MAXIMALLY TWISTED MASS LATTICE QCD AT THE PHYSICAL PION MASS

DISSERTATION

zur Erlangung des akademischen Grades

doctor rerum naturalium
(Dr. rer. nat.)

im Fach Physik

eingereicht an der
Mathematisch-Naturwissenschaflichen Fakultät
Humboldt-Universität zu Berlin

von
M. Sc. Bartosz Kostrzewa

Präsident der Humboldt-Universität zu Berlin
Prof. Dr. Jan-Hendrik Olbertz

Dekan der Mathematisch-Naturwissenschaftlichen Fakultät
Prof. Dr. Elmar Kulke

Gutachter/innen:
1. Dr. Karl Jansen
2. Prof. Dr. Rainer Sommer
3. Prof. Dr. Peter Boyle

Tag der mündlichen Prüfung: 21.06.2016

ABSTRACT

In computer simulations of Lattice Quantum Chromodynamics, the usage of unphysically large quark masses and the subsequent extrapolation of results to the physical value of the quark masses are major sources of systematic uncertainty. In this thesis, the feasibility and practicality of numerical simulations of Quantum Chromodynamics with physically light up and down quarks using the Wilson twisted mass quark discretisation are explored. Working in this regime is complicated firstly by the numerical expense of these simulations and secondly by the presence of potentially large lattice artefacts. The twisted mass discretisation is affected by an unphysical mass difference between the charged and neutral pions, rendering simulations at the physical charged pion mass infeasible if this mass splitting is too large.

With the aim of reducing it, the Sheikholeslami-Wohlert term is added to the twisted mass fermion action and simulations with mass degenerate up and down quarks are then performed as a proof of concept. It is demonstrated that these simulations are stable and that the parameters of the lattice theory can be successfully tuned to correspond to the physical charged pion mass. Subsequently, the parameter tuning for simulations with mass degenerate up and down quarks as well as strange and charm quarks is explored and it is shown that it can be carried out in steps.

As benchmark observables, the masses and decay constants of pseudoscalar mesons with light, strange and charm valence quarks are calculated and seen to largely reproduce their phenomenological values, even though continuum and infinite volume extrapolations are not performed. Light, strange and charm quark mass estimates are determined based on this data and also seen to coincide with phenomenological and other lattice determinations. In this analysis, a particular emphasis is placed on the systematic error due to the choice of fit range for pseudoscalar correlation functions and a weighting method is introduced which may become very useful on very large lattices. The pion mass splitting is studied as a function of the Sheikholeslami-Wohlert coefficient in simulations with four flavours and it is found to be approximately halved twisted mass quarks without this term. However, a dependence on the precise value of the coefficient cannot be identified within the large uncertainties and within the range of values studied.

To optimise the Hybrid Monte Carlo algorithm, mass preconditioning is explored empirically through simple fits to the magnitude of molecular dynamics forces generated by quark determinants and determinant ratios with a wide range of parameter values. Based on the functional form of these fits, mass preconditioning and integration schemes are proposed in which the relationships between all parameters are tuned simultaneously and which may allow more efficient simulations with predictable relative force magnitudes. As a complement to this work, a tentative study of the oscillation frequencies of these forces is performed with the finding that mass preconditioning seems to suppress large amplitude, high frequency oscillations in addition to reducing force magnitudes.

Crucial optimisations of the simulation software for twisted mass quarks are introduced. A multi-threading strategy based on OpenMP is devised and kernels which overlap communication and computation are developed and benchmarked on various architectures. Testing methodologies for the simulation code are presented and it is shown how they complement each other based on specific examples, providing a rather general set of integration tests.

ZUSAMMENFASSUNG

In der Gitterquantenchromodynamik sind der Einsatz von unphysikalisch großen Quarkmassen und die Extrapolation von Ergebnissen zu physikalischen Massen signifikante systematische Fehlerquellen. In dieser Arbeit wird die praktische Durchführbarkeit numerischer Simulationen der Quantenchromodynamik mit physikalisch leichten up und down Quarkmassen unter Verwendung der Wilson twisted mass Diskretisierung untersucht. Simulationen im Regime physikalisch leichter Quarkmassen sind jedoch einerseits numerisch sehr aufwendig, können andererseits aber auch durch das Auftreten großer Diskretisierungsartefakte nicht praktikabel sein. Fermionen des Wilson twisted mass Typs führen durch solche Gitterartefakte zu einer unphyskalischen Massendifferenz zwischen den geladenen und dem neutralen Pion welche, wenn sie zu groß ist, eine Simulation bei physikalisch leichter Quarkmasse verhindert.

Anhand von Simulationen mit massendegenerierten dynamischen up und down Quarks wird dargestellt dass die Erweiterung der twisted mass Fermionwirkung durch den Sheikholeslami-Wohlert Term es ermöglicht physikalisch leichte Quarkmassen zu erreichen. Es wird gezeigt, dass die Simulationen stabil sind und dass die Parameter der diskretisierten Theorie so gewählt werden können, dass das geladene Pion seine physikalische Masse annimmt. Ferner wird dargestellt, dass auch die Parameter für eine Simulation mit dynamischen massendegenerierten up und down quarks sowie nichtdegenerierten strange und charm Quarks schrittweise auf ihre physkalischen Werte gesetzt werden können.

Um das Verhalten von Observablen bei physikalischer Quarkmasse zu untersuchen, werden Massen und Zerfallskonstanten von pseudoskalaren Mesonen mit up, down sowie strange und charm Valenzquarks berechnet. Die Ergebnisse stimmen größtenteils überein mit den phänomenologischen Werten, obwohl weder Kontinuumslimes noch die Extrapolation zu unendlichem Volumen durchgeführt werden. Renormierte leichte, strange und charm Quarkmassen werden über Interpolationen in hadronischen Observablen berechnet und stimmen ebenso größtenteils mit phänomenologischen Werten und anderen Ergebnissen aus der Gitter-QCD überein. In diesen Analysen wird ein besonderes Augenmerk darauf gelegt den systematsischen Fehler zu bestimmen welcher durch die Mehrdeutigkeit des optimalen Kurvenanpassungsintervalls der Korrelationsfunktionen entsteht. Zur Abschätzung dieses Fehlers wird eine Wichtungsmethode eingeführt welche besonders auf sehr großen Gittern, wie sie für Simulationen mit physikalischer Quarkmassen von Bedarf sind, nützlich sein wird.

In Simulationen mit vier dynamischen Quarks wird die Abhängigkeit der Pionmassendifferenz vom Sheikholeslami-Wohlert Koeffizienten untersucht. Es wird gezeigt, dass die Massendifferenz im Vergleich zu Simulationen ohne Sheikholeslami-Wohlert Term ungefähr halbiert wird, eine genaue Abhängigkeit vom Koeffizienten kann aber innherhalb der Fehler und im untersuchten Intervall nicht festegestellt werden.

Zur Optimierung des zur Simulation genutzten Hybrid-Monte-Carlo-Algoritmus wird die Parameterabhängigkeit der Massenprökonditionierung über sehr große Parameterintervalle empirisch untersucht. Durch Kurvenanpassung werden diese Daten parametrisiert und es wird gezeigt, dass anhand der Kurven Verhältnisse zwichen den Parametern aufgestellt werden können, welche es erlauben Intergrationsschemas zu optimieren. Komplementär dazu wird versuchsweise das Oszillationsverhalten der Kräfte untersucht und es scheint, dass große Massenpräkonditionierugsparameter nicht nur die Größe der Kräfte stark beinflussen, sondern auch hochfrequente Oszillationen unterdrücken.

Es werden Optimierungen der Simulationsprogramme für twisted mass quarks vorgestellt. Insbesonere wird eine Strategie für Multithreading der Rechenkernel via OpenMP engeführt. Überlappende Kommunikation und Berechnung unter Zuhilfenahme von Threads werden zur Effizienzsteigerung genutzt und Benchmarkergebnisse auf verschiedenen Rechnerarchitekturen werden vorgestellt und interpretiert. Strategien für allgemeine Integrationstests des Simulationscodes werden vorgestellt und es wird gezeigt wie die verschiedenen gezeigten Techniken zueinander komplementär sind.

CONTENTS

INTRODUCTION

The physics of the early twentieth century led to the realisation that attempting a complete description of nature would require the reconciliation of the concepts of particles and classical fields. Building the connection between special relativity, classical field theory and quantum mechanics led to the development of causal quantum field theories (QFTs), which describe fundamental interactions in terms of local operator-valued random variables. When the Hilbert space acted on by these operators is constructed, particle-like states with well-defined quantum numbers emerge in the spectra of the resulting Hamiltonians. In this sense, quantum field theories appear to provide the requisite properties.

At the centre of these developments stand quantised gauge theories, the classical Lagrangian of which is invariant under local gauge transformations [1, 2] of some gauge group G. In the case of electromagnetism, this gauge group is $U(1)$, while more generally, G is some compact gauge group or direct product of compact groups. When these Yang-Mills theories [3] are combined with matter fields in a gauge-invariant manner, the latter transform locally under representations of G while the gauge connection A provides interactions through the gauge covariant derivative. The action contribution of the gauge fields and their self-interaction is in turn given by an invariant quadratic form on the Lie algebra of G, expressed in terms of the field strength tensor F. Matter fields are said to be *charged* if they transform non-trivially or acquire a phase under gauge transformations of G or its subgroups.

The Standard Model of particle physics (SM) is a gauge theory with $G = SU(3) \times SU(2) \times U(1)$ and describes the interaction of matter fields of half-integer spin through force carriers of integer spin, the gauge fields. It can be separated into two parts: the strongly interacting sector with colour gauge group $SU(3)$ on the one hand, and the electroweak sector with gauge group $SU(2) \times U(1)$. The six quarks and six anti-quarks are charged under colour $SU(3)$ and electroweak $SU(2) \times U(1)$. Three leptons and three anti-leptons are charged under $SU(2) \times U(1)$, while the respective neutrinos and anti-neutrinos are electrically neutral. The matter fields can be arranged into three *families* or *generations* of two quarks each and pairings of a charged lepton and its respective neutrino.

The masses of the quarks and leptons increase with generation number, but large mass differences also exist within the generations. Neutrinos are all very light while the respective electron-like leptons have a range of masses. The mass of the muon is about two hundred times larger than that of the electron and the mass of the tau is about another twenty times larger still. A discussion of quark masses is somewhat subtle, but the down and up quarks can be thought of as very light, with their masses differing by less than their average. The strange and charm quarks are about thirty and three hundred times heavier respectively and thus also have a large mass difference. Finally, the masses of the bottom and top quarks are around one thousand and forty thousand times larger than the average down/up quark mass, respectively.

The electroweak sector is described by the Weinberg-Salam-Glashow electroweak theory [4, 5, 6] of leptons (and quarks, in the SM) transforming under representations of the gauge group $SU(2) \times U(1)$. The Higgs

mechanism [7, 8, 9] provides a gauge-invariant process by which the gauge fields in this sector, coupled to a complex scalar field doublet, result in the weak interactions through three massive weak gauge bosons, $W^{\pm}$ and Z. The remaining combination of gauge fields is massless, providing the electromagnetic force mediator, the photon. An experimental confirmation of the existence of the Higgs field has been provided by the ATLAS and CMS experiments through the observation of a resonance consistent with a scalar particle with a mass of $125.09(21)(11)$ GeV/c^2, as given by the combination of data in Ref. [10]. Because a more complete discussion of the electroweak sector is beyond the scope of this exposition, interesting topics such as neutrino oscillations which were the basis of the 2015 Nobel Prize in physics, must be omitted.

The strongly interacting sector of quarks and gluons in the SM is described by the unbroken local $SU(3)$ colour gauge symmetry of Quantum Chromodynamics [11, 12, 13]. At high energy, QCD is weakly coupled and can be very successfully described perturbatively: the quarks are said to be *asymptotically free* [14]. At low energy, however, the coupling is strong, such that quarks and gluons cannot be observed as isolated free particles, a property also referred to as *confinement*. Instead, colour-neutral hadronic bound states of two and three quarks, mesons and baryons[1], populate the low energy spectrum of the theory. Despite its long history and importance, the many remaining open questions concerning the low energy properties of QCD are active areas of study.

In this regime, perturbation theory in terms of quark and gluon fields is inapplicable and non-perturbative approaches are required. The most developed and successful amongst these is lattice Quantum Chromodynamics (LQCD). This framework allows the theory to be explored non-perturbatively by discretising it on a four-dimensional regular lattice while fully preserving local gauge invariance. Particularly, doing so on a hypercubic Eucidean lattice allows the theory to be studied by numerical simulation. Its path integral can then be evaluated using Monte Carlo integration with importance sampling and observables become averages over statistical ensembles.

Hadron masses and meson decay constants can then be calculated directly from the asymptotic behaviour of appropriate correlation functions. Matrix elements can similarly be studied non-perturbatively, providing certain hadronic form factors or input for operator product expansions [16] of effective Hamiltonians. Unstable particles [17] and scattering [18, 19] are also accessible through the effects of the finite volume on the particle spectrum of the theory. Another major long-term goal within LQCD is the purely theoretical computation [20] of parton distribution functions which strongly affect the uncertainties of experiments at hadron colliders like the LHC [21]. Lattice QCD calculations can also contribute significantly to tests of the SM and searches for physics beyond. Examples are given by the hadronic contribution to the vacuum polarisation [22], various computations in the flavour sector [23] or determinations of the neutron electric dipole moment [24].

Lattice QCD offers an ab-initio approach in which all theoretical uncertainties can – in principle – be quantified and systematically improved upon. In general, these improvements come at significant computational cost, as a result of which numerical algorithms on supercomputers and their optimisation are areas of intense study within the LQCD community. The

[1] Experimental evidence for resonances of five bound quarks has recently been reported in Ref. [15].

simulations take place in finite volume at non-zero lattice spacing, both of which constitute *lattice artefacts* with particular consequences. These can be understood analytically through effective theories such as the Symanzik expansion [25, 26] and chiral perturbation theory (χPT) in finite volume. With this guidance, lattice artefacts can be studied and eliminated through simulations in a number of volumes and with a number of lattice spacings, although approaching both the infinite volume and continuum limits comes with significant computational cost.

Simulations of lattice QCD can be carried out at various levels of approximation with regards to how sea quarks contribute to the polarisation of the vacuum (equivalently, how they enter the importance sampling of of the LQCD path integral and thus affect observables). The crudest approximation takes the sea quarks to be infinitely massive, or *quenched*. The approximation is improved by including the effect of mass degenerate up and down quarks and improved further by also including the effect of the strange and even the charm quark. The most ambitious simulations lift the degeneracy of the light quark doublet and finally also add electromagnetic effects. These different approximations are often labelled as $N_f = 0$, $N_f = 2$, $N_f = 2+1$, $N_f = 2+1+1$ and $N_f = 1+1+1+1$. Increasing the number of active flavours can lead to the appearance of additional lattice artefacts which need to be checked. The computational effort is increased by the additional quark flavours and more significantly, when the masses of the light quarks are reduced towards their physical values.

As a consequence, for a long time, unquenched simulations of LQCD were carried out at a number of unphysically large values of the up and down quark masses. The results of these computations were then extrapolated towards the physical value, either using polynomial approaches or extrapolation functions guided by χPT. While this works rather well for many quantities, the extrapolation potentially introduces unquantified systematic errors. This is especially true in the baryon sector, where for example the axial charge of the nucleon shows some tension with experiment and it is unclear whether this is due to unquantified systematic uncertainties [27]. Another problematic area is the study of mesons containing very light and heavy quarks. The range of validity of the used effective theories is only known up to order of magnitude inequalities and large differences in mass scales can complicate judging this even further. Finally, even for quantities where effective theories provide very good guidance, the extrapolation can increase overall uncertainties and make it difficult to provide the kind of precision required by modern phenomenology for tests of the Standard Model. Eliminating these systematic errors requires simulations with physically light up and down quarks.

Lattice QCD is not unique: different types of discretisations will behave differently with respect to various criteria relevant for their practical usage. It is the aim of this thesis to investigate the feasibility of simulations using physically light quarks with a type of quark discretisation referred to as twisted mass lattice QCD (tmLQCD) with $N_f = 2$ and $N_f = 2+1+1$ active quark flavours. Twisted mass lattice QCD has a number of special properties which make it suitable for the computation of a very wide range of phenomenologically interesting quantities. However, certain lattice artefacts may make it especially difficult in tmLQCD to reach the regime of physical light quark masses at tolerable computational cost and without sacrificing the control over systematic errors offered by LQCD. A particular lattice artefact which appears in tmLQCD is an unphysical mass difference

between the neutral and the charged pion, subtly related to properties of the zero temperature phase structure of the lattice theory. When this mass difference is too large, simulations at the physical light quark mass are not possible. In the present work, a particular extension of tmLQCD is studied which is seen to significantly reduce this pion mass splitting. This allows simulations with two physically light, mass degenerate quarks to be performed and may continue to do so in simulations including strange and charm quarks. For the latter, the tuning of the numerous parameters is explored and proposals are made for how to proceed in practice.

Benchmark computations of a number of physical quantities are carried out based on these two and four flavour simulations. Firstly, the pion mass difference is calculated at different values of the parameters of the theory to check for the appearance of unexpected lattice artefacts. The results are also compared with those of previous simulations to quantify the level of improvement achieved. Secondly, the masses and decay constants of mesons containing light, strange and charm quarks are computed and compared to their phenomenological values as well as the results of other lattice calculations. Using ratios of meson masses as tuning conditions, estimates of the light, strange and charm quark masses as well as their ratios are computed. A particular emphasis is put on the study of systematic errors due to the presence of correlations in the data and a somewhat novel technique for their quantification is presented.

When the simulated quarks are physically light, the numerical algorithms used to sample the path integral for Monte Carlo integration suffer from a number of slowdowns. A significant portion of this thesis is therefore dedicated to their study and how to potentially overcome these issues. In addition, even once the simulation algorithms have been tuned to allow physically light quarks, the computational challenge is still formidable. For this reason, many optimisations were introduced into the tmLQCD simulation software and these as well as the resulting performance characteristics are reported on. Several general techniques for testing the correctness of LQCD codes are also presented and their applicability is demonstrated with specific examples.

The thesis is thus structured into four parts. Part 1 provides the relevant theoretical and algorithmic basis for the study of tmLQCD and frames the present results in the wider context of current simulations in LQCD. Part 2 is dedicated to a report on the simulations that were performed and how the parameters of these simulations were tuned, especially for the $N_f = 2+1+1$ case. Part 3 is concerned with the physical quantities that were computed on the basis of the simulations of Part 2. Finally, Part 4 collects the computational developments, testing methodologies and performance characteristics of the used simulation code as well as a number of novel techniques relating to these.

Part I

FUNDAMENTALS

1 THEORETICAL BACKGROUND

This thesis is concerned with the study of non-perturbative properties of Quantum Chromodynamics (QCD) using numerical simulations of Lattice QCD (LQCD) with twisted mass Wilson quarks in the regime of physically light quark masses. QCD will be introduced below, followed by a discussion of LQCD in Section 1.2 with a particular consideration of the computational challenge represented by numerical simulations of LQCD. The twisted mass lattice discretisation (tmLQCD) will be discussed in Section 1.3. Because the results presented here hinge on the size of certain lattice artefacts, the concept of improvement (the removal of lattice artefacts) and how it applies to tmLQCD will be covered. The Hybrid Monte Carlo algorithm for simulations of tmLQCD and some of its optimisations will be presented in Section 1.4. Finally, Section 2.2 closes the chapter with a discussion of the phase structure of Wilson lattice QCD. Specifically, its relationship to algorithmic difficulties of tmLQCD simulations in the regime of physically light quark masses is discussed and the addition of the so-called *clover* term to the twisted mass action is motivated.

QUANTUM CHROMODYNAMICS

The strongly interacting sector of quarks and gluons in the SM is described by the unbroken local SU(3) colour gauge symmetry of Quantum Chromodynamics (QCD) [11, 12, 13]. The gauge-invariant QCD Lagrange density, with all indices explicitly exposed, is given by

$$\begin{aligned}\mathcal{L}_{\mathrm{QCD}}(x) = &\sum_{f=1}^{n_f}\sum_{\alpha,\beta=1}^{4}\sum_{i,j=1}^{3}\sum_{\mu=0}^{3}\bar{\psi}^{j}_{f\alpha}(x)\left(i\gamma^{\mu}_{\alpha\beta}D^{ij}_{\mu} - m_f\delta_{\alpha\beta}\delta^{ij}\right)\psi^{j}_{f\beta}(x)\\ &-\frac{1}{4}\sum_{\mu,\nu=0}^{3}\sum_{a=1}^{3^2-1}F^{a}_{\mu\nu}(x)F^{\mu\nu a}(x)\,.\end{aligned} \tag{1.1}$$

The sums are over quark flavours f, Dirac spinor indicies α, β, colour indices i, j, a and four-vector indices μ, ν. The gauge covariant derivative is

$$D^{ij}_{\mu} \equiv \partial_{\mu}\delta^{ij} - igA^{a}_{\mu}t^{aij} \tag{1.2}$$

and the field strength tensor, defined in terms of the gluon vector field A^{a}_{μ}, is given by

$$F^{a}_{\mu\nu} \equiv \partial_{\mu}A^{a}_{\nu} - \partial_{\nu}A^{a}_{\mu} + gf^{abc}A^{b}_{\mu}A^{c}_{\nu}\,, \tag{1.3}$$

in both of which summation of repeated indices is now implicit. The dimensionless parameter g is called the *strong coupling constant* and f^{abc} are called the *structure constants* of SU(3). The latter are real numbers and define the Lie algebra via the commutation relations of the generators of the Lie group, 3×3 complex matrices t^{a}_{ij} for the fundamental representation of SU(3),

$$\left[t^{a}, t^{b}\right] = if^{abc}t^{c}\,. \tag{1.4}$$

The gauge covariant derivative acts on the quark fields ψ, which are labelled with the flavour index $f \in \{d, u, s, c, b, t\}$. In pure QCD, quark

masses m_f distinguish the quark flavours and their spinor fields transform under local gauge transformations in the fundamental representation of the gauge group,

$$\psi^i_{f\alpha}(x) \to V^{ij}(x)\psi^j_{f\alpha}(x)\,. \tag{1.5}$$

In the SM, the quark masses are furnished by the Higgs-Yukawa coupling. The transformations $V_{ij}(x)$ can be parametrised via

$$V_{ij}(x) = \exp\left\{ i \sum_{a=1}^{3^2-1} \alpha^a(x) t^a_{ij} \right\} \tag{1.6}$$

with $\alpha^a(x)$ real. The gluon field, expressed as a matrix

$$A^\mu_{ij}(x) \equiv -ig \sum_{a=1}^{3^2-1} A^{\mu a}(x) t^a_{ij}\,, \tag{1.7}$$

transforms as

$$A_\mu(x) \to V(x)A_\mu(x)V^{-1}(x) + \frac{i}{g}[\partial_\mu V(x)]V^{-1}(x)\,. \tag{1.8}$$

The fermionic part of Equation (1.1) can be written in a compact fashion by the introduction of the Dirac operator for the quark flavour f

$$M_f = i\gamma^\mu D_\mu - m_f\,, \tag{1.9}$$

which can be extended to act on all quark flavours in the obvious block-diagonal fashion.

QCD Path Integral

For the purpose of this thesis, QCD as introduced above is quantised via the path integral formulation [28, 29]. The expectation value of an observable $\hat{\mathcal{O}}$ is given by the functional integral

$$\langle\hat{\mathcal{O}}\rangle = \frac{1}{Z}\int \mathcal{D}A\mathcal{D}\psi\mathcal{D}\bar{\psi}\,\mathcal{O}(\psi,\bar{\psi},A)\exp\left\{ i\int d^4x\,\mathcal{L}_{QCD}(x,\psi,\bar{\psi},A)\right\}, \tag{1.10}$$

where $\mathcal{L}_{QCD}$ is given by Equation (1.1) and Z is referred to as the *partition function*

$$Z = \int \mathcal{D}A\mathcal{D}\psi\mathcal{D}\bar{\psi}\,\exp\left\{ i\int d^4x\,\mathcal{L}_{QCD}(x,\psi,\bar{\psi},A)\right\}. \tag{1.11}$$

In perturbative treatments, the theory further requires gauge fixing and, generally, the subsequent introduction of ghost fields. Since this is not necessary in the computation of gauge-invariant quantities in the lattice discretisation of QCD, it will not be elaborated upon further here.

Renormalisation, Scale Dependence, Confinement and Asymptotic Freedom

Like for any quantum field theory, the classical Lagrangian of QCD given above and the related tree-level amplitudes with bare parameters g and m_f, do not give a complete picture. The parameters of the theory can be fixed by matching an appropriate number of observables

$$\mathcal{O}^i_{th} = \mathcal{O}^i_{exp}\,, \tag{1.12}$$

computed in the theoretical framework and measured in experiment. In perturbation theory, this matching would be to some truncated asymptotic series in the bare parameters while non-perturbatively, the matching could be between ratios of hadron masses, for example.

At the classical level, the theory appears to be valid *as is* at all energy scales, but this is not the case. When the system is probed at different distance scales, the parameters of the theory change, or equivalently, which degrees of freedom are relevant in that regime. In QED, for example, the electric charge and the electron mass can be matched to the theory in the low energy limit where they can be measured extremely precisely. But for the theory to be predictive at all distance scales, they acquire an energy dependence. In addition, perturbatively, loop corrections give rise to ultraviolet divergences which have to be removed through a regularisation procedure.

In the sense of renormalisation, these divergences can be thought of as being due to overcounting how high energy degrees of freedom contribute to the description of the system at a particular distance scale. Their removal implicitly or explicitly introduces an energy scale, the so-called renormalisation scale μ, on which physical observables cannot depend. As a consequence, regularisation and renormalisation become intimately related through renormalisation conditions. In the simplest sense, the divergences accounted for by the former are systematically absorbed into redefinitions of the coupling constant, quark masses and fields in the latter.

Enforcing the requirement that physical results cannot depend on the renormalisation scale leads to the Callan-Symanzik renormalisation group (RG) equations [30, 31]. These describe how the parameters of the theory absorb the energy dependence and hence how the theory changes with the energy scale. For the coupling constant, this is given by the so-called β-function

$$\mu\frac{\partial g}{\partial \mu} = \beta(g) = -\beta_0 g^3 - \beta_1 g^5 + \dots , \tag{1.13}$$

where β_0 and β_1 depend on the number of active quark flavours. In QCD, at least in the regime where perturbation theory is applicable, $\beta(g)$ is negative and the coupling constant decreases with increasing μ. For small g, keeping only the leading term and integrating Equation (1.13), gives

$$\alpha_s(\mu) = \frac{g^2(\mu)}{4\pi} = \frac{1}{4\pi \ln(\mu/\Lambda_{\mathrm{QCD}})}, \tag{1.14}$$

where Λ_{QCD} is known as the QCD scale parameter. It is interesting to note how even in pure Yang-Mills theory, a dimensionful scale Λ can emerge without dimensionful parameters in the Lagrangian. In the full theory then, all the quark masses can be expressed as ratios with respect to this scale.

At low energy, the coupling constant of QCD is large and this is probably responsible for quark confinement, although the existence of the required mass gap has not been rigorously demonstrated yet. At high energy, on the other hand, the coupling constant becomes small and this *asymptotic freedom* [14] ensures the validity of perturbation theory and makes QCD consistent with the parton model of hadrons in this regime. Because quarks and gluons confined in hadronic bound states are not suitable as asymptotic states (even approximately), the strong coupling limit can only be studied perturbatively through effective theories like chiral perturbation theory (χPT) [32]. In order to study the low energy properties of QCD in terms of the fundamental degrees of freedom, it is thus necessary to do so

non-perturbatively. The most developed and successful approach for this involves lattice discretisations of the continuum QCD action. As will be seen further below, the masses of strongly bound objects, for example, can then be computed directly from the path integral.

When the parameters or matrix elements in this lattice theory are to be related to their continuum analogues, they need to be renormalised before they can be matched to a continuum scheme such as $\overline{\text{MS}}$. The computation of the required renormalisation constants can be done in lattice perturbation theory (LPT), in which the lattice theory is expanded in terms of A fields around the continuum limit. Because the theory in terms of U fields is gauge invariant at non-zero lattice spacing, many new interaction vertices appear, the Feynman rules are complicated and systematic errors difficult to quantify. In addition, since the coupling is not small, LPT may fail to account for potentially large uncontrolled systematic uncertainties. It is much better, and common practice now, to compute these factors non-perturbatively. For some renormalisation constants, this may be achieved through the application of lattice Ward identities [33]. More generally it is done through further lattice simulations in regularisation independent momentum [34] (RI-MOM) or Schrödinger functional [35] (SF) schemes.

The conventional reference scales at which quark masses, for example, are expressed in the $\overline{\text{MS}}$ scheme generally differ from the hadronic scales at which they were computed non-perturbatively in the lattice theory. As a result, the parameters need to be evolved to the appropriate scale, which can again in principle be done by matching lattice and continuum perturbation theory, but the validity of either at the involved energy scales is not a certainty. Hence, it is much better to perform the running non-perturbatively too, for example through SF finite size scaling techniques [36, 35], so that the matching can be performed safely at very high scales.

QCD in Euclidean Space-time

Following the seminal work of Wilson [37] in trying to understand quark confinement directly in QCD, the theory can be studied non-perturbatively by discretising it through the introduction of a regular four dimensional lattice. Lattice QCD (LQCD) introduces a gauge-invariant ultraviolet cut-off through the inverse of the lattice spacing a, and the continuum theory can be defined as the quantum field theory obtained when this cut-off is removed. As will be discussed further below, the discretisation makes the number of degrees of freedom countable and the path integral potentially amenable to evaluation by computer. However, the rapidly oscillating phase factor $\exp(iS)$ means that all modes need to be taken into account in order to provide the correct weighting. In addition, even if assumptions can be made about some steepest descent contribution, the fine cancellations required for its correct evaluation cannot be accounted for on a machine with finite-precision arithmetic. This is an example of a *sign problem* which also occurs, for example, in the numerical study of condensed matter systems with many fermionic degrees of freedom [38]. In the case of LQCD it can be overcome by working instead from the point of view of a Euclidean field theory.

Under the conditions outlined in the Osterwalder-Schrader theorem [39, 40], a Euclidean field theory can be analytically continued to a quantum field theory in Minkowski space as a result of the analyticity of the respective Green's functions. To the present author's knowledge, the

required positivity conditions have not been demonstrated for general gauge theories, including QCD, but this will be tacitly ignored in the following.

Practically, the analytical continuation proceeds via the replacement $x^0 \to -ix^4, \eta^{\mu\nu} \to \delta^{\mu\nu}$ and yields a path integral with a real weight factor. Compared to the theory in Minkowski space, the time evolution operator $T = e^{-Hx^4}$ suppresses excited states and hence, correlation functions in this theory can be expected to be especially sensitive to low energy properties. The expectation value of an observable $\hat{\mathcal{O}}$ is given by

$$\langle \hat{\mathcal{O}} \rangle_E = \frac{1}{Z} \int \mathcal{D}A \mathcal{D}\psi \mathcal{D}\bar{\psi}\, \mathcal{O}(\psi, \bar{\psi}, A)\, e^{-S^E_{QCD}[A,\psi,\bar{\psi}]}\,, \tag{1.15}$$

where the ψ stand for all quark flavours and the action is

$$S^E_{QCD}[\bar{\psi}, \psi, A] = \bar{\psi} M \psi + \frac{1}{4} \int d^4x\, F^a_{\mu\nu}(x) F^a_{\mu\nu}(x) \tag{1.16}$$

$$= \bar{\psi} M \psi - \frac{1}{2g^2} \operatorname{Tr} F^2[A]\,. \tag{1.17}$$

The integrals over the fermionic degrees of freedom can be carried out formally, resulting in the functional fermionic determinant $\det M[A]$ and thus yielding an effective action

$$S_{\text{eff}}[A] = \log \det M[A] - \frac{1}{2g^2} \operatorname{Tr} F^2[A]\,. \tag{1.18}$$

This reduces Equation (1.15) to a functional integral over A only.

LATTICE QCD

As will be seen below, Euclidean QCD *is* amenable to computer simulation such that the path integral can be evaluated directly and non-perturbatively. This section largely follows the standard development presented in Ref. [41], to which the reader is referred for a more detailed exposition. First, the Euclidean theory is discretised by restricting the support of the quark fields to a regular four-dimensional lattice with lattice spacing a, which will serve as an explicit *non-perturbative* ultraviolet regulator. On a lattice of finite volume L^4, momenta are thus restricted to

$$p = \frac{2\pi k}{La}\,, \qquad k = \{k_1, k_2, k_3, k_4\}, \qquad \pm k_\mu \in \{1, 2, \cdots, L/2\}. \tag{1.19}$$

The finite lattice is taken to be periodic in all directions and anti-periodic boundary conditions are enforced for the quark fields in the direction identified with time. The gauge field is introduced via line integrals of the gauge potential $A_\mu(x)$

$$U_\mu(x) \equiv U(x, x + a\hat{\mu}) = \exp\left\{ -\int_x^{x+a\hat{\mu}} dx_\mu\, A_\mu(x) \right\} \tag{1.20}$$

$$\simeq \exp\left\{ -a A_\mu(x + \frac{a}{2}\hat{\mu}) \right\}, \tag{1.21}$$

where no index summation was implied in the first line. For convenience of notation, $\hat{\mu}$ defines the unit four-vector pointing in direction μ. The $U_\mu(x)$ act as parallel transporters and are also referred to as *link variables*. They are elements of $SU(3)$ and transform as

$$U_\mu(x) \to V(x) U_\mu(x) V^\dagger(x + a\hat{\mu})\,, \tag{1.22}$$

where $V(x)$ are defined as in Equation (1.5). They further satisfy the property

$$U^{\dagger}(x, x + a\hat{\mu}) = U(x + a\hat{\mu}, x) = U_{-\mu}(x + a\hat{\mu}) . \tag{1.23}$$

Expanding Equation (1.21) to $\mathcal{O}(a)$ gives

$$U_{\mu}(x) = 1 - aA_{\mu}(x) + \mathcal{O}(a^2) . \tag{1.24}$$

A particular product of link variables in the plane spanned by $\hat{\mu}$ and $\hat{\nu}$, the *plaquette*

$$U_{\mu\nu}(x) = U_{\mu}(x)U_{\nu}(x + a\hat{\mu})U^{\dagger}_{\mu}(x + a\hat{\nu})U^{\dagger}_{\nu}(x) , \tag{1.25}$$

can be used to discretise the field strength tensor. The gauge part of the lattice action is then expressed as

$$S_G[U] = \frac{6}{g^2} \sum_x \sum_{\mu<\nu} \left\{ 1 - \frac{1}{6} \operatorname{Tr} \left[U_{\mu\nu}(x) + U^{\dagger}_{\mu\nu}(x) \right] \right\} , \tag{1.26}$$

where the inverse squared gauge coupling is defined as $\beta = 6/g^2$ to reproduce the form of the Yang-Mills gauge action in the continuum limit. It is straightforward to show that the trace of the plaquette is gauge invariant and that

$$\operatorname{Tr} \left[U_{\mu\nu}(x) + U^{\dagger}_{\mu\nu}(x) \right] = 2 \operatorname{Tr} \mathbb{1} + \frac{a^4}{2} \operatorname{Tr} \left[F_{\mu\nu}(x) F_{\mu\nu}(x) \right] + \mathcal{O}(a^6) . \tag{1.27}$$

The constant summand in Equation (1.26) is often dropped in practice because it has no impact on the dynamics or the expectation values of observables. It is intersting to remark that the positivity of physical states and the existence of a positive self-adjoint transfer matrix[1] for this discretisation of Yang-Mills theory have been demonstrated [42]. Further, the existence of a mass gap in the infinite volume limit of the lattice theory has also been shown [43] at strong coupling.

More complicated gauge actions can be constructed with the aim of reducing lattice artefacts. A particular choice which will be used in this thesis is the so-called *Iwasaki* [44] gauge action

$$S_G^{\text{Iwasaki}} = \beta \sum_x \Big[\sum_{\mu<\nu} b_0 \{ 1 - \frac{1}{3} \Re \operatorname{Tr} U_{\mu\nu}(x) \} + \sum_{\mu \neq \nu} b_1 \{ 1 - \frac{1}{3} \Re \operatorname{Tr} R_{\mu\nu}(x) \} \Big] , \tag{1.28}$$

where $R_{\mu\nu}(x)$ are rectangular plaquettes of size $2a$ by a and the parameters b_0 and b_1 take the values

$$b_0 = 1 - 8b_1 , \qquad b_1 = -0.331 . \tag{1.29}$$

The reasons for this choice of gauge action will be discussed in more detail in Section 2.2.

Naïve Lattice Fermions

Covariant lattice forward and backward derivatives acting on the lattice quark fields can now be defined as first order finite difference operators with appropriate parallel transporters

$$\nabla^{f}_{\mu}\psi(x) = a^{-1} \left[U_{\mu}(x)\psi(x + a\hat{\mu}) - \psi(x) \right] \tag{1.30}$$

$$\nabla^{b}_{\mu}\psi(x) = a^{-1} \left[\psi(x) - U^{\dagger}_{\mu}(x - a\hat{\mu})\psi(x - a\hat{\mu}) \right] . \tag{1.31}$$

[1] it has not been shown for the improved theory

The locally gauge invariant, so-called *naïve* fermion action is then given by

$$S_F[\psi,\bar{\psi},U] = a^4 \sum_x \bar{\psi}(x)\left[\frac{1}{2}\gamma_\mu(\nabla^f_\mu + \nabla^b_\mu) + m\right]\psi(x)\,. \tag{1.32}$$

Wilson Quarks

Unfortunately, in this form the theory exhibits the problem of fermionic species doubling as the lattice quark propagator has one pole at every corner of the Brouillin zone, leading to the propagation of 16 quarks in the continuum limit for each quark flavour. One way to cure this is to introduce the *Wilson term*, $-ar\nabla^f_\mu\nabla^b_\mu$, into the lattice action. Although this operator is naïvely irrelevant under renormalisation (it vanishes in the continuum limit), it furnishes the additional species with a divergent mass $2r/a$, thus decoupling them in the continuum limit. Setting $r = 1$ defines the Wilson Dirac operator as it is usually used in calculations with Wilson quarks

$$D_W(U) = \frac{1}{2}\left[\gamma_\mu(\nabla^f_\mu + \nabla^b_\mu) - a\nabla^f_\mu\nabla^b_\mu\right]. \tag{1.33}$$

It satisfies the property of γ_5-Hermiticity

$$D^\dagger_W = \gamma_5 D_W \gamma_5\,. \tag{1.34}$$

The action thus constructed is

$$S_W[U,\bar{\psi},\psi] = \sum_x \left\{a^4\bar{\psi}(x)\left[D_W[U] + m_W\right]\psi(x) + S_G[U(x)]\right\}. \tag{1.35}$$

A graphical representation of the plaquette, the lattice spinor fields and parallel transporters is shown in Figure 1.2. In practice, the fermionic part of the action employed in computer simulations is expressed in terms of the *hopping parameter* κ by expanding the lattice forward and backward derivatives, giving

$$S_F^{\text{hop}} = \sum_x \left\{\bar{\psi}(x)\psi(x) - \kappa\bar{\psi}(x)\sum_{\mu=1}^{4}\left[U_\mu(x)(1+\gamma_\mu)\psi(x+a\hat{\mu}) + U^\dagger_\mu(x-a\hat{\mu})(1-\gamma_\mu)\psi(x-a\hat{\mu})\right]\right\}. \tag{1.36}$$

The lattice spacing is absorbed completely into the field normalisations

$$\kappa = \frac{1}{2am_W + 8}\,, \qquad \frac{\sqrt{2\kappa}}{a^{3/2}}\psi\,, \qquad \frac{\sqrt{2\kappa}}{a^{3/2}}\bar{\psi}\,, \tag{1.37}$$

and it should be noted that a is not an input parameter. It can only be found à posteriori from a so-called *scale setting* procedure.

At non-zero lattice spacing, the Wilson term acts like a mass term (albeit not a divergent one) also for the remaining quark and breaks chiral symmetry even in the limit of $m_W \to 0$. This produces a leading lattice artefact of $\mathcal{O}(a)$ and leads to additive renormalisation m_{crit} of the quark mass on top of multiplicative renormalisation. The bare quark mass is thus given by

$$m_q = m_W - m_{\text{crit}}\,. \tag{1.38}$$

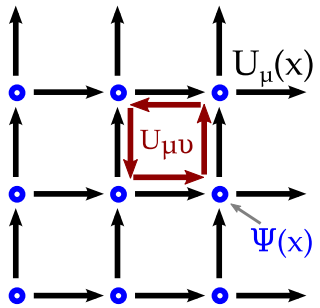

Figure 1.2: Graphical representation of the Wilson lattice action of Equations (1.26) and (1.35). Lattice spinors are shown in blue, link variables in black and a plaquette in dark red.

As will be discussed further below, computations using standard Wilson fermions generally require the implementation of involved *improvement* programmes in which special counterterms are added to the lattice action and composite operators to cancel various discretisation effects. The coefficients of these counterterms need to be computed by defining appropriate improvement conditions either in perturbation theory or, preferably, non-perturbatively.

It should be noted at this point that the discretisation of the action is not unique and different possibilities differentiate themselves through their symmetries (or the breaking thereof). These affect, for example, how the lattice theory behaves in the approach to the continuum limit or how demanding it is to simulate on a computer. In the following, however, only quark discretisations of the Wilson type will be discussed.

Lattice Path Integral

The lattice discretisation gives definite meaning to the notion of path integrals through

$$\int \mathcal{D}A \to \int \prod_{x,\mu} dU_\mu(x)\,, \quad \int \mathcal{D}\bar{\psi}(x)\mathcal{D}\psi(x) \to \int \prod_x d\bar{\psi}(x)\, d\psi(x)\,. \tag{1.39}$$

Functional integration turns into regular integration over a countable number of variables, a finite number in the case of finite volume. Unlike the integration measure over the A fields used in the continuum, the *Haar* measure [45] $dU_\mu(x)$ is fully gauge-invariant. The finite-dimensional path integral for an observable $\hat{\mathcal{O}}$ is then given by

$$\langle\hat{\mathcal{O}}\rangle = \frac{1}{Z}\int \prod_x \left[d\bar{\psi}(x)\, d\psi(x) \prod_\mu dU_\mu(x)\right] \mathcal{O}(U,\bar{\psi},\psi)\, e^{-S_W[U,\bar{\psi},\psi]}\,. \tag{1.40}$$

In order to evaluate the path integral numerically, one proceeds as follows. As in the continuum, the fermionic integrals can be carried out formally, yielding

$$\langle\hat{\mathcal{O}}\rangle = \frac{1}{Z}\int \prod_x \prod_\mu dU_\mu(x)\, \mathcal{O}(U, D^{-1}[U])\, \det(D[U])\, e^{-S_G[U]}\,, \tag{1.41}$$

where $D^{-1}[U]$ represents the quark propagators resulting from Wick contractions of the quark fields in the observable, if any. Again as in the continuum, the determinant contribution can be re-expressed as an integral over complex scalar fields ξ, with fermionic statistics, such that

$$\langle\hat{\mathcal{O}}\rangle = \frac{1}{Z}\int \prod_x \left[d\xi^\dagger(x) d\xi(x) \prod_\mu dU_\mu(x)\right] \mathcal{O}(U, D^{-1}[U]) \cdot \exp\left\{-S_G[U] - \xi^\dagger (D[U])^{-1}\, \xi\right\}\,,$$

but it should be noted that the distribution of the fields ξ is not straightforward to sample. Still, the quark determinant does not need to be evaluated directly anymore and the Grassmann variables and their combinatorics have been eliminated. However, D is a sparse matrix so large that it has to be computed on the fly, because it cannot be stored in memory for all but the smallest systems. As a result of this size, the computation of its inverse, which now appears in the action (and in the operator), is one of the most

demanding parts of a computation in LQCD. This can be achieved for a fixed right hand side by solving $Dx = b$ through the usage of various iterative subspace solvers. More approximately, the entire D^{-1} can be estimated stochastically by doing such inversions for a collection of vectors of random numbers.

For an action with two mass degenerate quarks, using the γ_5-hermiticity of D, the determinant factor

$$(\det D)^2 = \det\left(D^2\right) = \det(\gamma_5 D^\dagger \gamma_5 D) = \det(Q^\dagger Q)\,,$$

can be expressed directly in terms of just one ξ field

$$\det(Q^\dagger Q) = \int \prod_x d\xi^\dagger(x) d\xi(x)\, \exp\left\{-\xi^\dagger \frac{1}{Q^\dagger Q}\xi\right\}\,,$$

where $Q = \gamma_5 D$ is hermitian. An approximation to the ξ fields can now be given in terms of *pseudofermion* fields ϕ, which can be generated from complex random numbers R sampled from a Gaussian distribution, such that

$$\langle R^\dagger R\rangle \sim 1\,, \qquad \xi \sim \phi = QR\,,$$

and the ϕ have approximately the correct statistical distribution. Using K of these fields, this finally yields

$$\langle\hat{\mathcal{O}}\rangle = \frac{1}{Z}\int \prod_x \prod_\mu dU_\mu(x)\, \mathcal{O}(U, D^{-1}[U]) \cdot \exp\left\{-S_G[U] - \frac{1}{K}\sum_{k=1}^{K}\left[\phi_k^\dagger \frac{1}{Q^\dagger Q}\phi_k\right]\right\} + \mathcal{O}(1/\sqrt{K})\,, \quad (1.42)$$

such that that the fermionic contribution is now evaluated statistically as long as the inverse of $Q^\dagger Q$ can be computed. As will be discussed in Section 1.4, the average over pseudofermion fields is usually performed implicitly and not directly as indicated above. For actions with mass non-degenerate quarks or single quark flavours, the pseudofermion representation can also be used, but the construction has to rely on polynomial or rational approximations for the operator, the technical details of which are beyond the scope of the current exposition.

Although the integral can be carried out now in principle, a quick calculation shows that for a usual system, the number of degrees of freedom is still staggering and a direct evaluation is out of reach of current computational capacity. However, its form makes it amenable to a statistical interpretation over ensembles of field configurations $\{U\}$ which contribute with a Boltzmann weight $Z^{-1}\exp(-S_{\text{eff}}[U])$. Appropriate methods for the generation of these ensembles with their correct probability distribution

$$P(U) \propto \det D[U] \exp\left(-S_G[U]\right)$$

will be presented in Section 1.4. Even with these algorithms, the computational challenge of LQCD is formidable because multiple scales have to be respected simultaneously. In addition, the cost of computing D^{-1} depends inversely on the quark mass, one of the main reasons why simulations at the physical average light quark mass have only recently become possible. On current supercomputer architectures, the absolute maximum number of lattice sites that can realistically be simulated is around $V/a^4 = (L/a)^3 \cdot (T/a) =$

$100^3 \cdot 200 = 2 \cdot 10^8$. For simulations with physically light quarks, finite size effects are considered well under control when $M_\pi \cdot L \gtrsim 4$. In order to control discretisation errors, the lattice spacing should be taken as small as possible, especially when heavy quark flavours are to be included in the simulation.

However, when the lattice spacing is below around 0.04 fm, the system is close to criticality which manifests itself through long autocorrelation times. In the worst case, the algorithms even get stuck in topological sectors [46], putting the importance sampling of the path integral into question [2]. This property is called *critical slowing down* and one of its consequences is that longer simulations are required at small values of the lattice spacing. Fine lattice spacings also mean that the number of lattice points needs to be increased to compensate for the reduced size in physical units, making simulations more expensive. Finally, since all dimensionful quantities enter with an appropriate power of the lattice spacing, fine lattices also mean more expensive inversions.

As a result, the challenge of LQCD simulations is to approximately satisfy the following inequalities (with m_c the charm quark mass):

$$100a \gtrsim L \gtrsim \frac{4\hbar c}{M_\pi} > \frac{\hbar c}{m_c} > a > 0.04 \text{ fm}\,, \tag{1.43}$$

with multiple simulations at different L and a, such that the continuum and infinite volume limits can be taken. For $a \sim 0.1$ fm and $M_\pi \sim 130$ MeV, this means that $L/a \sim 60$. Hence, if the lattice spacing is to be halved to get a good estimate of discretisation effects, the maximum number of lattice points noted above will be exceeded.

The technical difficulties discussed above mean that the algorithms used for the computation of observables and the generation of gauge configurations are the subject of intense study. On the one hand, improved schemes for the computation of the matrix inverses and the sampling of the gauge ensembles are sought to reduce the overall computational effort. On the other hand, existing techniques and routines are meticulously optimised for particular supercomputer architectures and it is often the case that LQCD codes are at the cutting edge of exploiting these machines. Optimisations of this latter kind are presented in Chapters 8 and 10.

Continuum Limit

To end this section, it is insightful to look at the continuum limit of the lattice theory in terms of the bare coupling g. For simplicity this will be done considering only massless quarks, closely following the exposition in Ref.[41]. Physical quantities $\mathcal{P}(a, g)$ must become independent of the lattice spacing a as the continuum limit is taken. They thus satisfy the renormalisation group equation

$$\left[-a\frac{\partial}{\partial a} + \beta_{\text{LAT}}(g)\frac{\partial}{\partial g}\right]\mathcal{P}(a, g) = 0 + \mathcal{O}(a)\,, \tag{1.44}$$

[2] Simulations of LQCD with open boundary conditions in the time direction make it possible to solve this issue, at the cost of introducing a number of technical complications. [47]

where lattice artefacts have explicitly been allowed for via the term on the right hand side. β_{LAT} is the lattice β-function, integration of which defines Λ_{LAT} and gives the dependence of the lattice spacing on the bare coupling

$$a = \frac{1}{\Lambda_{\mathrm{LAT}}} (\beta_0 g^2)^{-\beta_1/(2\beta_0^2)} \exp\left\{-\frac{1}{2\beta_0 g^2}\right\} \left[1 + \mathcal{O}(g^2)\right] . \tag{1.45}$$

The coefficients β_0 and β_1 are universal, but Λ_{LAT} as well as higher order terms depend on the details of the lattice discretisation. It should be noted that this perturbative result is only a guide and the lattice spacing of given simulation is determined non-perturbatively through a scale-setting procedure. As an example, ratios of hadron masses can be computed in LQCD and the bare quark masses adjusted such that the ratios take their phenomenological values, fixing the bare quark masses. At this point in parameter space, one of the hadron masses can be compared to its experimental value to extract an estimate of the lattice spacing

$$\begin{aligned} aM_{\mathrm{lat}} &= a\left(M_{\mathrm{exp}} + \delta M\right) \\ a_{\mathrm{est.}} &\sim \frac{aM_{\mathrm{lat}}}{M_{\mathrm{exp}}} \end{aligned} \tag{1.46}$$

where δM is symbolic and represents all possible lattice artefacts. There are many methods for setting the scale and the running of $a(g^2)$ can also be computed non-perturbatively, but in any case the continuum limit is approached as $g \to 0$ or equivalently $\beta = 6/g^2 \to \infty$.

TWISTED MASS LATTICE QCD

Twisted mass lattice QCD (tmQCD) is an equivalent formulation of QCD which involves the addition of a mass term which is twisted in flavour space. In the continuum, the twisted formulation has the same properties as the standard formulation. However, the corresponding lattice theory exhibits a number of properties which make it interesting for practical non-perturbative calculations.

In this section, the continuum and lattice actions will be discussed, including an action for mass non-degenerate quark flavours. Automatic $\mathcal{O}(a)$-improvement will be demonstrated and the addition of a so-called *clover* term will be discussed briefly. The tmLQCD action in its familiar form was first used for lattice simulations with Wilson fermions to study the phase structure of the Wilson lattice theory in Ref. [48], where the twisted mass term was used to provide for explicit breaking of parity and flavour symmetries. In its modern form, the discretisation was first explored in the quenched approximation as a tool to avoid unphysical zero modes at small quark masses [49]. Quenched and dynamical twisted mass lattice QCD (tmLQCD) were then further studied in Refs. [50, 51, 52] until in Refs. [53, 54] its crucial properties were elaborated. Reviews can be found in Refs. [55, 56].

Continuum Twisted Mass QCD

The continuum QCD action for two degenerate quark flavours of mass m with fermionic Lagrangian density

$$\mathcal{L}_{N_f=2}(x) = \bar{\psi}(x)\left[\gamma_\mu D_\mu + m\right]\psi(x) , \tag{1.47}$$

is invariant under the global chiral rotation with *twist angle* ω

$$\psi(x) \to \chi(x) = e^{-i\omega/2\gamma_5\tau^3}\psi(x) , \qquad \bar{\psi}(x) \to \bar{\chi}(x) = \bar{\psi}(x)e^{-i\omega/2\gamma_5\tau^3} , \tag{1.48}$$

if the mass parameter is simultaneously transformed

$$m \to m \exp(i\omega\gamma_5\tau^3)\,. \tag{1.49}$$

To show this, it suffices to use the properties of the third Pauli matrix τ^3 acting in flavour space and γ_5 acting in Dirac space. Instead, if the action is expressed in terms of the $\chi(x)$ fields defined above, the twisted mass QCD (tmQCD) Lagrangian density for two mass degenerate flavours is obtained

$$\mathcal{L}_{F,tm}(x) = \bar{\chi}(x)\left[\not{D} + m_q + i\mu\gamma_5\tau^3\right]\chi(x)\,, \tag{1.50}$$

where the identifications

$$m_q = m\cos(\omega)\,, \qquad \mu = m\sin(\omega)\,, \qquad m = \sqrt{m_q^2 + \mu^2}\,, \tag{1.51}$$

have been made. This form is completely equivalent to the standard one, except that certain discrete symmetries look more complicated. The mass parameter μ is referred to as the *twisted mass*, m_q as the standard mass and m as the *polar* mass. Expressions in terms of the $\chi(x)$ fields are also referred to as being written in the *twisted basis*. The relationship between the rotation angle and the mass terms can also be inverted, defining

$$\omega = \arctan(\mu/m_q)\,. \tag{1.52}$$

In addition to the trivial case $\omega = 0$, a special value $\omega = \pi/2$ can be identified for which the quark mass is given entirely by the twisted mass. As will be seen further below, this particular value is referred to as *maximal twist* and is related to the property of automatic $O(a)$-improvement in the lattice theory.

Quark Currents

As given in Ref. [49], the usual isospin currents and densities can be defined in the twisted basis

$$A_\mu^a = \bar{\chi}\gamma_\mu\gamma_5\frac{\tau^a}{2}\chi\,, \tag{1.53}$$

$$V_\mu^a = \bar{\chi}\gamma_\mu\gamma_5\frac{\tau^a}{2}\chi\,, \tag{1.54}$$

$$P^a = \bar{\chi}\gamma_5\frac{\tau^a}{2}\chi\,. \tag{1.55}$$

The PCAC and PCVC relations take the form

$$\partial_\mu A_\mu^a = 2mP^a\,, \qquad a = 1,2\,, \tag{1.56}$$

$$\partial_\mu A_\mu^3 = 2mP^3 + i\mu\bar{\chi}\chi\,, \tag{1.57}$$

$$\partial_\mu V_\mu^a = -2\mu\epsilon_{ab}P^b\,, \qquad a,b = 1,2\,, \tag{1.58}$$

$$\partial_\mu V_\mu^3 = 0\,. \tag{1.59}$$

The currents in the usual basis[3] can be obtained by undoing the chiral rotation via

$$\psi(x) = e^{i\omega/2\gamma_5\tau^3}\chi(x)\,, \qquad \bar{\psi}(x) = \bar{\chi}(x)e^{i\omega/2\gamma_5\tau^3}\,, \tag{1.60}$$

[3] often called *physical* basis

with the twist angle as defined in Equation (1.52). This gives then

$$\tilde{A}_\mu^{1,2} = \cos(\omega) A_\mu^{1,2} + \sin(\omega)\epsilon_{12,21} V_\mu^{2,1}\,, \tag{1.61}$$
$$\tilde{A}_\mu^{3} = A_\mu^{3}\,, \tag{1.62}$$
$$\tilde{V}_\mu^{1,2} = \cos(\omega) V_\mu^{1,2} + \sin(\omega)\epsilon_{12,21} A_\mu^{2,1}\,, \tag{1.63}$$
$$\tilde{V}_\mu^{3} = V_\mu^{3}\,, \tag{1.64}$$
$$\tilde{P}^{1,2} = P^{1,2}\,, \tag{1.65}$$
$$\tilde{P}^{3} = \cos(\omega) P^{3} + \frac{i}{2}\sin(\omega)\bar{\chi}\chi\,, \tag{1.66}$$

for which the PCAC and PCVC relations take their usual form.

Lattice Twisted Mass QCD

Applying the chiral transformation to the fermionic part of the Wilson lattice action in Equation (1.35), results in a change of the Wilson Dirac operator

$$D_W \to D'_W = \frac{1}{2}\gamma_\mu\left(\nabla_\mu^f + \nabla_\mu^b\right) - ar e^{-i\omega\gamma_5\tau^3}\nabla_\mu^f\nabla_\mu^b\,. \tag{1.67}$$

The form D'_W is interesting in that the Wilson term is now twisted in flavour space and the invariance under Equations (1.48) and (1.49) is only recovered in the continuum limit. The properties of twisted mass lattice QCD (tmLQCD) can be thought of as originating from this twisting of the Wilson term. Alternatively, as in the continuum, the lattice action can instead be written in the twisted basis, giving

$$\begin{aligned} S_{tm} &= a^4 \sum_x \bar{\chi}(x)\left[D_W[U] + m_W + i\mu\gamma_5\tau^3\right]\chi(x) \\ &= a^4 \sum_x \bar{\chi}(x) D_{tm}[U]\chi(x) \end{aligned} \tag{1.68}$$

without changing D_W.

In the lattice theory, the twist angle is not uniquely defined anymore. One possible definition, ω_A, is given in terms of the PCAC quark mass, computed from a matrix element of the PCAC relation

$$Z_A\partial_\mu A_\mu^{1,2} = 2m_{PCAC}P^{1,2}\,, \qquad \omega_A = \arctan\left(\frac{\mu}{Z_A\, m_{PCAC}}\right)\,. \tag{1.69}$$

The renormalised masses are given by

$$m_R = Z_S^{-1}(m_W - m_{crit}) \quad \text{and} \quad \mu_R = Z_P^{-1}\mu\,, \tag{1.70}$$

defining the renormalised twist angle

$$\omega_R = \arctan\left(\frac{\mu_R}{m_R}\right)\,. \tag{1.71}$$

Other definitions of the twist angle are possible and will have different lattice artefacts, as discussed further below. It should also be noted that lattice artefacts can make a reliable determination of the twist angle difficult close to the chiral limit, which will be discussed in Section 2.2.

The case of $\omega = \pi/2$ has a special significance for the lattice theory. Firstly, the theory is invariant under flavour exchange and a simultaneous change of the sign of μ. This precludes additive renormalisation of the twisted

quark mass and so, when $m_W \to m_{crit}$, the Lagrangian quark mass μ is related to the physical quark mass by multiplicative renormalisation only (up to $\mathcal{O}(a^2)$). Secondly, the counterterms required to provide on-shell $\mathcal{O}(a)$-improvement (see Ref. [51]) for general values of ω become such that they affect physical results at most at $\mathcal{O}(a^2)$. These terms are thus irrelevant at $\mathcal{O}(a)$ in the sense of improvement (see Ref. [53]). This property is called *automatic* $\mathcal{O}(a)$*-improvement*.

It should be noted that while the continuum action of twisted mass QCD is completely equivalent to its usual form, the lattice theory is not. The twisted mass term explicitly breaks parity and flavour SU(2) symmetries present in the standard two flavour Wilson lattice theory. These are lattice artefacts of $\mathcal{O}(a^2)$ and the symmetries are restored in the continuum limit. In general, the effects are found to be very small with the exception of an unphysical mass difference between the neutral and charged pions which can be quite sizeable. This point will be discussed in more detail in Section 2.2.

Automatic $\mathcal{O}(a)$-improvement and the Sheikholeslami-Wohlert Term

Standard Wilson fermions and twisted mass fermions for general values of the twist angle ω exhibit discretisation artefacts at $\mathcal{O}(a)$, such that taking the continuum limit reliably may require very fine lattice spacings. Discretisation errors can be described in a systematic way using the Symanzik effective theory [25, 26]. Under the assumption that the lattice spacing is sufficiently fine that lattice artefacts appear only as asymptotically small corrections [55], the lattice expectation value of a multi-local operator $\hat{\mathcal{O}}$ without contact terms can be expressed as a path integral in the effective continuum theory

$$\langle \hat{\mathcal{O}} \rangle_{lat} = \frac{1}{Z} \int \mathcal{D}\bar{\psi}\mathcal{D}\psi\mathcal{D}A \, \mathcal{O}_{eff} \, e^{-S_{eff}} \tag{1.72}$$

$$S_{eff} = S_0 + aS_1 + a^2 S_2 + \cdots \tag{1.73}$$

$$\mathcal{O}_{eff} = \mathcal{O}_c + a\mathcal{O}_1 + a^2\mathcal{O}_2 + \cdots \tag{1.74}$$

Here, S_0 is the action of the target continuum theory (tmQCD in this case) and the higher order terms are linear combinations of operators of mass dimension five and higher. The dependence on the lattice spacing is now explicit and the notation $\mathcal{O}_c$ indicates that this operator is constructed out of continuum fields. The higher order terms in $\mathcal{O}_{eff}$ depend on the operator in question.

Expanding $\exp(-S_{eff})$ and $\mathcal{O}_{eff}$ simultaneously to $\mathcal{O}(a)$, gives

$$\langle \hat{\mathcal{O}} \rangle_{lat} = \langle \hat{\mathcal{O}}_c \rangle_{S_0} - a\langle S_1 \hat{\mathcal{O}}_c \rangle_{S_0} + a\langle \hat{\mathcal{O}}_1 \rangle_{S_0} + \mathcal{O}(a^2) \,, \tag{1.75}$$

where the expectation values are calculated with respect to the target continuum action S_0. In the standard Wilson lattice theory, S_1 can be shown to reduce to the Pauli term

$$S_1^{sw} = \int d^4x \, ic\bar{\psi}(x)\sigma_{\mu\nu}F_{\mu\nu}(x)\psi(x) \,, \tag{1.76}$$

where $\sigma_{\mu\nu} = i/2[\gamma_\mu, \gamma_\nu]$. $\mathcal{O}(a)$ discretisation artefacts in on-shell observables [57] can be removed by the introduction of a counterterm of the same form, the so-called Sheikholeslami-Wohlert (*clover*) term [58] into the lattice action. To do this, the coefficient function $c_{sw}(g^2)$ has to be tuned according

to an appropriate improvement scheme. For more general matrix elements, further operator-specific counterterms of the form $\hat{\mathcal{O}}_1$ need to be introduced and their coefficients tuned.

The property of automatic $\mathcal{O}(a)$-improvement of twisted mass Wilson fermions was first shown in Ref. [53]. A very simple alternative proof is given in in Refs. [59, 55] which does not rely on any properties of the critical mass m_{crit} which may be affected by lattice artefacts. When the Wilson quark mass is tuned to its critical value and $m_q \to 0$, the continuum action S_0 becomes invariant under the transformation

$$T_1 : \quad \chi \to i\gamma_5\tau^1\chi\,, \qquad \bar{\chi} \to \bar{\chi}i\gamma_5\tau^1\,, \tag{1.77}$$

while S_1 changes sign under T_1. Observables $\hat{\mathcal{O}}$ too can be odd or even under this transformation. For odd observables, the $\mathcal{O}(1)$ contribution $\langle\hat{\mathcal{O}}_c\rangle_{S_0}$ must vanish because because it is T_1-odd while S_0 is not. The behaviour of the $\mathcal{O}(a)$ contributions is, on the other hand

$$a\langle\hat{\mathcal{O}}_c S_1\rangle_{S_0} \to a\langle\hat{\mathcal{O}}_c S_1\rangle_{S_0}\,, \qquad a\langle\hat{\mathcal{O}}_1\rangle_{S_0} \to a\langle\hat{\mathcal{O}}_1\rangle_{S_0}\,, \tag{1.78}$$

and so

$$\langle\hat{\mathcal{O}}^{\text{odd}}\rangle_{\text{lat}} = -a\langle\hat{\mathcal{O}}_c S_1\rangle_{S_0} + a\langle\hat{\mathcal{O}}_1\rangle_{S_0} + \mathcal{O}(a^2)\,. \tag{1.79}$$

Hence, $\langle\hat{\mathcal{O}}^{\text{odd}}\rangle_{\text{lat}}$ vanishes like a in the continuum limit and is thus not $\mathcal{O}(a)$-improved. For physical observables which are even under T_1, on the other hand, the $\mathcal{O}(a)$ contributions must vanish by symmetry since they are T_1-odd overall while the $\mathcal{O}(1)$ contribution is even by definition. As a result, the lattice artefact terms in the Symanzik expansion for these observables start at $\mathcal{O}(a^2)$

$$\langle\hat{\mathcal{O}}^{\text{even}}\rangle_{\text{lat}} = \langle\hat{\mathcal{O}}_c\rangle_{S_0} + \mathcal{O}(a^2)\,, \tag{1.80}$$

and these physical observables are automatically $\mathcal{O}(a)$-improved if the Wilson quark mass has been tuned to its critical value. In fact, it can be shown that all odd powers of a are absent in the Symanzik expansion of T_1-even observables at maximal twist.

