuning of the application is hardware specific.

⁷ Double precision (DP) will be used throughout, so a real number is represented by 8 Bytes. In single precision (SP) the quoted sizes would be halved.

⁸ This number does take the analytic simplifications in the Dirac structure mentioned above into account. Also, the number may vary slightly depending on the way of counting.

⁹ In the scope of this thesis this means they have only (next-to) nearest neighbour interactions.

	\mathcal{D}	$F_{\text{plaquette}}^{\text{gauge}}$	$F_{\text{rectangles}}^{\text{gauge}}$	F_{fermion}	saxpy	heatbath	overrelax
RW-Size [Byte]	2880	2800	13168	4736	608	2880	2880
FLOPS	1632	2717	14813	1748	64	3912	3846

Table 2: Read-Write (RW) and FLOPS sizes for some LQCD functions. All numbers are per site (and direction). “saxpy” corresponds to the algebraic operation $\vec{z} = \alpha\vec{x} + \vec{y}$.

In terms of memory usage the gauge field is the biggest individual object to store (see Table 14 for an overview). However, depending on the algorithm one might need more than one instance of an object. It is possible to reduce the object size, e.g. even-odd preconditioning reduces the size of a spinor field by a factor of two. The N_c^2 complex entries of a link U are not independent of each other, in fact, it can be represented by only $N_c^2 - 1$ real numbers, one for each generator of $SU(N_c)$ (see Appendix A.2). As significant computational overhead is introduced by this representation, other methods are more feasible. Using the cross-product for $N_c = 3$, $\vec{c} = \vec{a} \times \vec{b}$, one column of U can be reconstructed from the other two exactly. This saves three complex numbers and U is said to be reconstructed by 12 real numbers (REC12). Two or four more numbers can be saved using additional restrictions on U (REC10 and REC8, respectively). However, REC8 is known to run into arithmetic problems during the reconstruction [Clark et al., 2010].

As was mentioned above, it is algorithmically natural to parallelise the LQCD functions. On the level of a single CPU this can be done by *Open Multi-Processing*¹⁰ (OpenMP), where for example each core operates only on a subset of the whole lattice. One of the parallel workflows is called a *thread*. Processors which do not share a common memory, e.g. are connected via a network, may be operated on by using *Message Passing Interface*¹¹ (MPI) functionality or the like. This is the situation found in computing clusters. There is a long history of clusters used for LQCD, for example the BLUE-GENE[Bhanot et al., 2005; Boyle, 2012; Doi, 2012] (see Appendix A.7 for more details of this machine and [Padua, 2011] for more general information). These are usually equipped with a huge amount of CPUs and come with quite high costs of acquisition. Additionally, parallelism always introduces a communication overhead. Therefore, doubling the number of computing elements does not necessarily result in doubled performance. Actually, for certain hardware optimal performance can be expected by combining OpenMP and MPI functionality [Abdel-Rehim et al., 2013b].

After these general considerations about the numerical operations in LQCD, they should be put into the context of a physical investigation. In general, it consist of two parts: The generation of ensembles of gauge configurations $\{U_n\}$ and the evaluation of observables $\mathcal{O}_i[U]$ on these (see (2.79) for an example). The first part depends on a set of parameters for the HMC, i.e. $(\beta, \kappa, \hat{\mu})$.¹² Each HMC chain is an explicitly serial process, i.e. only one instance of the application may be used per chain. On the contrary, the evaluation of $\mathcal{O}_i[U_j]$ on a given

¹⁰ See <http://openmp.org/wp/openmp-specifications/> .

¹¹ See <http://www.mcs.anl.gov/research/projects/mpi/> .

¹² Rather cheap observables like Plaquette or Polyakov loop are measured during the generation process already. They may also serve as monitoring tools.

	Chip	Peak SP [GFLOPS]	Peak DP [GFLOPS]	Peak BW [GByte/s]
AMD Radeon HD 5870	Cypress	2720	544	154
AMD FirePro V7800	Cypress	2016	403	128
AMD Radeon HD 6970	Cayman	2703	683	176
AMD Radeon HD 7970	Tahiti	3789	947	264
AMD FirePro W8000	Tahiti	3230	810	176
AMD FirePro FirePro S10000	Tahiti	6820	1700	480
NVIDIA Tesla C1060	Tesla	933	78	102
NVIDIA GeForce GTX 280	Tesla	933	78	142
NVIDIA GeForce GTX 480	Fermi	1345	132	177
NVIDIA GeForce GTX 580	Fermi	1581	198	192
NVIDIA Tesla M2090	Fermi	1331	665	177
NVIDIA GeForce GTX 680	Kepler	3090	258	192
NVIDIA Tesla K20	Kepler	3520	1170	208
AMD Opteron 6172	Magny-Cours	202	101	42.7
AMD Opteron 6278	Interlagos	307	154	51.2
Intel Xeon E5-2690	Sandy Bridge EP	371	186	51.2

Table 3: Theoretical peak performance of current GPUs and CPUs. Note that the AMD FirePro S10000 is a dual GPU.

configuration U_j is independent of the other U . Thus, this process can be naively parallelised using multiple instances of the application.

A further distinction has to be made between zero and finite temperature simulations. The first are usually done for spectroscopy studies, hence the focus lies more on a variety of observables computed on a few gauge ensembles. These even might have been created already in previous efforts, and typically at only a few different parameters. At finite temperature, one usually has only a few (and quite simple) observables, but wants to perform temperature scans, i.e. simulate at various β . In addition, one ideally wants various lattice volumes for the thermodynamic limit. This shows that the focus of finite temperature studies, which are covered in this thesis, in terms of computing time lies mainly on HMC runs. Of course, these may be naively parallelised for a given set of parameters.

Finally, one should emphasise that LQCD studies often fall into the category of High Performance Computing (HPC). As an example, one HMC trajectory in the Z12 setup described in Chapter 4 already takes approximately 32 CPU hours. However, compared to state-of-the-art zero temperature studies the lattice size used is rather small. These studies can be expected to have even higher costs, cf. (3.13).

3.4 GRAPHICS PROCESSING UNITS

In recent years, *Graphics Processing Units* (GPUs) have become a more and more important part in modern computers as commercial applications have increasingly high demands on graphical output, for example in video games. Table 3 shows that in fact GPUs surpass CPUs

by far in terms of peak performance as well as bandwidth. This evolution has also led to their usage in areas other than imaging. Today, modern GPUs can be used naturally for general calculations (*General Purpose GPU (GPGPU)*).¹³ In addition to their performance, GPUs are also very attractive in terms of acquisition and maintenance, as they have a better performance per Watt compared to equivalent CPU systems [Bach et al., 2013b]. Therefore, they have become an integral part of many modern clusters. Whereas traditional CPU clusters are quite homogeneous systems, GPU clusters are heterogeneous as they have a number of CPUs and GPUs per node.

At Goethe University in Frankfurt GPU clusters are accessible: The LOEWE-CSC [Bach et al., 2011] and SANAM supercomputers [Kalcher et al., 2013]. Both are equipped with *Advanced Micro Devices (AMD)* hardware. The LOEWE-CSC comes with an AMD Radeon HD 5870 GPU and has a rather balanced CPU and GPU performance. On SANAM, the two dual GPUs AMD FirePro S10000 provide more performance than the CPUs. Both clusters were developed aiming at high energy efficiency as well as reasonable costs of acquisition [Bach et al., 2013b]. For more details see Appendix A.7.

The GPU architecture is substantially different from the CPU one. On the latter, each core is connected to a large *cache*, whereas a GPU generally has smaller caches but more registers. A cache is a very fast storage for data re-usage. On the CPU a thread can operate on its registers in a vector-like fashion (*Single Instruction Multiple Data (SIMD)*) and usually there are as many threads as physical cores. As opposed to that, a GPU is optimised for high throughput and executes thousands of threads simultaneously (*Single Instruction Multiple Threads (SIMT)*). This allows for hiding the latency¹⁴ of one thread effectively. Here, each thread sees a scalar register. Those are dynamically allocated to the threads, depending on memory demands. Multiple threads are executed in lockstep. Thus, SIMT is quite like SIMD, but with a more dynamical register usage.

The GPU has a *global memory* (like *Random Access Memory (RAM)* in a CPU system) which is typically 1-6 GByte and may be accessed by all threads. This is an obvious limitation of GPU usage as it is much smaller than current RAM in CPU systems. A core usually has a local memory that can be accessed by all threads running on it. If the data a thread has to process does not fit into its registers, it can swap parts of it into so-called *scratch registers* located in global memory. To those the access is of course drastically slower.

3.4.1 GPU Programming

Another drawback of GPU usage is that existing software has to be modified or rewritten to be able to incorporate them. The *OpenACC* standard¹⁵ allows for OpenMP-like directives in existing code, signalling to the compiler which parts should be performed on the GPU. However, as LQCD needs low level optimisation for optimal performance, such high-level approaches are not really promising [Majumdar, 2013].

13 In fact, in the beginning professional GPGPUs come without a connection for graphical output.

14 The time a thread has to wait for a memory access.

15 See <http://www.openacc-standard.org/> .

<pre> 1 void saxpy(2 spinor * in1, 3 spinor * in2, 4 complex * c, 5 spinor * out) 6 { 7 8 9 for(int idx = 0; 10 idx < VTOT; 11 idx += 1) 12 { 13 spinor tmp = 14 spXCplx(in1[idx], c); 15 out[idx] = 16 spAdd(in2[idx], tmp); 17 } 18 } </pre>	<pre> 1 __kernel void saxpy(2 __global spinor * in1, 3 __global spinor * in2, 4 __global complex * c, 5 __global spinor * out) 6 { 7 int ID = get_global_id(0); 8 int GS = get_global_size(0); 9 for(int idx = ID; 10 idx < VTOT; 11 idx += GS) 12 { 13 spinor tmp = 14 spXCplx(in1[idx], c); 15 out[idx] = 16 spAdd(in2[idx], tmp); 17 } 18 } </pre>
--	---

(a) C version

(b) OpenCL version

Figure 7: The saxpy operation in C and OpenCL version. `spinor` denotes a $N_D * N_c$ -component complex vector. The functions `spXCplx` and `spAdd` denote spinor times complex and the addition of two spinors, respectively. See text for further explanations.

As its origin in graphical applications implies, a GPU is designed for parallel tasks, i.e. it is well-suited for LQCD, in particular regarding the substantially higher peak BW compared to a CPU. The first reported LQCD simulation on a GPU was *Lattice QCD as a video game* by Egri et al [Egri et al., 2007]. They directly programmed the GPU using *Open Graphics Library* (OpenGL)¹⁶. Today, the programming can be done much more user-friendly by C-like languages. There are two prominent frameworks which will be briefly introduced in the following.

CUDA by the NVIDIA corporation [NVIDIA, 2013] allows for the incorporation of GPU-code at a high-level in the programme. It also comes with a variety of well-tuned libraries for various scientific areas. In particular, there is a library for LQCD called *QUDA*.¹⁷ Consequently, almost all LQCD applications are based on CUDA [Babich et al., 2011; Babich et al., 2010; Bonati et al., 2011a; Bonati et al., 2010; Clark, 2009; Clark et al., 2010; Clark and Babich, 2012; Joó et al., 2013]. A significant drawback is that CUDA is destined to run only on NVIDIA hardware.

An open standard is defined by the *Open Computing Language* (OpenCL) [Khronos Working Group, 2013], which is consequently vendor, but also hardware independent, i.e. OpenCL code

¹⁶ See <http://www.opengl.org/> .

¹⁷ See <http://lattice.github.io/quda/> .

runs equally on CPUs and GPUs. An implementation of OpenCL can be found in AMD's *Accelerated Parallel Processing (APP)*¹⁸, formerly ATI Stream SDK, and as part of CUDA.

Both frameworks have slightly different nomenclature, the OpenCL one will be explained in more detail now. For more extensive introductions to OpenCL and CUDA see [AMD, 2013] and [NVIDIA, 2013], respectively. As OpenCL is hardware-agnostic, it executes code on a *computing device*, of which there can be numerous in the system. Therefore, OpenCL is principally well-suited for heterogeneous systems, in contrast to CUDA. All actions performed on the devices are controlled and scheduled via a *host* programme. Also the memory management is carried out via the host. A memory object is called a *buffer*.

The GPU architecture is reflected in the mapping of the threads to the hardware: The total number of threads, `global_size`, is equally divided into `num_groups` *working groups*, each consisting of `local_size` threads. A working group is executed on one core (also called *Compute Unit (CU)*). Within a workgroup, kernel execution can be synchronised. Each thread can be uniquely identified by a global *thread id*.

A general function executed on a device is called a *kernel*. It is basically an independent programme, which is executed by all threads on the device. All kernels to be executed are put in a *queue*. By default, a kernel is compiled at runtime of the programme.

OpenCL provides a subset of the C99 language standard to write the kernel code. Figure 7 contrasts a (serial) C version of the linear algebra operation `saxpy` to the OpenCL one (`saxpy` adds a spinor times a complex number to another spinor). Basically, both versions differ only in the loop over the lattice volume, where the parallel nature of the kernel becomes obvious. As the kernel operation should be applied only once to each site of the lattice, one has to map the threads to the sites in a unique fashion. Here, this is simply done by taking the number identifying each thread globally, its thread id, using `get_global_id(0)`. Consequently, the loop increment must be the total number of threads, given by `get_global_size(0)`.¹⁹ Note, however, that this looping is the optimal strategy on GPUs, while on a CPU generally the optimal cache usage is achieved if each thread works on consecutive memory [Bach et al., 2013a]. Other differences visible in Figure 7 are the OpenCL-specific qualifiers: `__kernel` signals to the compiler that the respective function is a kernel and `__global` shows that the arguments are pointers to buffers in global memory. There are more differences between C and OpenCL code which are not shown in this example. An important one should be mentioned here: For some operations, like the scalar product, a parallel reduction has to be performed, much like in MPI. As was mentioned before, threads can be synchronised inside a working group only. Therefore, the reduction needs two kernels: One which performs the reduction on all threads of a working group and one which collects these results. The latter is essentially a serial function.

Additional overhead comes along with OpenCL regarding kernel-calling. Suppose that it is desired to calculate $z=x*a+y$ on given memory objects. Whereas the function described in Figure 7a would be simply called as `saxpy(x,y,a,z)`, this is more tedious in OpenCL. Figure

¹⁸ See <http://developer.amd.com/SDKS/AMDAPPSDK> .

¹⁹ In this function, the 0 argument denotes that all threads are organised in one dimension. In general, it is possible to do this in up to three dimensions.

8 shows the principal procedure. First, a kernel object has to be initialised and the kernel must be compiled. To execute it, the arguments must be set as shown, as well as `local_size` and `global_size`. In the end, the kernel object has to be released again; of course the kernel can be executed multiple times before the release. Note, that the memory management is not shown: All buffer objects have to be properly initialised and lie in global memory. Using CUDA is similar, however, the provided high-level functionality mentioned above allows to hide overhead like in Figure 7b from the user. In addition, one can use C++ features inside kernels.

```

1 // pass total lattice volume as compile parameter
2 std::string compileOptions = "-D VTOT=256";
3 // collect source code files
4 std::string headerFiles = "operations_spinor.cl";
5 std::string sourceFile = "saxpy.cl";
6 // create kernel object and compile kernel code
7 cl_kernel saxpy =
8     createKernel(sourceFile, compileOptions, headerFiles);
9 // set kernel arguments
10 clSetKernelArg(saxpy, 0, sizeof(x), x);
11 clSetKernelArg(saxpy, 1, sizeof(y), y);
12 clSetKernelArg(saxpy, 2, sizeof(a), a);
13 clSetKernelArg(saxpy, 3, sizeof(z), z);
14 // execute kernel
15 clEnqueueNDRangeKernel
16     (queue, saxpy, numThreadsGlobal, numThreadsLocal);
17 // release kernel object
18 clReleaseKernel(saxpy);

```

Figure 8: Calling the kernel from Figure 7 given in `saxpy.cl` from a C++ environment on a 4^4 lattice. It is assumed that all objects used have been initialised properly and that `operations_spinor.cl` contains the functions needed inside the kernel. Note that `createKernel` denotes a wrapping function that passes the compile arguments to the actual OpenCL functions. See text for further explanations.

3.5 LATTICE QCD BASED ON OPENCL: CL^2 QCD

As is introduced in the previous section, modern clusters are often equipped with GPUs. Particularly such systems are available at Goethe University Frankfurt. Due to the NVIDIA bound LQCD applications available, utilising these AMD based systems requires new software. Since OpenCL constitutes a quite different approach to LQCD than existing software, we decided to start the new project CL^2 QCD from scratch instead of modifying an existing application.

It provides implementations of the HMC algorithm (Section 3.1.1) for twisted mass and pure Wilson fermions as well as the Heatbath algorithm (Section 3.1.2) for pure gauge the-

ory. Furthermore, the standard Krylov solvers CG and BiCGStab are supplied. In addition, executables for observable calculation are provided. More specifically, mesonic correlators for scalar, pseudo-scalar, vector and axial-vector channels and the chiral condensate using noisy estimators can be calculated (for both fermion discretisations, see Appendix A.5). Standard gauge observables can be computed. An application is the transport coefficient κ , which is a relevant observable for studies of the Quark-Gluon-Plasma, see [Philipsen and Schäfer, 2013] for further information. All parts of the simulation code are carried out in OpenCL and can be run on CPUs and GPUs. The software is highly tuned for AMD hardware (see Section 3.5.2) and has been utilised in the physics studies presented in Sections 4 and 5. File I/O for the gauge fields is provided in the *International Lattice DataGrid* (ILDG) format²⁰, allowing inter operation with other LQCD frameworks.

A report on the early stages and first results can be found in *Lattice QCD using OpenCL* [Philipsen et al., 2011]. Significant further development resulted in *Lattice QCD based on OpenCL* [Bach et al., 2013a], which is the main source of the presentation here. Thus, CL²QCD constitutes the first OpenCL based LQCD application for Wilson fermions. The latest proceedings can be found in [Bach et al., 2013b]. Here, the focus is on the utilisation of various GPUs in parallel (*Multi-GPU*).

As a reference implementation for the HMC we relied on the *tmlqcd* software suite [Jansen and Urbach, 2009], which represents the standard code for twisted mass fermions and was also used partly for the results presented in Section 4.²¹

In the following, details of the implementation are given. Afterwards, performance results are presented, in particular for the \mathcal{D} and HMC implementation.

3.5.1 Implementation Details

The host programme is written in C++11, allowing for independent work on different sections of the code. In addition, the object orientation of C++ is quite well suited for these purposes as there are many interdependencies between programme parts. Furthermore, future developments like additional fermion discretisations can be incorporated in a natural way.

Several levels of output have been implemented which can be controlled by compile- and runtime options (via command line). Thus, debugging information can be taken completely out of the code on the compiler level in case performance is crucial, i.e. production runs. Simulation parameter can be passed to the application either by input files or command line arguments. Details on the compilation procedure can be found in Appendix A.6.1. The aforementioned ILDG compatible I/O has been implemented utilising the *Lattice QCD Interchange Message Encapsulation*²² (LIME) scheme for C.

At several points, the implemented algorithms rely on random numbers. Numerically, these are generated using Pseudo Random Number Generators (PRNGs), of which RANLUX [Lüscher, 1994] is the standard choice, as it is designed to give reliable random numbers in

²⁰ See <http://cssm.sasr.edu.au/ildg/> .

²¹ See <https://github.com/etmc/tmLQCD> for the latest code version.

²² See <http://usqcd.jlab.org/usqcd-docs/c-lime/> .

parallel applications, too. The original implementation by Lüscher is used in the host code whereas an open-source OpenCL implementation of RANLUX, *RANLUXCL*²³, is used in the OpenCL parts. The initialisation follows the usual RANLUX rules on the host and the device. Each OpenCL thread runs on its own RANLUX PRNG state. In addition, the simpler NR3 [Press et al., 2007] generator is implemented, which was used for testing but not for simulations.

All mathematical functionality required for the algorithms has been implemented as OpenCL kernels. Kernel compilation is carried out as indicated in Figure 8. To simplify debugging, each kernel is built as stand-alone programme. To reduce the compilation time, the binary files created by the OpenCL compiler are saved for later re-use. Of course, this functionality ensures that current hardware and code version are not changed with respect to the former compilation. Already mentioned in Figure 7 was the possibility to pass compile time arguments in a C-like manner. This allows to save many kernel arguments (e.g. *NT*, *NS*, ...), which are “hard coded” into the kernel itself. To simplify kernel calling, the corresponding actions (see Figure 8) have been encapsulated in wrapper functions. All data types used in the kernels are implemented as structures [Bach et al., 2013a]. To access neighbouring sites, threads need to know the corresponding index based on their index (i.e. their global id). As especially with EO preconditioning this relation is non-trivial, traditional applications provide the indices by arrays. As on the GPU there is the limiting factor memory access, a functionality to calculate these indices directly was developed. This is explained in Appendix A.6.3.

CL^2 QCD went through a couple of major changes during the development process. At the beginning of the development, the target machine, the LOEWE-CSC, represented a quite heterogeneous system. Accordingly, one aim was to investigate possible hybrid applications of the given architecture, i.e. the usage of CPUs and GPUs together. As the most recent system used, SANAM, has a clear focus on GPU performance, the code was heavily restructured, setting aside hybrid usage for the moment. This was reasonable as the four AMD FirePro S10000 provided by each SANAM node are an ideal test ground for a Multi-GPU application. The two different structures will be presented in the following.

In the first code structure, the class `gaugefield` is the main object, holding the buffer of the gauge field and the OpenCL environment is initialised by it. The idea was that the physical gauge field is managed by one single object which in turn controls several devices in a possibly hybrid way. The devices are represented by `openc1_module`-objects that hold device related parts such as kernels and buffers. Problem-related functionality and algorithm logic is implemented by child classes of `gaugefield` and `openc1_module`, cf. Figure 9. For hybrid applications, “tasks” can be defined for specific problems, of which each may then contain various of device objects. Possible hybrid applications were reported in [Philipsen et al., 2011]. For example a splitting of fermionic observables like (2.79) into CPU (contraction) and GPU part (inversion) might be beneficial. However, this structural approach has some drawbacks: The fusion of many components (gauge field buffer, OpenCL environment...) into the `gaugefield` and `openc1_module` objects made the application inflexible, especially for Multi-GPU usage, where the gauge field has to be distributed among multiple devices. Also,

²³ See <https://bitbucket.org/ivarun/ranluxcl> .

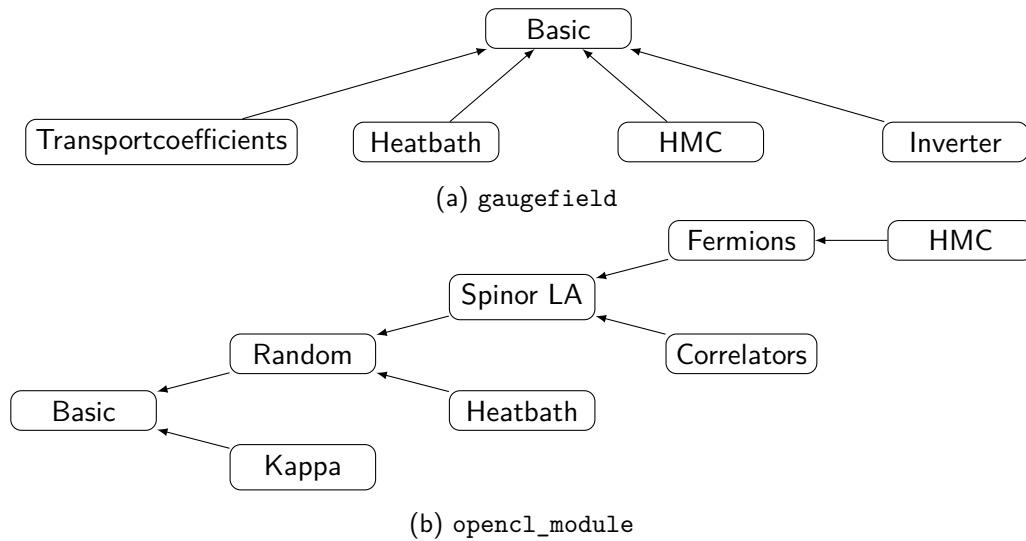


Figure 9: Structure of gaugefield and openc1_module class in the first version of CL²QCD.

the entanglement of the openc1_module objects, which came as a natural process during the development, is not a good way of coding as it misses a clean separation of different levels of the application. Furthermore, this hampered the programming process, as code changes were always affecting many parts of the programme immediately.

Accordingly, the code saw a major refactoring, which separated OpenCL initialisation, buffer management, and openc1_modules clearly into different objects. In addition, it introduced a clear distinction between high- and low-level functionality by introducing more abstraction, for example making the algorithm logic independent of the actual buffer type used and thus more versatile. This was achieved by creating three distinct packages: *Physics*, providing objects to represent gauge- or spinor fields (*Lattices*) and methods to work on (*Fermionmatrix*) or work with them (*Algorithm*). The *Hardware* package maps the objects used in the *Physics* package onto the actual hardware using the *Buffer* classes. The *System* object takes over the OpenCL initialisation, in particular it initiates as many *Device* instances as demanded. Similar to the first version, the kernels are combined problem-related in openc1_modules, however, stripped of the buffer and algorithmic content. These are collected in the *Code* namespace. Figure 10 shows the structure of the two packages. A third package, *meta*, holds objects that represent the parameters and provides a variety of helper functions.

In Section 3.3 it was argued that LQCD is always limited by memory bandwidth. Accordingly, for optimal performance on a GPU the memory characteristics must be somehow represented. By default, a memory object is stored as an *array of structures* (AoS), however, on a GPU a so-called *structure of arrays* (SoA) approach to data storage is favourable. The difference is depicted in Figure 11 using the example of a spinor object. This is a structure of twelve complex numbers. In the SoA pattern the elements of the structures lie consecutive in the memory, opposed to AoS, which holds one structure after the other. As GPUs execute multiple concurrent threads, these will simultaneously request the same element of a structure. However, the memory controller always schedules read access of a certain length. In the worst case, an




Figure 10: Simplified structure of the two main packages in the second version of CL²QCD. Note that not all objects are classes and possible subcontent is not shown.

AoS delivers only one structure element in such a read, needlessly throwing away the also-read data. This is improved in a SoA, where (given a suited alignment of the elements of the structure) all data per read is used. It turns out that this is crucial for high performance. On the CPU, in turn, the AoS is in general favourable against SoA. Both strategies are used in CL²QCD, depending on the used hardware.




Figure 11: Sketch of AoS against SoA pattern, showing a twelve component spinor object. Vertical lines indicate consecutive memory blocks. N denotes the number of elements of the vector of spinors.

On the GPU, the OpenCL compiler is part of the driver, and it is also responsible for optimisations of the kernel code. Thus, the compilation process is an important part of GPU usage. Since the beginning of the development, multiple miscompilations on various AMD driver generations have been observed. These would manifest themselves in compilation failure or incorrect code output, despite the fact that the correctness of the kernel code could be verified in other environments, e.g. on the CPU. The likelihood of such behaviour is enhanced with increasing code complexity. Furthermore, the compiler heavily influences the register usage of the compiled kernel code, which can often be much worse than what could be expected by kernel requirements and lead to usage of scratch registers. If this is the case, it can be extracted from the binary files produced by the compiler as these hold information about the compiled code, e.g. register usage statistics. As this influences performance massively, rewriting the kernel code is sometimes useful and may help to avoid such *register spilling*.

To ensure code correctness, we therefore incorporated regression tests for all kernels. These check kernel output against reference values for various sets of input parameters and compile options. In particular, these tests verify that execution on CPUs and GPUs yields the same results. In addition, similar tests for executables were added.

Apart from that, in particular the HMC and inverter executables were tested against the reference, tmlqcd. During code development, this was done by modifying the tmlqcd source

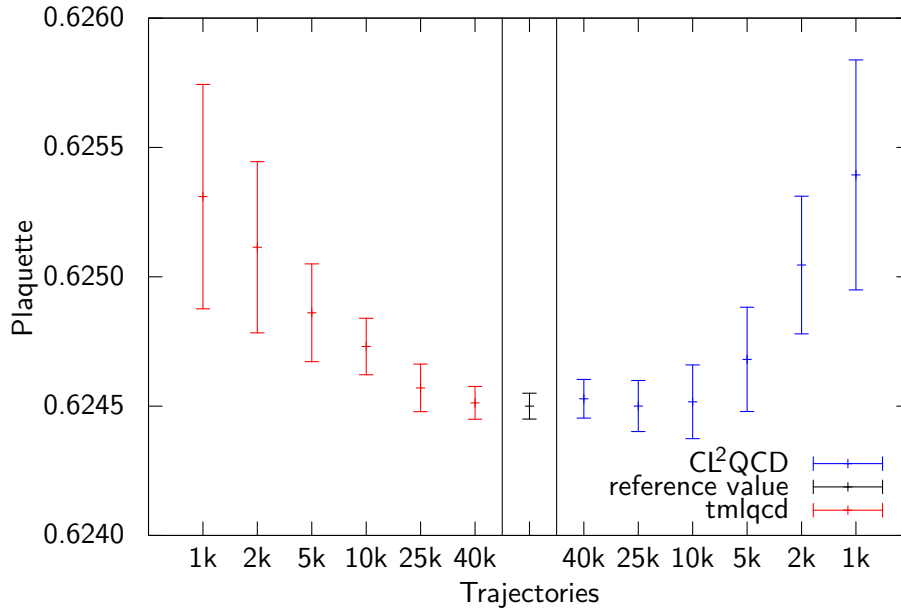


Figure 12: Plaquette as function of number of HMC trajectories, measured using the HMC implementations of tmlqcd and CL^2QCD , respectively. Parameters were used according to the sample input file `sample-hmc0.input` provided in the tmlqcd package, which constitutes a test setup with quoted reference value for the plaquette.

code itself to give required interim results. The full HMC of CL^2QCD was then tested against the corresponding one from tmlqcd using reference input provided by the latter. Figure 12 shows this exemplarily for one test case. One sees that both programmes give comparable results and converge to the reference value quoted for this setup in the tmlqcd package. The differences are due to the different random numbers used. The HMC without fermions can additionally be used to verify correctness of the heatbath implementation. Moreover, the inverter was tested, for example by comparing the chiral condensate calculated on a specific configuration. Agreeing results were obtained with reference values provided by my colleague Florian Burger.

For details on the installation process of CL^2QCD , see Appendix A.6.1.

3.5.2 Performance Results

In this section performance results for the implementations of the heatbath and HMC algorithm described above are given. For additional details see [Bach et al., 2013a] and [Bach et al., 2013b]. Special emphasis is placed on the \mathcal{D} as it is the crucial part for fermionic observables and the HMC itself. To show the versatility of OpenCL, different architectures are utilised and various lattice sizes are benchmarked. In terms of performance, single precision (SP) can be expected to give better results (see Table 3 or also [Gerhard et al., 2012]), but, for physical measurements a higher precision is desired. Hence, all benchmarks are done with DP, except

for the heatbath, where SP was used. In fact, since no summation over the whole lattice is needed here, the impact of the lower precision can be expected to be less severe.

Results were obtained on a variety of GPUs. In particular, the AMD Radeon HD 5870 and FirePro S10000 were benchmarked, as they are used in LOEWE-CSC and SANAM, respectively. Furthermore, the AMD GPUs Radeon HD 6970 and Radeon HD 7970 were also benchmarked. GPU vendors usually offer professional and consumer versions of the cards. The differences between these are different for specific GPUs. Professional cards are often clocked somewhat lower, equipped with more memory or, for example, provide error correction methods like ECC. This prevents data loss at transportation within the hardware at the price of lower data throughput. And, most importantly, professional GPUs are factors more costly than the consumer variants. For AMD, the Radeon cards are consumer GPUs, hence they are cheaper in acquisition than their professional counterparts, the FirePro GPUs. Results obtained on NVIDIA GTX 680 are also shown. In addition, the Intel Xeon E5520 and AMD Opteron 6278 CPUs were benchmarked. These are all server CPUs, which come with much higher costs of acquisition than some of the (consumer) GPUs used. Server CPUs, similar to the professional GPU variants, are designed for non-consumer targets like in supercomputing environments.

In [Bach et al., 2013a] all benchmarks on AMD GPUs were performed with the Catalyst 12.4 (Catalyst 12.2 for the AMD Radeon HD 7970). In [Bach et al., 2013b], a beta version of release 12.102 of the AMD FirePro variant of the Catalyst driver pre-released to my colleague Matthias Bach was used. This was used measuring the performances \mathcal{D} and HMC presented below. The driver is publicly available with the Catalyst 13.6 driver. The NVIDIA GPU was benchmarked using version 295.41 of NVIDIA's GPU driver.