In practice this means that when twisted mass quarks at maximal twist are used, there is no need for the involved improvement programmes required for standard Wilson quarks. This is especially true when observables with many quark fields are considered for which the pattern of counterterms can quickly become very difficult to handle. As a final remark it can be noted that the proof of automatic $\mathcal{O}(a)$-improvement has recently been extended to observables with contact terms in Refs. [60, 61].

Tuning Conditions

In the lattice theory, m_{crit} (and hence the twist angle) is not defined uniquely: different observables which vanish as $m_W \to m_{\text{crit}}$ will be subject to different lattice artefacts. Assuming that the behaviour of the theory as $m_q \to 0$ is smooth[4], the vanishing of the pion mass at $\mu = 0, m_q \to 0$ can in principle be used as a tuning condition for imposing maximal twist. However, it has been shown in Refs. [62, 63] that this results in substantial $\mathcal{O}(a^2)$ discretisation artefacts as small quark masses are approached along the line $\mu \to 0$. One way of understanding this is via the $\mathcal{O}(a)$ contribution S_1^μ, which arises in addition to S_1^{sw} at non-zero twisted quark mass

$$S_1^\mu = \int d^4x\, b_\mu \mu^2 \bar{\chi}(x)\chi(x)\,. \tag{1.81}$$

[4] deviations will be dicussed in Section 2.2

This can be interpreted as a μ-dependent $\mathcal{O}(a^2\Lambda^3)$ shift in m_{crit}, to which the twist angle becomes more and more sensitive as μ is decreased (even though the shift itself of course also becomes smaller). Consequently,

$$m_{\text{crit}}(M_\pi^\pm = 0, \mu = 0) \neq m_{\text{crit}}(\mu)\,, \tag{1.82}$$

and if the small quark mass limit is approached without taking the shift into account, large $\mathcal{O}(a^2)$ lattice artefacts appear. In the literature it is therefore said that the vanishing of the pion mass at $\mu = 0, m_q = 0$ is not an *optimal* tuning condition.

On the other hand, when the small quark mass limit is taken with the shift taken into account, no such increase of lattice artefacts is seen as μ is decreased. This can be achieved through a tuning condition which aims at the restoration of T_1 symmetry, provided by the vanishing of some T_1-odd lattice observable up to lattice artefacts. Conventionally this is a PCAC quark mass, extracted from the ratio of matrix elements

$$m_{\text{PCAC}} = \frac{\sum_{\mathbf{x}}\langle \partial_4 A_4^a(\mathbf{x},t) P^a(0)\rangle}{2\sum_{\mathbf{x}}\langle P^a(\mathbf{x},t) P^a(0)\rangle} \qquad a \in \{1,2\}\,, \tag{1.83}$$

where no renormalisation is necessary for this purpose. It should be noted that this holds for any t, but lattice artefacts reduce the range of t for which m_{PCAC} can be extracted cleanly. It is also possible to define the twist angle in terms of the quark currents of Section 1.3.1 in the physical basis and their relationships.

Different choices of T_1-odd observable will yield determinations of m_{crit} which differ by $\mathcal{O}(a)$ lattice artefacts. This is not a problem in practice as having $m_q = \mathcal{O}(a\Lambda^2)$ only leads to lattice artefacts at most[5] of $\mathcal{O}(a^2)$. As a further practical consideration, the condition $m_R \leqslant 0.1\mu_R$ is checked (up to $\mathcal{O}(a)$) as a criterion to ensure that the theory is $\mathcal{O}(a)$-improved. Hence, it is clear that for very small twisted quark masses, this requires high statistical precision.

The Sheikholeslami-Wohlert Term

To end this section, automatic $\mathcal{O}(a)$-improvement in the presence of the SW term in the lattice action must be discussed. Neither it nor the Wilson term are invariant under T_1, but in the sense of the Symanzik effective action, they do not affect the symmetries of the target continuum theory given by S_0. As a result, automatic $\mathcal{O}(a)$-improvement by the argument above (or the original argument of Ref. [53]) is not affected, whatever the value of the parameter c_{sw}. The clover term can thus be safely added to the twisted mass action in the form

$$\bar{\chi}(x)C(x)\chi(x) = c_{\text{sw}}\frac{i}{4}\bar{\chi}(x)\sigma_{\mu\nu}F_{\mu\nu}(x)\chi(x) \tag{1.84}$$

where $F_{\mu\nu}$ can in principle be any discretisation of the field strength tensor. Conventionally, the symmetric form [64]

$$\begin{aligned} F_{\mu\nu}(x) = \frac{1}{8}\Big[& U_\mu(x)U_\nu(x+\hat{\mu})U_\mu^\dagger(x+a\hat{\nu})U_\nu^\dagger(x) \\ & + U_\nu(x)U_\mu^\dagger(x+\hat{\nu}-a\hat{\mu})U_\nu^\dagger(x-a\hat{\mu})U_\mu(x-a\hat{\mu}) \\ & + U_\mu^\dagger(x-a\hat{\mu})U_\nu^\dagger(x-a\hat{\nu}-a\hat{\mu})U_\mu(x-a\hat{\nu}-a\hat{\mu})U_\nu(x-a\hat{\nu}) \\ & + U_\nu^\dagger(x-a\hat{\nu})U_\mu(x-a\hat{\nu})U_\nu(x-a\hat{\nu}+a\hat{\mu})U_\mu^\dagger(x) \\ & - \text{h.c.}\Big]\,, \end{aligned} \tag{1.85}$$

[5] this can be shown either by thinking of S_0^{lat} as having a standard mass term which vanishes like a or by adding a contribution of $\mathcal{O}(a^2)$ to the Symanzik expansion

is used. With c_{sw} tuned as for standard Wilson fermions, it is reasonable to expect reduced discretisation errors because the dominant contribution to single insertions of S_1 should be absent from the Symanzik expansion of odd observables. On the other hand, it might be possible to tune c_{sw} to change lattice artefacts at higher orders where multiple insertions of S_1 occur, or to account for Equation (1.81).

Mass Non-Degenerate Doublet

The twisted mass lattice discretisation can also be used for the simulation of mass non-degenerate quarks, such as strange and charm. A possible choice for the lattice action is given by

$$S_{F,tm}^{1+1} = a^4 \sum_x \bar{\chi}_h(x) \Big[D_W + m_W + i\mu_\sigma \gamma_5 \tau^3 - \mu_\delta \tau^1 \Big] \chi_h(x) \tag{1.86}$$

which is guaranteed to yield a real and positive quark determinant given $\sqrt{m_q^2 + \mu_\sigma^2} > \mu_\delta$. Different choices are possible, as long as the mass splitting μ_δ and the twisted mass μ_σ are orthogonal to each other in flavour space. With m_W tuned to m_{crit}, the bare strange and charm quark masses are then given by

$$\mu_s = \mu_\sigma - \frac{Z_P}{Z_S}\mu_\delta \,, \qquad \mu_c = \mu_\sigma + \frac{Z_P}{Z_S}\mu_\delta \,, \tag{1.87}$$

where Z_P and Z_S are the pseudoscalar and scalar current renormalisation factors and usually $Z_P/Z_S < 1$. This leads to the problematic situation of requiring $\mu_\delta > \mu_\sigma$ to give physical strange and charm quark masses. Positivity is now not guaranteed anymore and sign changes of the heavy quark determinant over the course of a simulation could potentially render the importance sampling of the path integral meaningless. However, eigenvalue measurements during $N_f = 2+1+1$ simulations [65] have consistently shown that $\mu_s \sim 100$ MeV provides a sufficient spectral gap to prevent this from happening, much like in $N_f = 2+1$ simulations with standard Wilson quarks. It will be shown in Chapter 4 that this remains true with the clover term in the action.

Proofs of automatic $\mathcal{O}(a)$-improvement for the non-degenerate doublet are given in Refs. [66, 56] and similar symmetry restoration arguments apply. From Equation (1.81), it should be clear that the value of κ corresponding to maximal twist will in general be different in the light and strange/charm sectors of the theory. However, it can be shown that $m_{PCAC}^{\ell} - m_{PCAC}^{h} \sim \mathcal{O}(a)$ (see the Appendix of Ref. [65]), such that setting $\kappa^h = \kappa^\ell$ will only affect $\mathcal{O}(a^2)$ lattice artefacts. In addition, since $\mu_c > \mu_s > \mu_\ell$, the condition of maximal twist in the strange and charm sector is easily satisfied by m_{crit} tuned in the light sector. As a result, the value of κ is only tuned in the light sector of the theory.

Osterwalder-Seiler Quarks

The twisted mass action for non-degenerate quarks of Equation (1.86) is not diagonal in flavour space. This leads to flavour mixings even in the comparatively simple pseudoscalar meson sector with strange and charm quarks [65, 67]. The kaon mass can be extracted as the ground state of the generalised eigenvalue problem of the two-flavour PP correlation function in mixed parity-flavour sector. The D meson, however, is a highly excited state.

In Ref. [67] it is argued that effective masses for states with definite parity and flavour content can be extracted using operators which project to states with the correct quantum numbers in the continuum limit. In the meson sector, the flavour separation is manageable, but for observables involving more quark fields the complexity quickly increases.

It has been shown in Ref. [54] that *Osterwalder-Seiler* (OS) quarks [43] can be employed as valence fermions in a mixed action setup with twisted mass quarks in the sea, fully preserving the property of $\mathcal{O}(a)$-improvement if the same Wilson quark mass is used in the valence and sea actions. These can then be used to compute observables involving strange and charm quarks without flavour mixings. The OS quark action is similar to the twisted mass action for a degenerate doublet, except that it is defined only for a single flavour (its quark determinant is therefore complex and cannot be used in the molecular dynamics Hamiltonian). In the valence sector then, a contribution

$$\mathcal{L}^{OS}_{F,val}(x) = \bar{\chi}^{OS}(x)\left\{D_W[U] + m_{crit} + i\gamma_5 \begin{bmatrix} \mu_c & 0 \\ 0 & -\mu_s \end{bmatrix}\right\}\chi^{OS}(x)\,, \tag{1.88}$$

is added to the action, where m_{crit} takes the same value as in the light sea sector. Formally, a ghost contribution with the same masses is required to cancel the quark determinant. In practice, OS quarks are implemented by adding two additional twisted mass valence quark doublets with degenerate masses μ_c and μ_s to the valence action with further (formal) cancellations from ghost fields.

Twisted Mass Clover Action for Light, Strange and Charm Quarks

The complete twisted mass clover action for $N_f = 2+1+1$ simulations at maximal twist is thus

$$\begin{aligned} S^{N_f=2+1+1}_{tm} = \beta \sum_x \Big[\sum_{\mu<\nu} b_0\{1 - \frac{1}{3}\Re\,\mathrm{Tr}\,U_{\mu\nu}(x)\} + \sum_{\mu\neq\nu} b_1\{1 - \frac{1}{3}\Re\,\mathrm{Tr}\,R_{\mu\nu}(x)\}\Big] \\ + \sum_x \bar{\chi}_\ell(x)\left[D_W[U] + m_{crit} + i\mu\gamma_5\tau^3 + C[U(x)]\right]\chi_\ell(x) \\ + \sum_x \bar{\chi}_h(x)\left[D_W[U] + m_{crit} + i\mu_\sigma\gamma_5\tau^3 - \mu_\delta\tau^1 + C[U(x)]\right]\chi_h(x)\,, \end{aligned} \tag{1.89}$$

where $C[U(x)]$ is the clover term as defined in Equation (1.84). In principle, c_{sw} and m_{crit} may be chosen different between the light and heavy sectors, but they are conventionally chosen to be the same. As discussed before, the twisted mass μ_σ in the heavy sector can also be chosen to be flavour off-diagonal while the mass splitting μ_δ is taken diagonal.

LATTICE QCD SIMULATIONS

The evaluation of observables in LQCD follows from Equation (1.42). Using a Markov chain Monte Carlo algorithm, collections of field configurations for the link variables can be generated with the correct probability distribution. Observables can then be evaluated as simple statistical averages on these *ensembles*, such that

$$\langle \hat{\mathcal{O}} \rangle = \frac{1}{N} \sum_{i=1}^{N} \mathcal{O}[U_i] + \mathcal{O}(1/\sqrt{N}) \,, \tag{1.90}$$

with a statistical error which decreases with the inverse square root of the number of configurations in the ensemble. When LQCD simulations were first attempted, it was clear that evaluating the contribution of the quark determinant was computationally intractable at the time. As a result, the quark contribution to the path integral weight was set to unity (which is equivalent to making the sea quarks infinitely massive) and calculations were performed in the so-called *quenched* approximation. The gauge field configurations with the correct probability distribution could be generated with simple local algorithms of the Metropolis type [68, 69].

When this approximation is lifted in so-called *dynamical* simulations, local algorithms which update one or a few link variables at a time become impractical because of the computational cost of inversions of the Dirac operator which are required to evaluate the action contribution of the quark determinant[6]. In current simulations using dynamical quarks, global updating algorithms are used instead, with most of them variants of the Hybrid Monte Carlo algorithm.

Hybrid Monte Carlo Algorithm

In the Hybrid Monte Carlo (HMC) algorithm [71, 72, 73], Markov chains of field configurations are produced by augmenting the action to be simulated with momenta conjugate to the link variables. The algorithm consists of the molecular dynamics (MD) step and a Metropolis accept-reject step. In the MD evolution, the equations of motion of this augmented system are integrated along a fictitious molecular dynamics time for a number of integration steps, called a *trajectory*. The Markov chain produced over the course of many such trajectories provides the ensemble of gauge configurations $\{U\}$ for the evaluation of observables. A discussion of the symplectic integrators used for updating the conjugate variables is given below in Section 1.4.2.

At the beginning of *each* trajectory, the momenta are sampled from a Gaussian distribution. Similarly, one pseudofermion field is generated for each non-local fermionic contribution as discussed in Section 1.2.3 and kept constant over the course of the integration. At the end of the trajectory, the total energy of the system, H_1, is compared to its energy at the beginning of the trajectory, H_0. The new field configuration is then accepted with a probability

$$P_{\text{acc}} = \max\{1, e^{-\delta H}\} \tag{1.91}$$

[6] Note that the fermionic contribution can also be represented as completely local bosonic field theory with n degrees of freedom which reproduces the quark determinant in the limit $n \to \infty$. A good approximation can be obtained with n not too large for local bosonic algorithms to be practical. [70]

where $\delta H = H_1 - H_0$. If a configuration is rejected, the initial configuration is added to the ensemble and the next trajectory is started from there. This ensures that finite step size errors in the integration do not bias the distribution of field configurations. The size of the energy violation depends on the volume and, in a complicated fashion, on the details of the integrator, its step-size and the action being simulated. In any case, the acceptance rate should be kept high because otherwise the ensemble will contain many copies of the same field configuration, increasing autocorrelations. Generally, there is a point where acceptance rate and computational expense balance out and it is desirable to perform simulations at this point.

The HMC algorithm respects the property of detailed balance which relates the transition probabilities between two field configurations

$$e^{-S_{\text{eff}}[U]} P(U \to U') = e^{-S_{\text{eff}}[U']} P(U' \to U) \tag{1.92}$$

and ensures that the algorithm samples the target probability distribution correctly. The inclusion of the conjugate momenta is inconsequential because they are random variables with a vanishing mean and their contribution to the Hamiltonian thus only generates a constant multiplicative factor which cancels in Equation (1.92). The pseudofermion fields are generated at the beginning of each trajectory ($K = 1$ in Equation (1.42)) and act like external fields over the course of it. It follows that the ensemble of gauge configurations thus effectively contains two stochastic estimates, with the average over pseudofermion fields implicit.

As an important side remark it should be noted that the molecular dynamics Hamiltonian H_{MD} and the target Hamiltonian H do not need to be exactly the same. If the former is an approximation of the latter, a non-zero acceptance probability can be obtained with the acceptance step ensuring that this does not introduce bias. In order to make the algorithm more efficient, it is common to make a number of approximations in the computation of H_{MD}. Because they are relevant to certain parts of this thesis, the details of the MD evolution will be discussed below.

Molecular Dynamics Integration

The MD Hamiltonian is given by.

$$H_{\text{MD}} = \frac{1}{2} \operatorname{Tr}\left[P_\mu(x) P_\mu(x)\right] + S_G[U] + \sum_i S_i[U]\,. \tag{1.93}$$

where the *monomials* $S_i[U]$ are the various, usually fermionic, contributions to the effective action. Hamilton's equations of motion take the usual form

$$\dot{U} = \frac{dH}{dP} = P \qquad \dot{P} = -\frac{dH}{dU} = -\frac{dS_{\text{eff}}}{dU}\,, \tag{1.94}$$

where the time derivatives are taken with respect to the fictitious MD time. For definiteness, the gauge field $U_\mu(x)$ is a member of the group $SU(3)$ in the fundamental matrix representation while the conjugate momenta $P_\mu(x)$ are in the Lie algebra $\mathfrak{su}(3)$ and thus traceless anti-Hermitian matrices (in this convention). They can be can be expressed in terms of the Gell-Mann matrices λ_a ($a = 1, 2, \ldots, 8$) as

$$P_\mu(x,\tau) = \sum_{a=1}^{8} i\lambda_a P_\mu^a(x,\tau)\,, \tag{1.95}$$

where the components $P_\mu^a(x)$ are real and the factor i is also conventional. At the beginning of each trajectory, they are sampled from a Gaussian distribution.

The derivative of a scalar function f with respect to group elements, can be defined via the exponential map (see e.g. Ref. [41])

$$\frac{\delta f[U]}{\delta U_\mu(x)} = \sum_{a=1}^{8} \frac{\partial}{\partial \sigma_a} f\left[e^{i\sigma_a \lambda_a} U_\mu(x)\right]_{\sigma_a=0}, \tag{1.96}$$

which is in the tangent space and thus in the algebra. Applying this to H_{MD}, the resulting derivative is called the force. It is easy to see that the computation of the derivative of fermionic contributions will produce factors containing gauge fields as well as products and inverses of the Dirac operator sandwiched between the pseudofermion fields. As a result, these contributions are the most computationally demanding part in the generation of gauge field configurations.

In the simplest *leapfrog* (LF) symplectic integration scheme with time-step $\delta\tau$, the phase space variables are alternately updated via

$$\delta P_\mu^a(x,\tau) = \left\{ \frac{\partial}{\partial \alpha} H_{MD}\left[e^{i\alpha\lambda_a} U_\mu(x,\tau)\right]\right\}_{\alpha=0} \tag{1.97}$$

$$P_\mu^a(x,\tau + 1/2\delta\tau) = P_\mu^a(x,\tau - 1/2\delta\tau) - \delta\tau \cdot \delta P_\mu^a(x,\tau) \tag{1.98}$$

$$U_\mu(x,\tau+\delta\tau) = \exp\left[\delta\tau \cdot P_\mu(x,\tau+1/2\delta\tau)\right] U_\mu(x,\tau), \tag{1.99}$$

with half-steps at the beginning and end of the trajectory. One step of the leapfrog algorithm can thus also be expressed as

$$\begin{aligned} T_{LF} &= T_P[\delta\tau/2] T_U[\delta\tau] T_P[\delta\tau/2] \\ T_P[\delta\tau] &: P \to P' = P - i\delta\tau \cdot \delta P \\ T_U[\delta\tau] &: U \to U' = \exp\left[i\delta\tau P\right] U, \end{aligned} \tag{1.100}$$

such that a trajectory of N steps is given by $(T_{LF})^N$ and in practice $T_P[\delta\tau/2]^2$ are of course combined into $T_P[\delta\tau]$. The exponential map ensures that the $U_\mu(x)$ remain in SU(3) up to round-off errors and these small deviations are corrected by group projections at regular intervals.

In simulations with light quarks, nested higher order integrators are used which reduce the finite step size errors and lead to higher acceptance rates at the same computational cost. A complete description of various symplectic integrators can be found in Ref. [74] and in this thesis the most commonly used one is the *second order minimal norm* (2MN) integrator. In terms of the update operators of Equation (1.100), it is given by

$$T_{2MN}[\delta\tau] = T_P[\lambda\delta\tau] T_U[\delta\tau/2] T_P[(1-2\lambda)\delta\tau] T_U[\delta\tau/2] T_P[\lambda\delta\tau], \tag{1.101}$$

where λ is a tunable parameter.

Multiple Time Scales and Mass Preconditioning

The integration of the equations of motion can be rendered more efficient by splitting it over multiple time scales, such that different contributions are integrated at different frequencies or, equivalently, with different step sizes. This was first proposed for QCD in [75] and can be generalised for arbitrary

numbers of time scales and any type of integrator. The different monomials S_j of the effective action are associated with different time scales t_i and integrated in a nested fashion. Multiple monomials can be integrated on the same time scale.

In the simplest scenario with two time scales t_0 and t_1, one fermionic contribution S_1 could be on time scale t_1 and the gauge contribution S_G on time scale t_0. The trajectory is split into N_1 steps $\delta\tau_1$, each step $\delta\tau_1$ is further split into N_0 steps of length $\delta\tau_0 = \delta\tau_1/N_0$. For the LF integrator, one proceeds in a nested fashion:

$$\begin{aligned} T &: \left(T_1[\delta\tau_1/2](T_0[\delta\tau_0])^{N_0} T_1[\delta\tau_1/2]\right)^{N_1} \\ T_1[\delta\tau] &: T_{P_1}[\delta\tau] \\ T_0[\delta\tau] &: T_{P_0}[\delta\tau/2] T_U[\delta\tau] T_{P_0}[\delta\tau/2] \end{aligned} \tag{1.102}$$

such that T_1 only involves the derivative of S_1 and T_0 contains the updates of the gauge field. Higher order integrators, such as the 2MN above, can be similarly generalised to multiple time scales by replacing the T_U in Equation (1.101) with $(T_{i-1})^{N_{i-1}}$ in a recursive fashion, such that only T_0 contains the gauge update. With this prescription, the generalisation to n time scales is straightforward. The efficiency gains of multiple time scales are greatest when the most demanding T_{P_i} can be computed least frequently. One way of optimising this computational cost is through mass preconditioning.

Mass preconditioning was initially introduced to speed up simulations with standard Wilson fermions [76, 77] by reducing the cost of computing the fermionic force. At the time, the quark determinant was split into a determinant and a determinant ratio (see below), such that the sum of the condition numbers of the respective operators in the pseudofermion action was lower than the condition number of the original operator, thus lowering the simulation cost. Later, the fact that the resulting terms have different force contributions was exploited by integrating them on different time scales. This was first done for two time scales in in Ref. [78] and generalised in Ref. [79].

The fermionic contribution to the path integral of a doublet of mass degenerate quarks is given by

$$\int \mathcal{D}\phi^\dagger \mathcal{D}\phi \, \exp\{-\phi^\dagger (Q^\dagger Q)^{-1} \phi\}, \tag{1.103}$$

in terms of pseudofermions and the Hermitian operator $Q = \gamma_5 M$. The operator Q can be shifted by a constant through the addition of a twisted mass

$$W_\pm = Q \pm i\rho\gamma_5 \,, \tag{1.104}$$

such that

$$W_+ W_- = Q^2 + \rho^2 \,. \tag{1.105}$$

It should be noted that Q can itself be a twisted mass operator with *target mass* μ.

Using this shifted operator, a fermionic contribution to the path integral, equivalent to Equation (1.103), is given by

$$\int \mathcal{D}\phi_1^\dagger \mathcal{D}\phi_1 \mathcal{D}\phi_2^\dagger \mathcal{D}\phi_2 \, \exp\{-\phi_1^\dagger (W_+ W_-)^{-1} \phi_1 - \phi_2^\dagger W_- \frac{1}{Q^2} W_+ \phi_2\}, \tag{1.106}$$

where the second factor is often referred to as a *determinant ratio*. This can be further extended with more preconditioning masses ρ_i.

Clearly, when W_+W_- and Q^2 are similar, $W_-Q^{-2}W_+ \sim 1$. This is therefore almost independent of the gauge field and the force of such a term in the action will have a small magnitude. As a result, the terms from determinants and the different determinant ratios can be integrated on different time scales. This distribution has been shown to increase acceptance rates compared to integration schemes with a single time scale and no mass preconditioning. In addition, it leads to significant efficiency gains because terms with large ρ_i are much cheaper to compute, contribute strongly to the derivative and are hence integrated on fine time scales. Terms with small ρ_i on the other hand contribute little to the derivative, are expensive to compute and can thus be integrated on coarser time scales.

To optimise this, preconditioning masses are arranged monotonically, guided by experience and force measurements during simulations to match the hierarchy of integration time scales. While Ref. [79] explains that the force of a single determinant ratio is inversely proportional to the square of preconditioning mass ρ, no clear functional dependence is discussed. Refs. [80, 81] suggest that the preconditioning masses should be chosen such that they are equally spaced on some logarithmic scale to provide an appropriate hierarchy of forces, but no clear guidance is given.

The Hybrid Monte Carlo algorithm with mass preconditioning and multiple time scales has been instrumental in enabling lattice QCD simulations with light dynamical quarks. It is probable that this will continue to be the case for simulations at the physical pion mass, but refinements in the tuning of the various parameters are likely necessary. An attempt at a purely empirical description of the parameter dependence of the force contributions and their optimisation will be given and tested in Chapter 5, in the hope that the simple fit functions can be useful in practice.

2 APPROACHING THE PHYSICAL POINT

Unphysically large light quark masses are one of the chief approximations of current simulations in lattice QCD. The discretisation of the QCD Lagrangian is not unique and various collaborations are currently working to simulate at or close to the physical pion mass. These efforts are not mutually exclusive: the usage of multiple discretisations provides an important test of the universality of the continuum limit of lattice QCD. In order to place the results of this thesis in a wider context, the "landscape" of various simulation efforts will be summarized below.

A MULTITUDE OF QUARK DISCRETISATIONS

The MILC collaboration is using highly improved staggared quarks [82] (HISQ) with a tuned smearing setup [83] using mass-degenerate light quarks as well as physical strange and charm quarks in the sea. Simulations using Wilson quarks with non-perturbatively computed improvement coefficients are performed widely. The PACS-CS collaboration pursued $N_f = 2+1$ simulations with the Iwasaki gauge action [44] while the CLS effort is currently performing [84] $N_f = 2+1$ simulations using the Lüscher-Weisz gauge action.

A second class of simulations with Wilson fermions are those employing some level of gauge field smearing in the discretisation of the covariant derivative in the Dirac operator. The QCDSF collaboration has performed a simulation programme [85] with non-perturbatively improved Wilson quarks with mild stout smearing in the covariant derivative. A similar approach with more aggressive smearing was used by the Hadron Spectrum collaboration on anisotropic lattices with tree-level tadpole improved values of the Sheikholeslami-Wohlert term. The BMW collaboration has performed $N_f = 2+1$ simulations with tree-level improved Wilson quarks, employing significant so-called *HEX* smearing [86, 87]. This programme was extended in two directions [88]. Firstly, $N_f = 1+1+1+1$ simulations with an added fourth quark were performed in which the up and down quarks were tuned to their physical masses (rather than being mass-degenerate) and the strange and charm quarks take on their physical mass values. Secondly, together with a simultaneous retuning of the bare quark masses, the unquenched electromagnetic interaction was added, providing the first ab initio calculation of the proton-neutron mass difference.

Lattice actions with a discrete realisation of chiral symmetry [89], so-called Ginsparg-Wilson fermions, are attractive because they offer the potential to study very small (even vanishing) quark masses and are bound to have very good scaling properties towards the continuum limit without the need for the expensive computation of improvement coefficients. Simulations using overlap [90] quarks were studied [91] by the χLF collaboration, and larger scale simulations were were out [92] by the JLQCD and TWQCD collaborations.

A different realisation of Ginsparg-Wilson relation is given by domain wall (DW) fermions [93]. These link a massive theory in $2n+1$ dimensions with a (potentially) massless theory in $2n$ dimensions in which all doubler

modes have large gauge invariant masses introduced by the usual Wilson term in the $(2n+1)^{\text{th}}$ dimension. At finite size of this extra dimension, a residual mass am_{res} persists which contributes to the total quark mass. As the extent of this extra dimension is increased, the Ginsparg-Wilson relation can be satisfied to essentially arbitrary precision. A large scale simulation programme using different implementations of DW fermions in five dimensions is being pursued by the RBC and UKQCD collaborations [94]. Similarly, the JLQCD and TWQCD collaborations are conducting [95] simulations with DW fermions.

Wilson twisted mass quarks provide a theoretically clean formulation of lattice QCD with symmetry based $\mathcal{O}(a)$-improvement. The simulation programmes of the ETMC for $N_f = 2$ [96] and $N_f = 2+1+1$ [97] span light and heavy flavour physics, studies of the phase structure of lattice QCD with Wilson fermions, topological properties, quantities related to the hadronic vacuum polarisation and studies of the baryon spectrum and baryonic form factors. The configurations generated by the ETMC are publicly available via the International Lattice Data Grid (ILDG) [98]. The simulation software is also freely available [79] and the development process is open to public scrutiny[1].

This wide applicability results chiefly from the property of automatic $\mathcal{O}(a)$-improvement at maximal twist as well as the relation [54] between twisted mass quarks and Osterwalder-Seiler (OS) quarks [43], which preserves $\mathcal{O}(a)$-improvement in a special mixed action setup. The first property is especially important for computations in the baryon sector where operator improvement can quickly make calculations intractable. The second property enables the non-perturbative computation of four quark operators, as they appear in the effective weak Hamiltonian, with simplified mixing under renormalisation. The usage of OS valence quarks also provides for a rather straightforward computation of strange and charmed meson masses and decay constants in the $N_f = 2+1+1$ setup, avoiding the flavour and parity mixings of the non-degenerate strange/charm sea quark action.

The $N_f = 2+1+1$ simulations encompass around 25 ensembles with charged pion masses ranging down to about 220 MeV and three lattice spacings. The strange and charm quark masses are tuned such that $2M_K^2 - M_{\pi^\pm}^2$ and the D meson mass take on approximately physical values. Any deviations can be accounted for by ensembles with slightly differently tuned strange and charm quark masses. The analysis of finite volume effects is enabled by a number of different volumes at constant light quark mass.

Simulations at the physical pion mass would eliminate the chiral extrapolation which currently dominates some systematic errors, as in the baryon sector for example [99]. In other situations, where effective theories provide better guidance, these same ensembles can give excellent control over the chiral extrapolation. This can can then be used to determine the low energy constants of the effective theory and even reduce the final uncertainties on the given physics results, relevant for instance in the determination of heavy-light meson decay constants.

As discussed in Section 1.3, the twisted mass action explicitly violates parity and flavour symmetry. The only severe manifestation of this lattice artefact observed so far is an unphysical mass difference between the charged and neutral pions at non-zero lattice spacing. This cut-off effect is the main drawback of the twisted mass discretisation as it can render simulations unstable, puts an effective lower limit on the light quark masses

[1] https://github.com/etmc

that can be simulated at a given lattice spacing and entails exponentially enhanced finite volume effects.

ISOSPIN SYMMETRY BREAKING

In practice, two measures of this mass splitting are considered. The first is the mass of an unphysical neutral pion computed from correlation functions not taking into account quark-line disconnected contributions. This state, which is heavier than the corresponding charged pion, can in some sense be thought of as a neutral pseudoscalar meson made up of purely valence Osterwalder-Seiler quarks with the same bare quark mass as the light quarks in the sea. The second is the full neutral pion, for which disconnected contributions have to be taken into account and which is lighter than the charged pions. In addition to affecting physics results through increased finite volume effects, the lightness of the neutral pion is related to the zero temperature phase structure of the lattice theory and can render simulations unstable as the light quark mass is lowered to its physical value.

The stability of simulations with the new twisted mass clover action will be discussed in Chapters 3 and 4. Measurements of the pion mass splitting in $N_f = 2$ simulations and how it is affected by the value of c_{sw} in $N_f = 2+1+1$ simulations will be presented in Chapter 6. To provide a framework for these discussions, the connection between the pion mass splitting and the zero temperature phase structure of the Wilson and Wilson twisted mass lattice theories (in terms of their bare parameters) and the stability of simulations is discussed below. The instabilities observed in simulations with twisted mass quarks have their counterpart in simulations with standard Wilson quarks and as a result, a discussion of the situation for standard Wilson fermions is insightful.

The discussion will be based on results from Wilson chiral perturbation theory (WχPT) [100, 101], an effective theory which adds to the chiral effective Lagrangian further spurionic terms proportional to powers of the lattice spacing, as they appear in the Symanzik effective action [25, 26] for Wilson fermions [58]. Like its continuum analogue, WχPT relies on the expansion of the effective Lagrangian around a saddle point in powers of momenta with spurionic quark mass terms breaking chiral symmetry explicitly.

In order order to account for discretisation artefacts of the lattice theory, a simultaneous expansion in powers of the lattice spacing is also done. This requires the definition of a suitable power counting scheme to fix the relative sizes of momenta, quark masses and lattice artefact terms. The choice of power counting scheme is equivalent to deciding whether the quark masses or lattice artefacts are the dominant source of explicit chiral symmetry breaking. The first of the two schemes used in practice is referred to as *generically small masses* (GSM), for which the (degenerate) light quark masses are considered to be of the order $m_q \sim a\Lambda_{QCD}^2$. The second is called *large cut-off effects* (LCE), for which $m_q \sim a^2\Lambda_{QCD}^3$ and lattice artefacts are on the same order as the quark masses.

An effective potential is constructed and minimized, parametrized by two combinations of low energy constants c_1 and c_2, the sign and magnitude of which affect the conclusions regarding the phase structure. The two coefficients in general depend on the details of the gauge and fermion actions of the lattice theory, but c_1 can be taken to be proportional to the bare quark mass. One of the consequences of this phase structure is that the squared

pion masses may depend differently on the quark masses than in continuum χPT, if lattice artefacts are sizeable compared to the quark masses.

The Phase Structure of Wilson Fermions

The importance of the power counting scheme becomes clear when the effective chiral Lagrangian is used to understand the phase structure of the Wilson lattice theory at zero temperature. This was done in Ref. [100], explaining the existence of so-called *Aoki phases* [102, 103, 48, 104]. These are regions in the β/m_q plane in which lattice artefacts dominate and break flavour and parity symmetries, making simulations unstable and the continuum limit unreliable.

When c_2 is positive, the Aoki phase exists in a region $-2c_2 < c_1 < 2c_2$, where $m_q \sim c_1$ is small. Here, the "charged" pions $\pi_{1,2}$ are the massless Goldstone bosons of the broken flavour symmetry and the "neutral" pion π_3 is massive and its mass depends on c_1 and c_2. At the phase boundary, all three pions are massless and outside of the Aoki phase, all squared pion masses are degenerate and depend linearly on the quark mass at this order of WχPT.

When c_2 is negative, a different, so-called *first-order* scenario takes place. All three squared pion masses depend linearly on the quark mass everywhere (at this order of WχPT), but there exists a minimum pion mass which depends on the magnitude of $c_2 a^2$. An interesting consequence of this last point is that any estimate of m_{crit} based on the extrapolated vanishing of the squared pion mass will likely be wrong by some offset when c_2 is large. This can be taken into account by approaching m_{crit} from both directions.

Ref. [100] claims that the introduction of the clover term and the subsequent improvement of the action at $\mathcal{O}(a)$ should not alter the sign of c_2, although this seems to be contradicted by inconsistencies in the literature, as will be discussed in Section 2.2.3 below. It is important, however, to note the limitations inherent in the description above. Firstly, the size of higher order terms is unknown and these could modify even the qualitative behaviour of the pion masses, *mimicking* a certain sign and value of c_2. Secondly, a comparison of the results from WχPT to simulation data is limited by the knowledge of the large number of low energy constants which enter the expressions for the pion masses, although some bounds are known. Still, the qualitative agreement of the picture given by WχPT and simulation data is very compelling.

Wilson Twisted Mass Fermions

Wilson twisted mass chiral perturbation theory [105, 106, 107] (WtmχPT) is an extension of the above description, taking into account the chiral transformation at non-zero lattice spacing which leads to the twisted mass formulation of lattice QCD, as discussed in Section 1.3. The phase structure at non-zero twisted mass was studied in this framework in Refs. [108, 109, 110] and the connection between the pion mass splitting and the existence of the Aoki and first order scenarios was elucidated. The situation is more complicated because in addition to the (subtracted) Wilson quark mass m_q, the twisted quark mass μ provides a further source of chiral symmetry breaking. In particular, Refs. [110, 62] provide a comprehensive understanding at the given order. When $c_2 > 0$, at non-zero twisted quark mass, the Aoki phase transition is washed out into a crossover

and the neutral pion is heavier than the charged pions. For the special case of maximal twist, the difference between the squared pion masses, $(M^2_{\pi_{1,2}} - M^2_{\pi_0})$, is proportional to $c_2 a^2$. For general values of the twist angle, the pion masses depend in a complicated fashion on the two quark masses.

The case of negative c_2 is the one so far observed in simulations with twisted mass quarks and is in line with the arguments of Ref. [62]. In the μ/m_q plane, there now exists a line of first order phase transition, curved slightly away from the twisted mass axis and delimited by some negative and positive critical twisted quark masses $\pm\mu_c$. Along this line, all three pion masses are non-zero but at its endpoints, the neutral pion mass vanishes while the charged pion mass assumes a minimal value proportional to $c_2 a^2$. Away from the phase transition line, all three pions are massive with the squared neutral pion mass is lower than that of the charged pions by a quantity proportional to $c_2 a^2$. Simulations with twisted mass quarks are generally performed at maximal twist, which to this order, is exactly on the twisted mass axis. If $|\mu_c|$ is large compared to the physical light quark mass, it is therefore not possible to simulate at maximal twist in the parameter region corresponding to the physical pion mass, unless the lattice spacing is reduced. An asymmetry and discontinuity in the proportionality of the squared pion mass to the Wilson quark mass can be observed, both of which are softened by the twisted quark mass. This is shown in the upper panel of Figure 2.2 which has been reproduced from Ref. [62].

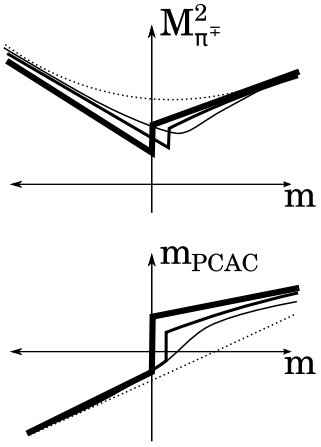

Figure 2.2: $M^2_{\pi^\pm}$ (top) and m_{PCAC} (bottom) as a function of the Wilson quark mass m in the LCE regime. The thickness of the line is inversely proportional to the twisted quark mass μ with the thickest line corresponding to $\mu = 0$. This graphic is a recreation of Figures 2 and 3 from Ref. [62].

In this figure, moving along the horizontal axis is equivalent to moving along the untwisted quark mass axis m_q in the μ/m_q plane, while the thickness of the lines indicates the dependence for different values of μ. The standard mass m shown here is the subtracted Wilson quark mass $m = m_W - m_{crit}$ with m_{crit} *not* including the μ-dependent shift discussed in Section 1.3.3.

Similarly, the PCAC quark mass shows a discontinuity when the Wilson quark mass passes from negative to positive values as depicted in the lower panel of Figure 2.2. As μ is increased, the discontinuities are softened and importantly, the point where m_{PCAC} vanishes and the theory is at maximal twist is displaced to the right. For small twisted quark masses and sizeable $c_2 a^2$, it is therefore very difficult to tune to maximal twist on account of the rapid evolution around $m_{PCAC} = 0$. This is another reason why reducing the magnitude of c_2 is thus essential for simulations with physically light quarks at reasonable lattice spacings.

From Ref. [111], in terms of the renormalised Wilson quark mass, the PCAC quark mass to lowest order is

$$m_{PCAC} = m - \frac{c_2 a^2}{B} \cos(\phi)\,, \tag{2.1}$$

where ϕ is the angle which parametrizes the chiral rotation of the ground state of the effective potential due to the twisted quark mass and is related to the usual twist angle. The most commonly used definition of maximal twist relies on tuning m such that $m_{PCAC} \sim 0$, but of course $\cos(\phi)$ depends on both m and μ. Following the argumentation of Ref. [112], when higher order terms are introduced as in Ref. [62], $m_{PCAC} = 0$ implies $\cos(\phi) = \mathcal{O}(a)$, the shift in the lower panel of Figure 2.2. At leading order, equating ϕ and the twist angle, the slope of m_{PCAC} is modified to

$$m_{PCAC} \sim \left(1 - \frac{c_2 a^2}{B\mu}\right) m + \ldots \tag{2.2}$$

such that a negative c_2 would imply a steeper slope and suggesting that the relation between the PCAC quark mass and the Wilson quark mass might be usable as a signal for the magnitude of c_2 when all renormalisation constants are taken into account.

Dependence of c_2 on the Details of the Lattice Discretisation

The lattice discretisation, both in the gauge and fermion sectors, as well as the number of dynamical quarks have a strong influence on discretisation artefacts. There is evidence that in particular the choice of gauge action has a significant effect on the phase structure of the lattice theory. In addition, there is limited evidence that the fermion discretisation also affects the magnitude of c_2.

Numerical results from Ref. [113] suggests that the Wilson plaquette gauge action combined with $N_f = 2$ Wilson quarks with a non-perturbatively tuned clover term exhibit $c_2 \geqslant 0$. This is in contrast to Refs. [114, 115], in which evidence for unimproved $N_f = 2$ Wilson quarks and plaquette gauge action points to $c_2 < 0$ and the strength of the phase transition is suppressed by the addition of a twisted quark mass. With the same action in the strong coupling regime ($\beta < 4.6$), Ref. [116] corroborates the indications from Refs. [102, 103, 104] for the existence of the Aoki phase in the limit $\mu = 0$ and hence $c_2 > 0$. It should be noted, however, that at the coarse lattice spacings considered in this last work, higher order contributions may not be negligible. A hypothetical phase diagram in the space of m_q, μ and β is given for example in Ref. [117][2], where the Aoki phase exists below some value of β and the first order scenario is realised for larger values of β.

The pion mass splitting with twisted mass fermions was studied in Ref. [118] in the quenched approximation. The addition of the clover term with a non-perturbatively tuned value of c_{sw} is seen to significantly reduce the splitting between the squared charged pion mass and that of the neutral pion computed from connected correlation functions. Further, an unpublished study by the ETMC suggests that in the quenched approximation, this mass splitting is increased if instead the tree-level value $c_{sw} = 1.0$ is used.

For $N_f = 3$ ($N_f = 2 + 1$), Ref. [119] provides strong evidence that perturbatively $\mathcal{O}(a)$-improved Wilson fermions exhibit strong signs of a first order phase transition and hence $c_2 < 0$ at values of β relevant for practical simulations. For this situation, the usage of improved gauge actions seems to significantly reduce the effects of the phase transition and the simulation programme of the PACS-CS collaboration therefore relied on the Iwasaki gauge action [44]. The same was observed in Ref. [120] with unimproved Wilson quarks and three different improved gauge actions. Mild stout smearing applied to $N_f = 2$ simulations with unimproved Wilson twisted mass quarks and tree-level Symanzik improved gauge action [121, 122] is shown in Ref. [123] to reduce the strength of the phase transition. However, it is unclear whether the influence of the stout smearing on the effective lattice spacing has been properly considered, although Ref. [124] also gives a measurement of the neutral pion mass which indicates a reduction of c_2.

$N_f = 2 + 1 + 1$ simulations with unimproved Wilson twisted mass fermions and the tree-level Symanzik improved gauge action in Ref. [65] show quite strong signals of $c_2 < 0$. Even with the Iwasaki gauge action, measurements on ETMC $N_f = 2 + 1 + 1$ ensembles consistently give [124] neutral pion masses lower than the charged ones. In addition, a direct

[2] Figure 4.8, page 97

computation of c_2 using methods proposed in Refs. [125, 126] provides a clear signal for $c_2 < 0$ [124]. Further, in the renormalisation programme of the ETMC, runs with $N_f = 4$ mass-degenerate unimproved Wilson twisted mass fermions were so strongly affected by the phase transition that simulations at maximal twist were not possible. As a result, $\mathcal{O}(a)$-improvement had to be carried out by averaging the requisite correlation functions from simulations with positive and negative Wilson quark masses.

Effects on Simulations and Physical Results

From the point of view of the simulations using Hybrid Monte Carlo techniques, the pion mass splitting can become problematic as the physical average up/down quark mass is approached if the first order scenario is realised in the relevant parameter region. Close to the endpoints of the phase transition line (approached from either $\mu = 0$ or $|\mu| > |\mu_c|$), the neutral pion becomes very light. The existence of a very light excitation in the spectrum of the lattice theory implies a large correlation length which may result in strong fluctuations and significantly increased autocorrelation times. In addition, closeness to a first order phase transition induces the approximate coexistence of the two phases at either side of the transition for a given set of parameters. This was demonstrated numerically for three flavours of perturbatively $\mathcal{O}(a)$-improved Wilson quarks with the plaquette gauge action in Ref. [119] and was also explored using two flavours of unimproved Wilson twisted mass quarks with the plaquette gauge action in Refs. [114, 115]. Tunnelling between these two phases is strongly suppressed and it is likely that configuration space is wrongly sampled.

In twisted mass lattice QCD, the analysis of observables for the extraction of the low energy properties of QCD is affected by the potential lightness of the neutral pion in two ways. Firstly, it has been shown [127] in WtmχPT that finite size artefacts (FSE) coming from π^0 are exponentially enhanced compared to those coming from $\pi^\pm$. This may easily double the FSE on observables such as the charged pion mass or charged pion decay constant if the pion mass splitting is large (compared to a situation with degenerate pion masses). In the LCE regime, it has been shown in Ref. [112] that chiral fits better describe the data when logarithms involving the neutral pion mass are included in the fit procedure. Modern chiral extrapolations of twisted mass data, such as in Ref. [128], include c_2 with priors in the fit procedure to account for the effect. In addition to enhanced finite volume effects and fits with more parameters, the extraction of physics results from twisted mass data with large pion mass splittings can be disturbed by lattice artefacts more directly. If the approach towards physical average up/down quark mass is affected by large $\mathcal{O}(a^2)$ lattice artefacts, it is unclear whether the continuum limit can be reliably taken. If a combined continuum and chiral extrapolation is to be used, for the smallest quark masses lattice artefacts may be a significant source of chiral symmetry breaking while being negligible for higher quark masses.

In WχPT, this can be thought of as passing from a regime where the GSM power counting is appropriate into one where the LCE power counting must be used. In principle, this transition could be incorporated into the fit procedure, but it would probably be quite complicated to reliably match the expressions from the two regimes. In order to use WχPT for chiral extrapolations, it is therefore in practice recommended that one regime or the other be enforced. At a lattice spacing of approximately 0.1 fm, this

means that the GSM regime is valid for $m_q \gtrsim 45$ MeV. The LCE regime on the other hand should be valid for $m_q \gtrsim 7$ MeV. Since both of these quark masses are above the physical average up/down quark mass, it is reasonable to assume that lattice artefacts will play a major role in simulations at the physical pion mass and that the combination of data from the two regimes should be done with care. It should be noted that taking the continuum limit at fixed quark mass, i.e. potentially passing from the LCE regime to the GSM regime because the lattice spacing is reduced, is not an issue as long as one is working at maximal twist throughout to the accuracy required [129]. Given stable simulations without signs of complicated phase structure at the physical pion mass, the continuum extrapolation should therefore be straightforward and automatic $\mathcal{O}(a)$-improvement should hold as expected.

The Pion Mass Splitting in Current Two and Four Flavour Simulations

Ref. [124] provides the most complete list[3] of pion mass measurements with Wilson twisted mass fermions for $N_f = 2$ [96] with tree-level Symanzik improved gauge action and $N_f = 2+1+1$ with Iwasaki gauge action [97]. For $N_f = 2$ at $\beta = 3.9$ ($a \sim 0.08$ fm), the pion mass splitting for the lightest simulated charged pion mass is around 60 MeV. At the finer lattice spacing of $a \sim 0.063$ fm at $\beta = 4.05$, this is reduced to about 35 MeV. For $N_f = 2+1+1$, with lattice spacings taken from Ref. [128], the pion mass splitting at $\beta = 1.90$ ($a \sim 0.089$ fm) at the lightest simulated quark mass for this lattice spacing is around 140 MeV and remains roughly constant as the light quark mass is increased. At the finest lattice spacing $a \sim 0.062$ fm at $\beta = 2.10$, a value of 40 MeV is obtained and this too remains roughly constant as the light quark mass is increased. One should take into account that the measurements of the pion mass splitting given above have statistical uncertainties between 10 to 30% and are subject to unquantified systematic errors.

Ref. [62] suggests a mild dependence of the difference of the squared pion masses on the (twisted) light quark mass and a linear extrapolation to the physical point. However, Ref. [124] finds that constant and linear extrapolations can barely be differentiated with the available data. On the other hand, the effect of the strange and charm quark masses on the pion mass splitting is substantial. For the coarser lattice spacing at a rather heavy light quark mass, a reduction in the strange quark mass of $\sim 20\%$ (while the charm quark mass is kept roughly constant) results in a reduction of the pion mass splitting of around 30%. For the finest lattice spacing, a reduction of the strange quark mass of around 12% with a simultaneous reduction of the charm quark mass by about 30% almost doubles the pion mass splitting. In the first case, only the parameter μ_δ was changed slightly while in the second case both the average bare quark mass μ_σ of the non-degenerate doublet as well ass μ_δ were changed significantly. In both cases the changes in the pion mass splitting are such that the indicated error bars just, or not quite, touch.

Although based on a very limited amount of data, this last point should be carefully considered in the tuning of strange and charm quark masses in production simulations. It seems to confirm the naïve expectation that the strange and charm sector has a strong effect on discretisation artefacts (and hence the pion mass splitting), even though common light quark low energy observables are largely unaffected by the presence of strange and

[3] Tables 3 and 1 in the reference, for $N_f = 2$ and $N_f = 2+1+1$ respectively

charm quarks in the sea [97]. Clearly, $N_f = 2+1+1$ simulations at the physical pion mass and a mass splitting of 140 MeV are not feasible. This would mean, therefore, that a reduction by at least a factor of 2 in the lattice spacing is required to control the pion mass splitting. The possibly resulting issues with topology freezing and the (currently) infeasibly large lattice volumes make proceeding with the unimproved Wilson twisted mass action impossible. This is the reason for exploring the twisted mass clover action to reduce the pion mass splitting and enable simulations at the physical pion mass at reasonable lattice spacings and manageable lattice volumes. The beneficial qualities of twisted mass quarks would then enable the computation of a wide range of physical quantities and a further test of universality.

Summary and Conclusions

The available theoretical arguments suggest that with Wilson quarks, simulations at or close to the physical pion mass should be possible as long as the magnitude of c_2 is not too large. The Iwasaki gauge action has proven to be effective in reducing the strength of the first order phase transition and it will therefore also be used in this thesis. It is hoped that the inclusion of the clover term with Wilson twisted mass quarks will further reduce the pion mass splitting in dynamical simulations with $N_f = 2$ and $N_f = 2+1+1$ twisted mass quarks and thus allow pion masses at or close to the physical pion mass to be reached.

Part II

PARAMETER TUNING AND ENSEMBLE GENERATION

The European Twisted Mass Collaboration has carried out two very successful research programmes with $N_f = 2$ and $N_f = 2+1+1$ Wilson twisted mass fermions. As discussed in Section 2.2, the actions so far employed did not allow simulations at the physical average light quark mass. In the following, the first simulations using Wilson twisted mass clover fermions at a physically light charged pion mass will be presented. In Chapter 3, $N_f = 2$ simulations will be discussed which serve as a proof of concept for a more involved $N_f = 2+1+1$ simulation programme discussed in Chapter 4. The latter discussion will focus on the involved parameter tuning for these four flavour simulations.

From the point of view of the simulation algorithms and their optimisation, working at physically light quark masses is especially challenging because severe fluctuations are seen in the molecular dynamics forces. These decrease acceptance rates and increase simulation cost beyond what is expected due to more costly inversions. Mass preconditioning as a strategy for dealing with these fluctuations will be discussed in Chapter 5. Specifically, the behaviour of the forces as a function of the mass preconditioning parameters is studied and empirical fits are used to suggest optimisation schemes which relate all the paramters to each other, allowing the forces to be changed in seemingly predictable ways. It is hoped that the proposed schemes will be useful in $N_f = 2+1+1$ simulations at the physical pion mass, as these are even more challenging because of the many scales involved.

3

TWO FLAVOUR SIMULATIONS

The push towards the physical charged pion mass using twisted mass clover quarks was begun with $N_f = 2$ simulations as a proof of concept. They provide a demonstration of the stability of the new action in molecular dynamics, all the way down to the physical pion mass at a lattice spacing of around 0.092 fm. Being mass degenerate, the $N_f = 2$ action has only four bare parameters that need to be tuned. In this chapter, these simulations will be presented and their salient features identified and compared to old $N_f = 2$ simulations without the clover term, where possible. Parts of this chapter were already presented in Refs. [P1, P5, P6].

The combination of Iwasaki gauge action and variously improved Wilson quarks, naturally without the twisted mass term, was studied by the CP-PACS [130] and PACS-CS/JLQCD [131] collaborations in two and three flavour simulations respectively. The bare inverse gauge coupling $\beta = 2.1$ was thus chosen from data tabulated in Ref. [132], while the Sheikholeslami-Wohlert [58] (SW) coefficient $c_{sw} = 1.57551$ was determined through Padé fits to the data presented there. The ensembles simulated based on this choice with $N_f = 2$ Wilson twisted mass clover quarks are listed in Table 3.1. Some basic measurements in lattice units are also provided. Because they will be used throughout this thesis, it is worthwhile to describe the meaning of the ensemble identifiers

$$\mathrm{ID}(c_{sw}, \beta, N_f, \mu_\ell, L) \equiv [c](A \ldots Z)(N_f)[a \ldots z].(a\mu_\ell \cdot 10^4).(L/a)\,, \qquad (3.1)$$

where the presence of a clover term is indicated by a *c* at the beginning of the identifier and different values of β are associated with upper-case letters of the alphabet. An optional lower-case letter following the number of flavours is used to differentiate between replica or between ensembles with the same physical parameters, but different algorithmic settings. In this notation it is implicitly assumed that ensembles with a clover term will use a value of c_{sw} that has been tuned according to some prescription. Ensembles without an identifier were used only for tuning purposes and are usually quite short, they are given here for reference only.

Tuning to maximal twist and the extrapolation towards the physical light quark mass will be described in Section 3.1. Noteworthy features of the simulations are discussed in Section 3.2 and the choice of algorithmic parameters for the HMC is presented in Section 3.3.

PARAMETER TUNING

In simulations using twisted mass quarks, the first parameter to be tuned is the hopping coefficient κ. For the purpose of automatic $\mathcal{O}(a)$-improvement, an estimate of the critical value of the hopping parameter can be determined from Ref. [132] with a relatively large error and without the requisite shift due to the twisted quark mass. This value was refined through a number of short simulations at $a\mu_\ell = 0.006, 0.003$, corresponding to charged pion masses of around 340 and 250 MeV respectively. This resulted in a new estimate $\kappa_c \sim 0.1373$, at which the PCAC quark mass was demonstrably sufficiently small for these light quark masses. This should be compared to

ID	L/a	T/a	κ	$a\mu_\ell$	am_{PCAC}	$aM_{\pi^\pm}$	af_π
cA2.60.24	24	48	0.1373	0.006	−0.00022(3)	0.1599(3)	0.0699(2)
cA2.60.32	32	64	0.1373	0.006	−0.00020(2)	0.1582(2)	0.0706(1)
cA2.30.24	24	48	0.1373	0.003	−0.00023(7)	0.113(2)	0.063(1)
cA2.09.48	48	96	0.13729	0.0009	0.00008(1)	0.06236(7)	0.06032(9)
cA2.09.64†	64	128	0.1372938	0.0009	0.00004(2)	0.0620(2)	0.0610(4)
–	24	48	0.13728	0.0009	−0.00013(30)	0.062(9)	NA
–	48	96	0.13728	0.0009	0.00045(10)	0.0646(6)	0.0566(8)
–	48	96	0.137295	0.0009	−0.00010(10)	0.063(2)	0.060(2)
–	64	128	0.13729	0.0009	0.00013(2)	0.0626(7)	0.0604(4)

Table 3.1: $N_f = 2$ twisted mass clover ensembles and measurements of a number of basic observables. The identifiers are explained in the body of the text around Equation (3.1). †: At the time of writing, this ensemble was very short and will not be discussed further, is is given here for reference and in order to discuss possible finite size issues with am_{PCAC}.

simulations with twisted mass quarks without the clover term in which, at comparative light quark masses, the fifth and sixth digits of κ_c had to be tuned.

Approaching the physical pion mass was carried out through only two $L/a = 24$ simulations with $a\mu_\ell = 0.006, 0.003$ at maximal twist. An extrapolation in $M_{\pi^\pm}/f_\pi$ to its physical value, corrected for finite size effects as per Ref. [133], was used to estimate the bare physical average light quark mass $a\mu_\ell = 0.0009$. Uncorrected measurements of $M_{\pi^\pm}/f_\pi$ on the different ensembles are given in Table 7.1 on page 101.

When the volume was increased and the quark mass reduced to $a\mu_\ell = 0.0009$, the estimate of κ_c had to be refined further through a linear interpolation in $1/2\kappa$, as shown in Figure 3.1 (at that time with fewer points and significantly larger statistical errors, of course). As was discussed in Ref. [134] and shown in subsequent simulation results, it is sufficient to tune the renormalised PCAC quark mass to be no larger than 10% of the renormalised twisted quark mass to ensure $\mathcal{O}(a)$ improvement in practice. It should be noted that for ensemble *cA2.09.48*, the value of $\kappa = 0.13729$ is only just acceptable, giving $am_{PCAC} \sim 0.00008(1)$ after several thousand trajectories.

As the volume was increased further for *cA2.09.64*, the measured $am_{PCAC} \sim 0.00013(2)$ was clearly too large, necessitating further adjustment. This apparent volume dependence is somewhat surprising given the properties of the PCAC relation and the lattice Ward-Takahashi identities derived from it and should be carefully monitored in future simulations. One should note that at the time of writing, however, ensemble *cA2.09.64* was very short and the statistical error on am_{PCAC} is likely underestimated. The ensemble will not be used further in the following but is provided for reference.

These observed irregularities might suggest that tuning to maximal twist needs to be performed very carefully at the physical light quark mass. From the practical point of view of the tuning procedure itself, however, it seems that the clover term improves the situation substantially. In past simulations without a clover term and at much higher light quark masses, it was seen that the linear behaviour shown in Figure 3.1 breaks down around κ_c. This made interpolations much more difficult than observed here and required many simulation points around the critical region. As a result, although it

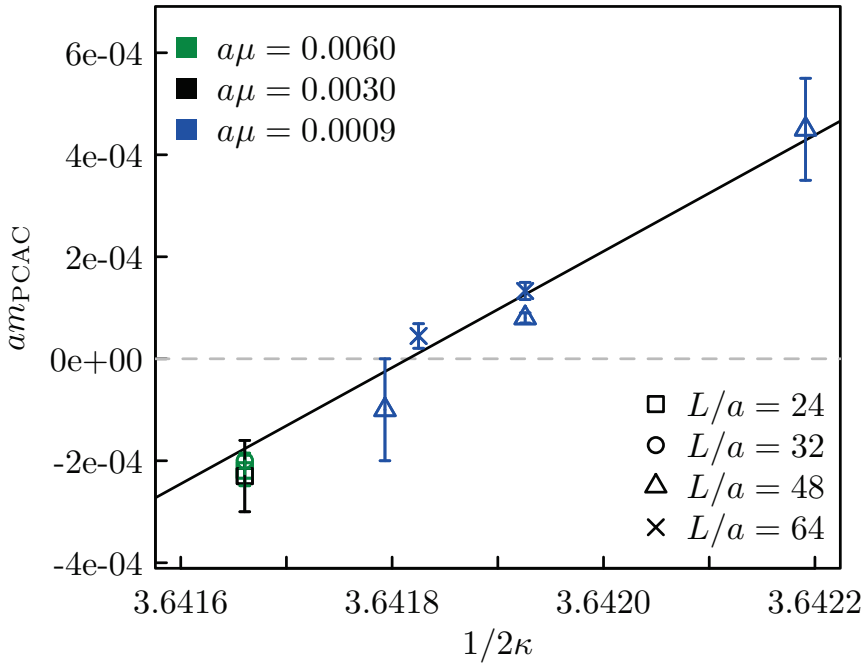

Figure 3.1: am_{PCAC} as a function of $1/2\kappa$ for the tuning of ensembles *cA2.09.48* with $\kappa = (2am_0 + 8)^{-1}$.

may be true that the tuning procedure at the physical pion mass is quite delicate and finite size effects might play a role, the tuning is practically much more straightforward, at least for the $N_f = 2$ case considered here.

Following the reasoning presented in Section 2.2.2, this absence of non-linearities and the slope of around 1 in Figure 3.1 may suggest that c_2 might be very close to 0, ignoring the value of Z_A for the moment. The indication that c_2 might be very small is further reinforced by the fact that the points with $a\mu_\ell = 0.006, 0.003$ fall on top of each other and on top of the best fit curve, suggesting that the μ_ℓ-dependent correction is almost or completely absent. Since the lattice spacing of $a \sim 0.092$ fm is likely to be the coarsest one used in the $N_f = 2$ simulation programme, the tuning procedure can only become easier as the lattice spacing is decreased. However, without simulations at different lattice spacings, it could well be that the observed behaviour is the result of some cancellations due to lattice artefacts of higher order. These features and how they relate to the pion mass splitting are discussed in some more detail in Chapter 6. Measurements of benchmark observables in the pseudoscalar meson sector on these ensembles and indirect observations on the size of lattice artefacts are presented in Chapter 7. In the next section, noteworthy features in the molecular dynamics histories of ensemble *cA2.09.48* will be discussed.