Heatbath performance

The performance of the heatbath implementation is addressed in Figure 13, which shows its main ingredients, the *heatbath* and *overrelax* kernels, as a function of the lattice size on various devices. First of all, a noticeable difference between GPU and CPU performance can be seen. The CPU performance shows that the code can be executed on different architectures. As there were no explicit optimisations, lower performance can be expected. Furthermore, only SP results are depicted, since in DP both kernels achieve $\mathcal{O}(10)$ GFLOPS only. This can be explained by the complexity of the functions, which means a large working set²⁴ and leads to huge register spillings in DP. It shows that the algorithmic density ρ_{alg} has to be taken with care, as it does not consider implementation details such as the size of the working set. In SP this is not such a problem and the performance is much better. Still, there is an obvious difference between the two kernels, which can be explained by the random numbers required for the heatbath kernel. The random numbers mean an increased code complexity. The difference in performance of the two kernels makes evident to what extent the code complexity is important for the compile process, as mentioned above. While the heatbath kernel can at least provide 30 GFLOPS for all lattice sizes, the overrelax (which does not need random numbers) achieves up to 160 GFLOPS. The different performance of different card generations is also visible.

²⁴ The amount of data needed throughout kernel execution.

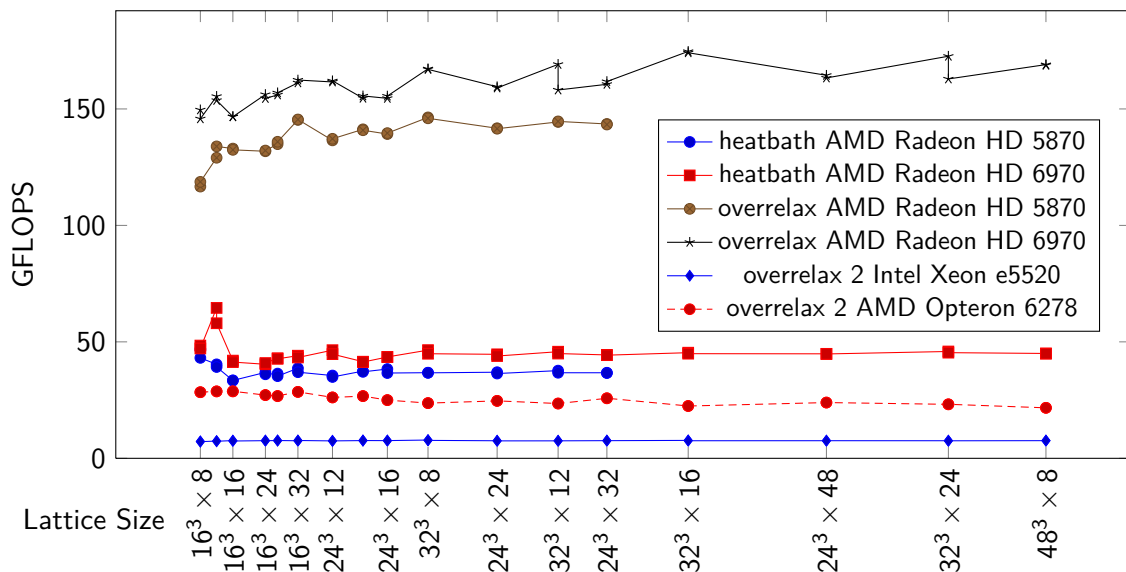


Figure 13: Performance of the SP heatbath and overrelax kernels in GFLOPS as a function of the lattice size. Not shown are CPU results of the heatbath kernel, which are a bit lower than those from overrelaxation, but of the same order of magnitude.

In literature, there are reported performances of a CUDA based implementation, given in [Cardoso and Bicudo, 2011]. Compared to these results, which were obtained on NVIDIA GeForce GTX 295 and NVIDIA GeForce GTX 580 GPUs, the heatbath kernel of CL^2QCD achieves only half of the performance. The overrelax kernel, in turn, actually performs slightly better on the AMD Radeon HD 6790 and achieves around 90% of the peak value of [Cardoso and Bicudo, 2011] for the AMD Radeon HD 5870. One has to mention that for the results presented here no memory bandwidth reducing techniques like REC12 are used as in [Cardoso and Bicudo, 2011] and that [Cardoso and Bicudo, 2011] relies on a highly improved CUDA library for random numbers (*cuRAND*), which, however, does not supply a RANLUX implementation. Additionally, [Cardoso and Bicudo, 2011] states that their implementation is limited in terms of lattice size by the maximal number of threads, whereas CL^2QCD is more flexible as the lattice size is limited by GPU memory only.

Despite the high performance, GPU memory is clearly limiting physical studies in pure gauge theory. One clearly sees the impact of the bigger memory of the Radeon HD 6970, however, the lattice extents reached are by far smaller than current pure gauge theory studies (see e.g. [Philipsen and Schäfer, 2013]). Further improvements are beyond the scope of this thesis. Nevertheless, this shows that all the tested GPUs can be used to efficiently run the algorithm in principle.

\mathcal{D} performance

As argued above, the performance of \mathcal{D} is crucial for the overall performance. Measurements as presented in [Bach et al., 2013b] are shown in Figure 14 for the Radeon HD 5870, Radeon

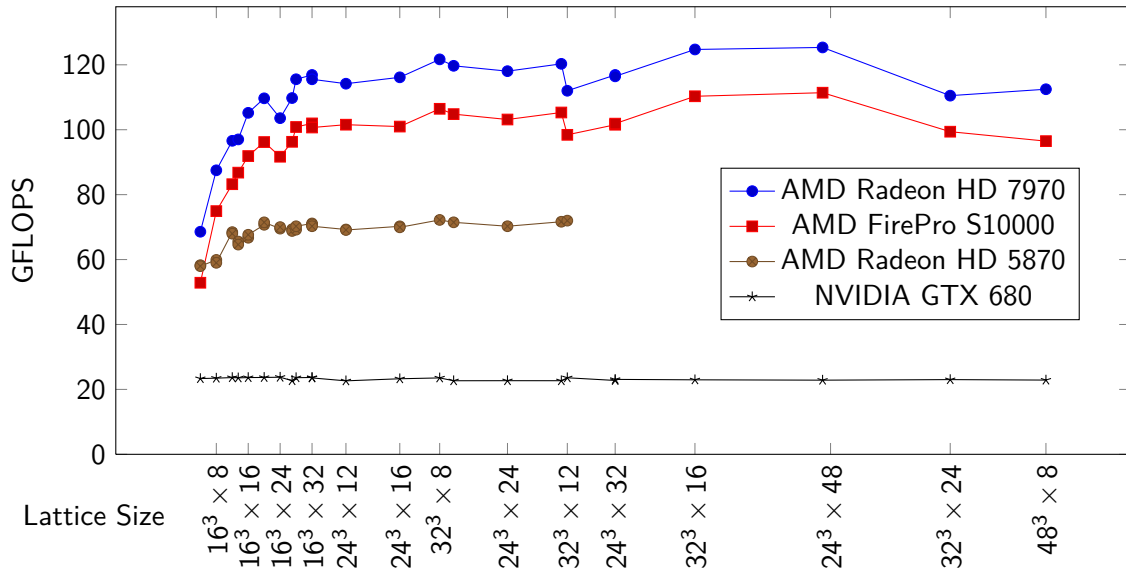


Figure 14: Performance of the \mathcal{D} kernel on several GPUs (DP).

HD 7970 and the FirePro S10000 by AMD as well as for the NVIDIA GTX 680. A somewhat lower performance on smaller lattices is visible. Apart from that, the Radeon HD 5870 achieves around 70 GFLOPS, the Radeon HD 7970 around 120 GFLOPS and the FirePro S10000 around 100 GFLOPS on a variety of lattice sizes. The better performance of the Radeon HD 7970 can be explained by its slightly better specifications compared to the dual GPU FirePro S10000, which is, however, of the same generation. Clearly, the newer GPUs perform better and the impact of bigger memory is visible. A better estimate of the performance is given by the bandwidth (BW) utilisation, as the \mathcal{D} kernel is bandwidth limited. It can be extracted from the performance easily, the conversion factor is just ρ_{alg} . The AMD GPUs achieve 70-80 % of their respective peak BW usage, which is a very satisfying result and shows the benefit of the optimisation techniques used [Bach et al., 2013a,b]. A remark has to be given with respect to [Bach et al., 2013a]: The results presented there are a bit lower for the Radeon HD 5870 and Radeon HD 7970. Whereas the \mathcal{D} on the newer Radeon HD 7970 GPU can be expected to benefit greatly from the newer driver, the increase in performance also on the older Radeon HD 5870 GPU makes explicit that the driver is an essential ingredient to performance. As expected, the impact of REC12 results in a slightly higher performance ($\approx 9\%$) [Bach et al., 2013a]. Nonetheless, the increased code complexity generally results in higher register usage and, more importantly, can lead to miscompilations (see above), so this technique has not been used heavily yet.

On the NVIDIA GPU, \mathcal{D} fails to reach the same performance. This is caused by massive register spilling, which has several reasons: This card generally has smaller registers than the AMD cards. Furthermore, the used OpenCL implementation provided by CUDA does not allow for a reduction of the maximal number of threads, which results in suboptimal register usage.

setup	A	B	C
$a\hat{\mu}$	0.0025	0.0035	0.1
m_π	260	310	520

Table 4: Setups used for benchmarking. $\beta = 3.9$ and $\kappa = \kappa_c = 0.160856$ were used throughout to simulate at maximal twist [Baron et al., 2010]. The value of m_π is only approximate. 10 HMC steps, each with 10 integration steps of the 2MN integrator, were performed, with $\tau = 0.1$ and 1 for setups A, B and C, respectively. CG was used throughout.

And, of course, NVIDIA was not the primary development platform, even though CL²QCD runs on this hardware, too.

\mathcal{D} DP performances can be found in literature in CUDA implementation, but only for older generations. Clark et. al [Clark et al., 2010] show DP results using REC12, which are around 75 % slower than those from CL²QCD. Yet, the used NVIDIA GTX 280 is one generation older than the AMD Radeon HD 5870. A fairer comparison can be drawn to [Alexandru et al., 2012], where a performance of 50 GFLOPS is reported on an NVIDIA GTX 480. This card is comparable to the Radeon HD 5870, the performance is yet 40 % slower. Hence, for similar generation GPUs the OpenCL based code on AMD GPUs outperforms CUDA based code from the literature. Results on modern NVIDIA GPUs are found only for *mixed-precision* solvers, where the majority of steps is carried out in SP [Clark and Babich, 2012; Joó et al., 2013]. Thus, a fair comparison is not really possible. However, the achieved BW usage shows that the \mathcal{D} is very well tuned, independent of the GPU generation. It can therefore be expected to perform equally or better to NVIDIA GPUs of the same generation.

HMC performance

The very good performance of the \mathcal{D} carries over to the full HMC. To study this, setups of different m_π have been used (see Table 4), starting from the same generated gauge configuration. Statistical errors on runtime were found to be negligible. Performance results are presented in [Bach et al., 2013a,b]. Reference values have been obtained with the aforementioned tmlqcd code on LOEWE-CSC, using the same setups and various numbers of cores. Note, that CL²QCD can be used on CPUs, too, showing again the versatility of OpenCL.

A LOEWE-CSC node consists of two CPUs, supplying a total of 24 cores per node (see also Appendix A.7). The reference code was run on one of these nodes, i.e. on 24 cores. The MPI used for communication in the reference code scales nicely on LOEWE-CSC and the communication overhead can be neglected. Figure 15 shows performance results for the AMD Radeon HD 5870 and the AMD FirePro S10000 compared to tmlqcd for one specific lattice size and varying pion masses. The observed behaviour is quite independent of the pion mass: The AMD Radeon HD 5870 is approximately twice as fast as the reference setup, while the newer AMD FirePro S10000 gives a speedup of four. Similar results were obtained on different lattice sizes and different GPUs [Bach et al., 2013a]. Putting this into a broader perspective, this means that on LOEWE-CSC, where nodes are equipped with the AMD Radeon HD 5870, the performance of four CPUs or two nodes can be achieved with one single GPU. On SANAM,

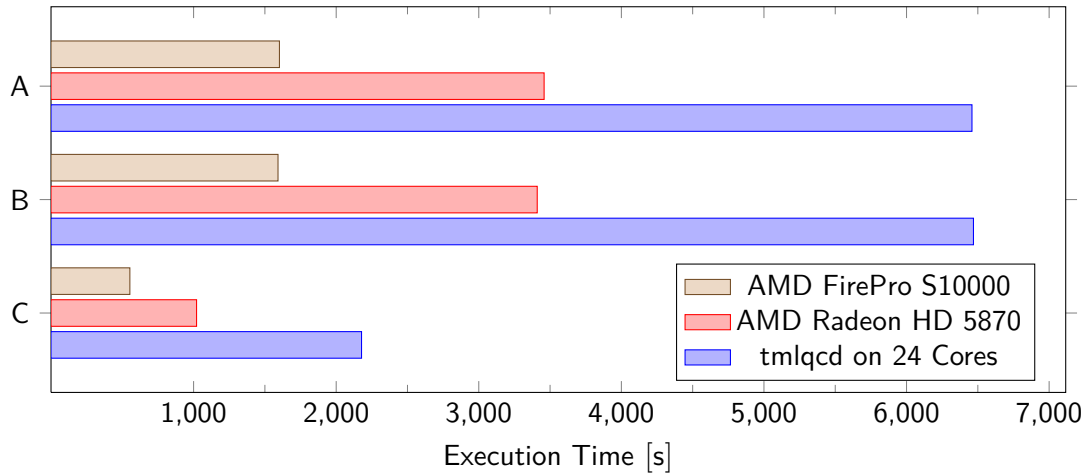


Figure 15: HMC runtimes in seconds for setups A, B, and C (see Table 4) for a lattice with $N_\tau = 8$ and $N_\sigma = 24$.

where one has four AMD FirePro S10000 per node, the speedup is even larger: Four nodes LOEWE-CSC CPU nodes are equivalent in performance to one GPU (or 16 LOEWE-CSC nodes to SANAM node). Given the acquisition costs of the individual hardware components, this underlines the power of GPU usage for LQCD simulations.

Note that these results should always be viewed keeping the memory limitations of the GPUs in mind (especially of the Radeon HD 5870, which has only 1 GByte of global memory). Furthermore, for higher pion masses the inverter becomes less important (Amdahl's Law), making the force computations (especially the gauge force) more important.

As was mentioned before, [Bach et al., 2013a] was the first report on a LQCD application for Wilson fermion using OpenCL. Thus, comparing the HMC to reference codes is somehow difficult, as there are no published numbers for Wilson fermions using CUDA. Extensive tests of a CUDA based HMC for staggered fermions are reported in [Bonati et al., 2011a]. They also report on OpenCL performance, and see that OpenCL version of the code performs slightly better on AMD than on NVIDIA hardware. However, compared to the code's CUDA version on NVIDIA GPUs, performance is worse. This might as well be explained by the choice of the primary development platform. In fact, the \mathcal{D} results reported above strongly suggest that the HMC of CL^2QCD would not perform equally on NVIDIA cards without hardware specific optimisations. Of course, a comparison is quite difficult because staggered fermions are fundamentally different, both, numerically and algorithmically. Actually, they have a much smaller working set, and therefore should work better on the NVIDIA GPUs used in [Bonati et al., 2011a] than Wilson fermions. As a further difference the HMC from [Bonati et al., 2011a] mainly uses SP within the HMC. Furthermore, it should be noted that [Bonati et al., 2011a] reports a decline in performance for larger lattices which is not seen in CL^2QCD .

3.6 SUMMARY AND DISCUSSION

In this chapter the focus has been laid on the numerical aspects of LQCD. They constitute an important aspect of each LQCD investigation. In particular, the numerical algorithms available (cf. Section 3.1) set certain limitations to the parameter ranges. For example, it is discussed how the infamous sign-problem hampers the exploration of QCD at finite (real) chemical potential μ . Another, more generic limitation is by means of computational costs, cf. (3.13). To this end, since recent years *Graphics Processing Units* GPUs have been used within the simulations. As is elaborated in Sections 3.4 and 3.3, the GPU architecture is well-suited for LQCD purposes. In particular, they outperform traditional CPUs in terms of memory bandwidth by far (cf. Table 3). Today the utilisation of GPUs is made possible by certain programming languages, most prominent NVIDIA's CUDA [NVIDIA, 2013]. Virtually all GPU-enabled LQCD applications are based on CUDA. However, it is restricted to NVIDIA hardware, which in particular does not allow to apply them on the AMD-based GPU clusters available in Frankfurt. An alternative is provided by the OpenCL standard [Khronos Working Group, 2013], which is not vendor-restricted. The utilisation of CPUs by the same code is also possible in this environment.

In the course of this thesis work, the LQCD application CL²QCD²⁵ has been developed and was reported on in [Bach et al., 2013a]. While the actual functions are realised in OpenCL, the algorithm logic, parameter handling and similar functionality is implemented in C++. CL²QCD constitutes implementations of the heatbath algorithm for $SU(3)$ pure gauge theory (cf. Section 3.1.2) as well as of the HMC algorithm for two flavour twisted mass and pure Wilson fermions (cf. Section 3.1.1). For this, standard simulation techniques like even-odd preconditioning, multiple timescale integration or mass preconditioning have been implemented. CL²QCD constitutes the first OpenCL based application for Wilson type fermions. Beside executables for the generation of gauge field configurations, executables for the measurement of various observables are available. In addition, ILDG (International Lattice Data Grid) compatible I/O is provided as well as RANLUX, the standard random number generator for LQCD. Thus, CL²QCD is well-equipped for LQCD studies of various kind. Details of the implementation are given in Section 3.5.

The implementation of the $SU(3)$ heatbath algorithm in single precision currently performs comparable to a CUDA-based implementation from the literature [Cardoso and Bicudo, 2011] for the overrelax kernel. In the heatbath kernel, the random numbers needed seem to hamper the performance due to the increased code complexity. This can be optimised by using simpler random number generators instead of RANLUX, which itself is quite complex. In double precision, its performance can also be optimised. However, the focus of performance optimisation was on the fermionic applications, in particular the HMC, as CL²QCD has been used for the studies presented in Chapters 4 and 5. Here, the \mathcal{D} implementation is crucial for overall performance. For a large variety of lattice sizes, the latter is shown to utilise 70 – 80% of the available bandwidth for those AMD GPUs used in the actual simulations. This is an excellent performance, outperforming published literature values [Alexandru et al., 2012; Clark et al.,

²⁵ See <http://code.compeng.uni-frankfurt.de/projects/clhmc> .

2010]. This translates nicely to a full HMC application. CL²QCD on one GPU outperforms the reference code running on two (server) CPUs by a factor of four in total runtime for the same simulation on the GPU used in LOEWE-CSC. The speedup is even higher for the SANAM GPU, where CL²QCD is four times faster than the reference code. This underlines the great potential of GPUs for LQCD applications. In addition, it is mentioned that GPUs are beneficial also in terms of acquisition.

 THE THERMAL TRANSITION IN TWISTED MASS LATTICE QCD

In this chapter the extension of a study by the tmft collaboration aiming at the two flavour chiral limit [Burger et al., 2013a] is presented. The order of the chiral transition in this limit has not been clarified yet, it might be first or second order. The study used twisted mass Wilson fermions at maximal twist with three ensembles at different pion masses. For each ensemble, the chiral transition was found to lie in the crossover region. Hence, the available pion masses m_π were found to be still too high to settle the issue. Consequently, a new ensemble at lower m_π was generated and added to the analysis. The data was gathered mainly relying on the LQCD application CL²QCD described in detail in Chapter 3. The chiral transition is addressed by means of its order parameter, the chiral condensate $\langle\bar{\psi}\psi\rangle$, as well as by investigating the behaviour of screening masses of suitable parity partners. These might indicate a restoration of the $U_A(1)$ symmetry, supporting the first order scenario.

This chapter is organised in the following way. At the beginning, the present status of LQCD investigations of the chiral transition is reviewed and the so-called Columbia plot is introduced. Then, the investigations of this transition by the tmft collaboration is discussed in Section 4.2. Details on the scale setting method and the simulation setup used in these studies are elaborated in Sections 4.3 and 4.4. The extension of the existing studies by a new ensemble at pion mass of approximately 270 MeV and implications for the chiral transition are presented in Sections 4.5 and 4.6, respectively. The chapter is concluded with a summary and discussion.

4.1 THE COLUMBIA PLOT

The strange quark is approximately thirty times heavier than the up and down quarks. Accordingly, it is justified to approximate the light quarks as mass-degenerate with mass m_l . In theory, the masses of the quarks can be varied, thus, allowing to study unphysical scenarios. These are usually called $N_f = 2, 3$ for the mass-degenerate case and $N_f = 2 + 1$ for the more physical case. Depending on the masses, the phase structure changes. The current knowledge about this at $\mu = 0$ is summarised in the so-called *Columbia Plot* (Figure 16a), combining deconfinement and chiral transitions. Their distinction is quite important, since both are only expected to be true phase transition in (regions around) the infinite mass and chiral limit, respectively. Elsewhere, the transitions are crossover. However, in the crossover region they are found to occur typically around the same temperature. Overviews are given in [Laermann and Philipsen, 2003; Levkova, 2011; Petreczky, 2012].




Figure 16: a) Columbia Plot. b) Alternative Columbia Plot for the first order scenario. See text for further details.

Definite statements can be made for the limiting cases of vanishing and infinite quark mass. In pure gauge theory (upper right corner in the plot), it is argued in Section 2.2.2 that centre symmetry is broken at some T_c as the temperature is increased. In simulations T_c is determined to be ≈ 270 MeV and the transition is first order [Boyd et al., 1996; Fukugita et al., 1990]. Consequently, there is a true order parameter, which is the Polyakov loop. For $N_f = 3$, the chiral transition in the chiral limit (lower left corner) must be first order [Pisarski and Wilczek, 1984]. With only one flavour of quarks ($N_f = 1$) the same is a crossover, as all symmetries are just phase factors and no breaking can occur [Pisarski and Wilczek, 1984]. Furthermore, it is well established that at the physical point the chiral/deconfinement transition is a rapid crossover [Aoki et al., 2006]. These first order and crossover regions have to be separated by second order lines, which are found to fall in the $Z(2)$ Ising universality class [Gavin et al., 1994; Karsch et al., 2001].



Figure 17: Possible scenarios for the transition in the chiral limit for $N_f = 2$. Figure 18: Chiral partner particles (following [Bernard et al., 1997]).

For the $N_f = 2$ chiral point (upper left corner) it is not clear if the (chiral) transition is first or second order [Pisarski and Wilczek, 1984]. This depends on the fate of the axial anomaly that breaks the $U_A(1)$ symmetry (see Section 2.2). This symmetry will eventually be restored effectively by thermal effects [Pisarski and Wilczek, 1984].

For the chiral transition this constitutes two scenarios. In the first case, $U_A(1)$ is still broken at T_c and the symmetry restored at this temperature is

$$SU_L(2) \times SU_R(2) \simeq O(4) . \quad (4.1)$$

Thus, the transition in this case is in the universality class of the 3D $O(4)$ model, which is of second order. This situation is widely believed to be the correct one and is supported by various studies (see e.g. [Brandt et al., 2013; Buchhoff et al., 2013; Halasz et al., 1998; Parisen Toldin et al., 2003; Rajagopal and Wilczek, 1993]). In the second case the $U_A(1)$ is restored or only weakly broken at the chiral transition. Then, the symmetry group reads:

$$U_L(2) \times U_R(2) \simeq U_{L+R}(4) \quad (4.2)$$

and the chiral transition is most likely first order [Parisen Toldin et al., 2003].¹ The restorations of both symmetries would appear in the particle spectrum as a mass degeneracy of the respective particles (see Figure 18). Similar to the indications of the other corners, the nature of this transition has an impact on the Columbia plot and the phase diagram of QCD if the physical point is close to the scaling region.

First, the $O(4)$ case is discussed. When moving away from the chiral limit for $N_f = 2$ by switching on a finite quark mass in the light sector, the transition is immediately a crossover (Figure 17) as the symmetry is broken explicitly. Lowering m_s from infinity does not have this effect, the symmetry is still intact. Consequently, there is an $O(4)$ line descending the $m_l = 0$ axis (left axis in the Columbia plot). Furthermore, the $Z(2)$ second order line enclosing the first order region around the $N_f = 3$ chiral limit must end on the same axis. Both lines meet in a tricritical point at strange quark mass m_s^{tric} . It is important to know if the physical point lies lower than the tricritical point or not, since then one might encounter a first order transition when going to the chiral limit, this is subject to ongoing studies.

In addition, considering the $O(4)$ scenario, predictions can be made for finite chemical potential μ , which does not break chiral symmetry. Thus, chiral QCD is similar to a metamagnet if h^\dagger is identified with μ (cf. Section 2.6). Indeed, this has led to the conjecture of a tricritical point in $N_f = 2$ QCD at $(m_l = 0, \mu \neq 0)$, which is located at finite μ and vanishing m , supported by model studies [Halasz et al., 1998]. In addition, this analogy implies a second order CEP at finite quark mass, because of the second order line emerging from the tricritical point extending in finite m direction. Finally, this has led to the conjecture of a CEP in physical QCD.

In general, at finite μ , the second order lines in the Columbia plot become surfaces. Thus, a second possibility for a CEP at physical mass to exist at non-zero μ is given if this surface bends towards the physical point, which, however, is not seen in LQCD studies [de Forcrand and Philipsen, 2008]. This will be commented on in the next section, see also Figure 47.

In the alternative first order scenario, the two flavour chiral transition persists to be of first order also at non-vanishing quark mass, eventually ending in a $Z(2)$ CEP at some nonzero value of the light quark mass (see Figure 17). This is necessary as the transition is known to be a crossover at finite values of the quark mass [Burger et al., 2013a]. At finite strange

¹ Other scenarios are also possible, depending on the strength of the anomaly, see [Butti et al., 2003].

quark mass, a $Z(2)$ line extends from this point and meets with the corresponding $Z(2)$ line originating from the $N_f = 3$ chiral region. Hence, there is no tricritical point in this scenario, instead the Columbia plot would look as depicted in Figure 16b.

To conclude, the situation in the $N_f = 2$ chiral limit is expected to have strong implications for the overall phase diagram at zero and non-zero chemical potential μ . The $O(4)$ scenario is favoured by many studies, however, it must be noted that the first order scenario is not ruled out. In fact, the latter is also favoured by recent studies, e.g. [Bonati et al., 2013].² Settling this issue is thus of great importance especially for ongoing experiments searching for a CEP in the QCD phase diagram.

4.2 INVESTIGATIONS OF THE THERMAL TRANSITION WITH $N_f = 2$ TWISTED MASS LQCD

An extensive study aiming at the chiral limit has been performed by the tmft collaboration [Burger et al., 2013a] using twisted mass Wilson fermions at maximal twist (see Section 2.5.2) and the tree level Symanzik improved gauge action (2.56). In particular, three different ensembles at $N_\tau = 12$ have been created, each at a fixed m_π in the range of approximately 316 to 470 MeV (see Table 5). The breaking of chiral symmetry is indicated by a non-vanishing

Ensemble	$N_\sigma^3 \times N_\tau$	β range	m_π [MeV]
A12	$32^3 \times 12$	$3.84 \leq \beta \leq 3.99$	316(16)
B12	$32^3 \times 12$	$3.86 \leq \beta \leq 4.35$	398(20)
C12	$32^3 \times 12$	$3.90 \leq \beta \leq 4.07$	469(24)

Table 5: $N_\tau = 12$ runs performed in the tmft study [Burger et al., 2013a].

chiral condensate $\langle \bar{\psi}\psi \rangle$. As was described in Section 2.6, the universal behaviour of chiral QCD can be described in terms of magnetic systems. In particular, $\langle \bar{\psi}\psi \rangle$ and its susceptibility $\chi(\langle \bar{\psi}\psi \rangle)$ are subject to universal scaling laws within the scaling regime, the magnetic equation of state (2.87) and (2.88). This can be used to deduce scaling relations for the critical coupling β_c ,

$$\beta_c(h) = \beta_{\text{chiral}} + Bh^{1/(\delta\tilde{\beta})}, \quad (4.3)$$

and the critical temperature T_c ,

$$T_c(m_\pi) = T_{\text{chiral}} + Am_\pi^{2/(\delta\tilde{\beta})}. \quad (4.4)$$

Here, A and B are constants and β_{chiral} and T_{chiral} represent the respective critical values in the chiral limit. m_π and the external field $h = 2a\hat{\mu}$ are related as $h \sim m_{\text{quark}} \sim m_\pi^2$ (2.26). The scaling is dictated by critical exponents δ and $\tilde{\beta}$. The transition point β_c and T_c , respectively, is determined for each ensemble from the peak of the susceptibility $\chi(\langle \bar{\psi}\psi \rangle)$.

² For details on this approach see Chapter 5.




Figure 19: Chiral extrapolation for $T_c(m_\pi)$ from [Burger et al., 2013a].

One of the main results of this study is the extrapolation of the critical temperature T_c to the chiral limit according to (4.4), as depicted in Figure 19. The first order scenario shown in Figure 17 can be included by replacing m_π^2 with $m_\pi^2 - m_{\pi,c}^2$. As can be seen in the figure, it is not possible to distinguish between the various second order scenarios. Numerically, this is caused by the similar values of the critical exponents for the different universality classes, see Table 6. For the transition temperature in the chiral limit, $T_{\text{chiral}} = 152(26)$ MeV is obtained for the $O(4)$ scenario. This is consistently extracted from a fit to β_c using 4.3, which yields a chiral value of $\beta_{\text{chiral}} = 3.73(9)$.

transition type	second order 3D Ising	second order O(4)	first order
δ	4.789(02)	4.851(22)	1
$\tilde{\beta}$	0.3265(03)	0.3836(46)	1

Table 6: Values for critical exponents δ and $\tilde{\beta}$ for different phase transitions [Burger et al., 2013a].

A similar behaviour is encountered when analysing the scaling of $\langle \bar{\psi}\psi \rangle$ directly according to (2.87). For this, possible scaling violations were taken into account [Ejiri et al., 2009]:

$$\langle \bar{\psi}\psi \rangle = c h^{1/\delta} f(dz) + a_t t h + b_1 h + b_3 h^3. \quad (4.5)$$

Here, t is defined as $\beta - \beta_{\text{chiral}}$. The observed data does not allow to discriminate between the scenarios, although the $O(4)$ scenario is favoured. Furthermore, the latter analysis indicates that the heaviest mass (ensemble C12) is not in the scaling regime yet. Therefore, it is suggested that lower pion masses are taken into account, which is the subject of the following.

β	3.8	3.9	4.05	4.2
$a(\beta)$ (after [Baron et al., 2010])	0.100(2)	0.080(1)	0.064(1)	0.051(8)
$a(\beta)$ (after [Blossier et al., 2010])	0.098(3)	0.085(2)	0.067(2)	0.054(1)

Table 7: Lattice spacing a from [Baron et al., 2010] (Table 7, column 3) and [Blossier et al., 2010]. The value for $\beta = 3.8$ from [Baron et al., 2010] was obtained from the fit to the first combination of data points mentioned in the text.

4.3 SCALE SETTING

In the line of the existing ensembles, the temporal extent is chosen to be $N_\tau = 12$. Consequently, temperature is varied by varying β , which varies the lattice spacing a (see (2.73)). For the extrapolation to the chiral limit it is crucial that each ensemble is created on a *line of constant physics* (LOC). Hence, m_π must be tuned for each simulation point as it depends on a . Finally, a crucial advantage of twisted mass fermions is $\mathcal{O}(a)$ improvement at maximal twist, which is realised by setting κ to κ_c . This value will also depend on β . In particular it was shown above that at maximal twist the mass of the fermions is determined solely by the twisted mass $\hat{\mu}$. Thus, for the ensemble generation it is required to know $a(\beta)$, $\kappa_c(\beta)$ and $m_\pi(\beta, \hat{\mu})$. To this end, data obtained by the *European Twisted Mass* collaboration (ETMC) in two consecutive publications [Baron et al., 2010] and [Blossier et al., 2010] is used.

The ETMC investigations constitute $T = 0$ studies at four different lattice couplings: $\beta = 3.8, 3.9, 4.05$ and 4.2 . At each of these β , κ is set to κ_c by tuning m_{PCAC} (2.65) to zero. For all β , a variety of mass parameters is simulated. These data is then fitted to χ PT formulae for m_π and the pion decay constant f_π :

$$m_\pi^2 = 2B_0m \left[1 + \frac{2B_0m}{16\pi^2f_0^2} \log \left(\frac{2B_0m}{16\pi^2f_0^2} \right) + P_1m + a^2 \left(P_2 + P_3 \log \left(\frac{2B_0m}{16\pi^2f_0^2} \right) \right) \right], \quad (4.6a)$$

$$f_\pi = f_0 \left[1 - 2 \frac{2B_0m}{16\pi^2f_0^2} \log \left(\frac{2B_0m}{16\pi^2f_0^2} \right) + P_4m + a^2 \left(P_5 + P_6 \log \left(\frac{2B_0m}{16\pi^2f_0^2} \right) \right) \right]. \quad (4.6b)$$

Here, B_0 and f_0 are low energy constants of χ PT. This allows also for the setting of the lattice spacing by the pion decay constant f_π . Values for the fit parameters can be found in [Baron et al., 2010] as well as a more detailed discussion. The fits were performed for three combinations of the simulation points: 3.8, 3.9 together with 4.05, 3.9 and 4.05 together with 4.2. These fits yield somewhat different values for the various parameters. They are collected in Table 7 in [Baron et al., 2010], which will be referred to in what follows. In the more recent publication [Blossier et al., 2010], updated values for $a(\beta)$ and some other parameters are given. $a(\beta)$ is given in Table 7. One sees that the more recent values are somewhat enlarged and a is slightly increased for the three higher β , yet decreased for $\beta = 3.8$. To obtain $\kappa_c(\beta)$, the ETMC values can be interpolated nicely using a Padé approximation, see Figure 20.