MOLECULAR DYNAMICS HISTORIES

As discussed in Section 2.2.4, in past simulations, lattice artefacts rendered simulations without a clover term meta-stable as the pion mass was lowered towards its physical value. As can be seen in Figure 3.2, the twisted mass clover action results in very stable molecular dynamics histories at the physical average up/down quark mass without any signs of meta-stability in the plaquette or the PCAC quark mass, despite the relatively coarse lattice spacing of around 0.092 fm. As expected at this lattice spacing, the topological charge in the field-theoretic definition, computed from gauge

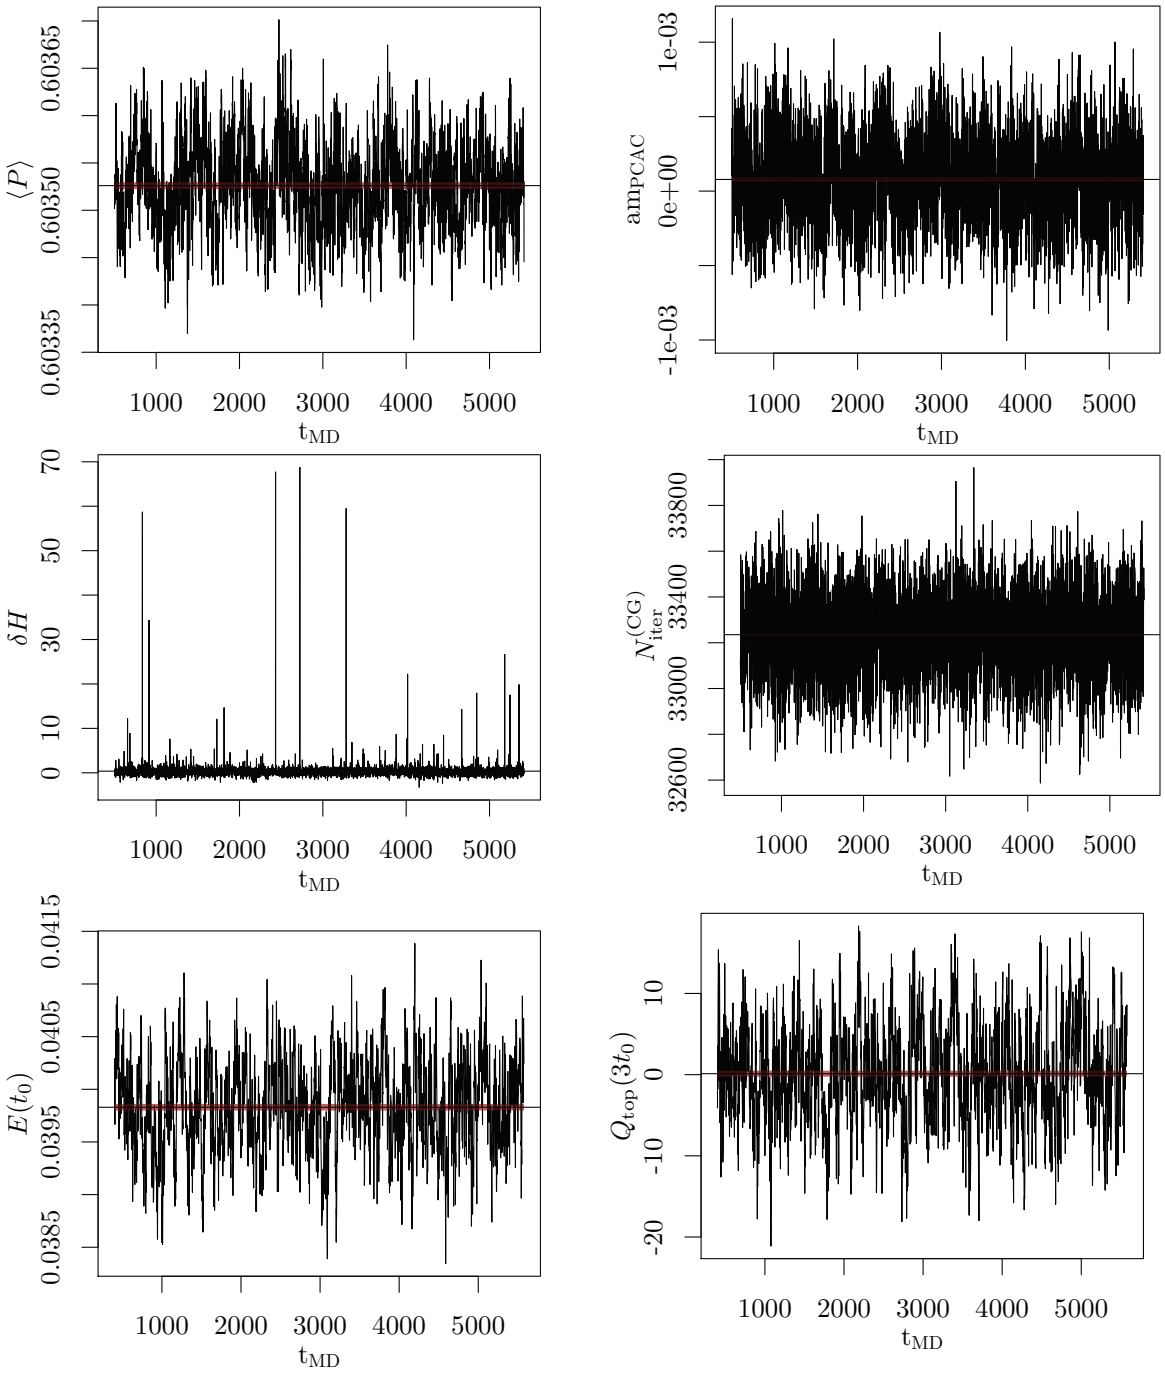

Figure 3.2: Molecular dynamics histories of various quantities on ensemble *cA2.09.48*.

observable	*cA2.09.48*	*cA2.30.24*	*cA2.60.24*	*cA2.60.32*
P_{acc}	0.726(6)	0.910(7)	0.771(5)	0.874(4)
$\langle P \rangle$	0.603526(4)	0.603562(9)	0.603535(5)	0.603533(2)
$\langle m_{PCAC} \rangle$	0.00008(1)	−0.00037(7)	−0.00026(3)	−0.00021(1)
$\langle \delta H \rangle$	0.37(3)	0.047(12)	0.177(8)	0.044(3)
$\langle \exp(-\delta H) \rangle$	1.00(1)	1.01(1)	1.003(7)	1.003(3)
$\langle Q(3t_0) \rangle$	0.1(4)	–	–	–
$\langle N_{iter}^{(CG)} \rangle$	33235(3)	10720(67)	5288(2)	5674(1)
$\tau_{int}\{P\}$	15(5)	3.2(8)	3.8(7)	2.9(5)
$\tau_{int}\{m_{PCAC}\}$	15(5)	1.6(4)	1.4(2)	1.2(1)
$\tau_{int}\{\delta H\}$	0.50(4)	0.50(2)	0.53(3)	0.50(1)
$\tau_{int}\{\exp(-\delta H)\}$	0.49(2)	0.48(2)	0.49(1)	0.50(1)
$\tau_{int}\{Q(3t_0)\}$	9(3)	–	–	–
$\tau_{int}\{E(t_0)\}$	14(4)	–	–	–
$\tau_{int}\{N_{iter}^{(CG)}\}$	0.83(9)	17(9)	1.8(2)	4.2(8)

Table 3.2: Expectation values and autocorrelation times of various observables for ensembles used in this study. $N_{iter}^{(CG)}$ refers to the number of CG iterations in the heat-bath and acceptance steps of the mass preconditioning determinant ratio which has the target light quark mass in the numerator. P_{acc} refers to the acceptance rate which should be used to scale the autocorrelation times which are given in units of trajectories.

fields smoothed by the Wilson flow [135] at flow time $3t_0$, shows a low autocorrelation time of

$$\tau^*_{int}\{Q(3t_0)\} = P_{acc}^{-1}\tau_{int}\{Q(3t_0)\} \sim 12(4) \tag{3.2}$$

trajectories. The plaquette expectation value, the PCAC quark mass and the energy density (at flow time t_0) are sampled equally well and show autocorrelation times well below one hundredth of the total simulation time. A complete listing of various algorithmic observables from the different ensembles is provided in Table 3.2 together with their autocorrelation times. The algorithmic parameters that were used in the HMC are discussed in Section 3.3. For ensemble *cA2.30.24*, it is notable that there are significant fluctuations in the number of CG iterations in the heatbath and acceptance steps for the mass preconditioning determinant ratio with the target quark mass in the numerator. These can be interpreted as an indication of an insufficient volume for the simulated pion mass.

The number of CG iterations in the simulation at the physical pion mass is very significant with $\sim 3 \cdot 10^4$ iterations in the heatbath and acceptance step of the determinant ratio with the target mass in the numerator. In the molecular dynamics, about $3 \cdot 10^5$ CG iterations are required in total for the computation of derivatives of this determinant ratio, although better mass preconditioning schemes can reduce this by a factor of two or more. While these numbers are substantial, the CG algorithm is very efficient for the twisted mass operator at maximal twist and it is unlikely that much would be gained through eigenvalue deflation. However, when lattices of spatial extent $L = 64$ or larger are considered, the practicality of simulating them certainly hinges on the availability of better linear solvers, an as yet unsolved problem for twisted mass quarks at maximal twist.

A feature of the simulation which deserves a special mention is the behaviour of the energy violation, denoted by δH. It seems that compared to twisted mass simulations without the clover term, large deviations from 0 occur quite frequently. However, they do not seem to affect the stability of the algorithm. There do not appear to be any measurable effects of these spikes in any of the observables that have been analysed on the simulated ensembles, but of course the acceptance probability P_{acc} is affected by their presence. Finally, these kinds of deviations are in line with what has been observed by other collaborations [86, 136]. The frequency of their occurence

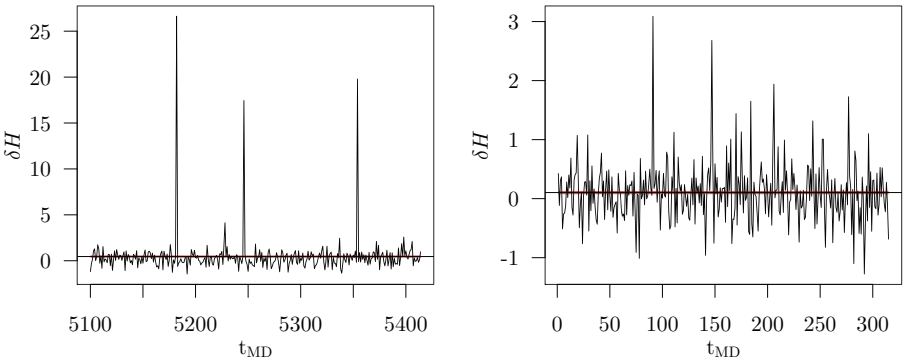

Figure 3.3: Detail of MD histories of the energy violation δH for runs *cA2.09.48*(left) and *cA2z.09.48*(right), indicating a reduction of large deviations from 0 as the number of integration steps is increased.

ensemble	N_t	$a\rho_t^{HB}$	$\log_{10}\left(\begin{smallmatrix} r_a^2 \\ r_f^2 \end{smallmatrix}\right)$
cA2.60.24	$\{1,2,2,7\}$	$\{-, 0.060, \begin{smallmatrix} 0.0110 & 0.0000 \\ 0.0600, & 0.0110 \end{smallmatrix}\}$	$\{-, \begin{smallmatrix} -22 & -22 & -22 \\ -14, & -14, & -14 \end{smallmatrix}\}$
cA2.60.32	$\{1,1,1,1,14\}$	$\{-, 0.800, \begin{smallmatrix} 0.0800 & 0.0080 & 0.0000 \\ 0.8000, & 0.0800, & 0.0080 \end{smallmatrix}\}$	$\{-, \begin{smallmatrix} -22 & -22 & -22 & -22 \\ -14, & -14, & -14, & -14 \end{smallmatrix}\}$
cA2.30.24	$\{1,2,2,10\}$	$\{-, 0.040, \begin{smallmatrix} 0.0080 & 0.0000 \\ 0.0400, & 0.0080 \end{smallmatrix}\}$	$\{-, \begin{smallmatrix} -22 & -22 & -22 \\ -14, & -14, & -14 \end{smallmatrix}\}$
cA2.09.48	$\{1,1,2,13\}$	$\{-, 0.030, \begin{smallmatrix} 0.0050 & [0.0013 & 0.0000] \\ 0.0300, & [0.0050, & 0.0013] \end{smallmatrix}\}$	$\{-, \begin{smallmatrix} -22 & -22 & \left[\begin{smallmatrix} -22 & -22 \\ -14, & -14 \end{smallmatrix}\right. \\ -14, & -14, & \end{smallmatrix}\}$
cA2x.09.48	$\{1,1,2,17\}$	$\{-, 0.030, \begin{smallmatrix} 0.0050 & [0.0013 & 0.0000] \\ 0.0300, & [0.0050, & 0.0013] \end{smallmatrix}\}$	$\{-, \begin{smallmatrix} -22 & -22 & \left[\begin{smallmatrix} -22 & -22 \\ -14, & -14 \end{smallmatrix}\right. \\ -14, & -14, & \end{smallmatrix}\}$
cA2y.09.48	$\{1,1,1,1,13\}$	$\{-, 0.250, \begin{smallmatrix} 0.0250 & 0.0025 & 0.0000 \\ 0.2500, & 0.0250, & 0.0025 \end{smallmatrix}\}$	$\{-, \begin{smallmatrix} -22 & -22 & -22 & -22 \\ -14, & -14, & -14, & -14 \end{smallmatrix}\}$
cA2z.09.48	$\{1,1,1,1,17\}$	$\{-, 0.250, \begin{smallmatrix} 0.0250 & 0.0025 & 0.0000 \\ 0.2500, & 0.0250, & 0.0025 \end{smallmatrix}\}$	$\{-, \begin{smallmatrix} -22 & -22 & -22 & -22 \\ -14, & -14, & -14, & -14 \end{smallmatrix}\}$

Table 3.3: Simulation parameters for the ensembles used in this work and three additional test ensembles. N_t: number of integration steps of second order minimal norm integrator on the various time-scales. $a\rho_t^{HB}$: Hasenbusch mass pre-conditioning parameters as in Ref. [77] but with multiple determinant ratios. $r_a^2(r_f^2)$: squared relative residual stopping criterion in the acceptance step (force calculation) in the conjugate gradients solver. Square brackets indicate that more than one monomial is placed on the same timescale.

depends on the target light quark mass and they seem to be strongly suppressed in the simulation of the ensembles with $a\mu_l = 0.006$. Similarly, away from maximal twist when the quark mass has a sizeable contribution from the standard mass term, large spikes in δH are much less frequent.

ALGORITHMIC PARAMETERS

The simulation parameters for the $N_f = 2$ ensembles used in the present work are listed in Table 3.3. The mass preconditioning parameters $a\rho_t^{HB}$ are given and these can be related to those defined Ref. [77]. The number of steps of the 2MN integrator on the various time scales is tabulated, with the rightmost number referring to the outermost time scale with the smallest number of total effective steps.

In order to determine the origin of the sizeable δH fluctuations observed in the molecular dynamics history of *cA2.09.48*, short simulations *cA2x.09.48*, *cA2y.09.48* and *cA2z.09.48* with more integration steps and more time scales were performed. As shown in Figure 3.3 it was found that this significantly reduces the magnitude of large energy violations at the price of increased simulation cost, although it should be kept in mind that the test was quite short. At the same time, the increase of the acceptance rate means that the effective increase in the simulation cost of ensemble *cA2z.09.48* compared to *cA2.09.48* is only about a factor of 1.2. It should be noted that none of

the observables that were determined on ensemble *cA2z.09.48* showed any deviation within errors compared to those on ensemble *cA2.09.48*.

The results of these checks influenced the ideas of Chapter 5 in the investigation of mass preconditioning and integration schemes with many time scales which may prove beneficial for simulations at the physical pion mass. The molecular dynamics (MD) forces of the different fermionic terms in the Hamiltonians for the simulation of ensembles *cA2.30.24* and *cA2.09.48* are investigated in Sections 5.3.2 and 5.3.3 on pages 76 and 79.

SUMMARY AND CONCLUSIONS

The suitability of the twisted mass clover action for simulations at the physical light quark mass was demonstrated with $N_f = 2$ simulations at a relatively coarse lattice spacing of $a \sim 0.092$ fm. It was seen that the tuning procedure is relatively straightforward and that the clover term seems to suppress certain lattice artefacts in the relation of the PCAC quark mass to the subtracted Wilson quark mass. At the same time however, a surprising volume dependence is seen in the value of the PCAC quark mass which should be kept in mind in future simulations.

4

FOUR FLAVOUR SIMULATIONS

The success of the twisted mass clover action in enabling $N_f = 2$ simulations at the physical pion mass provides an important feasibility test for the eventual goal of performing $N_f = 2+1+1$ simulations at physical quark masses. This is complicated by the fact that the pion mass splitting was observed in the past to increase with the number of dynamical flavours. This was so severe, in fact, that $N_f = 4$ simulations for the ETMC RI-MOM renormalisation programme could not be carried out at maximal twist and $\mathcal{O}(a)$-improvement had to be implemented out by an averaging procedure. While maximal twist could be reached in $N_f = 2+1+1$ simulations, the pion mass splitting was seen to be large, making simulations at the physical pion mass impossible at reasonable lattice spacings. It is a major goal of this thesis to investigate whether the clover term sufficiently reduces the pion mass splitting also in $N_f = 2+1+1$ simulations to allow the physical regime to be reached.

In contrast to previous four flavour simulations, the non-degenerate strange and charm quark doublet is simulated using a rational approximation rather than the polynomial approximation that was used in the past. The so-called RHMC algorithm was introduced in Refs. [137, 138] and the implementation used here closely follows that of Ref. [80]. The combination of the non-degenerate twisted mass operator, the clover term and even-odd preconditioning is described in Appendix A.2. Using the RHMC over the PHMC algorithm is beneficial because the partial fraction representation allows different terms of the approximation to be integrated on different time-scales, increasing efficiency. Further, being an optimal approximation, the accuracy is much higher for much lower orders of the employed polynomials. Therefore, the round-off errors which affected PHMC simulations and required very high polynomial orders and coefficients computed using arbitrary precision arithmetic in the accept-reject step are not an issue when using the RHMC.

One of the difficulties of $N_f = 2+1+1$ simulations with twisted mass clover quarks lies in the number of parameters that need to be tuned. In particular, whether all parameters can be tuned simultaneously while retaining control over the physical situation that these parameters correspond to is not guaranteed. In order to illustrate the reasoning behind the approach that was adopted, the tuning strategy will be laid out as it progressed in time. The tuning of β and c_{sw} will be discussed first, followed by two strategies for the tuning of strange and charm quark masses. Effective parametrisations, based on empirical fits, of the lattice spacing and the critical hopping parameter for the $N_f = 2+1+1$ twisted mass clover action will be given in terms of β and c_{sw}. Finally, a functional form for the tadpole-improved value of c_{sw} as a function of β will be introduced which was derived from these fits. To conclude the chapter, the first four flavour production simulations with these tuned parameters will be presented.

PARAMETER TUNING

Unlike in the $N_f = 2$ case, simulations with $N_f = 4$ or $N_f = 2+1+1$ Wilson clover quarks and Iwasaki gauge action have not yet been performed. Therefore, a non-perturbatively tuned value of c_{sw} is not available from the literature, nor are the values of the critical hopping parameter κ_c or the lattice spacing as a function of β . Because simulations with twisted mass quarks at maximal twist do not rely on the clover term for $\mathcal{O}(a)$-improvement of on-shell observables, a tadpole-improved estimate of c_{sw} is likely to be sufficient to stabilise simulations. This is supported by the fact that for the $N_f = 2$ case, the non-perturbatively tuned and the tadpole-improved values differed by less than 5% at a lattice spacing of $a \sim 0.092$ fm. Estimates of suitable values of β and c_{sw} for first $N_f = 2+1+1$ simulations were arrived at through exploratory runs, discussed below.

Exploratory Simulations

When the twisted mass clover project was started, the strategy for the tuning of c_{sw} in $N_f = 2+1+1$ simulations was outlined in Ref. [P1]. It involves reaching the tadpole-improved [139, 140] perturbative estimate for the Iwasaki gauge action via a number of tuning runs and the repeated application of Equation (4.1) below.

$$c_{sw} \sim 1 + 0.113(3)\frac{6}{\beta\langle P\rangle} \tag{4.1}$$

Tuning β, c_{sw} and the strange and charm quark masses in terms of the bare parameters μ_δ and μ_σ simultaneously, is likely to be excessively complicated. In addition, the stability of simulations should be checked before a comprehensive tuning programme is embarked upon. As discussed above, mass degenerate $N_f = 4$ simulations without the clover term showed very strong signs of meta-stabilities, to the point that simulations at maximal twist were not possible. It can therefore be argued that if these simulations can be successfully tuned to maximal twist with the clover term while remaining stable with small quark masses, so will $N_f = 2+1+1$ simulations.

In order to be closer to the $N_f = 2+1+1$ situation while still benefitting from the stability argument, the first exploratory four flavour simulations with the twisted mass clover action were chosen to be $N_f = 2+2$. These involved two light quarks corresponding to a pion mass of around 250 MeV and two heavier mass degenerate quarks, approximately in the region of the strange quark mass. The reasoning behind this choice was that if stability could be demonstrated, the estimates of κ_c, β and c_{sw} could likely be directly transferred to $N_f = 2+1+1$ simulations. Estimating the N_f-dependence, a value of $\beta = 1.85$ was chosen based on the lattice spacing from $N_f = 2$ simulations. Several simulations were performed, iteratively tuning the value of c_{sw} and κ. These exploratory runs showed that the action was indeed stable, but it was realised that the lattice spacing was much finer than the targeted 0.1 fm.

Further $N_f = 2+2$ simulations were carried out with $\beta = 1.7$ and a value of $c_{sw} = 1.85$ was arrived at, together with a successful tuning to maximal twist at a charge pion mass of around 250 MeV. The parameter pair

$$\beta = 1.7\,, \qquad c_{sw} = 1.85$$

would thus be used to proceed with tuning the strange and charm quark masses for true $N_f = 2+1+1$ simulations.

Tuning Strange and Charm Quark Masses

Tuning the strange and charm quark masses is another major challenge. As described in Section 1.3.4, unlike for other quark discretisations, the mass non-degenerate twisted mass doublet requires the simultaneous tuning of the bare mass parameters μ_σ and μ_δ. The effective bare strange and charm quark masses are then

$$\mu_s = \mu_\sigma - \frac{Z_P}{Z_S}\mu_\delta\,, \qquad \mu_c = \mu_\sigma + \frac{Z_P}{Z_S}\mu_\delta\,, \tag{4.2}$$

where the ratio of renormalisation constants $Z_P/Z_S < 1.0$ [67]. Clearly, tuning the strange and charm quark masses would be easiest at the physical light quark mass, but the computational expense of such an approach is currently excessive. The aim of the tuning strategy described in the following is thus to allow the tuning of the strange and charm sea quark masses (at any lattice spacing) with simulations at charged pion masses of around 250 MeV or even higher.

In the old $N_f = 2+1+1$ simulations of Ref. [97], the strange and charm quark masses were tuned to approximately reproduce the physical value of $2M_K^2 - M_{\pi^\pm}^2$ and D meson mass. This tuning strategy is somewhat complicated because of heavy flavour and parity mixing with the non-degenerate (ND) twisted mass quark action at finite lattice spacing [67]. The resulting strange-charm, scalar-pseudoscalar mixed flavour-parity sector is denoted by $(s/c, -/+)$. The kaon mass can be extracted easily as the lowest state of the generalized eigenvalue problem in this sector, but the D meson is a highly excited state, requiring a large correlator basis and even then the uncertainty is quite large. Despite the careful tuning of Ref. [97], some of the $N_f = 2+1+1$ ensembles without the clover term had sea strange quark masses too large by about 20% and this had to be accounted for in later analyses, adding further systematic errors.

Because the values of μ_σ and μ_δ are difficult to guess à priori, the tuning of the strange and charm quark masses for the present new set of simulations was begun on the $N_f = 2+2$ gauge configurations described above, with parameters and sea quark masses

$$\begin{aligned} \beta = 1.7\,, \quad c_{sw} = 1.85\,, \quad \kappa = 0.13872\,, \\ a\mu_\ell = 0.003\,, \quad a\mu_s = 0.025\,, \quad T/2a = L/a = 24\,. \end{aligned} \tag{4.3}$$

These were used to compute quark propagators with Osterwalder-Seiler (OS) valence quarks with masses in the strange ($a\mu \sim 0.025$) and charm ($a\mu \sim 0.27$) regions. From there, various pseudoscalar quantities involving light, strange and charm quarks were computed. The techniques for their computation are discussed in Chapter 7 and will not be outlined here. An estimate of the lattice spacing $a \sim 0.095$ fm, was determined from gluonic scales using literature inputs for $N_f = 2+1$.

The strategy now consisted of finding appropriate ratios of pseudoscalar quantities involving strange and charm quarks which do not depend strongly on the light quark mass. The phenomenological values of these could then be used to interpolate in the OS strange and charm quark masses in order to tune them. Then, the masses so determined immediately give a valence estimate of the average twisted quark mass μ_σ of the non-degenerate (ND) doublet

$$\mu_\sigma = \frac{\mu_c + \mu_s}{2}\,, \tag{4.4}$$

without any knowledge of Z_P/Z_S. Assuming a value of Z_P/Z_S around 0.75, ND quark propagators can now be computed at various values of μ_δ, constrained by the requirement $\mu_c/\mu_s \sim 12$. Since the kaon mass can be computed easily as the lowest energy state in the ND mixed flavour-parity sector, this can be used as a matching condition between the Osterwalder-Seiler and non-degenerate valence actions. Interpolating in μ_δ to match $M_K^{OS} = M_K^{ND}$ fixes the parameters of the ND action in the valence sector.

These parameter values can then be used in an $N_f = 2+1+1$ simulation with approximately physical strange and charm sea quark masses. On these new configurations, Osterwalder-Seiler and ND quark propagators are computed again and the tuning conditions are checked. If any deviations are found, the ND mass parameters are adjusted accordingly, leading to good estimates of the physical strange and charm quark masses after one or two iterations.

As an aside it should be noted that at this point, it may be necessary to adjust c_{sw} to satisfy Equation (4.1), because moving from $N_f = 2+2$ to $N_f = 2+1+1$ will affect the effective coupling. This will in turn affect the lattice spacing and complicate the tuning strategy. However, as will be shown in Section 4.2 below, this can be completely avoided once a sufficient number of $N_f = 2+1+1$ ensembles have been generated. For the purpose of the initial sea quark mass tuning, the parameters $\beta = 1.7, c_{sw} = 1.85$ were retained and two tuning conditions for the strange and charm quark masses were considered.

Tuning Condition 1

The first set of tuning conditions unfortunately did not lead to a satisfactory quark mass tuning, but it will be described here for reference and completeness. It involves the approximate charm to strange quark mass ratio, the mass M_{sc} of the D_s meson and the mass $M_{ss'}$ of a fictitious pseudoscalar meson with two "strange" constituent quarks. Interpolating in the OS quark masses to simultaneously satisfy the conditions

$$A_1 : \frac{\mu_c}{\mu_s} \sim 12\,, \qquad B_1 : \left(\frac{M_{ss'}}{M_{sc}}\right)^2 \sim 0.121(3)\,, \tag{4.5}$$

was expected to give a good estimate of the physical values of the strange and charm quark masses. It was hoped that lattice artefacts would cancel appreciably in the ratio B_1, which turned out not to be the case. This is presented in Figure 4.1, which shows a fit inspired by lowest order SU(3) χPT and heavy quark effective theory, of the form

$$\left(\frac{M_{ss'}}{M_{sc}}\right)^2 \sim \left(\frac{\alpha\mu_s}{\mu_c^2 + \beta}\right) + \gamma\,, \tag{4.6}$$

describing the data rather well ($\chi^2/\mathrm{df} \sim 1$).

Also indicated are the continuum limit estimate of the ratio $(M_{ss'}/M_{sc})^2$, as well as the fit result along the line $\mu_c = 11.85(16) \cdot \mu_s$, given by the charm to strange quark mass ratio of Ref. [141]. Along this latter line it is seen that within the region of interest, the ratio $(M_{ss'}/M_{sc})^2$ is rather insensitive to the strange quark mass and thus not very suitable for this purpose. The intersection point of these two lines gives the quark mass estimates

$$a\mu_s \sim 0.031\,, \qquad a\mu_c \sim 0.367\,,$$

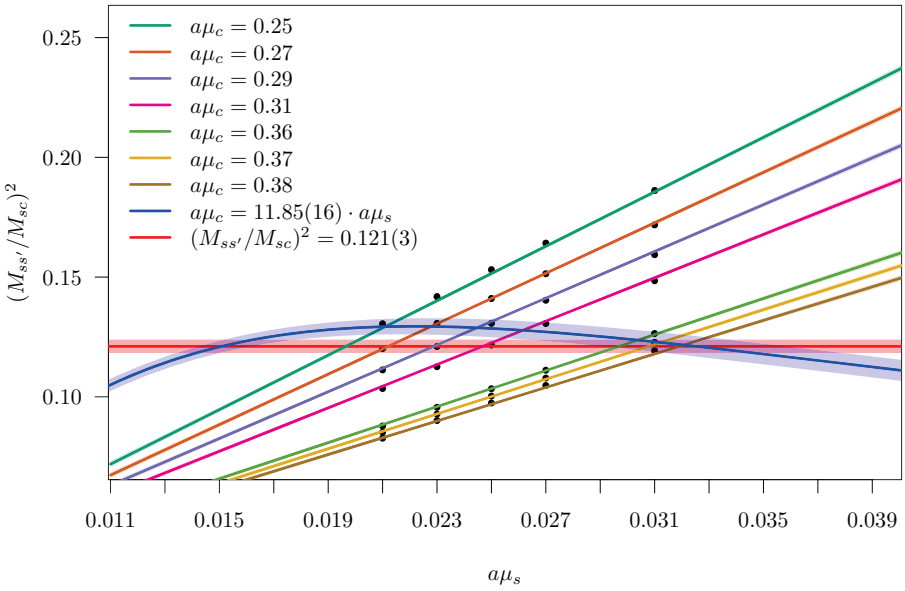

Figure 4.1: Fit of Equation (4.6) to lattice data for the purpose of tuning strange and charm quark masses by matching the strange to charm quark mass ratio of Ref. [141] and the continuum limit value of $(M_{ss'}/M_{sc})^2$ simultaneously.

both of which are far too large, with the strange quark mass likely corresponding to around 140 MeV or so. The second intersection point on the far left should be considered as an artefact of the model, as the fit is unlikely to reproduce the actual ratio there and the resulting quark masses are far too small. It was thus concluded that these tuning conditions do not offer the requisite sensitivity and may be subject to unwanted lattice artefacts in the ratio of pseudoscalar meson masses.

Tuning Condition 2

The second tuning condition is similar but uses the ratio of the D_s meson mass and decay constant

$$A_2 : \frac{\mu_c}{\mu_s} \sim 12 \qquad B_2 : \frac{M_{D_s}}{f_{D_s}} \sim 7.9(2)\,, \tag{4.7}$$

where it is hoped that discretisation errors in B_2 cancel to a larger degree than in B_1 above. This ratio is similarly insensitive to the light quark mass and both quantities can be extracted with small statistical errors in the twisted mass discretisation at maximal twist. Using these conditions, the mass estimates

$$a\mu_s \sim 0.021\,, \qquad a\mu_c \sim 0.25$$

were obtained, which give $a\mu_\sigma = 0.1355$. Matching the OS and ND kaon masses subsequently gave $a\mu_\delta = 0.145$, providing a first set of parameter values for $N_f = 2+1+1$ simulations

$$\beta = 1.7\,, \quad c_{sw} = 1.85\,, \quad a\mu_\sigma = 0.1355\,, \quad a\mu_\delta = 0.145\,. \tag{4.8}$$

First Four Flavour Simulation Parameters

With the strange and charm sea quark masses fixed to approximately physical values, simulations could now proceed in principle. As indicated above, however, the value of c_{sw} had to be adjusted to satisfy Equation (4.1). This affected the lattice spacing sufficiently strongly to also require an adjustment of β and subsequently μ_σ and μ_δ required retuning. This was

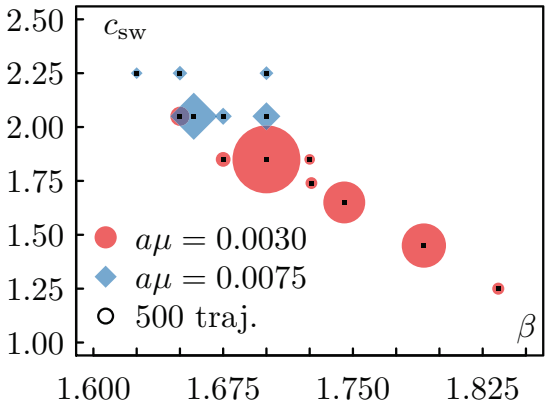

Figure 4.2: Overview of ensembles used for the determination of the functional dependence of the lattice spacing, critical hopping parameter and plaquette expectation value on β and c_{sw} as well as the tadpole-improved value of c_{sw} as a function of β.

done in the valence sector using $N_f = 2+1+1$ gauge configurations with parameter values

$$\beta = 1.726\,, \quad c_{sw} = 1.74\,, \quad a\mu_\sigma = 0.1355\,, \quad a\mu_\delta = 0.145\,, \tag{4.9}$$

and yielded updated values for

$$a\mu_\sigma = 0.1408\,, \qquad a\mu_\delta = 0.1521\,. \tag{4.10}$$

Simulations based on this parameter set will be further described in Section 4.3. The values of β and c_{sw} were derived from empirical fits to a large number of $N_f = 2+1+1$ simulations which had become available at this point. They were based on estimates of the lattice spacing computed in the context of the study presented in Section 6.2. These fits will be described below, before moving on to the first $N_f = 2+1+1$ twisted mass clover production simulations.

EMPIRICAL FITS OF BARE LATTICE PARAMETERS

In Chapter 6 the effect of the value of c_{sw} on the pion mass splitting is explored in $N_f = 2+1+1$ simulations with approximately constant physics. To proceed at constant lattice spacing and at maximal twist, β, c_{sw} and κ had to be tuned accordingly. For this purpose, 77 short simulations on $T/2a = L/a = 24$ lattices were carried out with various combinations of $\beta \in [1.625, 1.834]$ and $c_{sw} \in [1.25, 2.25]$. In order to keep the computational and tuning demands low, the bare parameters $a\mu_\ell$, $a\mu_\sigma$ and $a\mu_\delta$ were kept constant even though this of course resulted in simulations with vastly different dynamical quark content. It turns out that for the present purpose, the resulting systematic error appears to be surprisingly small. Most of the simulations were carried out with $a\mu_\ell = 0.0075$, corresponding to around 400 MeV pions at $a \sim 0.095$ fm. The four ensembles that were to be used for the measurement of the pion mass splitting were simulated with $a\mu_\ell = 0.003$, or equivalently a charged pion mass of around 250 MeV.

Tuned to maximal twist, 14 of these simulations were used to measure w_0/a [142] in the Wilson flow framework with the clover definition of the energy density[1]. The uncertainty on w_0/a was computed taking into account autocorrelations using the Gamma method [144]. This set of simulations is listed in Table 4.1 and a summary is shown in Figure 4.2, where the size of the symbol indicates the number of thermalised trajectories of unit

[1] A measurement routine for t_0/a^2 and w_0/a was added to the tmLQCD software suite [C1, 143] for this purpose while an analysis routine was added to [C2].

c_{sw}	β	$a\mu_\ell$	N_t	am_{PCAC}	a [fm]	w_0/a
1.25	1.834	0.0030	410	0.00020(40)	0.0987(14)	1.773(12)
1.45	1.791	0.0030	1500	0.00029(16)	0.0986(14)	1.777(10)
1.65	1.745	0.0030	1421	0.00027(13)	0.0993(15)	1.768(14)
1.74	1.726	0.0030	600	−0.00049(27)	0.0963(14)	1.823(13)
1.85	1.675	0.0075	500	0.00010(30)	0.1118(17)	1.569(12)
1.85	1.700	0.0030	2000	−0.00043(08)	0.0969(12)	1.810(08)
1.85	1.725	0.0075	360	0.00030(20)	0.0914(15)	1.920(17)
2.05	1.650	0.0075	650	−0.00043(28)	0.1010(15)	1.737(13)
2.05	1.658	0.0030	1400	−0.00039(14)	0.0944(13)	1.860(11)
2.05	1.675	0.0075	485	−0.00064(22)	0.0900(11)	1.951(07)
2.05	1.700	0.0030	950	−0.00043(15)	0.0771(10)	2.278(13)
2.25	1.625	0.0075	400	0.00048(26)	0.0923(11)	1.902(07)
2.25	1.650	0.0075	500	−0.00010(20)	0.0798(17)	2.200(30)
2.25	1.700	0.0030	470	−0.00040(10)	0.0586(20)	2.995(34)

Table 4.1: $N_f = 2+1+1$ ensembles used for the estimate of the lattice spacing a and critical hopping parameter κ_c as a function of c_{sw} and β. N_t is the number of thermalised trajectories.

length. The $N_f = 2+1$ value of w_0 from Ref. [142] was used to obtain an estimate of the lattice spacing, ignoring the N_f and pion mass dependence. In addition, the expectation value of the plaquette as a function of c_{sw} and β was recorded. The remaining simulations at different values of κ were used together with these 14 to interpolate am_{PCAC} to 0 for each combination of c_{sw} and β, resulting in rather good estimates of the critical hopping parameter at these points.

In the literature, $N_f = 2+1+1$ simulations with the Iwasaki gauge action and clover-improved Wilson fermions have not yet been performed. As a result, the dependence of the lattice spacing and κ_c on β and c_{sw} is presently unknown. In order to establish a programme for $N_f = 2+1+1$ simulations using the twisted mass clover action, these dependences should first be investigated. It will be shown that knowledge of $\langle P \rangle$ as a function of β and c_{sw} can be used to determine c_{sw} as a function of β directly to an apparently sufficient precision, without requiring further tuning. Of course, this is only a perturbative result which, in particular, only captures the N_f-dependence via the plaquette expectation value. However, $N_f = 2$ simulations with twisted mass clover fermions and Iwasaki gauge action seem to indicate that it deviates only by a few percent from the non-perturbative value for $\mathcal{O}(a)$-improvement of the standard Wilson theory and should thus be sufficient to stabilise simulations.

The functional forms in Equations (4.11) to (4.13) were fitted to the data from the ensembles of Table 4.1. The fits that were used can be motivated by considering the fact that c_{sw} can be perturbatively expanded as a function of g^2 and that the lattice β-function for the lattice spacing depends exponentially on $-\beta$ at leading order. As a result, the exponential dependence provides a non-perturbative estimate which is likely to be valid in the range of c_{sw} and β used for the fit and perhaps somewhat beyond. Errors and covariances of the parameters were estimated by repeating the fits on simulated datasets with Gaussian distributions. A fit to the interpolated values of κ_c leads to the following parametrisation

$$\ln(\kappa_c) = -0.12727(5)c_{sw} - 0.20339(20)\beta - 1.3931(5)\,, \tag{4.11}$$

in excellent agreement with the data as shown in Figure 4.3a. Of course, this fit does not reproduce the perturbative value in the limit $\beta \to \infty$, but in the relevant parameter region it is much better than a constrained fit which does.

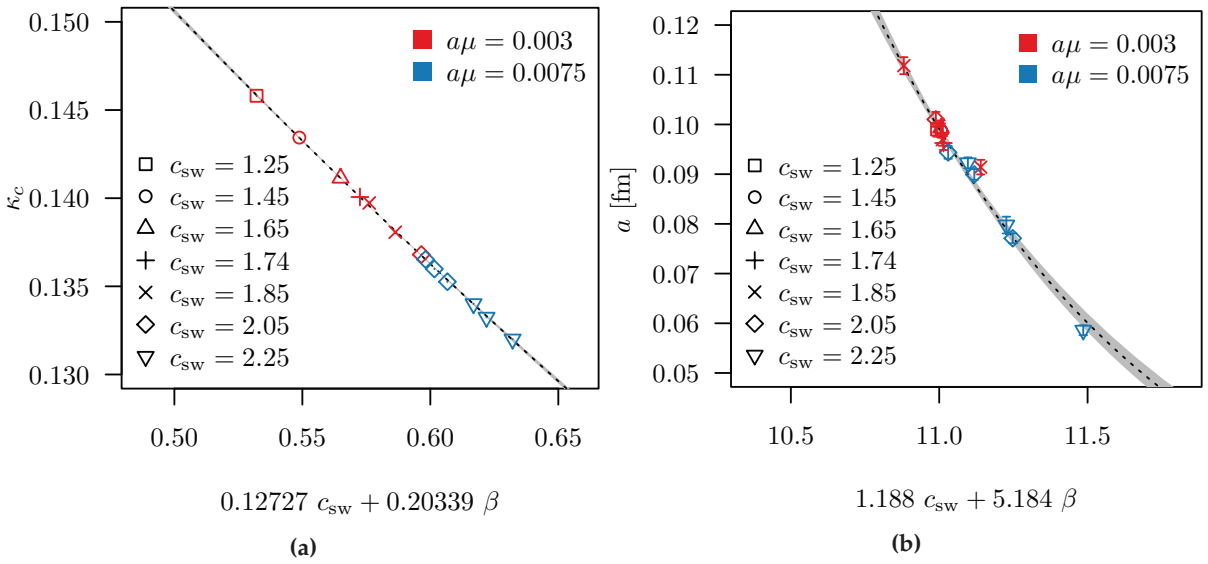

Figure 4.3: (a): Dependence of κ_c the parameters β and c_{sw}. **(b)**: Dependence of the lattice spacing determined from the Wilson flow on the parameters β and c_{sw}.

Similarly, the values of the lattice spacing derived from w_0 as explained above were fitted with the following function

$$\ln(a/\text{fm}) = -1.188(44)c_{sw} - 5.184(234)\beta + 8.68(46)\,, \tag{4.12}$$

giving very good agreement with the data as shown in Figure 4.3b. For the plaquette, the following three-parameter fit was used

$$\begin{aligned}\langle P\rangle &= p_0 + p_1 c_{sw} + p_2\beta \\ \langle P\rangle &= -0.045(1) + 0.3126(6)c_{sw} + 0.0265(1)\beta\,,\end{aligned} \tag{4.13}$$

where it should be kept in mind that it will only be valid in relatively small range of β and c_{sw} in which it agrees very well with the data, as is clear from Figure 4.4a. This linear approximation was used in order to derive the results below.

Using Equation (4.13) in Equation (4.1) above, an estimate for the tadpole-improved value of c_{sw} as a function of β can be obtained by solving the resulting quadratic equation

$$p_1\beta c_{sw}^2 + (p_2\beta^2 + p_0\beta - p_1\beta)c_{sw} - (p_2\beta^2 + p_0\beta + 6\cdot 0.113(3)) = 0\,.$$

The solutions are

$$\begin{aligned}c_{sw} = \frac{1}{2p_1\beta}\Big(&-p_0\beta - \beta(p_2\beta - p_1) \pm \\ &\sqrt{(\beta(p_2\beta - p_1) + p_0\beta)^2 - 4p_1\beta(-p_2\beta^2 - \beta p_0 - 0.678(18))}\Big)\,,\end{aligned} \tag{4.14}$$

of which only the positive branch makes sense. Without propagating the errors on the various fit parameters, this is

$$\begin{aligned}c_{sw} = 0.428 - 0.0424\beta + \\ \sqrt{2.169\beta^{-1} + 0.3272 + 0.04849\beta + 0.001797\beta^2}\,,\end{aligned} \tag{4.15}$$

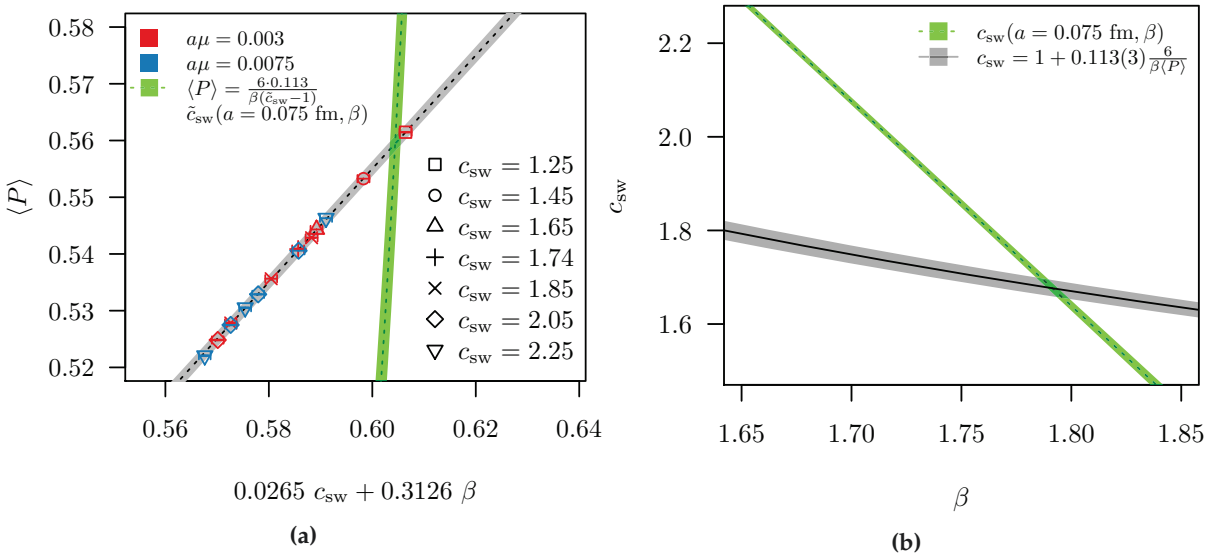

Figure 4.4: (a): Dependence of the plaquette expectation value on the parameters β and c_{sw}. The green band follows the curve in β when c_{sw} is fixed such as to give $a = 0.075$ fm according to the fit in Equation (4.12). **(b)**: The green band from (a) in the β-c_{sw} plane intersected the line of solutions in Equation (4.15).

which approaches 1 for $\beta \gg 1$, but diverges to $-\infty$ in the limit $\beta \to \infty$. The range of β, c_{sw} relevant for practical simulations should lie well within the range of data that was used for the fit. Hence, the value given by Equation (4.15) should be a good estimate for the tadpole-improved value of c_{sw} for practical simulations with $N_f = 2+1+1$ Wilson clover fermions with the Iwasaki gauge action. The line of solutions for c_{sw} is shown in Figure 4.4b with the propagated uncertainty, the meaning of the intersecting green line will be clarified below. In Figure 4.5 c_{sw} is shown for a range β far outside the original fit range to give an idea of the asymptotic behaviour of Equation (4.15) (the estimate of the error band is too small to be visible). Finally, Equation (4.15) can be used in Equations (4.11) and (4.12) to give the approximate lattice spacing and critical hopping parameter with this value of c_{sw} as a function of β only.

It is clear that these functions have limited accuracy and correctness, given that they were not fitted to data on lines of constant physics. However, the fact that data with vastly different dynamical quark content falls so well on the fitted curves suggests that the choice of c_{sw} and β and the tuning of μ_σ and μ_δ can be carried out largely independently. The rather substantial N_f and pion mass dependence of w_0 is well known, but it is reasonable to expect that the estimate of the lattice spacing is good to within a few percent, sufficient for choosing potential simulation points. The estimates of κ_c have proven accurate to within less than a per-mille, suggesting that the tuning effort will be substantially reduced for future simulation points if the presented formulae are used as first guesses. The uncertainties and (unquoted) systematic errors entering Equation (4.15) via Equation (4.13) appear to be smaller than or similar to the 3% theoretical uncertainty on the perturbative factor in Equation (4.1). It should thus be possible to use the presented formulae to suggest values of β, c_{sw} and κ for the entire $N_f = 2+1+1$ twisted mass clover simulation programme of the ETMC.

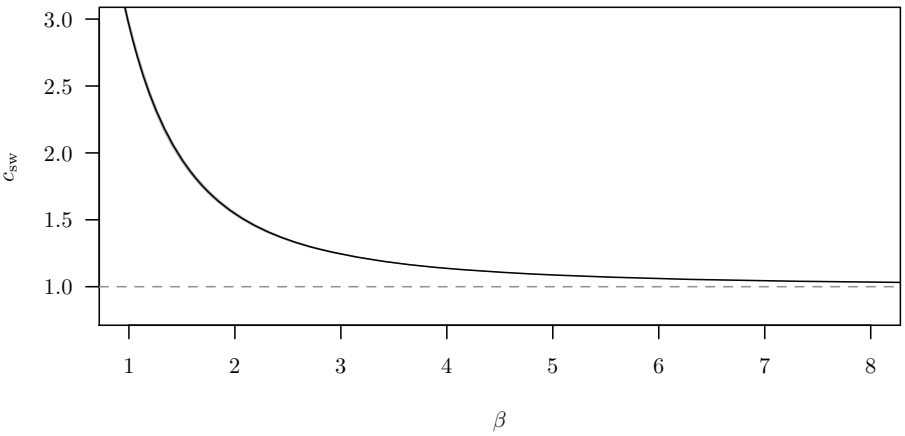

Figure 4.5: Illustration of the asymptotic behaviour of Equation (4.15) when extended far beyond the original fit range $c_{sw} \in [1.25, 2.25]$ and $\beta \in [1.625, 1.834]$.

As an example of a candidate simulation point, in Figures 4.4a and 4.4b, the green band indicates c_{sw} as a function of β such that the lattice spacing is approximately 0.075 fm. The intersection point

$$c_{sw} \sim 1.69\,, \quad \beta \sim 1.79\,, \tag{4.16}$$

thus gives a pair of parameter values for a simulation at this lattice spacing. The results of Section 4.1.2 and the present section can be taken as a sufficient demonstration of the fact that tuning the six parameters involved in $N_f = 2+1+1$ simulations with twisted mass clover quarks in indeed possible.

It would be very interesting to see how Equation (4.15) compares to what is found in other works with different gauge actions, by matching the results through some non-perturbative coupling. In Ref. [145] for $N_f = 4$, c_{sw} is given for Wilson fermions with the plaquette gauge action in a massless improvement scheme. In Ref. [146], it is argued that the presence of the charm quark requires working in a massive $N_f = 3+1$ improvement scheme along a line of constant physics. A numerical result is given for the tree-level improved Lüscher-Weisz gauge action at one value of β and it would be insightful to see how the eventual parametrisation will differ from the result in the massless scheme. Comparison to Equation (4.15) would also be interesting since in some sense, the effect of the charm has been included implicitly via the plaquette expectation value, although of course the meaning of Equation (4.1) is unclear in this sense.

FIRST PRODUCTION SIMULATIONS WITH FOUR FLAVOURS

First $N_f = 2+1+1$ production simulations with the parameters of Equation (4.10) have been started and the bare physical light quark mass has been found through the same kind of extrapolation explained in Chapter 3 for the $N_f = 2$ case. An overview of the tuning to maximal twist of these simulations is shown in Figure 4.6. The estimate of $\kappa_c \sim 0.140072$ of Equation (4.11) was refined with three $T/2a = L/a = 16$ simulations with $a\mu_\ell = 0.01$ to yield $\kappa_c \sim 0.140066$. Using this value, further simulations were conducted with $a\mu_\ell = 0.007, 0.005, 0.003$ for the purpose of extrapolating to the physical light quark mass. The estimate $a\mu_\ell = 0.0009$ is obtained from this extrapolation.

As can be seen, the slope of am_{PCAC} is steeper than for the $N_f = 2$ case, but less pronounced than in past $N_f = 2+1+1$ simulations without the

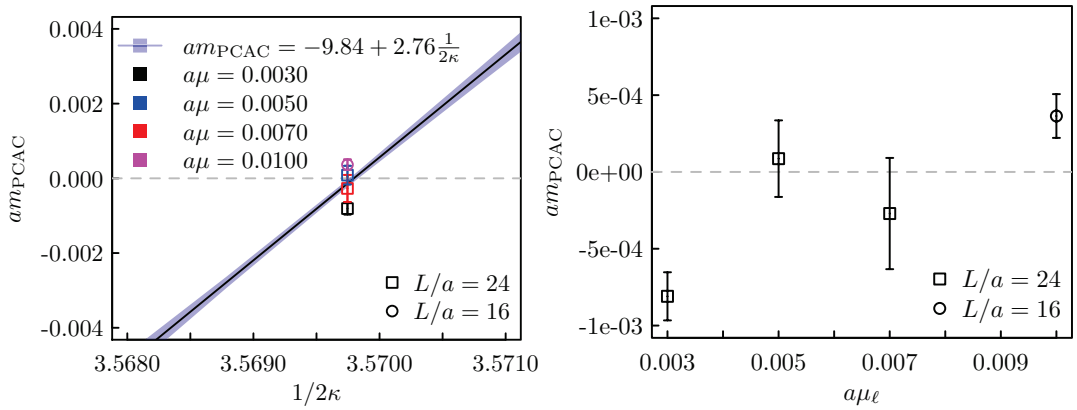

Figure 4.6: (left): Tuning to maximal twist in first $N_f = 2+1+1$ production simulations (the slope is due to points outside of the shown $1/2\kappa$ interval). (right): μ_ℓ dependence of am_{PCAC} as at $\kappa = 0.140066$.

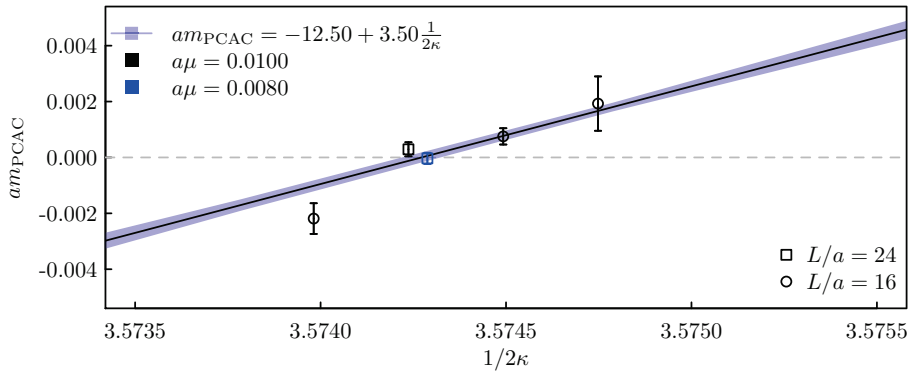

Figure 4.7: Tuning to maximal twist in first $N_f = 4$ production simulations.

clover term. A particular feature that deserves mention is that also the μ-dependent shift in m_{crit} is visible in the data, although the uncertainties are such that they largely overlap. If these shifts are truly present and not just finite size effects, the non-linearities are small, hinting at a small value of c_2, compared to simulations without the clover term. Simulations at the physical light quark mass are in the thermalisation stage at the time of writing of this thesis and a complete discussion is therefore not possible. However, the MD histories look encouraging and tuning to maximal twist is progressing well.

At the same time, the generation of $N_f = 4$ ensembles for RI-MOM renormalisation was begun. In contrast to the problematic situation without the clover term, tuning to maximal twist was possible without issues and has so far been completed for $a\mu_\ell = 0.01, 0.008$. A summary of the tuning procedure is given in Figure 4.7. The slope here is slightly higher than in the $N_f = 2+1+1$ case and a μ-dependence can also be seen, but it should be pointed out that for the purpose of the chiral limit for renormalisation, sea quark masses below around $a\mu_\ell = 0.005$ are certainly not required. Finally, it should be remembered that the lattice spacing is $a \sim 0.095$ fm or even a bit higher and that therefore lattice artefacts affecting the phase structure of the theory and the pion mass splitting, will diminish as the continuum limit is approached.

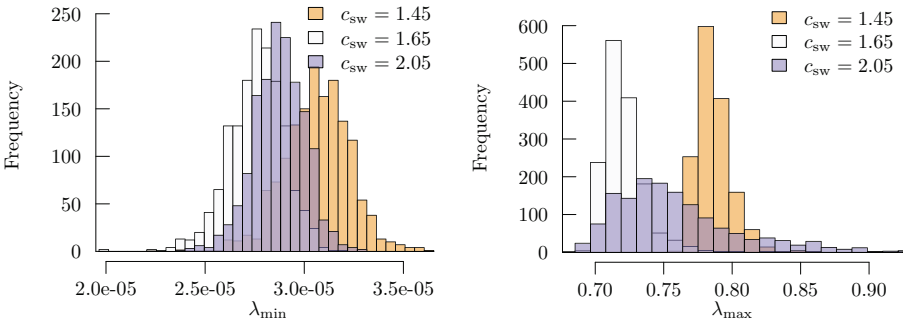

Figure 4.8: Normalised minimal and maximal eigenvalues of $Q_h^\dagger Q_h$ as a function of c_{sw} in $N_f = 2+1+1$ simulations at approximately constant physical situation.

	$c_{sw} = 1.45$	$c_{sw} = 1.65$	$c_{sw} = 2.05$
$[S_{min}, S_{max}]$	$[3.24 \cdot 10^{-5}, 3.6]$	$[4.0 \cdot 10^{-5}, 4.0]$	$[4.0 \cdot 10^{-5}, 4.8]$
$\langle \lambda_{min} \rangle$	$3.06(1) \cdot 10^{-5}$	$2.76(1) \cdot 10^{-5}$	$2.88(1) \cdot 10^{-5}$
$\langle \lambda_{max} \rangle$	0.7863(4)	0.7203(5)	0.765(2)
$\langle \delta H \rangle$	0.051(9)	0.081(10)	0.5(3)
$\langle \exp(-\delta H) \rangle$	1.001(9)	0.994(10)	0.99(2)
P_{acc}	0.873(9	0.830(10)	0.79(1)

Table 4.2: Approximation intervals for the rational approximation of the heavy doublet as well as *normalised* eigenvalues in $N_f = 2+1+1$ simulations with approximately constant physical situation. Also shown are expectation values of algorithmic observables. For more detail, see Section 6.2.

Effect of the Clover Term on Simulations

To close this section, the effect of the clover term on algorithmic aspects of the simulation will be discussed. Three of the four simulations for the computation of the pion mass splitting as a function of c_{sw} at approximately constant physical situation in Section 6.2 were also simulated with practically identical algorithmic parameters. These are ensembles *P145*, *P165* and *P205* with $c_{sw} = 1.45, 1.65, 2.05$ respectively. Clearly β, c_{sw} and κ were different between the three, but all mass preconditioning parameters were essentially identical. Further, the bounds of the eigenvalue interval for the rational approximation of the non-degenerate doublet differed significantly between the three simulations as the clover term has a substantial effect on minimal and especially maximal eigenvalues. It should be noted that the choices for the approximation intervals were not made systematically, making the eigenvalues somewhat difficult to compare. As a side remark it can also be mentioned that the value of c_{sw} increases the inversion cost, which rises with rising c_{sw}.

The Clover Term and The Positivity of the ND Quark Determinant

As discussed in Section 1.3.4, simulating physical strange and charm quarks using the mass non-degenerate twisted mass quark action requires $\mu_\delta > \mu_\sigma$ because $Z_P/Z_S < 1.0$. In this situation, the positivity of the quark determinant is not guaranteed. However, monitoring the minimum eigenvalue of $Q_h^\dagger Q_h$ in the course of a simulation has consistently shown that the strange quark mass is sufficiently large to prevent sign changes of the heavy quark determinant. For the three ensembles here, the kaon masses range from around 550 MeV at $c_{sw} = 2.05$ to around 490 MeV at $c_{sw} = 1.45$, likely owing to the fact that Z_P/Z_S depends substantially on c_{sw} and that the lattice spacings are not perfectly matched.

The unnormalised bounds of the approximation interval are given in Table 4.2, together with measurements of the *normalised* minimal and

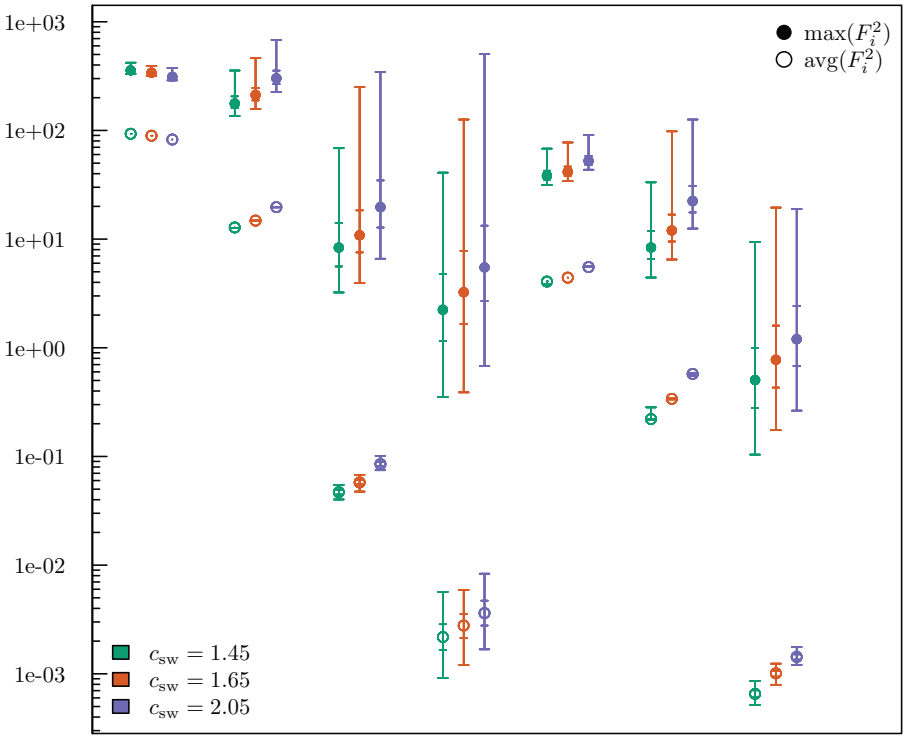

Figure 4.9: Box-and-whiskers type plot of the distributions of average and maximal forces in $N_f = 2+1+1$ simulations at roughly constant physical situation as a function of c_{sw}. The quantiles are $0.0, 0.1573, 0.5, 0.8427, 1.0$. From left to right, the groups of three are:
(1) gauge,
(2) light degenerate doublet determinant with largest preconditioning mass $2a\kappa\rho = 0.1$,
(3) determinant ratio with $2a\kappa\rho_i = \{0.01, 0.1\}$,
(4) determinant ratio with target quark mass in numerator,
(5) partial fractions $k = 0, 1, 2$ of non-degenerate strange-charm doublet,
(6) $k = 3, 4, 5$,
(7) $k = 6, 7, 8$.

maximal eigenvalues. For the rational approximation to be valid, these normalised eigenvalues should be well within the interval $[S_{min}/S_{max}, 1]$. Because the strange quark masses differ somewhat and the eigenvalue bounds were not selected systematically with this comparison in mind, the effect of the clover term on the eigenvalues $Q_h^\dagger Q_h$ is not visible in the expectation values. Looking at the histograms in Figure 4.8, however, shows that the maximal eigenvalues fluctuate very strongly for $c_{sw} = 2.05$ while they are quite stable for the two other values. It should be noted that in simulations without the clover term, the maximal eigenvalues were seen to be extremely stable with a much smaller variance than seen here. Because of the observed fluctuations with the clover term, the approximation intervals have to be chosen to be larger, making the simulation more expensive. In any case, however, the minimal eigenvalues are well separated from 0 and it appears that there is thus no problem with regards to the positivity of the heavy quark determinant.

The Clover Term and Molecular Dynamics Forces

Another consequence of an increasing c_{sw} is the well known increase of the molecular dynamics (MD) forces. The availability of the data presented here gives a unique view on how the distribution changes as a function of c_{sw}, although much like in the case of the eigenvalues, the mass preconditioning parameters are not exactly comparable. A summary of the force distribution for the different monomials in the MD Hamiltonian for the three ensembles is shown in Figure 4.9 by means of a quantile representation. The forces were measured at each trajectory and histograms of the average and maximal forces are shown in Appendix C for completeness. Definitions of the MD force as well as the meaning of *average* and *maximal* are given Chapter 5.

The differences in the forces for the gauge monomial can clearly be attributed to the reduction in β and are likely only consequential. Looking at

the fermionic contributions to the force, the average forces are simply shifted upwards without affecting their distribution substantially. For the maximal forces however, the central value of the distribution increases with c_{sw} and the distributions become significantly wider in the positive direction. This is also reflected in δH and the acceptance probability of these three simulations, as listed in Table 4.2.

For twisted mass fermions at maximal twist, which do not require c_{sw} to be tuned for the purpose of $\mathcal{O}(a)$-improvement, this may suggest an interesting possibility. If simulations are stable and the pion mass splitting is not affected by reducing c_{sw} below the value suggested by Equation (4.15), one could reduce the simulation cost by choosing smaller values of c_{sw}. As will be shown in Section 6.2 within the large uncertainties there, there does indeed not be a difference between the pion mass splitting at $c_{sw} = 1.45$ and $c_{sw} = 2.05$. The cost savings would be further accentuated by the reduction in the inversion cost and the stricter bounds that could be used for the eigenvalue interval of the rational approximation.

SUMMARY AND CONCLUSIONS

The features so far observed in $N_f = 2+1+1$ and $N_f = 4$ simulations with the twisted mass clover action suggest that it is suitable for simulating also at the physical light quark mass. There are certainly no signs of meta-stabilities in simulations at charged pion masses around 250 MeV and a lattice spacing around 0.095 fm. Tuning runs at the physical light quark mass are in the thermalisation stage and show encouraging molecular dynamics histories, although a detailed analysis is not yet possible.

In this chapter it was shown that a tuning strategy has been devised which allows all six parameters of the four flavour theory to be tuned in stages without prior information, starting from simulations with two mass degenerate light quarks and two further mass degenerate quarks in the region of the strange quark mass. Using this starting point, values for the μ_σ and μ_δ bare parameters were found which are expected to correspond very closely to the physical strange and charm quark masses.