In the following, two such interpolations will be presented. One is based on the data from [Baron et al., 2010], Table 7, Column 3 (first row of Table 7), which will be denoted as $\{a(\beta)\}_I$.

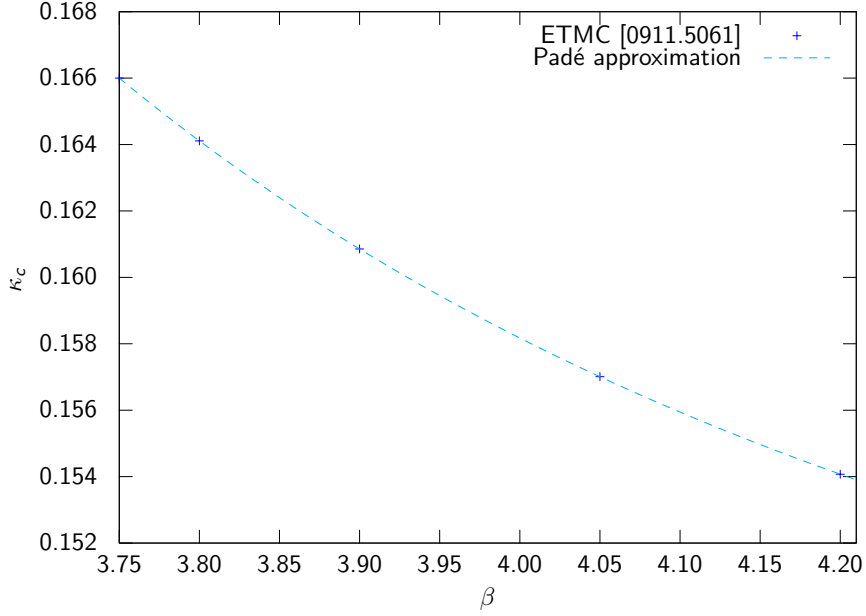


Figure 20: Interpolation for $\kappa_c(\beta)$ based on ETMC data, after [Burger et al., 2013a].

Furthermore, the values presented in [Blossier et al., 2010] (second row in Table 7) can be used, too. As not all parameters necessary for (4.6b) and (4.6a) are updated in here, one has to rely on those values given in Table 7, column 1 in [Baron et al., 2010], too. This will be denoted as $\{a(\beta)\}_{II}$.

The data of $a(\beta)$ is fitted to a quadratic polynomial and the results are shown in Figure 21. Generally, $\{a(\beta)\}_{II}$ shifts $a(\beta)$ to higher values (around 5%). Using $\{a(\beta)\}_I$, the lowest β value 3.8 can also only be included in the fit within errors. For the tmft studies, $a(\beta)$ from $\{a(\beta)\}_I$ was employed.

Using the parameters provided by ETMC, the χ PT relation (4.6a) can be used to determine $m_\pi(\hat{\mu})$ at one of the couplings simulated at by ETMC. Now, to obtain $a\hat{\mu}(\beta)$, the one loop RG group inspired ansatz

$$a\hat{\mu}(\beta) = C \exp\left(-\frac{\beta}{12\beta_0}\right) \quad (4.7)$$

was used (β_0 is defined in (2.44)).

4.4 SIMULATION SETUP

For the new ensemble, the same lattice extent was chosen as for the existing ones: $N_\sigma^3 \times N_\tau = 32^3 \times 12$. Following the previous study [Burger et al., 2013a], the ensemble was named Z12 and $\{a(\beta)\}_I$ was used for the scale setting. The new pion mass was chosen to be $m_\pi \approx 270$ MeV, which gives roughly the same Δm_π as for the prior ensembles. This was done by (4.6a), which takes on the value 270.483 MeV at $a\hat{\mu} = 0.0026099$. The constant C in the ansatz for $a\hat{\mu}(\beta)$ (4.7) was accordingly estimated to be 0.527(3). For this pion mass, T_c can be estimated to be around 190 MeV based on the existing extrapolation, cf. Figure 19. To scan the transition

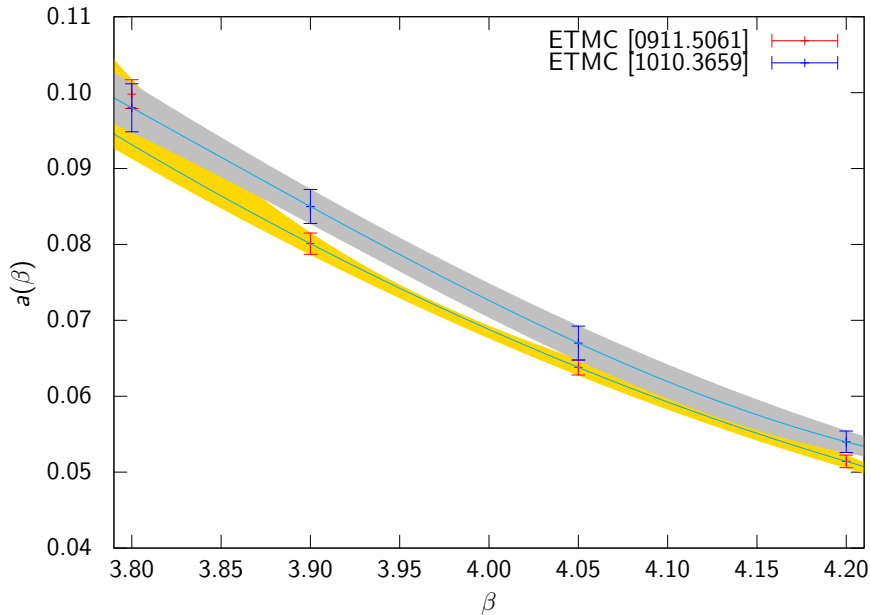


Figure 21: Interpolation for $a(\beta)$ based on ETMC data.

region, the β range given in Table 8 was chosen, which corresponds to a temperature range of roughly 175 to 200 MeV with $\Delta T \approx 2$ MeV. Within these values, the critical β can be estimated to be $\beta_c \approx 3.850(75)$, which was sampled with $\Delta\beta = 0.0075$, while this resolution was slightly enlarged for the other values. The simulation points 3.78, 3.8025, 3.8175, 3.821 and 3.83 were inserted later on.³ Details on the parameters and the simulations are given in Tables 8 and 16, respectively.

As the simulated pion mass is quite light, the generation of gauge configurations is expensive in terms of computer time (see e.g. (3.13)). It was shown in Section 3.5 that CL^2QCD is faster on a single GPU than the reference code tmLQCD on a whole CPU node. Therefore, computing resources can be used most efficiently with CL^2QCD . However, in terms of real time this approach is of course slower compared to an application like tmLQCD running on hundreds of CPU cores. To use both possibilities advantageously, the following strategy was employed. First, thermalisation runs as well as the generation of the first number of configurations for each β -point was done using tmLQCD on LOEWE-CSC and JUQUEEN⁴. Then, $\mathcal{O}(15)$ thermalised configurations were chosen as starting points for new HMC chains (i.e. each with a different PRNG seed). These chains were then run in parallel on SANAM. In this way, the statistic gain could be improved a lot. It was ensured that every single chain has enough trajectories in order to sample the phase space sufficiently. In this *pedestrian parallelism* approach, $\mathcal{O}(60k)$ trajectories were produced for most of the β . It should be emphasised that this huge amount of statistics would have taken significantly more time produced with a single chain per β than in this approach. 60k trajectories on LOEWE-CSC, for instance, would take around 1.2

³ The values of $a\hat{\mu}(\beta)$ were obtained from a slightly (0.18%) different value of m_π of 269.987 MeV, which is, however, well inside any error estimates influencing $a\hat{\mu}$ at the older points.

⁴ See Appendix A.7 for more information on these clusters.

β	T [MeV]	κ_c	$a\hat{\mu}$	traj.	acc. rate
3.7800	171(20)	0.164844	0.003082	10588	50.0
3.8025	177(16)	0.164022	0.002980	66940	59.5
3.8100	179(15)	0.163756	0.002950	58576	59.7
3.8175	181(13)	0.163494	0.002920	65924	61.6
3.8210	182(13)	0.163375	0.002899	19986	54.4
3.8250	183(12)	0.163236	0.002890	60024	62.5
3.8300	184(11)	0.163066	0.002860	19840	68.8
3.8350	186(10)	0.162898	0.002851	67634	77.1
3.8425	188(09)	0.162649	0.002822	71004	65.8
3.8500	190(08)	0.162403	0.002794	69140	47.3
3.8575	192(07)	0.162161	0.002765	74404	47.6
3.8650	195(07)	0.161923	0.002737	52940	70.8
3.8750	198(06)	0.161610	0.002700	61336	49.3

Table 8: Z12 simulation parameters. The trajectories are at $\tau = 0.5$.

years in wall time, based on the good estimate of ten minutes per trajectory using 384 cores. Opposed to that, the production runs on SANAM were performed roughly in three months, as one trajectory needs approximately 45 minutes running on a single GPU.⁵ Details about the simulations are given in A.8.

It was mentioned above that the used scale setting, especially $a(\beta)$, becomes more uncertain the lower β . As the estimated β_c lies in this region, the scale setting could be redone with $\{a(\beta)\}_{II}$. The resulting m_π according to (4.6a) is shown in Figure 22. In both estimates, the same $a\hat{\mu}$ values have been used, obtained from the fit to (4.7). One sees that for $\beta = 3.9$, m_π agrees within errors, although the estimate is a bit lower with the new parametrisation ($m_\pi \approx 255$ MeV). At $\beta = 3.8$, where m_π can now be estimated, too, the value of $m_\pi \approx 260$ MeV is again lower than the original estimate $m_\pi \approx 270$ MeV, but it is within both error estimates. The error on m_π is, however, greatly enlarged. This is due to the uncertainty of the parameters given by ETMC, which are increased as well for this dataset. Generally, the new scale does not seem to allow a better estimation on m_π for $\beta < 3.9$. For $\beta \gtrsim 3.9$, however, the error on m_π is slightly decreased. For the temperature estimates, the shift in $a(\beta)$ (see Figure 21) would cause a quite constant shift to lower temperatures. This is in agreement with the shift to lower m_π that would come with this scale setting. Still, the larger error bar on $m_\pi(\beta = 3.8)$ might also imply that the tuning to a LOC is not reliably done for lower β .

⁵ These estimates reflect the ideal case of continuous running of the simulations and do not take into account breaks like queue time, data analysis or technical issues.

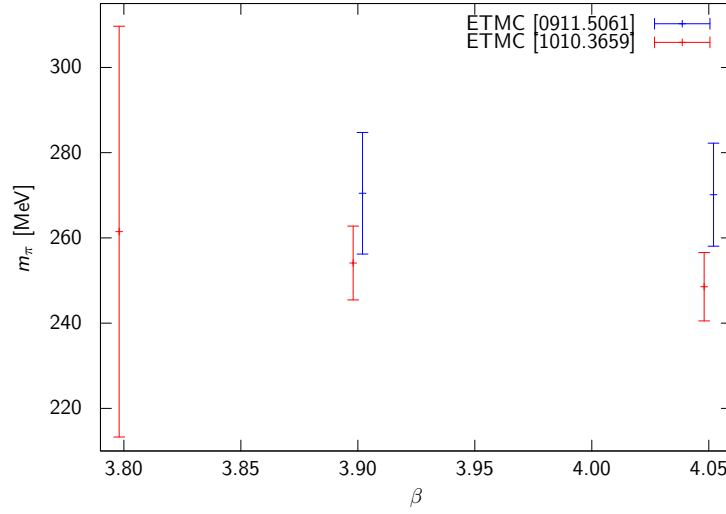


Figure 22: Comparison of m_π obtained from (4.6a) according to the to parameter sets $\{a(\beta)_I$ and $\{a(\beta)_{II}$.

4.5 THE THERMAL TRANSITION FOR THE Z12 ENSEMBLE

The chiral condensate $\langle\bar{\psi}\psi\rangle$ (2.24) has been evaluated on every 16th trajectory using 24 Gaussian volume sources as stochastic estimators (see Appendix A.5 and Table 17 therein) using CL²QCD. This means in particular that $\mathcal{O}(4k)$ values are obtained for most of the simulation points, whereas those points with less statistics have $\mathcal{O}(1k)$ $\langle\bar{\psi}\psi\rangle$ values. For more details on these results see Appendix 17. The transition point can be identified by the maximum of the susceptibility $\chi(\langle\bar{\psi}\psi\rangle)$, which is defined as

$$\chi_{\langle\bar{\psi}\psi\rangle} = \frac{\partial\langle\bar{\psi}\psi\rangle}{\partial m}. \quad (4.8)$$

It consists of a *connected* and a *disconnected* part [Bernard et al., 1997]:

$$\begin{aligned} \chi_{\langle\bar{\psi}\psi\rangle} &= -N_f \frac{T}{V} \langle\text{Tr} D^{-2}\rangle - \frac{V}{T} \left(\langle\bar{\psi}\psi\rangle^2 - \langle\bar{\psi}\psi\rangle \right)^2 \\ &\sim \chi_{\text{disconnected}} + \chi_{\text{connected}}. \end{aligned} \quad (4.9)$$

In accordance to [Burger et al., 2013a] the disconnected part is neglected. The connected part is then simply the variance per configuration $\sigma_{\langle\bar{\psi}\psi\rangle}^2$ (3.17). $\langle\bar{\psi}\psi\rangle$ and its susceptibility is shown in Figures 23 and 24, respectively. Against expectation, the susceptibility does not show a distinct peak, on the contrary, it rises for lower β . This means in particular that the signal does not show a peak at the estimated value of β_c based on the results of [Burger et al., 2013a], but a plateau (grey area in the figure). Possible implications of this will be discussed at the end of this section. First, the data will be examined in more detail.

The rising signal might have the simple explanation that β_c is not covered in the present range and one sees a broad shoulder instead. To this end, $\beta = 3.78$ was added to the original simulation points. However, even at this value quite far away from the estimated β_c




Figure 23: The chiral condensate $\langle \bar{\psi}\psi \rangle$ for the Z12 ensemble.

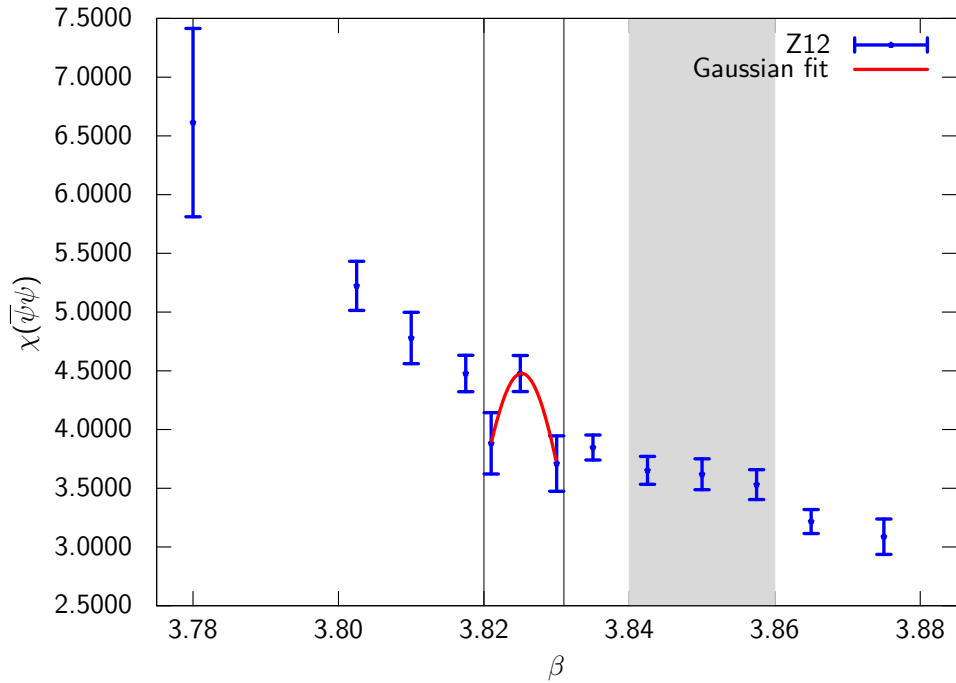


Figure 24: Susceptibility of the chiral condensate. The grey area indicates the estimate of the critical β based on [Burger et al., 2013a]. The vertical lines indicate the range for the Gaussian fit.

the susceptibility rises, although the statistics are lower here. This suggests that the rising

behaviour is rather a feature. Nevertheless, it is necessary to point out again that the scale setting might not be trustworthy the lower β . This will be commented on further below.

A peak is visible in the susceptibility around $\beta = 3.825$, which features a higher value than its direct neighbours well outside errors. However, the points at $\beta = 3.821$ and 3.83 have a significantly lower amount of data, which might cause the structure of the signal after all. This can also be seen in their mean values which seem to be a little bit off from the general functional form, cf. Figure 23. The peak is indicated in the figure by a Gaussian fit,

$$f(\beta) = A \exp\left(-(\beta - \beta_c)^2 / B\right), \quad (4.10)$$

to the three points, yielding $\beta_c = 3.82519$.

To rule out autocorrelation effects, the susceptibility of $\langle \bar{\psi}\psi \rangle$ obtained from every 64th trajectory instead of every 16th is depicted in Figure 25. Once again, no sign of a maximum of $\chi(\langle \bar{\psi}\psi \rangle)$ can be clearly seen within the estimated range. The peak at $\beta \approx 3.825$ gets washed out a little, its position can be fitted to a slightly lower $\beta_c = 3.82445$. However, it should be noted that with the reduced statistics, the points at 3.821 and 3.83 are now based on $\mathcal{O}(250)$ data points only.

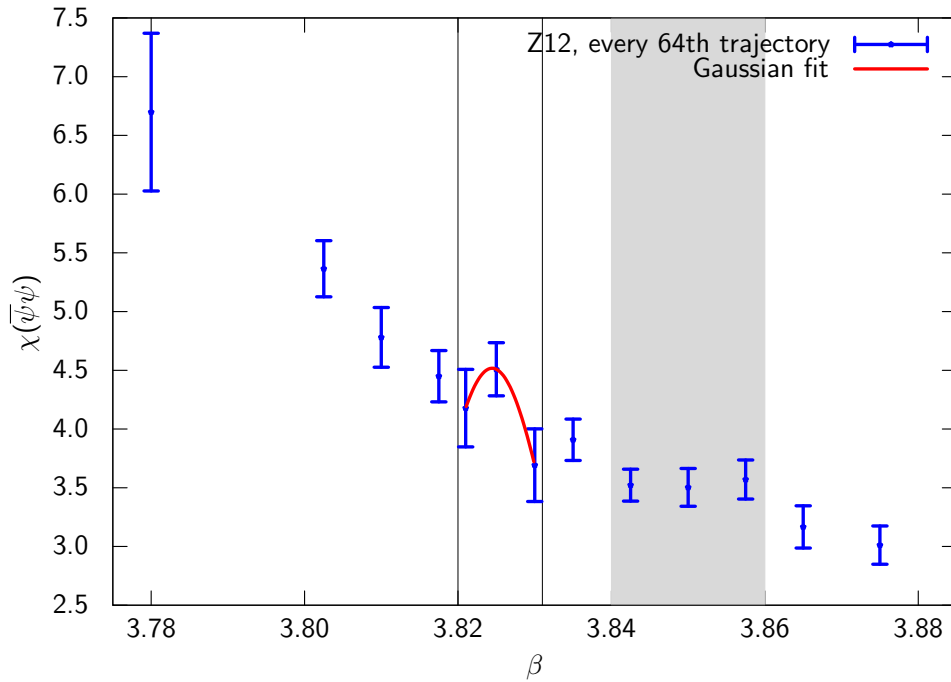


Figure 25: Same as Figure 24, but obtained from every 64th trajectory.

Keeping these possible drawbacks in mind, an estimate of β_c based on the available data is

$$\beta_c^{Peak} = 3.825(10). \quad (4.11)$$

The error is chosen such that β_c lies within simulation points of high statistics. This corresponds to a critical temperature of $183(12)$ MeV.

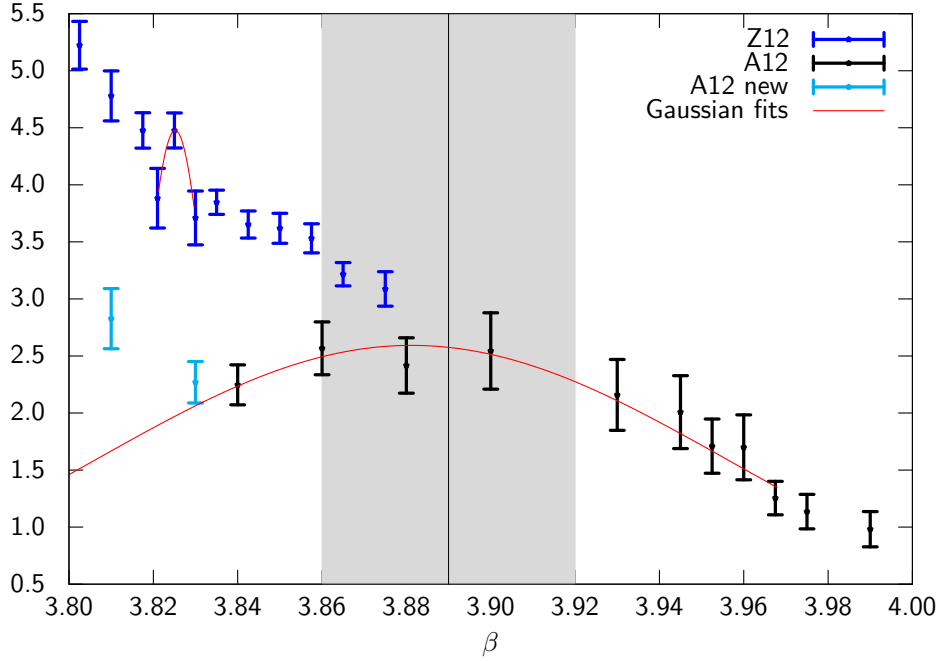


Figure 26: Susceptibility of the chiral condensate for ensembles A12 and Z12. Also shown are Gaussian fits to the data. The vertical line indicates β_c^{A12} (4.12), the grey area its error.

Insight to the behaviour of the data can be gained by comparing it to the data obtained for [Burger et al., 2013a]. There, a Gaussian was fit to the susceptibility data for each m_π . As an example, the A12 data is shown in Figure 26, whose value for β_c was determined to be

$$\beta_c^{A12} = 3.89(3) . \quad (4.12)$$

The increase in the total values of the susceptibility with decreasing pion mass is consistent with the data presented in [Burger et al., 2013a].⁶ In the direct comparison of both ensembles, one can clearly see the impact of the extended statistics of the Z12 ensemble in the decreased error bars: For most of the points, the statistic is an order of magnitude higher than for A12.⁷ In addition, the region around β_c is sampled quite coarsely for A12 ($\Delta\beta = 0.02$), which can be seen clearly in the figure. Thus, β_c is determined in a region where the signal could also be seen as a plateau. In particular, no clear peak can be obtained by eye. Nevertheless, the data also declines for higher β , which makes a fit to a Gaussian reasonable. This is supported by the two heavier ensembles, B12 and C12, which show a clear peak structure.

In order to see the behaviour of the signal for lower β , additional runs in the A12 setup were performed at $\beta = 3.81$ and 3.83 . They were carried out on JUQUEEN using the `tmLQCD` software. After thermalisation, $\mathcal{O}(7k)$ and $\mathcal{O}(5k)$ trajectories were run, respectively, which is comparable to the preceding points. The susceptibility is also shown in Figure 26. Again, the signal is increased for lower β , yet not as steep as for the Z12 ensemble. The exact position of

⁶ Note that a different normalisation of $\chi(\langle\psi\psi\rangle)$ has been used in [Burger et al., 2013a].

⁷ For the simulation points of A12 at most $\mathcal{O}(6k)$ trajectory were performed.

the additional points may be subject to statistical fluctuations, however, it is clear that they do not supply the Gaussian fit.

Thus, a rise of the susceptibility for lower β is observed consistently at two different m_π . If this is of physical origin, β_c can be extracted from the Z12 data, (4.11). Similar signals have been reported on by a recent study with Wilson clover fermions [Brandt et al., 2013] at similar m_π . Also here, the susceptibility rises strongly for lower β and shows a peak in a relative small temperature region. However, the rise may also be caused indirectly by the scale setting procedure. If m_π would not be on a LOC for all simulation points, the rising would correspond to (constantly) lowering of m_π . Then, also β_c would decrease and the total values of the susceptibility would increase. If so, the effect seems to be more or less independent of the actual value of m_π and seems to show up around $\beta \lesssim 3.84$.⁸




Figure 27: Fit to the β_c scaling form (4.3) according to $O(4)$ and $Z(2)$ critical exponents.

Despite these discussions the extracted β_c^{Peak} can be compared to the previous estimates and fitted to (4.3), as shown in Figure 27. β_c^{Peak} lies significantly lower than the others and β_c from A12 can not really be included in the fit. The fits to the scaling form are indistinguishable in the covered parameter range. Nevertheless, the obtained chiral values for the $O(4)$ and $Z(2)$ scenarios differ. They read $\beta_{chiral}^{O(4)} = 3.61(4)$ and $\beta_{chiral}^{Z(2)} = 3.65(4)$. These values are significantly lower than $\beta_{chiral} = 3.73(9)$ from [Burger et al., 2013a], which can of course be expected as β_c is lower than expected. They correspond to critical temperatures in the chiral limit of 134(59) MeV and 142(49) MeV, respectively. Here, the large errors of approximately 30% are caused by the large error in $a(\beta)$. In the same fashion, the critical temperature can be fitted to (4.4), cf. Figure 28. As can be expected, the lower β_c causes a general decrease in the extracted critical values. The critical temperatures in the chiral limit

⁸ Note that B12 and C12 have $\beta \gtrsim 3.86$




Figure 28: Chiral extrapolation for $T_c(m_\pi)$ including β_c^{Peak} from Z12 (4.11).

read $T_{\text{chiral}}^{O(4)} = 144(30)$ MeV and $T_{\text{chiral}}^{Z(2)} = 155(23)$ MeV. Here, errors are due to the fit only. These values coincide with the estimate from the scaling of β_c above within errors, as they are huge on the latter. Leaving out the C12 ensemble from the analysis does not alter the values significantly. Comparing the various values of T_c , one sees that the growing uncertainty in the scale setting prevents further restrictions on the phase transition in the chiral limit. This is despite the fact that the error on the critical β itself is quite small (see Figure 27) due to the good resolution in β .

So far the analysis has focused on the extraction of β_c from a peak in the susceptibility. As an alternative approach, the plateau around the β_c estimate from [Burger et al., 2013a] can be investigated. This would correspond to a critical β of

$$\beta_c^{Plateau} = 3.850(15) . \quad (4.13)$$

The plateau may indicate that the maximum of the susceptibility at finite temperature is washed out by the zero temperature background, as this quantity is not renormalised. This can be seen in a line with the ensembles at heavier pion mass. The B12 and C12 ensembles show distinct peaks in their bare susceptibilities (see Figures 4 and 5 in [Burger et al., 2013a]), whereas the A12 ensemble already has a rather plateau-like behaviour, especially with the additional points taken into account (Figure 26). Thus, as m_π is lowered, the finite temperature signal generally seems to be affected by a large background, which would speak against β_c^{Peak} , obtained from a rather narrow β region.

To this end, one can look at the renormalised chiral condensate, which can be estimated following [Burger et al., 2013a] as a ratio

$$\langle \bar{\psi}\psi \rangle_{ren} = \frac{\langle \bar{\psi}\psi \rangle(T, \hat{\mu}) - \langle \bar{\psi}\psi \rangle(0, \hat{\mu}) + \langle \bar{\psi}\psi \rangle(0, 0)}{\langle \bar{\psi}\psi \rangle(0, 0)} . \quad (4.14)$$




Figure 29: The renormalised chiral condensate $\langle \bar{\psi}\psi \rangle_{ren}$ for the Z12 ensemble. The grey area indicates the estimate of $\beta_c^{Plateau}$ based on [Burger et al., 2013a].

The values at $T = 0$ and in the chiral limit are obtained from extrapolations to ETMC data in the way described in [Burger et al., 2013a].⁹ It approaches one for $T \rightarrow 0$ and zero for $T \rightarrow \infty$. β_c may be extracted from the inflection point of this quantity. $\langle \bar{\psi}\psi \rangle_{ren}$ is shown in Figure 29. As expected it decreases as the temperature is increased. No clear sign of an inflection point around $\beta_c^{Plateau}$ is visible. This is mainly due to the large error bars originating from the chiral and zero temperature extrapolations. However, no such point can be seen around β_c^{Peak} , too. Nevertheless, numerically $\langle \bar{\psi}\psi \rangle_{ren} \approx 0.65$ around $\beta_c^{Plateau}$ agrees with the signal obtained for the B12 ensemble in [Burger et al., 2013a].


In order to make this analysis more coherent, additional runs at β ranging from 3.8825 to 3.9650 are considered, details are given in Table 9. It must be noted that these points have significantly less statistics compared to the points used in the previous discussion.¹⁰ Yet, $\langle \bar{\psi}\psi \rangle$ is consistent with the prior ones, see Figure 30. The signal in the bare susceptibility is not changed by these new runs, i.e. no additional peak arises, albeit the plateau in Figure 24 is slightly enhanced. With the additional runs, the renormalised condensate, Figure 31, shows a clear decrease for high temperatures. Yet again no clear inflection point can be seen. Moreover, it is again obvious in the huge error bars that for lower temperatures the scale setting is getting unreliable.

In conclusion, the renormalised chiral condensate $\langle \bar{\psi}\psi \rangle_{ren}$ does not allow to clearly locate β_c at the plateau of the susceptibility, yet favours this scenario. After all, β_c^{Peak} (4.11) is not

⁹ We thank Florian Burger for very useful discussions and valuable support on this topic.

¹⁰ A further point at $\beta = 3.9$ is left out of the analysis as it was fairly off.

β	T [MeV]	κ_c	$a\hat{\mu}$	traj.	acc. rate
3.8825	200(05)	0.161380	0.002642	7789	85.4
3.8900	202(04)	0.161153	0.002612	5979	82.5
3.9100	208(04)	0.160564	0.002575	12724	59.6
3.9175	211(04)	0.160349	0.002548	11880	57.3
3.9250	213(04)	0.160137	0.002523	10667	52.5
3.9400	219(04)	0.159722	0.002472	8386	62.3
3.9550	223(04)	0.159319	0.002422	7920	63.6
3.9650	226(04)	0.159056	0.002389	7444	64.7

Table 9: Parameters for simulations at higher β . The trajectories are at $\tau = 0.5$.Figure 30: The chiral condensate $\langle \bar{\psi}\psi \rangle$ for the Z12 ensemble including the additional β points.

based on an unambiguous signal as well. In addition, the $\langle \bar{\psi}\psi \rangle$ signal seems to be affected by a larger background as m_π is lowered. It is therefore reasonable to look at the implications of $\beta_c^{Plateau}$, see Figures 27 and 32. Here, the fits to the scaling forms yield $\beta_{chiral}^{O(4)} = 3.67(4)$ and $\beta_{chiral}^{Z(2)} = 3.70(4)$, leaving out the C12 ensemble. These results are comparable to the estimate 3.73(9) from [Burger et al., 2013a]. Hence, $\beta_c^{Plateau}$ is a better estimate for the critical β than β_c^{Peak} . It confirms the analysis of T_{chiral} from the older three ensembles from [Burger et al., 2013a] within errors.




Figure 31: The renormalised chiral condensate $\langle \bar{\psi}\psi \rangle_{ren}$ for the Z12 ensemble as a function of temperature.




Figure 32: Chiral extrapolation for $T_c(m_\pi)$ including $\beta_c^{Plateau}$ from Z12 (4.11).

As was mentioned above, the scaling behaviour of $\langle \bar{\psi}\psi \rangle$ can be investigated by means of the Magnetic Equation of State (2.87), including scaling violations (4.5). The analysis presented in [Burger et al., 2013a] was extended by including the Z12 ensemble, the fit results from the former are shown in Table 10 for completeness. Various combinations of fits were performed,

Nr.	Set	β_{chiral}	c	d	a_t	b_1	b_3	χ^2/dof
1	A12	3.57(4)	0.4(2)	0.367(7)	0	0	0	0.43
2	B12	3.40(5)	0.22(4)	0.36(2)	0	0	0	0.64
3	C12	3.12(2)	0.42(3)	0.39(2)	0	0	0	2.42
4	1	3.368(6)	0.257(6)	0.383(5)	0	0	0	3.31
5	1	3.48(2)	0.225(6)	0.48(2)	0.7(1)	0	0	2.2
6	1	3.57(2)	0.152(7)	0.53(2)	0	0.90(6)	0	1.75
7	1	3.82(4)	0.028(9)	1.1(2)	-2.2(2)	2.49(8)	0	0.42
8	1	3.73	0.1279(8)	0.825(8)	4.01(4)	0	0	76
9	1	3.73	0.0759(7)	0.81(2)	0	1.61(2)	0	7.2
10	1	3.73	0.053(2)	0.74(2)	-1.8(2)	2.23(6)	0	0.63
11	2	3.76(2)	0.047(6)	0.83(6)	-1.5(2)	2.20(6)	50(11)	1.8

Table 10: Fits to (4.5) from [Burger et al., 2013a]. Bold entries were fixed during fits. Datasets 1 and 2 correspond to ensembles A12 + B12 and A12 + B12 + C12, respectively.

Nr.	Set	β_{chiral}	c	d	a_t	b_1	b_3	χ^2/dof
1	1	3.67(2)	0.10(1)	0.39(1)	0	0	0	2.07
2	2	3.44(1)	0.22(1)	0.40(1)	0	0	0	2.55
3	2	3.32(3)	0.28(2)	0.34(1)	-0.81(16)	0	0	1.96
4	2	3.38(5)	0.26(4)	0.37(2)	0	-0.40(30)	0	2.03
5	2	3.44(9)	0.11(3)	0.32(8)	-3.85(67)	2.87(45)	0	1.67
6	2	3.30(3)	0.34(4)	0.34(1)	-0.10(21)	0	0.70(10)	2.23
7	2	3.34(5)	0.20(3)	0.31(2)	-2.38(61)	1.50(52)	-1.45(12)	1.94
8	2	3.73	0.029(3)	0.60(3)	-5.00(46)	3.20(12)	0	1.27
9	2	3.67	0.040(5)	0.47(3)	-5.00(59)	3.31(21)	0	1.39
10	3	3.73	0.034(1)	0.69(1)	-2.23(9)	2.46(3)	0	2.02
11	3	3.73	0.045(1)	0.69(2)	-2.31(11)	2.47(6)	-9(11)	2.07
12	4	3.76(3)	0.04(2)	0.79(7)	-2.09(35)	2.43(22)	41(32)	3.81
13	4	3.73	0.051(9)	0.70(12)	-1.73(7)	2.24(3)	170(43)	3.25

Table 11: Overview over fits to (4.5). Bold entries were fixed during fits. Datasets 1, 2, 3 and 4 correspond to ensembles Z12, Z12 + A12, Z12 + A12 + B12 and Z12 + A12 + B12 + C12, respectively. Combinations of fit parameters not shown have $\chi^2 \gg 1$.

Nr.	Set	c	d	a_t	b_1	b_3	z_1	z_2	χ^2/dof
14	2	2.7(3)	5.6(5)	-7.0(2.0)	3.44(25)	0	4.5(4.5)	0	1.31
15	2	2.8(3)	5.9(3)	-5.72(91)	3.32(18)	0	0	3.7(4.2)	1.30
16	2	2.7(3)	5.7(5)	-6.5(1.9)	3.40(21)	0	2.6(6.3)	2.1(6.7)	1.39
17	2	2.7(3)	5.7(6)	-6.4(2.0)	3.39(23)	0.58(2)	2.3(6.6)	2.2(6.8)	1.50
18	3	3.6(2)	6.1(2)	-5.86(64)	3.00(10)	0	6.4(1.2)	0	1.44
19	3	4.3(1)	6.8(1)	-3.00(26)	2.61(6)	0	0	3.3(1.1)	1.84
20	3	3.4(2)	5.6(3)	-7.42(86)	3.17(1)	0	11(2)	-4.8(2.2)	1.36
21	3	3.3(2)	5.6(3)	-7.4(1.0)	3.2(1.5)	-2(14)	11(3)	-4.9(2.6)	1.40
22	4	3.7(1)	6.2(1)	-6.11(45)	3.00(7)	0	7.3(8)	0	2.15
23	4	5.2(1)	7.3(1)	-1.40(15)	2.22(4)	0	0	-2.0(1)	3.36
24	4	3.5(2)	5.8(2)	-6.66(53)	3.05(9)	0	10(1)	-4.6(8)	1.42
25	4	3.4(2)	5.8(2)	-6.90(53)	3.12(9)	-146(73)	11(1)	-6.4(1.2)	1.35

Table 12: Overview over fits to (4.15). β_{chiral} is fixed to 3.73 throughout. Datasets 1, 2, 3 and 4 correspond to ensembles Z12, Z12 + A12, Z12 + A12 + B12 and Z12 + A12 + B12 + C12, respectively. Bold entries were fixed during fits. Combinations of fit parameters not shown have $\chi^2 \gg 1$. Given values for c and z_2 have been weighted by a factor 100 and d by a factor of 10.

see Table 11.¹¹ The fit range for the Z12 ensemble was [3.835, 3.875], for the other ensembles according to [Burger et al., 2013a]. β -values 3.821 and 3.83 were left out of the analysis. Fitting the Z12 ensemble alone (Fit Nr. 1) confirms the general trend as β_{chiral} goes up, c goes down and d stays constant. Note that β_{chiral} agrees with the extrapolations from β_c^{Plateau} . Yet, the quality of the fit, indicated by χ^2 , suggests to include scaling corrections. A similar behaviour is seen including the A12 ensemble, fit Nr. 2. Including the corrections (Fits 3-7) yields reasonable fits, but with a somewhat lower $\beta_{\text{chiral}} \approx 3.4$. This is consistent with the reports in [Burger et al., 2013a]. Hence, β_{chiral} was set to the fixed value 3.73 from the independent analysis of β_c . This yields reasonable fits (Nr. 8, 10 and 11), confirming the β_c analysis. Varying the value of β_{chiral} does not change the results much (Fit Nr. 9). As the B12 and C12 ensembles are included in the fits, the fit quality worsens, in particular the C12 ensemble can not be included to a satisfying extent.

To investigate this further, higher order scaling violations can be considered:

$$\langle \bar{\psi}\psi \rangle = c h^{1/\delta} f(dz) + a_t t h + b_1 h + b_3 h^3 + z_1 t^2 h + z_2 t^3. \quad (4.15)$$

Results to the fits are given in Table 12. For the Z12 and A12 ensembles, the higher order coefficients are zero within errors (Fits 14-17). This changes when B12 and C12 are included (Fits 18-25). For both combinations, χ^2 is smaller than in the fits to 4.5, although the improvement is most pronounced in fits 22-25. This suggests that C12 is indeed not in the

¹¹ The fits were carried out with an R-program provided by Florian Burger.

scaling region, as was concluded in [Burger et al., 2013a]. The same might be true for the B12 ensemble: Looking at the scaling of the data, shown in Figures 33 and 34, one sees that the B12 and C12 data is well described by the fit only if the B12 ensemble is included. This indicates different characteristics of the two subgroups of ensembles.

In conclusion, the Z12 ensemble was used to fully enlarge the analysis presented in [Burger et al., 2013a] to smaller m_π . It is concluded that the C12 ensemble is not within the scaling region, which may also be true for the B12 ensemble. Hence, even lower pion masses are needed to clarify the nature of the $N_f = 2$ chiral transition.




Figure 33: Scaling of $\langle \bar{\psi}\psi \rangle$ according to fit Nr. 8 in Table 11.




Figure 34: Scaling of $\langle \bar{\psi}\psi \rangle$ according to fit Nr. 10 in Table 11.

4.6 STRENGTH OF THE AXIAL ANOMALY

The situation in the chiral limit can also be addressed differently than in the previous section. The axial anomaly may be effectively restored at the same temperature scale where the chiral transition happens. As Figure 18 shows, this can be seen in a mass degeneracy of parity partners. This means that one should be able to see the restoration of either the axial anomaly or chiral symmetry by a degeneration of the masses of the appropriate particles. This will be done in this section. In particular, the analysis will focus on the scalar (SC) and pseudo scalar (PS) channels, results on vector (VE) and axial vector (AV) channels will be presented, too. A similar analysis for the A12, B12 and C12 ensembles was carried out in [Zeidlewicz, 2011].

As pointed out in Section 2.5.5, it is far more convenient to calculate flavour-multiplet than flavour-singlet propagators because in the latter there are disconnected diagrams. Therefore, the masses of π^\pm and a^\pm will be analysed, which have $\Gamma(\pi) = \gamma_5$, $\Gamma(a) = \mathbb{1}$. Using (2.78), the correlators for these particles read

$$C_\pi = -\text{Tr} \left[|D_u^{-1}(n, m)|^2 \right] , \quad (4.16)$$

$$C_a = -\text{Tr} \left[D_u^{-1}(n, m) \gamma_5 (D_u^{-1})^\dagger(n, m) \gamma_5 \right] . \quad (4.17)$$

As was explained in Section 2.5.5, the effective masses m_{eff} are extracted from fits to the correlator data for large separation.¹² As will be elucidated below, this is not always possible due to noise in the propagator signals. Therefore, in addition, integrated correlators are evaluated and compared to the previous analysis:

$$\chi_\Gamma = \sum_{z \neq 0} C_\Gamma(z) . \quad (4.18)$$

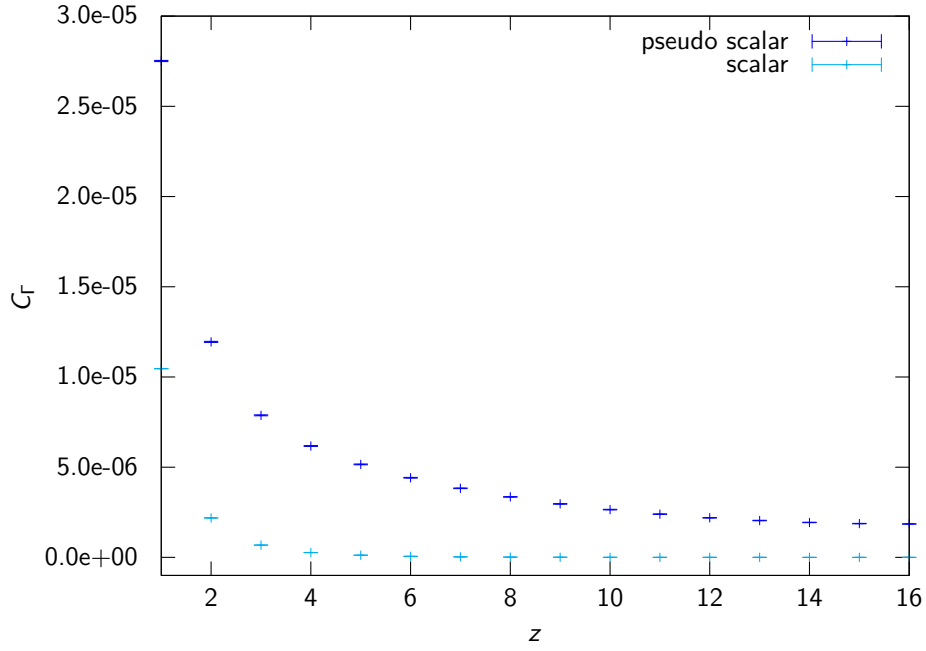

Note that zero separation terms, dominated by contact terms, are neglected for consistency with [Zeidlewicz, 2011].

The correlators have been measured on every 16th trajectory for those Z12 simulation points with high statistics. That is, $\mathcal{O}(4k)$ correlators were obtained per β . Up to 16 point sources have been used in each case. Details on the results can be found in Table 18 in the Appendix.

As an example, the scalar and pseudo scalar correlators are shown in Figure 35. The pseudo scalar correlator generally takes on higher values than its scalar counterpart. In particular, the scalar correlator has a quite small signal. This behaviour is persistent for all β . Throughout this analysis, the periodicity of the lattice will be used, accordingly N_σ is effectively halved.

From this data, m_{eff} can be extracted, which is depicted in Figure 36. In the pseudo scalar channel a plateau can be easily identified, indicated by the lines in the figure. In particular, errors from the fit are negligible compared to the scalar channel. In this channel, in turn, a plateau can be seen starting at $z = 6$, which then starts to fluctuate wildly for larger separations. The effective masses for $\beta = 3.85$ and 3.8575 are shown in Figures 37 and 38, respectively. While for $\beta = 3.8575$ the data looks similar to $\beta = 3.825$, at $\beta = 3.85$ a plateau at large separations z is clearly visible, which would correspond to a much lower m_{eff} than coming from the fit of $\beta = 3.8575$. However, generally m_{eff} should decrease with temperature. Therefore,

¹² For the actual fit of m_{eff} a program provided by Marc Wagner was used.

Figure 35: Scalar and pseudo scalar correlator at $\beta = 3.825$.Figure 36: m_{eff} for scalar and pseudo scalar channels at $\beta = 3.825$. Lines indicate fit regions.

the plateau seen at $\beta = 3.85$ may be misleading. After all, as was mentioned in Section 2.5.2, the scalar particle is not protected against decay into the pseudo scalar channel in twisted mass LQCD. Thus this plateau, which is close to the pseudo scalar data, might not indicate the effective mass of the scalar particle. In fact, a second plateau at intermediate separations may

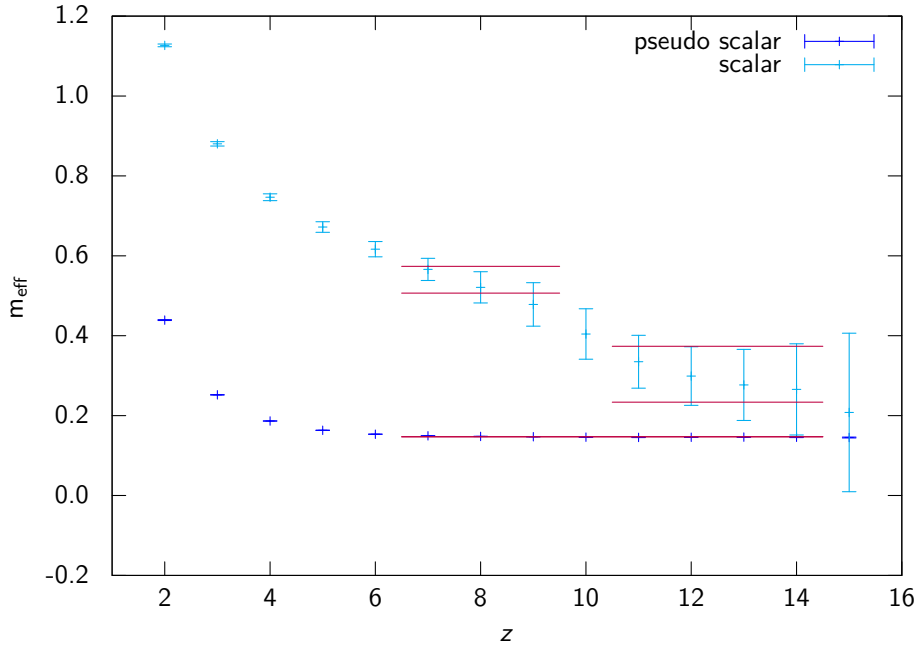


Figure 37: m_{eff} for scalar and pseudo scalar channels at $\beta = 3.85$. Lines indicate fit regions.




Figure 38: m_{eff} for scalar and pseudo scalar channels at $\beta = 3.8575$. Lines indicate fit regions.

be seen in the data, as indicated in the figure. Also for the highest temperature at $\beta = 3.875$, this structure can be observed, cf. Figure 39. However, the lower plateaus are more concise, therefore these will be used further on.




Figure 39: m_{eff} for scalar and pseudo scalar channels at $\beta = 3.875$. Lines indicate fit regions.

By contrast, it may also be the case that the scalar mass cannot be clearly dissolved if there is only one plateau (as in Figure 36). This could be caused by too few statistics, as the scalar channel is usually affected with more noise than the pseudo scalar one. But this seems unlikely regarding the given amount of data. And, of course, a clear plateau can be seen in the given data.

The differences of the extracted screening masses are shown in Figure 40. On the one hand, there is clearly no degeneracy of the masses around β_c^{Peak} (grey area in the plot), hence no restoration of $U_A(1)$. In fact, the mass differences are quite constant here. On the other hand, around β_c^{Plateau} some of the scalar masses are very close to the pseudo scalar one (blue area in the plot). In particular, this is the case for $\beta = 3.85$ and 3.865 . As was mentioned above, however, this must not necessarily indicate an actual lowering of the scalar mass. In particular, these points are put into question by the general trend of the data, which is rather constant. Especially $\beta = 3.8575$ does not show any signs of degeneracy (see Figure 38).

To investigate this further, one can look at the integrated correlator (4.18). In Figure 41, the ratio of this observable for the scalar and pseudo scalar channels is shown.¹³ This ratio should approach one as the symmetry gets restored. Its value is constantly decreasing, but in the investigated β range it is far off from one. This clear signal particularly indicates that the effective masses also decrease in a similar way. Thus, the observed huge differences in the effective mass differences of neighbouring simulation points, Figure 40, are most likely due to decays of the scalar particle into pseudo scalar ones.

For the vector and axial vector channels, $\Gamma(\text{VE}) = \gamma_0 \gamma_i$, $\Gamma(\text{AV}) = \gamma_5 \gamma_i$, $i = 1, 2$, the integrated correlator is shown in Figure 42. The ratio is constant over the whole temperature

¹³ Similar signals are obtained for the difference of these observables.

range, no signs of restoration are seen in this channel, too. Moreover, the effective masses have a constant difference for most of the simulation points, as seen in Figure 43. At $\beta = 3.835$ and 3.85, the mass of the heavier axial vector particle drops. In the respective effective mass plots, one can once more see the structure mentioned above, as again a decay between the two channels is not forbidden by parity for twisted-mass fermions.




Figure 40: Strength of the axial anomaly extracted from screening mass analysis of Z12 data. The grey area indicates β_c^{Peak} , the blue one $\beta_c^{Plateau}$.




Figure 41: Ratio of integrated correlators for scalar and pseudo scalar channels.




Figure 42: Ratio of integrated correlators for vector and axial vector channels.




Figure 43: Effective masses extracted from the correlators of the vector and axial vector channels.

In conclusion, the analysis of screening masses for the Z12 ensemble does not yield a restoration of $U_A(1)$ around the chiral transition. For the chiral limit, this behaviour would mean that the chiral transition is $O(4)$. However, as the simulations are performed at finite m_π , an extrapolation to this limit is necessary. In the analysis of the preceding ensembles, this was carried out in [Zeidlewicz, 2011]. There, a rising of the anomaly strength with decreasing m_π was observed. For instance, for the A12 ensemble the ratio of the integrated correlators for the scalar and pseudo scalar channels takes on values < 5 , which is lower than all values

obtained for Z12 (cf. Figure 41). Yet, this analysis was carried out for $\beta \gtrsim 3.9$, so not all data presented in [Burger et al., 2013a] was covered. Nevertheless, as the anomaly should decrease with temperature, it can not be weaker at lower β . Thus, it is fair to say that the anomaly does not weaken as m_π is decreased.

This is in accordance with recent, independent studies. In [Brandt et al., 2013] the issue was addressed with clover improved Wilson fermions at $m_\pi = 200$ and 290 MeV. The obtained differences in $M_{\text{scr}}^{\text{sc}} - M_{\text{scr}}^{\text{ps}}$ were extrapolated to be non-vanishing also in the chiral limit. In an approach utilising Domain Wall fermions [Buchoff et al., 2013], which have much better chiral properties than Wilson fermions, a persistent anomaly is seen for 200 MeV and even physical, 135 MeV, pion masses.

An important remark on these considerations has to be made: It is not clear how “strong” the anomaly has to be in the chiral limit in order to drive the transition to be $O(4)$, see e.g. [Brandt et al., 2013].

4.7 SUMMARY AND DISCUSSION

In this section the nature of the $N_f = 2$ chiral transition was addressed, which is not settled yet. This was done within the framework of maximally twisted mass Wilson fermions, introduced in Sections 2.3 and 2.5.2. In particular, the presented study constitutes an extension of an investigation by the tmft collaboration [Burger et al., 2013a]. This was built on the analysis of the chiral transition at different pion masses m_π . More specifically, three ensembles at 315, 380 and 470 MeV pion mass were studied (labelled A12, B12, C12, respectively). For each ensemble, the chiral transition was identified as a crossover by the peak of its susceptibility $\chi(\langle\bar{\psi}\psi\rangle)$. The observed critical temperatures were analysed in terms of universal scaling behaviour for the different universality classes in question. However, these scenarios were found to be indistinguishable within the simulated m_π . Accordingly, it was suggested to add an additional ensemble at lower m_π . The new ensemble, labelled Z12, was generated at $m_\pi \approx 270$ MeV and the analysis presented in [Burger et al., 2013a] was extended in detail including this new ensemble.

The ensemble generation was done in a pedestrian parallelism approach, i.e. several chains per β have been simulated. In this way, the OpenCL-based LQCD application CL²QCD presented in Chapter 3 could be employed, yielding a much higher efficiency running on several GPUs than a single chain on a comparable machine with a multi CPU application. For instance, the statistics gained for most simulation points are very high, $\mathcal{O}(60k)$. A similar effort pursuing the standard approach would have taken at least three times more time. Thus, the development of CL²QCD was mandatory to allow this amount of statistics in a relative short time. In particular, it should be emphasised that the resolution in β for Z12 is much higher than in the previous studies. In fact, the number of accumulated trajectories for the Z12 ensemble (approximately 800k, see Tables 8 and 9) is nearly twice as large as those from ensembles A12, B12 and C12 together (approximately 400k, see [Burger et al., 2013a]).

All these considerations strongly rely on the conversion of lattice results to physical scales. To this end, data obtained by the ETMC was used [Baron et al., 2010]. They provide values

for κ_c , $a(\beta)$ for four β values, 3.8, 3.9, 4.05 and 4.2 as well as χ PT formulae to estimate m_π at a given $a\hat{\mu}$. These values can be interpolated by reasonable ansätze to obtain the necessary parameters for intermediate β .

However, as was elaborated in Section 4.3, this procedure may be untrustworthy for $\beta \lesssim 3.84$. This is caused by the increased uncertainty of $a(\beta)$ in this region. Additionally, the ansatz for $a\hat{\mu}(\beta)$ may not be valid anymore far away from the actual determination point of m_π . A more recent publication of ETMC [Blossier et al., 2010] was investigated for possible improvements for this β region. While $a(\beta)$ improves significantly, the m_π estimate is unreliable at $\beta = 3.8$ due to huge errors caused by the available parameters by ETMC for the χ PT formulae.

Based on the extrapolation to $m_\pi = 270$ MeV, T_c could be estimated to be around 190 MeV. This translates to an estimate for β_c of around 3.85. Surprisingly, $\langle\bar{\psi}\psi\rangle$ of Z12 does not show signs of a transition around this point. Instead, only a plateau can be seen in the susceptibility. Against expectation, this plateau does not decline for lower β , but rises strongly. It has been ruled out that β_c is located at lower β by adding an additional point at $\beta = 3.78$. Also here, the susceptibility is still higher than at the former points. Insight into this issue can be gained by comparing the data for Z12 and A12. First of all, this makes explicit that β_c^{A12} is relying on a much smaller resolution in β . Secondly, the A12 range was enlarged to lower β by two additional points. In agreement with Z12, the susceptibility rises, too. Thus, this behaviour is a systematic one.

Possible effects of the unphysical Sharpe-Singleton phase on the signal can be expected to be negligible. No signs of this transition have been found in a former tmft study [Ilgenfritz et al., 2009] for $\beta > 3.75$ although at higher $\hat{\mu}$. Nevertheless, extending N_τ to 16 would reduce possibly artifacts further.

A peak can be seen in a narrow region around $\beta_c^{Peak} = 3.825(10)$. Nevertheless, it should be noted that its neighbouring points have less statistics. In the scaling analysis, this β_c lowers the predictions for the chiral limit. Yet, $O(4)$ and $Z(2)$ scenarios can not be distinguished. Furthermore, the aforementioned huge uncertainties of $a(\beta)$ yield even higher uncertainties on the T_{chiral} estimate. Alternatively, a possible β_c located at the plateau in the susceptibility data, $\beta_c^{Plateau} = 3.850(15)$, can be investigated by means of the renormalised chiral condensate $\langle\bar{\psi}\psi\rangle_{ren}$ (4.14). This analysis does not yield clear signals for this scenario, yet additional simulation points enhance the plateau. In this scenario, the T_{chiral} extrapolation is in agreement with the analysis of the previous ensembles, which comes as no surprise. Despite the high resolution in β in the Z12 ensemble, improvement on the chiral extrapolation for T_{chiral} are difficult since the error on T is enlarged for smaller m_π due to the scale setting. On the contrary, the presented data emphasises that the analysis is getting exceedingly difficult as m_π is lowered, both in required computation time as well as in the quality of the (bare) signal. In addition, it was shown that β_c^{A12} is based on a rather coarse set of β points and that also for this point, no clear peak can be seen in the bare susceptibility. This supports that $\beta_c^{Plateau}$ is the better β_c -estimate. The scaling of $\langle\bar{\psi}\psi\rangle$ has been investigated using all available ensembles. This analysis clearly showed that the C12 ensemble is not within the scaling region yet, which may also be true for the B12 ensemble. The Z12 and A12 ensemble can be nicely fitted to an

ansatz including scaling violations. This analysis confirms the estimate of the critical β in the chiral limit from [Burger et al., 2013a].

A different approach was followed by studying the $U_A(1)$ anomaly in terms of mesonic screening masses. If the mass differences of corresponding partner channels vanished around T_c , this would signal a restoration of this symmetry, favouring a first order transition in the chiral limit. No such clear signal of a restoration can be seen in the data. However, the analysis is hampered as parity is no longer a symmetry for twisted mass fermions. This seems to spoil the extraction of the screening masses. But also the integrated correlators did not show signs of restoration. This is in agreement with current investigations by other groups. A rough estimate of the strength of the anomaly with varying pion mass could be done based on results shown in [Zeidlewicz, 2011]. There is no sign that the anomaly weakens for lower pion masses.

Hence, keeping the discussion about the possible caveats in mind, both analysis presented in this work suggest that the transition in the chiral limit is indeed second order $O(4)$ for $N_f = 2$. Yet, this has to be confirmed by studies at even lower pion masses.

 THE ROBERGE-WEISS TRANSITION WITH WILSON FERMIONS

In this chapter the Roberge-Weiss (RW) transition is addressed, particularly the nature of the RW endpoint is investigated. This is done using Wilson fermions and mainly by using the OpenCL-based application CL²QCD described in Chapter 3 that allows for an acceleration of the simulations by the use of Graphics Processing Units (GPUs). The presentation is based on [Philipsen and Pinke, 2014].

This study constitutes one of the first studies with Wilson fermions at imaginary μ using the setup elucidated below, together with [Wu and Meng, 2013; Wu et al., 2007], of which the latter has been published during the development of this thesis. Studying the region of imaginary chemical potential μ_I of QCD is interesting as it can be used to constrain the QCD phase diagram at finite μ , which is currently not directly accessible by LQCD simulations due to the sign-problem. At μ_I , the sign-problem is absent and standard simulation techniques can be applied. In addition, QCD shows a rich phase structure in this region, in particular there is the RW transition between adjacent $\mathbb{Z}(N_c)$ sectors as μ_I is increased and passes through critical values μ_I^c . Relating the findings of this effort to existing results based on simulations with staggered fermions is of particular interest, as the latter discretisation is under discussion because of the necessary rooting procedure [Creutz, 2007; Sharpe, 2006]. The expectation for the RW endpoint is to find two tricritical points as the fermion mass is varied, $\kappa_{\text{heavy}}^{\text{tric}}$ and $\kappa_{\text{light}}^{\text{tric}}$, compare Figure 44b. Additionally, a comparison to a heavy quark effective theory for LQCD with Wilson fermions is made, which can predict the value of $\kappa_{\text{heavy}}^{\text{tric}}$. This is of notable importance as the nature of the $N_f = 2$ chiral phase transition, which is yet unresolved (see Chapter 4), can be addressed following the tricritical line emerging from this point.

In the next section, the phase structure of QCD at imaginary μ is discussed in detail as well as the status of LQCD investigations in this region, especially how the phase structure changes as the fermion mass is varied. In Section 5.2 the simulation setup and considered observables are introduced, respectively. The last two sections are devoted to the presentation of the results as well as a summary and discussion.

5.1 THE QCD PHASE DIAGRAM AT IMAGINARY μ

The sign problem prevents direct LQCD simulations at nonzero μ . Nevertheless, there is a lot of effort carried out in order to make statements about this region using $\mu = 0$ informations (see [Fukushima and Hatsuda, 2011; Levkova, 2011] for recent overviews). Since fully accessible

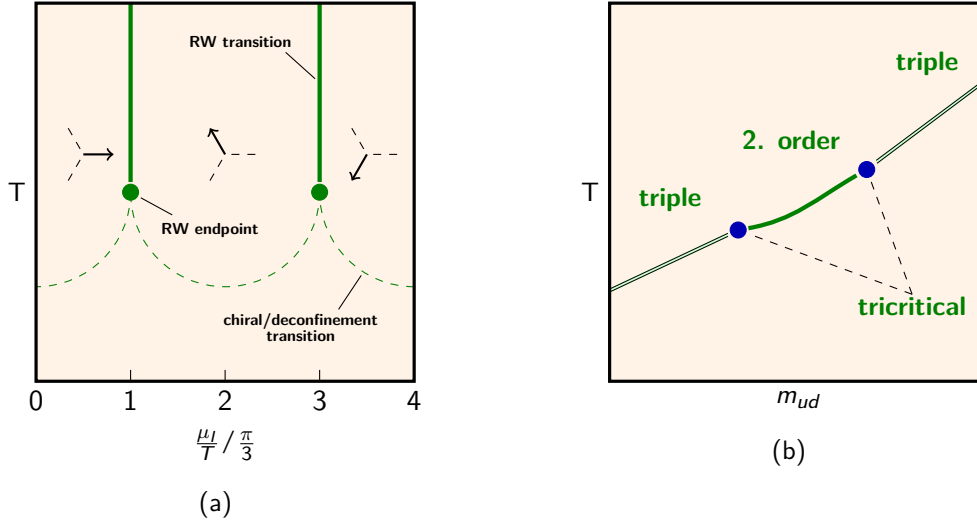


Figure 44: a) Schematic phase diagram for QCD at imaginary chemical potential (following [Bonati et al., 2011b]). The dashed line shows the chiral/deconfinement phase transition, the solid ones the first order RW-transitions at μ_I^c . Below T_c , the RW transitions are crossover. b) RW endpoint as function of mass, following [Bonati et al., 2011c].

on the lattice, purely imaginary μ , denoted as μ_I , is investigated in many studies. These have been predominantly based on staggered fermions [Bonati et al., 2011b,c; Cea et al., 2012; de Forcrand and Philipsen, 2002, 2008, 2010; D'Elia and Lombardo, 2003; Falcone et al., 2010; Takaishi et al., 2009]. Recently, studies using Wilson fermions have been reported [Alexandru and Li, 2013; Nagata and Nakamura, 2011, 2012; Nakamura et al., 2011; Wu and Meng, 2013; Wu et al., 2007].

As was mentioned in Section 2.2, there is a periodicity of $2\pi/N_c$ in this region for physical observables. The general phase structure was worked out in [Roberge and Weiss, 1986] and is shown in Figure 44a. At critical values of μ_I ,

$$\mu_I^c = i (2k + 1) \pi T / N_c \quad (k = 0, 1 \dots N_c - 1), \quad (5.1)$$

there is a phase transition between adjacent $\mathbb{Z}(N_c)$ sectors. This is indicated by the aforementioned phase of the (modified) Polyakov loop shown in the figure. At low temperatures the transition is a crossover, at high temperatures first order. Consequently, these two regions meet at some point, the *Roberge-Weiss (RW) endpoint*. This point is connected to the deconfinement and chiral transitions at $\mu = 0$ (dashed lines in the figure). In $N_f = 2$ and 3, staggered simulations found that these transitions extend into the μ_I region and meet the first order RW line in the RW endpoint. Thus, the nature of the endpoint is non-trivial and depends on N_f and the fermion mass.