Sets of $N_f = 2+1+1$ ensembles at different β and c_{sw} values were used for empirical fits to parametrise the lattice spacing and the critical hopping parameter in terms of β and c_{sw}. A similar parametrisation of the plaquette expectation value was used to give a functional form for the tadpole-improved value of c_{sw} as a function of β, thus significantly reducing the tuning burden of future simulations. It was firstly seen that – to the precision required for estimating a target lattice spacing – the exact mass of the strange and charm quarks does not affect the dependence of the lattice spacing on β and c_{sw}. Further, it was seen that even rather significant changes of around 10% in μ_δ and μ_σ do not induce a retuning of c_{sw}, to within the precision of Equation (4.1). As a result, the proposed functions can likely provide starting values of β and c_{sw} for the entire set of $N_f = 2+1+1$ twisted mass clover simulations. Since the fits ignored many sources of systematic error and a number of approximations were made, deviations are likely to appear. However, at least in the parameter region so far studied, the estimate of κ_c, for example, appears to be valid to within around 0.4 per-mille.

It was further shown that the RHMC algorithm is well suited for the simulation of the non-degenerate strange and charm quark doublet. The rational approximation behaves predictably and there are no problems with

the positivity of the quark determinant. From eigenvalue measurements it seems that the maximal eigenvalue of $Q_h^\dagger Q_h$ fluctuates quite strongly compared to simulations without the clover term and the effect increases with increasing c_{sw}. This makes simulations more expensive because the eigenvalue interval for the rational approximation needs to be enlarged. Similarly, it was shown that the molecular dynamics forces are affected by the value of c_{sw} and the possibility of reducing costs by using smaller values of c_{sw} in simulations was considered.

5 STUDY OF MASS PRECONDITIONING

Mass preconditioning and multiple time scale integration have been discussed in Section 1.4.3 and their parameter dependence will be studied in the present chapter. The average and maximal forces from the different terms in the effective action span many orders of magnitude and their distributions are generally not simple Gaussians. The preconditioning masses are continuous parameters, they have a very strong effect on the acceptance probability and simulations at or close to the physical light quark mass may have three or more of them. Because tuning a large number of parameters independently is costly and impractical, it would be highly beneficial if it were possible to reduce the parameter space by establishing simple relations amongst all of the masses based on simple measurements or a theoretical understanding of the parameter dependence. Together with knowledge of the number of integration steps on different time scales, it may then be possible to tune all of the parameters simultaneously by tuning their relationships. It will be shown that these kinds of relations can be found and that such a tuning is not only possible, but also works well in practice to reduce total simulation cost.

Despite its significance in allowing simulations at small light quark masses, the relationship between the observed forces and the preconditioning parameters remains poorly understood. Expansions around the free field limit may allow this relationship to be explored theoretically, but there is no guarantee that the non-perturbative corrections could be captured appropriately. To begin, therefore, a purely *empirical* approach will be proposed in which the (possibly) complex parameter dependence of the forces of the even-odd preconditioned twisted mass clover determinant and determinant ratios will be reduced to simple functional forms with only two free fit parameters. These functional forms are not theoretically motivated, but their leading dependence on the mass preconditioning parameters is quite reasonable. The dependence is studied entirely in lattice units because the definition of the forces in the continuum limit is unclear. This approach also captures the parameter dependence in the way it is encountered in actual simulations as pure numbers depending on other pure numbers.

The formulae are fitted to measurements on a test ensemble with forces and preconditioning parameters spanning many orders of magnitude and are shown to describe the data surprisingly well. In addition, it is demonstrated that for similar physical situations (target quark mass, lattice spacing, volume), they provide qualitatively correct predictions of the force measurements for different simulations. When the physical situation is very different, the fit parameters determined on the test ensemble fail to reflect the force hierarchy even qualitatively. However, the hope is that in a given simulation, these effective functional forms can be fitted to just a few data points, fixing the fit parameters and allowing the mass preconditioning scheme to be suitably optimised.

In the last part of this chapter, the functional forms are assumed to be meaningful and the consequences of the leading dependence as well as their full form are used to derive several simplified relations between mass preconditioning parameters and force hierarchies. In particular, a mass precon-

ditioning and integration scheme is explored which is shown to produce good acceptance at significantly reduced simulation cost compared to current production simulations of twisted mass clover ensembles. The idea is then extended to suggest a mass preconditioning and integration scheme for simulations at the physical pion mass and the resulting simulation cost is estimated. A very tentative discussion of the frequency spectra of the force oscillations based on fast Fourier transforms of long molecular dynamics trajectories is given in Chapter 11.

FUNCTIONAL FORMS FOR THE FORCES

In this section, the magnitude of the gauge derivative (or force) coming from a given term in the effective action will be studied and an attempt will be made to describe the dependence on the preconditioning parameters more explicitly. In order to establish an appropriate norm for the force, the notation of Ref. [79] will be used for the square norm of an element X in the Lie algebra of SU(3)

$$\|X\|^2 = -2\,\mathrm{Tr}\,X^2\,,$$

but unlike there, the square root will *not* be taken. Two measures of the force $\|F\|^2$ due to a term in the effective action will be considered. Firstly, the average

$$\|F\|^2_{\mathrm{av}} = \frac{1}{4V}\sum_{x,\nu}\|F\|^2(x,\nu) \tag{5.1}$$

and secondly the maximal force

$$\|F\|^2_{\mathrm{max}} = \max_{x,\nu}\|F\|^2(x,\nu)\,, \tag{5.2}$$

where F corresponds to δP^a_μ of Equation (1.97) and where V is the four-dimensional lattice volume, x is a site index and ν is a lattice direction, running from 1 to 4.

For the purpose of this study, a small $N_f = 2$ twisted mass clover ensemble was generated with parameters

$$\beta = 2.1\,,\ c_{\mathrm{sw}} = 1.57551\,,\ \kappa = 0.1373\,,\ a\mu_\ell = 0.01\,,\ T/2a = L/a = 16\,,$$

corresponding to a pion mass of around 450 MeV, $M_\pi L \sim 3.3$ and a lattice spacing $a \sim 0.09$ fm.

Subsequently, around 100 measurements of the average and maximal $\|F\|^2$ due to even-odd preconditioned determinants and determinant ratios with a wide range of mass parameters were carried out on these configurations. To do this, appropriate pseudofermion fields (see Section 1.4) for the different terms were generated for each gauge configuration and the analytical expressions for the forces were evaluated using these fields on the given gauge background, without actually performing further molecular dynamics integration. It should be noted that the used configurations are not very well separated and that autocorrelation effects have not been taken into account. However, this approximates the situation in actual simulations rather well, where the forces are of course computed on strongly correlated configurations. In order make direct contact with simulations, the notation

$$\tilde{\mu}_i = 2a\kappa\mu_i$$

will be adopted throughout.

Forces of the Quark Determinant

Hybrid Monte Carlo simulations of QCD with a pair of mass degenerate light quarks generally rely on the pseudofermion representation of the light quark determinant. For Wilson twisted mass lattice QCD at maximal twist, this takes the form

$$\det\left(Q^2+\mu_\ell^2\right) \propto \int \mathcal{D}\phi^\dagger \mathcal{D}\phi \exp\left(-\phi^\dagger \frac{1}{Q^2+\mu_\ell^2}\phi\right), \tag{5.3}$$

where $Q = \gamma^5 M$ is the Hermitian Wilson Dirac operator with the Wilson quark mass tuned to its critical value up to $\mathcal{O}(a)$ shifts. Q may also contain a clover term and μ_ℓ is the (light) twisted quark mass. In practice, even-odd preconditioning is used and the determinant splits into two parts as described in detail in Appendix A. For the purpose of mass preconditioning, the parameter $\tilde{\rho} = 2a\kappa\rho$ is introduced. Unlike for standard Wilson clover fermions, the diagonal components of the Dirac operator, M_{ee} and M_{oo}, now depend on the twisted quark mass. In order to avoid inverting $M_{ee} = (1 + T_{ee} \pm i\gamma_5\tilde{\mu})$ too frequently, the Hermitian even-odd preconditioned twisted mass clover operator is implemented as follows

$$\begin{aligned}\hat{W}^{\pm}(\tilde{\mu}_\ell,\tilde{\rho}) =& \gamma_5\left[1+T_{oo}\pm i\gamma_5(\tilde{\mu}_\ell+\tilde{\rho}) - M_{oe}(1+T_{ee}\pm i\gamma_5\tilde{\mu}_\ell)^{-1}M_{eo}\right]\\ =& \hat{Q}^{\pm}(\tilde{\mu}_\ell) \pm i\tilde{\rho}\,,\end{aligned} \tag{5.4}$$

such that M_{ee} and M_{oo} depend only on the target mass $\tilde{\mu}_\ell$. A more symmetric implementation is possible, but this has not yet been done in the tmLQCD software suite. The determinant term now takes the form

$$\|\det(1+T_{ee}+i\tilde{\mu}_\ell\gamma_5)\|^2 \int \mathcal{D}\phi_o^\dagger \mathcal{D}\phi_o \exp\left(-\phi_o^\dagger \frac{1}{\hat{W}^+\hat{W}^-}\phi_o\right), \tag{5.5}$$

which is real and positive as required and where the ϕ_o are defined only on the odd sites. In order to parametrise the results of this section, the notation

$$\tilde{\mu} = \tilde{\mu}_\ell + \tilde{\rho} \tag{5.6}$$

will be adopted.

In the molecular dynamics integration of the equations of motion, the derivative with respect to the gauge field of the determinant term needs to be computed. As has been demonstrated for standard Wilson fermions in Ref. [147], for example, the magnitude of this derivative grows substantially as the quark mass is lowered. This is complicated further by increased fluctuations in the force magnitude, with a particular increase in the occurrence of very large maximal forces as given by Equation (5.2). In consequence, simulations can become unstable (at constant trajectory length), unless the number of integration steps is increased substantially, making them computationally expensive.

It has been shown in Refs. [77, 78, 79] that mass preconditioning can be used to lower the cost of simulations with light dynamical quarks through the introduction of determinant ratios which distribute the forces over multiple terms in the effective action. The mass parameter in the remaining determinant term can thus be chosen to be sufficiently large to reduce the forces to an acceptable level and most importantly, dampen the strong fluctuations.

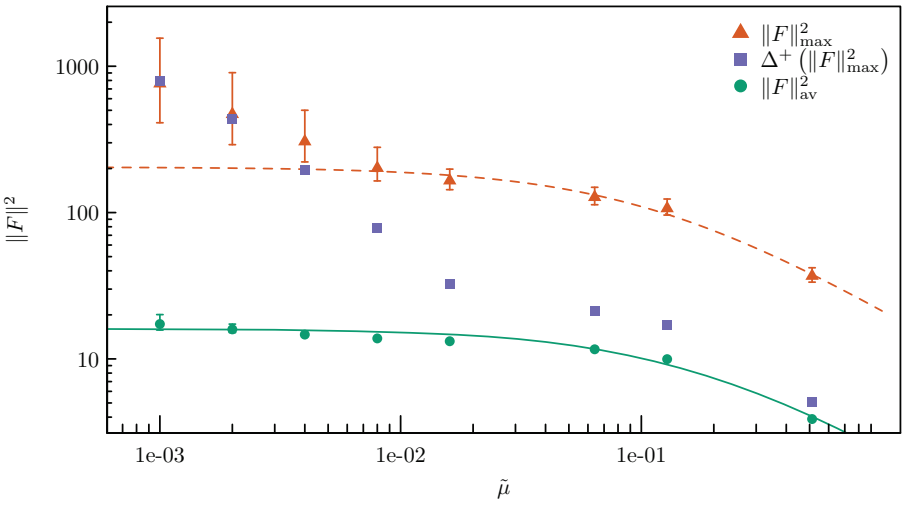

Figure 5.1: Square norm of the forces from an even-odd preconditioned determinant for two flavours of mass degenerate twisted mass quarks at maximal twist as a function of the mass parameter $\tilde{\mu} = 2a\kappa\mu$. The notation Δ^+ indicates the difference between the median and 84.27 percentile of the maximal square norm force. The line corresponds to Equation (5.7) while the dashed line is Equation (5.8).

The average and maximal square norm forces were measured on the test ensemble introduced above for a wide range of mass parameters $\tilde{\mu}$ (setting $\tilde{\rho} = 0.0$) to determine the expected forces with just a determinant term in the effective action for different $\tilde{\mu}_\ell$. The results of these measurements are shown in Figure 5.1. The lines in the figure represent fits of the form

$$\|F\|^2_{\text{av}}(\tilde{\mu}) = \frac{2.74(17)}{\tilde{\mu} + 0.17(2)} \tag{5.7}$$

$$\|F\|^2_{\max}(\tilde{\mu}) = \frac{23(2)}{\tilde{\mu} + 0.11(2)}\,, \tag{5.8}$$

which was guessed based on experience and confirmed by trial and error to best describe the data. It is clear that this is only an effective empirical description. In particular, the constant term in the denominator helps to align the model with the data for larger values of $\tilde{\mu}$. The inverse dependence on the quark mass was already described in Ref. [148] and similar empirical fits for Wilson fermions were attempted in Ref. [147]. There, the quark masses were much higher, the dependence on the quark mass was milder and the validity of the fit functions seemed to be limited to a much smaller range of quark masses than here.

The average force is described rather well by the fitted function, but the maximal force deviates strongly from the proposed form for $\tilde{\mu} < 0.01$. The notation Δ^+ in the figure indicates the difference between the median and the 84.27 percentile of the maximal square norm force and represents a very good measure of the magnitude and frequency of exceptionally strong fluctuations. This grows substantially below this quark mass, even reaching the same order of magnitude as the median maximal force. It is likely that this is the reason why simulations with very light quarks are not feasible with just a quark determinant. These results suggest that to ensure stable simulations, mass preconditioning should be used to set the mass parameter of the determinant $\tilde{\mu} \gtrsim 0.2$ as this reduces the average and maximal forces and significantly dampens the fluctuations in the maximal force.

In Appendix D.1, the fits of this section are repeated for determinants where $\tilde{\rho}$ is varied for $\tilde{\mu}_\ell = 0.0, \tilde{\mu}_\ell = 0.0002746$ and $\tilde{\mu}_\ell = 0.0010984$ kept constant, corresponding to the situation encountered when mass preconditioning is employed. The results suggest that the behaviour is exactly the same as a function of $\tilde{\mu}$, except that the factor $(1 + T_{ee} \pm i\gamma_5\tilde{\mu}_\ell)^{-1}$ amplifies the forces as $\tilde{\mu}_\ell \to 0$. The values of the parameters of the fits

in the appendix are consistent amongst each other, suggesting that the dominant dependence is truly on $\tilde{\mu} = \tilde{\mu}_\ell + \tilde{\rho}$. In addition, although the parameter values are different from Equations (5.7) and (5.8), they give the same constant term when $\tilde{\mu} = 0.0$. Of course, the model is not expected to be valid there.

In conclusion, it seems that as long as the mass parameter $\tilde{\mu} = \tilde{\mu}_\ell + \tilde{\rho}$ is sufficiently large, both the average and the maximal square norm force can be described by a simple function of $\tilde{\mu}$. When mass preconditioning is not used ($\tilde{\rho} = 0.0$) and only a determinant is present, this means that the simulated quark mass must indeed be quite large for the fluctuations to be appropriately dampened. When it is used, on the other hand, setting $\tilde{\rho} \gtrsim 0.2$ significantly reduces the average and maximal forces and, most importantly, the fluctuations in the maximal force. Although this gives a target value for the largest mass preconditioning parameter, a similar model for the forces coming from determinant ratios is required to inspire a complete mass preconditioning scheme. This will be the goal of the following sections.

Forces of Quark Determinant Ratios

Mass preconditioning results in a hierarchy of forces because, for a determinant ratio R of the form

$$R = \frac{\det\left(Q^2 + (\mu_\ell + \rho_1)^2\right)}{\det\left(Q^2 + (\mu_\ell + \rho_2)^2\right)} = \frac{\det\left(Q^2 + \mu_1^2\right)}{\det\left(Q^2 + \mu_2^2\right)},$$

it is clear that $R \to 1$ as $\mu_1 \to \mu_2$. Because a constant cannot depend on the gauge field, $\|F\|^2 \to 0$ in the same limit. In addition, there should be some level of symmetry whether μ_2 approaches μ_1 from above or from below. When even-odd preconditioning is used, the pseudofermion part of the determinant ratio is

$$S_{PF_R} = \phi_o^\dagger \hat{W}^-(\tilde{\mu}_\ell, \tilde{\rho}_2) \frac{1}{\hat{W}^+(\tilde{\mu}_\ell, \tilde{\rho}_1)\hat{W}^-(\tilde{\mu}_\ell, \tilde{\rho}_1)} \hat{W}^+(\tilde{\mu}_\ell, \tilde{\rho}_2)\, \phi_o\,, \tag{5.9}$$

where $\hat{W}^\pm$ are as defined in Equation (5.4). This suggests a more complicated behaviour for the forces, but the requirement for them to vanish in the limit $\mu_2 \to \mu_1$ remains the same. The combination of twisted mass, clover term and even-odd preconditioning is described in detail in Appendix A. It should be noted that due to the form of Equation (5.4), the terms corresponding to the first factor in Equation (5.5) cancel for determinant ratios.

A possible functional dependence on $\tilde{\mu}_2$ and $\tilde{\mu}_1$ can be motivated as follows:

- There cannot be a constant term because of the aforementioned behaviour of R in the limit $\tilde{\mu}_2 \to \tilde{\mu}_1$.
- The function must be positive because $\|F\|^2$ is positive, suggesting a dependence on an absolute value or a term with an even exponent.
- There should be an approximate symmetry under interchange of $\tilde{\mu}_1$ and $\tilde{\mu}_2$ and the positivity requirement suggests $f \propto |\tilde{\mu}_2 - \tilde{\mu}_1|^b$.
- The determinants depend quadratically on $\tilde{\mu}_1$ and $\tilde{\mu}_2$ respectively, this hints at $f \propto (\tilde{\mu}_2 - \tilde{\mu}_1)^2$.

- There is likely some further dependence on the relative sizes of the various mass parameters.

In fact, it will be shown below that the average force $\|F\|^2_{av}$ as well as the maximal force $\|F\|^2_{max}$ can be modelled by

$$\|F\|^2_{av} \sim a_{av} \, (\tilde{\mu}_2 - \tilde{\mu}_1)^2 \left(\frac{\tilde{\mu}_2}{\tilde{\mu}_1} \right)^{b_{av}} \tag{5.10}$$

$$\|F\|^2_{max} \sim a_{max} \frac{(\mu_2 - \mu_1)^2}{(\mu_1 \mu_2)^{b_{max}}} \, , \tag{5.11}$$

with a_{av}, a_{max}, b_{av} and b_{max} distinct real parameters. The empirical terms which multiply $(\mu_2 - \mu_1)^2$ were guided by the behaviour of the data in the regions $\tilde{\mu}_1 \gg \tilde{\mu}_2$ and $\tilde{\mu}_2 \gg \tilde{\mu}_1$. In addition, amongst all the possibilities that were tried, this particular form was the only one to follow the data in the ratio $\|F\|^2_{max}/\|F\|^2_{av}$. This is an important consideration: when only the leading behaviour is considered or the average and maximal forces are fitted with functions of the same form, this ratio would reduce to a constant in clear disagreement with the data. When polynomial fits or Padé approximants are attempted to take a more model-independent approach, the ratio of the fits for $\|F\|^2_{max}$ and $\|F\|^2_{av}$ produces singularities and further does not follow the data at all, despite $\|F\|^2_{av}$ and $\|F\|^2_{max}$ being well described individually. In addition, the number of fit parameters required for polynomial or Padé fits makes them unstable and certainly not amenable to being fitted with only a few data points, at odds with the aims of this chapter. Finally, the simple form of the fitted functions also interesting consequences to be derived, as will be seen below.

Of course, it should be remembered that the parameters are effective and certainly depend on the volume and the dynamical quark content. In particular, it was hoped that the data could be described with two or at most three free parameters, in a way which absorbs the dependence on $\tilde{\mu}_\ell$ and is well-behaved under scale transformations of the mass parameters, thereby allowing the functions to be fitted with just a few data points in the tuning stage of a given simulation.

A dependence on $|\tilde{\mu}_2^2 - \tilde{\mu}_1^2|$ did not describe the data well at all, suggesting that the cross terms are indeed essential. As indicated above, the dependence on $(\mu_2 - \mu_1)^2$ seems reasonable, but presently there seems to be no explanation for the value and form of the other factors. The presence of these terms can probably be understood by considering the bounds of the eigenvalue spectrum of the respective operators since these are bounded from below by the preconditioning masses. It may be possible to derive better motivated forms for the forces by expanding the derivative for the even-odd preconditioned determinant ratio in the free field limit.

Fitting the Functional Forms

In order to fit the proposed functional forms over a wide range of $\tilde{\mu}_1$ and $\tilde{\mu}_2$, $\tilde{\rho}_1$ and $\tilde{\rho}_2$ were varied while keeping $\tilde{\mu}_\ell = 0.0002476$, which corresponds approximately physically light quarks for this lattice spacing and κ value. The median as a central value as well as the 15.73 and 84.27 percentiles as estimates of the spread were determined and fits of the two functional forms in Equations (5.10) and (5.11) to the medians were carried out. In these fits, correlations between the data were ignored but they were weighted by the inverse of an effective variance determined from the percentiles. The data

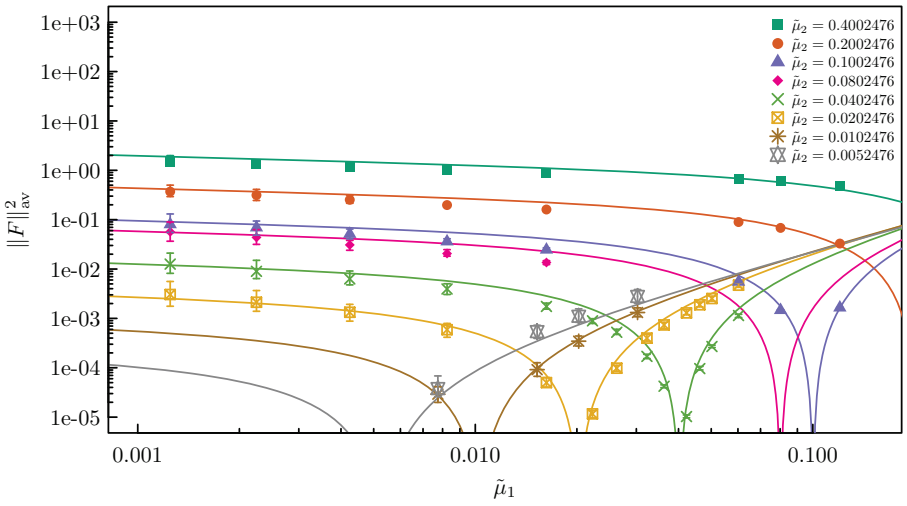

Figure 5.2: Average forces coming from a determinant ratio as a function of the mass preconditioning parameters $\tilde{\mu}_1$ and $\tilde{\mu}_2$ on gauge configurations with the $N_f = 2$ twisted mass clover action. The lines correspond to the fit of the form Equation (5.12).

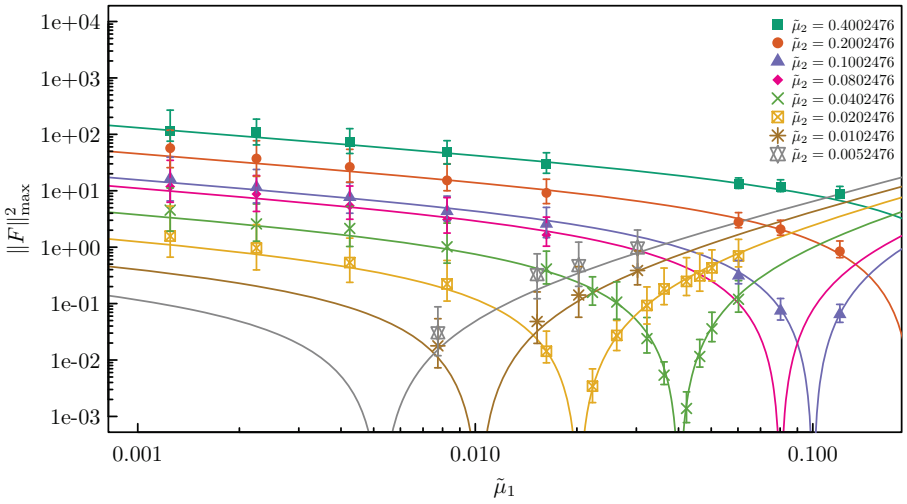

Figure 5.3: Maximal forces coming from an even-odd preconditioned determinant ratio as a function of the mass preconditioning parameters $\tilde{\mu}_1$ and $\tilde{\mu}_2$ on gauge configurations with the $N_f = 2$ twisted mass clover action. The lines correspond to the fit in Equation (5.13).

as well as the fit result are shown in Figures 5.2 and 5.3 and the parameters take the values

$$\|F\|^2_{\mathrm{av}}(\tilde{\mu}_1, \tilde{\mu}_2) = 4.2(2)\,(\tilde{\mu}_2 - \tilde{\mu}_1)^2 \left|\frac{\tilde{\mu}_2}{\tilde{\mu}_1}\right|^{0.18(3)} \tag{5.12}$$

$$\|F\|^2_{\mathrm{max}}(\tilde{\mu}_1, \tilde{\mu}_2) = 21(1)\frac{(\tilde{\mu}_2 - \tilde{\mu}_1)^2}{|\tilde{\mu}_1\tilde{\mu}_2|^{0.46(1)}}\,. \tag{5.13}$$

Both fits reproduce the qualitative aspects of the data rather well, although the fit for $\|F\|^2_{\mathrm{av}}$ has $\chi^2/\mathrm{df} \gg 1$ due to deviations in the region $\mu_1 \sim \mu_2$. In Appendix D.2, the figures and fits are also shown with $\tilde{\mu}_2$ on the abscissa for completeness. Additionally, the analysis was repeated with a tenfold larger target mass $\tilde{\mu}_\ell$ to ascertain the effect of the way even-odd preconditioning is implemented for the twisted mass clover action. The results in Appendix D.2 suggest that the proposed functional forms remain useful and that the residual dependence is indeed absorbed into the fit parameters (which now of course take different values).

Maximal Forces and their Distribution

An issue of particular importance in the tuning of the mass preconditioning parameters is the difference between the average and the maximal forces

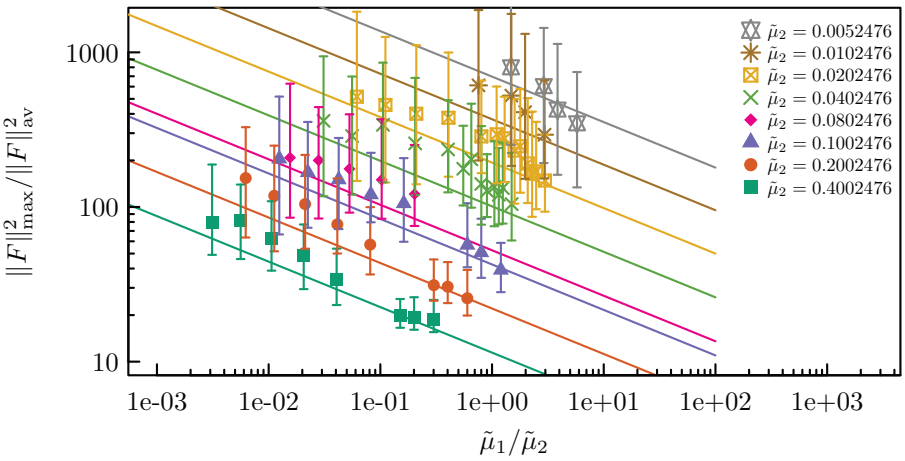

Figure 5.4: Ratio of the maximal and average forces from even-odd preconditioned determinant ratios with different parameters. The lines give the ratio of the model in Equations (5.10) and (5.11) for constant $\tilde{\mu}_2$ as a function of $\tilde{\mu}_1/\tilde{\mu}_2$, as indicated by the colours.

coming from the different terms in the effective action. In particular, it is known that large spikes in the maximal forces can trigger integrator instabilities, see Ref. [147] for example. As a consequence, it would be of value if it were possible to model not only the median of the maximal force as done in the previous section but also its variance, especially in the positive direction. Figure 5.4 shows the ratio of the maximal force and the average force for different parameter combinations as measured on the test ensemble. In addition, the lines depict the prediction from the ratio of Equations (5.12) and (5.13) overlaid over the data points (i.e. the lines are not lines of best fit).

It seems that within the large errors, $\|F\|^2_{max}/\|F\|^2_{av}$ is linearly proportional to some power of $\tilde{\mu}_1/\tilde{\mu}_2$ (or equivalently $\tilde{\mu}_1$ at constant $\tilde{\mu}_2$ as depicted in the figure). Some of the points show slightly different behaviour around $\tilde{\mu}_1 \sim \tilde{\mu}_2$, but it should be remembered that the number of measurements is quite low and it is likely that the distribution of maximal forces has not been sampled sufficiently well. Further, if the functional forms can truly be trusted, it is clear that the dependence should be

$$\frac{\|F\|^2_{max}}{\|F\|^2_{av}} \sim \frac{a_{max}}{a_{av}} \left| \mu_1^{(b_{max}-b_{av})} \mu_2^{-(b_{max}+b_{av})} \right| . \tag{5.14}$$

Despite the mild fluctuations in the data, the simple dependence indicates that for any given $\tilde{\mu}_1$, a $\tilde{\mu}_2$ can be chosen predictably to reduce the difference between the average and the maximal force. The indication that the ratio between maximal and average force is reduced for $\tilde{\mu}_2 < \tilde{\mu}_1$ is of little practical value because such a term would in turn require a compensating term for which the ratio would be larger. Finally, for small values of $\tilde{\mu}_2$, the ratio between the average and the maximal forces becomes very large, although it should be kept in mind that in those situations, $\|F\|^2_{av}$ is usually small because $\tilde{\mu}_2 \sim \tilde{\mu}_1$.

As mentioned above, the most severe complication in the integration of the equations of motion is the occurrence of large spikes in the maximal forces. As discussed further below, these spikes are seen especially in the simulation of very light fermions. Figure 5.5 shows the difference between the 84.27 percentiles and the median of the maximal forces measured on the test ensemble (referred to as $\Delta^+(\|F\|^2_{max})$) as a function of the mass preconditioning parameters. At present, no errors have yet been computed

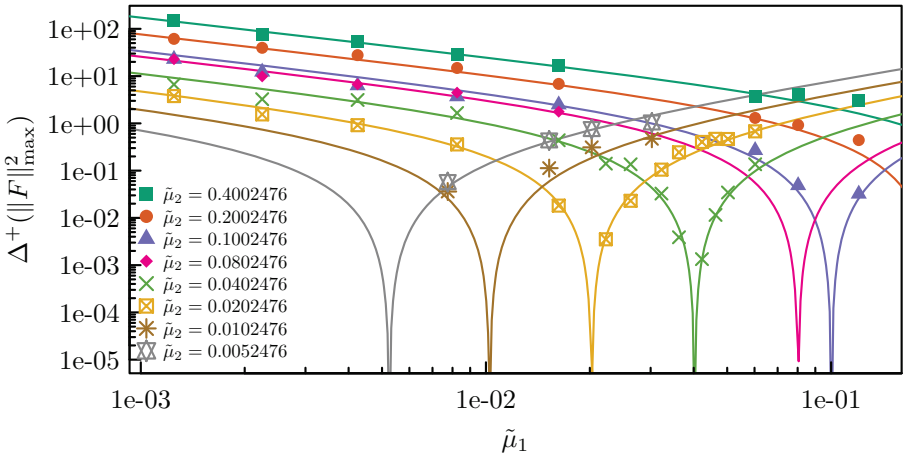

Figure 5.5: Difference between 84.27 percentiles and the median of the maximal forces of even-odd preconditioned determinant ratios with different parameter combinations, the lines correspond to the functional form in Equation (5.15)

on the percentiles and as such the data has to be taken at face value. The lines in the figure correspond to a fit of the form:

$$\Delta^+\left(\|F\|^2_{\max}(\tilde{\mu}_1,\tilde{\mu}_2)\right) = 1.69(23)\frac{(\tilde{\mu}_2-\tilde{\mu}_1)^2}{|\tilde{\mu}_1\tilde{\mu}_2|^{0.83(2)}}\,, \tag{5.15}$$

where the errors on the parameters are estimates given by fits which remove sets of points.

The last quantity to be investigated was the relative size of the positive variance, $\Delta^+(\|F\|^2_{\max})/\|F\|^2_{\max}$, which is a quantity that one would like to minimize, perhaps above all others. From Figure 5.5 and the result of the fit (Equation (5.15)), it would seem that $\tilde{\mu}_1$ and $\tilde{\mu}_2$ can be chosen such that the variance is minimised. While this is true, it is of limited usefulness because the points where it is minimised correspond to the points where the forces themselves vanish. In particular, the functional forms can be combined as above to obtain

$$\frac{\Delta^+(\|F\|^2_{\max})}{\|F\|^2_{\max}} \sim \frac{1.69(23)}{21(2)}\,|\tilde{\mu}_1\tilde{\mu}_2|^{[-0.37(3)]}\,. \tag{5.16}$$

Still, at the very least the model above can be used to make qualitative predictions about the size and/or frequency of large spikes in the maximal forces, potentially aiding in the tuning effort.

FORCE HIERARCHY RELATIONS

The empirical model with the determined fit parameters presented above is of course not generally applicable, because the phenomenological terms $|\tilde{\mu}_1/\tilde{\mu}_2|^b$ and $|\tilde{\mu}_1\tilde{\mu}_2|^{b'}$ and the fit parameters carry implicit dependences. However, because the leading dependence seems to be $(\tilde{\mu}_2-\tilde{\mu}_1)^2$, it is very likely that an approximate force hierarchy, applicable anywhere, can be established by ignoring these terms. It is further interesting to explore the consequences of the functional forms and the kinds of relations that can be derived based on these.

Mass Preconditioning Recurrence Relations

The following ideas are based on arguments from Ref. [138], which suggests using the RHMC to integrate degenerate flavours on multiple time scales by

splitting the fermion determinant into equivalent factors of the n-th root. The factors are then implemented using the partial fraction representation of the Zolotarev approximation, the individual terms of which are very much like determinant ratios. It is argued that this splitting is optimal, but simulations of light dynamical quarks using the RHMC can be very expensive. In the case of mass preconditioning, a similar argument can be made which may allow the force contributions to be arranged predictably and then distributed over multiple time scales. One particularly interesting possibility is to choose the parameters such that the forces from two or more determinant ratios are of the same order and can thus be integrated on the same time scale.

Given a pair of determinant ratios

$$R_2 R_1 = \frac{\det\left(Q^2 + \tilde{\mu}_2^2\right)}{\det\left(Q^2 + \tilde{\mu}_3^2\right)} \cdot \frac{\det\left(Q^2 + \tilde{\mu}_1^2\right)}{\det\left(Q^2 + \tilde{\mu}_2^2\right)}$$

and the target mass $\tilde{\mu}_1 = \tilde{\mu}_\ell$ (hence $\tilde{\rho}_1 = 0.0$), the ratio of their squared forces is, to leading approximation,

$$\frac{\|F\|^2(\tilde{\mu}_3, \tilde{\mu}_2)}{\|F\|^2(\tilde{\mu}_2, \tilde{\mu}_1)} = \frac{(\tilde{\mu}_3 - \tilde{\mu}_2)^2}{(\tilde{\mu}_2 - \tilde{\mu}_1)^2} .$$

The two forces of two determinant ratios can thus be made approximately equal by requiring

$$\frac{\tilde{\mu}_3 - \tilde{\mu}_2}{\tilde{\mu}_2 - \tilde{\mu}_1} = 1 ,$$

where the square root has been taken and the positive branch is chosen. This leads to

$$\tilde{\mu}_3 = 2\tilde{\mu}_2 - \tilde{\mu}_1 ,$$

while the negative branch would have given $\tilde{\mu}_3 = \tilde{\mu}_1$ and thus $R_1 R_2 = 1$, which is not a useful result.

This can be extended to more determinant ratios by considering the ratios of the differences of their preconditioning masses

$$\frac{\tilde{\mu}_4 - \tilde{\mu}_3}{\tilde{\mu}_3 - \tilde{\mu}_2} = 1 ,$$

giving immediately

$$\tilde{\mu}_4 = 3\tilde{\mu}_2 - 2\tilde{\mu}_1 ,$$

until finally a recursion relation can be obtained

$$\tilde{\mu}_n = (n-1)\tilde{\mu}_2 - (n-2)\tilde{\mu}_1 . \tag{5.17}$$

In addition to the target mass $\tilde{\mu}_1 = \tilde{\mu}_\ell$, the whole ladder of determinant ratios now depends only on $\tilde{\mu}_2$, which can be used to tune the force contribution of all these terms to some acceptable level. Since the inversion cost can be presumed known, this allows the simulation cost to be reliably predicted. The pattern can be broken at some given $\tilde{\mu}_m$ and the force of the resulting ratio can then be integrated on a finer time scale. If required, multiple levels of equal forces can be established following the recursion relation above with $\tilde{\mu}_m$ as the new target mass.

As an alternative to the proposal above, the force ratios can also be tuned by requiring the ratios of the squared mass differences to be a certain value different from 1, chosen such that the force hierarchy matches the hierarchy

of integration steps on the different time scales. Assuming that the squared forces should increase by a factor of f^2 from the coarsest time scale onwards,

$$\frac{(\tilde{\mu}_3 - \tilde{\mu}_2)}{(\tilde{\mu}_2 - \tilde{\mu}_1)} = f\,,$$

leads to the recurrence relation

$$\tilde{\mu}_n = \frac{f^n(\tilde{\mu}_2 - \tilde{\mu}_1) + f(f\tilde{\mu}_1 - \tilde{\mu}_2)}{f(f-1)}\,, \tag{5.18}$$

which, despite appearances, goes to Equation (5.17) in the limit $f \to 1$.

The final preconditioning mass $\tilde{\mu}_N$, which will be in the determinant term, can be chosen by following the results of Section 5.1.1. It should of course be kept in mind that the procedure presented here ignores the contribution $\left|\tilde{\mu}_n/\tilde{\mu}_{(n-1)}\right|$ and the resulting average forces will thus not scale exactly by the factor f. In addition, the maximal forces depend more strongly on $|\tilde{\mu}_n\tilde{\mu}_{(n-1)}|$ and will thus certainly not scale as expected, limiting the number of contributions that can be integrated on one time scale in the case $f = 1$. In principle, however, these effects could also be taken into account in the construction of the ladder of preconditioning masses. Of course, in that case knowledge of the phenomenological parameters b_{av} and b_{max} is required, defeating the purpose of this simplification.

Simplified Full Formulae for the Forces

It has been shown above how the choice of the preconditioning masses can be simplified significantly by ignoring the phenomenological terms. The recurrence relation in Equation (5.18) gives a rule for determining the mass hierarchy with a specific ladder of forces in mind. A further simplification of all the functional forms presented in the previous sections occurs when all of the preconditioning masses are specified in terms of $\tilde{\mu}_1$. In general for $\tilde{\mu}_n = x_n\tilde{\mu}_1$, Equations (5.10), (5.11) and (5.15) become

$$\begin{aligned}
\|F\|^2_{av}(x_n, x_{n-1}) &\sim a_{av}\left|\frac{x_n}{x_{n-1}}\right|^{b_{av}}(x_n - x_{n-1})^2\,\tilde{\mu}_1^2\\
\|F\|^2_{max}(x_n, x_{n-1}) &\sim a_{max}\,|x_n x_{n-1}|^{-b_{max}}(x_n - x_{n-1})^2\,\tilde{\mu}_1^{2(1-b_{max})}\\
\Delta^+\left(\|F\|^2_{max}(x_n, x_{n-1})\right) &\sim a_{\Delta^+}\,|x_n x_{n-1}|^{-b_{\Delta^+}}(x_n - x_{n-1})^2\,\tilde{\mu}_1^{2(1-b_{\Delta^+})}\,.
\end{aligned} \tag{5.19}$$

A final simplification can be obtained for $\tilde{\mu}_n = x^{n-1}\tilde{\mu}_1$, which leads to

$$\begin{aligned}
\|F\|^2_{av}(n) &\sim a_{av}\,|x|^{b_{av}}\,x^{2n-4}\,(x-1)^2\,\tilde{\mu}_1^2\\
\|F\|^2_{max}(n) &\sim a_{max}\frac{x^{2n-4}}{\left|x^{2n-3}\right|^{b_{max}}}(x-1)^2\,\tilde{\mu}_1^{2(1-b_{max})}\\
\Delta^+\left(\|F\|^2_{max}(n)\right) &\sim a_{\Delta^+}\frac{x^{2n-4}}{\left|x^{2n-3}\right|^{-b_{\Delta^+}}}(x-1)^2\,\tilde{\mu}_1^{2(1-b_{\Delta^+})}\,.
\end{aligned} \tag{5.20}$$

In this notation, $\|F\|^2(n)$ corresponds to to the determinant ratio with $\tilde{\mu}_n$ in the denominator, so the ratio with the target mass in the numerator would be $n = 2$. The average forces coming from the different determinant ratios now follow a very simple pattern because the phenomenological terms cancel in their ratios

$$\sqrt{\frac{\|F\|^2_{av}(n)}{\|F\|^2_{av}(n-1)}} = x \tag{5.21}$$

and similar simplified relationships hold for all of Equation (5.20), softened by b_{max} and b_{Δ^+}. This also explains the unwritten rule of placing the preconditioning masses at equal intervals on a logarithmic scale because $\sqrt{\|F\|^2_{av}}$ and the masses now scale with the same factor. Alternatively, to get a different scaling for the masses and the forces, the general recurrence relation may be used

$$\begin{aligned}\tilde{\mu}_2 &= x\tilde{\mu}_1 \\ \tilde{\mu}_n &= \frac{f^n(x-1)+f(x-1)}{f(f-1)}\tilde{\mu}_1\,,\end{aligned}$$

keeping in mind that unlike Equation (5.20), this again ignores the contribution from the phenomenological term $|\tilde{\mu}_n/\tilde{\mu}_{n-1}|$.

MODEL VALIDATION AND APPLICATION

In this section, a validation of the model established in the previous section is attempted by extrapolating the fit functions to the simulation points of a number of production simulations and comparing the predictions for the forces to actual measurements. Based on the target parameters of the $N_f = 2$ twisted mass clover ensemble *cA2.30.24*, a high efficiency simulation is attempted using the ideas from the previous section and it is shown that the scheme is indeed much more efficient than the one that was used to produce this ensemble. Finally, a proposal is made for simulations at the physical pion mass.

Extrapolations

To validate the model, it was extrapolated to the simulation points of three $N_f = 2$ ensembles from Section 6.1 and two $N_f = 2+1+1$ ensembles from Section 6.2.1. The extrapolations are shown in Figure 5.6, compared to measurements of the forces from the actual simulations. It is clear that for larger values of $\tilde{\mu}_1$ and $\tilde{\mu}_2$, the model reasonably predicts the simulation points and provides at least a qualitative idea of the force hierarchy. However, it seems that the measured forces scale with the volume despite being local quantities and thus the model strongly undershoots the measurements for the larger lattices. This is not surprising because the eigenvalue density is different in larger volumes. In addition, as has already been noted in Section 4.3.1, the average and maximal forces depend on c_{sw} and this likely explains why the predictions for ensemble *P125* are quite good while undershooting the measurements for *P205*. Finally, it should be understood that the force measurements are *partially quenched* observables. As a result, much like for other quantities such as meson masses, their values will depend on the effective dynamical field content and hence the target light quark mass. The values determined for the phenomenological parameters a_{av}, a_{max}, b_{av} and b_{max} on the test ensemble are thus certainly inappropriate for other physical situations.

For new (unfitted) simulation points on the test ensemble introduced in this section, the model describes the dependence quite perfectly. This suggests that for wider applicability the model would just need to be extended by terms which capture the remaining dependences. It should be noted, however, that the functional forms obtained above are really quite restrictive due to the leading dependence on $(\tilde{\mu}_2 - \tilde{\mu}_1)^2$. As a result, it seems

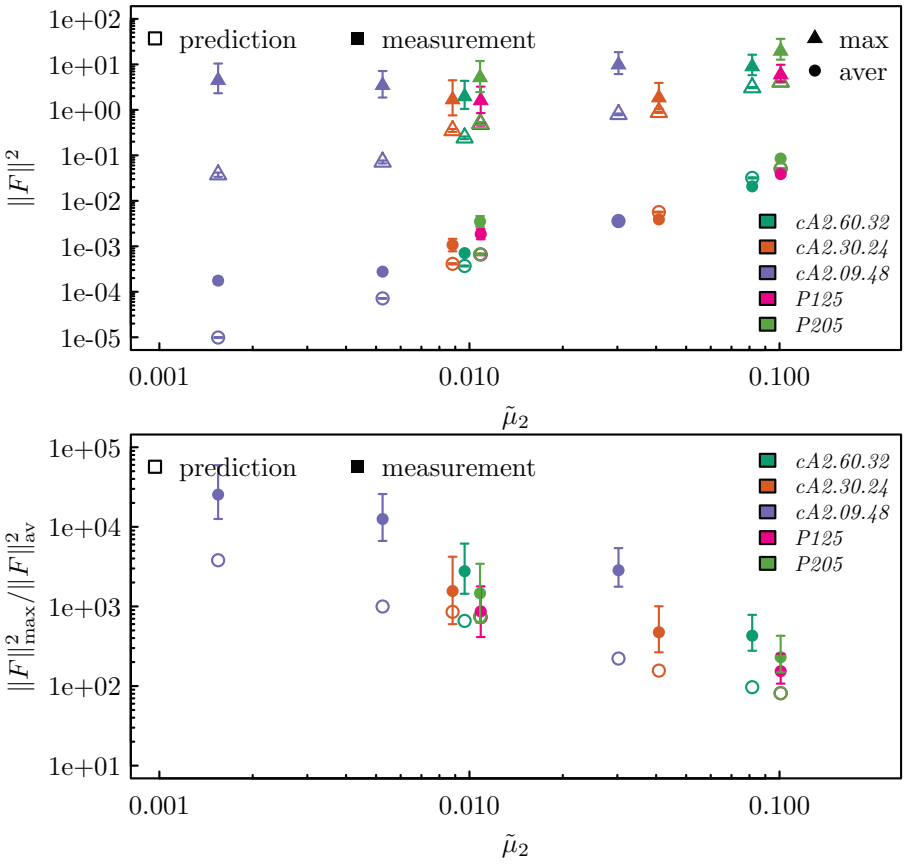

Figure 5.6: Average and maximal forces in $N_f = 2$ (Section 6.1) and $N_f = 2+1+1$ (Section 6.2.1) simulations with the twisted mass clover action. Empty symbols correspond to the predictions of the model in Equations (5.10) and (5.11) with $\tilde{\mu}_1$ and $\tilde{\mu}_2$ chosen as in the respective simulations, while filled symbols are measurements from the actual simulations. *cA2.30.24* corresponds to a $N_f = 2$ simulation with $M_{\pi^\pm} \sim 250$ MeV on a lattice of spatial size $L/a = 24$ while *cA2.09.48* is at the physical pion mass on a $L/a = 48$ lattice. The *P* ensembles are $N_f = 2+1+1$, have $M_{\pi^\pm} \sim 250$ MeV, $L/a = 24$ and different values of c_{sw} at approximately constant physics.

reasonable to expect that for any given set of fundamental parameters (lattice spacing, volume and target mass), it should be possible to obtain reasonable values for the effective parameters a_{av}, a_{max}, b_{av} and b_{max} from just a few measurement points. This should be sufficient to provide a very good description of the dependence of the forces on the mass preconditioning parameters and allow the simulation to be continued with optimized ones. In this way, the functional forms obtained here can be of real practical use.

High Efficiency Simulation of Ensemble cA2.30.24

The results of the previous sections suggest that the forces from the various determinant ratios can be controlled quite predictably through the preconditioning masses. Especially the relations in Equation (5.20) allow the whole tower of forces to be controlled with just one parameter. It will now be shown that the $N_f = 2$ ensemble *cA2.30.24* (see Chapter 3) was simulated with a reasonably efficient choice of parameters, but an excessive number of integration steps. As a metric of simulation cost, the number of CG iterations in the molecular dynamics, heatbath and acceptance steps will be used. This is in contrast to other works which make use of the more accurate total number of matrix vector products. Consequently, the cost estimates are lower bounds and terms should be added to account for the cost of the derivatives of the gauge term as well as the contraction of the force terms of the quark contributions. It should also be remembered that the number of CG iterations can only be compared meaningfully for equal volumes.

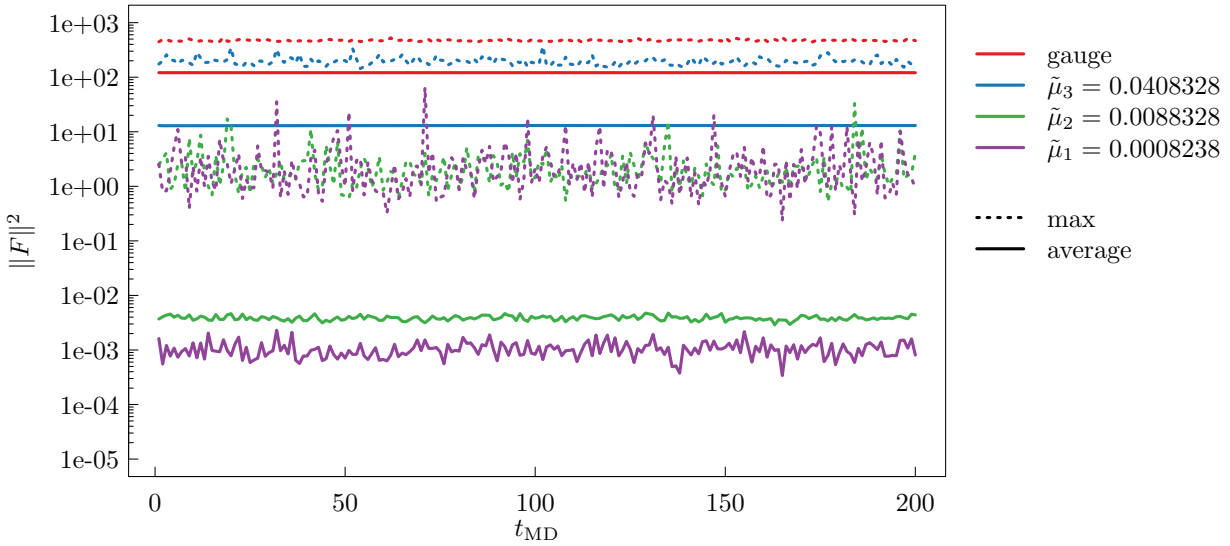

Figure 5.7: Molecular dynamics time histories of the four monomials in the simulation of the *cA2.30.24* ensemble. The force labelled $\tilde{\mu}_3 = 0.0408328$ is a determinant while the two others are determinant ratios.

Table 5.1: Square norm forces in simulation of ensemble *cA2.30.24*. The central value is the median while the positive and negative deviations indicate the 84.27 and 15.73 percentiles respectively. †: This line is for a determinant rather than a determinant ratio.

time scale t	$\tilde{\mu}_{4-t}$	$\|F\|^2_{\mathrm{av}}$	$\|F\|^2_{\mathrm{max}}$
3	0.0008238	$1.1\left(^{+0.4}_{-0.3}\right)\cdot 10^{-3}$	$1.68\left(^{+2.80}_{-0.92}\right)\cdot 10^{0}$
2	0.0088238	$3.9\left(^{+0.4}_{-0.3}\right)\cdot 10^{-3}$	$1.84\left(^{+2.06}_{-0.80}\right)\cdot 10^{0}$
1	0.0408238†	$1.30\left(^{+0.01}_{-0.01}\right)\cdot 10^{1}$	$1.91\left(^{+0.40}_{-0.23}\right)\cdot 10^{2}$
0 (gauge)	–	$1.21\left(^{+0.01}_{-0.01}\right)\cdot 10^{2}$	$4.67\left(^{+0.17}_{-0.13}\right)\cdot 10^{2}$

As shown in Table 3.3 on page 44, the ensemble was simulated on four time scales. For convenience, the mass parameters are listed in Table 5.1. The outermost time scale was integrated with 10 steps of the 2MN integrator and $N_t = \{1, 2, 2, 10\}$ integration steps overall. The acceptance rate was around 91%, but the simulation cost was excessive with a total of about $2.3 \cdot 10^5$ CG iterations in the molecular dynamics step per trajectory. A more efficient simulation with almost the same acceptance rate can be obtained by changing the integration steps to be $N_t = \{1, 1, 1, 13\}$, but the cost per trajectory is still around $1.75 \cdot 10^5$ CG iterations because so many steps are necessary on the outermost time scale.

Time histories of the square norm forces in the original simulation are shown in Figure 5.7, where the average forces are indicated by the lines while the maximal forces are given by the dotted lines. It should be noted how the maximal forces of the two determinant ratios are almost the same while their average square norm forces differ by around a factor of 5. The maximal force of the monomial with the target mass shows fluctuations over two orders of magnitude, thus requiring many integration steps on the coarsest time scale to retain good acceptance. The distribution of the forces is listed in Table 5.1 for the different monomials.

Exploiting what has been learned in the previous sections an alternative mass preconditioning scheme can be established whose primary aim is to provide a real hierarchy of maximal forces and to minimse their fluctuations. To this end, the simulation was carried out on five time scales with three determinant ratios on the three outermost ones. In order to keep the average

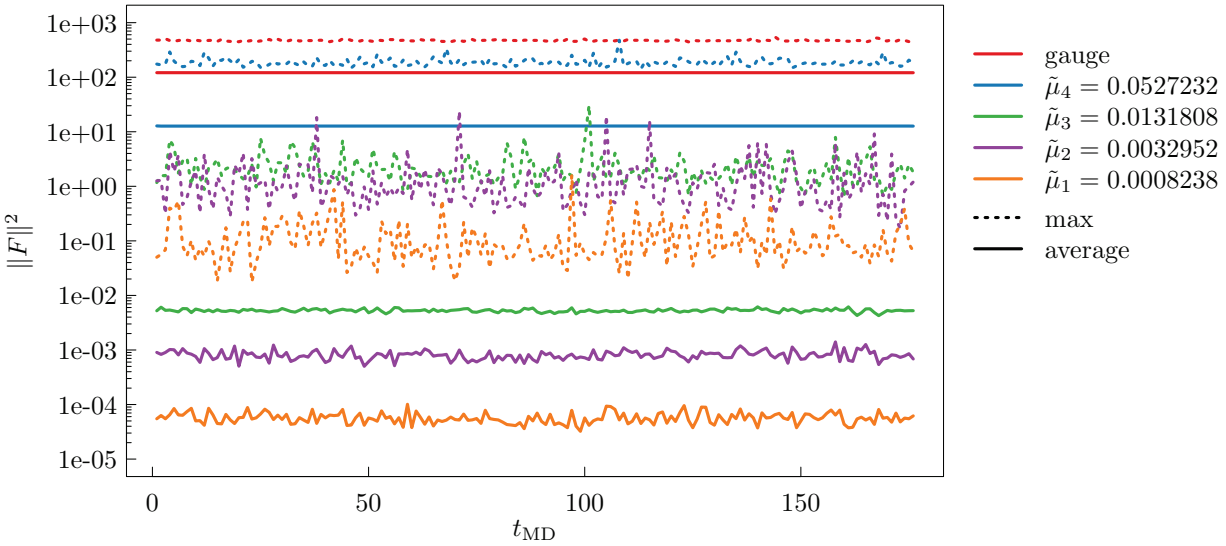

Figure 5.8: Molecular dynamics time histories of the four monomials in the highly optimised simulation of the *cA2.30.24* ensemble. The force labelled $\tilde{\mu}_4 = 0.0527232$ is a determinant while the three others are determinant ratios.

time scale t	$\tilde{\mu}_{5-t}$	$\|F\|^2_{av}$	$\|F\|^2_{max}$
4	0.0008238	$5.54\left(^{+1.54}_{-1.17}\right)\cdot 10^{-5}$	$0.80\left(^{+1.41}_{-0.35}\right)\cdot 10^{-1}$
3	0.0032952	$8.05\left(^{+1.74}_{-1.39}\right)\cdot 10^{-4}$	$0.90\left(^{+1.36}_{-0.68}\right)\cdot 10^{0}$
2	0.0131808	$5.25\left(^{+0.39}_{-0.29}\right)\cdot 10^{-3}$	$1.82\left(^{+1.54}_{-0.67}\right)\cdot 10^{0}$
1	0.0527232†	$1.270\left(^{+0.097}_{-0.002}\right)\cdot 10^{1}$	$1.78\left(^{+0.34}_{-0.17}\right)\cdot 10^{2}$
0 (gauge)		$1.21\left(^{+0.01}_{-0.01}\right)\cdot 10^{2}$	$4.69\left(^{+0.13}_{-0.13}\right)\cdot 10^{2}$

Table 5.2: Square norm forces in the highly optimised simulation of ensemble *cA2.30.24*. The central value is the median while the positive and negative deviations indicate the 84.27 and 15.73 percentiles respectively. †: This line is for a determinant rather than a determinant ratio.

and maximal forces on the outermost time scale small, the masses were chosen as

$$\tilde{\mu}_n = 4^{n-1}\tilde{\mu}_1\,,$$

such that

$$\tilde{\mu}_1 = 0.0008238, \tilde{\mu}_2 = 0.0032952, \tilde{\mu}_3 = 0.0131808, \tilde{\mu}_4 = 0.0527232\,.$$

The number of integration steps was set to be $N_t = \{1,1,1,1,6\}$ with the 2MN integrator on all time scales, providing 192 effective integration steps on the finest one.

The acceptance rate of this scheme was found to be around 90% at a simulation cost of $1.1\cdot 10^5$ CG iterations in the molecular dynamics step per trajectory, a reduction by more than a factor of 2 compared to the original simulation. In particular, the number of CG iterations on the outermost time scale is reduced by a factor of around 5 compared to the optimised original scheme with $N_t = \{1,1,1,13\}$. Because the cost estimate ignores several contributions, the actual cost reduction factor is closer to 2.5. One would expect that $x = 4 < 10$ for the scaling of preconditioning masses would result in significantly more expensive heatbath and acceptance steps. However, this cost in all three schemes discussed here is around $1.1\cdot 10^4$ CG iterations.

The force histories for this highly optimized mass preconditioning and integration scheme are shown in Figure 5.8 and statistical data is listed in Table 5.2. First of all it is notable that a real force hierarchy now exists and the square norm forces on time scales 4 and 3 differ by about a factor of 16, as expected. The force on time scale 2 does not quite fit the prediction of Equation (5.21), but the deviation is small and covered by the distributions of the two average forces. The maximal forces also display a monotonic ordering and their fluctuations are somewhat reduced compared to the original scheme. The maximal force on time scale 2 seems to be quite a bit lower than naively expected. Taking into account the complete x-dependence from Equation (5.20) gives

$$\frac{\|F\|^2_{\max}(4)}{\|F\|^2_{\max}(3)} = \frac{x^4}{|x^5|^{b_{\max}}} \cdot \frac{|x^3|^{b_{\max}}}{x^2} = |x|^{2(1-b_{\max})} = 8.64\,,$$

while the measured ratio is only around 2, again suggesting a strong dependence of $b_{\max}$ on the target quark mass and volume, as the test ensemble was simulated with rather heavy pions.

An extension of the scheme with a sixth time scale was attempted, significantly reducing the force magnitude coming from the determinant factor as suggested in Section 5.1.1. This, however, resulted only in a marginal increase in acceptance while increasing the simulation cost by about 20%. It is likely, however, that for simulations with even lighter quark masses, the guideline suggested in Section 5.1.1 should be followed even though it was not necessary here.

Scheme Proposal for the Physical Pion Mass

The chief aim of this study is to make simulations at the physical pion mass computationally cheaper. In this final section, the information from above will thus be used to propose a scheme for simulations with very small target quark mass. The simulation of ensemble *cA2.09.48* had an acceptance rate of only around 73% and a molecular dynamics cost of around $5.5 \cdot 10^5$ CG iterations per trajectory. In addition, around $4.6 \cdot 10^4$ CG iterations were required for the heatbath and acceptance steps. It was simulated on four time scales with two determinant ratios on the outermost time scale and $N_t = \{1, 1, 2, 13\}$ integration steps with the 2MN integrator overall. This corresponds to 416 effective steps on the finest time scale. The masses are listed in Table 5.3 and the resulting square norm forces are shown in Figure 5.9.

The two determinant ratios on time scale 3 have very similar maximal and average forces, as required for this scheme. However, the maximal forces are both as large as the maximal force coming from the determinant ratio on time scale 2, suggesting that for a better acceptance rate a more hierarchical distribution would be required. Because the target mass $\tilde{\mu}_1$ is about a factor of 3.3 lower than in the simulation of *cA2.30.24*, six time scales are likely to be required to obtain a proper force hierarchy and a sufficiently low force from the determinant factor. The mass scaling factor $x = 4$ would probably work well, although a value $10 > x > 4$ may be beneficial in terms of total simulation cost. This is because $(x-1)^2\tilde{\mu}_1^2$ will be smaller than for *cA2.30.24* even for $x > 4$ and thus the force on the coarsest time scale is expected to be lower as well. On the other hand, the fluctuations in the maximal forces are clearly larger than for *cA2.30.24*, so keeping $x = 4$ might provide

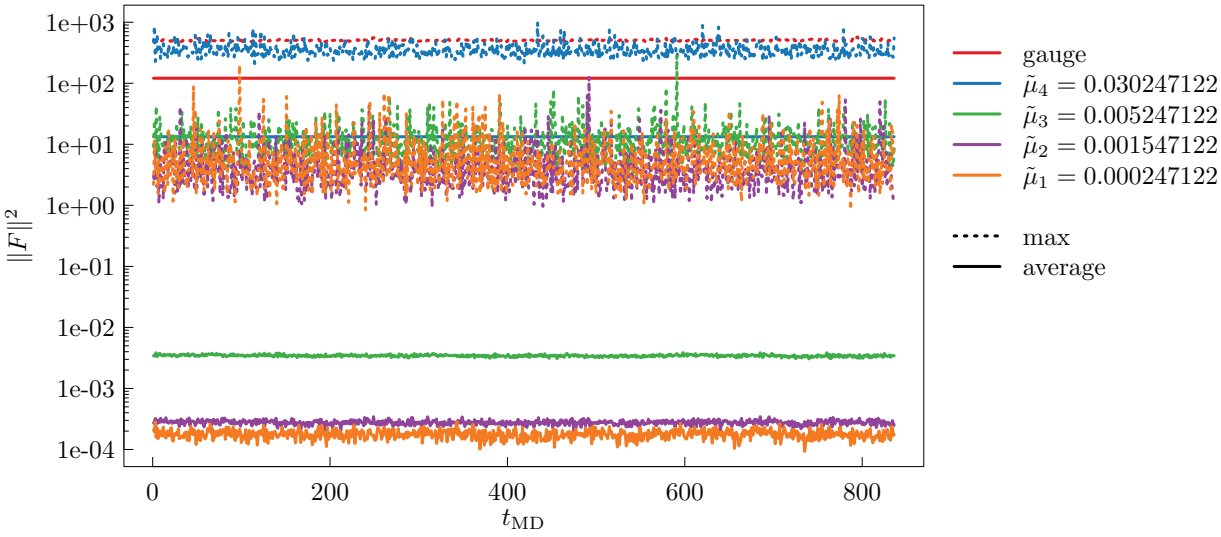

Figure 5.9: Molecular dynamics time histories of the four monomials in the simulation of the *cA2.09.48* ensemble. The force labelled $\tilde{\mu}_4 = 0.030247122$ is a determinant while the two others are determinant ratios.

time scale t	$\tilde{\mu}_n$	$\|F\|^2_{\mathrm{av}}$	$\|F\|^2_{\mathrm{max}}$
3	0.000247122	$1.76\left(^{+0.31}_{-0.28}\right)\cdot 10^{-4}$	$4.47\left(^{+5.92}_{-2.15}\right)\cdot 10^{0}$
3	0.001547122	$2.77\left(^{+0.24}_{-0.23}\right)\cdot 10^{-4}$	$3.47\left(^{+3.67}_{-1.60}\right)\cdot 10^{0}$
2	0.005247122	$3.44\left(^{+0.13}_{-0.13}\right)\cdot 10^{-3}$	$9.79\left(^{+8.68}_{-3.68}\right)\cdot 10^{0}$
1	0.030247122†	$1.33\left(^{+0.01}_{-0.01}\right)\cdot 10^{1}$	$3.29\left(^{+1.07}_{-0.54}\right)\cdot 10^{2}$
0 (gauge)	–	$1.33\left(^{+0.01}_{-0.01}\right)\cdot 10^{2}$	$5.03\left(^{+0.16}_{-0.13}\right)\cdot 10^{2}$

Table 5.3: Square norm forces in the simulation of ensemble *cA2.09.48*. The central value is the median while the positive and negative deviations indicate the 84.27 and 15.73 percentiles respectively. †: This line is for a determinant rather than a determinant ratio.

the right balance between intrinsic simulation cost and a dampening of the fluctuations in the maximal forces.