For the high mass region, this is shown in Figure 45. For very high masses, the deconfinement transition at $\mu = 0$ is first order (the upper right corner in the Columbia plot, cf. Figure 16a), which extends into the μ_I region, i.e. the dashed line in Figure 44a is a first order line. Thus, at the RW endpoint there is a triple point, as the RW transition and two first order deconfinement transitions meet there (Figure 45a). As the mass is lowered, the $\mu = 0$ transition passes through

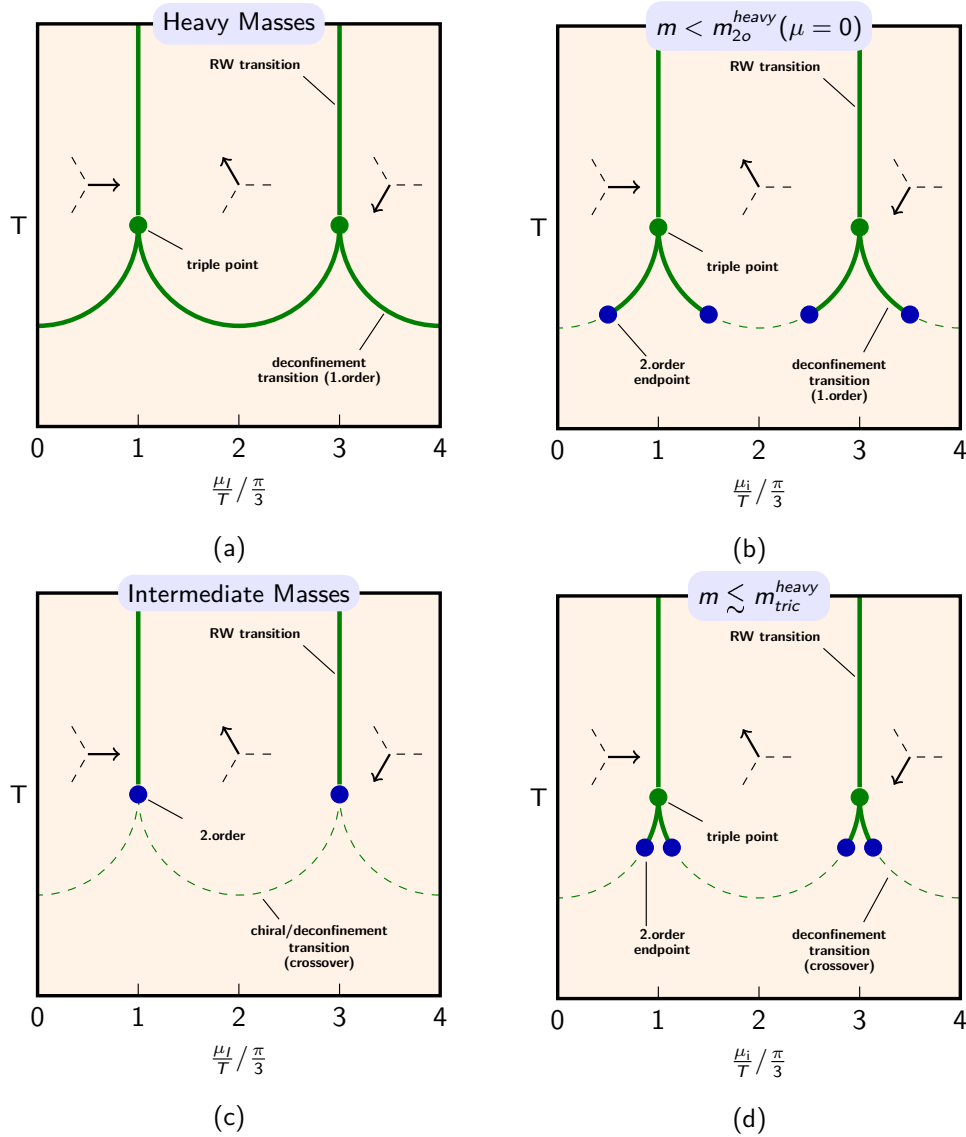


Figure 45: The phase diagram of QCD at imaginary μ for different values of the quark mass. The quark mass is varied clockwise, starting in 45a.

the second order line and becomes a crossover. This carries over to the μ_I region, Figure 45b. The second order point approaches the RW endpoint from $\mu = 0$, ending the first order line. The latter is shortened (Figure 45d), until it eventually meets the RW point (Figure 45c). The same mechanism happens when coming from the chiral limit, increasing the mass, at least for $N_f = 3$. Looking only at the nature of the transition in the RW endpoint, one has a triple point for high masses, which for some critical mass becomes a second order endpoint. At this critical mass one has a tricritical point. At sufficiently low masses, the endpoint becomes a triple point again, where the first order RW line now meets with two first order chiral transitions. This is depicted in Figure 44b. A picture similar to the Columbia Plot emerges naturally, Figure

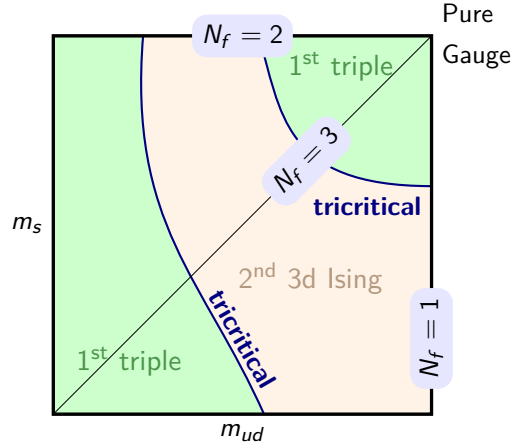


Figure 46: Phase diagram at $\mu = i\pi T$ (following [Bonati et al., 2011b]).

46. For low and high masses there are regions of triple points, which are bounded from a second order region by tricritical lines.

The two planes at $\mu = 0$ and $\mu = i\pi/N_c/T$ must be connected analytically, in particular, the $Z(2)$ lines in the Columbia Plot become surfaces, meeting the tricritical lines at the RW value (Figure 47). The curvature of the surfaces allows for conclusions about the physical point at real μ . It was found [Bonati et al., 2011b; de Forcrand and Philipsen, 2008] that both the chiral and deconfinement critical surfaces will bend towards the corners of the Columbia plot for real μ , which speaks against a CEP at physical quark masses.

However, the yet unresolved nature of the transition in the $N_f = 2$ chiral limit may change this picture. The scenario would be settled by determining the meeting point of the $N_f = 2$ tricritical line with the μ axis (light-blue line and point in the northern part of Figure 47): If it meets below the $\mu = 0$ plane, the transition must be second order at $\mu = 0$, otherwise first order [Bonati et al., 2011b]. Indeed, ongoing studies [Bonati et al., 2013] support the first order scenario, at least on coarse lattices.

After this discussion of the phase structure of QCD at imaginary values of μ , some studies relevant to the one carried out within this thesis are explained in more detail. $N_f = 2$ simulations with staggered fermions aiming also at the nature of the RW endpoint have been reported in [Bonati et al., 2011c]. Systems with temporal extent $N_\tau = 4$ at $\mu_l = i\pi T$ have been investigated for various spatial volumes. This study mainly relies on a scaling analysis of $\chi(|L_{lm}|)$, which can be expected to scale according to

$$\chi = N_\sigma^{\gamma/\nu} f(tN_\sigma^{1/\nu}) \quad (5.2)$$

in the vicinity of the transition point. Here, f is a universal scaling function and t is the reduced temperature. γ and ν are critical exponents specific to the universality class of the transition (see Table 13).¹ f does not need to be known, the critical exponents can be estimated by looking at $\chi/N_\sigma^{\gamma/\nu}$ against $tN_\sigma^{1/\nu}$ for multiple spatial volumes. These curves should coincide

¹ Strictly speaking, critical exponents can be defined for second order transitions only. However, similar considerations can be carried out for the first order case, too, cf. [Fisher and Berker, 1982].




Figure 47: Columbia plot extended to real imaginary μ regions following [Bonati et al., 2011b]. Blue lines indicate tricritical lines. Red lines indicate the $Z(2)$ critical lines in the $\mu = 0$ plane and the red surfaces display their continuation. Tricritical points are indicated by blue dots. At real μ , the critical surfaces bend away from a possible CEP at physical masses. It is currently under investigation what position the endpoint of the light-blue line has. Going to lower μ^2 , the phase diagram is periodic.

for the correct values of ν and γ (*collapse plot*).² However, as noted in [Bonati et al., 2011c], this estimate may be inconclusive, as different scenarios have similar numerical values of ν and γ and are therefore difficult to distinguish. In addition, histograms of various observables (L_{Re} , L_{Im} , plaquette) reweighted to β_c are considered. The authors state values for both tricritical masses: $am = 0.043(5)$ and $am = 0.72(8)$. In particular, they give an estimate of the lower tricritical mass, which should correspond to a pion mass of $m_\pi \approx 400$ MeV.

Relevant studies with pure Wilson fermions have been reported in two consecutive papers [Wu and Meng, 2013; Wu et al., 2007]. In both studies $N_\tau = 4$ systems are investigated with various spatial extents. As the first study actually aims at an analytic continuation to real μ , various values of μ_l below μ_l^c are considered and the transition of interest is the deconfinement/chiral transition. The authors investigate the behaviour of plaquette, chiral condensate and $|L|$ and report that for $\kappa = 0$ and 0.005 the transition is first order while it is a crossover for $\kappa = 0.165$. In perspective of the RW transition this means that the RW endpoint should be a first order triple point for the first two values of κ and a second order endpoint for the latter. The second study mentioned was published during the development of this thesis. Here, the authors investigate essentially the same setup as will be described in Section 5.2. Their studies cover nine values of κ , ranging from 0.155 to 0.198 at the critical

² This holds for other observables which fulfil similar scaling relations as well.

RW value $\mu_I = i\pi T$. Simulations were carried out on $N_\sigma = 8, 12, 16$ lattices, sometimes supported by $N_\sigma = 10$ or 20 . A resolution in β is not directly given, it is stated that four to six β values have been simulated for each κ . Figures 2 and 9, however, suggest that the resolution was $\Delta\beta = 0.01$. The simulated data was reweighted using Ferrenberg-Swendsen reweighting. Regarding statistics, 20k trajectories were performed with $\tau = 2$. In their analysis, they followed the one from [Bonati et al., 2011c] as described above, i.e. they looked at the scaling of $\chi(|L_{lm}|)$ according to (5.2) and the histograms of the HMC runs.³ In their conclusion, the authors stated to find first order RW endpoints for all investigated values of κ . Nonetheless, they noted that the distinction between the first order and second order scenarios is not always unambiguous. This is in accordance with the statement given in [Bonati et al., 2011c] mentioned above. Furthermore, they did not cite any results for a tricritical scaling analysis of their data. To support their findings, a scaling analysis of B_4 has been performed according to (5.9) for $\kappa = 0.165$ and 0.168 , yielding values of the critical exponent ν of 0.3661 and 0.3594 , respectively (no errors are given). This is in contradiction to the results obtained in [Wu et al., 2007] which imply a second order endpoint for $\kappa = 0.165$. Nevertheless, this is not commented on. The lattice spacing is estimated to be $0.12 - 0.25$ fm, based on literature values [Bitar et al., 1994; Iwasaki et al., 1996]. A pion mass at $\kappa = 0.16$ and $\beta = 3.28$ is estimated to be ≈ 2 GeV based on data from [Bitar et al., 1991]. However, this is not at the β_c at this κ , which is $\beta \approx 5.35$.

A very interesting and promising approach to LQCD was reported in [Fromm et al., 2012], where a dimensionally reduced effective theory for LQCD with Wilson fermions was presented. It is based on the *hopping expansion*:

$$D_{\text{Wilson}}^{-1}(n, m) = (1 - \kappa H(n, m))^{-1} = \sum_{j=0}^{\infty} \kappa^j H^j(n, m). \quad (5.3)$$

Here, the *hopping term* $H = \sum_{\pm\mu} (1 - \gamma_\mu)_{\alpha\beta} U_\mu(n)_{ab} \delta_{n+\hat{\mu}, m}$ makes the dependence of \mathcal{D} on κ explicit (compare (2.60)). For heavy quarks and small κ , respectively, the expansion may be truncated at a finite order in κ . In [Fromm et al., 2012], κ^2 terms are taken into account. This theory is then reduced to a three dimensional one by integrating out the spatial link variables. The resulting effective theory is thus much cheaper numerically than full LQCD simulations. In addition, the sign problem was found to be mild.

In particular, the value of the tricritical mass in the heavy mass region, $\kappa_{\text{heavy}}^{\text{tric}}$, is extracted from a fit to tricritical scaling, yielding (equation (3.13) in the paper):

$$M_{\text{tric}}/T = 6.25(3) \text{ for } N_f = 2. \quad (5.4)$$

This can be converted to κ by inserting equations (3.7) and (2.25) into (3.6) in [Fromm et al., 2012]⁴. For $N_\tau = 4$, which will be considered below, this yields:

$$\kappa_{\text{heavy}}^{\text{tric}} = 0.1048 \pm 0.0008. \quad (5.5)$$

Full LQCD results offer a direct check of the predictive power of the effective theory.

³ In the actual publication it is written that $\chi(L_{lm})$ is analysed. However, this is unlikely to signal a transition as is mentioned in the text below.

⁴ The applicability of (3.6) has been checked with (2.21).

transition type	crossover	2o 3D Ising	1o triple point	tricritical
ν	-	0.6301(4)	1/3	1/2
γ	-	1.2372(5)	1	1
$B_4(X)$	3	1.604	1.5	2

Table 13: Values for critical exponents ν and γ [Pelissetto and Vicari, 2002] and the binder cumulant $B_4(X)$ [de Forcrand and Philipsen, 2010] for different phase transitions.

5.2 SIMULATION SETUP & OBSERVABLES

This presentation is based on simulations using two flavours of standard Wilson fermions together with the Wilson gauge action and close to the analysis of simulations with three flavours of staggered fermions described in [de Forcrand and Philipsen, 2010]. To investigate the nature of the RW endpoint, simulations were carried out on one of the critical values μ_f^c . For convenience, this value was chosen to be $\mu_f^c = i\pi T$ (see below). The simulations were carried out using CL²QCD on LOEWE-CSC and SANAM. A fixed temporal extent, $N_\tau = 4$ was chosen. That means that temperature was varied by varying the lattice coupling β , (2.73).⁵ Simulations at various masses ranging from $\kappa = 0.03$ to 0.165 were performed. For each κ , at least three, sometimes four or five spatial volumes were simulated. The critical temperature T_c was mapped out approximately on the smallest lattice. Then, simulations at at least ten β -values around this estimate were performed for all lattice sizes with $\Delta\beta = 0.001$ in order to sample the transition region sufficiently. In each β run, 35k HMC trajectories of unit length ($\tau = 1$) were performed for the majority of the lattices after 5k trajectories of thermalisation. In some cases (e.g. $N_\sigma = 20$), this number was extended to 75k. The acceptance rate in each run was of the order of 75%. The data was analysed using jackknife estimators of the error (see Section 3.2). In here, an optimal blocksize was chosen from plateau behaviour of the jackknife error as the blocksize M was varied.⁶ Additional β -points were filled in using Ferrenberg-Swendsen reweighting.⁷ An overview of the simulated systems and obtained results can be found in Appendix A.9.

As was shown before, the Polyakov Loop (2.28) is closely related to the RW symmetry and is thus the observable of interest. At some spatial coordinate \mathbf{n} on the lattice, it is the trace over the temporal links:

$$L(\mathbf{n}) = \frac{1}{V} \text{Tr}_c \prod_{x_0=0}^{N_\tau-1} U_0(x_0, \mathbf{n}). \quad (5.6)$$

A temperature scan at $\kappa = 0.1$ and $N_\sigma = 12$ is shown in Figure 48. The raw data is interpolated nicely by the reweighted data. As was mentioned before, the depicted $|L|$ serves

⁵ As $a(\beta) \rightarrow 0$ for $\beta \rightarrow 0$, (2.50), high/small β means high/small T .

⁶ A single blocksize M was chosen for all β of one lattice volume.

⁷ This procedure was carried out using a programme provided by Owe Philipsen. It also calculates the peak position for susceptibilities.

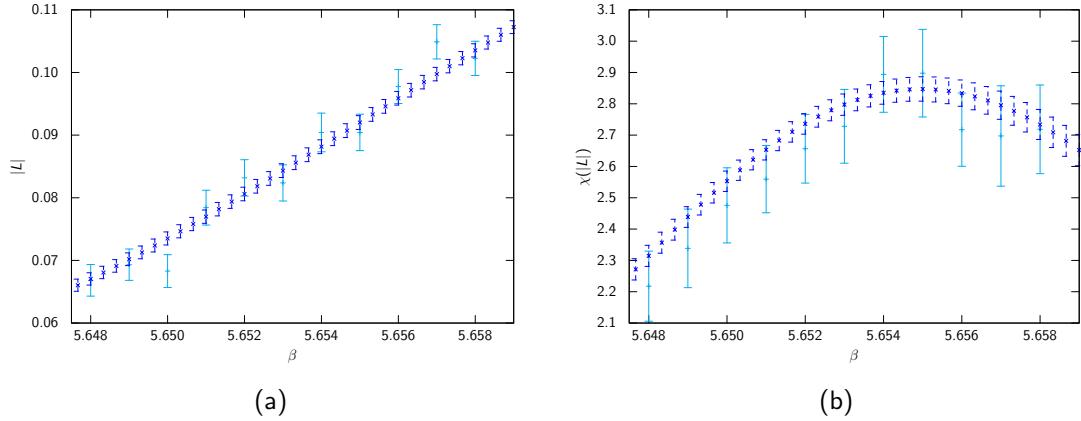


Figure 48: Temperature scans at $\kappa = 0.1$ and $N_\sigma = 12$. Data obtained from HMC runs is interpolated by reweighting. a) Absolute value $|L|$ of the Polyakov Loop. b) Susceptibility of $|L|$.

as an indicator for deconfinement. The transition point $\beta_c = 5.6550(3)$ is extracted from the peak of its susceptibility


$$\chi(|L|) = V \langle (|L| - \langle |L| \rangle)^2 \rangle, \quad (5.7)$$

which can be seen nicely in the figure.


For the six lowest masses, the chiral condensate $\langle \bar{\psi}\psi \rangle$ (2.63) was evaluated. This was done using stochastic estimators (Appendix A.5), using 96 volume sources of Gaussian noise on each configuration, taking into account every fiftieth generated configuration to rule out autocorrelation.

At the chosen value of μ_f^c , the RW transition is between the $\mathbb{Z}(3)$ sectors $\exp(\pm i2\pi/3)$. This allows for a clear identification of the $\mathbb{Z}(3)$ sectors by the phase of the Polyakov Loop (2.28), or, more practically, simply by its imaginary part L_{Im} , which takes on the values $\pm\sqrt{3}/2 |L|$ (see also Figure 44a). At temperatures above the RW endpoint, L_{Im} will hence jump between two distinct values as $|L| \gg 0$, whereas for $T \ll T_c$ the data will fluctuate around zero (as here $|L|$ is small). It should be noted that in both scenarios one has $\langle L_{\text{Im}} \rangle = 0$ for sufficiently large statistics. Only in the thermodynamic limit the system will eventually take on one of the possible values. Formally, this can be identified with the spontaneous breaking of the symmetry of the system under charge conjugation [Wu and Meng, 2013].

The described behaviour is shown in Figure 49a, which depicts L_{Im} above and below T_c . One sees that the high temperature phase shows characteristics clearly different from the low temperature one, in accordance to the expectations. In particular, one can see that in the low temperature phase, L_{Im} can reside at zero for multiple trajectories, where in the high temperature phase this value is only taken on shortly as the system tunnels to the other $\mathbb{Z}(3)$ sector. Figure 49b shows the susceptibility of the absolute value of L_{Im} , whose peak indicates the passing through the RW endpoint. The absolute value is mandatory here: As β is increased beyond T_c the susceptibility of L_{Im} will increase, too, due to the tunnelling between the two $\mathbb{Z}(3)$ sectors. The peak position at $5.6546(3)$ coincides with the one obtained from $|L|$. Thus, judging from these two observables, deconfinement and RW transition happen at the same



(a)



(b)

Figure 49: L_{lm} at $\kappa = 0.1$ and $N_\sigma = 12$. a) HMC history above and below T_c . b) Susceptibility of $|L_{lm}|$, data obtained from HMC runs is interpolated by reweighting.

temperature, as expected from previous studies. However, it should be noted that $|L|$ and $|L_{lm}|$ are, of course, not independent of each other.

An indicator for the order of the transition is the histogram of L_{lm} . As the temperature passes through T_c , the distribution changes from a peak around zero to a two peak structure, one for each of the two $\mathbb{Z}(3)$ sectors, located at non-zero values of L_{lm} . In the picture of an effective potential, this corresponds to the change of the global minimum, see also Figure 4. In the vicinity of the transition point, L_{lm} fluctuates between the three values. For a first order transition, the histogram would then show a three peak structure, as the system jumps between




Figure 50: Histogram of $|L_{lm}|$ at $\kappa = 0.04$ (a) and $\kappa = 0.12$ (b) ($N_\sigma = 16$). Both simulations have the same statistics.

the minima of the effective potential. In a second order transition, in turn, there would be a rather smooth, plateau-like behaviour, as the transition is continuous. This difference is shown in Figure 50, which depicts the histograms for two different κ . Note that for visualisation purposes the absolute value of L_{lm} is plotted, hence the two peaks around non-zero values of L_{lm} are merged to one. One sees that for the lower κ , the histogram shows two peaks around the transition point (light blue curve in the figure), while below and above there are distinct peaks at zero and non-zero values of $|L_{lm}|$, respectively. While the latter is also true for the higher κ , a plateau is visible around the transition temperature here. This indicates a second order transition. However, it is important to note that this can only be an estimate of the order of the transition. If the volume is too small, the system can tunnel between the minima of the potential in a quite large region around T_c . For larger volumes, the tunnelling is suppressed, so one might sample the transition region insufficiently and not see it in the histogram. And, after all, on the lattice there is always a crossover due to the finite system extents.

The transition at the RW endpoint is a true phase transition only in the thermodynamic limit $V \rightarrow \infty$. Thus, to extract it from finite volume LQCD simulations, an extrapolation to the thermodynamic limit must be applied (*finite-size scaling*), where the spatial volume is enlarged keeping the other parameters fixed. A particular useful quantity to extract the order of the phase transition from is the so-called *Binder cumulant* [Binder, 1981]. For a general observable X , it is defined as

$$B_4(X) = \langle (X - \langle X \rangle)^4 \rangle / \langle (X - \langle X \rangle)^2 \rangle^2. \quad (5.8)$$

B_4 indicates the order of the transition by its value in the thermodynamic limit (see Table 13). Loosely speaking, B_4 counts the number of transitions between the different minima of the effective potential. In the vicinity of the transition point β_c , the infinite volume value $B_4(\beta, \infty)$ will be affected by finite-size corrections. These can be described by the ansatz [de Forcrand and Philipsen, 2010]

$$B_4(\beta, N_\sigma) = B_4(\beta, \infty) + a_1(\beta - \beta_c)N_\sigma^{1/\nu} + a_2((\beta - \beta_c)N_\sigma^{1/\nu})^2 + \dots \quad (5.9)$$




Figure 51: Reweighted Binder Cumulant of L_{fm} at $\kappa = 0.07$ for various N_σ . Also shown is the fitted scaling function (5.9).

Figure 51 shows the functional behaviour of B_4 as the spatial volume is increased. One sees that B_4 takes on higher values at lower β , going to lower values at higher β , and gets steeper as the volume is increased. This is expected as below and above β_c a crossover and first order region is located, which has a B_4 value of 3 and 1.5 in the thermodynamic limit, respectively. In fact, approaching the thermodynamic limit B_4 should become a step function. The intersection of the three curves at β_c is the location of the RW endpoint.

5.3 ROBERGE-WEISS TRANSITION AS A FUNCTION OF THE FERMION MASS

A finite size scaling analysis for B_4 has been carried out for all simulated values of κ . The scaling form (5.9) is fitted to the B_4 data in the vicinity of β_c , see also Figure 51. An overview of these fits can be found in Table 24 in Appendix A.9. In addition, for the majority of the simulation points, β_c was extracted from the susceptibility peak for $|L|$, $|L_{\text{fm}}|$ and $\langle \bar{\psi}\psi \rangle$. These values are listed in Appendix A.9.

The obtained value for $B_4(\beta, \infty)$ is found to be slightly higher than the universal values, in accordance with [de Forcrand and Philipsen, 2010]. The critical exponent ν , however, can be extracted quite well. The result of this effort can be seen in Figure 52. The extracted value of ν takes on first order values for small κ (high m). When κ is increased, it passes through the tricritical value until it takes on second order values. As the mass is decreased further, a




Figure 52: Fitted critical exponent ν as a function of κ . Also shown are values of ν for certain universality classes as well as the prediction for the tricritical mass from the effective theory [Fromm et al., 2012].

reverse behaviour can be seen: ν approaches the first order value again, passing through the tricritical value at $\kappa = 0.155$. Thus, the two tricritical κ are determined to be:

$$\kappa_{\text{heavy}}^{\text{tric}} = 0.1000 \pm 0.090, \quad (5.10)$$

$$\kappa_{\text{light}}^{\text{tric}} = 0.1550 \pm 0.050. \quad (5.11)$$

Note that the errors on these values are rather conservative and chosen in a way that the neighbouring simulation points which clearly fall into the first and second order scenario, respectively, are taken as a boundary for the tricritical masses. In particular, the fitted ν for $\kappa = 0.1575$ is found to be a little higher than the one for $\kappa = 0.1550$.

These findings can be supported by additional observables. First of all, one can compare the transition points from the fit to the B_4 data to those from the other observables. Deviations are to be expected as $\beta_c(B_4)$ is determined from data of multiple N_σ and thus it should not be affected by large finite-size corrections. The ratios of the various β_c with regard to $\beta_c(B_4)$ are depicted in Figures 53, 54 and 55. It is clearly visible that the biggest deviations are found on $N_\sigma = 8$ lattices. An important observation is that the deviations are always below 0.1%, that is within the β resolution of the simulations (which is $\Delta\beta = 0.001$). This shows that the obtained results are self-consistent.

In the small κ region the RW endpoint is a triple point, thus, the deconfinement transition has to be first order. One can see signals supporting this scenario in the histograms of $|L|$, L_{Im} and L_{Re} (which are similar to Figure 50a). Actually, the influence of μ_I is still low. L_{Re} takes on positive values in the deconfinement phase, which one would not expect as both $\mathbb{Z}(3)$




Figure 53: Ratio of β_c as determined from the peak of the susceptibility of $|L_{Im}|$ and β_c extracted from the fit to the Binder cumulant of L_{Im} according to (5.9).




Figure 54: Same as Figure 53 with β_c from $|L|$.




Figure 55: Same as Figure 53 with β_c from $\langle\bar{\psi}\psi\rangle$.



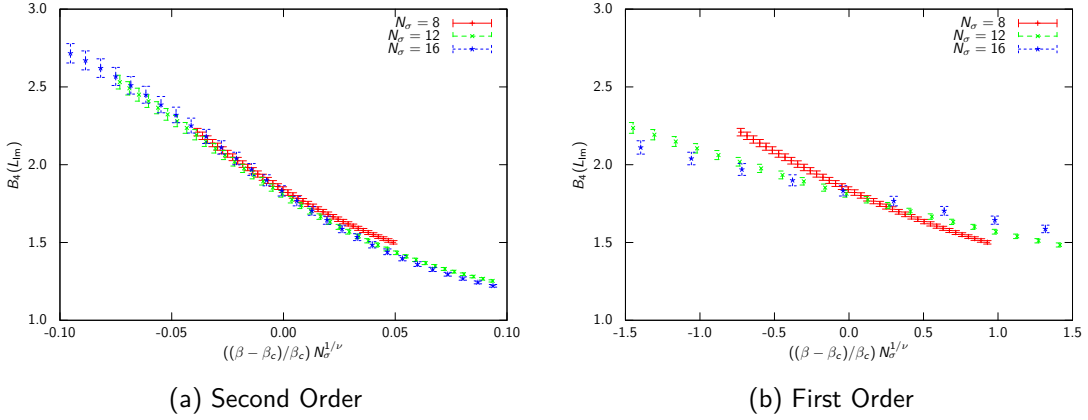


Figure 56: β_c as a function of κ , extracted from fits to B_4 data according to (5.9). Also shown is a fit of the data to a fifth grade polynomial.

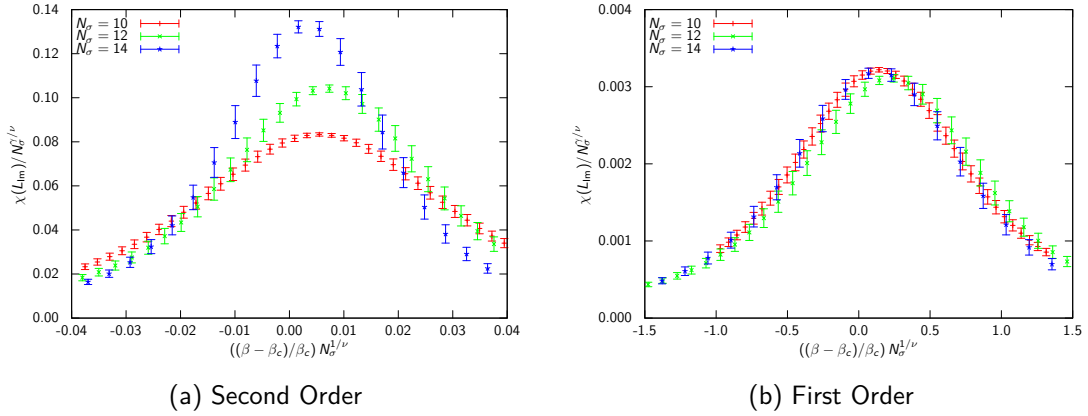
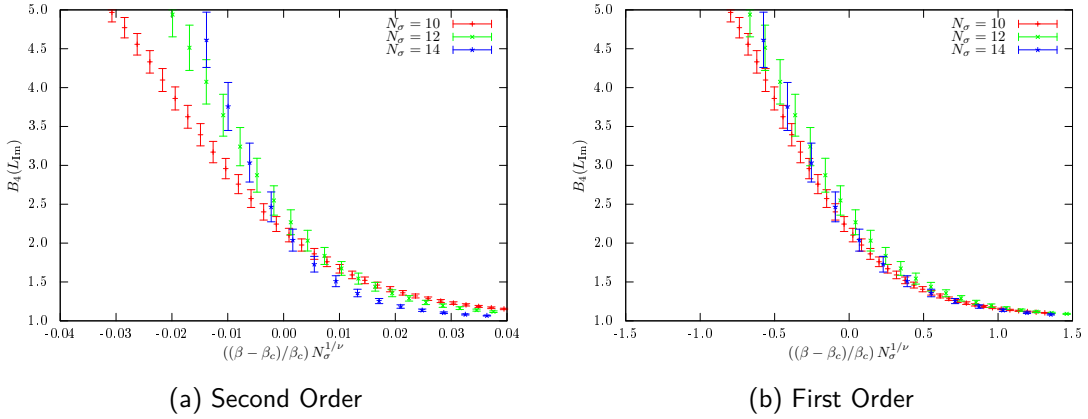
Figure 57: Collapse plots of $\chi(|L_{\text{Im}}|)$ for different critical exponents at $\kappa = 0.13$.Figure 58: Collapse plots of $B_4(L_{\text{Im}})$ for different critical exponents at $\kappa = 0.13$.

sectors selected by μ_I have negative L_{Re} (see Figure 44a). This is simply because the mass is very high and it dominates the fermionic contribution to the system. In fact, the overall influence of the fermions is still small and the phase transition is of first order as in pure gauge theory. This picture is supported by the value of β_c extracted from the fit to B_4 , (5.9), which shows a significant fermionic influence as of $\kappa \gtrsim 0.085$ only, see Figure 56.⁸ This effect also shows up in Figure 53: For the smaller κ , no β_c could be extracted from $|L_{\text{Im}}|$, as there was no peak in its susceptibility (instead in that of L_{Re}).

The aforementioned collapse plots can also be used to support the found values of ν . Figures 57 and 58 show the appropriate depiction of $\chi(|L_{\text{Im}}|)$ and $B_4(L_{\text{Im}})$. Here, the reduced temperature t has been replaced by a “reduced β ” $(\beta - \beta_c)/\beta_c$. Both figures show that the data collapses if second order exponents are used, whereas first order exponents do not lead to this result.⁹ As was mentioned before, this method can be inconclusive if different exponents take on similar values. This can be seen in Figure 57: The collapse plot with tricritical exponents

⁸ The fit shown in Figure 56 has been used to estimate β_c for new simulation points.



⁹ The equivalent of Figure 57 with first order exponents is not shown, but the data does not collapse, similar to Figure 58b.

Figure 59: Collapse plots of $\chi(|L_{\text{Im}}|)$ for different critical exponents at $\kappa = 0.165$.Figure 60: Collapse plots of $B_4(L_{\text{Im}})$ for different critical exponents at $\kappa = 0.165$.

is very similar to the second order one, since γ/ν is numerically similar (1.963 and 2). Thus, a discrimination of the two scenarios can clearly be carried out better by the direct fit of ν .

The RW endpoint should also be a triple point in the chiral region, where the RW transition meets with a first order chiral transition. The latter is indicated by the melting of the chiral condensate $\langle\bar{\psi}\psi\rangle$ and the peak of its susceptibility, respectively (see also Section 2.2.1). This is shown exemplarily in Figure 61, which shows $\langle\bar{\psi}\psi\rangle$ at $\kappa = 0.165$ for different N_σ . When the volume of the system is increased, the region, where $\langle\bar{\psi}\psi\rangle$ melts, shrinks and the gradient gets steeper. This is what is expected to happen in a first order transition. Similar behaviour has been observed at $\kappa = 0.1600$ and 0.1625 . Here, the melting is not as steep as in the presented figure, in agreement with a strengthening of the first order transition as κ is increased. In addition, the collapse plot of $\chi(\langle\bar{\psi}\psi\rangle)$ in Figure 62 clearly shows first order behaviour. It should be noted that the β_c extracted from $\langle\bar{\psi}\psi\rangle$ is in very good agreement with that from $B_4(L_{\text{Im}})$, Figure 55. Next to the chiral transition, also the RW transition shows clear first order signals, as the collapse plots in Figures 60 and 59 reveal.

In order to estimate the m_π at the tricritical point at high κ , $T = 0$ simulations at $\kappa = 0.1575$ on a $16^3 \times 32$ lattice were performed on JUQUEEN. This was done at the critical β of

Figure 61: $\langle \bar{\psi}\psi \rangle$ at $\kappa = 0.165$ as N_σ is increased.Figure 62: Collapse plot of $\chi(\langle \bar{\psi}\psi \rangle)$ at $\kappa = 0.165$ according to first order exponents.

5.355 and at $\mu = 0$. In particular, 4200 trajectories were performed after 500 trajectories of thermalisation. On these, the effective masses for the pion and rho particle were estimated to be $am_\pi = 1.1426(17)$ and $am_\rho = 1.2147(25)$,¹⁰ which corresponds to a ratio $m_\pi / m_\rho = 0.94064$. Compared to the physical value of 0.18003,¹¹ one sees that the pion is rather heavy. This can

¹⁰ This analysis was carried out by my colleague Georg Bergner.

¹¹ According to the PDG, the physical values are $m_\pi = 139.57028(35)$ MeV and $m_\rho = 775.26(29)$ MeV.

be estimated by setting am_ρ equal to the physical value, yielding $m_\pi \approx 729(2)$ MeV and a lattice spacing of $a \approx 0.3$ fm.

5.4 SUMMARY AND DISCUSSION

In this chapter, the nature of the RW endpoint in $N_f = 2$ LQCD with Wilson fermions was addressed. This was done at the critical value of the chemical potential, μ_f^c , on $N_\tau = 4$ lattices for various mass parameters κ , ranging from 0.03 to 0.165 and 3-5 spatial lattice extents for each κ . On each point, a huge amount of HMC trajectories was run, scanning the critical region in β with a quite fine resolution of $\Delta\beta = 0.001$. The simulations were carried out using CL²QCD, which is explained in Chapter 3. The main observable was the Binder cumulant B_4 of the imaginary part of the Polyakov Loop, (2.28). The nature of the RW endpoint was investigated by means of finite-size scaling analysis of $B_4(L_{\text{Im}})$ over the whole range of κ according to the ansatz given in (5.9). From this the critical exponent ν has been extracted, which takes on distinct values in the different phase transition scenarios. Two first order regimes at high and low masses were mapped out, separated by a second order region. These findings have been supported by different suited observables and collapse plots. At the boundaries, tricritical points were found, $\kappa_{\text{heavy}}^{\text{tric}} = 0.1000 \pm 0.090$ and $\kappa_{\text{light}}^{\text{tric}} = 0.1550 \pm 0.050$.

This completely confirms the picture based on staggered studies, see Figure 44b. In addition, the prediction of $\kappa_{\text{heavy}}^{\text{tric}}$ from an effective theory for heavy quarks [Fromm et al., 2012] was found to be quite accurate. This encourages further studies using this approach. Compared to a very similar study with Wilson fermion that was published during the development of this project [Wu and Meng, 2013], a generally consistent behaviour is found. In the latter, the light mass region is investigated, where a first order behaviour is seen up to $\kappa = 0.198$, i.e. the chiral first order region seems to extend beyond the κ range investigated in this thesis. Nevertheless, at the lowest κ studied in [Wu and Meng, 2013], 0.155, a first order behaviour is seen instead of a tricritical one. However, their classification is mostly based on collapse plots, which has been reported to be inconclusive [Bonati et al., 2011c] and may cause the different findings. In general, the study presented in this thesis relies on much more data compared to [Wu and Meng, 2013] and has a finer resolution in β . This is most likely the explanation for the observed deviations.

The pion mass at the tricritical for lower masses was estimated roughly to be 729(2) MeV. The much higher m_π in the estimation of [Wu and Meng, 2013] may be caused by their interpolation of literature data, which is not explained in detail. The estimation presented here is also higher than the estimation from the staggered study, which saw 400 MeV pion mass. Of course, this study relies on a completely different fermion discretisation, which might be a simple explanation for this discrepancy. Nevertheless, the estimation of the lattice spacing is the same, $a \approx 0.3$ fm. Thus, both results suggest that at physical masses the RW transition is first order, although a continuum limit is yet missing.

CONCLUSIONS AND RESEARCH PERSPECTIVES

In this thesis, the theory of the strong interactions, QCD, has been investigated on the lattice as a very interesting interplay of theoretical physics and computer science. A new LQCD application was successfully developed and used in physics studies with two flavours of Wilson fermions. To conclude, each topic is briefly summarised in the following and research perspectives are given. More extensive summaries are given at the ends of Chapters 3, 4 and 5.

The numerical aspects of LQCD are a vital part of this research field and display a very interesting mixture of physics and computer science. In particular, it is discussed in Chapter 3 why LQCD is numerically challenging and the use of Graphics Processing Units (GPUs) is motivated. The OpenCL-based LQCD application CL²QCD¹ was developed within this thesis [Bach et al., 2013a] as a very fruitful collaboration between the Institute for Theoretical Physics and the department of High Performance Computer Architecture² in Frankfurt. The project has instantly drawn interest within the lattice community.³ During the development the focus was laid on the HMC application and the excellent performance results obtained on AMD hardware. It underlines the great advantages GPUs can offer for LQCD, but also provides an alternative to existing vendor-restricted applications based on NVIDIA's CUDA. As a real-life application, CL²QCD was successfully applied as the main tool for the physics studies presented in this thesis, which clearly could not have been carried out to this extent without the computational power CL²QCD can supply.

The functionality provided by CL²QCD has been implemented mainly for the purposes of these physics studies. It has become quite complex, at the time of writing CL²QCD consists of $\mathcal{O}(100k)$ lines of code. Yet, it can easily be extended, especially due to a clear separation of high- and low-level functionality. For example, a working test implementation of the Langevin algorithm, which possibly might have a milder sign problem at finite μ than the HMC [Sexty, 2013], has been achieved within a day. For Wilson fermions, a variety of extensions to the functionality already available is possible, e.g. a clover term is implemented by means of a pure gauge observable, simplifying the potential adding of clover Wilson fermions. Also, basic functionality for stout-smearing [Morningstar and Peardon, 2004] is available, which might be relevant for future applications. Furthermore, a currently on-going project by my colleague

1 See <http://code.compeng.uni-frankfurt.de/projects/clhmc> .

2 See <http://www.compeng.de/> .

3 See for example the plenary talk by B. Joó at the Lattice Conference 2011, <http://tsailab.chem.pacific.edu/lat11/plenary/joo/JooGPUs.pdf> .

Alessandro Sciarrà consists of implementing staggered fermions, which adds an additional fermion discretisation. With this at hand, $N_f = 2 + 1$ Wilson fermions can be added quite easily, as this requires partially the same functionality as the algorithms for staggered fermions. In addition, Multi-GPU usage was added recently [Bach et al., 2013b]. This is mandatory for studies on larger lattices as the on-board memory of the GPU is the strongest restriction for the lattice size in the simulation. This implementation currently parts the lattice in time direction only, which can be generalised in the future. This can be advantageous for finite temperature studies, where the spatial extent of the lattice is larger than the temporal one. On the algorithmic side, the solver performance can be optimised further by implementing a mixed-precision solver [Clark et al., 2010]. Here, performance is accelerated by performing the majority of iterations in single precision and repeated correction steps in double precision. Another interesting possibility would be to investigate performance optimisations on CPUs. First attempts to use CPUs and GPUs in a hybrid fashion have been reported in [Philipsen et al., 2011]. Furthermore, in principle it is possible to use CL²QCD on large-scale CPU systems like JUQUEEN by replacing the low-level functionality. In this way, the versatility of CL²QCD would be enlarged tremendously. Heading in the same direction, one might extend CL²QCD to use multiple compute nodes inside a cluster, e.g. via MPI. This would also be relevant for Multi-GPU usage.

It is not known if the $N_f = 2$ chiral transition in QCD is of first or second order. This important question was addressed in Chapter 4 by fully extending a study by the tmft collaboration [Burger et al., 2013a] using Twisted Mass Wilson fermions on $N_\tau = 12$ lattices. The Z12 ensemble at $m_\pi \approx 270$ MeV was added to the existing ensembles and analysed by means of the chiral condensate $\langle \bar{\psi}\psi \rangle$ and screening masses. Both analyses support the second order $O(4)$ scenario favoured in the previous study, even so different sources of uncertainties are present. In general, the analysis gets exceedingly difficult the smaller m_π . The signal of the susceptibility of $\langle \bar{\psi}\psi \rangle$ does not show a distinct peak at lower m_π but a plateau. This is also seen in the lightest ensemble studied in [Burger et al., 2013a], the A12 ensemble at $m_\pi \approx 316$ MeV, which was enlarged for this purpose. Yet, the Z12 ensemble exceeds the A12 ensemble in terms of statistics as well as resolution in β . This suggests to add more points to A12 in order to refine the β_c determination here. In addition, the scale setting seems to be problematic at the investigated β range. It should be checked that this procedure is safe, i.e. that the rise of the susceptibility signal is physical. This can be done by repeating the scan of the Z12 temperature region with a larger N_τ , which would then move the β range to higher values for the same temperature scope. There seems to be no scale setting problems around $\beta \approx 3.9$, judging from the signals reported in [Burger et al., 2013a]. In addition, using an updated scale setting described in Chapter 4.3 might be beneficial, as the errors are decreased for most of the parameters and temperature is constantly shifted to lower values, compared to the used scale setting. This, similar to enlarging N_τ , would have the advantage to shift simulations to higher β . Indeed, the large errors of the current scale setting in the Z12 β range prevent more precise constraints on the scaling of T_c towards the chiral limit (cf. Figure 28), despite the fact that the resolution in β is better than for the preceding ensembles.

An obvious extension of this project would be to go to even lower pion masses. This should be done at least with $N_\tau = 16$ due to the arguments given above. In a broader picture, it is possible to extend the twisted mass formulation to the more physical case of $N_f = 2 + 1 + 1$, see [Jansen and Urbach, 2009] for details on this topic and further references. Currently, this setup is investigated by the ETM collaboration around the physical point [Abdel-Rehim et al., 2013a], allowing for future thermodynamical studies. A first approach in this direction has been reported in [Burger et al., 2013b].

The development of CL²QCD and its usage in the simulations has been crucial to produce the high statistics for the Z12 ensemble. Despite the uncertainties in the analysis of the data, this has been a successful test of the application in a large-scale investigation.

The region of imaginary chemical potential was investigated in Chapter 5, in particular the nature of the Roberge-Weiss endpoint was studied. Using two flavours of Wilson fermions on $N_\tau = 4$ lattices, the picture based on staggered studies [Bonati et al., 2011c; de Forcrand and Philipsen, 2010] was clearly confirmed. For large and small masses, the endpoint is a triple point and changes to a second order $Z(2)$ point for intermediate masses. At the boundaries, tricritical points are located, which were determined from the simulations by means of the Binder cumulant B_4 . This behaviour could be nicely extracted from the critical exponent ν .

The confirmation of the mass of the tricritical point in the heavy mass region predicted by a heavy quark effective lattice theory [Fromm et al., 2012] has been particularly interesting, as the latter is numerically tremendously less expensive than full LQCD simulations. As the tricritical mass in the light mass region was estimated to be rather heavy, around 700 MeV, it would be interesting to study the model for these mass region as well. A possible confirmation of the results obtained in this thesis by the effective theory would open up new possibilities to explore the phase diagram of QCD, especially as the sign-problem in the effective theory is quite mild [Philipsen et al., 2012].

The project was particularly suited to run on GPUs as the lattice sizes are rather small, hence the available GPU memory is not an issue. In particular, LOEWE-CSC and SANAM offered enough resources to trivially parallelise the simulations. Therefore, this study relies on a huge amount of data generated with CL²QCD mentioned above, allowing for a more thorough analysis than the similar study [Wu and Meng, 2013].

For two flavours, possible cutoff effects can be studied by increasing N_τ . This would in turn correspond to a tremendous increase in numerical costs of the simulations, which were not negligible already in this study. As a continuation of the results presented, it might be interesting to go to even lower masses. In particular, investigating the unphysical Aoki phase in the region of imaginary μ and how it possibly affects the RW transition would be interesting. Furthermore, it is tempting to address the $N_f = 2$ chiral limit as discussed in Chapter 5.1 by following the critical line emerging from the tricritical point, as the nature of the chiral transition is of high interest and was also addressed in the study discussed above. This investigation would not require additional algorithmic functionality. Additionally, adding one flavour and study the Roberge-Weiss transition in $N_f = 2 + 1, 3$ would allow to compare to the staggered results in these limits.

A

APPENDIX

In this Appendix, after fixing notations and conventions (Section A.1), certain rather technical aspects of the topics described in this thesis are explained. For instance, group theoretical aspects of QCD are elucidated in Section A.2. In Section A.3, Grassmann calculus will be introduced. On the algorithmic level, even-odd preconditioning is addressed in Section A.4 while details on the calculation of the inverse fermion matrix D^{-1} are given in Section A.5. Details of the implementation of CL²QCD and of computing resources used during this thesis are given in Sections A.6 and A.7. In the last two sections, A.8 and A.9, simulation details on the studies described in Chapters 4 and 5 are provided.

A.1 NOTATIONS AND CONVENTIONS

Throughout this thesis, *natural units* will be used (See e.g. [Peskin and Schroeder, 1995]). In these, Planck's constant \hbar , the speed of light c and the Boltzmann constant k_b are all set to unity. All quantities are then given in units of (inverse) energy, which is denoted by electron volt (eV). To convert into standard units, one can use for example:

$$\hbar c = 197.3 \text{ fm MeV} . \quad (\text{A.1})$$

Standard notation for covariant (lower index) and contravariant (upper index) vectors is used. Greek and Latin letters denote Dirac and colour indices, respectively. When unambiguous, indices will be suppressed and summation over repeated indices is understood.

In *Quantum Field Theory* the dynamics of a field ϕ is described by its action

$$\mathcal{S}[\phi] = \int d^4x \mathcal{L}(\partial_\mu \phi, \phi) , \quad (\text{A.2})$$

which is the integral over the Lagrangian density $\mathcal{L}(\partial_\mu \phi, \phi)$ (See e.g. [Peskin and Schroeder, 1995]). Via a Legendre transformation one can switch to the Hamiltonian formulation with density \mathcal{H} . The Hamiltonian H is then $\int d^3x \mathcal{H}$.

In Minkowski space time, the metric tensor $g_{\mu\nu}$ is $\text{diag}(1, -1, -1, -1)$ and $\text{diag}(1, 1, 1, 1)$ in Euclidean space time. Thus, in the latter co- and contravariant vectors are formally identifiable. Throughout this thesis, time direction is denoted as zero component of vectors and spatial directions x , y and z as the first, second and third component, respectively.

Commonly used matrices throughout the text are the Pauli and Dirac matrices, which will therefore be given explicitly here. The Pauli matrices are defined as

$$\tau_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tau_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \tau_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (\text{A.3})$$

The Euclidean Dirac matrices γ_μ satisfy

$$\{\gamma_\mu, \gamma_\nu\} \equiv \gamma_\mu \gamma_\nu + \gamma_\nu \gamma_\mu = 2\delta_{\mu\nu}. \quad (\text{A.4})$$

They are chosen to be

$$\begin{aligned} \gamma_0 &= \begin{pmatrix} 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix}, & \gamma_1 &= \begin{pmatrix} 0 & 0 & 0 & -i \\ 0 & 0 & -i & 0 \\ 0 & +i & 0 & 0 \\ +i & 0 & 0 & 0 \end{pmatrix}, \\ \gamma_2 &= \begin{pmatrix} 0 & 0 & 0 & -1 \\ 0 & 0 & +1 & 0 \\ 0 & +1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix}, & \gamma_3 &= \begin{pmatrix} 0 & 0 & -i & 0 \\ 0 & 0 & 0 & +i \\ +i & 0 & 0 & 0 \\ 0 & -i & 0 & 0 \end{pmatrix}, \end{aligned} \quad (\text{A.5})$$

according to [Jansen and Urbach, 2009] (*chiral representation*). The product of all γ_μ is the so-called γ_5 :

$$\gamma_5 \equiv \gamma_0 \gamma_1 \gamma_2 \gamma_3 = \begin{pmatrix} +1 & 0 & 0 & 0 \\ 0 & +1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \quad (\text{A.6})$$

It anti commutes with all other γ_μ . All γ -matrices are Hermitian and obey $\gamma_\mu^2 = \mathbb{1}$.

A.2 $SU(N)$ AND GROUP THEORY

As QCD is formulated as a $SU(N_c)$ gauge theory, some of the properties of $SU(N)$ groups are introduced in the following. For a more fundamental introduction see [Georgi, 1982] or the dedicated sections in [Gattringer and Lang, 2010; Montvay and Münster, 1994; Peskin and Schroeder, 1995].

Within the scope of this thesis, $SU(N)$ is a group of complex $N \times N$ matrices U , obeying

$$\det U = 1 \quad \text{and} \quad U^\dagger = U^{-1}. \quad (\text{A.7})$$

Thus, they have $N^2 - 1$ independent, real entries. For $N > 1$, $SU(N)$ is non-abelian. In addition, it is not closed under the addition of matrices.

Furthermore, $SU(N)$ is a *Lie group*, and its elements can be represented by elements of the corresponding *Lie algebra* $\mathfrak{su}(N)$:

$$U = \exp \left(i \sum_a \omega^a T_a \right). \quad (\text{A.8})$$

The algebra is spanned by the *generators* T_a , $a = 1, \dots, N^2 - 1$. These are traceless and Hermitian matrices which are usually normalised to $\text{Tr}([T_a, T_b]) = \frac{1}{2} \delta_{a,b}$. They obey

$$[T_a, T_b] = i f_{abc} T_c. \quad (\text{A.9})$$

Elements of the algebra are linear combinations of the generators, $\sum_a \omega^a T_a$, with coefficients $\omega^a \in \mathbb{R}$.

Important examples are $SU(2)$ (Isospin) and $SU(3)$ (QCD). Their generators are usually chosen to be (up to normalisation) the *Pauli matrices* τ_a and *Gell-Mann matrices* λ_a , respectively, see for example [Gattringer and Lang, 2010].

Derivatives of $SU(N)$ elements live in the algebra $\mathfrak{su}(N)$. For a direction a in algebra space, it can be defined as

$$(\partial f[U] / \partial U)^a \equiv \frac{\partial f[U]}{\partial \omega^a} = \frac{\partial}{\partial \omega} f[e^{i\omega T_a} U] \Big|_{\omega=0}. \quad (\text{A.10})$$

An example would be the force term $F_{\text{plquette}}^{\text{gauge}}$ (3.11). Suppose that the link U lives at a specific lattice site n . Then the non-trivial part of the gauge action (2.55) is proportional to

$$\text{Tr}(U\tilde{P} + \tilde{P}^\dagger U^\dagger).$$

\tilde{P} denotes the staple product of links. The derivative at site n in direction a then reads:

$$\begin{aligned} \frac{\partial \mathcal{S}_{\text{gauge}}}{\partial \omega^a(n)} &\sim \frac{\partial}{\partial \omega} \text{Tr} \left(e^{i\omega T_a} U\tilde{P} + \tilde{P}^\dagger e^{-i\omega T_a} U^\dagger \right) \Big|_{\omega=0} \\ &= i \text{Tr} \left[T_a \left(U\tilde{P} - \tilde{P}^\dagger U^\dagger \right) \right]. \end{aligned} \quad (\text{A.11})$$

A.3 GRASSMANN CALCULUS

The fermion fields ψ must obey Fermi statistics. This means that they must be antisymmetric under the exchange of (fermionic) quantum numbers. To this end, *Grassmann numbers* η were introduced (see for example [Gattringer and Lang, 2010]), which obey

$$\eta_i \eta_j = -\eta_j \eta_i. \quad (\text{A.12})$$

In particular, they are nilpotent, $\eta^2 = 0$. These two properties have severe implications for derivatives and integrals. In particular for a fermionic Gaussian integral with complex matrix M one has (*Matthews-Salam formula*) [Matthews and Salam, 1954, 1955]

$$\int d\eta_N d\bar{\eta}_N \dots d\eta_1 d\bar{\eta}_1 \exp \left(\sum_{i,j=1}^N \bar{\eta}_i M_{ij} \eta_j \right) = \det M. \quad (\text{A.13})$$

Here, the $\bar{\eta}$ denote anti-fermions and are independent of the η . This formula can be generalised to give *Wick's theorem* (see [Gattringer and Lang, 2010]), which relates fermionic n -point functions to fermion propagators:

$$\langle \eta_{i_1} \eta_{j_1} \dots \eta_{i_n} \eta_{j_n} \rangle = (-1)^n \sum_P \text{sign}(P) (M^{-1})_{i_1 j_{P_1}} \dots (M^{-1})_{i_n j_{P_n}} . \quad (\text{A.14})$$

Here, P denotes permutations over the indices $1, \dots, n$ and $\text{sign}(P)$ is its signum. Note that the n -point function vanishes if there is not the same number of η and $\bar{\eta}$.

A.4 EVEN-ODD PRECONDITIONING

In this section, even-odd preconditioning [Degrand and Rossi, 1990] will be explained. It is a technique to precondition matrix inversions like

$$M\mathbf{x} = \mathbf{b} \Rightarrow \mathbf{x} = M^{-1}\mathbf{b} . \quad (\text{A.15})$$

The lattice is split into even and odd sites, depending on their parity $(-1)^{x_0+x_1+x_2+x_3}$, i.e. $\mathbf{x} = (\mathbf{x}_e, \mathbf{x}_o)$ (see also Figure 63b). The matrix M can then be divided into parts that act on even and odd sites only, respectively, and simplified by decomposing it into three parts (*Schur's Lemma*):

$$M = \begin{pmatrix} M_{ee} & M_{eo} \\ M_{oe} & M_{oo} \end{pmatrix} = \begin{pmatrix} 1 & M_{eo}M_{ee}^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} M_{ee} - M_{eo}M_{oo}^{-1}M_{oe} & 0 \\ 0 & M_{oo} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ M_{oo}^{-1}M_{oe} & 1 \end{pmatrix} \\ \equiv L \tilde{M} U \quad (\text{A.16})$$

The triangular matrices L and U can be inverted trivially. (A.15) then becomes:

$$L \tilde{M} U\mathbf{x} = \mathbf{b} \Leftrightarrow \tilde{M} U\mathbf{x} = L^{-1}\mathbf{b} \Leftrightarrow \tilde{M}\tilde{\mathbf{x}} = \tilde{\mathbf{b}} . \quad (\text{A.17})$$

with $\tilde{\mathbf{x}} = L^{-1}\mathbf{x}$ and $\tilde{\mathbf{b}} = U\mathbf{b}$. In QCD, the diagonal parts of M can be inverted analytically and thus, it is sufficient to solve only the system $(M_{ee} - M_{eo}M_{oo}^{-1}M_{oe})\tilde{\mathbf{x}}_e = \tilde{\mathbf{b}}_e$ numerically. Then, the full solution \mathbf{x} can be constructed easily.

A.5 ESTIMATION OF D^{-1}

In this section, details on the calculation of the inverse of the fermion matrix, D^{-1} , from different types of sources are given.

(2.62) states that each entry of D^{-1} corresponds to the propagator from one definite site and definite colour and Dirac indices, (n_0, α_0, a_0) , to another such pair, (m_0, β_0, b_0) . This is also called *source* and *sink*. In the same fashion, one column of D^{-1} collects all propagation from one site, given by the column index, to all other sites (*point-to-all* propagator). Denote this column as a vector $\phi[n_0, \alpha_0, a_0]$. To project the specific column out of D^{-1} , define a vector

$$S_0[n_0, \alpha_0, a_0] = S_0[n_0, \alpha_0, a_0]_a^\alpha(n) \equiv \delta(n - n_0)\delta(\alpha - \alpha_0)\delta(a - a_0) , \quad (\text{A.18})$$

which naturally gives:

$$(D^{-1})_{ba}^{\beta\alpha}(m, n) S_0[n_0, \alpha_0, a_0]_a^\alpha(n) = \phi[n_0, \alpha_0, a_0]_b^\beta(m). \quad (\text{A.19})$$

However, D^{-1} is not known explicitly. Therefore, multiplying (A.19) with D yields:

$$\begin{aligned} (D)_{cb}^{\gamma\beta}(l, m) \phi[n_0, \alpha_0, a_0]_b^\beta(m) &= (D)_{cb}^{\gamma\beta}(l, m) (D^{-1})_{ba}^{\beta\alpha}(n, m) S_0[n_0, \alpha_0, a_0]_a^\alpha(n) \\ &= (D)_{cb}^{\gamma\beta}(l, m) (D^{-1})_{ba}^{\beta\alpha}(n, m) \delta(n - n_0) \delta(\alpha - \alpha_0) \delta(a - a_0) \\ &= (D)_{cb}^{\gamma\beta}(l, m) (D^{-1})_{ba_0}^{\beta\alpha_0}(y, n_0) \\ &= \delta(l - n_0) \delta(\gamma - \alpha_0) \delta(c - a_0) \\ (M\phi[n_0, \alpha_0, a_0])_c^\gamma(l) &= S_0[n_0, \alpha_0, a_0]_c^\gamma(l) \\ &\equiv M\phi[n_0, \alpha_0, a_0] = S_0[n_0, \alpha_0, a_0]. \end{aligned} \quad (\text{A.20})$$

Thus, solving the linear equation $M\phi = S_0[x_0, \alpha_0, A_0]$ yields the point-to-all propagator. S_0 is called a *point source*. It can be performed using Krylov-space based algorithms, see Section 3.1.3.

One important remark has to be made: One is often not interested in one particular flavour or spin particle in the propagator, but only in the space time dependency. Thus, one does not need to have not one column, but $N_c * N_{\text{Dirac}} = 12$ columns. Therefore, calculating the “full” point-to-all propagator $\phi[n_0]$ consists of 12 inversions of the fermion matrix with 12 different point sources, one for each spin-flavour combination at n_0 .

Point sources constitute the simplest source type possible. It has several disadvantages, as it provides only the minimal amount of information extracted from a given gauge configuration. Also, translational invariance on the lattice is lost, which can be used efficiently if one is interested for example in a correlator at fixed time-like distances only. Yet, the extracted point-to-all propagator is exactly calculated. To extract more information, one would have to add more point sources, which is very expensive numerically. Therefore, it is more convenient to rely on stochastic methods.

This can be done for example by *noisy estimators* [Boucaud et al., 2008]. These are stochastic sources fulfilling¹

$$\delta_{n,m} \delta_{a,b} \delta_{\alpha,\beta} = \lim_{r \rightarrow \infty} \frac{1}{r} \sum_r \xi_{a,\alpha}^r(n) (\xi_{b,\beta}^r(m))^\dagger. \quad (\text{A.21})$$

r denotes the number of sources. This allows for the estimation of the propagator as

$$D^{-1}(n, m)_{a,b,\alpha,\beta} = \lim_{r \rightarrow \infty} \frac{1}{r} \sum_r \phi_{a,\alpha}^r(n) (\xi_{b,\beta}^r(m))^\dagger, \quad (\text{A.22})$$

with $\phi^r = D^{-1} \xi^r$.

The approximation using noisy estimators is strictly valid for an infinite number of sources only. Therefore, the quantities shown here will be afflicted with two sources of noise in an actual simulation, from the statistics of the gauge ensemble and from the use of noisy estimators.

¹ The lattice spacing a will be suppressed in this section.

A.6 CL²QCD DETAILS

A.6.1 *Compiling CL²QCD*

In this section, an example of the installation process of CL²QCD will be given. The CL²QCD development makes use of the version control system `git`.² For the configuration of the system, `cmake`³ is used. It is supposed that access to the CL²QCD git repository is granted and that the required libraries are provided on the installation system (see `INSTALL` within the git repository):

```

1 # clone the git repository
2 > cd <CL2QCD_INSTALL_DIR>
3 > git clone <PATH_TO_GIT_REPOSITORY>
4 # make a build directory
5 > cd <CL2QCD_INSTALL_DIR>/cl2qcd
6 > mkdir build
7 > cd build/
8 # run cmake
9 # if not found automatically, cmake variables can be set
10 # in the command line as "-D<CMAKE_VARIABLE>=<VALUE>"
11 > cmake ..
12 # Build all executables:
13 > make -j

```

To see if the installation gives correct results, run regression tests via `make test`. In case there are multiple GPUs in the system, use `clinfo` to show the available OpenCL devices. `export GPU_DEVICE_ORDINAL=X` selects the device with number `X` and CL²QCD will use only this one.

A.6.2 *Memory Requirements*

As was stated above, LQCD simulations are usually limited by the memory available on a system. This is particularly true for GPUs, which generally provide much less memory compared to a standard CPU system. The requirements can be estimated by counting the number of fields of various types the application needs. Some of these are given in Table 14. For example, an HMC algorithm with even-odd preconditioning and a CG solver needs two gauge fields, two gauge momentum objects, one force object and 5 spinor field objects. This accumulates to 2400 Bytes per lattice site. For a GPU equipped with 3 GByte of main memory, this would mean an HMC could be run on a lattice with each extent approximately 32 sites.

² <http://git-scm.com/>

³ <http://www.cmake.org/>

	general size [V_{tot}]	Bytes [V_{tot}]
ϕ	$N_{\text{Dirac}} \times N_c \times C$	192
ϕ_{eo}	$(N_{\text{Dirac}} \times N_c \times C)/2$	96
U	$N_c^2 \times N_D \times C$	576
U_{REC12}	$2N_c \times N_D \times C$	384
U_{REC10}	$(2N_c - 1) \times N_D \times C$	320
U_{REC8}	$(2N_c - 2) \times N_D \times C$	256
P, F	$(N_c^2 - 1) \times N_D$	256

Table 14: Overview over memory requirements of LQCD quantities in double precision per site (and direction). C denotes the size of one complex number (2 real numbers).

A.6.3 Even-odd Mapping

In this section, methods which can be used to map among super indices and lattice coordinates are presented. If the induced computational overhead is negligible, this provides an alternative to simply storing all neighbour indices for each super index in memory. As LQCD calculations are always memory bound, these methods are used in CL^2 QCD. The coordinates of site n in the specific directions will be denoted as x, y, z, t . It is assumed that all spatial directions have the same extent N_σ .

```

1 t = [ n / NS^3 ]
2 acc = t
3 z = [ n / NS^2 ] - NS * acc
4 acc = NS * acc + z
5 y = [ n / NS ] - NS * acc
6 acc = NS * acc + y
7 x = n - NS * acc

```

(a)

x/y	0	1	2	3
0	e	o	e	o
1	o	e	o	e
2	e	o	e	o
3	o	e	o	e

(b)

Figure 63: a) Sketch of dissection of super index n into its components. [...] denotes the floor-function or integer-division. b) Sketch of EO preconditioning applied to a $(N = 4)^2$ lattice. Directions are labelled x and y . e and o denote even and odd sites, respectively.

On the lattice, all sites can be labelled by one super index n , which can be defined as

$$n = x + y * N_\sigma + z * N_\sigma^2 + t * N_\sigma^3. \quad (\text{A.23})$$

The super index is also used to loop over all lattice sites. Thus, to calculate a neighbour in direction $\hat{\mu}$ of a given site n , n needs to be dissected into its components again, see Figure 63a. The super index of the neighbour can be calculated straightforwardly.

If possible, EO preconditioning is used to quicken the matrix inversion (see Section 3.1.3). Essentially, it means that one lattice extent is halved. The super index runs n is then mapped to two new super indices n_{even} and n_{odd} , both running from $0, \dots, V_{\text{tot}}/2 - 1$. For instance, this can be done simply by $n\%2$ and the parity of n . $n_{\text{even/odd}}$ can be dissected into its components $\tilde{x}, \tilde{y}, \tilde{z}, \tilde{t}$ similar to n , except that one index runs over half the lattice extent only. The mapping is depicted in Figure 63b as a two dimensional example. The problem is now that, given an EO super index \tilde{n} , it is not trivial to get the corresponding super index of a neighbouring site. In fact, assuming that $n = x + N * y$, the even sites in the figure are $\{0, 2, 5, 7, 8, 10, 13, 15\}$ and the odd sites $\{1, 3, 4, 6, 9, 11, 12, 14\}$. This pattern can be described via the functions

$$f_{\text{even}}(\tilde{x}, \tilde{y}) = N * 2 * \tilde{x} + 2 * \tilde{y} + [2 * \tilde{y} / N] \quad (\text{A.24})$$

$$f_{\text{odd}}(\tilde{x}, \tilde{y}) = N * 2 * \tilde{x} + 2 * \tilde{y} - [2 * \tilde{y} / N] + 1. \quad (\text{A.25})$$

In here, $[...]$ denotes the floor-function. These functions reproduce the observed even-odd pattern in the super index n . This can be extended to a four dimensional lattice straightforwardly: The pattern in Figure 63b must be shifted if additional dimensions are considered such that an even site only has odd neighbours. This can be achieved by taking f_{even} or f_{odd} depending on the parity of the additional coordinates \tilde{z} and \tilde{t} . An even or odd super index can then be calculated as:

$$n_{\text{even}}(\tilde{x}, \tilde{y}, \tilde{z}, \tilde{t}) = ((\tilde{z} + \tilde{t})\%2) * f_{\text{odd}}(\tilde{x}, \tilde{y}) + ((\tilde{z} + \tilde{t} + 1)\%2) * f_{\text{even}}(\tilde{x}, \tilde{y}) \quad (\text{A.26})$$

and

$$n_{\text{odd}}(\tilde{x}, \tilde{y}, \tilde{z}, \tilde{t}) = ((\tilde{z} + \tilde{t} + 1)\%2) * f_{\text{odd}}(\tilde{x}, \tilde{y}) + ((\tilde{z} + \tilde{t})\%2) * f_{\text{even}}(\tilde{x}, \tilde{y}). \quad (\text{A.27})$$

From this, neighbours can be obtained in the way described above.