The number of integration steps on the outermost time scale will probably have to be no smaller than 9, on account of the larger volume, but of course this can always be adjusted downwards if the acceptance rate is good. The proposed scheme would thus be

- $N_t = \{1, 1, 1, 1, 1, 9\}$
- $\tilde{\mu}_1 = 0.000247122, \tilde{\mu}_2 = 0.000988488, \tilde{\mu}_3 = 0.003953952,$ $\tilde{\mu}_4 = 0.015815808, \tilde{\mu}_5 = 0.063263232, \tilde{\mu}_6 = 0.253052928$

The resulting simulation cost can be predicted quite easily. The number of CG iterations for a given volume, set of fundamental parameters and termination criterion can be fitted reliably to a function of the form

$$N^{\mathrm{CG}}_{\mathrm{iter}}(\tilde{\mu}) \sim \frac{a}{\tilde{\mu} + b},$$

with a and b positive real parameters.

The resulting function is valid up to $\tilde{\mu}$ slightly smaller than the lightest $\tilde{\mu}$ used in the fit. It works very well for larger $\tilde{\mu}$, as long as the same stopping criterion is used for all $\tilde{\mu}$ considered. Doing so for the simulation data from ensemble *cA2.09.48*, gives $a = 7.32$, $b = 0.00042$. Using the number of effective integration steps on each time scale and the number

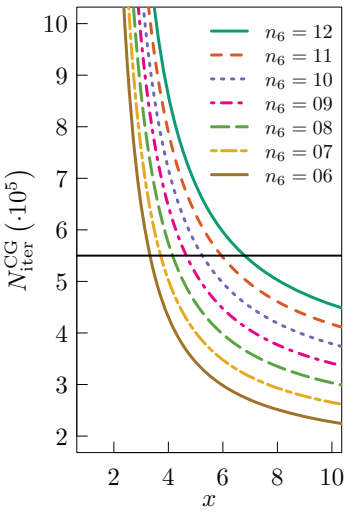

Figure 5.10: Simulation cost per trajectory in terms of number of CG iterations in the molecular dynamics step with an optimised mass preconditioning scheme with 6 time scales, n_6 steps on the outermost time scale and a mass scaling factor x. The horizontal black line indicates the simulation cost of the original scheme for the *cA2.09.48* ensemble.

of CG iterations predicted by the formula above, gives a cost of $5.66 \cdot 10^5$ CG iterations for the original scheme. This is in line with what was observed in the simulation and thus validates this cost model. For the proposed scheme, a total cost of $6.3 \cdot 10^5$ CG iterations is expected for $x = 4$ and $4.2 \cdot 10^5$ for $x = 6$. Figure 5.10 shows the dependence of the simulation cost on x and the number of integration steps n_6 on the outermost time scale, with the cost of the original scheme indicated by the black horizontal line. It is clear that for $n_6 = 9$ as proposed, $x \gtrsim 5$ is required to obtain the same or a lower total simulation cost. One should keep in mind that for $x < 10$, a higher acceptance rate might be expected than with the original scheme, justifying the extra cost if the right balance is reached. It was unfortunately not possible to test this scheme in production, but it will certainly be tested in the future in $N_f = 2+1+1$ simulations at the physical pion mass.

OPTIMIZED MASS PRECONDITIONING

The results of the previous sections can now be used to propose a simple general algorithm for an optimized mass preconditioning scheme depending only on one parameter and the number of integration steps on the outermost time scale. Given a target quark mass $\tilde{\mu}_1 = \tilde{\mu}_\ell$, a mass scaling factor $x < 10$ is chosen and N_{ts} time scales are used in the integration such that

$$\tilde{\mu}_{N_{ts}} = x^{N_{ts}-1}\tilde{\mu}_1 \gtrsim 0.2$$

in order to minimize the force from the determinant factor as suggested by the results of Section 5.1.1. The preconditioning masses are thus given by

$$\tilde{\mu}_n = x^{n-1}\tilde{\mu}_1$$

and the results from above suggest that the average forces will scale roughly by the factor x from time scale to time scale. Note that in practice, at least in the tmLQCD software suite when simulating the twisted mass clover action, the preconditioning masses are specified via the $\tilde{\rho}_i$, in which case

$$\tilde{\rho}_n = \left(x^{n-1} - 1\right)\tilde{\mu}_1$$

for the scheme proposed here.

The scaling factor determines the forces on the coarsest time scale, and these in turn determine how many integration steps need to be carried out there. On all other time scales the number of integration steps is set to 1, which effectively corresponds to 2 or 4 when a second or fourth order integrator is used. In this way the effective doubling or quadrupling of integration steps from time scale to time scale provides a close to optimal integration scheme, when x is chosen appropriately. Although there is no guarantee that the proposed scheme will be efficient and provide a high acceptance rate, it certainly produces a predictable hierarchy of forces and can be used as a starting point. Refinements can be made to the parameters if a given force turns out to be much larger or smaller than expected. The full formalism in Equations (5.7) and (5.8) can then be used to predict the resulting changes.

SUMMARY AND DISCUSSION

The results of this chapter show that the molecular dynamics forces coming from the mass preconditioning determinant splitting can be fitted to simple

empirical functions of the mass preconditioning parameters. While the used fit formulae are not motivated theoretically, the functional dependence on the preconditioning parameters is strongly constrained by the requirement that the forces from determinant ratios vanish in the limit $\tilde{\mu}_1 \to \tilde{\mu}_2$. The used functional forms thus have only two free parameters each for the average and maximal square norm of the forces for any combination of $\tilde{\mu}_1$ and $\tilde{\mu}_2$.

These functional forms were fitted to measurements of the forces in a small volume at a relatively heavy pion mass and the functions were found to reproduce the data rather well, where it should be noted that both the forces and the mass preconditioning parameters span several orders of magnitude. In order to take into account the way even-odd preconditioning is implemented for the twisted mass clover action, the fits were repeated for two target masses $\tilde{\mu}_\ell$. The model predictions were checked against simulation data with different volumes, different pion masses, different N_f and different β and c_{sw}. As expeted, it was found that the model deviates significantly for simulations with different fundamental parameters, but the simplicity of the functional forms suggests that fitting these for a given physical situation can be achieved with just a few measurements. In practice it is hoped that this will provide empirical guidance for the tuning of the mass preconditioning parameters.

It has been shown that to leading order, the forces of determinant ratios are proportional to $(\mu_2 - \mu_1)^2$ and this can be exploited rather easily in the tuning of the force contributions in practice. In particular, when the masses are scaled as $\tilde{\mu}_n = \varkappa^{n-1}\tilde{\mu}_1$, a significant simplification of the parameter dependence occurs. Assuming that the fitted functional forms offer a correct effective description of the forces, this simplification allows the proposal of a general algorithm for the choice of a reasonable mass preconditioning and integration scheme.

It has been shown that for ensemble *cA2.30.24*, a mass preconditioning and integration scheme of this type increases simulation efficiency by a factor of 2.5 compared to the original scheme used for the simulation. For the physical pion mass, a similar scheme was devised which, it is hoped, should provide a higher acceptance rate at a very similar cost. Unfortunately, this choice of simulation parameters could not yet be tested, but an analysis of the forces in the original *cA2.09.48* simulation suggests that it should be effective. More generally, the presented recurrence relations can be used to scale the forces and the mass parameters with different factors which may allow further efficiency gains.

The results should also be valid for standard Wilson fermions (with and without clover term), as long as the hopping parameter κ is not used for additional preconditioning purposes. Specifically, for the simulations carried out in the CLS effort [84], which employ small twisted quark masses to stabilise simulations, the simplified formulae may be used to establish mass preconditioning and integration schemes directly.

An understanding of the forces coming from different terms in the effective action is not sufficient to fully predict the efficiency of a given simulation. In particular, the presented results do not allow conclusions to be drawn regarding expected acceptance rates. Refs. [149, 150, 151, 152] advance an argument based on *Shadow Hamiltonians*. These are conserved exactly and can be represented by asymptotic expansions of the molecular dynamics Hamiltonian with correction terms proportional to powers of the step sizes $\delta\tau_i$. In these analyses it is suggested that the acceptance rate does not depend on the first derivatives with respect to the gauge field of the

effective action (the forces), but rather on the second derivatives. In some sense, the attempt in the present chapter to focus on reducing $\Delta^+ \left(\|F\|^2_{\max} \right)$ is a reflection of this argumentation. A final aspect which has been ignored here is the interplay between multiple time scales, the magnitude of the forces and the tuning parameters λ_i present in second and higher order symplectic integrators [74].

It would be highly insightful to attempt a theoretical motivation of the proposed functional forms such as through expansions of the analytical expressions for the force around the free field limit, although it is unclear whether possible non-perturbative corrections can be captured in this way. This may in turn allow the results of Refs. [149, 150, 151, 152] to be used in order to draw conclusions on the expected acceptance rate and find sets of parameters which are optimal with respect to some well-motivated criterion. In any case even at this stage, it is the opinion of the present author that the proposed schemes can serve as a very good starting point and refinements can be made based on the acceptance rate during the initial stages after thermalisation of a given simulation.

Part III

PHYSICS RESULTS

The $N_f = 2$ and $N_f = 2+1+1$ twisted mass clover ensembles generated as discussed in Chapters 3 and 4 were used to determine a number of benchmark quantities. As an issue of central importance to simulations with twisted mass quarks at small quark masses, the effects of the explicit isospin breaking are discussed in Chapter 6. The pion mass splitting is measured on the $N_f = 2$ ensembles and attempts are made to extrapolate it to the physical light quark mass, as shown in Section 6.1. A more extensive analysis is done on $N_f = 2+1+1$ ensembles, for which approximately constant physical situations were simulated at different values of c_{sw}. The pion mass splitting is also measured there and the effect of the value of c_{sw} is discussed in Section 6.2.

In order to test the suitability of the twisted mass clover action for the computation of physical quantities, a study of the pseudoscalar meson sector was carried out on the $N_f = 2$ ensembles. Focussing on the *cA2.09.48* ensemble at the physical charged pion mass, pseudoscalar meson masses and decay constants and their many ratios are computed in Chapter 7 with a careful comparison to previous lattice and phenomenological determinations. Values for the average up/down, strange and charm quark masses are also provided, computed directly from the bare quark masses used in the analysis. A particular focus is put on systematic errors stemming from the ambiguity of the fit range, which might play an increasing role in calculations on large lattices and an analysis technique is introduced for this purpose.

6 ISOSPIN SYMMETRY BREAKING

The connection between the pion mass splitting, the zero temperature phase structure for Wilson fermions and the stability of simulations has been discussed in Section 2.2. In this chapter, measurements of the pion mass splitting using the new twisted mass clover action in $N_f = 2$ and $N_f = 2+1+1$ simulations will be presented.

The mass of the neutral pion can be determined from correlation functions of the interpolating operator P^0 in the twisted basis

$$P^0 = \sum_{\mathbf{x}} \bar{\chi}_\ell(\mathbf{x}, t)\, \mathbb{1}_f\, \chi_\ell(\mathbf{x}, t)\,, \tag{6.1}$$

where $\mathbb{1}_f$ denotes the unit matrix in flavour space. The corresponding correlation function $C_{P^0}(t)$ has connected and disconnected contributions and is, therefore, noisy. In order to extract the mass of the neutral pion, ground state exponentials with forward and backward propagating states are fitted to the correlation function in a region where the related effective mass shows an approximate plateau and where the signal quality is still sufficiently high for the fits to converge. The details of the analysis will not be discussed here as they can be found in Section 7.1 and Section 7.4. As will be discussed further below, at large Euclidean times the connected contribution to $C_{P^0}(t)$ is exponentially suppressed because the connected neutral pion is significantly heavier than the full neutral pion. In order to benefit from the improved signal quality when the connected piece is ignored, the neutral pion mass is extracted from the disconnected piece only.

TWO FLAVOUR TWISTED MASS CLOVER PION MASS SPLITTING

The motivation for twisted mass clover fermions stems from a quenched study [118], in which it was shown that the mass splitting between the charged and neutral connected pions is significantly reduced compared to standard quenched twisted mass fermions. In this section, it will be shown that this continues to be the case for simulations with two flavours of mass-degenerate twisted mass clover quarks all the way down to the physical pion mass. In addition, there are strong indications that also the mass splitting between the charged pion and the full neutral pion is reduced significantly. Parts of these results have been published in Ref. [P6].

It has already been shown in Chapter 3 that no meta-stabilities are observed with the new action at the lattice spacing of $a \sim 0.092$ fm at the physical pion mass. Figure 3.1 on page 41 further shows that the slope of am_{PCAC} as a function of the untwisted quark mass is very close to 1. In addition, although the number of points is quite limited, there does not seem to be a significant μ-dependent shift as discussed at the end of Section 2.2.2, indicating that c_2 must be very small. Although the gauge action is different, this might be a reflection of results from Ref. [113] (with the plaquette gauge action), that $c_2 \gtrsim 0$ for two flavours of non-perturbatively improved Wilson fermions.

Table 6.1 lists measurements of the charged, neutral connected and full neutral pion masses on the twisted mass clover $N_f = 2$ ensembles discussed

Table 6.1: Charged, neutral connected and full neutral pion masses in lattice units determined on $N_f = 2$ twisted mass clover ensembles. †: For *cA2.09.48*, disconnected diagrams have not been computed yet. ⋆: For the full neutral pion, a study of systematic effects from excited states was not possible due to the poor signal. N_{meas} represents the number of measurements for meson quantities on the corresponding ensemble.

	cA2.09.48	*cA2.30.24*
$aM_{\pi^\pm}$	$0.06196(09)(^{+12}_{-05})$	$0.1147(7)(^{+4}_{-7})$
$aM_{\pi^{(0,c)}}$	$0.1191(05)(^{+07}_{-10})$	$0.1541(13)(^{+05}_{-05})$
aM_{π^0}	–†	$0.09(1)^\star$
N_{meas}	1457	728
	cA2.60.24	*cA2.60.32*
$aM_{\pi^\pm}$	$0.15941(38)(^{+15}_{-21})$	$0.15769(26)(^{+15}_{-14})$
$aM_{\pi^{(0,c)}}$	$0.18981(61)(^{+21}_{-25})$	$0.18840(44)(^{+46}_{-29})$
aM_{π^0}	$0.11(2)^\star$	$0.13(2)^\star$
N_{meas}	1351	670

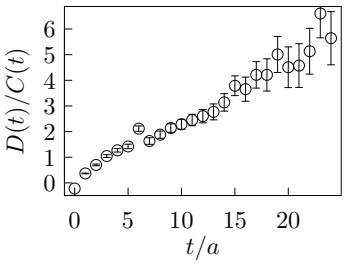

Figure 6.1: Disconnected piece $D(t)$ of the neutral pion correlation function compared to its connected part $C(t)$ on ensemble *cA2.30.24*.

in Chapter 3. The data is presented in Figures 6.3a and 6.3b in units of r_0, in the form which allows it to be compared to the results of Ref. [124]. In particular, it can be seen in Figure 6.3a that the mass splitting between the squared charged and neutral connected pion masses is about a factor of three smaller than with the old action without the clover term. It also seems that a constant chiral extrapolation is appropriate for this quantity. The difference between the squared neutral connected and full neutral pion masses is shown in Figure 6.3b, indicating that this too is significantly reduced.

A measurement of the full neutral pion mass on ensemble *cA2.09.48* is complicated by the poor signal to noise ratio in large volume as well as the expense of computing the disconnected contributions at this light quark mass. It must hence be deferred to a future publication, employing some variance reduction technique such as the generalized hopping parameter expansion [153]. However, even the data at larger than physical pion masses should be considered with care because the plateau quality is low and no proper assessment of systematic errors could be carried out given the poor signal to noise ratio. In addition, the data was analysed under the assumption that the neutral pion is in fact lighter than the charged ones and as a result also unphysical connected contribution to the neutral pion correlation function. Hence, the full neutral pion mass was extracted *only* from the disconnected contribution. This approach is justified because the connected part of the neutral pion correlation function is exponentially suppressed compared to the disconnected contribution, as shown in Figure 6.1. Although this gives a better signal quality, it may have biased the analysis to give lower values for the full neutral pion mass, especially on the ensembles with $a\mu_\ell = 0.006$ where the relative suppression is not as strong.

While the measurements above are useful for extracting individual contributions to c_2, it is the mass of the full neutral pion which affects the stability of simulations and exponentially enhanced finite size effects. Figure 6.3c shows constant and linear extrapolations of $a^2(M^2_{\pi^\pm} - M^2_{\pi^0})$ to the physical light quark mass. In the constant extrapolation, the value is strongly constrained by the (underestimated) uncertainty on the point at $a\mu_\ell = 0.003$ while the linear extrapolation would likely require at least one more point to be trustworthy. In any case, plugging in $M^2_{\pi^\pm}$ at $a\mu_\ell = 0.0009$,

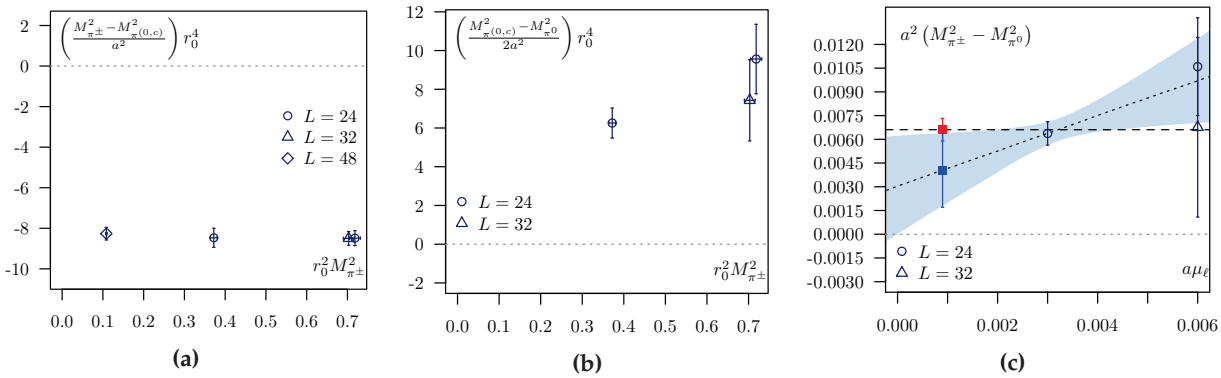

Figure 6.3: Charged and neutral pion mass splittings **(a)**: $(M^2_{\pi^\pm} - M^2_{\pi^{(0,c)}})/a^2$ and **(b)**: $(M^2_{\pi^{(0,c)}} - M^2_{\pi^\pm})/2a^2$ as a function of $M^2_{\pi^\pm}$ in units of r_0 **(c)**: Chiral extrapolation to the physical light quark mass of $a^2(M^2_{\pi^\pm} - M^2_{\pi^0})$ with a constant fit (■) and a linear fit (■).

both extrapolations give negative values for the squared neutral pion mass with (likely underestimated) uncertainties.

$$M^2_{\pi^0}(a\mu_\ell = 0.0009)_{\text{cst}} = -[112(15)\ \text{MeV}]^2$$
$$M^2_{\pi^0}(a\mu_\ell = 0.0009)_{\text{lin}} = -[25(180)\ \text{MeV}]^2$$

Clearly, this is not only unphysical, it is also in contradiction with the indications from stable simulations and should be checked with an explicit calculation at the physical light quark mass. Aside from the underestimated uncertainties and the issues with the signal for the neutral pion mass, this may also be a sign of chiral logarithms [112] which may enhance the pion mass splitting in a range of intermediate quark masses. Ignoring these logarithms would then clearly lead to the wrong extrapolation.

For direct comparison to simulations with the old $N_f = 2$ and $N_f = 2+1+1$ actions, the pion mass splitting on ensemble *cA2.30.24* is around 70(10) MeV, where it should be remembered that the uncertainty is underestimated. Given that the lattice spacing is around 15% higher than the coarsest lattice spacing employed in $N_f = 2$ simulations with the old action, this compares quite favourably to the 60 MeV measured there at a slightly higher pion mass. However, given all the caveats involved in these measurements, firm conclusions can simply not be drawn at present.

STUDY OF THE PION MASS SPLITTING AS A FUNCTION OF THE SW COEFFICIENT

The results of the previous section and the successful simulation at the physical pion mass with two flavours of twisted mass clover quarks at the physical charged pion mass suggest that the clover term successfully reduces the pion mass splitting. However, as discussed in Section 2.2.3, the pion mass splitting and the phase structure of the lattice theory have shown a strong dependence on the number of active flavours in the past. This means that while simulations at the physical pion mass may be feasible for $N_f = 2$, the pion mass splitting may turn out to be too large in $N_f = 2+1+1$ simulations.

Unpublished studies in the quenched approximation by the ETMC suggest that the magnitude of the pion mass splitting depends on the value of c_{sw}. In particular, it was shown that the neutral connected to charged pion mass splitting is reduced when going from the tree-level value $c_{sw} = 1.0$ to

	P205	*P185*	*P165*	*P145*
β	1.658	1.700	1.745	1.791
c_{sw}	2.05	1.85	1.65	1.45
κ	0.136530	0.138857	0.141140	0.143431
$\langle P \rangle$	0.52749(7)	0.53563(3)	0.54431(3)	0.55327(3)
$a m_{PCAC}$	−0.00038(14)	−0.00040(14)	0.00027(12)	0.00029(15)
$aM_{\pi^\pm}$	0.1298(9)	0.1305(10)	0.1281(8)	0.1241(9)
$aM_{\pi^{(0,c)}}$	0.196(4)	0.203(5)	0.200(4)	0.191(3)
aM_{π^0}	0.102(10)	0.086(10)	0.097(10)	0.091(14)
aM_K	0.2639(19)	0.2590(14)	0.2513(14)	0.2427(13)
w_0/a	1.860(11)	1.810(8)	1.768(14)	1.777(12)
a [fm]	0.0944(12)	0.0969(12)	0.0993(15)	0.0988(14)
$M_{\pi^\pm}$ [MeV]	∼ 270	∼ 260	∼ 250	∼ 250
N_{traj}	1600	2300	1400	1500
N_{meas}	250	250	250	250

Table 6.2: Parameter values and numbers of thermalised trajectories of the $N_f = 2+1+1$ ensembles used for the determination of the c_{sw} dependence of the pion mass splitting. All ensembles have $L/a = 24$, $a\mu_\ell = 0.003$, $a\mu_\sigma = 0.1355$ and $a\mu_\delta = 0.145$. N_{traj} is the number of thermalised trajectories while N_{meas} is the number of measurements for the pseudoscalar meson correlators. Lattice spacings from $w_0 = 0.1755(19)$ fm [142].

the non-perturbatively tuned value for the plaquette gauge action. Based on this, it can be imagined that there may exist *optimal* values of c_{sw} at each lattice spacing, different from the non-perturbatively determined value for $\mathcal{O}(a)$-improvement of standard Wilson quarks, which minimize the pion mass splitting.

As long as they behave smoothly under variations of the lattice spacing, these optimal values could be used consistently in order to reduce or even eliminate the pion mass splitting. Lacking clear guidance from the Symanzik expansion, a numerical check of the c_{sw} dependence of the pion mass splitting was carried out. For this purpose, four $N_f = 2+1+1$ ensembles with values of c_{sw} in the range 1.45 to 2.05 were tuned to approximately constant physical situation, with a particular focus on matching their lattice spacings.

Ensembles and Constant Physics

In order to match the lattice spacing for different values of c_{sw}, the tuning in Chapter 4 was performed and the functional form in Equation (4.12) was determined. With this information, four $N_f = 2+1+1$ simulations with relatively light pion masses of around 250 MeV at four values of c_{sw} were tuned to maximal twist and the charged, neutral connected and full neutral pion masses were determined on 250 well-separated gauge configurations each. A list of these ensembles is given in Table 6.2, hereinafter referred to as the *P* ensembles.

It should be noted that the lattice spacings are relatively well matched, differing at a level of at most 5%, as far as can be determined from the scale $w_0 = 0.1755(19)$ fm computed in Ref. [142] for $N_f = 2+1$. It has to be kept in mind that w_0/a may be subject to differing discretization artefacts between the four ensembles and w_0 differs between $N_f = 2+1$ and $N_f = 2+1+1$, as summarized recently in Table 9 of Ref. [154]. The apparent mismatch in lattice spacings means that the observed pion mass splitting may differ by up to 10% just by virtue of the lattice spacing and this should be kept in mind in the following.

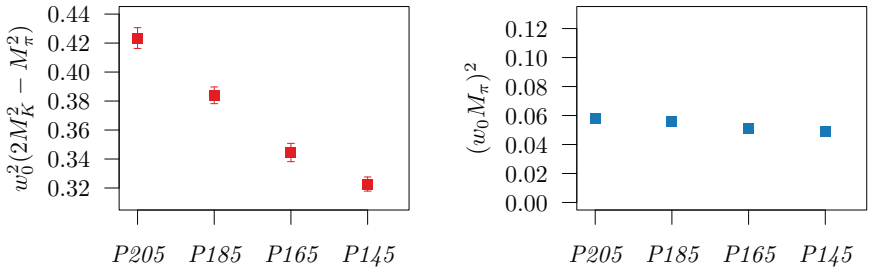

Figure 6.4: $2M_K^2 - M_\pi^2$ (**top**) and M_π^2 (**bottom**) in units of w_0 for the four ensembles from Table 6.2. Note that the vertical ranges are the same.

For the purpose of computing the pion mass splitting it is important to check that the dynamical quark content is consistent between the four ensembles. Firstly, it is known that lattice artefacts are affected by how close the charm quark mass is to the cut-off scale $1/a$ and secondly the value of Z_P/Z_S, on which the effective strange and charm quark masses depend, may be affected by the value of c_{sw}. The importance of these facts becomes clear when comparing the neutral pion mass measurements in Ref. [124] for ensembles *A80.24* and *A80.24s*, the latter of which has a strange quark mass lower by about 15-20%. The mass splitting between the charged and neutral pion is reduced markedly from *A80.24* to *A80.24s*, while the mass splitting between the purely connected contributions is unaffected.

The bare mass parameters $a\mu_\sigma$ and $a\mu_\delta$ were checked on ensemble *P185* to reproduce the phenomenological value of $M_{D_s}/f_{D_s} \sim 8$ and the quark mass ratio $m_c/m_s \sim 12$. This was done by computing aM_K and M_{D_s}/f_{D_s} using Osterwalder-Seiler (OS) valence quarks with a number of different strange and charm quark masses. Interpolating to the point $M_{D_s}/f_{D_s} \sim 8$ in the OS charm and strange quark masses along the lines where $m_c/m_s \sim 12$ fixes the strange and charm quark masses and gives $a\mu_\sigma = (a\mu_s + a\mu_c)/2 = 0.1355$. The parameter $a\mu_\delta$ can then be fixed by matching the OS and unitary kaon masses through interpolations in $a\mu_\delta$ at fixed $a\mu_\sigma$, giving $a\mu_\delta = 0.145$. This procedure is explained in more detail in Section 4.1.2.

In leading order SU(3) chiral perturbation theory, the quantity $2M_K^2 - M_\pi^2$ is proportional to the strange quark mass. It is shown in Figure 6.4, together with the squared pion mass in units of w_0 for the four *P* ensembles. From this it can be seen that the average up/down quark masses are well matched, the slight deviation of up to 5% corresponds to the mismatch in lattice spacings. The strange quark mass, however, seems to differ by up to 25% between the four ensembles. Since this difference cannot be attributed to the mismatch in lattice spacings, one is drawn to conclude that the value of Z_P/Z_S has a strong dependence on c_{sw}. For the measurement of the pion mass splitting, the deviation should be kept in mind. However, unlike in the quoted example from Ref. [124], the strange quark mass is at least a factor 2.7 or so heavier than the light quark mass, which may reduce the impact of the mismatch. Finally, the relative effect of a deviation in Z_P/Z_S on the charm quark mass is much milder (since $a\mu_\sigma$ is unaffected), suggesting that all four charm quark masses are similarly close to the cut-off scale.

Mass Splitting

The absolute differences between the squared connected neutral pion mass $M^2_{\pi^{(0,c)}}$ and the charged pion mass $M^2_{\pi^\pm}$ are shown in the left panel in the first row of Figure 6.5 for the four *P* ensembles. The difference between the squared charged pion mass and the squared (full) neutral pion mass

	A1	*A*1*	*A*2*	*B1*	*B*1*	*B*2*	*C1*	*C2*
L/a	32	32	32	32	48	48	24	24
N_f	2	2	2	$2+1+1$	$2+1+1$	$2+1+1$	2	2
β	3.9	4.05	4.05	1.9	2.1	2.1	2.1	2.1
c_{sw}	–	–	–	–	–	–	1.57551	1.57551
$a\mu_\ell$	0.004	0.003	0.006	0.003	0.0015	0.003	0.003	0.006
$aM_{\pi^\pm}$	0.1338(02)	0.1038(6)	0.1432(6)	0.1234(03)	0.0695(3)	0.0978(4)	0.1147(10)	0.1594(4)
$aM_{\pi^{(0,c)}}$	0.2080(30)	0.150(3)	0.180(2)	0.2111(35)	0.1124(15)	0.1296(15)	0.1541(14)	0.1898(7)
aM_{π^0}	0.1100(80)	0.090(6)	0.123(6)	0.0811(50)	0.056(3)	0.086(5)	0.09(1)	0.11(2)
w_0/a	1.715(5)	2.21(1)	2.192(5)	1.701(13)	2.609(7)	2.544(6)	1.8517(55)	1.8142(41)
a [fm]	0.079(3)	0.063(2)	0.063(2)	0.0885(36)	0.0619(19)	0.0619(19)	0.0914(23)	0.0914(23)
$M_{\pi^\pm}$ [MeV]	~ 340	~ 330	~ 450	~ 280	~ 220	~ 310	~ 250	~ 340

Table 6.3: $N_f = 2$ and $N_f = 2+1+1$ reference ensembles for the pion mass splitting, relabelled for brevity. ***Ax,Bx***: pion mass measurements taken from Ref. [124], w_0/a values should not be seen as final as they are taken from an unpublished ETMC report, as measured at these quark mass mass values (rather than in the chiral limit). ***Ax***: Lattice spacing estimates are from Table 1, Ref. [96]. ***Bx***: Lattice spacing estimates are from Equation 31, Ref .[128]. ***C1,C2***: w_0/a, pion masses and lattice spacing estimate as published in Ref. [P6].

$M^2_{\pi^0}$ are shown in the left panel in the second row of Figure 6.5. The data is presented in units of w_0^2 to account for the apparent slight mismatch in lattice spacings, although it should be kept in mind that the four values of w_0/a may be subject to different lattice artefacts on account of the value of c_{sw} as well as the mismatch in strange and charm quark masses. The following two rows of the same figure give the linear mass splittings relative to the charged pion mass on the given ensembles, thus avoiding potential ambiguities coming from w_0/a. Although this quantity diverges in the chiral limit at finite lattice spacing, it is a useful scale-independent measure when the pion masses are well matched, which is the case for the *P* ensembles.

It seems that within the present precision, all mass splittings are compatible within errors and that the value of c_{sw} has no effect on their magnitude. Further, the compatibility of the relative mass splittings suggests that the effect of c_{sw} on w_0/a is small, at least as far as can be seen in these combinations and within these rather large uncertainties. Finally, very much unlike the situation in Ref. [124], the value of the strange quark mass does not seem to affect the mass splitting within errors. It should be kept in mind, however, that the bare mass parameters in lattice units are identical between the four *P* ensembles which was not the case for the ensembles in Ref. [124].

As a comparison, the mass splittings from $N_f = 2$ and $N_f = 2+1+1$ simulations without the clover term as well as the mass splittings already presented in Section 6.1 with $N_f = 2$ twisted mass clover fermions are shown in the right panels of Figure 6.5. The data for these reference ensembles is listed in Table 6.3, where for the ensembles without the clover term, the values for w_0/a stem from an unpublished internal ETMC report[1] and should not be taken as final. For consistency, the quoted values of w_0/a are given as measured at the respective light quark masses rather than in the chiral limit, which would be higher by a few percent. In addition, one should note the N_f, gauge and fermion action dependence of w_0/a, which becomes apparent when the quoted values are related to the quoted lattice spacings which were determined in Refs. [96, 128, P6] by other means. It seems that for similar lattice spacings the values of w_0/a may differ by up to 10% depending on the details of the discretisation used and the dynamical

[1] Dated 2014.03.19. The present author would like to express his gratitude to U. Wenger for providing it.

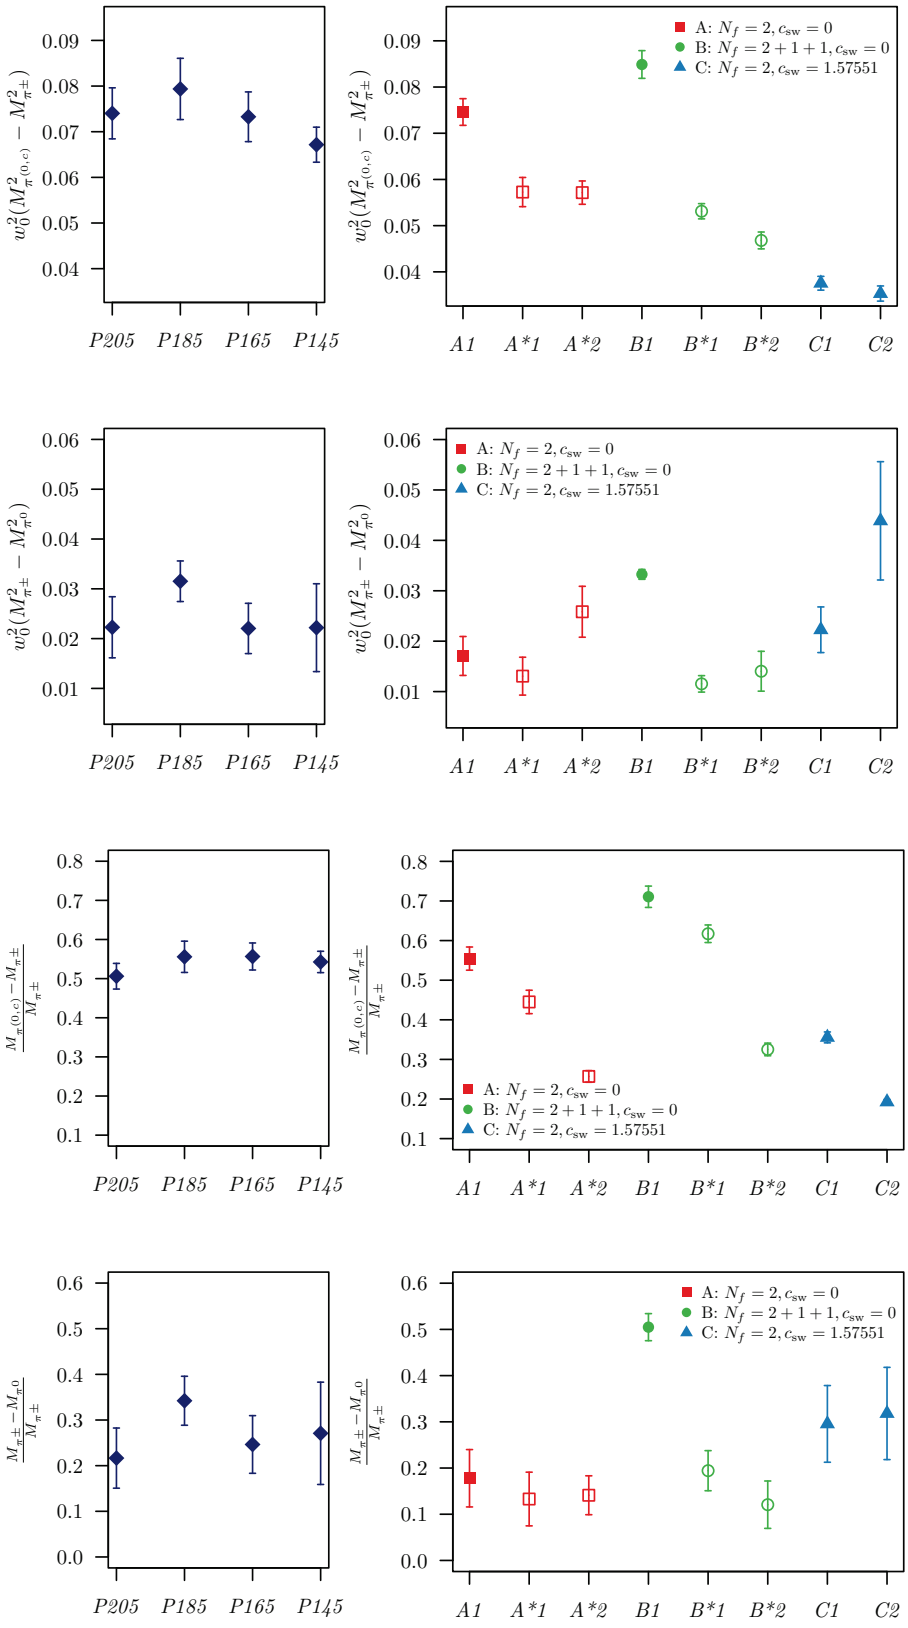

Figure 6.5: $\left(M^2_{\pi^{(0,c)}} - M^2_{\pi^\pm}\right)$ and $\left(M^2_{\pi^\pm} - M^2_{\pi^0}\right)$ in units of w_0 as well as the linear mass splitting relative to the charged pion mass for the four ensembles from Table 6.2 (left) and the reference ensembles from Table 6.3 (right). For the reference ensembles, starred ensemble names and empty symbols indicate finer lattice spacings while higher numbers in the ensemble names indicate heavier light quark masses. The *A* ensembles are $N_f = 2$ twisted mass w/o clover, the *B* ensembles are $N_f = 2+1+1$ twisted mass w/o clover and the *C* ensmbles are $N_f = 2$ twisted mass clover one of Chapter 3.

Label	a (fm)	$M_{\pi^\pm}$ (MeV)
P	0.095	250
A1	0.079	340
*A*1*	0.063	330
*A*2*	0.063	450
B1	0.089	280
*B*1*	0.062	220
*B*2*	0.062	310
C1	0.092	250
C2	0.092	340

quark content. Although r_0 has not been computed on the *P* ensembles, it would be possible to express the low energy constants (LECs) of Wilson chiral perturbation theory in units of w_0, similar to what is done in Ref. [124] in units of r_0. However, the LECs would carry an intrinsic systematic uncertainty of up to 40% when expressed in units of w_0, on account of w_0/a entering the expressions to the fourth power. For this reason, this kind of comparison was abandoned.

The reference ensembles in Table 6.3 have been selected such that for *A1* and *B1*, the lattice spacings and pion masses are as closely matched as possible to the *P* ensembles given the further condition that measurements of the charged, neutral connected and full neutral pion masses are available. The starred reference ensembles have been chosen to have significantly lower lattice spacings and two values of the pion mass, the lighter one of which should match the *P* ensembles as closely as possible. For these ensembles, the first two rows of Figure 6.5 confirm that the mass splitting is only very mildly dependent on the light quark mass. The uncertainties shown in quantities involving M_{π^0} should be taken as lower bounds on account of the difficulty of controlling the signal to noise ratio in the quark line disconnected contribution to the correlation function. The mass splittings scale with the lattice spacing squared and the deviation seen in $\left(M^2_{\pi^\pm} - M^2_{\pi^0}\right)$ for ensemble *A*2* is likely to be a statistical fluctuation on account of the fact that the signal to noise ratio in the neutral pion correlation function worsens as the quark mass is increased. It is interesting to see that for ensemble *C2*, the same increase is observed.

It should be kept in mind that the *P* ensembles have lattice spacings larger by 10% to 20% compared to the coarsest reference ensembles. The neutral connected to charged pion mass splitting $\left(M^2_{\pi^{(0,c)}} - M^2_{\pi^\pm}\right)$ is seen to be on the level of ensemble *A1*, while being slightly lower than for ensemble *B1*. Comparing the *P* ensembles to the $N_f = 2$ twisted mass clover *C* ensembles, the mass splitting is substantially larger. In addition, the difference between $N_f = 2$ and $N_f = 2+1+1$ twisted mass clover ensembles significantly exceeds the difference between the respective ensembles without the clover term. This might suggest that the clover term suppresses the mass splitting for $N_f = 2$, it does not do so efficiently for $N_f = 2+1+1$.

The picture is not quite so clear for $\left(M^2_{\pi^\pm} - M^2_{\pi^0}\right)$, which seems to be quite compatible with $N_f = 2$ twisted mass ensembles *A1*, *A*1*, $N_f = 2+1+1$ twisted mass ensembles *B*1*, *B*2* and the $N_f = 2$ twisted mass clover ensemble *C1*, although the lack of a variance reduction method for the *P* and *C* ensembles should be kept in mind. In the relative mass splittings, comparisons with *B1*, *B*1* and *C1* are the most useful since the pion masses are similar. Here, the value for *B1* indicates that both mass splittings are reduced due to the clover term. On the other hand, compared to *C1*, the neutral connected pion is heavier while the full neutral pion seems to have a compatible mass. Within the substantial uncertainties and in the tested range, the value of c_{sw} does not seem to play a role with respect to the pion mass splitting. It would be interesting to extend the range of c_{sw} values in this analysis to around and below 1.0.

Future studies including a variance reduction technique for the disconnected contribution to the neutral pion correlation function as well as multiple lattice spacings should give a clearer picture. Unlike in simulations without the clover term, a dependence on the strange quark mass of the mass splitting between the charged and full neutral pion is apparently not

	P205	*P185*	*P165*	*P145*
$\partial(a m_{PCAC})/\partial(1/2\kappa)$	3.04(9)	3.22(10)	3.08(20)	2.73(27)

Table 6.4: Slopes of $a m_{PCAC}$ vs. $1/2\kappa$ for $N_f = 2+1+1$ ensembles of approximately fixed physical situation with different values of c_{sw}.

visible within errors. This could be seen as an indication of reduced lattice artefacts from the heavy quark sector in light quark observables and it will be interesting to see if this is reproduced in scaling studies. It should be kept in mind, however, that the bare parameters μ_σ and μ_δ were not changed here. By contrast, the differences tabulated in Ref. [124] were due to changes in the strange and charm quark masses under variation of μ_δ or both μ_σ and μ_δ.

In conclusion it can be said that the pion mass splitting is reduced by the clover term also for $N_f = 2+1+1$ twisted mass fermions compared to simulations without it. At a charged pion mass of around 250 MeV, a weighted average of the four values gives

$$M_{\pi^\pm} - M_{\pi^0} \sim 73(18)\ \text{MeV}\,, \tag{6.2}$$

which is subject to possibly significant uncontrolled systematic uncertainties due to the poor signal for aM_{π^0}. As a result, this value and its error should be taken as indicative only, also because the estimates of the lattice spacings are almost certainly a little too high. By comparison, for the *B1* ($N_f = 2+1+1$) reference ensemble at a comparable charged pion mass, the splitting was

$$M_{\pi^\pm} - M_{\pi^0} \sim 139(10)\ \text{MeV}\,,$$

where the uncertainties and central value are under much better control but questions remain with regards to the effect of strange and charm masses on the pion mass splitting.

Given the results of this chapter, it is difficult to conclude whether the pion mass splitting is sufficiently low with the $N_f = 2+1+1$ twisted mass clover action. However, the measured value agrees surprisingly well with that measured in Section 6.1 for $N_f = 2$, perhaps indicating that the N_f-dependence is much less severe with the twisted mass clover action than with standard twisted mass fermions.

The PCAC Quark Mass as a Proxy for the Pion Mass Splitting

The results of the previous section suggest that within the range $1.45 \leqslant c_{sw} \leqslant 2.05$, the value of c_{sw} has no influence (within the large uncertainties) on the size of the pion mass splitting, contradicting evidence from the quenched approximation as described in Section 2.2.3. Because the signal for the neutral pion mass is extremely poor, it is helpful to extract as much information as possible from the available data. As discussed in Section 2.2.2, c_2 as well as the twisted quark mass enter the WtmχPT expression for the PCAC quark mass at leading order in the LCE regime. The bare Wilson quark mass am_0 can be related to the hopping parameter κ via $\kappa = (2am_0 + 8)^{-1}$. Measurements of am_{PCAC} as a function of $1/2\kappa$ of the four *P* ensembles presented in the previous section are shown in Figure 6.6. The slopes of the linear fits shown are listed in Table 6.4.

The data may be interpreted as confirming the result of the previous section that varying c_{sw} in the range $1.45 \leqslant c_{sw} \leqslant 2.05$ has no significant

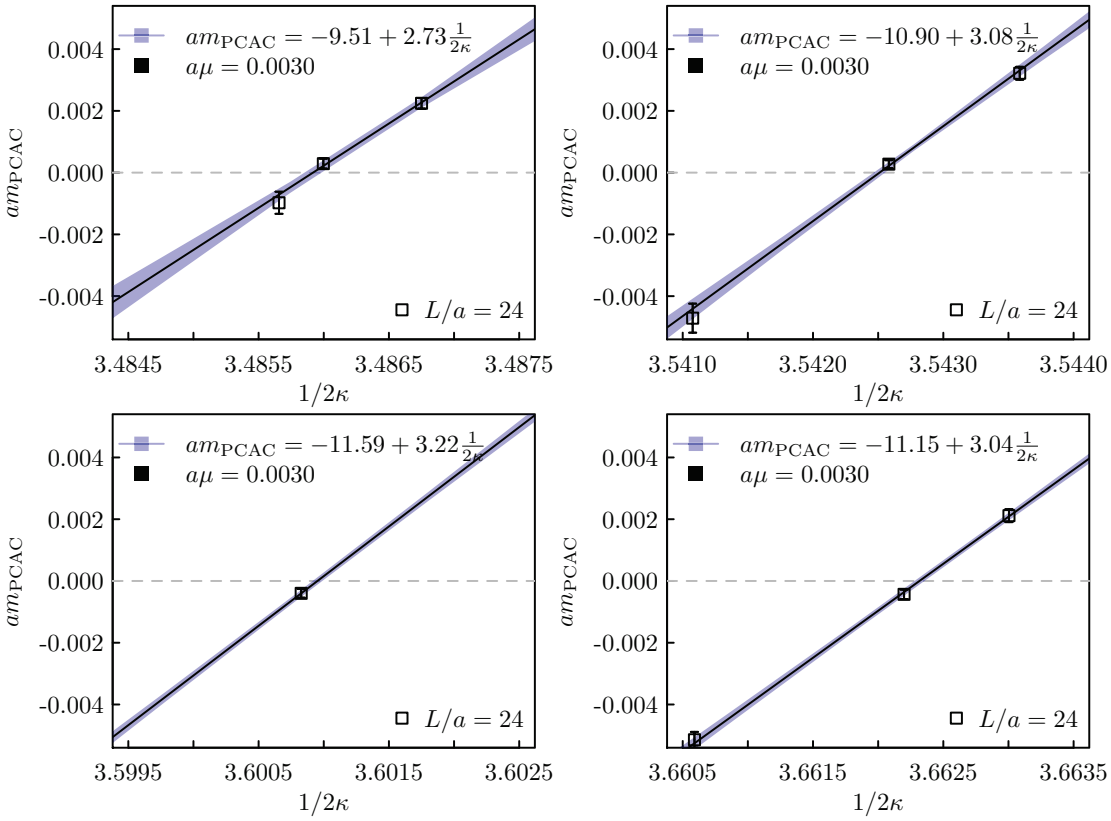

Figure 6.6: am_{PCAC} as a function of $1/2\kappa$ on $N_f = 2+1+1$ twisted mass clover ensembles with different values of c_{sw} but approximately matched physical situation. From top-left to bottom-right: $c_{sw} = 1.45, 1.65, 1.85, 2.05$.

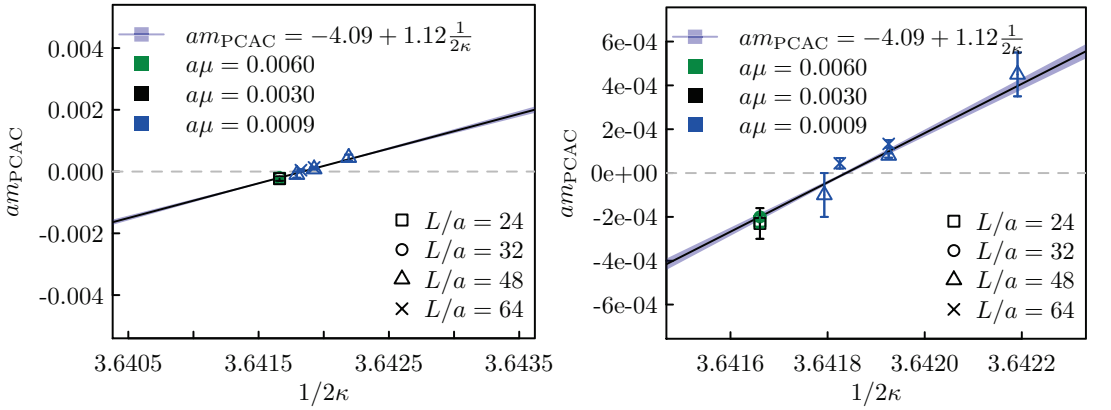

Figure 6.7: am_{PCAC} as a funtion of $1/2\kappa$ on $N_f = 2$ twisted mass clover production ensembles. The right panel is a zoom of the left panel.

effect on the value of c_2, assuming that Z_q/Z_A is largely independent of c_{sw}. Comparing the situation to the one with two flavours of twisted mass clover fermions, as shown in Figure 6.7, would suggest that c_2 is significantly larger in the $N_f = 2+1+1$ case, since the lattice spacing is comparable (a difference of at most 10% could be accounted to the lattice spacing). In fact, the slope for the $N_f = 2$ case is very close to 1, which may even indicate that $c_2 \sim 0$, although direct measurements of the pion mass splitting suggest otherwise. It should also be remembered that the LECs entering the expression should differ between $N_f = 2$ and $N_f = 2+1+1$, making any conclusions difficult.

By contrast, for $N_f = 2+1+1$ simulations with the old action at the coarsest lattice spacing and a light quark mass corresponding to a heavier charged pion mass than here (ensemble *A40*)[2], the slope was around 5. In addition, $a m_{PCAC}$ showed very strong signs of the μ-dependent shift discussed in Section 2.2.2. It may be interesting to carry out an analysis along the lines of Ref. [155] using the full WtmχPT formulae of Ref. [62] in an attempt to determine c_2 indirectly for $N_f = 2$ and $N_f = 2+1+1$ twisted mass clover ensembles. It is clear, however, that this type of analysis will require the computation of a number of renormalisation constants as well as simulations with at least one more lattice spacing and perhaps three different twisted quark masses. It would be particularly interesting to see the complete behaviour of $a m_{PCAC}$ for a wide range of untwisted quark masses for different values of μ.

[2] Talk by G. Herdoiza at GDR "Physique subatomique et calculs sur réseau", Marseille, June 25-27 2008, `http://gdr-lqcd.in2p3.fr/reunion08/talks/Herdoiza.pdf`

7

PSEUDOSCALAR MESON MASSES AND DECAY CONSTANTS

The $N_f = 2$ ensembles presented in Chapter 3 were used to compute the masses and decay constants of pseudoscalar mesons with light, strange and charm quarks in a partially quenched lattice setup using Osterwalder-Seiler valence quarks. These computations are an important proof of concept for calculations at the physical pion mass and will be discussed in the present chapter. In Section 7.1, the theoretical background for the lattice computation will be presented, the results will be shown in Section 7.2 and a summary will be given in Section 7.3. The details of the analysis, especially the quantification of uncertainties, will be given at the end of the chapter in Section 7.4. A significant portion of this chapter is based on Ref. [P6].

The decay constant f_P of a pseudoscalar meson P parametrises its leptonic decays and, to lowest order, the decay width is given by [156]

$$\Gamma(P \to \ell\nu) = \frac{G_F^2}{8\pi} f_P^2 m_\ell^2 M_P \left(1 - \frac{m_\ell^2}{M_P^2}\right)^2 \left|V_{q_1 q_2}\right|^2 , \tag{7.1}$$

where G_F is the Fermi constant, m_ℓ is the lepton mass and M_P is the mass of the pseudoscalar meson. $\left|V_{q_1 q_2}\right|$ is an element of the Cabibbo-Kobayashi-Maskawa (CKM) matrix [157, 158] between the flavours of the constituent quarks of P. As a result, the decay constant computed in LQCD and knowledge of the decay width from experiment allows the determination of $\left|V_{q_1 q_2}\right|$. This is especially important for the elements between heavier quarks, where *isolated* experimental determinations are poor (e.g. f_{D_s}) or not possible (e.g. f_{B_s} [159]).

The CKM matrix parametrises the strength of flavour-changing weak decays between quarks and thus constitutes an important ingredient in the quantification of charge-parity (CP) violation in the Standard Model. Besides its innate importance, an accurate understanding of CP violation is required to explain the matter-antimatter imbalance in light of the fact that the universe appears to be matter-dominated.

A further important quantity which can be obtained from lattice data is the pion decay constant in the chiral limit, which is an essential low energy constant (LEC) of chiral perturbation theory. Simulations directly at the physical charged pion mass should prove useful to strongly constrain the uncertainties of the requisite chiral extrapolation to vanishing quark mass.

The masses of the pseudoscalar mesons are efficient hadronic quantities for the tuning of bare quark masses in the lattice theory. As will be shown below, ratios of the masses of the kaon and D meson over the physical pion mass in the isospin-symmetric limit of pure QCD were used to tune the strange and charm valence quark masses in the present analysis. The resulting light, strange and charm quark masses and their ratios are seen to match their phenomenological values rather well.

The ratio M_π/f_π can further be used to set the lattice scale by extrapolating to the point where it takes its physical value[1]. The experimentally measured value of f_π can the be used to set the scale at this point through

$$a \sim \frac{a f_\pi^{\text{lat}}(\mu_{\text{phys}})}{f_\pi^{\text{exp}}}\,, \tag{7.2}$$

where care has to be taken in regards to lattice artefacts, especially finite volume effects. In the present chapter finite volume corrections are largely omitted because the small number of ensembles makes it difficult to estimate the correction factors. Further, it is interesting to see what kind of results can be achieved from the raw lattice data.

LATTICE COMPUTATION

The extraction of pseudoscalar meson masses and decay constants in the twisted mass formulation employs interpolating pseudoscalar operators based on the quark currents of Section 1.3.1 defined in the twisted basis. The pseudoscalar interpolating operator for flavours $f, f' \in \{\ell, s, c\}$ is defined as

$$P^{\pm}_{f,f'}(t) \;=\; \sum_{\mathbf{x}} \bar{\chi}_f(\mathbf{x},t)\, i\gamma_5\, \tau^{\pm}\, \chi_{f'}(\mathbf{x},t)\,, \qquad \tau^{\pm} = \frac{\tau^1 \pm i\tau^2}{2}\,, \tag{7.3}$$

where the χ fields are two-flavour spinors. The pseudoscalar correlation function then follows as

$$C_{PS}^{f,f'}(t) \;=\; \langle\, P^{\pm}_{f,f'}(t)\; P^{\pm}_{f,f'}(0)^{\dagger}\,\rangle\,, \tag{7.4}$$

which projects onto the charged pion states in the light sector as the twisted mass parameters of the flavours f and f' enter with opposite signs. For the kaon and D meson, the same combination is used because it can be shown that in the case of opposite signs, leading discretisation artefacts are reduced compared to the same sign correlation function. The details of the lattice approximation used to compute Equation (7.3) and Equation (7.4) are given in Section 7.4.

The pseudoscalar meson mass is extracted from $C_{PS}^{f,f'}(t)$ through a ground state fit involving forward and backward propagating states, as suggested by the spectral decomposition in the limit $t \to \infty$

$$\lim_{t\to\infty} C_{PS}^{f,f'}(t) = \frac{|\langle 0|P^{\pm}_{f,f'}(t)|PS\rangle|^2}{2M_P^{f,f'}}\,\left(e^{-M_P^{f,f'}t} + e^{-M_P^{f,f'}(T-t)}\right), \tag{7.5}$$

where T is the lattice extent in the temporal direction. Hence, $M_P^{\ell,\ell}$, $M_P^{\ell,s}$ and $M_P^{\ell,c}$ correspond to the charged pion, the kaon and the D meson masses, respectively. In practice, the correlation function is fitted in a region where excited states have decayed sufficiently for the ground state to dominate while the signal to noise ratio of the lattice correlation function, which worsens with increasing t, is still good. As described in Section 7.4, for the results presented here, a very careful analysis of the systematic error due to the fit range dependence was carried out. However, a first estimate of the appropriate fit range can be obtained by observing a plateau in the effective mass

$$M_{\text{eff}}(t) = -\log\left(\frac{C(t)}{C(t-1)}\right)\,, \tag{7.6}$$

[1] or simulating directly at this point, as done for ensemble *cA2.09.48*

given here for general correlation functions $C(t)$. This can also be utilised to determine hadron masses through a constant fit in a region where it assumes a plateau. In practice, Equation (7.6) is not directly useful for the pion at its physical mass because the backward-propagating contribution is too large. Instead of taking the logarithm, then, all the effective masses shown in this chapter are determined from the relation

$$\frac{C(t)}{C(t-1)} = \frac{\cosh\left[-(t)\cdot M_{\mathrm{eff}}(t)\right]}{\cosh\left[-(t-1)\cdot M_{\mathrm{eff}}(t)\right]}\,, \tag{7.7}$$

which is solved numerically for $M_{\mathrm{eff}}(t)$ at each t.

The determination of pseudoscalar decay constants in tmLQCD at maximal twist can be approached as in a theory with exact chiral symmetry: f_P is directly related to the matrix element $\langle 0|P^{\pm}_{f,f'}(t)|PS\rangle$ via

$$f_P^{f,f'} \;=\; (\mu_f + \mu_{f'})\frac{\langle 0|P^{\pm}_{f,f'}(t)|PS\rangle}{(M_P^{f,f'})^2} + \mathcal{O}(a^2)\,, \tag{7.8}$$

which follows from the PCVC relation in tmLQCD at maximal twist [53]. The lattice dispersion relation for mesons can be taken into account by exchanging $(M_P^{f,f'})^2$ in Eq. 7.8 for $M_P^{f,f'}\sinh M_P^{f,f'}$. In the following, the former will be referred to as "Continuum Definition" (CD) and the latter as "Lattice Definition" (LD). The two definitions are clearly equivalent in the continuum limit ($\sinh aM_P \to aM_P$ as $a \to 0$), but the latter is expected to somewhat reduce scaling violations, especially for heavy mesons. Since the continuum limit will not be taken in the present analysis, the two definitions will be used to interpret the possible consequences of the size of lattice artefacts in the decay constants. It should be noted though that there is no argument suggesting that these are indicative of the true lattice artefacts.

In order to determine the ground state masses more reliably, 2×2 matrices of correlation functions are built using local and fuzzed [160] interpolators. By comparing the effective masses of the three types of correlation functions, the decay of excited states can be better estimated. The masses are then extracted by solving the corresponding GEVP [161, 19, 162] or using a constrained matrix fit, the details of which are given in Ref. [134].

ANALYSIS RESULTS

In the following, the results of the determination of meson masses and decay constants as well as their ratios will be presented. This section will be limited mostly to determinations on ensemble *cA2.09.48* at the physical charged pion mass. Although all quantities have also been computed on the ensembles with larger than physical pion mass, they are only listed and the light quark mass dependence is not studied explicitly. The closeness to the physical pion mass of ensemble *cA2.09.48* will first be discussed, followed by a presentation of scale-independent ratios of pseudoscalar quantities in Section 7.2.1. Dimensionful quantities will be determined in Section 7.2.2, using the value of the lattice spacing given in Ref. [P6]. Finally, estimates of light, strange and charm quark masses in physical units will be given in Section 7.2.3 using the bare quark masses determined in this analysis.

As an example for the mass determinations, Figure 7.1 shows the effective mass of the nucleon as determined in Ref. [P7], the pion and the kaon, the latter with $a\mu_s = 0.0245$. In this figure, the nucleon effective mass has been computed from a correlation function smeared at both source and sink to

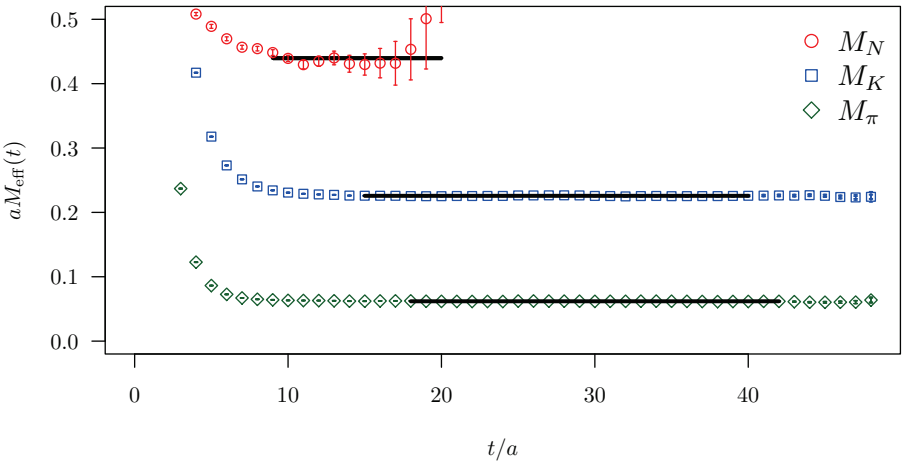

Figure 7.1: Effective masses as a function of t/a for nucleon, kaon and pion for the physical ensemble *cA2.09.48*.

improve the signal, while the meson effective masses shown here are from local-local correlation functions only.

For the nucleon, the expected exponential error growth is observed and effective masses for $t/a > 21$ are thus not used in the analysis. The fit range as well as the fit result and its statistical error are indicated by the solid line and shaded region. For the pion and kaon, the plateau is visible up to $t = T/2$ as expected. As already mentioned above, the final values for the pseudoscalar masses are the result of a weighted average over many fit ranges, as described in Section 7.4.1. The indicated fit range is thus only representative. It should be noted that the errors are too small to be visible on this scale, but details are shown in Figure 7.4 on page 109.

The Physical Pion Mass and Scale-Independent Quantities

The first set of quantities to be discussed here concerns explicitly scale-independent quantities such as mass ratios, ratios of masses and decay constants and ratios of decay constants. Working at the physical pion mass and with appropriate tuning conditions for the strange and charm quark masses, these ratios can be compared directly to their phenomenological determinations and should give important indications of residual lattice artefacts. The ratio f_K/f_π can be used to compute the ratio $|V_{us}|/|V_{ud}|$ and the precisely measured experimental value of $|V_{ud}|$ can in turn be used to compute $|V_{us}|$. Since neither continuum nor infinite volume limits are taken here, the determination of these derived quantities is omitted, however.

The left panel of Figure 7.2 shows the ratio $r_0 M^2_{\pi^\pm}/f_{\pi^\pm}$ as a function of $(r_0 M^2_{\pi^\pm})$, comparing measurements on the $N_f = 2$ ensembles of Chapter 3 to published values from Ref. [96] of $N_f = 2$ twisted mass data without a clover term. In this particular ratio, a cancellation of lattice artefacts seems to take place compared to alternative choices. The data is presented with only Gasser-Leutwyler finite size corrections [163] applied and the experimental value is indicated by the green star with $r_0 = 0.4907(86)$ fm as determined in Ref. [P6]. The dashed line shows a combined fit of the continuum NLO χPT expression [164, 165]

$$\frac{M^2_{\pi^\pm}}{f_{\pi^\pm}} = \frac{M^2_{\pi^\pm}}{f_0}\left(1 + 2\frac{M^2_{\pi^\pm}}{(4\pi f_0)^2}\log\frac{M^2_{\pi^\pm}}{\Lambda_4^2}\right), \tag{7.9}$$

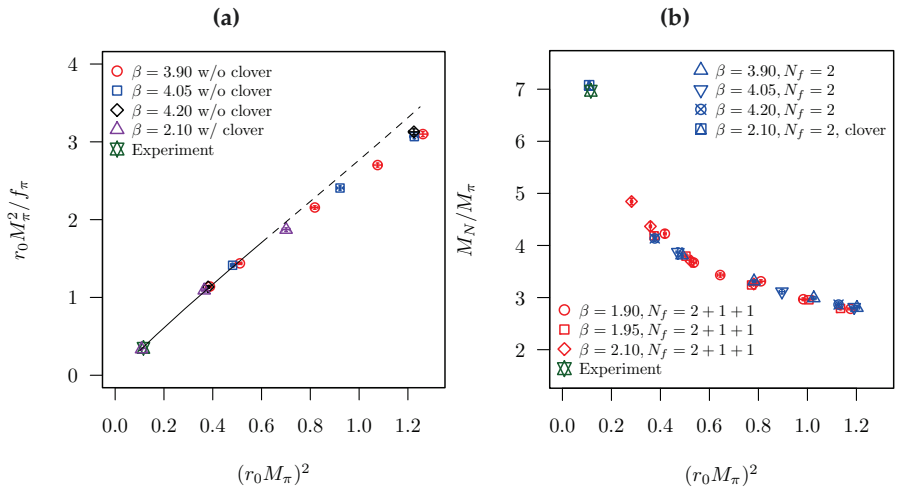

Figure 7.2: (a) $r_0 M_\pi^2 / f_\pi$ as a function of $(r_0 M_\pi)^2$ comparing $N_f = 2$ results w/o clover term [96] with the results presented in this chapter. The solid ($M_{\pi^\pm} < 300$ MeV) and dashed lines are NLO χPT fits of Equation (7.9) to the data as explained in the text. **(b)** Ratio of the nucleon mass to the pion mass as a function of the pion mass squared in units of r_0. Data for $N_f = 2$ and $N_f = 2+1+1$ w/o clover term and the result for physical pion mass ensemble *cA2.09.48* is shown.

	cA2.09.48	*cA2.30.24*	*cA2.60.24*	*cA2.60.32*
$M_{\pi^\pm}/f_{\pi^\pm}^{(CD)}$	$1.0254(31)(^{+26}_{-12})$	$1.879(22)(^{+08}_{-17})$	$2.30(11)(^{+02}_{-03})$	$2.2395(76)(^{+39}_{-24})$
$M_{\pi^\pm}/f_{\pi^\pm}^{(LD)}$	$1.0260(31)(^{+26}_{-12})$	$1.884(22)(^{+09}_{-17})$	$2.31(11)(^{+02}_{-03})$	$2.2489(77)(^{+38}_{-23})$

Table 7.1: $M_{\pi^\pm}/f_{\pi^\pm}$ determined on different $N_f = 2$ twisted mass clover ensembles without any finite size corrections.

as a function of $M_{\pi^\pm}^2$ to all the lattice data in units of r_0, neglecting lattice artefacts. The solid line indicates a fit restricted to $M_{\pi^\pm} < 300$ MeV, giving $f_0 = 0.122(4)$ GeV and $\bar{\ell}_4 = 3.3(4)$. It should be noted that these fit results are completely in line with those of Ref. [166]. The p-value of this fit is 0.49 and the inclusion of the chiral logarithm is favoured over the linear fit. Data for $M_{\pi^\pm}/f_{\pi^\pm}$ is given in Table 7.1 for the four different $N_f = 2$ twisted mass clover ensembles used in this analysis. If this is fitted instead of the ratio used above, the results do not change but the p-value worsens significantly, suggesting residual finite volume and lattice artefacts. Although the statistical uncertainty is unfortunately large, a comparison of the values on ensembles *cA2.60.24* and *cA2.60.32* gives an indication of the finite volume effects on this ratio. On the other hand, the individual measurements given in Table 7.4 indicate a much smaller effect.

As discussed in Chapter 3 and Chapter 6, there are good indications that the clover term reduces the pion mass splitting. As a result, it can be expected that finite size effects (FSE) are smaller than for ensembles without the clover term. In addition, at the physical pion mass, the FSE are further suppressed by a factor of M_π^2. Hence, corrections which take into account the pion mass splitting, as discussed in Refs. [127, 112], are likely to give only small contributions. Still, once more data becomes available, these kinds of fits will have to be repeated to quantify the additional FSE from the neutral pion.

The mass ratio of the nucleon to the pion for $N_f = 2$ and $N_f = 2+1+1$ ETMC ensembles is shown in the right panel of Figure 7.2 as a function of the pion mass squared in units of r_0. The nucleon mass on ensemble *cA2.09.48* at the physical pion mass has been measured on 96 independent configurations with 16 sources per configuration as discussed in Ref. [P6]. The masses for the $N_f = 2$ ensembles without clover term have been taken from Ref. [167] while those for the $N_f = 2+1+1$ ensembles are given in Ref. [99]. The values for r_0/a were taken from Ref. [128] for $N_f = 2+1+1$ and from Ref. [168] for $N_f = 2$ without the clover term. As can be seen, the lattice results follow a universal curve, suggesting that cut-off effects are

Table 7.2: Bare quark masses in lattice units and their ratios as determined by matching M_K/M_π and M_D/M_π to their phenomenological values. For μ_c/μ_s, the asymmetric error corresponds to the maximum spreads of the dividend and divisor while for the other ratios it derives directly from M_K/M_π and M_D/M_π.

$a\mu_l$	$a\mu_s$	$a\mu_c$
0.0009	$0.02485(7)(^{+4}_{-3})$	$0.3075(15)(^{+14}_{-14})$
μ_s/μ_l	μ_c/μ_l	μ_c/μ_s
$27.61(8)(^{+4}_{-4})$	$342.1(1.8)(^{+1.6}_{-1.6})$	$12.39(8)(^{+6}_{-9})$

small in this ratio. Moreover, differences between $N_f = 2$ and $N_f = 2+1+1$ are smaller than the statistical uncertainties.

For quantities involving strange and charm quarks the valence quark mass needs to be tuned. This was done by matching the phenomenological values of the pseudoscalar meson mass ratios M_K/M_π and M_D/M_π through linear interpolations of the lattice data, resulting in the bare quark masses and their ratios given in Table 7.2. The details of this procedure are discussed in Section 7.4.2 and bare strange and charm quark masses using different quantities for tuning are provided. M_K/M_π and M_D/M_π were chosen because they seem to produce the smallest uncertainties with the systematic errors currently considered. In addition, this choice retains predictivity for f_K, f_D and f_{D_s} as well as their ratios.

Ratios of masses and decay constants resulting from this analysis are given in Table 7.3. It is clear from Figure 7.2 and the value of M_π/f_π that the ensemble *cA2.09.48* is at the physical pion mass within errors. For the other quantities, agreement with phenomenological determinations and continuum limit lattice averages is seen to be quite good. As discussed above in Section 7.1, a first estimate of the residual $\mathcal{O}(a^2)$ artefacts can be obtained by comparing the two definitions (see Equation (7.8)) of the decay constant in quantities involving f_D and f_{D_s}. It seems that these effects should be no larger than about 15%, indicating that a well-behaved continuum limit is certainly achievable once more ensembles become available. Finally, the present results are seen to be in full agreement with previous ETMC determinations from Ref. [169] for $N_f = 2$ twisted mass fermions in the infinite volume and continuum limit which gave $f_K/f_\pi = 1.210(18)$ and $f_{Ds}/f_D = 1.24(3)$. A summary of the comparison to phenomenological determinations is given in Figure 7.3 on page 106.

Dimensionful Quantities

All the dimensionful quantities from this analysis are listed in Table 7.4 in lattice units. Results from ensembles with larger than physical charged pion mass were computed in the same way as on ensemble *cA2.09.48*. In particular, observables involving strange and charm quarks were interpolated to the strange and charm quark masses given in Table 7.2. This approach may be an alternative to the one usually employed in partially quenched computations at larger than physical pion masses. In these, for example, sums of hadronic masses are held constant as the physical pion mass is approached whereas here it would be the strange and charm quark masses which are held fixed at what are likely to be good estimates of their physical values.

In Ref. [P6], the lattice spacing for $\beta = 2.1$ is determined from weighted, correlated averages with χ^2/df-stretched statistical errors following the

	CD	LD	PDG [156]	FLAG [170]
M_π/f_π	$1.0254(31)(^{+26}_{-12})^\dagger$	$1.0262(30)(^{+33}_{-18})^\dagger$	$1.0337(28)^\star$	$1.035(11)^\star$
M_K/f_K	$3.1404(55)(^{+13}_{-11})^\dagger$	$3.1675(56)(^{+13}_{-11})^\dagger$	$3.164(14)^\star$	$3.162(18)^\star$
M_D/f_D	$8.395(64)(^{+41}_{-16})^\dagger$	$9.466(71)(^{+41}_{-17})^\dagger$	9.11(22)	–
M_{D_s}/f_{D_s}	$7.474(21)(^{+03}_{-03})$	$8.531(28)(^{+04}_{-03})$	7.64(14)	–
M_{D_s}/M_π	$14.564(54)(^{+03}_{-03})^\dagger$	–	$14.603(33)^\star$	–
f_K/f_π	$1.1976(21)(^{+06}_{-07})$	$1.1881(21)(^{+06}_{-07})$	1.1979(57)	1.200(15)
f_D/f_π	$1.694(14)(^{+04}_{-10})$	$1.503(12)(^{+04}_{-07})$	1.569(38)	1.61(3)
f_{D_s}/f_π	$1.998(6)(^{+1}_{-1})$	$1.751(5)(^{+1}_{-1})$	1.975(35)	1.91(3)
f_D/f_K	$1.413(12)(^{+02}_{-03})$	$1.264(10)(^{+02}_{-02})$	1.309(33)	1.34(2)
f_{D_s}/f_D	$1.206(23)(^{+04}_{-04})$	$1.190(22)(^{+04}_{-04})$	1.258(38)	1.19(2)

Table 7.3: Ratios of pseudoscalar meson observables calculated on ensemble *cA2.09.48* at the physical pion mass, interpolated to the strange and charm valence quark masses from the matching procedure described in Section 7.4.2. The values for the decay constants are computed taking into account continuum (CD) and lattice (LD) dispersion relations for mesons (see Equation (7.8)). All starred reference ratios involving M_π or M_K use the isospin symmetric values of these quantities taken from Ref. [170]. Daggered quantities are not independent and given for reference only. "FLAG" refers to the $N_f = 2+1$ averages presented therein.

	cA2.09.48	*cA2.30.24*	*cA2.60.24*	*cA2.60.32*
$aM_{\pi^\pm}$	$0.06196(09)(^{+12}_{-05})$	$0.1147(7)(^{+4}_{-7})$	$0.15941(38)(^{+15}_{-21})$	$0.15769(26)(^{+15}_{-14})$
$aM_{\pi^{(0,c)}}$	$0.1191(05)(^{+07}_{-10})$	$0.1541(13)(^{+05}_{-05})$	$0.18981(61)(^{+21}_{-25})$	$0.18840(44)(^{+46}_{-29})$
aM_{π^0}	$-^\dagger$	$0.09(1)^\star$	$0.11(2)^\star$	$0.13(2)^\star$
$af^{(CD)}_{\pi^\pm}$	$0.06042(11)(^{+07}_{-03})$	$0.06104(43)(^{+15}_{-14})$	$0.06946(22)(^{+03}_{-05})$	$0.07043(19)(^{+06}_{-05})$
$af^{(LD)}_{\pi^\pm}$	$0.06038(11)(^{+07}_{-03})$	$0.06090(43)(^{+16}_{-14})$	$0.06917(22)(^{+03}_{-05})$	$0.07013(19)(^{+06}_{-05})$
$af^{(CD)}_K$	$0.07235(9)(^{+2}_{-2})$	$0.07265(31)(^{+06}_{-06})$	$0.07774(19)(^{+07}_{-07})$	$0.07816(16)(^{+09}_{-07})$
$af^{(LD}_K)$	$0.07173(9)(^{+2}_{-2})$	$0.07197(32)(^{+06}_{-05})$	$0.07692(19)(^{+07}_{-06})$	$0.07734(16)(^{+09}_{-07})$
$af^{(CD)}_D$	$0.1022(9)(^{+3}_{-7})$	$0.1087(14)(^{+09}_{-13})$	$0.1127(7)(^{+5}_{-7})$	$0.1110(10)(^{+05}_{-06})$
$af^{(LD)}_D$	$0.0906(8)(^{+2}_{-6})$	$0.0960(12)(^{+07}_{-11})$	$0.0994(6)(^{+4}_{-6})$	$0.0980(8)(^{+4}_{-5})$
$af^{(CD)}_{D_s}$	$0.1207(2)(^{+1}_{-1})$	$0.1220(7)(^{+1}_{-1})$	$0.1237(5)(^{+1}_{-2})$	$0.1219(5)(^{+1}_{-1})$
$af^{(LD)}_{D_s}$	$0.1058(2)(^{+1}_{-1})$	$0.1068(5)(^{+1}_{-1})$	$0.1082(4)(^{+1}_{-1})$	$0.1067(5)(^{+1}_{-1})$
aM_{D_s}	$0.9022(27)(^{+06}_{-07})$	$0.905(3)(^{+1}_{-1})$	$0.9062(27)(^{+02}_{-02})$	$0.9034(26)(^{+01}_{-01})$
N_{meas}	1457	728	1351	670

Table 7.4: Dimensionful pseudoscalar quantities in lattice units determined on the four $N_f = 2$ twisted mass clover ensembles in this analysis. N_{meas} refers to the number of configurations used and the uncertainties have been determined as discussed in Section 7.4.

Table 7.5: Reference values for the dimensional quantities calcuated in this study. FLAG refers to $N_f = 2+1$ averages from Ref. [170], ETM'09 refers to Ref. [169] ($N_f = 2$) and ETM'15 to Ref. [173] ($N_f = 2+1+1$).

	f_K [MeV]	f_D [MeV]	f_{D_s} [MeV]	M_{D_s} [MeV]
ETM'09	158.1(2.4)	197(9)	244(8)	–
ETM'15	155.0(1.9)	207.4(3.8)	247.2(4.1)	–
PDG	156.2(7)	204.6(5.0)	257.5(4.6)	1968.50(32)
FLAG	156.3(9)	209.2(3.3)	248.6(2.7)	–

procedure of Ref. [171]. The lattice spacing is determined from the gluonic scales t_0 [135], w_0 [142] and r_0 [172], giving

$$a_{\text{gluonic}} = 0.0931(10)(15)\,\text{fm}\,. \tag{7.10}$$

Similarly, the phenomenological values of M_N, $M_{\pi^\pm}$, f_π and f_K are used to give a *hadronic* estimate of the lattice spacing

$$a_{\text{hadronic}} = 0.0913(2)(11)\,\text{fm}\,. \tag{7.11}$$

Finally, an uncertainty-weighted average is calculated, giving

$$a = 0.0914(3)(15)\,\text{fm}\,, \tag{7.12}$$

where the combined relative uncertainty is around 1.7%. In this way it is hoped that the effect of different lattice artefacts is captured while not overestimating the related systematic error.

In order to provide values in physical units for the independent dimensionful quantities in this analysis, a_{gluonic} will be used which carries a slightly larger uncertainty of around 2%. The resulting values for f_K, f_D, f_{D_s} and M_{D_s} are

$$\begin{aligned}
f_K^{(CD)} &= 153.35(0.18)(^{+0.04}_{-0.04})(2.96)\ \text{MeV}\,,\\
f_D^{(CD)} &= 216.71(1.99)(^{+0.59}_{-1.47})(4.19)\ \text{MeV}\,,\\
f_{D_s}^{(CD)} &= 255.85(0.49)(^{+0.10}_{-0.14})(4.95)\ \text{MeV}\,,\\
M_{D_s} &= 1912.3(5.73)(^{+0.13}_{-0.15})(37.0)\ \text{MeV}\,,
\end{aligned}$$

where the first uncertainty is statistical, the second from the fit range ambiguity and the third from the estimate of the lattice spacing. These results can be compared to various lattice and phenomenological world averages as well as ETMC determinations listed in Table 7.5. It is interesting that within the quoted uncertainties, despite the lack of infinite volume and continuum limits, values largely compatible with those in the literature are obtained. When the scale ambiguity is not considered, however, it becomes clear that at least in the charm sector, sizeable lattice artefacts are likely to be present. A summary of the comparison to phenomenological determinations is given in Figure 7.3 on page 106.

Quark Masses

At maximal twist the Lagrangian twisted quark mass can be directly related to the physical renormalised quark mass in some renormalisation scheme through multiplicative renormalisation. The relevant renormalisation factor in twisted mass lattice QCD is Z_P, since $m_q^R = Z_P^{-1}\mu_\ell$. Z_P is determined in Refs. [174, P7] using the RI'-MOM scheme employing the momentum source technique [175]. The value of Z_P at 2 GeV in the $\overline{\text{MS}}$ scheme reads

$$Z_P = 0.501(8)(26)(12)\,, \tag{7.13}$$

where the first error is statistical and the second is a systematic error due to the extrapolation towards $(ap)^2 = 0$, also including the perturbative subtraction of leading lattice artefacts [176, 177]. The third error stems from the conversion of RI' to $\overline{\text{MS}}$ at 2 GeV.

As both $M_{\pi^\pm}/f_{\pi^\pm}$ and $M_{\pi^\pm}/M_N$ indicate that ensemble *cA2.09.48* with $a\mu_\ell = 0.0009$ is at or very close to the physical average up/down quark mass, this can be used to give an estimate in physical units without further extrapolation. With the Z_P-value given above and the estimate of the lattice spacing from Equation (7.12), the average up/down quark mass takes the value

$$m_{ud}^{\overline{\text{MS}}}(2\text{ GeV}) = 3.88(6)(21)(10)\text{ MeV}, \tag{7.14}$$

where the first error is statistical, the second from the combined systematic errors of the lattice scale and Z_P, and the third from the conversion of RI'-MOM to $\overline{\text{MS}}$ at 2 GeV.

From the quark mass ratios μ_s/μ_ℓ and μ_c/μ_ℓ given in Table 7.2, one can then further compute estimates for the strange and charm quark masses

$$\begin{aligned} m_s^{\overline{\text{MS}}}(2\text{ GeV}) &= 107(2)(6)(3)\text{ MeV}, \\ m_c^{\overline{\text{MS}}}(2\text{ GeV}) &= 1.33(3)(7)(3)\text{ GeV}. \end{aligned} \tag{7.15}$$

Both, m_{ud} and m_s compare well to the quark mass values determined on the $N_f = 2$ ETMC ensembles without clover term [178]. The estimates of m_{ud} and m_s can also be compared to the $N_f = 2$ FLAG [170] averages of Refs. [168, 179, 180, 87, 86], which read

$$m_{ud} = 3.6(2)\text{ MeV}, \qquad m_s = 101(3)\text{ MeV},$$

giving agreement in both cases. Of course, the values above and especially their uncertainties should be considered with care since neither a continuum nor an infinite volume limit were performed. The dependence of the quark masses and their ratios on the matching conditions is explored in explored in Table 7.7 in Section 7.4. Clearly, using different ratios of observables to tune the strange and charm quark masses significantly affects their central values and uncertainties.

SUMMARY AND CONCLUSIONS

In this chapter the computation of observables in the pseudoscalar meson sector on twisted mass clover ensemble *cA2.09.48* was presented. The ratios $M_{\pi^\pm}/M_N$ and $M_{\pi^\pm}/f_{\pi^\pm}$ indicate this ensemble is at or very close to the physical charged pion mass. From the technical point of view, a careful analysis of the systematic error due to the fit range ambiguity was introduced, as discussed in detail in Section 7.4.1. It was seen that for some quantities, most notably the pion mass and decay constant, the resulting systematic error is on the level of the statistical uncertainty and should certainly be taken into account.

A summary of the calculated quantities is given in Figure 7.3, where the lattice result is divided by the phenomenological value of the given quantity. The contribution to the total uncertainty from the uncertainty of the phenomenological quantity is indicated separately by the red shaded band, thus allowing lattice and phenomenological uncertainties to be compared. Given that no continuum and infinite volume limits were taken, the observed deviations are modest and may suggest that the scaling properties of the

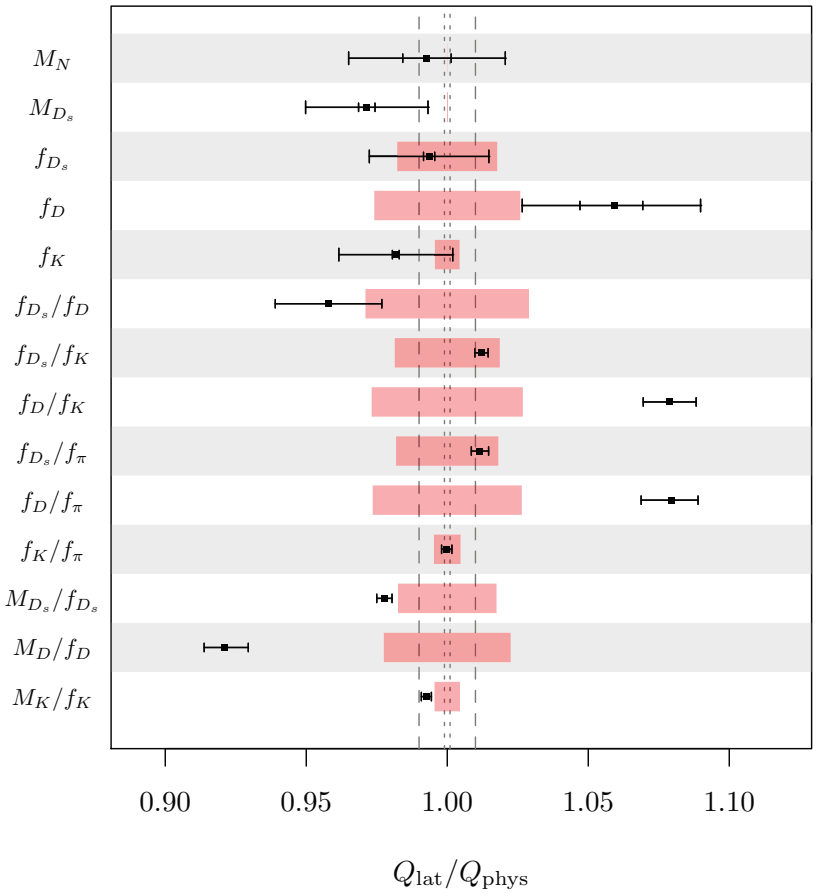

Figure 7.3: Ratios of lattice results and phenomenological values of the quantities in the legend with lattice decay constants computed via the continuum definition. For dimensionful quantities, the inner error bar combines the statistical and systematic errors in quadrature while the outer error bar stems from the estimate of the lattice spacing from gluonic scales (Equation (7.10)). The red bands show the contribution of Q_{phys} to the total uncertainty on $Q_{\mathrm{lat}}/Q_{\mathrm{phys}}$ separately (the respective experimental errors on M_N and M_{D_s} are too small to be visible). The dotted and dashed lines indicate per-mille and per-cent deviations from 1.0 respectively.

twisted mass clover action are quite favourable. Of course, this will need to be confirmed by explicit simulations on larger lattices and at finer lattice spacings. Deviations of up to 10% in observables involving charm quarks indicate that at this coarse lattice spacing, cut-off artefacts can indeed be substantial, although the deviations may also be related to the tuning procedure for the strange and charm valence quark masses.

It was shown that strange and charm valence quark masses can be tuned by matching ratios of appropriate quantities to their phenomenological values through simple linear interpolations. In this matching procedure, the presence of sizeable lattice artefacts in the charm sector should be carefully considered and Table 7.7 shows that using different quantities can have a significant impact on the value of the tuned quark masses as well as their uncertainties. If quantities involving decay constants are used as matching conditions, the lattice and continuum definitions of these may result in significantly different quark masses. These deviations could be quantified as an additional systematic error. By using multiple matching conditions in parallel, better estimates of the quark masses may be arrived at but doing so, one would of course sacrifice predictivity for a number of quantities. With the choice of $M_K/M_{\pi^\pm}$ and $M_D/M_{\pi^\pm}$ as matching conditions, quark mass ratios were computed and it was shown that these largely agree with the available literature. It should also be noted that the uncertainties on the ratios are quite small and are likely to increase only moderately with continuum and infinite volume limits. The values of the quark masses, converted to the $\overline{\mathrm{MS}}$ scheme via appropriate renormalisation constants and matching factors, are a little high however.

In dimensionful quantities, the uncertainty is dominated by the ambiguity of the lattice scale. Quantities in the charm sector, especially f_D also come with significant statistical error due to poor plateau quality and systematic error due to the resulting fit range ambiguity. Thus, for many quantities, reaching per-mille level uncertainties is likely impossible in this kind of *naïve* analyis and procedures involving (heavy meson) WχPT may be necessary, similar to the ones used in Refs. [173] and [181]. By combining data at larger than physical pion mass in this way, it may be possible to significantly reduce the eventual uncertainties on final results. However, as discussed in Section 2.2, the chiral extrapolation of lattice data at fixed lattice spacing may be complicated by the fact that one may be working in different power counting regimes, such that lattice artefacts cannot be smoothly connected between lattice data at different light quark masses. In addtion, this kind of analysis will certainly have to be done on $N_f = 2+1+1$ ensembles such that the effect of additional dynamical flavours, especially the strange quark, is taken into account correctly.

In any case, the results of this chapter serve as a proof of concept that calculations in the pseudoscalar meson sector at the physical charged pion mass are possible using twisted mass quarks. The apparent absence of any surprising strong lattice artefacts suggests that this kind of analysis can be extended to larger lattice volumes and a continuum limit performed. It is likely that for this purpose a more involved analysis procedure must be used, perhaps through the introduction of intermediate scales and the explicit consideration of lattice artefacts and finite volume corrections. From the point of view of this analysis, no difficulties are foreseen in the extension to $N_f = 2+1+1$ ensembles.

ANALYSIS DETAILS

In the final section of this chapter, details of the analysis methods used for the computation of the results of Section 7.2 will be provided. First, a somewhat novel technique for quantifying the uncertainty due to the choice of fit range for correlation functions is introduced. This is followed by a discussion of the tuning of valence strange and charm quark masses and the effect of various matching conditions on their central values and uncertainties. Given these quark mass estimates, the linear interpolations that were done to give the final results will be discussed. The section will close with a discussion of further uncertainties which have not been considered such as lattice artefacts and finite volume effects.

One stochastic $\mathbb{Z}/2$ time-slice source was generated for each gauge configuration from which propagators were computed for all quark masses. The correlation function was then constructed using the one-end trick, as discussed in Ref. [134]. The computation of the correlation functions made use of the software packages [C5] and [C6] and the analysis techniques are implemented in [C3] using also routines from [C2].

Local-local, fuzzed-local and fuzzed-fuzzed correlation functions were computed to improve the isolation of the ground state mass. This was extracted from a constrained 2×2 matrix fit, while the decay constant extracted from the local-local correlation function. To profit from correlations in the data, the complete analysis was carried out in a stationary blocked bootstrap [182] framework with block lengths tuned to accommodate the short autocorrelations in the data as determined from the Gamma method [144]. All observables were bootstrapped with the same bootstrap

L/a		bare valence quark masses					
	$a\mu_l$		0.003	0.006			
24, 32	$a\mu_s$	0.0224	0.0231	0.0238	0.0245	0.0252	0.0259
	$a\mu_c$	0.2586	0.2704	0.2822	0.294	0.3058	0.3176
	$a\mu_l$	0.0009					
48	$a\mu_s$		0.0231	0.0238	0.0245	0.0252	
	$a\mu_c$		0.2704	0.2822	0.294	0.3058	

L/a	fit range minimum and maximum time-slices						
	$\pi^{\pm}$	$\pi^{(0,c)\star}$	π^0	K	D	D_s	$(\Delta t)_{min}$
24	[9, 23]	[9, 23]	[7, 15]	[9, 23]	[12, 23]	[18, 23]	6
32	[9, 28]	[9, 28]	[7, 15]	[9, 28]	[13, 27]	[15, 27]	6
48	[9, 47]	[11, 45]	–†	[9, 47]	[11, 35]	[11, 35]	6

Table 7.6: Bare valence quark mass parameters and fit range restrictions for the computation of pseudo-scalar meson correlators used in this analysis. $\star$: $(0, c)$ refers to the connected part of the neutral pion; †: disconnected contributions have not yet been computed on the physical point ensemble

samples, preserving all correlations. This reduces the statistical error in ratios and other combinations of the various computed quantities. It also allows data to be interpolated such that the result of the interpolation carries virtually the same statistical uncertainty as the data points the interpolating curve was fitted to, giving realistic error estimates.

Fit Range Dependence and Reliable Central Values

The choice of fit range for a correlation function is quite ambiguous as excited state contamination as well as random oscillations in the data can move the apparent onset of the plateau in effective masses by multiple time-slices. In addition, correlations between the time-slices can cause data at several successive source sink separations to rise and fall together, delaying or expediting the onset of an apparent plateau. This kind of correlation can be seen in the effective masses of the pion and kaon shown in Figure 7.4, for example. Both of these effects have been studied to a limited extent as early as in Ref. [183], but even modern analyses often only take into account variations of the fit ranges by a few time-slices in either direction, concluding that the resulting effect is covered by the statistical error. Although this is true in most cases, for certain quantities, most notably the pion mass and decay constants, the spread of values determined using different fit ranges is larger than the statistical error.

It is further observed that round-off errors stemming from the computation of heavy quark propagators can induce unwanted systematic biases in the extracted masses and amplitudes. This is especially true with the twisted mass clover action because the clover term increases the number of matrix vector products involved in the inversion. The effect can be clearly seen in the right-most panel of Figure 7.4 showing the effective mass from the D_s meson correlation function. The effect is also present in the correlation function of the D meson, but there the statistical error grows much faster and the deviation cannot be observed.

An additional effect which may bias fit results is the low signal to noise ratio for large source sink separations. Since the fits in this chapter make use of the full inverse variance-covariance matrix, enhanced decorrelation at large source sink separations can increase the relative contribution of these data to the correlated χ^2 function, possibly pushing the central values up or

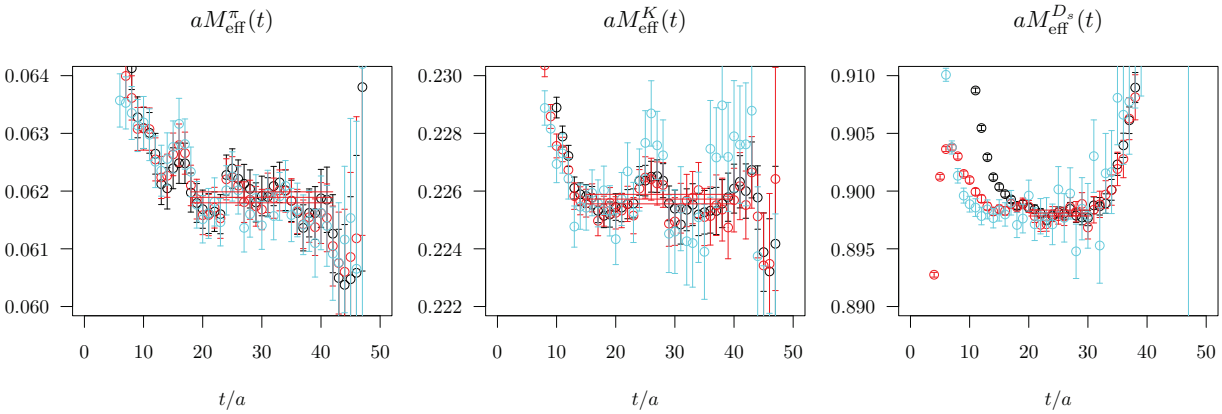

Figure 7.4: Effective masses of the pion, kaon and D_s meson computed on ensemble *cA2.09.48* from local-local (black), local-fuzzed (red) and fuzzed-fuzzed (cyan) correlation functions. Oscillations and correlations between time-slices can clearly be seen here, affecting the apparent locations of the onset of the respective plateaus. The three lines indicate one possible choice of fit range and the resulting effective mass and statistical errors from a constant fit, taking into account the full variance-covariance matrix. For the D_s, the deviation for $t/a \gtrsim 32$ is likely a sign of round-off errors which appear because the correlation function is numerically small in this region and falls by many orders of magnitude over the shown range.

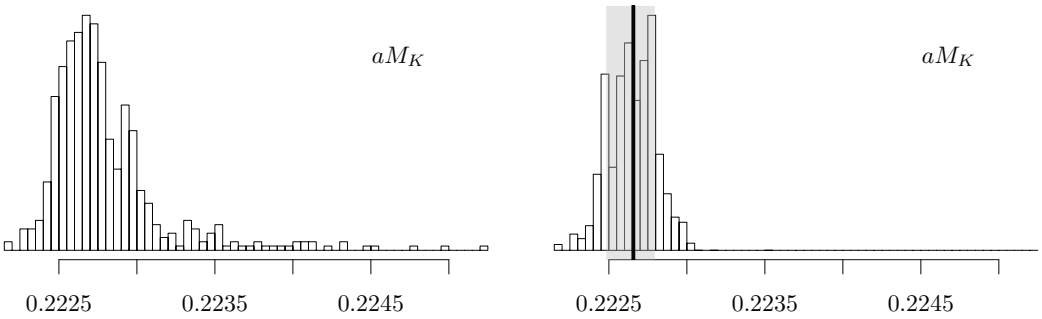

Figure 7.5: (left) Distribution of masses extracted from constrained matrix fit to kaon correlation function at $a\mu_s = 0.0245$ on the set of all "reasonable" fit ranges as described in the body of the text. **(right)** The same distribution after weighting according to Eq. 7.16 with the weighted median indicated by the thick vertical line and 34.27 percentiles around the median shown by the shaded area.

down. Combined with a bias due to round-off errors, this can have a similar effect to excited state contamination.

In order to quantify the ensuing ambiguities coming from these effects without overestimating them, an analysis technique was implemented which takes into account all *reasonable* fit ranges through a weighted average. The method was subsequently also employed by the authors of Ref. [184]. A somewhat arbitrary choice of about 0.5 fm is made for the minimum length of a fit range, denoted $(\Delta t)_{\mathrm{min}}$ (which corresponds to 6 successive time-slices). The meaning of "reasonable" is that fit ranges are only included in the analysis for a given quantity if all relevant fits converge on all bootstrap samples. A full listing of the minimum and maximum source-sink separations this requirement entails for the various quantities in this analysis is given Table 7.6. In the case of the kaon on the *cA2.09.48* ensemble at the physical pion mass, for example, this results in 561 fits with a distribution of fitted masses as shown in the left panel of Figure 7.5. Subsequently, the

Table 7.7: Bare quark masses resulting from matching the quantity in the leftmost column. The labels (LD) and (CD) correspond to f_K (f_D) extracted according to the two definitions given in Eq. (7.8). The starred $a\mu_c$ are derived from the corresponding $a\mu_s$ and the HPQCD charm to strange ratio. The bold values are the strange and charm quark masses used for the final results of the analysis.

	$a\mu_s$	$a\mu_c$
FLAG/HPQCD	0.0247(4)	0.293(6)*
$M_K/f_K^{(CD)}$	$0.02536(10)(^{+05}_{-05})$	$0.3005(42)(^{+06}_{-06})^\star$
$M_K/f_K^{(LD)}$	$0.02480(10)(^{+04}_{-04})$	$0.2938(41)(^{+04}_{-05})^\star$
M_K/M_π	$\mathbf{0.02485(7)}(^{+\mathbf{4}}_{-\mathbf{3}})$	$0.2940(40)(^{+04}_{-04})^\star$
$M_D/f_D^{(CD)}$	–	$0.3629(66)(^{+70}_{-96})$
$M_D/f_D^{(LD)}$	–	$0.2902(26)(^{+09}_{-17})$
M_D/M_π	–	$\mathbf{0.3075(15)}(^{+\mathbf{14}}_{-\mathbf{14}})$

fits are weighted according to their p-values and statistical errors Δ by the weight

$$w = \left(\frac{1}{\Delta}\left(1 - 2\cdot|p - 0.5|\right)\right)^2 , \tag{7.16}$$

resulting in the distribution in the right panel of Figure 7.5.

The same approach is taken for ratios, with the difference that only those fit ranges are considered which have been analysed for both the dividend and the divisor and the weight factor is modified to

$$w = \left(\frac{1}{\Delta_{12}}\right)^2 \left(1 - 2\cdot|p_1 - 0.5|\right)\left(1 - 2\cdot|p_2 - 0.5|\right) . \tag{7.17}$$

The median of this weighted distribution is taken as the central value and its statistical error is computed on the bootstrap samples. The estimate of the systematic error is given by the 34.27 percentiles around the median of this weighted distribution. As an example, the median and systematic error in the determination of aM_K at $a\mu_s = 0.0245$ are shown in Figure 7.5. This quantifies the ambiguity on the raw data which enters the interpolations and the subsequent analysis. Clearly, this error needs to be propagated and this will be discussed below after the tuning of the strange and charm valence quark masses has been presented.

Tuning the Strange and Charm Valence Quark Masses

In very early versions of this analysis (Ref. [P1], for example), the bare strange and charm valence quark masses were fixed using the $N_f = 2$ strange to light quark mass ratio from Ref. [170] and the charm to strange quark mass ratio given in Ref. [141]. These quark mass ratios will be respectively referred to as the *FLAG* and *HPQCD* ratios in the following. Although this is in principle valid because the *cA2.09.48* ensemble is demonstrably at the physical pion mass, the quark mass ratios in the literature have significant uncertainties and simply using their central values to set the valence masses does not propagate the uncertainty to observables.

In order to obtain a reliable error estimate and to allow comparison with an analysis performed using quark mass ratios, all quantities in this analysis were computed at four values each of the bare strange and charm quark masses. This gives up to 16 different combinations for observables depending on both, as listed in Table 7.6. At heavier than physical light quark mass, a total of 36 combinations of strange and charm masses were used to make more reliable interpolations.

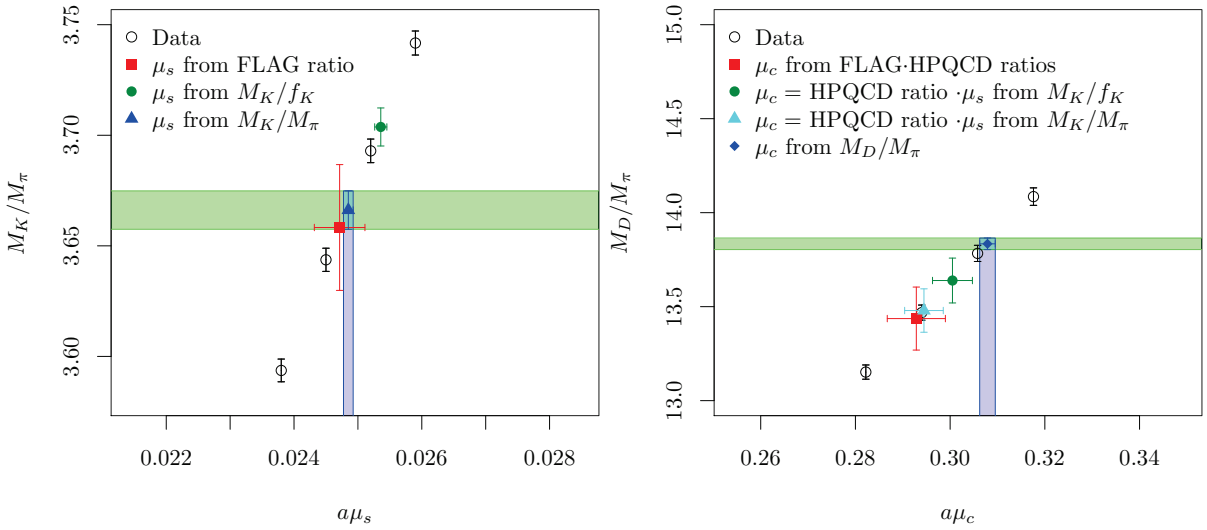

Figure 7.6: Linear interpolation in the strange and charm quark masses and matching with the phenomenological values of the ratios M_K/M_π and M_D/M_π. The other points are given for reference and the strange quark mass from M_K/f_K is determined using the "continuum definition" of af_K with the phenomenological value indicated by the green band.

The mass dependence was fitted linearly in the valence quark masses, with independent fits for all quantities in this analysis. The various quantities can now be interpolated to the bare quark masses given by the FLAG and HPQCD quark mass ratios and the error can be properly propagated. Alternatively, the strange and charm quark masses can now also be fixed by matching to the phenomenological values of ratios of mesonic quantities. Here, this is done for $M_K/M_{\pi^\pm}$ and $M_D/M_{\pi^\pm}$, as shown in Figure 7.6. This results in much smaller errors on the estimates of the strange and charm quark masses for the rest of the analysis than using the quark mass ratios or by matching other quantities, such as M_K/f_K, for example.

Quark mass estimates extracted from different matching conditions or by partial application of the quark mass ratios are given in Table 7.7. It is clear that especially in the charm sector, the choice of matching condition can significantly affect the central value as well as the uncertainty of the quark mass estimate, hinting at the presence of various lattice artefacts. All the final results are quoted with $a\mu_s$ and $a\mu_c$ as determined from the matching to the phenomenological values of M_K/M_π and M_D/M_π, with the other values given for comparison only.

It is interesting to note that for the strange quark mass the usage of the lattice definition of the decay constant in the ratio M_K/f_K results in good agreement with the value of $a\mu_s$ as given by the $N_f = 2$ FLAG strange to light quark mass ratio and the one determined from M_K/M_π. In the charm sector, usage of the lattice definition in M_D/f_D results in a charm quark mass which agrees with those determined via the HPQCD charm to strange quark mass ratio and the three strange quark masses discussed above. The statistical and systematic errors on $a\mu_c$ derived from the lattice definition of M_D/f_D are quite small because in this definition, the charm quark mass dependence of af_D is suppressed, giving M_D/f_D a substantial slope. The large value of $a\mu_c$ and the associated uncertainties derived from the continuum definition of M_D/f_D are just a reflection of how flat the behaviour of this ratio becomes as a function of $a\mu_c$, which can be seen as an indication for discretisation errors. The charm quark mass

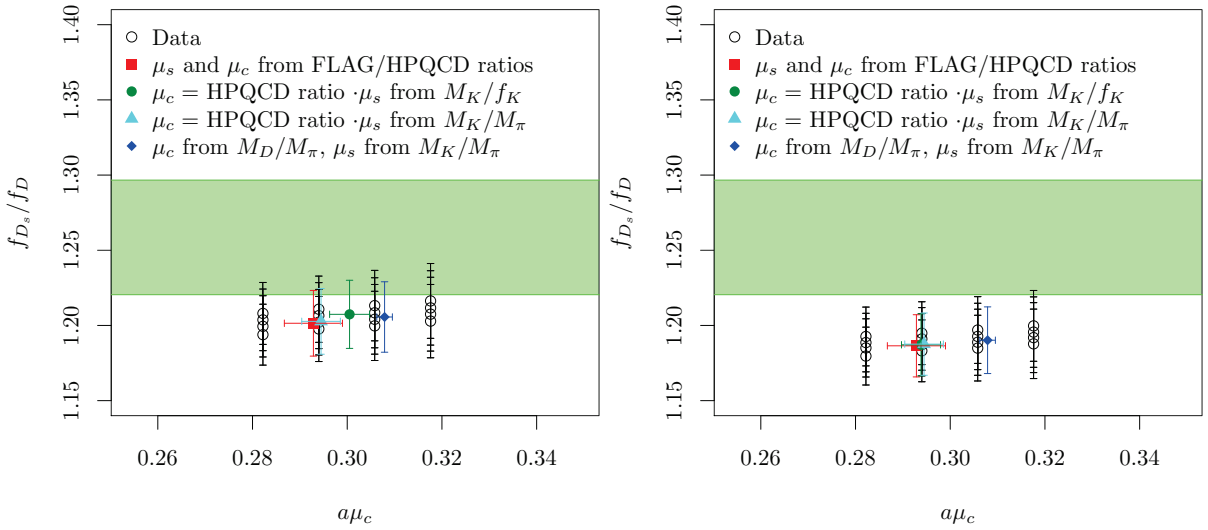

Figure 7.7: Quark mass interpolation of the ratio f_{D_s}/f_D on the physical pion mass ensemble *cA2.09.48* with the phenomenological value indicated by the green band.

determined from M_D/M_π has a statistical uncertainty lower by a factor of two or three compared to the other estimates but disagrees with the other values. In addition to the possibly sizeable lattice artefacts in M_D, finite size corrections on M_π are likely to be at the few percent level which means that without the necessary corrections, the current uncertainties are likely to be strongly underestimated.

In principle, at the cost of losing predictivity for f_K and f_D, estimates for the physical strange and charm quark masses could be derived from weighted averages of some or all of the lattice determinations given in Table 7.7 or combined fits in which several matching conditions would be used simultaneously. The spread of the different values could then be taken as a first estimate of systematic uncertainties due to discretisation and finite volume artefacts. For the ensembles at larger than physical pion mass the values of the bare strange and charm quark masses determined on ensemble *cA2.09.48* were used. Thus, simulations at the physical pion mass give an alternative approach also for the chiral extrapolation with the strange and charm quark masses held constant at what are presumably good estimates of their physical values.

Interpolations

After reliable central values and statistical errors have been determined for the data as described above, independent linear interpolations were performed in all quantities under study towards the values of the strange and charm quark masses listed in Table 7.7. In principle, other approaches could be used for the interpolations, such as using the squares of the quantities or forms inspired by chiral perturbation theory, but for reasons of simplicity and consistency, all the data is interpolated linearly. This seems to be very well justified by the shortness of the interpolations and the shape of the quark mass dependences. Here, only the statistical error is used as a weight for the linear fits because of the difficulties involved in defining a sum of squared residuals with asymmetric weights. The statistical uncertainties in the values of the quark masses are propagated by

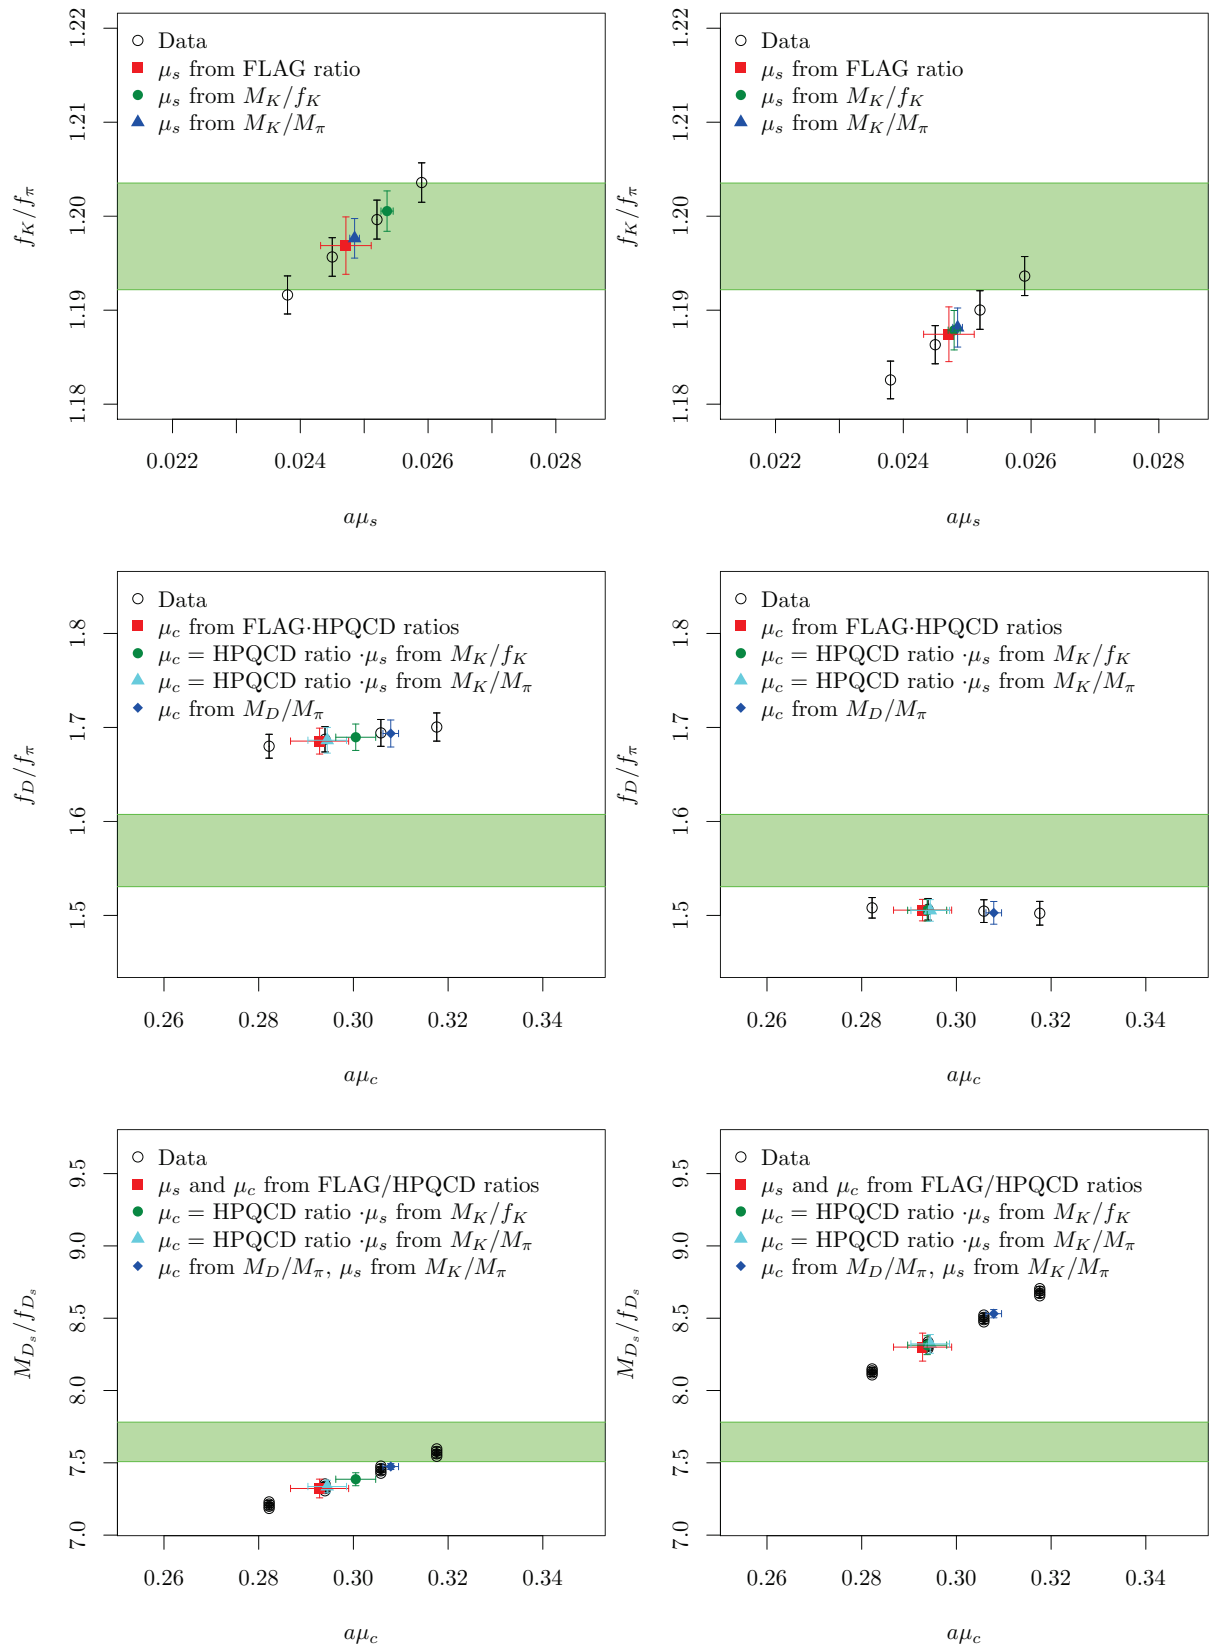

Figure 7.8: Representative choice of interpolations of various quantities involving decay constants using the continuum definition **(left)** and the lattice definition **(right)** with the phenomenological value indicated by the green band.

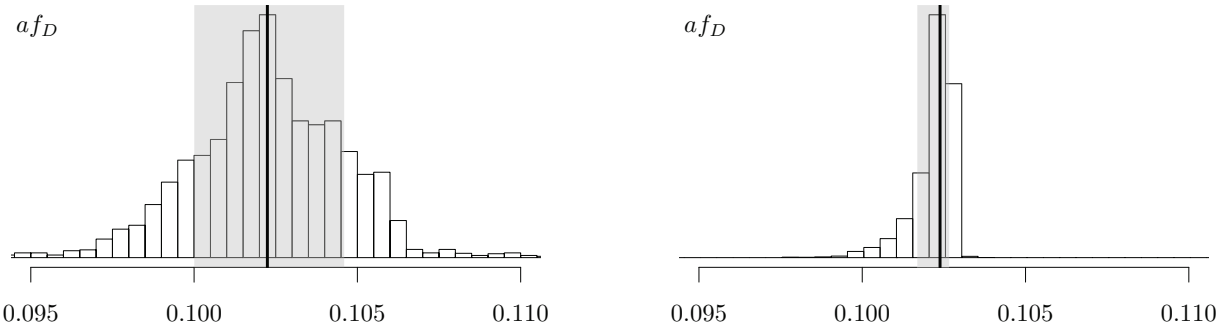

Figure 7.9: Propagation of the systematic error due to fit range ambiguity to the interpolation of af_D. **(left):** Distribution resulting from uniform sampling of data from different fit ranges. **(right):** Distribution resulting from sampling which accounts for the weights of different data points as relative sampling probabilities. The median and the 34.27 percentiles around the median, our estimate of the systematic error, are indicated by the thick vertical line and the grey rectangle.

Taylor expansion, contributing to the total statistical error of the interpolation results in quadrature. A representative set of these interpolations is shown in Figure 7.8 with the continuum definition of the decay constant shown in the left panels and the lattice counterpart in the right panels.

A number of features seen in Figure 7.8 deserve discussion. First of all, for many of the quantities that were analysed, the quark mass dependence is so weak that compatible values are obtained over the whole range of strange and charm quark masses. Consequently, the error is also largely independent of which set of strange and charm quark masses is used, as exemplified by the ratio f_{D_s}/f_D in Figure 7.7. This is of course not unexpected for decay constants and it shows that for many quantities, a slight mis-tuning of the valence strange and charm quark masses does not lead to significant biases when working at the physical light quark mass. This might of course change with quantities computed on $N_f = 2+1+1$ ensembles with a range of sea and valence strange and charm quark masses.

For a quantity like M_{D_s}/f_{D_s} which has a noticeable slope in both the strange and charm mass, which interpolation point is chosen has a strong effect on the central value as well as the errors, which differ by up to a factor of 4 as shown in the bottom-most panels of Figure 7.8. The next notable feature concerns the (unsurprisingly) rather large effect of the definition of the decay constant on ratios involving f_D and f_{D_s}, except when both are involved simultaneously. Clearly, the two definitions agree in the continuum limit and these differences show that discretisation effects could be at the level of 15% to 20% for quantities involving charm quarks. It should be kept in mind, however, that the different dispersion relations do not hint at the size of actual lattice artefacts. Finally, quantities involving M_π and f_π are expected to be subject to finite volume corrections at the few-percent level which will be accounted for in a future study with finer lattice spacings and larger volumes.

To propagate the systematic error to the quark mass estimates and the results of interpolations, 5000 random samples of the data points involved in a given interpolation were generated by randomly drawing from the various fit ranges for each combination of quark masses. In order to obtain a reliable estimate of the resulting error, rather than sampling uniformly, the weights of Equation (7.16) and Equation (7.17) were used as relative sampling probabilities, such that data with large weights occurs more frequently in the set. The effect of this choice can be rather profound and correctly propagates

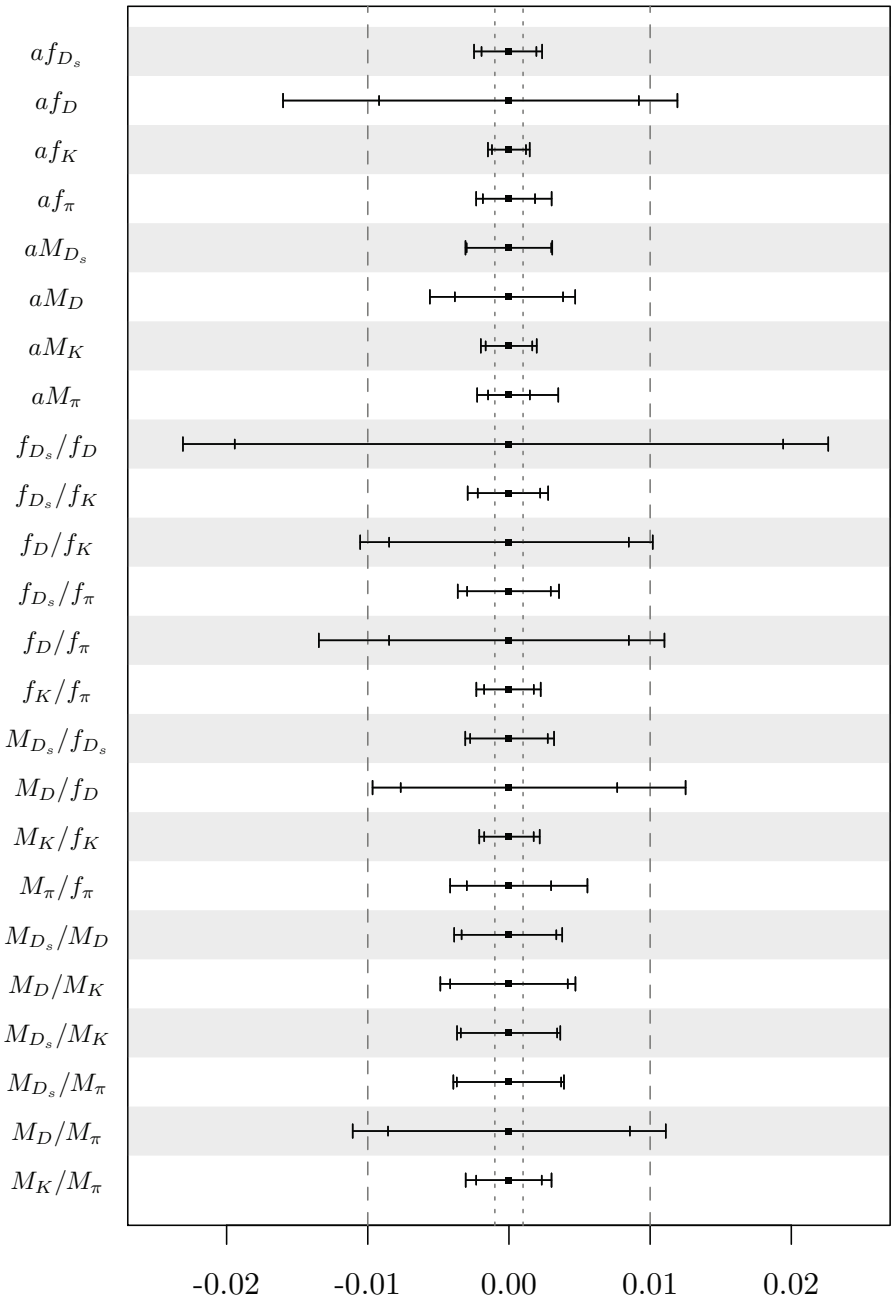

Figure 7.10: Error budget for various mesonic observables on ensemble *cA2.09.48* relative to their central values. Those involving strange and charm quarks have been interpolated to physical strange and charm quark masses as described in Section 7.4.3. The inner error bar is statistical and includes contributions from the fitting procedure as well as the error propagated from the uncertainty in the quark mass estimates. The outer error bar indicates the systematic error due to the ambiguity in the choice of fit range. The dotted and dashed lines show the per-mille and percent error boundaries respectively. The errors are shown cumulatively and would add in quadrature if combined.

the systematic error to derived quantities which is shown in Figure 7.9 for af_D. The left panel corresponds to the distribution when the data from different fit ranges is sampled uniformly and the right panel shows the distribution when the weights are taken into account, approximately reproducing the distribution of the raw data at the four different charm quark masses.

A summary of the statistical and systematic errors is given in Figure 7.10 for the 24 quantities from this analysis, normalised by their respective central values and including those that are technically not independent. It is clear that for most quantities, the choice of fit range has a very limited effect on the total uncertainty and past analyses were probably well-justified in using only one or a few fit ranges. This is despite the fact that the features of the effective masses discussed above suggest significant ambiguities in the estimate of the fit range. For the pion mass and decay constant, however, the systematic error is on the order of the statistical error and must be taken into account.

Quantities involving the D meson show significant spread which might increase even further as the volume is enlarged and more fit ranges become available. On the one hand this is caused by the lightness of the pions which limit the signal to noise ratio for large source-sink separations. On the other hand, this originates from the discussed round-off errors in the computation of charm quark propagators which cause the effective mass to deviate from a plateau at a level exceeding the statistical error for large source-sink separations. Because this limits the maximum source-sink separations that can be taken into account, future studies should be performed with more robust solvers, possibly with a number of iterations in quadruple precision. This should allow to significantly extend the source-sink separation for which a reasonable plateau can be observed, as suggested in Ref. [185]. It would probably especially benefit the analysis of quantities related to the D_s meson.

Part IV

COMPUTATIONAL ASPECTS

A major focus in the preparation of this thesis was on computational aspects of Lattice QCD through numerous changes to the tmLQCD software suite [143, C1]. This part is dedicated to a report on the work to enable simulations at the physical pion mass on modern supercomputer architectures as well as a number of additional developments which resulted from this.

In Chapter 8, the development of multi-threading in modern hardware architectures will be discussed and how its advantages and challenges relate to LQCD codes. This will be followed by an overview of the techniques that were adopted in the development of a multi-threaded implementation of the tmLQCD software package. For this purpose, the OpenMP application programming interface will be introduced.

Based on a number of particular use-cases, the reasoning for the particular approach taken in tmLQCD will be elucidated with respect to considerations of performance, safety, correctness, maintainability, readability and the retention of single-thread optimisations. The value of sampling-based performance profiling in this process is demonstrated and a visual representation is introduced which aids in the interpretation of these profiles.

As a result of the significant changes to the code-base, a testing programme was carried out to ensure the continued correctness of the various algorithms and Chapter 9 outlines a number of subtle issues which were discovered in this process. A number of general techniques are presented which were used for these correctness tests and their optimisation and usefulness is demonstrated with specific examples.

Subsequently, Chapter 10 gives an overview of the performance characteristics of the tmLQCD code using these new developments with a focus on understanding remaining overheads. Parts of this Chapter 8 and Chapter 8 were presented in Refs. [P2] and [P3].

Finally, a potentially interesting observation regarding the tuning of step sizes for integrators with multiple time-scales is presented in Chapter 11.

8

MULTI-THREADING WITH OPENMP

Computationally demanding applications such as Lattice QCD make use of multiple forms of parallelism to scale problems over tens or hundreds of thousands of computational units on modern supercomputers. For the past two decades, distributed memory parallel computers have consistently outranked their shared memory counterparts both in number and in absolute performance.[1] Applications for these systems generally exploit the parallelism using some implementation of the MPI library [186], which provides functions and data types for efficient inter-process communication locally or across network interconnects.

The increase in the performance of supercomputers can be related back not only to the increased parallelism of these machines, but also to the increase in the number of components in microprocessors. In 1965, based on only a few data points, Gordon E. Moore predicted [187] that the number of components of single integrated circuits would double roughly every 18 to 24 months. This scaling, termed "Moore's law", has been satisfied by modern microprocessors until about the turn of the millennium. In the first few years of the twenty-first century, however, it became clear that silicon-based semiconductors would not be able to scale much beyond this point without increases in on-chip parallelism. For this reason, IBM® invested significant resources into the POWER4™ [188] and POWER5™ [189] microarchitectures, two of the first general purpose microprocessors with multiple cores etched into the same die and capable of executing multiple threads simultaneously.

With the number of functional units on modern processors increasing into the triple-digit range such as on the Intel ® Xeon ™ Phi, the overheads involved in multi-process parallelism are beginning to limit the scaling ability of applications. As a consequence, combinations of distributed memory and shared memory parallelism such as multi-threading are becoming increasingly important. More concretely, personal communication with Juelich Supercomuting Centre (JSC) suggested that highest performance on JuQueen, the new IBM® BlueGene/Q™ installation there, could only be reached by hybrid codes employing both thread-level parallelism and inter-process communication. Simulations at the physical pion mass involve extraordinary amounts of computation and the number of functional units on modern supercomputers is bound to increase significantly in the next few years. In addition, the BlueGene/Q™ architecture constitutes a primary provider of computational resources for the ETMC. For these reasons, a multi-threaded implementation of the tmLQCD code-base was considered a necessity.