A.7 CLUSTERS & MACHINES

The development and benchmarking of CL^2QCD took place mainly on development systems provided by the *Department of High Performance Computer Architecture* of Goethe University Frankfurt⁴. Production runs were performed on three clusters in Frankfurt and Jülich. Details on their specifications will be given in the following and in Table 15.

cluster	Total CPU cores	Total GPUs	Peak performance [PFLOPS]
LOEWE-CSC	20928	778	0.599 (DP) / 2.1 (SP)
SANAM	4468	1216	1.113 (DP) / 4.296 (SP)
JUQUEEN	458752	/	5.9 (DP)

Table 15: Specifications of used clusters.

⁴ <http://compeng.uni-frankfurt.de/>

LOEWE-CSC

The LOEWE-CSC supercomputer⁵ was developed as a general purpose machine at Goethe University Frankfurt [Bach et al., 2011]. It consists of 786 nodes equipped with an *Advanced Micro Devices* (AMD) Radeon HD 5870 and two AMD Opteron 6172 each. In addition, it provides 40 pure CPU nodes. In November 2010, it was ranked 22nd in the TOP500⁶ and 8th in the Green500 list⁷, which lists the TOP500 clusters in terms of energy efficiency.

SANAM

The SANAM supercomputer [Kalcher et al., 2013] was build as an international collaboration of the *Frankfurt Institute of Advanced Studies* (FIAS)⁸, *King Abdulaziz City for Science and Technology*⁹ and *GSI Helmholtzzentrum für Schwerionenforschung*¹⁰. It consists of 304 nodes, each equipped with two AMD FirePro S10000 (which itself is a double GPU, i.e. effectively one has four GPUs per node) and two Intel Xeon E5-2650 CPUs. It was ranked 52nd in the TOP500¹¹ and 2nd in the Green500¹² in November 2012.

JUQUEEN

The JUQUEEN supercomputer¹³ was installed in the *Jülich Supercomputing Centre* in 2012. It is based on the 16 core BLUE-GENE-Q machine by the *International Business Machines Corporation* (IBM) and provides currently 28672 compute nodes. It was ranked seventh in the TOP500 list of June 2013¹⁴. In contrast to the two clusters at Frankfurt, the JUQUEEN is a pure CPU cluster, developed for highly parallelised applications which utilise thousands of cores.

5 <http://csc.uni-frankfurt.de/?51>

6 <http://www.top500.org/list/2010/11/>

7 <http://www.green500.org/list/2010/11/>

8 <http://fias.uni-frankfurt.de/>

9 <http://www.kacst.edu.sa>

10 <http://www.gsi.de/>

11 <http://www.top500.org/list/2012/11/>

12 <http://www.green500.org/list/2012/11/>

13 http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JUQUEEN/JUQUEEN_node.html

14 <http://www.top500.org/list/2012/06/>

A.8 DETAILS OF SIMULATIONS WITH TWISTED MASS FERMIONS

An overview of parameter sets used in the Z12 simulations described in Chapter 4 is given in Table 8. Results for the plaquette and $\langle \bar{\psi}\psi \rangle$ are given in Tables 16 and 17, respectively. The effective masses extracted from the different correlators are given in Table 18.

β	plaquette	$\chi(\text{plaquette})$	acc-rate
3.7800	0.566534(47)	2.90(32)e-07	50.0
3.8025	0.568784(50)	3.97(36)e-07	59.5
3.8100	0.569957(45)	3.27(26)e-07	59.7
3.8175	0.570876(39)	2.58(24)e-07	61.6
3.8210	0.571609(43)	2.54(20)e-07	54.4
3.8250	0.572061(40)	2.73(24)e-07	62.5
3.8300	0.573064(38)	2.17(19)e-07	68.8
3.8350	0.573441(28)	1.97(15)e-07	77.1
3.8425	0.574512(27)	1.71(12)e-07	65.8
3.8500	0.575716(30)	1.83(13)e-07	47.3
3.8575	0.576695(26)	1.53(10)e-07	47.6
3.8650	0.577782(23)	1.35(08)e-07	70.8
3.8750	0.579161(25)	1.34(09)e-07	49.3
3.8825	0.580043(17)	8.07(42)e-08	85.4
3.8900	0.581120(18)	8.30(46)e-08	82.5
3.9100	0.583899(20)	8.71(62)e-08	59.6
3.9175	0.584939(20)	8.77(62)e-08	57.3
3.9250	0.585891(17)	7.17(47)e-08	52.5
3.9400	0.587793(12)	5.92(29)e-08	62.3
3.9550	0.589779(15)	6.24(36)e-08	63.6
3.9650	0.591033(12)	5.54(28)e-08	64.7

Table 16: Mean plaquette value, its susceptibility and acceptance rate for the Z12 runs. The errors are statistical.

β	$\langle\bar{\psi}\psi\rangle$	$\chi(\langle\bar{\psi}\psi\rangle)$	statistics
3.7800	0.00229(18)	2.920(354)e-08	502
3.8025	0.00206(11)	2.306(92)e-08	3969
3.8100	0.00203(09)	2.110(97)e-08	3661
3.8175	0.00195(09)	1.977(69)e-08	4017
3.8210	0.00196(10)	1.714(115)e-08	1090
3.8250	0.00190(07)	1.977(68)e-08	3747
3.8300	0.00192(08)	1.638(104)e-08	1122
3.8350	0.00183(06)	1.699(47)e-08	4271
3.8425	0.00178(05)	1.612(52)e-08	4321
3.8500	0.00173(06)	1.598(58)e-08	4323
3.8575	0.00168(07)	1.559(56)e-08	4651
3.8650	0.00164(05)	1.420(45)e-08	3306
3.8750	0.00155(06)	1.363(67)e-08	3833
3.8825	0.00151(09)	1.265(97)e-08	439
3.8900	0.00147(09)	1.147(103)e-08	338
3.9100	0.00134(10)	1.153(102)e-08	721
3.9175	0.00129(12)	1.278(140)e-08	689
3.9250	0.00125(10)	0.973(105)e-08	606
3.9400	0.00114(05)	0.488(40)e-08	481
3.9550	0.00108(10)	0.659(94)e-08	451
3.9650	0.00101(07)	0.346(87)e-08	449

Table 17: Mean $\langle\bar{\psi}\psi\rangle$, its susceptibility and number of data points for the Z12 runs. The errors are statistical. Note that $\chi(\langle\bar{\psi}\psi\rangle)$ has to be normalised by a factor of V/T to match (4.9).

β	$m_{\text{eff}}^{\text{ps}}$	χ^2	$m_{\text{eff}}^{\text{sc}}$	χ^2	$m_{\text{eff}}^{\text{av}}$	χ^2	$m_{\text{eff}}^{\text{ve}}$	χ^2	statistics
3.8025	0.1529(4)	3.59	0.514(45)	0.23	0.7517(253)	0.02	0.4697(64)	3.16	3976
3.8100	0.1503(4)	1.96	0.634(39)	0.34	0.8108(391)	0.04	0.4356(112)	1.79	3668
3.8175	0.1520(4)	0.62	0.674(88)	0.12	0.8438(182)	0.17	0.4427(100)	0.83	4067
3.8250	0.1507(3)	1.45	0.701(36)	0.03	0.7456(210)	0.42	0.4682(58)	0.76	3755
3.8350	0.1492(4)	2.72	0.645(69)	0.49	0.5558(1469)	0.13	0.4382(89)	0.82	4274
3.8425	0.1478(3)	11.60	0.660(25)	0.13	0.7629(173)	0.62	0.4178(125)	0.02	4435
3.8500	0.1471(3)	12.67	0.303(70)	0.14	0.5780(784)	0.31	0.4453(131)	0.59	4327
3.8575	0.1462(3)	6.24	0.672(29)	0.13	0.7447(155)	0.04	0.4196(95)	3.01	4655
3.8650	0.1458(5)	3.38	0.210(41)	0.42	0.7415(232)	0.38	0.4244(108)	0.77	3312
3.8750	0.1499(4)	2.60	0.404(95)	0.25	0.7500(162)	1.57	0.4518(84)	2.80	3838

Table 18: Values for the effective masses extracted from the correlators of the Z12 ensemble.

A.9 DETAILS OF SIMULATIONS WITH WILSON FERMIONS

In this section, details about the results carried out in the setup described in Chapter 5 will be given.

An overview about the simulated systems can be seen in Table 19. Analysis details are given in Tables 20,22,23 and 21. The results of the fits of the Binder cumulant to (5.9) are given in Table 24.

κ	β -range	$N_\sigma=8$	$N_\sigma=10$	$N_\sigma=12$	$N_\sigma=14$	$N_\sigma=16$	$N_\sigma=20$
0.0300	5.685-5.696	-	40k	40k	40k	40k	-
0.0400	5.685-5.695	-	40k	40k	40k	40k	-
0.0500	5.683-5.695	-	40k	40k	40k	40k	-
0.0600	5.681-5.695	40k	40k	40k	40k	40k	(40k)
0.0650	5.676-5.689	40k	40k	80k	-	40k	-
0.0700	5.676-5.688	40k	-	40k	-	60k	80k
0.0865	5.662-5.678	40k	-	40k	-	60k	(80k)
0.0910	5.659-5.673	40k	-	40k	-	40k	80k
0.1000	5.647-5.658	40k	-	40k	-	40k	80k
0.1040	5.640-5.655	40k	-	40k	-	40k	-
0.1050	5.638-5.650	40k	-	40k	-	40k	-
0.1060	5.638-5.650	40k	-	40k	-	40k	-
0.1100	5.629-5.640	40k	-	40k	-	40k	80k
0.1200	5.602-5.613	40k	-	40k	-	40k	80k
0.1300	5.562-5.578	40k	-	40k	-	40k	-
0.1400	5.508-5.520	40k	-	40k	-	40k	-
0.1450	5.474-5.485	40k	-	40k	40k	40k	-
0.1500	5.431-5.441	40k	-	40k	-	40k	80k
0.1525	5.407-5.417	40k	40k	40k	40k	40k	-
0.1550	5.380-5.389	40k	40k	40k	40k	40k	-
0.1575	5.350-5.361	40k	40k	40k	40k	40k	-
0.1600	5.319-5.330	40k	40k	40k	40k	40k	-
0.1625	5.284-5.294	40k	40k	40k	40k	40k	-
0.1650	5.246-5.256	40k	40k	40k	40k	40k	-

Table 19: Overview of simulations carried out at $\mu_l = i\pi T$ and $N_\tau = 4$. The numbers given denote the statistics produced on each β point. A given β -range was scanned with $\Delta\beta = 0.001$ for each N_σ . Numbers in brackets indicate that some β values have smaller statistics.

κ	β -range	$N_{\sigma=8}$	$N_{\sigma=10}$	$N_{\sigma=12}$	$N_{\sigma=14}$	$N_{\sigma=16}$	$N_{\sigma=20}$
0.0300	5.685-5.696	-	70	70	60	50	-
0.0400	5.685-5.695	-	70	50	40	50	-
0.0500	5.683-5.695	-	60	50	50	40	-
0.0600	5.681-5.695	70	70	50	40	60	(90)
0.0650	5.676-5.689	70	50	60	-	50	-
0.0700	5.676-5.688	70	-	50	-	50	75
0.0865	5.662-5.678	90	-	90	-	60	(100)
0.0910	5.659-5.673	60	-	60	-	50	75
0.1000	5.647-5.658	90	-	90	-	70	100
0.1040	5.640-5.655	90	-	90	-	50	-
0.1050	5.638-5.650	90	-	90	-	70	-
0.1060	5.638-5.650	90	-	90	-	60	-
0.1100	5.629-5.640	90	-	60	-	60	100
0.1200	5.602-5.613	70	-	70	-	50	100
0.1300	5.562-5.578	90	-	60	-	60	-
0.1400	5.508-5.520	90	-	70	-	60	-
0.1450	5.474-5.485	90	-	70	60	60	-
0.1500	5.431-5.441	60	-	50	-	60	75
0.1525	5.407-5.417	70	70	60	60	60	-
0.1550	5.380-5.389	60	70	50	50	60	-
0.1575	5.350-5.361	70	70	60	50	60	-
0.1600	5.319-5.330	70	70	60	60	60	-
0.1625	5.284-5.294	60	60	60	50	60	-
0.1650	5.246-5.256	70	60	50	50	50	-

Table 20: Overview of the number of blocks used in the jackknife analysis of the simulations in Table 19.

κ	β -range	$N_\sigma=8$	$N_\sigma=10$	$N_\sigma=12$	$N_\sigma=14$	$N_\sigma=16$	$N_\sigma=20$
0.0300	5.685-5.696	-	5.6903(6)	5.6905(4)	5.6906(3)	5.6915(2)	-
0.0400	5.685-5.695	-	5.6903(4)	5.6910(4)	5.6906(3)	5.6904(2)	-
0.0500	5.683-5.695	-	5.6892(4)	5.6898(3)	5.6904(3)	5.6897(2)	-
0.0600	5.681-5.695	5.6878(4)	5.6873(5)	5.6883(4)	5.6881(3)	5.6873(2)	-
0.0650	5.676-5.689	5.6869(6)	5.6857(4)	5.6856(2)	-	5.6855(2)	-
0.0700	5.676-5.688	5.6856(5)	-	5.6844(4)	-	5.6833(2)	5.6832(2)
0.0865	5.662-5.678	5.6745(5)	-	5.6723(3)	-	5.6714(2)	5.6716(1)
0.0910	5.659-5.673	5.6705(5)	-	5.6673(4)	-	5.6668(2)	5.6666(2)
0.1000	5.647-5.658	5.6565(6)	-	5.6550(3)	-	5.6545(3)	5.6547(1)
0.1040	5.640-5.655	5.6498(6)	-	5.6489(3)	-	5.6481(3)	-
0.1050	5.638-5.650	5.6488(5)	-	5.6469(3)	-	5.6460(3)	-
0.1060	5.638-5.650	5.6471(5)	-	5.6448(3)	-	5.6444(3)	-
0.1100	5.629-5.640	5.6385(6)	-	5.6365(4)	-	5.6360(3)	5.6360(2)
0.1200	5.602-5.613	5.6113(5)	-	5.6104(3)	-	5.6093(2)	5.6092(2)
0.1300	5.562-5.578	5.5743(5)	-	5.5725(2)	-	5.5715(2)	-
0.1400	5.508-5.520	5.5193(4)	-	5.5182(3)	-	5.5166(3)	-
0.1450	5.474-5.485	5.4828(4)	-	5.4817(3)	5.4812(2)	5.4807(2)	-
0.1500	5.431-5.441	5.4384(4)	-	5.4377(3)	-	5.4376(2)	5.4373(1)
0.1525	5.407-5.417	5.4137(3)	5.4131(3)	5.4129(2)	5.4127(2)	5.4124(2)	-
0.1550	5.380-5.389	5.3854(3)	5.3852(3)	5.3855(3)	5.3852(2)	5.3851(1)	-
0.1575	5.350-5.361	5.3560(3)	5.3555(3)	5.3555(2)	5.3555(1)	5.3554(2)	-
0.1600	5.319-5.330	5.3230(3)	5.3237(2)	5.3227(2)	5.3229(1)	5.3231(1)	-
0.1625	5.284-5.294	5.2876(3)	5.2877(2)	5.2880(2)	5.2881(1)	5.2880(1)	-
0.1650	5.246-5.256	5.2493(3)	5.2510(2)	5.2509(2)	5.2504(1)	5.2510(1)	-

Table 21: Overview of β_c for the deconfinement transition obtained from the peak of the susceptibility of $|L|$ from the simulations in Table 19.

κ	β -range	$N_\sigma=8$	$N_\sigma=10$	$N_\sigma=12$	$N_\sigma=14$	$N_\sigma=16$	$N_\sigma=20$
0.0300	5.685-5.696	-	-	-	-	-	-
0.0400	5.685-5.695	-	-	-	-	5.6920(4)	-
0.0500	5.683-5.695	-	-	-	5.6928(6)	5.6914(5)	-
0.0600	5.681-5.695	-	-	5.6897(5)	5.6888(3)	5.6873(3)	-
0.0650	5.676-5.689	-	-	5.6862(3)	-	5.6854(2)	-
0.0700	5.676-5.688	-	-	5.6847(4)	-	5.6832(2)	5.6831(2)
0.0865	5.662-5.678	5.6765(7)	-	5.6720(3)	-	5.6712(2)	5.6715(1)
0.0910	5.659-5.673	5.6707(6)	-	5.6671(4)	-	5.6666(2)	5.6665(2)
0.1000	5.647-5.658	5.6561(6)	-	5.6546(3)	-	5.6541(3)	5.6543(2)
0.1040	5.640-5.655	5.6493(6)	-	5.6484(3)	-	5.6476(2)	-
0.1050	5.638-5.650	5.6483(5)	-	5.6464(3)	-	5.6455(2)	-
0.1060	5.638-5.650	5.6465(5)	-	5.6442(3)	-	5.6439(3)	-
0.1100	5.629-5.640	5.6378(6)	-	5.6358(4)	-	5.6354(3)	5.6354(2)
0.1200	5.602-5.613	5.6106(5)	-	5.6094(3)	-	5.6085(3)	5.6085(2)
0.1300	5.562-5.578	5.5734(5)	-	5.5713(2)	-	5.5706(2)	-
0.1400	5.508-5.520	5.5181(4)	-	5.5169(3)	-	5.5155(3)	-
0.1450	5.474-5.485	5.4815(4)	-	5.4806(3)	5.4803(2)	5.4799(2)	-
0.1500	5.431-5.441	5.4372(4)	-	5.4367(2)	-	5.4369(1)	5.4368(1)
0.1525	5.407-5.417	5.4127(3)	5.4122(3)	5.4121(2)	5.4120(2)	5.4118(2)	-
0.1550	5.380-5.389	5.3845(3)	5.3844(3)	5.3848(3)	5.3847(2)	5.3847(1)	-
0.1575	5.350-5.361	5.3553(3)	5.3549(3)	5.3550(2)	5.3551(2)	5.3551(2)	-
0.1600	5.319-5.330	5.3224(3)	5.3232(3)	5.3224(2)	5.3226(1)	5.3229(2)	-
0.1625	5.284-5.294	5.2872(3)	5.2874(2)	5.2878(2)	5.2880(1)	5.2879(1)	-
0.1650	5.246-5.256	5.2490(3)	5.2509(2)	5.2508(2)	5.2504(1)	5.2510(1)	-

Table 22: Overview of β_c for the deconfinement transition obtained from the peak of the susceptibility of $|L_m|$ from the simulations in Table 19.

κ	β -range	$N_\sigma=8$	$N_\sigma=10$	$N_\sigma=12$	$N_\sigma=14$	$N_\sigma=16$	$N_\sigma=20$
0.1525	5.407-5.417	5.4120(6)	5.4132(3)	5.4128(2)	5.4128(2)	5.4126(1)	-
0.1550	5.380-5.389	-	5.3851(4)	-	5.3852(2)	-	-
0.1575	5.350-5.361	5.3552(5)	5.3553(3)	5.3554(2)	5.3555(1)	5.3553(1)	-
0.1600	5.319-5.330	5.3220(4)	5.3229(3)	5.3226(2)	5.3229(1)	5.3231(2)	-
0.1625	5.284-5.294	5.2867(4)	5.2875(2)	5.2879(2)	5.2881(1)	5.2882(2)	-
0.1650	5.246-5.256	5.2488(3)	5.2510(2)	5.2508(2)	5.2504(1)	5.2510(1)	-

Table 23: Overview of β_c for the deconfinement transition obtained from the peak of the susceptibility of $\langle \bar{\psi}\psi \rangle$ from the simulations in Table 19.

κ	N_σ	β_c	ν	$B_4(\beta, \infty)$	a_1	a_2	χ^2
0.0300	12,14,16	5.6921(3)	0.289(57)	2.37(4)	-0.016(29)	-	0.968
0.0400	12,14,16	5.6883(2)	0.305(70)	2.63(4)	-0.036(73)	-	0.983
0.0500	10,12,14,16	5.6891(1)	0.351(34)	2.18(2)	-0.089(66)	-	0.965
0.0600	10,12,14,16	5.6862(1)	0.369(14)	2.04(1)	-0.126(50)	0.005(3)	0.773
0.0650	10,12,16	5.6844(1)	0.375(13)	1.92(1)	-0.13(30)	0.007(3)	1.408
0.0700	12,16,20	5.6829(1)	0.423(39)	1.86(2)	-0.28(18)	-	1.029
0.0865	8,12,16	5.6704(1)	0.450(18)	1.89(1)	-0.378(09)	-	1.026
0.0910	8,12,16,20	5.6655(1)	0.458(22)	1.85(1)	-0.38(11)	0.062(34)	1.173
0.1000	8,12,16,20	5.6539(1)	0.501(19)	1.74(1)	-0.56(12)	-	0.952
0.1040	8,12,16	5.6469(1)	0.547(17)	1.77(1)	-0.79(11)	-	0.991
0.1050	8,12,16	5.6438(1)	0.650(23)	1.85(1)	-1.52(21)	-	1.019
0.1060	8,12,16	5.6425(1)	0.589(23)	1.82(1)	-1.10(18)	-	1.015
0.1100	8,12,16, 20	5.6341(1)	0.582(43)	1.80(1)	-1.08(38)	-	1.064
0.1200	12,16,20	5.6075(1)	0.598(21)	1.75(1)	-1.167(20)	0.40(14)	0.996
0.1300	8, 12, 16	5.5689(1)	0.637(22)	1.83(1)	-1.64(22)	-	0.860
0.1400	8, 12, 16	5.5146(1)	0.612(20)	1.83(1)	-1.59(21)	-	0.821
0.1450	8, 12, 16	5.4790(1)	0.588(23)	1.80(1)	-1.50(26)	-	1.003
0.1500	12, 16, 20	5.4367(1)	0.611(42)	1.66(2)	-2.12(69)	-	0.950
0.1525	10, 12, 14, 16	5.4114(1)	0.620(39)	1.76(2)	-2.51(66)	-	1.029
0.1550	10, 12, 14, 16	5.3849(1)	0.512(37)	1.67(1)	-1.26(46)	-	0.984
0.1575	12, 14, 16	5.3548(1)	0.555(28)	1.80(2)	-2.42(60)	2.15(1.05)	1.015
0.1600	8, 10, 12	5.3225(1)	0.376(37)	1.77(2)	-0.43(26)	-	0.997
0.1625	10, 12, 14	5.2886(1)	0.331(18)	1.58(1)	-0.16(07)	0.015(12)	0.983
0.1650	10, 12, 14	5.2501(1)	0.364(47)	2.15(6)	-0.67(61)	0.11(20)	0.981

Table 24: Overview of fits to $B_4(L_{\text{Im}})$ according to (5.9). The N_σ column indicates which datasets have been used in the fit, see also Table 19. If not value for a_2 is given, the fit has been performed with the linear ansatz.

BIBLIOGRAPHY

- Aad, G. et al. (2012). "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". *Phys.Lett.* B716, pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv:1207.7214 [hep-ex] (cited on p. 1).
- Abdel-Rehim, A., Boucaud, P., Carrasco, N., Deuzeman, A., Dimopoulos, P., et al. (2013a). "A first look at maximally twisted mass lattice QCD calculations at the physical point". *PoS, LATTICE2013*, p. 264. arXiv:1311.4522 [hep-lat] (cited on p. 105).
- Abdel-Rehim, A., Burger, F., Deuzeman, A., Jansen, K., Kostrzewa, B., et al. (2013b). "Recent developments in the tmLQCD software suite". arXiv:1311.5495 [hep-lat] (cited on p. 37).
- Alexandru, A., Pelissier, C., Gamari, B., and Lee, F. (2012). "Multi-mass solvers for lattice QCD on GPUs". *J.Comput.Phys.* 231, pp. 1866–1878. arXiv:1103.5103 [hep-lat] (cited on pp. 51, 53).
- Alexandru, A. and Li, A. (2013). "QCD at imaginary chemical potential with Wilson fermions". *PoS, LATTICE2013*, p. 208. arXiv:1312.1201 [hep-lat] (cited on p. 86).
- Alles, B., D'Elia, M., and Lombardo, M. (2007). "Topology, chiral and screening transitions at finite density in two colour QCD". *Nucl.Phys.Proc.Suppl.* 174, pp. 225–228. DOI: 10.1016/j.nuclphysbps.2007.08.130. arXiv:hep-lat/0609017 [hep-lat] (cited on p. 34).
- AMD (2013). *AMD Accelerated Parallel Processing OpenCL Programming Guide*. <http://developer.amd.com/sdks/AMDAPPSDK/documentation> (cited on p. 41).
- Amdahl, G. M. (1967). "Validity of the single processor approach to achieving large scale computing capabilities". *Proceedings of the April 18-20, 1967, spring joint computer conference on - AFIPS '67 (Spring)*. DOI: 10.1145/1465482.1465560 (cited on p. 36).
- Aoki, Y., Endrodi, G., Fodor, Z., Katz, S., and Szabo, K. (2006). "The Order of the quantum chromodynamics transition predicted by the standard model of particle physics". *Nature*, 443, pp. 675–678. DOI: 10.1038/nature05120. arXiv:hep-lat/0611014 [hep-lat] (cited on pp. 3, 26, 56).
- Babich, R. et al. (2011). "Scaling Lattice QCD beyond 100 GPUs". arXiv:1109.2935 [hep-lat] (cited on p. 40).
- Babich, R., Clark, M. A., and Joo, B. (2010). "Parallelizing the QUDA Library for Multi-GPU Calculations in Lattice Quantum Chromodynamics". arXiv:1011.0024 [hep-lat] (cited on p. 40).

- Bach, M., Kretz, M., Lindenstruth, V., and Rohr, D. (2011). “Optimized HPL for AMD GPU and multi-core CPU usage”. *Computer Science - Research and Development*, pp. 1–12–12. DOI: 10.1007/s00450-011-0161-5 (cited on pp. 5, 39, 115).
- Bach, M., Lindenstruth, V., Philipsen, O., and Pinke, C. (2013a). “Lattice QCD based on OpenCL”. *Computer Physics Communications*, p. 19. DOI: 10.1016/j.cpc.2013.03.020. arXiv:1209.5942 (cited on pp. 29, 41, 43, 44, 47, 48, 50–53, 103).
- Bach, M., Philipsen, O., and Pinke, C. (2013b). “Twisted-Mass Lattice QCD using OpenCL - an Update”. In: *The XXXI International Symposium on Lattice Field Theory, Lattice2013* (cited on pp. 39, 43, 47–51, 104).
- Baron, R. et al. (2010). “Light Meson Physics from Maximally Twisted Mass Lattice QCD”. *JHEP*, 1008, p. 097. DOI: 10.1007/JHEP08(2010)097. arXiv:0911.5061 [hep-lat] (cited on pp. 51, 60, 61, 82).
- Beringer, J. et al. (2012). “Review of Particle Physics (RPP)”. *Phys.Rev.* D86, p. 010001. DOI: 10.1103/PhysRevD.86.010001 (cited on p. 1).
- Bernard, C. W., Blum, T., Detar, C. E., Gottlieb, S. A., Heller, U. M., et al. (1997). “Which chiral symmetry is restored in high temperature QCD?” *Phys.Rev.Lett.* 78, pp. 598–601. DOI: 10.1103/PhysRevLett.78.598. arXiv:hep-lat/9611031 [hep-lat] (cited on pp. 56, 64).
- Bernard, V. and Meissner, U.-G. (2007). “Chiral perturbation theory”. *Ann.Rev.Nucl.Part.Sci.* 57, pp. 33–60. DOI: 10.1146/annurev.nucl.56.080805.140449. arXiv:hep-ph/0611231 [hep-ph] (cited on p. 11).
- Bhanot, G., Chen, D., Gara, A., Sexton, J., and Vranas, P. (2005). “QCD on the BlueGene/L Supercomputer”. *Nuclear Physics B - Proceedings Supplements*, 140, pp. 823–825. DOI: 10.1016/j.nuclphysbps.2004.11.153 (cited on p. 37).
- Binder, K. (1981). “Finite size scaling analysis of Ising model block distribution functions”. *Z.Phys.* B43, pp. 119–140. DOI: 10.1007/BF01293604 (cited on p. 94).
- Bitar, K. M., DeGrand, T. A., Edwards, R., Gottlieb, S. A., Heller, U. M., et al. (1991). “Hadron thermodynamics with Wilson quarks”. *Phys.Rev.* D43, pp. 2396–2400. DOI: 10.1103/PhysRevD.43.2396 (cited on p. 90).
- (1994). “Hadron spectrum and matrix elements in QCD with dynamical Wilson fermions at $6/g^{*2} = 5.3$ ”. *Phys.Rev.* D49, pp. 3546–3562. DOI: 10.1103/PhysRevD.49.3546. arXiv:hep-lat/9309011 [hep-lat] (cited on p. 90).
- Blossier, B. et al. (2010). “Average up/down, strange and charm quark masses with $N_f=2$ twisted mass lattice QCD”. *Phys.Rev.* D82, p. 114513. DOI: 10.1103/PhysRevD.82.114513. arXiv:1010.3659 [hep-lat] (cited on pp. 60, 61, 83).
- Bonati, C., Cossu, G., D’Elia, M., and Incardona, P. (2011a). “QCD simulations with staggered fermions on GPUs”. arXiv:1106.5673 [hep-lat] (cited on pp. 40, 52).

- Bonati, C., Cossu, G., D'Elia, M., and Di Giacomo, A. (2010). "Staggered fermions simulations on GPUs". *PoS, LATTICE2010*, p. 324. arXiv:1010.5433 [hep-lat] (cited on p. 40).
- Bonati, C., de Forcrand, P., D'Elia, M., Philipsen, O., and Sanfilippo, F. (2011b). "Constraints on the two-flavor QCD phase diagram from imaginary chemical potential". *PoS, LATTICE2011*, p. 189. arXiv:1201.2769 [hep-lat] (cited on pp. 25, 86, 88, 89).
- Bonati, C., Cossu, G., D'Elia, M., and Sanfilippo, F. (2011c). "The Roberge-Weiss endpoint in $N_f = 2$ QCD". *Phys.Rev. D83*, p. 054505. DOI: 10.1103/PhysRevD.83.054505. arXiv:1011.4515 [hep-lat] (cited on pp. 5, 86, 88–90, 102, 105).
- Bonati, C., D'Elia, M., de Forcrand, P., Philipsen, O., and Sanfillippo, F. (2013). "The chiral phase transition for two-flavour QCD at imaginary and zero chemical potential". arXiv:1311.0473 [hep-lat] (cited on pp. 5, 58, 88).
- Borsanyi, S. et al. (2010). "Is there still any T_c mystery in lattice QCD? Results with physical masses in the continuum limit III". *JHEP*, 1009, p. 073. DOI: 10.1007/JHEP09(2010)073. arXiv:1005.3508 [hep-lat] (cited on p. 26).
- Boucaud, P. et al. (2008). "Dynamical Twisted Mass Fermions with Light Quarks: Simulation and Analysis Details". *Computer Physics Communication*, 179, pp. 695–715. DOI: 10.1016/j.cpc.2008.06.013. arXiv:0803.0224 [hep-lat] (cited on pp. 20, 111).
- Boyd, G., Engels, J., Karsch, F., Laermann, E., Legeland, C., et al. (1996). "Thermodynamics of SU(3) lattice gauge theory". *Nucl.Phys. B469*, pp. 419–444. DOI: 10.1016/0550-3213(96)00170-8. arXiv:hep-lat/9602007 [hep-lat] (cited on p. 56).
- Boyle, P. (2012). "The BlueGene/Q supercomputer". In: *The 30 International Symposium on Lattice Field ...* (Cited on p. 37).
- Brandt, B. B., Francis, A., Meyer, H. B., Philipsen, O., and Wittig, H. (2013). "QCD thermodynamics with $O(a)$ improved Wilson fermions at $N_f=2$ ". arXiv:1310.8326 [hep-lat] (cited on pp. 57, 68, 82).
- Braun-Munzinger, P. and Wambach, J. (2009). "The Phase Diagram of Strongly-Interacting Matter". *Rev.Mod.Phys.* 81, pp. 1031–1050. DOI: 10.1103/RevModPhys.81.1031. arXiv:0801.4256 [hep-ph] (cited on pp. 2, 24).
- Buchhoff, M. I., Cheng, M., Christ, N. H., Ding, H. T., Jung, C., et al. (2013). "The QCD chiral transition, $U_A(1)$ symmetry and the Dirac spectrum using domain wall fermions". arXiv:1309.4149 [hep-lat] (cited on pp. 57, 82).
- Burger, F., Ilgenfritz, E.-M., Kirchner, M., Lombardo, M. P., Müller-Preussker, M., Philipsen, O., Pinke, C., and Zeidlewicz, L. (2013a). "The thermal QCD transition with two flavours of twisted mass fermions". *Phys.Rev. D87*, p. 074508. DOI: 10.1103/PhysRevD.87.074508. arXiv:1102.4530 [hep-lat] (cited on pp. 4, 55, 57–59, 61, 64, 65, 67–75, 82, 84, 104).
- Burger, F., Hotzel, G., Müller-Preussker, M., Ilgenfritz, E.-M., and Lombardo, M. P. (2013b). "Towards thermodynamics with $N_f = 2 + 1 + 1$ twisted mass quarks". arXiv:1311.1631 [hep-lat] (cited on p. 105).

- Butti, A., Pelissetto, A., and Vicari, E. (2003). "On the nature of the finite temperature transition in QCD". *JHEP*, 0308, p. 029. DOI: 10.1088/1126-6708/2003/08/029. arXiv:hep-ph/0307036 [hep-ph] (cited on p. 57).
- Cabibbo, N. and Marinari, E. (1982). "A new Method for updating SU(N) Matrices in Computer Simulations of Gauge Theories". *Physical Letters, Volume 119B, number 4,5,6* (cited on p. 32).
- Cardoso, N. and Bicudo, P. (2011). "Generating SU(Nc) pure gauge lattice QCD configurations on GPUs with CUDA and OpenMP". arXiv:1112.4533 [hep-lat] (cited on pp. 49, 53).
- Cea, P., Cosmai, L., D'Elia, M., Papa, A., and Sanfilippo, F. (2012). "The critical line of two-flavor QCD at finite isospin or baryon densities from imaginary chemical potentials". *Phys.Rev. D* 85, p. 094512. DOI: 10.1103/PhysRevD.85.094512. arXiv:1202.5700 [hep-lat] (cited on p. 86).
- Chatrchyan, S. et al. (2012). "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC". *Phys.Lett. B* 716, pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021. arXiv:1207.7235 [hep-ex] (cited on p. 1).
- Clark, M. A. (2009). "QCD on GPUs: cost effective supercomputing". *PoS, LAT2009*, p. 003. arXiv:0912.2268 [hep-lat] (cited on p. 40).
- Clark, M. A., Babich, R., Barros, K., Brower, R. C., and Rebbi, C. (2010). "Solving Lattice QCD systems of equations using mixed precision solvers on GPUs". *Comput. Phys. Commun.* 181, pp. 1517–1528. DOI: 10.1016/j.cpc.2010.05.002. arXiv:0911.3191 [hep-lat] (cited on pp. 5, 37, 40, 51, 53, 104).
- Clark, M. A. and Babich, R. (2012). "High-efficiency Lattice QCD computations on the Fermi architecture". In: *2012 Innovative Parallel Computing (InPar)*. IEEE, pp. 1–9. DOI: 10.1109/InPar.2012.6339591 (cited on pp. 40, 51).
- Creutz, M. (1980). "Monte Carlo Study of Quantized SU(2) Gauge Theory". *Phys. Rev. D* 21, pp. 2308–2315. DOI: 10.1103/PhysRevD.21.2308 (cited on p. 32).
- Creutz, M. (1998). "Evaluating Grassmann integrals". *Phys.Rev.Lett.* 81, pp. 3555–3558. DOI: 10.1103/PhysRevLett.81.3555. arXiv:hep-lat/9806037 [hep-lat] (cited on p. 29).
- (2007). "Chiral anomalies and rooted staggered fermions". *Phys.Lett. B* 649, pp. 230–234. DOI: 10.1016/j.physletb.2007.03.065. arXiv:hep-lat/0701018 [hep-lat] (cited on pp. 5, 21, 85).
- (2011). "Confinement, chiral symmetry, and the lattice". *Acta Phys.Slov.* 61, pp. 1–127. arXiv:1103.3304 [hep-lat] (cited on pp. 2, 11, 13, 14, 18, 20).
- de Forcrand, P. and Philipsen, O. (2002). "The QCD phase diagram for small densities from imaginary chemical potential". *Nucl.Phys. B* 642, pp. 290–306. DOI: 10.1016/S0550-3213(02)00626-0. arXiv:hep-lat/0205016 [hep-lat] (cited on pp. 34, 86).

- (2008). “The Chiral critical point of $N(f) = 3$ QCD at finite density to the order $(\mu/T)^{**4}$ ”. *JHEP*, 0811, p. 012. DOI: 10.1088/1126-6708/2008/11/012. arXiv:0808.1096 [hep-lat] (cited on pp. 57, 86, 88).
 - (2010). “Constraining the QCD phase diagram by tricritical lines at imaginary chemical potential”. *Phys.Rev.Lett.* 105, p. 152001. DOI: 10.1103/PhysRevLett.105.152001. arXiv:1004.3144 [hep-lat] (cited on pp. 5, 86, 91, 94, 95, 105).
- DeGrand, T. and DeTar, C. (2006). *Lattice Methods for Quantum Chromodynamics*. World Scientific Publishing Company (cited on pp. 7, 22).
- Degrand, T. and Rossi, P. (1990). “Conditioning techniques for dynamical fermions”. *Computer Physics Communications*, 60(2), pp. 211–214. DOI: 10.1016/0010-4655(90)90006-M (cited on pp. 33, 110).
- D’Elia, M. and Lombardo, M.-P. (2003). “Finite density QCD via imaginary chemical potential”. *Phys.Rev. D67*, p. 014505. DOI: 10.1103/PhysRevD.67.014505. arXiv:hep-lat/0209146 [hep-lat] (cited on p. 86).
- Doi, J. (2012). “Peta-scale Lattice Quantum Chromodynamics on a Blue Gene/Q supercomputer”. In: *2012 International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE. DOI: 10.1109/SC.2012.96 (cited on p. 37).
- Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). “Hybrid Monte Carlo”. *Phys. Lett.* B195, pp. 216–222. DOI: 10.1016/0370-2693(87)91197-X (cited on p. 31).
- Durr, S., Fodor, Z., Frison, J., Hoelbling, C., Hoffmann, R., Katz, S. D., Krieg, S., Kurth, T., Lellouch, L., and Lippert, T. (2008). “Ab-Initio Determination of Light Hadron Masses”. *Science*, 322, pp. 1224–1227. DOI: 10.1126/science.1163233. arXiv:0906.3599 [hep-lat] (cited on p. 2).
- Egri, G. I., Fodor, Z., Hoelbling, C., Katz, S. D., Nogradi, D., et al. (2007). “Lattice QCD as a video game”. *Comput.Phys.Commun.* 177, pp. 631–639. DOI: 10.1016/j.cpc.2007.06.005. arXiv:hep-lat/0611022 [hep-lat] (cited on p. 40).
- Ejiri, S., Karsch, F., Laermann, E., Miao, C., Mukherjee, S., et al. (2009). “On the magnetic equation of state in (2+1)-flavor QCD”. *Phys.Rev. D80*, p. 094505. DOI: 10.1103/PhysRevD.80.094505. arXiv:0909.5122 [hep-lat] (cited on p. 59).
- Englert, F. and Brout, R. (1964). “Broken Symmetry and the Mass of Gauge Vector Mesons”. *Phys.Rev.Lett.* 13, pp. 321–323. DOI: 10.1103/PhysRevLett.13.321 (cited on p. 1).
- Eshraim, W. I., Janowski, S., Giacosa, F., and Rischke, D. H. (2013). “Decay of the pseudoscalar glueball into scalar and pseudoscalar mesons”. *Phys.Rev. D87*, p. 054036. DOI: 10.1103/PhysRevD.87.054036. arXiv:1208.6474 [hep-ph] (cited on p. 1).
- Falcone, R., Laermann, E., and Lombardo, M. P. (2010). “Study of finite temperature QCD with 2+1 flavors via Taylor expansion and imaginary chemical potential”. *PoS, LATTICE2010*, p. 183. arXiv:1012.4694 [hep-lat] (cited on p. 86).

- Ferrenberg, A. M. and Swendsen, R. H. (1989). “Optimized Monte Carlo analysis”. *Phys.Rev.Lett.* 63, pp. 1195–1198. DOI: 10.1103/PhysRevLett.63.1195 (cited on p. 34).
- Fisher, M. and Berker, A. (1982). “Scaling for first-order transitions in thermodynamic and finite systems”. *Phys.Rev.* B26, pp. 2507–2513. DOI: 10.1103/PhysRevB.26.2507 (cited on p. 88).
- Frezzotti, R. and Rossi, G. (2004). “Chirally improving Wilson fermions. 1. O(a) improvement”. *JHEP*, 0408, p. 007. DOI: 10.1088/1126-6708/2004/08/007. arXiv:hep-lat/0306014 [hep-lat] (cited on pp. 4, 13, 14, 19).
- Fromm, M., Langelage, J., Lottini, S., and Philipsen, O. (2012). “The QCD deconfinement transition for heavy quarks and all baryon chemical potentials”. *JHEP*, 1201, p. 042. DOI: 10.1007/JHEP01(2012)042. arXiv:1111.4953 [hep-lat] (cited on pp. 90, 96, 102, 105).
- Fukugita, M., Mino, H., Okawa, M., and Ukawa, A. (1990). “Effective Z(3) Description of the SU(3) Deconfining Phase Transition”. *Nucl.Phys.Proc.Suppl.* 17, pp. 180–184. DOI: 10.1016/0920-5632(90)90233-K (cited on p. 56).
- Fukushima, K. and Hatsuda, T. (2011). “The phase diagram of dense QCD”. *Rept.Prog.Phys.* 74, p. 014001. DOI: 10.1088/0034-4885/74/1/014001. arXiv:1005.4814 [hep-ph] (cited on pp. 2, 24, 85).
- Gattringer, C. and Lang, C. (2010). *Quantum Chromodynamics on the Lattice: An Introductory Presentation*. Springer (cited on pp. 2, 7, 22, 23, 30, 35, 108–110).
- Gavin, S., Gocksch, A., and Pisarski, R. D. (1994). “QCD and the chiral critical point”. *Phys.Rev.* D49, pp. 3079–3082. DOI: 10.1103/PhysRevD.49.R3079. arXiv:hep-ph/9311350 [hep-ph] (cited on p. 56).
- Gell-Mann, M., Oakes, R., and Renner, B. (1968). “Behavior of current divergences under SU(3) × SU(3)”. *Phys.Rev.* 175, pp. 2195–2199. DOI: 10.1103/PhysRev.175.2195 (cited on p. 11).
- Georgi, H. (1982). “Lie algebras in particle physics. From Isospin to unified theories.” *Front.Phys.* 54, pp. 1–255 (cited on p. 108).
- Gerhard, J., Lindenstruth, V., and Bleicher, M. (2012). “Relativistic Hydrodynamics on Graphic Cards”, p. 9. arXiv:1206.0919 (cited on p. 47).
- Ginsparg, P. H. and Wilson, K. G. (1982). “A Remnant of Chiral Symmetry on the Lattice”. *Phys.Rev.* D25, p. 2649. DOI: 10.1103/PhysRevD.25.2649 (cited on pp. 4, 21).
- Goldenfeld, N. (1992). “Lectures on phase transitions and the renormalization group” (cited on p. 24).
- Gross, D. and Wilczek, F. (1973a). “Asymptotically Free Gauge Theories. 1”. *Phys.Rev.* D8, pp. 3633–3652. DOI: 10.1103/PhysRevD.8.3633 (cited on pp. 1, 14).
- (1973b). “Ultraviolet Behavior of Nonabelian Gauge Theories”. *Phys.Rev.Lett.* 30, pp. 1343–1346. DOI: 10.1103/PhysRevLett.30.1343 (cited on pp. 1, 14).

- Hagedorn, R. (1965). "Statistical thermodynamics of strong interactions at high-energies". *Nuovo Cim.Suppl.* 3, pp. 147–186 (cited on p. 26).
- Halasz, A. M., Jackson, A., Shrock, R., Stephanov, M. A., and Verbaarschot, J. (1998). "On the phase diagram of QCD". *Phys.Rev.* D58, p. 096007. DOI: 10.1103/PhysRevD.58.096007. arXiv:hep-ph/9804290 [hep-ph] (cited on p. 57).
- Hasenbusch, M. (2001). "Speeding up the hybrid Monte Carlo algorithm for dynamical fermions". *Phys.Lett.* B519, pp. 177–182. DOI: 10.1016/S0370-2693(01)01102-9. arXiv:hep-lat/0107019 [hep-lat] (cited on p. 32).
- Hasenfratz, P. and Karsch, F. (1983). "Chemical Potential on the Lattice". *Phys.Lett.* B125, p. 308. DOI: 10.1016/0370-2693(83)91290-X (cited on p. 21).
- Higgs, P. W. (1964). "Broken Symmetries and the Masses of Gauge Bosons". *Phys.Rev.Lett.* 13, pp. 508–509. DOI: 10.1103/PhysRevLett.13.508 (cited on p. 1).
- Hooft, G. 't (1976). "Symmetry Breaking Through Bell-Jackiw Anomalies". *Phys.Rev.Lett.* 37, pp. 8–11. DOI: 10.1103/PhysRevLett.37.8 (cited on p. 11).
- Ilgenfritz, E.-M., Jansen, K., Lombardo, M., Muller-Preussker, M., Petschlies, M., et al. (2009). "Phase structure of thermal lattice QCD with $N(f) = 2$ twisted mass Wilson fermions". *Phys.Rev.* D80, p. 094502. DOI: 10.1103/PhysRevD.80.094502. arXiv:0905.3112 [hep-lat] (cited on pp. 18–20, 83).
- Iwasaki, Y., Kanaya, K., Kaya, S., Sakai, S., and Yoshie, T. (1996). "Finite temperature transitions in lattice QCD with Wilson quarks: Chiral transitions and the influence of the strange quark". *Phys.Rev.* D54, pp. 7010–7031. DOI: 10.1103/PhysRevD.54.7010. arXiv:hep-lat/9605030 [hep-lat] (cited on p. 90).
- Jansen, K. and Urbach, C. (2009). "tmLQCD: a program suite to simulate Wilson Twisted mass Lattice QCD". *Comput. Phys. Commun.* 180, pp. 2717–2738. DOI: 10.1016/j.cpc.2009.05.016. arXiv:0905.3331 [hep-lat] (cited on pp. 22, 32, 36, 43, 105, 108).
- Joó, B., Kalamkar, D. D., Vaidyanathan, K., Smelyanskiy, M., Pamnany, K., Lee, V. W., Dubey, P., and Watson, W. I. (2013). "Lattice QCD on Intel Xeon Phi Coprocessors". In: *Supercomputing*. Ed. by J. M. Kunkel, T. Ludwig, and H. W. Meuer. Springer Berlin Heidelberg, pp. 40–54. DOI: 10.1007/978-3-642-38750-0_4 (cited on pp. 40, 51).
- Kalcher, S., Rohr, D., Bach, M., Alaqeeli, A. A., Alzaid, H. M., Lindenstruth, V., Alkhereyf, S. B., Alharthi, A., Almubarak, A., Alqwaiz, I., Bin Suliman, R., and Eschweiler, D. (2013). "SANAM: An Energy- and Cost-Efficient Multi-GPU Supercomputer". *in preparation*. See also <http://www.amd.com/us/press-releases/Pages/amd-firepro-server-2012nov14.aspx> or <http://www.top500.org/system/177996> (cited on pp. 5, 39, 115).
- Kaplan, D. B. (1992). "A Method for simulating chiral fermions on the lattice". *Phys.Lett.* B288, pp. 342–347. DOI: 10.1016/0370-2693(92)91112-M. arXiv:hep-lat/9206013 [hep-lat] (cited on pp. 4, 21).

- Kapusta, J. and Gale, C. (2006). “Finite-temperature field theory: Principles and applications” (cited on pp. 7, 9, 24).
- Karsch, F., Laermann, E., and Schmidt, C. (2001). “The Chiral critical point in three-flavor QCD”. *Phys.Lett.* B520, pp. 41–49. DOI: 10.1016/S0370-2693(01)01114-5. arXiv:hep-lat/0107020 [hep-lat] (cited on p. 56).
- Kennedy, A. D. and Pendleton, B. J. (1985). “Improved Heatbath Method for Monte Carlo Calculations in Lattice Gauge Theories”. *Physical Letters, Volume 156B, number 5,6* (cited on p. 32).
- Khronos Working Group (2013). *The OpenCL Specification*. <http://www.khronos.org/registry/cl/> (cited on pp. 5, 40, 53).
- Kogut, J. B. and Susskind, L. (1975). “Hamiltonian Formulation of Wilson’s Lattice Gauge Theories”. *Phys.Rev.* D11, p. 395. DOI: 10.1103/PhysRevD.11.395 (cited on pp. 4, 21).
- Laermann, E. and Philipsen, O. (2003). “The Status of lattice QCD at finite temperature”. *Ann.Rev.Nucl.Part.Sci.* 53, pp. 163–198. DOI: 10.1146/annurev.nucl.53.041002.110609. arXiv:hep-ph/0303042 [hep-ph] (cited on p. 55).
- Levkova, L. (2011). “QCD at nonzero temperature and density”. *PoS, LATTICE2011*, p. 011. arXiv:1201.1516 [hep-lat] (cited on pp. 34, 55, 85).
- Linde, A. D. (1980). “Infrared Problem in Thermodynamics of the Yang-Mills Gas”. *Phys.Lett.* B96, p. 289. DOI: 10.1016/0370-2693(80)90769-8 (cited on p. 2).
- Lüscher, M. (1994). “A portable high-quality random number generator for lattice field theory simulations”. *Computer Physics Communications*, 79(1), pp. 100–110. DOI: 10.1016/0010-4655(94)90232-1 (cited on p. 43).
- Majumdar, P. (2013). “Lattice Simulations using OpenACC compilers”. In: *The XXXI International Symposium on Lattice Field Theory, Lattice2013* (cited on p. 39).
- Matthews, P. and Salam, A. (1954). “The Green’s functions of quantized fields”. *Nuovo Cim.* 12, pp. 563–565. DOI: 10.1007/BF02781302 (cited on p. 109).
- (1955). “Propagators of quantized field”. *Nuovo Cim.* 2, pp. 120–134. DOI: 10.1007/BF02856011 (cited on p. 109).
- Meister, A. (2008). *Numerik linearer Gleichungssysteme*. Wiesbaden: Vieweg. DOI: 10.1007/978-3-8348-8100-7 (cited on p. 32).
- Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A., and Teller, E. (1953). “Equation of state calculations by fast computing machines”. *J.Chem.Phys.* 21, pp. 1087–1092. DOI: 10.1063/1.1699114 (cited on p. 31).
- Montvay, I. and Münster, G. (1994). “Quantum fields on a lattice” (cited on pp. 2, 7, 22, 30, 35, 108).

- Morningstar, C. and Peardon, M. (2004). “Analytic smearing of SU(3) link variables in lattice QCD”. *Phys. Rev. D* 69, p. 054501. DOI: 10.1103/PhysRevD.69.054501. arXiv:hep-lat/0311018 (cited on p. 103).
- Nagata, K. and Nakamura, A. (2011). “Imaginary Chemical Potential Approach for the Pseudo-Critical Line in the QCD Phase Diagram with Clover-Improved Wilson Fermions”. *Phys.Rev. D* 83, p. 114507. DOI: 10.1103/PhysRevD.83.114507. arXiv:1104.2142 [hep-lat] (cited on p. 86).
- (2012). “EoS of finite density QCD with Wilson fermions by Multi-Parameter Reweighting and Taylor expansion”. *JHEP*, 1204, p. 092. DOI: 10.1007/JHEP04(2012)092. arXiv:1201.2765 [hep-lat] (cited on p. 86).
- Nakamura, A., Nagata, K., Motoki, S., Nakagawa, Y., and Saito, T. (2011). “Lattice QCD and High Baryon Density State”. *AIP Conf.Proc.* 1388, pp. 509–515. DOI: 10.1063/1.3647441. arXiv:1109.0490 [hep-lat] (cited on p. 86).
- Neuberger, H. (1998a). “A Practical implementation of the overlap Dirac operator”. *Phys.Rev.Lett.* 81, pp. 4060–4062. DOI: 10.1103/PhysRevLett.81.4060. arXiv:hep-lat/9806025 [hep-lat] (cited on pp. 4, 21).
- (1998b). “Exactly massless quarks on the lattice”. *Phys.Lett.* B417, pp. 141–144. DOI: 10.1016/S0370-2693(97)01368-3. arXiv:hep-lat/9707022 [hep-lat] (cited on pp. 4, 21).
- Nielsen, H. B. and Ninomiya, M. (1981a). “Absence of Neutrinos on a Lattice. 1. Proof by Homotopy Theory”. *Nucl.Phys.* B185, p. 20. DOI: 10.1016/0550-3213(81)90361-8 (cited on p. 18).
- (1981b). “Absence of Neutrinos on a Lattice. 2. Intuitive Topological Proof”. *Nucl.Phys.* B193, p. 173. DOI: 10.1016/0550-3213(81)90524-1 (cited on p. 18).
- (1981c). “No Go Theorem for Regularizing Chiral Fermions”. *Phys.Lett.* B105, p. 219. DOI: 10.1016/0370-2693(81)91026-1 (cited on p. 18).
- NVIDIA (2013). *NVIDIA CUDA C Programming Guide*. <http://developer.nvidia.com/nvidia-gpu-computing-documentation> (cited on pp. 5, 40, 41, 53).
- Padua, D., ed. (2011). *Encyclopedia of Parallel Computing*. Boston, MA: Springer US. DOI: 10.1007/978-0-387-09766-4 (cited on p. 37).
- Parisen Toldin, F., Pelissetto, A., and Vicari, E. (2003). “The 3-D O(4) universality class and the phase transition in two flavor QCD”. *JHEP*, 0307, p. 029. arXiv:hep-ph/0305264 [hep-ph] (cited on pp. 26, 57).
- Pelissetto, A. and Vicari, E. (2002). “Critical phenomena and renormalization group theory”. *Phys.Rept.* 368, pp. 549–727. DOI: 10.1016/S0370-1573(02)00219-3. arXiv:cond-mat/0012164 [cond-mat] (cited on pp. 26, 91).
- Perkins, D. (1982). “Introduction to high energy physics” (cited on p. 1).

- Peskin, M. E. and Schroeder, D. V. (1995). “An Introduction to quantum field theory” (cited on pp. 7, 107, 108).
- Petersen, H., Steinheimer, J., Burau, G., Bleicher, M., and Stocker, H. (2008). “A Fully Integrated Transport Approach to Heavy Ion Reactions with an Intermediate Hydrodynamic Stage”. *Phys.Rev.* C78, p. 044901. DOI: 10.1103/PhysRevC.78.044901. arXiv:0806.1695 [nucl-th] (cited on p. 26).
- Petreczky, P. (2012). “Review of recent highlights in lattice calculations at finite temperature and finite density”. *PoS, ConfinementX*, p. 028. arXiv:1301.6188 [hep-lat] (cited on p. 55).
- Petronzio, R. and Vicari, E. (1990). “An overrelaxed Monte Carlo algorithm for SU (3) lattice gauge theories”. *Physics Letters B*, 248, pp. 159–162. DOI: 10.1016/0370-2693(90)90032-2 (cited on p. 32).
- Philipsen, O. (2013). “The QCD equation of state from the lattice”. *Prog.Part.Nucl.Phys.* 70, pp. 55–107. DOI: 10.1016/j.pnpnp.2012.09.003. arXiv:1207.5999 [hep-lat] (cited on p. 26).
- Philipsen, O. and Pinke, C. (2014). “The nature of the Roberge-Weiss transition in $N_f = 2$ QCD with Wilson fermions”. arXiv:1402.0838 [hep-lat] (cited on p. 85).
- Philipsen, O. and Schäfer, C. (2013). “The second order hydrodynamic transport coefficient κ for the gluon plasma from the lattice”. arXiv:1311.6618 [hep-lat] (cited on pp. 43, 49).
- Philipsen, O., Pinke, C., Schäfer, C., Zeidlewicz, L., and Bach, M. (2011). “LatticeQCD using OpenCL”. *PoS, LATTICE2011*, p. 044. arXiv:1112.5280 [hep-lat] (cited on pp. 43, 44, 104).
- Philipsen, O., Langelage, J., Lottini, S., and Neuman, M. (2012). “Finite density QCD from an effective lattice theory”. *PoS, ConfinementX*, p. 218 (cited on p. 105).
- Pisarski, R. D. and Wilczek, F. (1984). “Remarks on the Chiral Phase Transition in Chromodynamics”. *Phys.Rev.* D29, pp. 338–341. DOI: 10.1103/PhysRevD.29.338 (cited on pp. 11, 56).
- Politzer, H. D. (1973). “Reliable Perturbative Results for Strong Interactions?” *Phys.Rev.Lett.* 30, pp. 1346–1349. DOI: 10.1103/PhysRevLett.30.1346 (cited on pp. 1, 14).
- Povh, B., Scholz, C., Rith, K., and Zetsche, F. (1993). “Particles and nuclei: An introduction to the physical conceptions. (In German)” (cited on p. 1).
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical Recipes*. Cambridge University Press (cited on p. 44).
- Rajagopal, K. and Wilczek, F. (1993). “Static and dynamic critical phenomena at a second order QCD phase transition”. *Nucl.Phys.* B399, pp. 395–425. DOI: 10.1016/0550-3213(93)90502-G. arXiv:hep-ph/9210253 [hep-ph] (cited on p. 57).

- Roberge, A. and Weiss, N. (1986). "Gauge Theories with imaginary chemical potential and the phases of QCD". *Nucl.Phys.* B275, p. 734. DOI: 10.1016/0550-3213(86)90582-1 (cited on pp. 4, 13, 86).
- Sexty, D. (2013). "Simulating full QCD at nonzero density using the complex Langevin equation". arXiv:1307.7748 [hep-lat] (cited on p. 103).
- Shamir, Y. (1993). "Chiral fermions from lattice boundaries". *Nucl.Phys.* B406, pp. 90-106. DOI: 10.1016/0550-3213(93)90162-I. arXiv:hep-lat/9303005 [hep-lat] (cited on pp. 4, 21).
- Sharpe, S. R. (2006). "Rooted staggered fermions: Good, bad or ugly?" *PoS, LAT2006*, p. 022. arXiv:hep-lat/0610094 [hep-lat] (cited on pp. 5, 21, 85).
- Sheikholeslami, B. and Wohlert, R. (1985). "Improved Continuum Limit Lattice Action for QCD with Wilson Fermions". *Nucl.Phys.* B259, p. 572. DOI: 10.1016/0550-3213(85)90002-1 (cited on pp. 4, 21).
- Shindler, A. (2008). "Twisted mass lattice QCD". *Phys. Rept.* 461, pp. 37-110. DOI: 10.1016/j.physrep.2008.03.001. arXiv:0707.4093 [hep-lat] (cited on pp. 14, 19, 20).
- 't Hooft, G. (1999). "Monopoles, instantons and confinement". arXiv:hep-th/0010225 [hep-th] (cited on p. 2).
- Takaishi, T. and de Forcrand, P. (2006). "Testing and tuning new symplectic integrators for hybrid Monte Carlo algorithm in lattice QCD". *Phys.Rev.* E73, p. 036706. DOI: 10.1103/PhysRevE.73.036706. arXiv:hep-lat/0505020 [hep-lat] (cited on p. 31).
- Takaishi, T., de Forcrand, P., and Nakamura, A. (2009). "Equation of State at Finite Density from Imaginary Chemical Potential". *PoS, LAT2009*, p. 198. arXiv:1002.0890 [hep-lat] (cited on p. 86).
- Ukawa, A. (2002). "Computational cost of full QCD simulations experienced by CP-PACS and JLQCD Collaborations". *Nuclear Physics B - Proceedings Supplements*, 106-107(0). LATTICE 2001 Proceedings of the XIXth International Symposium on Lattice Field Theory, pp. 195-196. DOI: [http://dx.doi.org/10.1016/S0920-5632\(01\)01662-0](http://dx.doi.org/10.1016/S0920-5632(01)01662-0) (cited on p. 33).
- Urbach, C., Jansen, K., Shindler, A., and Wenger, U. (2006). "HMC algorithm with multiple time scale integration and mass preconditioning". *Comput.Phys.Commun.* 174, pp. 87-98. DOI: 10.1016/j.cpc.2005.08.006. arXiv:hep-lat/0506011 [hep-lat] (cited on p. 32).
- Weisz, P. (1983). "Continuum Limit Improved Lattice Action for Pure Yang-Mills Theory. 1." *Nucl.Phys.* B212, p. 1. DOI: 10.1016/0550-3213(83)90595-3 (cited on p. 17).
- Weisz, P. and Wohlert, R. (1984). "Continuum Limit Improved Lattice Action for Pure Yang-Mills Theory. 2." *Nucl.Phys.* B236, p. 397. DOI: 10.1016/0550-3213(84)90543-1 (cited on p. 17).

- Wilson, K. G. (1974). "Confinement of Quarks". *Phys.Rev.* D10, pp. 2445–2459. DOI: 10.1103/PhysRevD.10.2445 (cited on pp. 4, 17).
- Wolff, U. (2004). "Monte Carlo errors with less errors". *Comput.Phys.Commun.* 156, pp. 143–153. DOI: 10.1016/S0010-4655(03)00467-3, 10.1016/j.cpc.2006.12.001. arXiv:hep-lat/0306017 [hep-lat] (cited on p. 35).
- Wu, L.-K. and Meng, X.-F. (2013). "Nature of the Roberge-Weiss transition end points in two-flavor lattice QCD with Wilson quarks". *Phys.Rev.* D87, p. 094508. DOI: 10.1103/PhysRevD.87.094508. arXiv:1303.0336 [hep-lat] (cited on pp. 85, 86, 89, 92, 102, 105).
- Wu, L.-K., Luo, X.-Q., and Chen, H.-S. (2007). "Phase structure of lattice QCD with two flavors of Wilson quarks at finite temperature and chemical potential". *Phys.Rev.* D76, p. 034505. DOI: 10.1103/PhysRevD.76.034505. arXiv:hep-lat/0611035 [hep-lat] (cited on pp. 85, 86, 89, 90).
- Zeidlewicz, L. (2011). *The Thermal Transition of Quantum Chromodynamics with Twisted Mass Fermions (Disputation)*. See also <http://www.dart-europe.eu/full.php?id=820507>. Univ.-Bibliothek Frankfurt am Main (cited on pp. 76, 81, 84).

DANKSAGUNG

Danke, danke, danke!

— Christopher Pinke

Diesen Worten möchte ich mich anschließen und mich an dieser Stelle bei allen Leuten bedanken, ohne deren Unterstützung ich die vorliegende Arbeit nicht hätte fertigstellen können.

An erster Stelle gebührt Owe Philipsen ein großer Dank, welcher mir nach meiner Diplomandenzeit in Münster die Möglichkeit bot mit ihm nach Frankfurt zu kommen. Es war sehr spannend, ein ganz neues Umfeld zu erleben und auch, den Aufbau einer Arbeitsgruppe so nah mitzubekommen. Weiterhin möchte ich mich für das große Vertrauen bedanken, das Owe mir entgegenbringt und welches das Arbeiten mit ihm sehr angenehm gestaltet. Gleichermäßen konnte ich bei Problemen jeglicher Art stets auf seine kompetente Hilfe zählen.

Ich möchte mich besonders bei denjenigen bedanken, mit denen ich an den verschiedenen Projekten gearbeitet habe. Auf Lars Zeidlewicz und Florian Burger konnte ich immer zählen wenn es einmal (oder auch öfter) Probleme beim Simulieren oder Analysieren gab. Hervorheben möchte ich Matthias Bach, ohne den CL²QCD in der Form nicht möglich gewesen wäre. Wir können, denke ich, zurecht stolz sein auf das, was aus dem Projekt geworden ist. Ich habe sehr viel bei der Arbeit mit ihm gelernt; sei es über C++, GPUs oder ein beliebiges anderes Thema, ich habe stets schnelle und kompetente Antworten auf meine Fragen erhalten.

Einen großen Dank möchte ich an alle Mitglieder der AG Philipsen richten. Ich habe das Arbeitsklima immer als ausgesprochen angenehm und freundschaftlich empfunden, sei es auf Weihnachtsfeiern, Ausflügen oder ganz einfach im Büroalltag. Nicht vergessen darf man an dieser Stelle Andrea, ohne deren Zutun vieles nicht so reibungslos funktionieren würde und bei der man stets eine offene Tür vorfindet. Mit viel Freude denke ich an die Zeit in der Riedberg-WG zurück, an Pizza, Pasta und Zwiebelkuchen. Danke an euch, Lars und Steo.

Weiterhin möchte ich mich bei den Verantwortlichen und Mitarbeitern von HGS-HIRe bedanken. Ich habe die Angebote, welche die Graduiertenschule bietet, gerne angenommen und weiß sie zu schätzen.

Für das Durchsehen und Korrigieren dieser Arbeit bedanke ich mich bei Georg Bergner, Alessandro Sciarra, Christian Schäfer, Matthias Bach und Stephanie Merker.

Zu guter Letzt möchte ich mich bei meiner Familie bedanken, auf die ich immer zählen kann und werde. Und natürlich hätte ich das alles nicht schaffen können ohne meine Steffi. Danke dass es dich gibt!