There is a high level of debate on whether threads are the most suitable mechanism to increase application-level parallelism for deterministic problems such as scientific computations. For example, Ref. [190] makes a very strong argument that the inherent non-determinism of threads needlessly complicates the introduction of concurrency for this class of computational problem. In particular, the loss of sequential ordering of operations, the (possibly) countably infinite set of thread interleavings and the related is-

[1] A very interesting visualisation of this can be found here: `http://www.top500.org/statistics/overtime/` (visited 2015/05/20)

sues of memory locking and data races can entail very subtle defects which hide even in extremely well-tested applications, as in the code discussed in Ref. [190] by the same authors. The conclusion is that threads are a fine way of implementing parallelism at a very low level, but at the application level, except for server or networking code which is by definition non-deterministic, they make reasoning about the execution of a program so complicated that faults are almost guaranteed.

Despite these very serious shortcomings, introducing thread-level parallelism into an existing code-base is generally syntactically simple and there is very good industry and language support for this kind of augmentation. In addition, rewriting a substantial application using a more suitable framework is out of reach of most practitioners of scientific computation and the tmLQCD software suite is no exception to this unwritten rule. Therefore, keeping in mind that threads significantly affect the semantics of a program, the decision was taken to implement thread-level parallelism across the tmLQCD code-base using OpenMP on account of its ease of use, wide adoption and portability [191].

In this chapter the challenges that were faced in this process as well as the methodologies that were adopted will be described. Section 8.1 provides some background on OpenMP and introduces important nomenclature, Section 8.2 gives implementation details specific to the tmLQCD software suite, Section 8.3 discusses the essential issue of overhead reduction while Section 8.4 deals with the problematic topic of locking and data races. Finally, the interaction of inter-process communication and multi-threading is discussed in Section 8.5 and some summarizing remarks are collected in Section 8.6.

OPENMP

In order to introduce some important nomenclature and for completeness, this section gives a short review of the basic functionality of OpenMP. It is designed around special directives introduced into the source code of a program which instruct the compiler to execute the parts they enclose (or precede) in parallel. This is achieved by spawning (*forking*) multiple threads and either having each thread execute the same instructions or distributing different parts of some loop construct amongst the different threads. When a given work-load has been processed, the threads are *joined* and program control is returned to the master thread. More fine-grained parallelism is also possible, such as forcing a certain part of a routine to be executed by one thread only and modern versions of OpenMP even provide advanced constructs for delayed execution.

The details of how these directives are used depend on the type of application and in the following the focus will be on a structure that is most relevant for Lattice QCD. This often consists of a *driver* routine which manages the computation and context (this could be the `main` function, for example) and a multitude of *worker* routines which typically contain loops to be iterated over for some given work-load, such as a matrix-vector multiplication.

When OpenMP is added to such an application, thread-level parallelism is usually introduced at the level of the worker routines in order to distribute the loop iterations amongst threads. This procedure is generally quite straightforward, barring complications such as ensuring that the loop iterations are independent and the possible requirement of `critical`

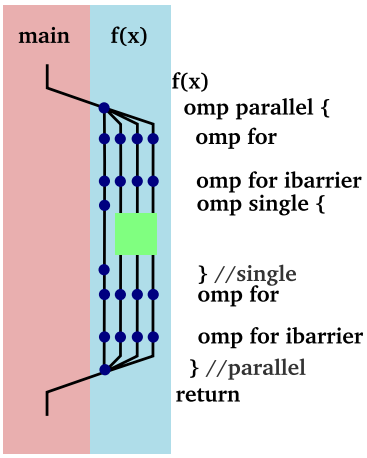

Figure 8.1: Graphical representation of a multi-threaded worker function using fine-grained parallelism. The threads are launched when the function is entered and two work-sets in for loops are separated by a single-threaded section (for MPI communication, for example). The threads are joined before the function returns.

sections or `atomic` directives which enforce mutual exclusion to prevent data races. Figure 8.1 shows a graphical representation of this kind of implementation with two parallelized loops separated by a single-threaded section. In order to keep overheads minimal, rather than using multiple `omp parallel for` directives, the threads are spawned with a `omp parallel` directive as the worker function is entered and the work-load is distributed using two `omp for` directives. Several of the shown instructions have implicit *barriers*, these directives instruct the compiler to synchronise data between threads and can have significant overhead. With the basic nomenclature in place, the next section will focus on the strategies that were developed for the deployment of OpenMP in the tmLQCD software suite, based on explicit examples from the code-base and general considerations such as maintainability, performance and correctness.

IMPLEMENTATION

The basic ingredients laid out in Figure 8.1 were used to implement thread-level parallelism in the tmLQCD software suite and this procedure is detailed in this section with a particular focus on how functionality was added while trying to avoid removing any existing optimisations. Further, the topic of data locality in OpenMP is discussed and its relation to code maintainability is emphasized, a problem which was found to be simplified by exploiting variable scoping rules of the C programming language and the OpenMP standard for C-style languages. Finally, a number of performance-critical decisions are investigated and possible improvements are discussed.

Exploiting Variable Scoping Rules

A chief concern during the deployment of OpenMP directives in the tmLQCD code-base was that of maintainability. As will be shown, this aspect is intertwined with the localisation of variables in the *private* thread-local context or the *shared* thread-global context. This discussion is best done though a simple example like the function given in Listing 8.1. There, the qualifier `private` indicates that thread-local copies of the variables `temp1` and `temp2` should be allocated as the parallel section is entered. In addition, the variables `result` and `x` are indicated to be thread-global shared variables. Other qualifiers exist which allow for more fine-grained control, especially regarding the initialisation of private and shared variables and the

copying of their final values from the thread-local context. Although quite logical and simple, in a more complicated function the list of variables with qualifiers can have many more members and this makes it a potential source for difficult to trace bugs. In particular, when the algorithm is changed and more variables are added or the meaning of some variables is changed, it could be that updating the qualifier list is forgotten, especially if there are multiple parallel sections in a function. If one of the added variables was intended as private but turns out to be shared because it is omitted from the qualifier, very subtle issues may develop that depend on the execution environment.

Listing 8.1 thread-local and thread-global variables using qualifiers

```
function(double * x) {
  double temp1, temp2;
  double result = 0.0;
  #pragma omp parallel for private(temp1, temp2) shared(result, x)
  for(int i = 0; 0 < N; ++i){
    temp1 = x[i]*x[i+2];
    temp2 = temp1*temp1;
    #pragma omp atomic
    result += temp2;
  }
  return result;
}
```

Listing 8.2 thread-local and thread-global variables relying on variable scope in a C-style language

```
function(double * x) {
  double result = 0.0:
  #pragma omp parallel
  {
  double temp1, temp2;
  #omp parallel for
  for(int i = 0; 0 < N; ++i){
    temp1 = x[i]*x[i+2];
    temp2 = temp1*temp1;
    #pragma omp atomic
    result += temp2;
  }
  } // OpenMP parallel section closing brace
  return result;
}
```

Because tmLQCD is implemented in the C programming language in which variable scope is a central language concept, the decision was made to largely forego the usage of qualifiers and rely on the variable scoping rules instead in the hope that this potential source of bugs could be mostly avoided. OpenMP treats variables declared inside parallel sections as private (and unitialised) by default while those declared outside parallel sections are automatically shared. With this in mind, Listing 8.2 gives an alternative implementation of the function without qualifiers in which variable scope

determines which variable is shared and which is private. In this kind of implementation changes to the algorithms are less likely to produce difficult to find bugs because as long as the simple rule regarding variable scope is followed, the data-localisation is automatic. Although this is quite beneficial, it will be shown in Section 8.3.1 that it likely complicates the introduction of so-called *orphaned* directives for the purpose of reducing overheads.

Conditional Directives

For performance reasons many parts of the tmLQCD code-base use functions with static variables and loops with dependent iterations. Of course, in order to parallelise these using OpenMP, they need to be modified to remove the loop dependencies and static variables. In order to retain the existing single-thread performance, however, it was decided to implement all OpenMP directives in conditional blocks depending on the pre-processor constant `OMP`.

In Listing 8.3 from `derivSb.c`, for example, the `static` keyword is temporarily redefined to nothing and a parallel section is started. Thus, when the code is compiled without support for OpenMP, the static variables are still declared as such and their performance benefit it retained.

Listing 8.3 Redefinition of the `static` keyword to nothing when OpenMP is used.

```
#ifdef OMP
#define static
#pragma omp parallel
  {
#endif
  [...]
  static su3 v1,v2;
  static su3_vector psia,psib,phia,phib;
  static spinor rr;
  [...]
#ifdef OMP
#undef static
#endif
```

Similarly, in the implementation of `operator/D_psi.c`, given in Listing 8.4, a loop dependency is used to increment the integer `iy` as well as the pointers `sp` and `up` as the loop progresses. When multiple threads execute this section, however, their values need to be set explicitly at the beginning of each iteration to avoid this loop dependency. This is again achieved at compile-time through reliance on the pre-processor constant `OMP`.

While these conditional directives have allowed most of the original structure in tmLQCD to be retained, they tend to affect the readability of the source-code because the scope of pre-processor conditionals is often hard to see, even if indentation and comments are used to highlight it. Still, it is the author's belief that the usage of conditional directives allows better code-paths to be chosen at compile-time without code duplication when multi-threading is not to be used. At the time of writing, it seems that this is quite beneficial on most Intel® machines with the current code-base.

Listing 8.4 Loop dependency on `iy, sp` and `up`. Without OpenMP their values carry over from one iteration to the other, otherwise they need to be set explicitly for each thread.

```
[...]
#ifndef OMP
  iy=g_iup[0][0];
  sp=(spinor *) Q + iy;
  up=&g_gauge_field[0][0];
#else
#pragma omp for
#endif
  for (int ix=0;ix<VOLUME;ix++){
#ifdef OMP
    iy=g_iup[ix][0];
    up=&g_gauge_field[ix][0];
    sp=(spinor *) Q + iy;
#endif
[...]
```

Retaining Kahan Summations

Rounding errors are an unavoidable part of any numerical computation which may or may not affect the quantitative and qualitative properties of a result depending on the type of problem under study. In lattice QCD, the equations of motion integrated in the molecular dynamics step of the HMC algorithm are chaotic [192] and rounding errors at various stages in the simulation may affect physical observables unexpectedly strongly as a result. This was studied to some extent in Ref. [193] with very heavy dynamical quarks and small volumes using two integration time-scales. Although no significant effect on observables could be detected, round-off errors were predicted to play a measurable role on production lattice volumes in the energy violation which determines the acceptance probability. The issue was studied further in Ref. [147] for more realistic volumes and quark masses and some tests were done for Refs. [79] and [117], but the author is unaware of any recent studies on present production volumes and physical quark masses. In general, it is desirable to reduce round-off errors to a minimum, especially in summation operations in which numbers of vastly different magnitude may be added together, resulting in significant loss of precision because of this difference. These operations occur for instance in the computation of the scalar vector products required in the conjugate gradients (CG) algorithm or the energy contributions of the different monomials in the acceptance step of the HMC algorithm. A simple yet effective way of reducing round-off errors is the *Kahan* summation [194], in which the sum and the round-off error are added up concurrently and finally subtracted from each other.

Unfortunately when a program is parallelised, the summation (*reduction*) algorithms offered by MPI and OpenMP do not natively support this kind of improvement. On the level of inter-process communication, this is probably not a big issue because the individual contributions to the global sum coming from different processes can arguably be expected to be of similar magnitude if the process-local lattice volume is sufficiently large. In addition, there is nothing which prevents the sum over the process-local volume to be a

Listing 8.5 Multi-threaded sum retaining Kahan summation for each thread.

```
double kahan_sum(const double * const input, const unsigned int N){
  double res;
#ifdef OMP
#pragma omp parallel
{
  int thread_num = omp_get_thread_num()
#endif

  double sum=0,kc=0,ks,tr,tt;

#ifdef OMP
#pragma omp for
#endif
  for (unsigned int ix = 0; ix < N; ++ix) {
    sum += *(input+ix);
    /* each thread does its own Kahan summation */
    tr=sum+kc;
    ts=tr+ks;
    tt=ts-ks;
    ks=ts;
    kc=tr-tt;
  }
  kc=ks+kc;

#ifdef OMP
  /* the results from each thread are stored in a global helper array */
  g_omp_acc_re[thread_num] = kc;

  } /* OpenMP parallel section closing brace */
  /* and added up after leaving the parallel section */
  res = 0.0;
  for(int i = 0; i < omp_num_threads; ++i)
    res += g_omp_acc_re[i];
#else
  res = kc;
#endif
  return(res);
}
```

Kahan sum, thus taking into account the most significant source of rounding errors because variations from lattice site to lattice site are probably more significant than from process to process.

In OpenMP, however, summations are usually done using a construct such as `omp parallel for reduction(+:x)`, where the operand in parentheses specifies the reduction operation and the variable name `x` indicates which object should be summed up. In a Kahan summation, the last operation to be carried out for *each* thread would be the subtraction of the sum and the rounding error and it is clear that within the confines of OpenMP doing so is not possible. This is because keeping track of the round-off error introduces a special kind of loop dependency which not only depends on the iteration, but also which part of the loop is being executed.

OpenMP doesn't guarantee the order in which reductions are carried and because rounding errors can make identifying subtle bugs extremely difficult, especially in a mutli-threaded program. As a result, it was decided to retain the Kahan sums by foregoing OpenMP reduction directives in favour of the manual approach given in Listing 8.5. Here, each thread performs a Kahan summation for its part of the sum over some loop (over lattice sites, for example) and the results for each thread are stored in a global pre-allocated array (of which there are two: one for real and one for complex double precision variables). The final reduction is carried out by the master thread once the parallel section has been exited, resulting in a similar situation as with pure MPI because the summands in the final sum are likely to be of a similar order of magnitude. Because the implementation here adds a small overhead, it would be interesting to see what effect, if any, the removal of Kahan summations would have on physical observables or the number of CG iterations to reach a certain accumulated residual and the deviation of this from the real residual, but this has not been attempted yet.

Summary

The deployment of OpenMP in the tmLQCD software package stands out through a number of particular considerations. First of all, keeping optimisations in the single-threaded versions of many functions was considered important and thus OpenMP directives were made dependent on the preprocessor constant `OMP`. In the author's opinion, even though this makes the code less readable, it results in a code-base which is at least theoretically more optimal than an implementation which attempts to compromise between the requirements of multi-threaded and single-threaded execution. In addition, the fact that all the OpenMP directives can be reliably disabled allows for checking serial and multi-threaded codes against each other, functionality which may be quite useful for testing purposes. This is especially true if bugs in the compiler or how it is set up on some machine make the compiler always interpret OpenMP directives, even if a flag to disable them is passed.

A further important focus was that of maintainability and here the decision was made to avoid data-locality qualifiers wherever possible and instead rely on the variable scoping rules of the C programming language. Although this makes for a more readable code-base and certainly reduces the probability of bugs related to keeping the qualifier lists up to date, it is likely that it interferes with the deployment of so-called *orphaned* directives. Finally, it was considered important to retain Kahan summations in important parts of the code-base, foregoing the OpenMP `reduction` directive for a

manual approach which allows multi-threaded Kahan summations to be implemented. Moving on from details of the implementation to more practical matters, the issue of overheads related to multi-threading and their reduction will be discussed in the next section.

REDUCING OVERHEADS

Although the standard implementation procedure shown in Figure 8.1 is generally effective, it will be shown below that some of the largest overheads still stem from thread-management in `parallel` directives. This experience agrees with conclusions from the literature such as Ref. [195] or, for a larger scale shared-memory machine, Ref. [196]. Even on machines like BG/Q, where shared-memory parallelism is local to one node, hardware threads are kept alive between parallel sections and thread sleeping is rather light-weight, resetting the thread-local stack and ensuring proper synchronization of shared variables induces significant costs.

In this section a number of overheads specific to multi-threaded code will be discussed and the strategies that are available for reducing these overheads as well as the choices that were made during the work on the tmLQCD software package. In particular, two mitigation strategies for thread management overheads will be considered and the issues of thread contention and false sharing as well as thread idling from cache misses will be discussed.

Many of the strategies described in this section were guided by measurements using SCALASCA [197], an automatic instrumentation-based toolkit which uses profiling for performance measurements of highly parallel applications. It must be noted here that tools relying on instrumentation are often seen critically because they are known to have large overhead, thus potentially skewing the measurements. In the author's experience, however, the relative timings given by SCALASCSA of various sections of a given routine seem to reflect manual timings sampled from hardware counters. Further, relative reductions in overheads from optimisations seem to agree between profiling and time-to-solution measurements without instrumentation. Thus, since the main focus is on measuring and minimizing relative cost, one can argue that for this purpose SCALASCA profiles provide a convenient and appropriate measurement framework, although of course absolute timings of an instrumented application can be orders of magnitude larger than the uninstrumented version.

Thread-Management Overhead

As mentioned above it was observed that in the tmLQCD software suite, the main overhead related to OpenMP originates from `omp parallel` directives. The relative effect of this kind of overhead on performance depends very strongly on the complexity of the work-load in a worker function and it is important to keep this in mind when looking for optimisations. Generally, when a function is computationally simple, only a small fraction of time will be spent on the actual work-load with the rest going towards thread management, synchronization and communication overheads. In order to understand this in more detail, the `scalar_prod_r` routine which computes the scalar product of two vectors provides a good example and a version without Kahan summation is given in Listing 8.6.

Listing 8.6 Simplified version of `scalar_prod_r` function.

```
double scalar_prod_r(const spinor * const S, const spinor * const R,
                     const int N, const int parallel)
{
  double ALIGN res = 0.0;
#ifdef MPI
  double ALIGN mres;
#endif

#ifdef OMP
#pragma omp parallel
  {
  int thread_num = omp_get_thread_num();
#endif
  double ALIGN ds=0.0;
  const spinor *s,*r;

#ifdef OMP
#pragma omp for
#endif
  for (int ix = 0; ix < N; ++ix) {
    s = S + ix;
    r = R + ix;

    ds += sp(s,r);
  }

#ifdef OMP
  g_omp_acc_re[thread_num] = ds;

  } /* OpenMP closing brace */

  for(int i = 0; i < omp_num_threads; ++i)
    res += g_omp_acc_re[i];
#else
  res = ds;
#endif

#if defined MPI
  if(parallel)
  {
    MPI_Allreduce(&res, &mres, 1, MPI_DOUBLE, MPI_SUM, MPI_COMM_WORLD);
    return mres;
  }
#endif
  return res;
}
```

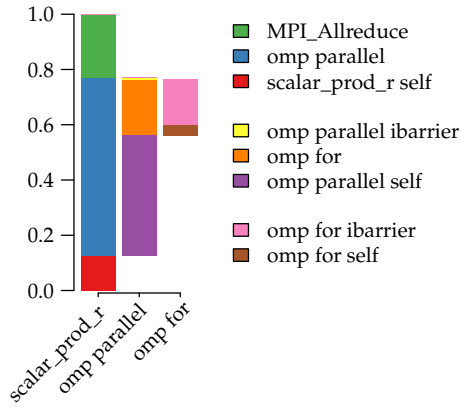

Figure 8.2: Profile of relative time spent in different parts of the `scalar_prod_r` routine in the tmLQCD software suite on 512 nodes of BlueGene/Q during an inversion of the Dirac operator on a realistic $48^3 \cdot 96$ lattice. The three parts of the legend, top to bottom, each correspond to one bar from left to right. The colouring of the bars is mirrored in the colouring of the legend, although some parts of the bars can be so small as to be invisible. The specifier "self" denotes the time spent in a function or section before any of its subroutines are called. For `omp parallel`, this corresponds to the time it takes to spawn the threads.

Figure 8.2, produced by profiling the CG algorithm on a realistic $48^3 \cdot 96$ volume using SCALASCA on 512 nodes of IBM® BlueGene/Q™, shows the relative time spent in the scalar product function and its subroutines from left to right. Because this type of figure is quite complicated and will occur frequently in what follows, it is worth to understand how to best read it. The labels on the horizontal axis give the name of the function, subroutine or directive block under consideration. In the bars, going from bottom to top matches the order of execution in the program and the elements in the three blocks in the legend are arranged in the same order as the parts of the bars, so that both the colour code and their ordering can be used to identify which part is which. The different fractions of the bar give the share of time spent in the sub-parts, the names of which are given in the legend in as many groups as there are bars in the plot. Here, `scalar_prod_r` has three sub-parts: `scalar_prod_r_self`, `omp parallel` and `MPI_Allreduce`. The qualifier *self* refers to time spent in the routine or block under consideration but not spent in any of its sub-blocks. For this function, this may be the time needed to allocate memory for any local variables, for example. The next bar in the middle shows the time spent in the `omp parallel` block and its sub-parts while the rightmost bar gives the time spent computing the actual work-load in the `omp for` work-sharing construct.

Analysing the figure in more detail, the routine spawns threads in a `omp parallel` section and then iterates over the lattice volume in a `omp for` section. After the scalar product for the local volume has been computed, `MPI_Allreduce` carries out the global sum. Keeping in mind the meaning of the *self* qualifier, it can be seen that the block `omp parallel self` is very closely related to the overhead for spawning threads in `scalar_prod_r`. It is clear from the figure that almost 50% of the total time spent in the function is lost on thread management overhead with another 20% spent in the global sum. It will be seen below that this confirms the view that computationally simple functions are dominated by thread management overhead because this overhead is much smaller (yet still significant) in a computationally more complex construct and a prime candidate for optimisations.

It must be noted that the types of overheads seen and the amount of time spent in different parts of functions and their subroutines can vary substan-

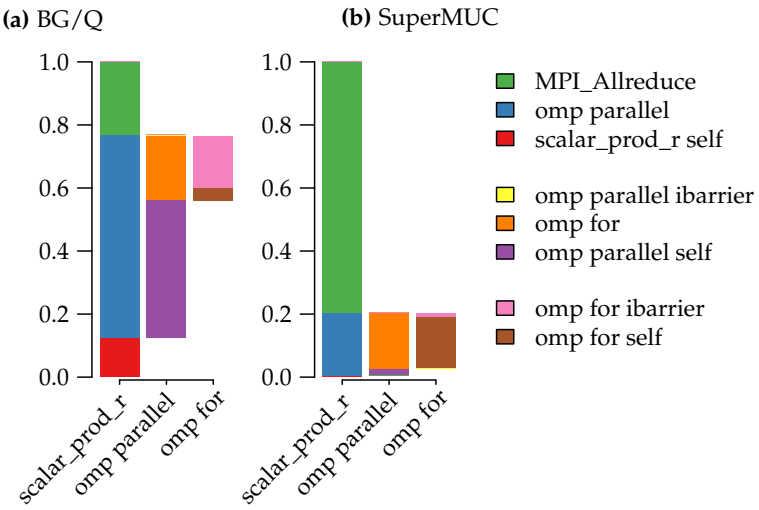

Figure 8.3: SCALASCA profiles of the sclar_prod_r function on 512 (256) nodes of BG/Q (SuperMUC), using 1 (2) tasks per node and 64 (8) threads per process during an inversion on a lattice volume of $48^3 \cdot 96$.

tially from architecture to architecture, further complicating the problem of achieving ideal or close to ideal performance and possibly necessitating different specialised implementations. To illustrate this fact, Figure 8.3 shows a comparison between SCALASCA profiles of the `sclar_prod_r` function given the exact same work-load on BG/Q and SuperMUC, an Intel®-based supercomputer with a fat tree network topology up to 512 nodes equipped with two eight-core Xeon™ processors each. It is clear that on this machine, the collective MPI communication at the end of the computation of the scalar product is a far more important target for optimisation, although this situation may change as the number of cores on this type of architecture approaches that of the BG/Q. Still, this serves to underline the importance of carefully measuring the performance characteristics on different architectures in order to identify hot spots and determine optimisation strategies.

Combining Operations

One obvious way of reducing or eliminating thread management overhead is to combine operations which frequently occur in sequence into a single `omp parallel` section. In tmLQCD, this strategy was adopted for the application of the even-odd preconditioned twisted mass Dirac operator where the combination $(M_{ee}^{\pm})^{-1} M_{eo}$ occurs frequently and was usually done in two subroutines with a volume loop each. First, the even-odd hopping matrix was applied to a spinor and then the even-even inverse part was applied to the result, now both of these operations are done in the same volume loop.

This kind of combination can introduce code duplication because either operation can still be required on its own in other routines, thus requiring multiple implementations. For the case above, this was solved using preprocessor constants and directives which are used to compile several versions of the hopping matrix combined with different simple operations which occur frequently. For definiteness, an extract of the relevant code sections is given in Listing 8.7. While this minimizes code duplication, it does produce a larger binary and it may reduce readability of the code. Finally, the usage of `include` directives is not compatible with automatic instrumentation using performance tools like SCALASCA [197] which fail to instrument the included code sections (which are the most important ones). As a result, if one wishes to use such a tool, it becomes necessary to produce a version of the code-base in which the included files are actually copied into

the code before even the pre-processor is run, adding more development overhead.

Listing 8.7 Source code extract from the half-spinor hopping matrix with simple operations `_MUL_G5_CMPLX` and `_TM_SUB_HOP` combined with the hopping matrix through `include` statements and preprocessor constants.

```
operator/halfspinor_body.c
--------------------------
[...]
for(unsigned int i = 0; i < bodyV; i++){
  [...]
#ifdef _TM_SUB_HOP
  pn=p+i;
#endif
  [...]
#ifdef _MUL_G5_CMPLX
  _hop_mul_g5_cmplx_and_store(s);
#elif defined _TM_SUB_HOP
  _g5_cmplx_sub_hop_and_g5store(s);
#else
  _hop_store_post(s);
#endif
}
[...]

operator/Hopping_Matrix.c
-------------------------
void Hopping_Matrix
     (const int ieo, spinor * const l, spinor * const k) {
  [...]
  #include "operator/halfspinor_body.c"
  [...]
}

operator/tm_times_Hopping_Matrix.c
----------------------------------
void tm_times_Hopping_Matrix
     (const int ieo, spinor * const l, spinor * const k,
      complex double const cfactor) {
  [...]
  #define _MUL_G5_CMPLX
  #include "operator/halfspinor_body.c"
  [...]
}
```

Orphaned Directives

Thread-management overhead can be reduced more comprehensively by using more coarse-grained parallelism. In particular, this can be accomplished through so-called *orphaned* directives which have been shown [198] to provide a significant performance benefit. These are instances of work-sharing directives such as `omp for` not syntactically enclosed by an `omp parallel` section. When such a directive is encountered by multiple threads[2], work sharing functions as expected if the compiler supports orphaned directives. When encountered outside of the execution context of a `parallel` section, the directive is simply ignored and the routine is executed single-threaded. For definiteness, Figure 8.4 illustrates the usage of orphaned directives. Here,

[2] which must have been launched by a `parallel` directive somewhere above the current execution context in the call hierarchy

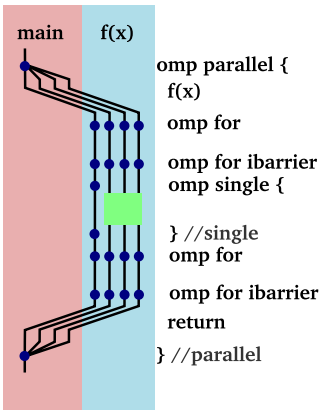

Figure 8.4: Graphical representation of a multi-threaded worker function using coarse-grained parallelism though orphaned directives. The threads are launched further up in the function hierarchy and all of them enter the function $f(x)$. Inside the function, two orphaned `omp for` sections distribute the work-load and code inside an `omp single` section is executed single-threaded. All threads return to the calling function.

the threads are spawned in the calling function and enter $f(x)$ in parallel where the work-sharing constructs are executed as expected and all threads return to the calling function once they complete their workload. This should be compared to Figure 8.1 in which the threads are spawned only once $f(x)$ has been entered.

As will shown in Chapter 10 below, trial benchmarks show that orphaned directives can result in a significant reduction of overheads but they entail several significant complications. Firstly, at least in principle, this kind of orphaning cannot be easily implemented incrementally because wherever an orphaned function is called, it must be entered by multiple threads. As a consequence, in order to avoid leaving parts of the code-base unparallelized and suffering severe penalties via Amdahl's law [199], one some function has been equipped with orphaned directives, every single function calling it must set up the required parallel section. In addition, dealing correctly with nested parallel sections is quite complicated, requiring therefore that if one sub-routine called by some function has orphaned directives, all of them must in order to avoid this nesting of parallel sections. Finally and most importantly, the burden of managing data-locality and coherence moves up the function hierarchy, entailing that the calling function needs to prepare private and shared memory before the sub-routine is entered. This severely complicates the writing of multi-threaded software and necessitates significant changes to function signatures in order to satisfy these requirements.

At first sight, the severeness of this problem of this problem seems quite shocking because it breaks a fundamental idea in software modularisation: the independence of interface and implementation. It can be envisioned, however, that in most cases function signatures simply need to be modified, exposing more of the internal workings to the calling code without signficantly breaking the modularisation paradigm, although of course some maintanability and readability is necessarily lost. In fact, in many situations in the tmLQCD software suite this is already done by relying on the calling code to allocate and free temporary working memory, for instance. Less severely, any maintainability benefits gained from relying on scoping rules to ensure proper locality of variables as discussed in Section 8.2.1 will almost certainly be lost, at least as far as thread-global variables are concerned. Finally, some operations would be more difficult to implement, a multi-threaded Kahan summation, for example, would require the final sum to be carried out in a `omp single` section, possibly inducing more overhead than the performance gain achieved from using orphaned directives in the first place.

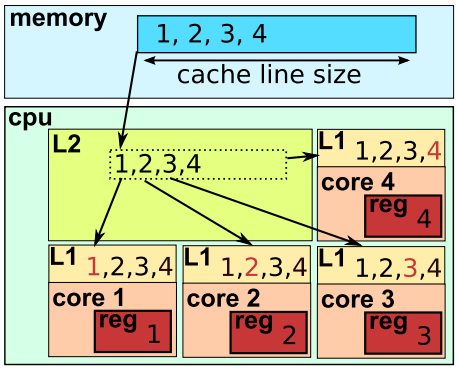

Figure 8.5: Illustration of false sharing. Even though none of the four threads shares any of the values in the cache line, all of the threads must load the cache line. If one of the threads were to change a value, the result would be written back to memory and all threads would be forced to refetch the entire cache line even tough only one of the values had changed.

For these reasons, orphaned directives have only been tested to a limited extent in tmLQCD, although a potential strategy for a full implementation has been established. OpenMP provides mechanisms for determining whether the execution context is within a `parallel` section which should enable the design of two versions of each multi-threaded routine with the correct code-path chosen based on the execution context. It is envisioned that code duplication can be avoided with the same strategy used for the combination of multiple routines into single `parallel` sections through `include` statements. There remain, however, a number of questions on how data locality should be handled in these situations. Given that increased on-chip parallelism and increases in multi-threading by orders of magnitude are the cornerstone of current concepts towards exascale supercomputers [200], a dedicated effort towards solving this issue in tmLQCD will be necessary in the very near future. As a final conclusion one can also imagine the important role that performance modelling and measurement will play in any such endeavor because the performance characteristics of different approaches will be difficult or impossible to predict, requiring either accurate models or the implementation of multiple solutions and their subsequent comparison.

Avoiding Cache Contention and False Sharing

In multi-threaded programming two closely related issues which can easily produce performance bottlenecks are cache contention and false sharing [201]. For performance reasons, reading from memory on most architectures can only be accomplished in blocks of a certain size, referred to as the size or length of the *cache line*. In addition, the number of these cache lines which can concurrently be *pre-fetched* from memory is also limited, on the BG/Q, this number is lower than the number of hardware threads. This means that when many threads read from memory at the same time, the pre-fetching algorithms in the compiler and hardware will have to contend with competition between threads. In particular, when too many threads attempt to issue pre-fetch instructions, performance is markedly reduced. This agrees with experience with pre-fetching instructions in tmLQCD on BlueGene/Q™ in the hopping matrix kernel, for which these instructions seem to provide only very small benefits and even then only in well-chosen places. One is forced to conclude that the memory access patterns in multi-threaded programs must be carefully studied if this kind of overhead is to be reduced.

False sharing is a related issue, illustrated in figure Figure 8.5. Here, four threads running on four cores load sequential values from memory into their respective registers and as discussed above, when a datum is requested

from main memory, an entire block of sequential data must be loaded rather than the single value. As an additional complication, whether this block begins exactly at the requested value or at some coarser byte-level boundary depends on the architecture in question. When multiple threads read from the same or overlapping cache lines and one of the threads commits a new value to memory, the entire cache line is *invalidated* to ensure consistency. This forces all cores to reload the cache line, even though only one value was changed and their respective data remains unaffected. Further, the existence of multiple levels of cache complicates the matter even more because an invalidated cache line must be propagated from memory along the different cache levels, some of which may be shared by all cores of a given processor, like the level 2 cache in this example. The issue can become most severe when all threads are reading from and writing to locations that are very close in memory. In Lattice QCD and in tmLQCD in particular this occurs when a datum at a given lattice site affects another at a different lattice site, such as in the application of the hopping matrix or any derivatives that are computed during the molecular dynamics step of Hybrid Monte Carlo. Looking at these two examples, it will be shown below that the issue is intertwined with the necessity of mutual exclusion for many operations in a multi-threaded program.

There there exist strategies to reduce or eliminate both cache contention and false sharing, however, time constraints have allowed only a small subset to be implemented in tmLQCD because the required structural changes were simply excessive at this time. Competition between threads can be eliminated once it is understood exactly how many cycles each operation for every given thread consumes and by optimizing the movement of data with respect to the ordering of any other operations. This has been achieved, for example, in the BAGEL [202] assembler generation library which generates highly optimized assembler code for computational kernels common in Lattice QCD. False sharing can be eliminated by ensuring that multiple threads always write to memory that is sufficiently well separated to avoid invalidations. This can be achieved by using a thread-optimized data layout with domain decomposition (in terms of data layout, not in terms of multi-grid iterative solvers) depending on the number of threads as well as the addition of padding for data structures which do not fill a cache line completely. Both solutions suffer from added complications: the data layout is affected by the neighbourhood relationships of different lattice sites as well as the boundary conditions of the problem and too much padding can reduce the total memory bandwidth and interfere with optimisations for increasingly long single instruction multiple data vector units on modern processors. An implementation of Lattice QCD kernels which takes into account the very complicated restrictions on the Intel® Xeon Phi™ architecture has been presented in Ref. [203]. In order to offload the burden of measuring and understanding various architecture dependent overheads, work is currently under way to leverage the computational kernels from these code generators in the tmLQCD software suite, although the work is currently still in its early stages.

Reducing Cache Misses and Thread Idling

A cache miss occurs when the prediction algorithms on a processor are unable to load a datum into a low-level cache before it is actually required, thus incurring the full performance penalty for loading data directly from main memory. When a cache miss occurs in a multi-threaded program, one

or multiple threads may idle waiting for data from memory and depending on the operations to be carried out, no other threads may be available to take over at this moment. Reducing cache misses is important also in single-threaded programs, but the above-mentioned cache contention and false sharing potential make it crucial and difficult to deal with in multi-threaded ones. Finally, optimisations which may not be required for single-threaded execution can become important when multiple threads are operating, as will be discussed now.

A particular example from the tmLQCD code-base concerns the implementation of the full spinor hopping matrix in which the communication of the boundary terms of the input spinors is done before any computation and all four spin components are communicated. While this restricts performance on distributed-memory machines, the fact that only one volume loop is required potentially allows for relatively small OpenMP overhead. This property makes the full spinor hopping matrix an ideal candidate for execution on a single high-powered multi-core processor or a single accelerator card, when communication overhead is not an issue because no communication is done.

While introducing OpenMP into this function it was noticed that performance was far below expectations, reaching only about 20% of peak on a dual socket, four core Intel™ Xeon™ X5560 machine while running with 8 threads, significantly slower than the corresponding half spinor version with on-node communication between 8 MPI tasks. In tmLQCD, the neighbourhood relationships between different lattice sites are stored in a global array and careful analysis using the Intel® VTune™ performance tool revealed that cache misses for this array were responsible for extensive thread idling. When the lattice index of, say, the neighbour into the negative "2" direction of lattice site x was required, the compiler was able to correctly pre-fetch the necessary datum into cache. However, when the neighbour of a neighbour was required, the pre-fetching mechanism failed.

To resolve this issue a new global array `g_hi` was introduced which holds the neighbourhood relationships required for all 16 operations per lattice site in the full-spinor hopping matrix in sequential order. By using this array, all cache misses were removed and performance increased to almost 50% of peak which is quite exceptional and outperforms the corresponding half-spinor implementation by about 25%, as can be seen by comparing the darker bar (1) with the light bar (8) in Figure 10.1 in Chapter 10. This example shows how important the careful study of overheads can be to achieve good performance and as mentioned before, the situation will only become more complicated as parallelism increases.

Summary

This section dealt with the identification and reduction of overheads related to the multi-threaded execution of various parts of the tmLQCD software suite. Particular attention was devoted to thread management overheads which were found to be generically large through profiling using SCALASCA and other performance tools, especially on BG/Q. First of all, the usage of `include` statements and pre-processor directives allows the generation of multiple versions of certain routines such as the hopping matrix combined with various simpler functions in combinations that occur frequently, such as the addition of the twisted mass term or multiplication with a gamma matrix. In this way, the significant overhead of computation-

ally simple routines can be reduced without code duplication. Secondly, by employing data layouts which take into account the number of threads and the underlying architecture, overheads stemming from cache contention and the potential for false sharing can be significantly reduced. Finally, the topic of cache misses and resulting thread idling was discussed with the specific example of the full spinor hopping matrix, for which the addition of a precomputed set of pointers was able to significantly increase performance on a particular architecture.

It can be concluded that multi-threading entails many overheads which must be properly identified and maximally reduced to obtain good performance, evidence for which will be given in Chapter 10. The relative importance of various overheads depends on the architecture and as such, lattice QCD codes on modern supercomputers are likely to require many specialised versions in the medium to long term future, complicating an already difficult part of the computational aspect of LQCD. One way of achieving many of the overhead reductions discussed above in general is via the usage of code generators, as presented in Refs. [202] and [203]. Because the burden of overhead reduction is going to become worse as supercomputers become more complicated, an effort is currently under way to utilize the optimized computational kernels from these code generators in the tmLQCD software suite but this is still at a very early stage.

LOCKING AND DATA RACES

The most difficult problem encountered when dealing with multi-threading is that of data locality as well as synchronization and the resulting need for mutual exclusion in concurrent operations. It is easy to imagine that when multiple threads read and write from and to the same memory location, the order of the operations can affect the result since a reading thread might inadvertently read an intermediate value of a given memory location. For this reason, it is necessary to prevent concurrent access to memory which can potentially be accessed bidirectionally by multiple threads. A further classification of the problem is whether multiple threads read while one thread may write or multiple threads read and write at the same time. The issue is particularly complicated by the fact that it may not always be syntactically obvious that a given code segment will produce reads and writes to the same memory location. In addition, the detection of the consequences of these conflicts as deviations from some expected result is complicated by the fact that they occur probabilistically: their occurrence increases with the amount parallelism and so they may go completely unnoticed even in comprehensive tests if only a small number of threads is operating.

The neighbourhood relationship of a four dimensional lattice implies that one site may have up to eight neighbours depending on the boundary conditions. In operations such as the next-neighbour hopping matrix or the computation of derivatives with respect to the gauge field in the molecular dynamics step of Hybrid Monte Carlo, loops over the volume produce the most problematic pattern of multiple threads reading and writing to the same memory locations from multiple instructions in different source-code lines. In the sections below the available mutual exlusion mechanisms of OpenMP will be discussed and their usage in tmLQCD presented.

Critical and Atomic Sections

The most important mutual exclusion mechanisms provided by OpenMP are the `omp critical` and `omp atomic` directives. The former forces an enclosed code section to be executed by at most one thread at the same time and is generally safe and suited for longer code sections. The latter is intended for *atomic* updates of a single memory location given in the line immediately following the directive, where atomic means that the operation must be reducible to a single hardware instruction on an elementary type. An example of such operations is given in Listing 8.8, where it must be noted that for the second operation, the pointers `real_field` and `other_field` must refer to *distinct* memory regions.

Listing 8.8 Atomic updates of memory locations in OpenMP.

```
[...]
  extern unsigned int * indices;
  extern double * real_field;
  extern double * other_field;
  unsigned int count;
#pragma omp parallel for
for(unsigned int x = 0; x < N; ++x) {
  #pragma omp atomic
  ++count;
  #pragma omp atomic
  real_field[indices[x]] += other_field[x];
}
[...]
```

The most prominent examples in tmLQCD where such operations are necessary are the derivatives with respect to the gauge field in the molecular dynamics step of the Hybrid Monte Carlo simulation. The next-neighbour nature of the (twisted mass) Wilson Dirac matrix (Equation (1.36)) implies that the derivative with respect to the gauge field $U_\mu(x)$ has contributions to the force at lattice sites x and $x - a\hat{\mu}$. When a clover term is present, the force contribution is even more non-local.

In order to maximize cache re-use, the force computation iterates over the gauge field rather than the derivative field, necessitating mutual exclusion when the latter is updated. The derivative field is an anti-hermitian traceless 3×3 complex matrix for each lattice site and direction, but it is stored in the form of eight real parameters which are exponentiated with the Gell-Mann matrices when the gauge field is integrated at each molecular dynamics step. At first, in order to guarantee safety, the updates of the components of the force field were implemented as critical sections updating either all eight components or each component separately. Unfortunately, this resulted in intolerable overhead and cannot be recommended as a strategy in a lattice QCD codes. As a consequence, the mutual exclusion was thus implemented using the `omp atomic` directive in the form given in Listing 8.9.

There are two reasons why this form may not be thread-safe. Firstly, the usage of functions on the right hand side is not strictly allowed in this construct, although it seems that none of the compilers on any of the architectures that were tested produce conflicts here. The second reason is more subtle and debated: when multiple reads and writes occur to the same memory location from multiple threads and originate from atomic

Listing 8.9 Atomic updates of the derivative field components 1, 2 and 3 at lattice site x in direction μ. u is a 3×3 complex matrix and c is a real constant.

```
[...]
# pragma omp atomic
derivative[x][mu].d1 += c*(-cimag(u.c10)-cimag(u.c01));
# pragma omp atomic
derivative[x][mu].d2 += c*(+creal(u.c10)-creal(u.c01));
# pragma omp atomic
derivative[x][mu].d3 += c*(-cimag(u.c00)+cimag(u.c11));
[...]
```

statements on multiple source lines, the absence of data races cannot be guaranteed [204], despite the fact that the OpenMP specification seems to imply otherwise. As will be shown in Chapter 9, however, even very long simulations did not display any indication of conflicts within errors, at least for the observables that were tested. Finally, it should be noted that the chosen implementation clearly does not produce an atomic update of the entire derivative field for a given lattice point and direction. It is therefore imperative that the derivative is not used while being updated in this manner, which is luckily the case in tmLQCD but which might no longer be true if some clever new form of domain decomposition is used for integrating the equations of motion.

Push-style Algorithms with a Halo Buffer

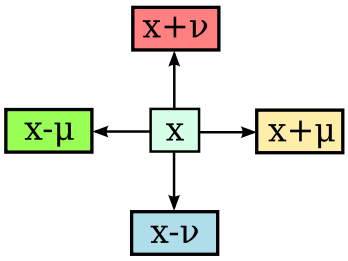

Figure 8.6: Pictorial representation of the concept of a *halo buffer*. Some function $f(x)$ maps x into its neighborhood, $x \pm \mu$. In 2 dimensions, each lattice point x can be equipped with $n = 4$ exclusive regions of memory to store the results of $f(x)$, ensuring that this memory is only ever written to once and foregoing the need for mutual exclusion in multi-threaded applications.

The above implementation of the derivative of a pseudo-fermion term is an example of what will be reffered to as a "push-style algorithm" in which data from one memory location affects data at one or more other memory locations. Because it is often computationally beneficial to load this data only once, multiple writes to the data in the other memory locations will occur. As was discussed before, multi-threading forces a balance between the overhead of mutual exclusions in this type of algorithm and any performance benefits that are gained from this non-locality through cache re-use. Although using the `omp atomic` directive seems to produce correct code with the considerations outlined at the end of the last section, there is still some overhead associated with it.

An alternative to locking which can be attractive for multi-threaded code is the introduction of a halo buffer, which is exclusive memory allocated for all data affected by a given datum, a two-dimensional illustration of this is given in Figure 8.6. As an actual example and because the nomenclature will be used in the next section, the half-spinor implementation of the hopping matrix will be discussed in the following. With the Wilson parameter set to $r = 1$, the hopping part of the twisted mass Dirac operator is given by Equation (8.1)

$$\phi(x) = \sum_{\mu=1}^{4} \left[U_{x,\mu}(1+\gamma_\mu)\psi(x+a\hat{\mu}) + U^{\dagger}_{x-a\hat{\mu},\mu}(1-\gamma_\mu)\psi(x-a\hat{\mu}) \right]. \quad (8.1)$$

The projection $p^{\pm}(x + a\hat{\mu}) = (1 \pm \gamma_\mu)\psi(x + a\hat{\mu})$, for example, results in a four-component spinor with only two independent components, as $p^{\pm}_3(x + a\hat{1}) = \pm i p^{\pm}_2(x + \hat{1})$, for instance. Consequently, significant savings in

communication can be made by communicating only two components and carrying out the reconstruction to four components and multiplication by the gauge field after the communication. In the implementation in tmLQCD, eight two-component complex double halo buffers for each lattice point are allocated to store the result of what is called the *pre* or *projection* step which involves one loop over the lattice volume. The contents of these buffers are communicated if they are on the boundaries of the local lattice, after which a second loop over the volume serves to reconstruct the full neighbourhood spinors and to calculate final result of the application of the hopping matrix in the *post* or *reconstruction* step.

Since all the buffers are non-overlapping, it is not possible for multiple threads to read from and write to the same memory location, thus removing the need for locking in a multi-threaded implementation of the function. The overhead coming from the second volume loop is small compared to the gain and this continues to be true in a multi-threaded implementation, as will be shown in Chapter 10. There are several functions in the tmLQCD software suite for which the multi-threaded implementation could benefit from the introduction of halo buffers, but the performance impact of these potential changes has not been investigated yet. An example are the derivatives of the pseudo-fermion parts for the various fermionic monomials discussed in the previous section. However, unlike for the half-spinor hopping matrix, the changes would require code duplication because when the routine is executed without multi-threading, the resulting overhead would be unacceptable.

Summary

The issue of overheads is a very important consideration when dealing with multi-threaded code because making it perform well requires considerable effort. During the development process a major overhead was found to come from thread management at the beginning of parallel sections. This resulted in a strategy which employs pre-processor constants to generate multiple versions of the hopping matrix function combined with a number of simple operations with which it frequently appears in sequence and which were normally done in separate function calls, significantly reducing this overhead. Further, the usage of orphaned directives was explored and some preliminary benchmarks are presented in Chapter 10 which indicate that their usage is a very effective technique for reducing overheads, but their deployment requires careful planning and may interfere with some of the design ideas laid out in Section 8.2. In addition, the role of cache misses was studied and a simple modification was introduced to the way that indices are stored in the full spinor hopping matrix, leading to a significant performance improvement on certain hardware, showing just how important it is to measure and understand hot spots in numeric applications like lattice QCD. It must also be understood that these are entirely dependent on the underlying architecture, as was demonstrated by looking at the `scalar_prod_r` function on BlueGene/Q™ and SuperMUC. This view is reinforced in the next section which deals with the interplay of shared and distributed memory parallelism and the prospect of overlapping communication and computation for the purpose of hiding or reducing the communication overhead.

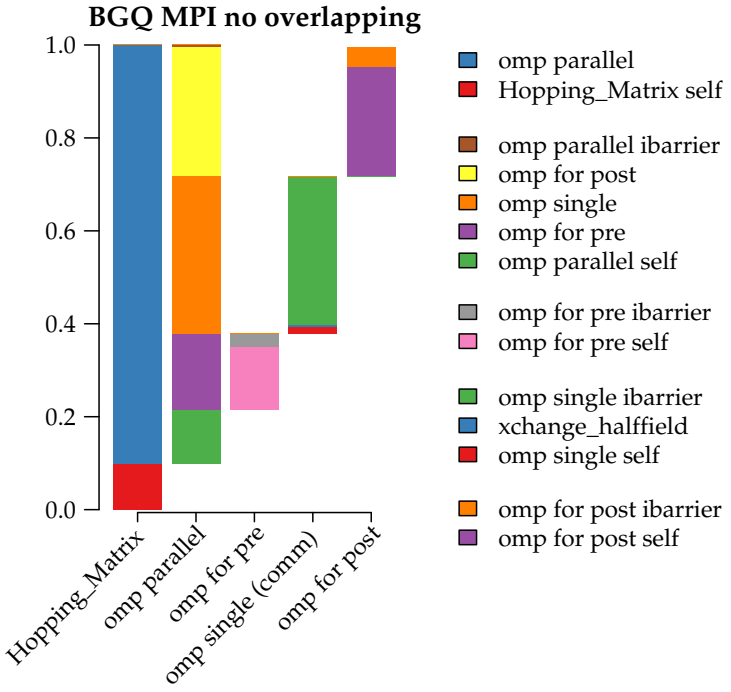

Figure 8.8: SCALASCA profile of the non-overlapping half-spinor implementation of the hopping matrix on BlueGene/Q™ on 32 nodes using 1 MPI process per node and 64 threads per process with the ideal local volume of $V_l = 12^4$ lattice sites. Communication is done entirely using portable MPI.

OVERLAPPING COMMUNICATION AND COMPUTATION

As a final and very important topic, the overlapping of communication and computation will be discussed in this section. Distributed memory parallelism requires explicit communication of the boundary or *surface* terms while operations on the body or *bulk* of the domain are fully local to one process. The size of the surface depends on the details of the algorithm and in the case of Lattice QCD with Wilson fermions, this usually corresponds to one or two lattice points in every lattice dimension plus some overlapping ones from the corners. The overhead of the communication step in a parallel computation depends on the available bandwidth and the amount of data that needs to be communicated. In order to minimize this overhead, it would be beneficial to compute the boundary terms, start the communication and at the same time do computations which are independent of the boundary terms. When these are completed, the program would wait until the communication completes and compute any remaining work-load. This approach is even more attractive in light of the fact that on modern machines a lot of the actual communication is offloaded to the networking hardware, thus freeing up the CPU to be otherwise occupied.

Unfortunately, it seems that many MPI implementations do not natively support overlapping communication and computation in single-threaded applications. Even though they offer non-blocking point-to-point communicators such as `MPI_Isend` and `MPI_Irecv` and some benefit is derived from being able to queue communications in all directions in one step, the actual communication still takes place in the subsequent `MPI_Wait` or `MPI_Waitall` calls which are of course blocking. This is shown in Figure 8.8 which shows the relative time spent in different parts of the half-spinor hopping matrix function (see Equation (8.1) and discussion there) on BlueGene/Q™ in a benchmark on 32 nodes using 1 MPI process per node and 64 threads per MPI process. A description of how the figure is to be read is given in Section 8.3.1.

The hopping matrix function is entered and a parallel section launches the threads after which about 20% of the total time has already been spent on tasks unrelated to the actual computation (`Hopping_Matrix self` and `omp parallel self`). When the threads have been spawned, the `pre` loop computes the data to be communicated (including the bulk contribution which doesn't pass through MPI) and the communication is launched in the following `omp single` section. This single thread queues all the communications and runs `MPI_Waitall` in the function `xchange_halffield`. During this time, most threads idle waiting and in the end only about 40% of the total time is spent in actual computations.

In a multi-threaded implementation it is clear that this is not ideal because the thread idling can be avoided. After changing the data layout to separate the body from the surface, an implementation with three loops and two single sections can be constructed as shown in Figure 8.9a. The function is entered and the threads are launched after which the data to be communicated is computed in the `omp for surface` section. Subsequently the communication is queued by `xchange_halffield` in a `omp single nowait` section which has no implicit barrier and thus very low overhead, allowing $n-1$ threads to be otherwise occupied. The next loop computes the bulk contribution, after which the next `single` section calls `MPI_Waitall`. Unfortunately, as becomes clear comparing Figure 8.8 and Figure 8.9a, a lot of thread idling occurs in this single section, indicating that the communication is actually only started once `MPI_Waitall` is called. Although the relative share of time spent in computation has increased to about 50%, the added overhead from having three loops make this only marginally faster, as will be seen in Chapter 10 in Figure 10.4.

An alternative implementation which attempts to eliminate this thread idling is shown in Figure 8.9b. Here, after the data to be communicated has been computed, the `omp single nowait` section queues all communication and immediately calls `MPI_Waitall`, thus actually starting the communication. In this time, the remaining threads already work on the bulk computation in `omp for body` and the communicating thread is joined by the implicit barrier in this loop before the `post` step is executed. This has the benefit of avoiding the overhead of a second `omp single` section, but it is insufficient to really improve performance. Unfortunately, while having the lowest overhead, the default `static` scheduling for the `omp for` work-sharing directive means that the communicating thread is assigned the same work-load as the other threads, and his results in a significant amount of time spent waiting for this thread after it has completed the communication.

A final improvement can be made to this implementation by modifying the scheduling for the bulk computation loop. OpenMP offers four types of work scheduling, a review with measurements of overheads on different architectures is given in Refs. [195, 196].

- `static` scheduling in which each thread is given the same amount of work with the lowest overhead
- `static,n` scheduling in which each thread is given n iterations of a loop from a first-come first-serve queue, the overhead of which increases with decreasing chunksize
- `dynamic,n` scheduling which allocates chunks of size n dynamically to any threads which are not occupied and has the highest overhead

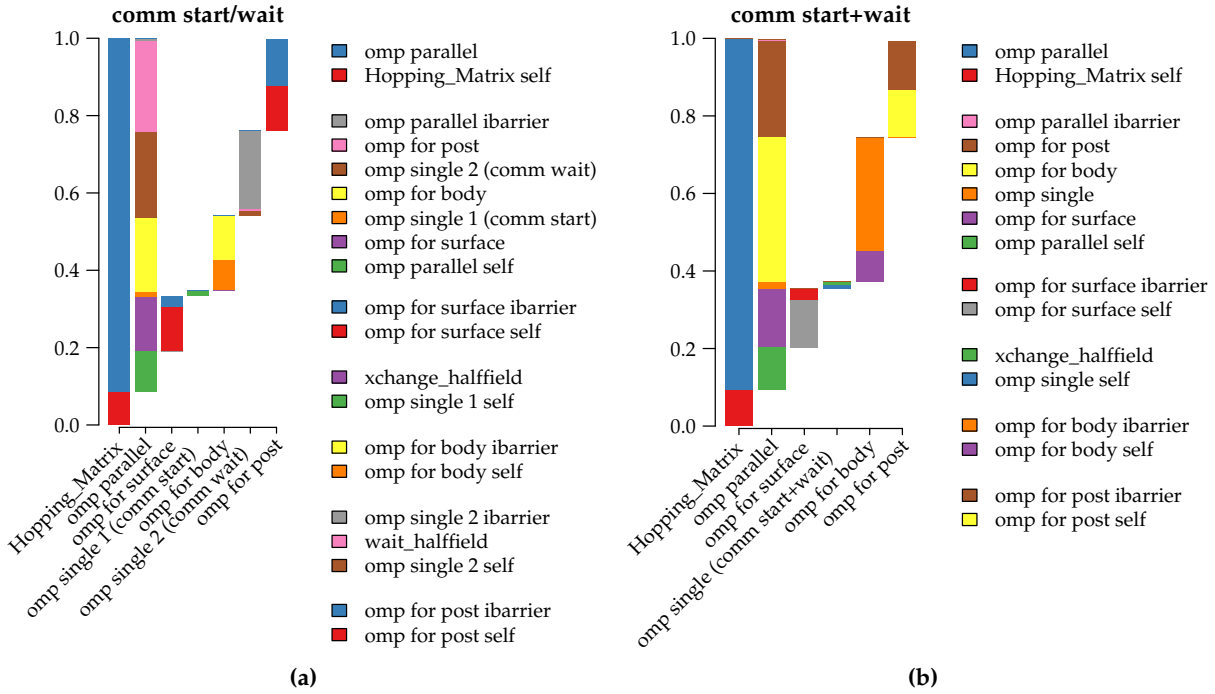

Figure 8.9: SCALASCA profile of time spent in the sub-parts of the half-spinor hopping matrix in the tmLQCD benchmark application. The legend gives the division of each bar where the qualifier "self" refers to time spent in a given part which is not spent in any of its sub-parts. Panel (a) shows the distribution in the function where non-blocking communicators are called before and `MPI_Waitall` is called after the body loop. Panel (b) shows the situation when the same thread is used to start and wait for communication, with implicit synchronization at the end of the body loop.

- `guided,n` scheduling which allocates work in an exponentially decaying fashion up to a minimum chunksize n on a first-come first-serve basis, implying the first thread gets the most work while subsequent threads are given less

After some amount of testing it was found that guided scheduling with $n = 32$ delivered the best performance for the situation given in Figure 8.9b when MPI is used for communication.

On BlueGene/Q™ there is an alternative communication library in the form of the IBM® MU SPI, which can be used to write directly to memory on a remote node (also called remote direct memory access (RDMA)). Although this can also be achieved in a portable way with MPI using `MPI_Get` and `MPI_Put`, the structure of MPI RDMA makes it hard to use for overlapping communication and computation and the overheads are rather large. It will be shown in Chapter 10 that SPI, unlike the non-blocking MPI communicators discussed above, starts communication immediately and has very low overheads. These two properties mean that when SPI is used for communication, the implementation with two `omp single` sections is faster than the one using `guided,32` scheduling in the body computation and one `omp single` section for communication. It is clear that all these options depend strongly on the problem and the architecture and therefore a judicious choice is only possible after careful benchmarks.

Finally it should be mentioned that another alternative implementation is possible, some ideas of which were presented in Ref. [205]. Here the communcations are started as soon as the data becomes available, rather

than computing the whole surface before the communication is started. The implementation in Ref. [205] is incomplete, however, in that it also only uses `MPI_Isend` and `MPI_Irecv` which, as was shown above, only queue the communications. It would be interesting to pursue this kind of strategy with SPI to see if the asynchronous queueing of smaller amounts of data and the larger number of `opemmp single nowait` sections would be able to reduce the significant amount of time spent waiting for all threads to rejoin at the implicit barrier in the body computation in Figure 8.9b.

SUMMARY AND CONCLUDING REMARKS

In this chapter the addition of multi-threading via OpenMP into the tmLQCD software suite was discussed. Specific implementation details such as conditional directives and multi-threaded Kahan summations stand out compared to other lattice QCD codes which provide OpenMP support such as the Columbia Physics System [206] or the MILC suite [207]. Also in contrast to these codes, almost the entire tmLQCD code-base has been multi-threaded and thus prepared for the current and next generation of supercomputers with ever-increasing numbers of functional units.

In this work, overheads associated with the use of threads were considered on multiple architectures with a specific focus on IBM® BlueGene/Q™ and improvements were made where possible, guided by measurements using SCALASCA and other performance tools. The most prominent potential improvement, the usage of orphaned OpenMP directives, was tested, but a full deployment has not been carried out on account of the significant complications with regards to data locality and the potential loss of decoupling between function interfaces and their implementations. Further, the issues of locking and data races were studied carefully and working high performance solutions were found for the computation of derivatives in the molecular dynamics and the full and half spinor hopping matrix implementations. In addition, the interplay of shared and distributed memory parallelism was considered with a specific focus on the possibility of implementing true overlapping communication and computation in the hopping matrix kernel in the hope of overcoming some limitations of non-blocking MPI communicators and in order to hide communication overheads. Several implementations of an overlapping hopping matrix were tested and two were deployed. The first uses MPI communicators and tuned thread scheduling to allow some communication overhead to be hidden on all kinds of MPI machines while the second uses a more naïve approach with SPI communication on IBM® BlueGene/Q™.

The significant number of changes to the code-base and the potential for subtle bugs from multi-threading as a result of the work described in this chapter and other development work required careful testing procedures to be established and carried out. These tests and their results will be discussed in Chapter 9 while the performance of the new implementations of various parts of the code-base will be laid out in Chapter 10.

9 CORRECTNESS TESTS

In the preparation of this thesis, changes to about 60% of the source code of the tmLQCD software suite were made. First, a transition from a hand-crafted complex data-type to the standard C99 `complex` type affected almost every single code line. Implementations of the twisted mass clover operators and actions as described in Appendix A were added. In addition, a version of the RHMC [138] for all the actions available in tmLQCD was introduced, chiefly however to simulate the non-degenerate twisted mass clover doublet. Finally, significant changes were caused by the addition of specialised code for BlueGene/Q and the introduction of OpenMP, but especially by the latter because almost every routine was parallelised.

As a consequence of these developments, a large-scale testing programme was embarked upon to ensure continued correctness. It is the aim of this chapter to present the methods that were developed for this testing programme and to demonstrate their wide applicability. Because a general discussion of possible problems in Lattice QCD codes is far beyond the scope of this work, this will be done on the basis of a number of specific examples. A particular emphasis is placed on the fact that the techniques presented here can uncover issues which may otherwise be very difficult or completely impossible to detect. It should be noted that the tests that were conducted go well beyond what is usually done for LQCD codes. This is especially important in light of the precision desired for new LQCD results to be phenomenologically useful.

Early on in the testing programme, it was realised how valuable high statistics Hybrid Monte Carlo runs can really be, despite their apparent simplicity. Insights about the correctness of the HMC implementation can be gained through short sample trajectories by tracking the energy violation using different code paths to simulate the same physics. However, the chaotic nature of the equations of motion and the added non-determinism when threads are used necessitates a more rigorous testing framework, lest severe but otherwise undetectable bugs be missed. This will will be shown in Section 9.1, demonstrating how a number of issues were found even though they passed every other test and did not produce any suspicious behaviour.

The implementation of the operators and RHMC for the non-degenerate twisted mass clover doublet was tested rigorously. Even though the above-mentioned high statistics runs did not point to any issues, first large-scale simulations with $N_f = 2+1+1$ twisted mass clover quarks had puzzling acceptance problems, culminating in a complete loss of acceptance as the volume was increased further. This prompted a further set of tests described in Section 9.2, chiefly an approach in which the derivative of the action contribution of the non-degenerate doublet was computed numerically and compared to the implementation of the analytical derivative. This was complemented by a further check presented in Section 9.3, in which the behaviour of the plaquette expectation value as a function of the mass splitting parameter μ_δ was studied. Specifically, the expected slope $\partial\langle P\rangle/\partial(\mu_\delta^2)$ was computed as a valence observable and compared to that obtained from direct simulations with a number of different μ_δ.

HIGH STATISTICS RUNS

The tmLQCD package is a highly versatile code-base for simulations with Wilson fermions which contains support for many different types of mass-degenerate and mass non-degenerate fermion monomials with and without twisted mass as well as with and without a clover term. It contains a number of different simulation algorithms such as HMC, RHMC and PHMC for configuration generation and supports parallelisation via MPI and OpenMP. The parallelisation is quite flexible as the lattice volume can be distributed in up to four dimensions and a number of optimisations for different computer architectures are available. As a result, tmLQCD can be run in many different ways and the interconnectedness and interdependence of the various parts of the code-base lead to a large number of potential points of failure.

There exist a number of simple ways for testing large portions of the code-base for correctness. First of all, the implementations of the Dirac operators and their parallelisation can be tested by doing inversions and comparing the results against other code-bases, if available. The different parallelisations can be compared against each other until one is convinced that any remaining differences are due to round-off errors only. In addition, different Dirac operators can be used in the limits where they are equivalent, thus allowing a comparison when no separate code-base exists. Of course, this kind of test may hide issues because they may only occur outside of those limits. A strong test of this kind is the computation of the real residual between the even/odd preconditioned Dirac operator and the non-preconditioned one.

For configuration generation, the acceptance step of the HMC algorithm provides a relatively strong test of the consistency of the action and the derivatives that are necessary for molecular dynamics. While this is true, it must be noted that the energy violation is an extensive quantity and is often insufficient in magnitude to show problems on the small lattice volumes that are usually used in testing, an issue which will be discussed in more detail in Section 9.2. In addition, subtle issues with the parallelisation or random numbers may not show even in production, despite having possibly significant effects on the expectation values and statistical errors of certain observables. This is because the biases introduced by such issues may only be visible if results from simulations with different parallelisations are compared against each other in the high statistics limit, which of course is not usually done.

One very effective way for checking for these sorts of problems is the usage of high statistics tests. These consist of very long runs of the HMC algorithm ($\mathcal{O}(10^4)$ trajectories or more) for different physical situations, different mass preconditioning and different parallelisations on moderate lattice volumes (8^4 or larger, say). They can be run on commodity hardware and allow the determination of benchmark observables, the plaquette expectation value $\langle P \rangle$ for example, to very high precision. By comparing different situations which should give identical results up to round-off errors and statistical fluctuations, systematic effects can be identified. A plethora of high statistics runs was carried out after the major changes to the tmLQCD software were nearing completion. A complete discussion of all tests and the different parameters for the tests is beyond the scope of this thesis, suffice it to say that the complete set of runs contains around 130 different combinations of monomials, mass preconditionings and parallelisations. The

chief aim of this section is to use the given examples to advance the argument that high statistics runs should be an integral and regular part of the development process for LQCD codes.

Three examples of faults which may hide even in production situations were uncovered in this testing programme. Although it is clear that the exact same issues are unlikely to be uncovered in different codes, they do serve as interesting case studies of very subtle faults which can be uncovered with high statistics runs. For this reason, they are described in detail in Appendix E while only the main results will be summarized here.

The non-determinism involved in the usage of threads makes it difficult to differentiate between genuine faults and differences due to round-off errors and execution order. When data races are introduced by threading, their effect may not show at all in normal testing procedures. This is because their occurence is probabilistic and depends on the current execution context, the architecture, the number of concurrent threads and the given work load. One such issue was encountered in the computation of the MD derivative of the clover term. Only through very high statistics runs with different numbers of threads was it realised that the acceptance rate, but not the expectation value, was affected by the fault. This was well after the code-base had already been thoroughly tested.

The remaining two issues were related to the parallelisation logic of the pseudo random number generator (PRNG) employed in tmLQCD, as described in Appendix E.2. In the first case, after it was discovered, it was realised that the bug had affected two ETMC $N_f = 2 + 1 + 1$ production ensembles which had to be regenerated as a consequence. During production, it would have perhaps been possible to see the issue, had the parallelisation or architecture been changed in the middle of the production process. In the worst case scenario, however, it could have been simply put off as a thermalisation effect. In certain situations, different MPI processes would reuse random numbers that had already been used elsewhere for a different purpose. The only easily visible effect of the bug was to significantly increase autocorrelations between the generated gauge configurations as the number of processes was increased. After it was fixed, it was realized that the bug also introduced a significant bias (in the high statistics limit) into the expectation value of the plaquette. This kind of problem is completely undetectable without high statistics runs with different parallelisations, which explains why it was able to remain undiscovered for so long. The second issue of this type was actually introduced when the above problem was fixed. It involved the computation of the PRNG seed for each MPI process and only very mildly affected the high statistics runs done for the previous problem. A specially designed high statistics run with 8 quark flavours significantly amplified the effect that it had on the expectation values. This even allowed the exact point in the logic to be identified because a pattern emerged in the bias that the bug introduced.

Based on experience with the case studies discussed in Appendix E, it has become the present author's opinion that high statistics runs are an excellent type of *integration test* which should be a central part of the software development process for LQCD applications. Special purpose setups can be used to pinpoint suspected issues, as was done for the problem involving the PRNG seed. More generally, high statistics runs should be run thoroughly and regularly and checked against known results as modifications are made to a code-base or new features are added. In this process, it is especially

important to have as wide a variety of tests as possible because more often than not, differences between situations which should be identical hint at problems with the code-base. Because of the sheer number of combinations of different parameters, parallelisations and architectures, the process should be automated. To a limited extent, this kind of automation was attempted for tmLQCD by the present author in [C7, C8], although the framework is currently implemented only for one type of machine and should be generalised at some point in the future. Sometimes high statistics runs may hint at problems, but it may be unclear where exactly these originate. In the next section another type of testing framework will be presented which focusses on ensuring that derivatives in molecular dynamics and the simulated actions are implemented consistently.

NUMERICAL DERIVATIVE

The HMC algorithm requires the integration of the equations of motion of the gauge field and conjugate momenta in its molecular dynamics step and this generally constitutes the most complicated part of configuration generation in lattice QCD, both conceptually as well as computationally. Two features of the system and its equations of motion make testing the correctness of the integration algorithm quite complicated. Firstly, one is dealing with a chaotic system [192] in the sense that small differences in initial conditions quickly lead to divergent integration trajectories. This makes it difficult to differentiate between bugs and the consequences of simple round-off errors, therefore requiring testing to be done in conditions which minimize round-off. Secondly, the energy of the system and the energy violation from finite step-size integration are extensive quantities, making the acceptance test an unreliable measure of correctness in small volumes on which test runs are usually performed.

In light of these aspects, it is important to have a framework which allows the implementation of the integration algorithm to at least be checked for consistency and potential bugs to be traced to specific sections in the code. This can be achieved practically by computing the derivative of the action with respect to the gauge field numerically and comparing it against what can be derived analytically and computed directly. It will be shown that although the idea is straightforward, there are some subtle issues which need to be considered. After the theoretical background of the method and its implementation have been presented, a particular bug will be discussed which was identified using this method after being mistaken for a fundamental problem of the simulated system.

Motivation & Background

Although the details of the HMC algorithm have already been discussed in Section 1.4.1, it is useful to review the updating scheme of the phase space variables in the leapfrog symplectic integration scheme, which alternates between updates of the conjugate momenta P and the link variables U.

$$\begin{aligned} \delta P_\mu^a(x,\tau) &= \left\{ \frac{\partial}{\partial \alpha} H_{\mathrm{MD}} \left[e^{i\alpha\lambda_a} U_\mu(x,\tau) \right] \right\}_{\alpha=0} \\ P_\mu^a(x,\tau+1/2\delta\tau) &= P_\mu^a(x,\tau-1/2\delta\tau) + \delta\tau \cdot \delta P_\mu^a(x,\tau) \\ U_\mu(x,\tau+\delta\tau) &= \exp\left[\delta\tau \cdot P_\mu(x,\tau+1/2\delta\tau)\right] U_\mu(x,\tau)\,, \end{aligned} \tag{9.1}$$

where H_{MD} is the molecular dynamics Hamiltonian and where half-steps of length $\delta\tau/2$ are done at the beginning and end of the trajectory. The exponential map ensures that the $U_\mu(x)$ remain in SU(3) up to round-off errors and in practice these small deviations are corrected by group projections at regular intervals.

In modern simulations, the evolution is somewhat more complicated because multiple time-scales are used for different parts of the effective action and higher order nested integrators lead to more involved updating schemes, but the basic features are the same. It is useful to note that the derivative in Equation (9.1) splits into the sum of the derivatives of the individual contributions S_i in Equation (1.93). The following discussion will focus on these individual terms. When the definition in Equation (1.96) is followed through, one obtains an expression in terms of products of gauge fields and inverses as well as derivatives of the Dirac operator. This can of course be computed directly and will be referred to as the *analytical* derivative δP with components $\delta P^a_\mu(x,\tau)$ of a given *monomial* S_i.

A second order numerical approximation to $\delta P^a_\mu(x,\tau)$ can be computed from the difference in action obtained when the gauge field at one lattice point x and in one direction μ is rotated by small angles $\pm\epsilon$ for one generator λ_a

$$\delta\mathcal{P}^a_\mu(x,\tau) = \frac{S_i[e^{i\lambda_a\epsilon}U_\mu(x,\tau)] - S_i[e^{-i\lambda_a\epsilon}U_\mu(x,\tau)]}{2\epsilon} = \delta P^a_\mu(x,\tau) + \mathcal{O}(\epsilon^3)\,. \tag{9.2}$$

Comparison to the analytical derivative can then be used to demonstrate consistency between the implementation of the action contribution S_i and the analytical derivative for a given monomial. In practice, all the numerical approximations involved in the computation of either version of $\delta P^a_\mu(x,\tau)$ must be checked first and the implementations can then be considered consistent if they differ by no more than some small number ω

$$\Delta P = \left|\delta\mathcal{P}^a_\mu(x,\tau) - \delta P^a_\mu(x,\tau)\right| < \omega\,,$$

which will be seen to depend strongly on the action being simulated and the aforementioned numerical approximations.

Implementation and Tests

In tmLQCD, the numerical derivative was implemented in a separate codebase which supports some advanced memory management features and can be found in branch `ndcloverrat_num_deriv` in the present author's public git repository[1]. The implementations of the different monomials were extended with the addition of a function for the computation of the numerical derivative as well as the optional parameters `numderiv=yes/no` and `decouple=yes/no`. The meaning of the former should be obvious while the latter can be used to turn the monomial into a *spectator*, meaning that although the derivative and action contributions are computed, the monomial does not participate in the dynamics of the system.

In this way, the derivative implementations can be compared even over many integration steps while ensuring that any bugs in the implementation do not cause instabilities in the dynamics which could be misinterpreted. The code was further instrumented with a facility for writing to disk the values of $\delta P^a_\mu(x,\tau)$ and $\delta\mathcal{P}^a_\mu(x,\tau)$ as the trajectory progresses, either for the whole volume or a randomly selected subset of lattice points.

Table 9.1: Parameters and scan ranges for tests of numerical derivative. σ_s: absolute squared CG residue in computation of energy contribution for numerical derivative. σ_f: absolute squared CG residue in computation of analytical derivative. ϵ: rotation angle for numerical derivative. All parameters were stepped through by factors of 10.

name	V	gauge	β	κ	$2\kappa a\mu$	c_{sw}	τ	N_{int}
TM	2^4	tlSym	3.8	0.17	0.01	0	1.0	100
cTM						1.1		
*c*TM*						2.5		

parameter	σ_s	σ_f	ϵ
range	$10^{-2} - 10^{-26}$	$10^{-2} - 10^{-26}$	$10^{-2} - 10^{-10}$

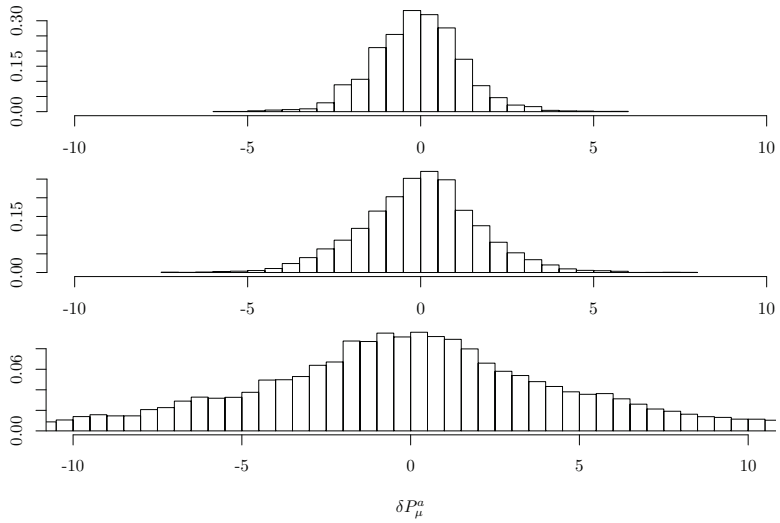

Figure 9.1: Histograms of $\delta P^a_\mu(x,\tau)$ for all x, μ, a and τ. From top to bottom: *TM*, *cTM* and *c*TM*. Note how strongly the width of the distribution depends on the value of c_{sw}.

In order to establish under which conditions the analytical and numerical derivatives can be trusted and thus compared, a number of tests were performed using monomials of degenerate twisted mass fermions with and without a clover term[2]. In addition to testing different situations, the tests also involved scans in three parameters:

- σ_f, the (accumulated) absolute squared target residual of the conjugate gradient solver in the analytical derivative computation;
- σ_s, the (accumulated) absolute squared target residual of the conjugate gradient solver in the computation of the action contributions entering in Equation (9.2);
- ϵ, the rotation angle in Equation (9.2).

The tests were carried out reproducibly along identical trajectories in which only the gauge monomial participated in the dynamics of the system. For simplicity, a leapfrog integrator was used and both the gauge and fermionic monomials were put onto the same time-scale. The parameters and the ranges of σ_s, σ_f and ϵ for the different tests are listed in Table 9.1.

Three physical situations $TM(c_{sw} = 0.0)$, $cTM(c_{sw} = 1.1)$ and $c^*TM(c_{sw} = 2.5)$ were tested and Figure 9.2 shows typical trajectories of $\delta P^a_\mu(x,\tau)$ for *TM* and *c*TM*. Both trajectories are smooth but the magnitude of the derivative is increased markedly with the value of c_{sw}. To demonstrate this and to

[1] `https://github.com/kostrzewa/tmLQCD/tree/ndcloverrat_num_deriv`
[2] these are considered to be implemented correctly

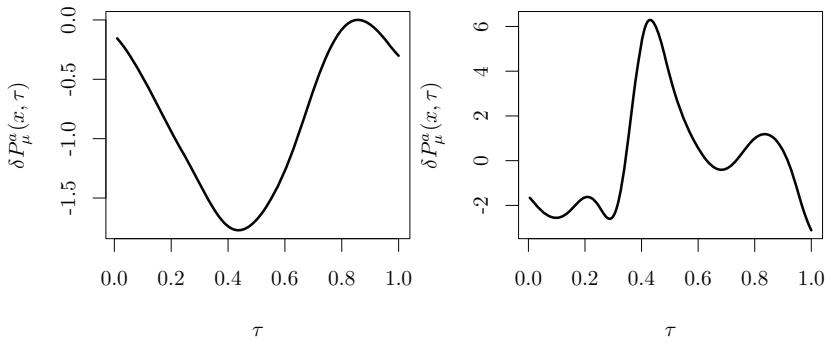

Figure 9.2: $\delta P^a_\mu(x,\tau)$ for one lattice site x, one generator a and one direction μ for two physical situations *TM* (top) and *c*TM* (bottom).

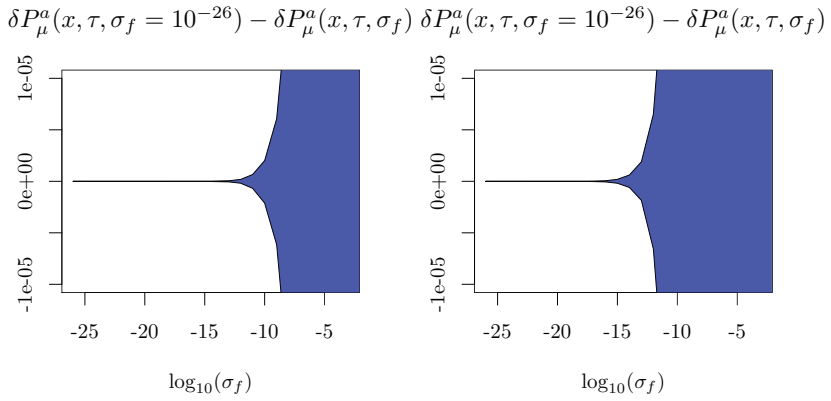

Figure 9.3: 34.27 percentiles around the median of the difference between the most precise computation of the analytical derivative ($\sigma_f = 10^{-26}$) and the analytical derivative at a given value of σ_f evaluated for all x, μ, a and τ for the two physical situations *TM* (top) and *c*TM* (bottom).

show the typical range of values, the histograms in Figure 9.1 give the distribution of the values of the derivative for all lattice points, all directions, all generator indices and all time-steps. It is clear that the width of the distribution depends crucially on the value of the clover coefficient and that it could be expected that large values of c_{sw} may lead to instabilities unless the time-step is reduced appropriately.

The results of the scans in σ_f are shown in Figure 9.3 for physical situations *TM*($c_{sw} = 0.0$) and *c*TM*($c_{sw} = 2.5$). The blue band depicts the 15.73 and 84.27 percentiles of the distribution of $\delta P^a_\mu(x,\tau,\sigma_f = 10^{-26}) - \delta \mathcal{P}^a_\mu(x,\tau,\sigma_f)$ for all x, μ, a and τ. It can be seen that increasing the solver precision leads to a convergence of the median value of the analytical derivative. In addition, it appears that when the clover term is added to the action, the solver precision must be increased compared to the standard twisted mass case to obtain the same accuracy. Finally, one can conclude that for this test, the highest possible precision can be used for the analytical derivative without worrying about accidental divergences due to round-off errors, at least for the physical situations considered here.

Scans in the two remaining parameters σ_s and ϵ relating to the numerical derivative are depicted in Figure 9.4 as the maximum absolute difference between $\delta P^a_\mu(x,\tau,\sigma_f = 10^{-26})$ and $\delta \mathcal{P}^a_\mu(x,\tau,\epsilon,\sigma_s)$. A notable feature is that differences on the order 10^{-4} can be observed for the *TM* case when the solver precision is insufficient. Much more striking, however, is the fact that for *c*TM* the maximum absolute difference between the numerical and analytical derivatives can diverge severely as ϵ is reduced below around 10^{-6}, even for the highest solver precision. Given that $\sigma_s = 10^{-26}$ is on the edge of what is possible in double precision arithmetic, it seems that when c_{sw} is comparatively large, disagreements between the numerical and analytical derivatives on the order of 10^{-5} may be observed without actually

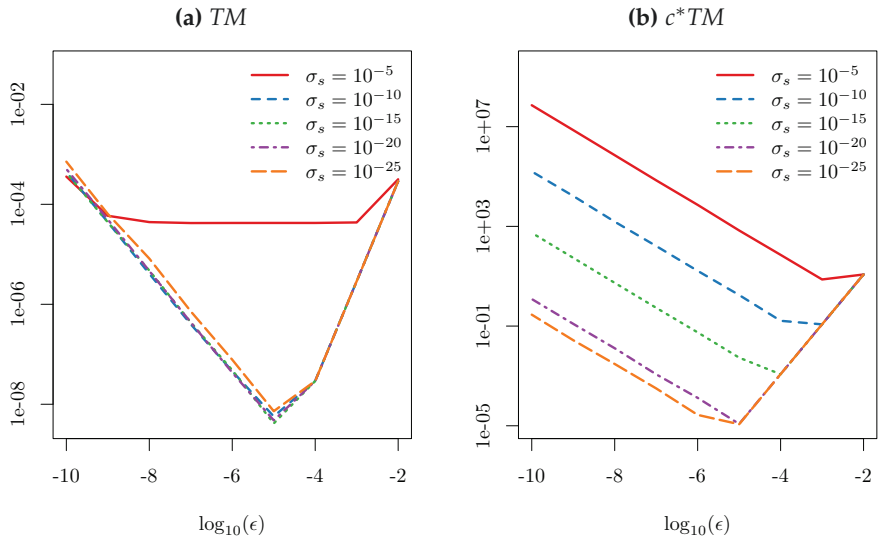

Figure 9.4: $\max[|\delta P^a_\mu(x,\tau,10^{-26}) - \delta\mathcal{P}^a_\mu(x,\tau,\epsilon,\sigma_s)|]$. Maximum absolute difference between $\delta P^a_\mu(x,\tau)$ at the highest precision and $\delta\mathcal{P}^a_\mu(x,\tau)$ as a function of σ_s and ϵ for two different physical situations.

indicating a discrepancy. In addition, if ϵ and σ_s are chosen unwisely, large differences can appear and suggest faults where there are none.

Having said that, an optimum value of $\epsilon = 10^{-5}$ seems to exist and it appears to be advisable for the solver precision to be maximized both for the analytical and the numerical derivative computations. The test of the non-degenerate clover twisted mass rational approximation monomial was thus carried out using the highest practical precision in the multiple right hand side solver and $\epsilon = 10^{-5}$. The expected tolerance of the test was taken to be around $\omega \sim 10^{-6}$, given that no unusually large clover coefficient would be used.

Results

After optimal parameters for the numerical derivative and the comparison to the analytical derivative had been established, the framework was applied to the test of the non-degenerate twisted mass clover rational approximation monomial. The test was carried out using $N_f = 1+1$ flavours decoupled from the dynamics of the system which consisted of a tree-level Symanzik-improved gauge action on a 2^4 volume with 200 leapfrog steps and a trajectory length of $\tau = 1.0$. The remaining parameters were set to the following values:

- $\beta = 3.3$, $\kappa = 0.17$, $c_{sw} = 1.0$, $a\mu_\sigma = 0.1105$, $a\mu_\delta = 0.0935$, $\sigma_f = 10^{-19}$, $\sigma_s = 10^{-19}$, $\epsilon = 10^{-5}$.

Note that the twisted mass $a\mu_\sigma$ and the mass splitting $a\mu_\delta$ were chosen such as to guarantee that the lighter of the two quarks would retain a positive bare mass parameter to not induce any unnecessary complications due to a quark determinant with a fluctuating sign.

The result of the test for one trajectory is shown in Figure 9.5 as the difference between the numerical and analytical derivatives with the bug shown in the left panel while the situation without the bug is depicted in the right one. It is clear that the inconsistency seemed to induce smooth systematic differences which disappeared into random noise once the bug was fixed. The distributions of the deviations for both situations are shown for all lattice sites and directions as well as all generator indices and all time-steps in the histograms in Figure 9.6. It is notable that despite the total loss of

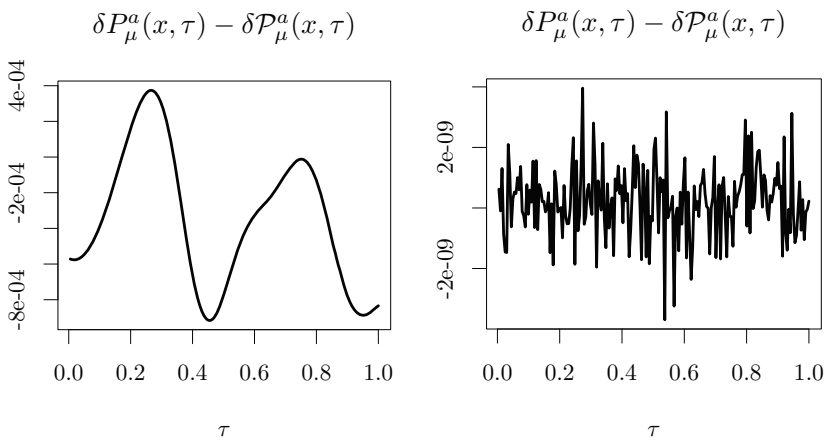

Figure 9.5: Difference between the numerical and analytical derivatives along a trajectory for non-degenerate twisted mass clover fermions. Shown at the left is the situation with the bug present while the right shows the situation after the bug had been fixed.

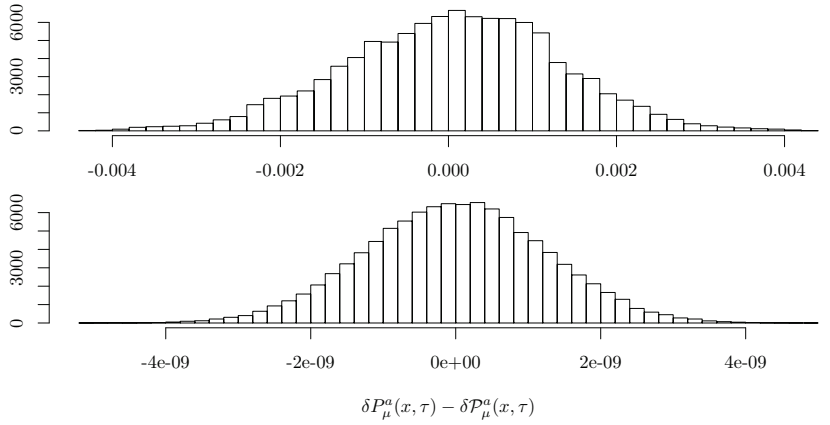

Figure 9.6: Histograms of of the difference between the numerical and analytical derivatives with (top) and without (bottom) the bug in the implementation of the non-degenerate clover twisted mass rational approximation monomial.

acceptance on large volumes with the bug present, the deviations between the analytical and numerical derivatives can be seen to be on the order of 10^{-3} or smaller, which means that an unwise choice of ϵ and σ_s could have masked the problem completely as the tolerance ω could then have been of the same order. A posteriori, the difference in the horizontal scales in the top and bottom panels of Figure 9.6 is evident and the distribution with the bug also seems to deviate mildly from a Gaussian shape.

The bug that was observed in the implementation was related to even-odd preconditioning, which splits the quark determinant into a local part which can be calculated exactly (up to round-off errors) and a part which requires evaluation via pseudo-fermions, as given in Appendix A.2. It turned out that the energy contribution of the local part for the acceptance step was implemented incorrectly and this was thus accepting and rejecting configurations with the wrong probability distribution, resulting in the loss of acceptance on large volumes. In the present test, the observed deviation between the numerical and analytical derivatives thus resulted from the numerical derivative being computed incorrectly (since S_i was not implemented correctly).

Summary and Conclusions

In this section, the numerical implementation of the derivative of the action contribution S_i of a given monomial was presented. It was shown how this can be used to check the consistency of the calculation of the energy contribution and the analytical derivative for the molecular dynamics

step of the HMC algorithm. In particular, a bug was uncovered in the implementation of the action of the non-degenerate twisted mass clover monomial, making it inconsistent with its correctly implemented analytical derivative.

It was shown that this required some level of tuning of the various numerical approximations and that differences between the analytical and numerical derivatives must be considered across many integration steps to reveal systematic effects. From Figure 9.4 it should be clear that single observations of the analytical and numerical derivatives may be misleading, because round-off errors can appear even at the order 10^{-4} for certain physical situations. In addition, the difference may even diverge if the different parameters are not chosen correctly. Further, deviations from Gaussianity in the distribution of the differences as in Figure 9.6 can serve as an indication for inconsistencies. Finally, it must be stressed again that the test described in this section only demonstrates that the analytical derivative and the corresponding action contribution are implemented consistently, but not whether the desired distribution is being simulated. In the next section a method will be presented which allows this to be done at least in principle and this will be used to demonstrate that non-degenerate twisted mass clover action is now indeed implemented correctly.

DERIVATIVES OF THE PLAQUETTE

As demonstrated in the previous section, the initial implementation of the non-degenerate twisted mass clover doublet via the rational approximation was incorrect. After the issue was identified and fixed, the test via the numerical derivative shows that the simulated action and the molecular dynamics derivatives are now implemented consistently. However, this does not mean that the action has actually been implemented correctly. In order to ensure that this is indeed the case, a further test was carried out, based on principles from Ref. [64].

Theoretical Background

The non-degenerate twisted mass clover action in is given by

$$S_h = \sum_x \bar{\chi}(x) \left[D_W[U] + m_W + i\mu_\sigma \gamma_5 \tau^3 - \mu_\delta \tau^1 + \frac{i}{4} c_{sw} \sigma_{\mu\nu} F^{\mu\nu}[U] \right] \chi(x) \,, \tag{9.3}$$

where $\bar{\chi}$ and χ are flavour doublets, τ^1 and τ^3 act in flavour space while γ^5 and $\sigma_{\mu\nu} = i/2\,[\gamma_\mu, \gamma_\nu]$ act in Dirac space. D_W is the massless Wilson Dirac operator, m_W is usually tuned to m_{crit} and $\mathcal{F}^{\mu\nu}$ is the field strength tensor in the clover definition.

This can be rewritten as

$$S_h = S_{\mu_\sigma} + \mu_\delta S_\delta \,, \tag{9.4}$$

$$\text{with } S_\delta = -\sum_x \bar{\chi}(x) \tau^1 \chi(x) \tag{9.5}$$

and where S_{μ_σ} is the degenerate twisted mass clover action with twisted quark mass μ_σ.

Adopting the notations $\int \mathcal{D}u = \int \mathcal{D}U\mathcal{D}\bar{\chi}\mathcal{D}\chi$ and $S_{deg} = S_G + S_{\mu_\sigma}$, the expectation value of the plaquette $\langle P \rangle$ is then given by:

$$\langle P \rangle = \frac{\int \mathcal{D}u\, P\, e^{-S_{deg} - \mu_\delta S_\delta}}{\int \mathcal{D}u\, e^{-S_{deg} - \mu_\delta S_\delta}}, \tag{9.6}$$

which can be expanded around $\mu_\delta = 0$ to give

$$\langle P \rangle \simeq \frac{\int \mathcal{D}u\, P \left[1 - \mu_\delta S_\delta + \frac{\mu_\delta^2}{2} S_\delta^2 + \ldots\right] e^{-S_{deg}}}{\int \mathcal{D}u \left[1 - \mu_\delta S_\delta + \frac{\mu_\delta^2}{2} S_\delta^2 + \ldots\right] e^{-S_{deg}}}. \tag{9.7}$$

Truncating the series at order μ_δ^2 and evaluating the denominator, one obtains

$$\langle P \rangle \simeq \frac{\int \mathcal{D}u\, P \left[1 - \mu_\delta S_\delta + \frac{\mu_\delta^2}{2} S_\delta^2\right] \left[1 + \mu_\delta \langle S_\delta \rangle_{\mu_\delta = 0} - \frac{\mu_\delta^2}{2} \langle S_\delta^2 \rangle_{\mu_\delta = 0}\right] e^{-S_{deg}}}{\int \mathcal{D}u\, e^{-S_{deg}}}, \tag{9.8}$$

where the notation $\langle \rangle_{\mu_\delta = 0}$ means that the expectation value is calculated using the degenerate fermion effective action. Then, evaluation of the remaining integrals gives

$$\langle P \rangle \simeq \langle P \rangle_{\mu_\delta = 0} + \frac{\mu_\delta^2}{2} \left(\langle P S_\delta^2 \rangle_{\mu_\delta = 0} - \langle P \rangle_{\mu_\delta = 0} \langle S_\delta^2 \rangle_{\mu_\delta = 0} \right), \tag{9.9}$$

where $\langle S_\delta \rangle_{\mu_\delta = 0} = 0$ and $\langle P S_\delta \rangle_{\mu_\delta = 0} = 0$ have been used. Based on this, the derivative of the plaquette expectation value with respect to μ_δ^2 is given by

$$\frac{\partial \langle P \rangle}{\partial \mu_\delta^2} = \frac{1}{2} \Bigg\{ \langle P \sum_{x,y} (\bar{\chi}\tau^1\chi)(x)(\bar{\chi}\tau^1\chi)(y) \rangle_{\mu_\delta = 0} - \langle P \rangle_{\mu_\delta = 0} \langle \sum_{x,y} (\bar{\chi}\tau^1\chi)(x)(\bar{\chi}\tau^1\chi)(y) \rangle_{\mu_\delta = 0} \Bigg\}. \tag{9.10}$$

This can now be evaluated on degenerate twisted mass clover gauge configurations as

$$\frac{\partial \langle P \rangle}{\partial \mu_\delta^2} \simeq \Bigg\{ \frac{1}{N_c} \sum_{i=1}^{N_c} (-1) \left[P \sum_{x,y} \mathrm{Tr}[S_u(x,y) S_d(y,x)] \right]_{U_i} - \frac{1}{N_c} \sum_{j=1}^{N_c} [P]_{U_j} \frac{1}{N_c} \sum_{i=1}^{N_c} (-1) \left[\sum_{x,y} \mathrm{Tr}[S_u(x,y) S_d(y,x)] \right]_{U_i} \Bigg\}, \tag{9.11}$$

where N_c is the number of configurations. The traces can be estimated using the one-end trick from the $\langle SS^\dagger \rangle$ correlation function, keeping in mind the field normalisation used in the tmLQCD software suite.

$$\sum_{x,y} \mathrm{Tr}\left[S_u(x,y) S_d(y,x)\right] \sim (1/2)(\sqrt{\kappa})^2 \cdot L^3 \cdot T \cdot \langle SS^\dagger \rangle \tag{9.12}$$

Numerical Setup and Results

The plaquette expectation value must be computed on a number of ensembles with μ_δ^2 small but non-zero, giving $\langle P \rangle$ as a function of μ_δ^2. The slope

Table 9.2: Simulation parameters for the computation of the plaquette as a function of μ_δ^2. A rational approximation of order 6 was used for the non-degenerate doublet.

code	$L^3 \cdot T$	$a\mu_\sigma$	$a\mu_\delta$	N_c
$N_f = 2$	8^4	0.2	0.0	1500
$N_f = 1+1$	8^4	0.2	$\{0.03, 0.05, 0.07, 0.1, 0.12\}$	1500

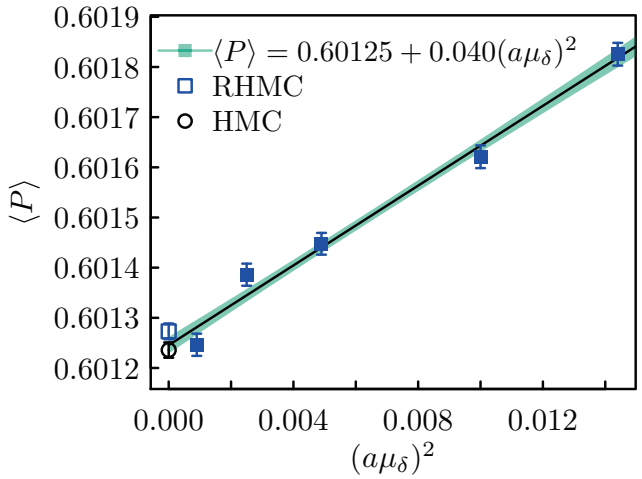

Figure 9.7: Dependence of $\langle P \rangle$ on μ_δ^2 in $N_f = 1+1$ simulations with $\mu_\sigma = 0.2$ using the RHMC. Only the filled points have been used in the fit and the empty RHMC point is a simulation with $\mu_\delta = 0.0$. The lattice size is 8^4.

thus determined should be equal to that in Equation (9.10), for which the correlation functions are evaluated on degenerate twisted mass clover ensembles. These are generated using the $N_f = 2$ twisted mass clover code-base which is well tested and which doesn't use the rational approximation. In order to guarantee that the non-degenerate determinant be positive-definite, the condition $\mu_\delta < \mu_\sigma = 0.2$ was followed and in order to ensure that P and S are sufficiently correlated with a large signal to noise ratio, a small volume $L^3 \cdot T = 8^4$ was used. The simulation parameters are listed in Table 9.2.

A simple linear fit of the form $\langle P \rangle = c_0 + c_1 (a\mu_\delta)^2$ was used to extract the slope from the $N_f = 1+1$ configurations. The error on the fit parameters is computed using simulated datasets with Gaussian distributions and the fit result

$$\langle P \rangle = 0.60125(2) + 0.040(2)(a\mu_\delta)^2 \,, \tag{9.13}$$

is shown with the data in Figure 9.7. The numerical evaluation of Equation (9.10) gives

$$\frac{\partial \langle P \rangle}{\partial (a\mu_\delta)^2} \sim 0.038(2) \,, \tag{9.14}$$

where the correlations in the data have been taken into account in a stationary blocked bootstrap analysis. This is clearly compatible with the value in Equation (9.13) and together with the results of Section 9.2 confirms that the non-degenerate twisted mass clover action is simulated correctly using the RHMC.

TEST OF THE RHMC FOR FOUR FLAVOUR SIMULATIONS

Before concluding this chapter, a peculiar observation will be checked in this section to ensure that using the RHMC as implemented in tmLQCD does not introduce unexpected biases in observables. In Figure 9.7, the simulations at $\mu_\delta = 0.0$ using the HMC and RHMC respectively are just compatible within errors. This is well within the statistical variation that can be expected from using different algorithms and different random numbers.

	$\langle P \rangle$	am_{PCAC}	$aM_{\pi^\pm}$	af_π	$\langle \delta H \rangle$	$\langle \exp(-\delta H) \rangle$	$\lambda_{min}(10^{-5})$	λ_{max}
A	0.54067(2)	0.00036(14)	0.2369(20)	0.0755(3)	0.027(3)	1.001(3)	2.904(8)	0.7616(3)
B	0.54065(2)	0.00041(16)	0.2363(9)	0.0761(3)	0.030(3)	0.998(3)	3.236(9)	0.8463(3)
C	0.54065(2)	0.00042(14)	0.2344(9)	0.0768(3)	0.025(3)	1.003(3)	3.236(9)	0.8458(3)

Table 9.3: Expectation values of observables in $N_f = 2+1+1$ simulations with different eigenvalue intervals and orders for the rational approximation. The unnormalised approximation intervals are **A**: $\lambda \in [4 \cdot 10^{-5}, 4.0]$, **B**(order 9), **C**(order 12): $\lambda \in [3.6 \cdot 10^{-5}, 3.6]$.

However, in some $N_f = 2+1+1$ test runs not discussed here, it seemed that the details of how the approximation interval for the rational approximation of the non-degenerate heavy doublet were chosen, affected the value of am_{PCAC} measured in the simulation. In particular, when the eigenvalue interval was set such that many of the measured maximal eigenvalues of the non-degenerate twisted mass clover operator would be just below the upper bound, the measured value of am_{PCAC} seemed to be higher by about two standard deviations compared to a run with a wider approximation interval.

During this particular simulation, a very small number of normalised maximal eigenvalues ($\mathcal{O}(10)$ out of several thousand) were measured to be larger than unity, violating the approximation interval. Some bias could probably be expected if this happened frequently, but the effect seemed to be too substantial, prompting three test simulations conducted on smaller lattices. This aspect is particularly important considering the fact that with the clover term present, the maximal eigenvalues seem to fluctuate much more strongly than they do without, as discussed already in Chapter 4.

The simulations **A**, **B** and **C** were set up with parameters chosen as in the first $N_f = 2+1+1$ production simulations discussed in Chapter 4 with a relatively heavy light quark mass of $a\mu_\ell = 0.01$. The parameters were

- $\beta = 1.726$, $c_{sw} = 1.74$, $\kappa = 0.140066$, $a\mu_\ell = 0.01$, $a\mu_\sigma = 0.1408$ and $a\mu_\delta = 0.1521$,

corresponding to a pion mass of around 450 MeV giving $M_{\pi^\pm} L \sim 3.4$ on a lattice of size $T/2a = L/a = 16$ and a lattice spacing around 0.1 fm. The Kaon mass is expected to be somewhat heavier than physical while the D meson is likely to be only very slightly heavier. The light quark doublet was simulated with two determinant ratios with preconditioning masses $\tilde{\rho}_3 = 0.2$ and $\tilde{\rho}_2 = 0.02$ and the simulations were carried out on four time scales with $N_t = \{1, 1, 1, 12\}$ steps of the 2MN integrator (12 steps on the coarsest time scale). The determinant was placed on time scale 1 and the determinant ratios on time scales 2 and 3 respectively while the gauge field was integrated on time scale 0. For the simulation of the non-degenerate doublet, rational approximations of order 9 (**A,B**) and 12 (**C**) were chosen and partial fractions $\{0, 1, 2, (3)_{\mathbf{C}}\}$ were integrated on time scale 1, partial fractions $\{3, 4, 5, (6)_{\mathbf{C}}\}$ on time scale 2 and the remaining partial fractions on the outermost time scale. The approximation intervals were chosen as

- **A**: $\lambda \in \left[4 \cdot 10^{-5}, 4.0\right]$ **B,C**: $\lambda \in \left[3.6 \cdot 10^{-5}, 3.6\right]$,

such that in both cases the measured *normalised* eigenvalues should be in the interval $\left[1 \cdot 10^{-5}, 1.0\right]$ for the approximation to be valid.

The mean values for a number of observables from simulations **A**, **B** and **C** from 5200 thermalised trajectories each are given in Table 9.3. In runs **B** and **C**, 5 and 3 normalised maximal eigenvalues respectively were

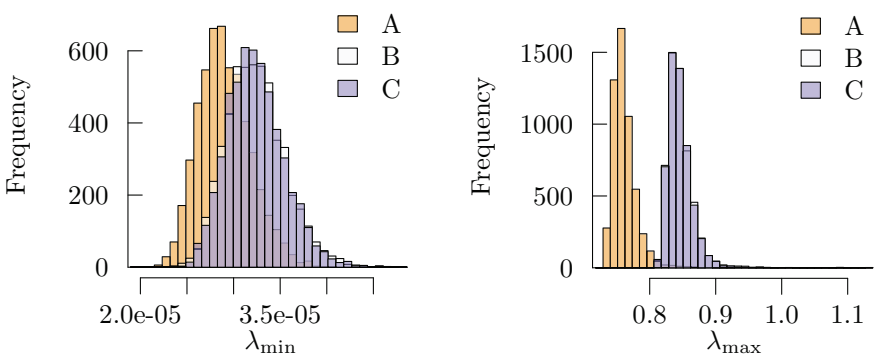

Figure 9.8: Normalised minimal and maximal eigenvalues of the non-degenerate twisted mass clover operator in $N_f = 2+1+1$ simulations with parameters as described in the text and two unnormalised eigenvalue ranges for the rational approximation. **A** corresponds to $\lambda \in [4 \cdot 10^{-5}, 4.0]$ and **B**(order 9) as well as **C**(order 12) to $\lambda \in [3.6 \cdot 10^{-5}, 3.6]$. **B** and **C** are seen to overlap, as they should.

measured to be larger than 1.0. The full distributions of the normalised minimal and maximal eigenvalues of the non-degenerate twisted mass clover operator for the given μ_σ and μ_δ are given in Figure 9.8. It is encouraging to see that the observation mentioned above was probably just a statistical fluctuation and that the approximation interval and polynomial order do not seem to introduce a bias in the measured observables, even if a small number of maximal eigenvalues falls outside the approximation interval. One should keep in mind, however, that this test was carried out on a relatively small volume and with only three sets of parameters. In addition, the approximation interval may affect the Kaon mass and as such it might be prudent to carry out a comparison on these configurations. Further, simple residual tests could be done where the inverse of the non-degenerate Dirac operator is calculated using the rational approximation and $\mathbf{r} = (1 - DD^{-1})\mathbf{b}$ is computed. As a final test, ensemble averages of the action contributions and perhaps gauge derivatives of monomials with different intervals and polynomial orders could be compared.

CONCLUSIONS

In this chapter, it was shown how the tmLQCD code-base was rigorously tested after very substantial changes were made to it. The value of high statistics HMC runs was demonstrated in Section 9.1 by explaining how three subtle but severe bugs were uncovered. In addition, it was shown how specially designed tests can be used to amplify suspected bugs and even pinpoint their most likely location in the code. The importance of these kinds of tests becomes especially apparent when the non-determinism from multi-threading makes it difficult to differentiate between real issues and differences which arise only because of changes in execution order or round-off errors.

The technique of using a numerical implementation of the derivative to check the consistency of a given action contribution and its analytical derivative was introduced in Section 9.2. It was shown how single observations of the difference between numerical and analytical derivatives may be misleading and how the parameters controlling the various approximations need to be carefully chosen to avoid drawing wrong conclusions. An application of the technique demonstrated that a particular bug in the implementation of the non-degenerate twisted mass clover action had been successfully fixed.

In Section 9.3 it was shown how a partially quenched calculation of the contribution of some term in the effective action can be used to calculate the derivative with respect to the coefficient of this term. In particular,

the derivative of the plaquette expectation value with respect to the mass splitting parameter μ_δ was computed using this technique. This was compared to the direct evaluation from simulations with a range of different μ_δ, where the RHMC was used to simulate the non-degenerate twisted mass clover doublet. The observed agreement further supports the observation that the non-degenerate twisted mass clover action is implemented correctly in the tmLQCD software package.

Finally, in Section 9.4, a further test of the RHMC algorithm implemented in tmLQCD showed that the choice of approximation interval of the rational approximation for the non-degenerate twisted mass clover doublet does not seem to affect the expectation vales of several observables. Simulations of $N_f = 2 + 1 + 1$ twisted mass clover fermions can thus proceed with confidence, although it would be prudent to conduct further tests of the rational approximation by varying more parameters.

Modern algorithmic features as well as the large number of optimisations for modern architectures have increased the complexity of lattice QCD codes significantly over the last decade or so. This level of complexity introduces significant interconnectedness of, and interactions between, different parts of the code-base. Coupled with the non-determinism introduced by multi-threading, this has made it very difficult to reason about the execution order and side-effects of the various parts of a given application. As a result, integration tests meant to check large parts of the code-base in production-level conditions should become an integral part of the software development process. The methods presented in this chapter provide a good set of such tests and will certainly complement the more specific but limited unit testing framework which has also been implemented in tmLQCD during the preparation of this thesis, but which has not been discussed here.

10 PERFORMANCE

The complexity of lattice QCD codes like tmLQCD has steadily increased with the introduction of more involved algorithms over the past ten years or so. As a result, writing special-purpose implementations of whole LQCD codes for specific machines is generally not possible anymore. This in turn entails that a code-base like tmLQCD needs optimisations for many different architectures which are often at odds with each other. Succeeding with a hybrid parallelisation, as discussed in Chapter 8 for example, is a challenge of balancing various overheads. As a result, the performance characteristics of the computational kernels, various algorithms and programs as a whole must be understood and bottlenecks identified.

Traditionally, this is done by establishing a performance model of a given algorithm on a given architecture, possibly with parameters that need to be determined empirically, and trying to optimize the features suggested by the model. Unfortunately, optimising in this way is extremely time consuming and beyond the scope of many practitioners of LQCD. Therefore, as discussed in Chapter 8, the optimisation of the various additions to tmLQCD was performed through sampling profiles using the SCALASCA toolkit [197]. In this chapter, the performance characteristics of the optimised routines will be discussed based on benchmarks of the hopping matrix kernel.

THE HOPPING MATRIX

In LQCD simulations which do not employ any gauge field smearing, the dominant expenditure of compute cycles occurs in the application of the so-called hopping matrix H. This involves the multiplication of fermion spinors with gamma matrices and gauge field components

$$H_{ij}^{\alpha\beta}(x,y) = \kappa \sum_{\mu=1}^{4} \Big[(U_\mu(x)(1+\gamma_\mu)\delta(x+a\hat{\mu},y))_{ij}^{\alpha\beta} \tag{10.1}$$

$$+ \left(U_\mu^\dagger(y)(1-\gamma_\mu)\delta(x-a\hat{\mu},y)\right)_{ij}^{\alpha\beta} \Big], \tag{10.2}$$

where $\delta(x,y)$ gives the neighbourhood relationship between lattice sites. In tmLQCD, the quark field boundary conditions are implemented by multiplying κ with a complex phase and as a result, the application of the hopping matrix requires a total of 1608 floating point operations (FLOPs) per lattice site.

The next-neighbour nature of the LQCD action means that when domain decomposition with inter-process communication is used for parallelisation, the processes need to communicate boundary terms of the gauge and spinor fields to a depth which depends on the algorithm under consideration. For the hopping matrix, this corresponds to one lattice site in each boundary direction for both spinor and gauge fields. The structure of H means that the amount of data that needs to be communicated can be halved by exploiting the fact that

$$\rho^{\pm}(x+a\hat{\mu}) = (1 \pm \gamma_\mu)\psi(x+a\hat{\mu}) \tag{10.3}$$

is a spinor with only two independent components. The so-called *half-spinor* hopping matrix first computes this projection, communicates the result on the boundaries, projects back to a four-component spinor and only then performs the remaining multiplications. This comes at the cost of having two loops over the volume compared to the *full-spinor* matrix, which communicates the unprojected spinors before looping over the volume once. Quite generally, whenever inter-process communication is used, the half-spinor matrix is faster because usually performance is very much bound by the communication latency and bandwidth. However, multiple volume loops lead to larger multi-threading overheads, and these need to be carefully balanced.

As has already been discussed in Section 8.5, the hopping matrix can be optimised further by *overlapping communication and computation*. The process-local lattice domain naturally splits into a *surface* and a *body* contribution, depending on whether or not the given lattice site is on the boundary. On modern architectures, the communication process is offloaded largely onto the networking hardware and as a result, the CPU idles while the communication of the boundary terms is in progress. In the half-spinor implementation of the hopping matrix, the projection and multiplication stages for lattice sites in the body of the local domain can be done right away, since no data needs to be communicated. Consequently, the communication overhead can be reduced or hidden if this is done while the communication is in progress. As a downside, a further overhead is introduced because a third (partial) volume loop becomes necessary. Consequently, as described in Section 8.5, the computation must be balanced with care if the optimisation is to bear fruit.

In the following measurements, the nomenclature introduced here will be used frequently without further exposition.

PERFORMANCE MEASUREMENTS

The measurements in this section are based on results from the tmLQCD benchmark application which applies the hopping matrix H many times, computes the number of theoretical FLOPs executed and divides by the elapsed time. Many other routines were optimised during the preparation of this thesis, but since the hopping matrix is dominant almost everywhere, it is essential for it to be as fast as possible. The following measurements also feature a separate indication with *communication disabled*. This means that the computation is done as usual, but dummy routines which return immediately are called instead of the communication routines. Of course, the results of these computations are meaningless, but the measurement allows the communication overhead to be monitored. As will be seen, this is especially insightful when the performance is very sensitive to the local lattice volume.

Commodity Intel Clusters

Baseline performance was established on a commodity Infiniband cluster with two, quad-core Intel Xeon X5560 processors per node with Hyper-threading disabled. The full- and half-spinor matrices were benchmarked in pure MPI, pure OpenMP and hybrid OpenMP/MPI parallelisation using a lattice volume of $16 \cdot 8^3$ sites per node. This optimally utilises the L2 cache when two double precision spinor fields and the double precision gauge

	N_T, N_X, N_Y	$N_{threads}$	V
MPI, 8 nodes	4, 4, 4	–	$32 \cdot 16^3$
MPI, 1 nodes	2, 2, 2	–	$16 \cdot 8^3$
Hybrid, 8 nodes	2, 2, 2	4	$32 \cdot 16^3$
OpenMP, 1 nodes	–	8	$16 \cdot 8^3$

Table 10.1: Parallelisation parameters for baseline performance benchmarks shown in Figure 10.1.

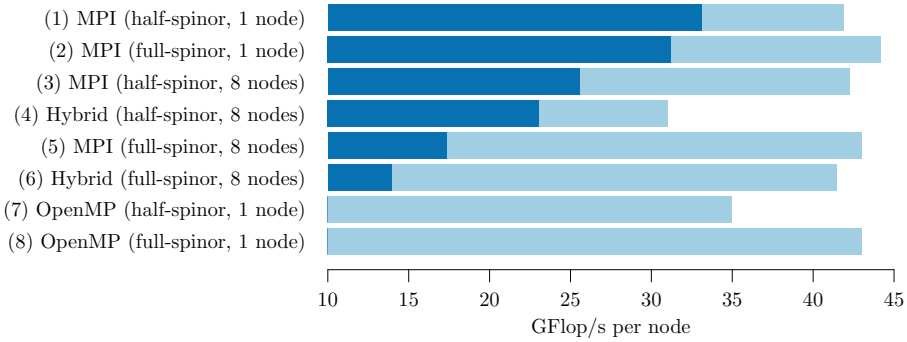

Figure 10.1: tmLQCD hopping matrix benchmarks on a dual-socket Intel Xeon X5560 Infiniband cluster. The lighter bars show performance with communication routines explicitly disabled. Parallelisation details in Table 10.1.

field are considered as the working set. Measurements were done on one and eight nodes of this machine, chiefly to study the performance characteristics of the multi-threading that had been added to the code-base.

The OpenMP implementation in tmLQCD is designed such that the code can be compiled with all OpenMP instructions explicitly disabled if they are not going to be used. In addition, this mode also enables some manual optimisations which are not possible when multi-threading is used because they lead to loop dependencies. This version was used for pure MPI measurements.

The different parallelisations are characterised by the number of MPI processes in each direction and the number of OpenMP threads per process (if used). In tmLQCD, the slowest running array dimension is T, followed by X, Y, Z, spin (if present), colour and finally the real and imaginary parts. MPI parallelistation was done in the three dimensions T, X and Y and four threads per process were used in the hybrid mode with a kernel which attempted to overlap communication and computation. Pure OpenMP was naturally only run on only one node with 8 threads. The details are summarized in Table 10.1. The code was compiled with version *14.0.3.174* of the Intel compiler suite, with only `-O3 -axSSE4.2` as compiler flags beyond the defaults.

The benchmark results are shown numbered in Figure 10.1 in units of GFlop/s per node. The lighter bar indicates performance with communication routines disabled. The fastest performance on one node is achieved by the pure OpenMP (8) with the full-spinor implementation of the hopping matrix and the optimisation of the neighbourhood look-up array described in Section 8.3.1. Performance on this architecture is surprisingly good with the OpenMP result corresponding to an efficiency of over 50% in terms of the theoretical peak performance[1]. It is interesting to note the overhead stemming from the half-spinor implementation of H with pure OpenMP (7),

[1] The full-spinor OpenMP implementation is likely to be a very good starting point for optimizing tmLQCD for high performance on Intel Xeon Phi accelerator cards with communication disabled, which would be ideal for trivially parallelisable workloads such as inversions. In its current state, however, the efficiency on Xeon Phi is only around 3% on account of the poor auto-optimisation for the large SIMD unit. Reaching high performance on this achitecture involves a much larger number of optimisations [208].

likely a result of the second volume loop. The next fastest result on a single node is pure MPI (1) with the half-spinor implementation, suffering a substantial communication overhead of almost 10 GFlop/s compared to (8).

For the larger total volumes, the half-spinor pure MPI (3) implementation is fastest and the communication overhead grows by about another 6 GFlop/s compared to only node-local communication. Comparing hybrid parallelisation (4) and pure MPI (3) with communication routines disabled, one sees that the threading overhead is around 10 GFlop/s while the communication overhead is significantly smaller than in the pure MPI case. As a result, despite overlapping communication and computation hiding some of the overhead and the apparent good performance of OpenMP on this architecture, the hybrid code is slower than the pure MPI code. It should be noted though that when many nodes are used on this kind of machine, there may come a point where the hybrid implementation scales better than the pure MPI counterpart because of inter-process overheads.

This is in fact observed on larger machines such as the Cray XC30/XC40. In addition, with many architectures now equipped with twelve, fourteen or sixteen core CPUs, usually with two hardware threads per core, the limitations of pure MPI parallelisation are being reached. Especially when the number of cores contains factors of three, five or seven, the usual domain decomposition of lattice volumes containing or requiring factors of two is not possible, thus ruling out pure MPI parallelisation utilising all available cores. In these cases, the usage of a hybrid parallelisation can increase efficiency by as much as double digit percentages.

BlueGene/Q

The main production architecture for the ETMC at the time of writing of this thesis was IBM BlueGene/Q (BG/Q). Consequently, a significant amount of time was invested into understanding the performance characteristics on machines of this type and introducing various optimisations. This was discussed in Section 8.5 and benchmark results on this architecture will be presented here. The BG/Q compute chip [209] has seventeen cores, sixteen of which are available to compute applications, each capable of four hardware threads with very fast switching between them. Peak performance is quoted as 204.8 GFlop/s per compute chip. Nodes are networked in a five dimensional torus using dedicated network controllers, assisted by the seventeenth CPU core. As a result, there is in principle almost no communication overhead as far as the compute cores are concerned. This means that overlapping communication and computation should be highly efficient and this is in fact observed, with the hybrid implementation of the code being almost a factor of two faster than the pure MPI one, as will bee seen below. The details of the implementation of the hopping matrix kernel were already discussed in Section 8.5 and the reader is referred there for various nomenclature. All the measurements in this section were done using double precision floating point arithmetic.

A particular feature of performance on BG/Q, at least with tmLQCD, is a significant dependence of the efficiency on the node-local lattice volume. This is likely due to inefficiencies in the way data is read from and written to memory in the computational kernels in tmLQCD as it is less apparent in more advanced implementations, such as that of Ref. [202].

The performance characteristics were studied as a function of the thread-local volume L^4 and the distribution of threads N_{th} and MPI processes

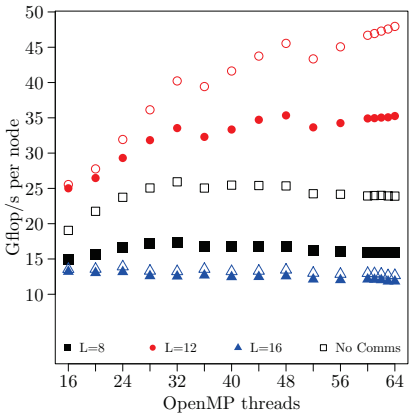

(a) 1 process per node and increasing numbers of threads.

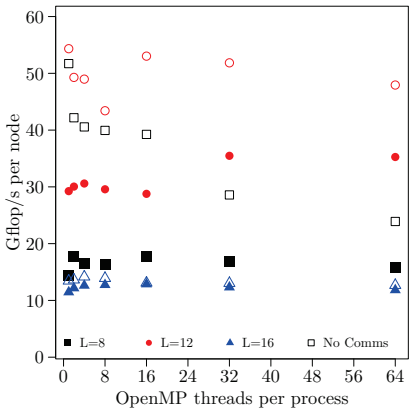

(b) The number of threads times the number of processes per node is kept constant at 64.

Figure 10.2: Performance on 32 nodes of BG/Q in GFlop/s per node using an overlapping OpenMP/MPI hopping matrix for different node-local lattice volumes L^4. Empty symbols correspond to inter-process communication being disabled.

N_{MPI} (using MPI rather than IBM MU SPI communication), the result of which can be found in Figure 10.2. The scaling with the number of threads with one process per node is shown in Figure 10.2a. The empty symbols show performance with communication disabled and it is clear that going from $L = 8$ to $L = 16$, the optimal local lattice volume is $V = 12^4$ where cache usage is maximised. Dropping out of cache ($L = 16$) induces an extreme performance regression while not fully utilising cache results in a performance plateau from $N_{th} = 32$ onwards. For the optimal $L = 12$ and sub-optimal $L = 8$ lattice extents, performance grows linearly to 64 and 32 threads respectively, which shows that when cache utilisation is optimal, OpenMP overheads are suppressed. This is true with and without communication, but evidently the MPI overhead on BG/Q is quite severe, as will be seen further below.

The performance of the hopping matrix when the product $N_{MPI} \cdot N_{th} = 64$ is held constant is shown in Figure 10.2b. Again, looking first at the empty symbols, with communication switched off, performance is maximsed for $N_{th} = 1$, but this result is not meaningful as soon as communication is enabled, as will be seen further below. More realistically, it is seen that for the optimal lattice extent $L = 12$, maximizing the number of threads and having one process per node is advisable as the increased OpenMP overhead seems to be amortized well. On the other hand, for the sub-optimal lattice extent, there is a significant regression above $N_{th} = 16$ onwards. This may suggest that for volumes which do not fully utilise the L2 cache, working with fewer threads per process is beneficial. When communication is switched on, for the optimal volume maximum performance is reached from $N_{th} = 32$ onwards while the non-optimal volume is rather flat over the whole range.

The effect of the optimisation strategies on the hopping matrix kernel is shown is shown in Figure 10.3a. Before optimisation and using MPI communication only, the black diamonds give an efficiency of below 5%. Utilising the full SIMD through QPX intrinsic functions, as far as possible with the current memory layout in tmLQCD, leads to the red squares (with communication disabled). Enabling MPI communication results in the magenta circles, showing a significant performance regression. Switching to SPI communication routines (cyan triangles) improves this somewhat but only when the processes are mapped correctly to the network is maximum performance *with* communication reached (dark blue downward pointing

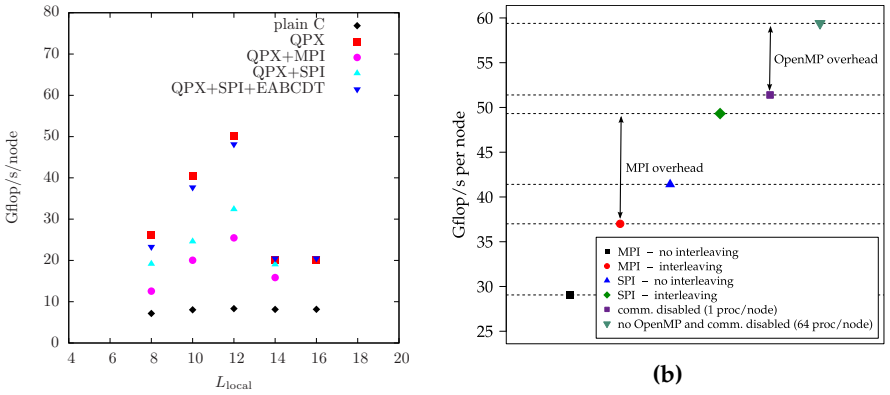

Figure 10.3: (a): Different optimisation levels of the hopping matrix computational kernel in tmLQCD on BG/Q with different node-local lattice volumes L^4. **(b)**: Overhead estimates in hopping matrix for optimal node-local lattice volume $V = 12^4$.

triangles). This last measurement at $L = 12$ represents about 24% of peak performance and the strong dependence on the node-local lattice volume is apparent.

An estimate of the various overheads which are still present and explain the differences between the various implementations can be gathered from Figure 10.3b which can be read from top to bottom. The measurements stem from 32 nodes with a node-local lattice volume of $V = 12^4$ and full utilisation of QPX. When all communication is disabled and no OpenMP is used, the hopping matrix performs at just below 60 GFlop/s per node, giving the current maximum performance that can be reached with tmLQCD. Enabling OpenMP (which is necessary for overlapping communication and computation), has an overhead of around 8 GFlop/s. Enabling SPI communication with this overlapping code leads to a small performance regression of about 2 GFlop/s, showing that the overlapping kernel with SPI communication is very good at hiding the communication latency. In fact, when overlapping is disabled, performance drops by another 8 GFlop/s. Switching back to the overlapping kernel but using MPI for communication instead, produces even lower performance, indicating that the MPI overhead is over 12 GFlop/s or so, which is very substantial. Switching off overlapping with MPI communication deteriorates performance further and shows the significant impact that hiding communication latency through overlapping communication and computation can have, even when using MPI communicators.

A complete overview of the performance and the various overheads is shown in Figure 10.4, including the differences between the different ways in which overlapping communication and computation can be implemented, as discussed in Section 8.5. As should be clear from Figure 10.3b, a remaining significant overhead is the one due to threading with OpenMP. Following the argument in Section 8.3.1, most of this stems from thread management at the entry and exit points of multi-threaded routines and can be reduced or eliminated through the usage of orphaned directives. When this is done for the hopping matrix kernel, performance with communication increases by about 3 GFlop/s (comparing the dark bars of 1 and 2). Performance with communication disabled, however, increases substantially more, almost all the way to the performance without OpenMP, in excess of 60 GFlop/s per node. This confirms that most of the OpenMP overhead is indeed from thread management. It also indicates that there exist additional communication overheads even with SPI, because the difference in (1) between performance with and without communication is much greater than

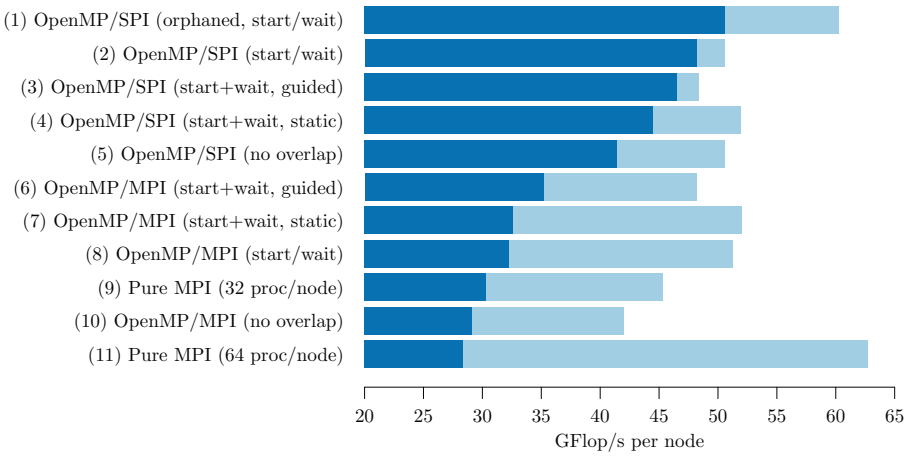

Figure 10.4: Performance measurements for different versions of the hopping matrix. The lighter bar shows measurements without communication. Unless noted otherwise, the hybrid kernel attempts to overlap communication and computation. QPX intrinsics are used throughout.

when orphaned directives were not used (2). In principle this suggests that if orphaned directives were deployed and further tuning were done to the SPI communication, a performance just below 60 GFlop/s per node would be attainable with the current tmLQCD memory layout in double precision.

CONCLUSIONS

It has been shown in this chapter that the optimisations of the tmLQCD code-base in the process of preparing this thesis were largely effective at increasing performance. On BlueGene/Q, a combined performance increase of around a factor of 6 was attained compared to the unoptimised code and this was only possible through a combination of multi-threading, utilising the available SIMD units and mapping the problem to the network correctly. Many of the various overheads on this crucial architecture were understood and catalogued, allowing good choices to be made with regards to how computations for different lattice volumes are parallelised. It has been shown that a good choice can impact performance almost by around a factor of 2 and that computations which spill out of cache should be avoided at all costs on this architecture.

The performance modelling done using SCALASCA in Chapter 8 has turned out to be largely representative of real world performance, correctly explaining for instance the origin of the observed OpenMP overheads. It was also shown that well-chosen benchmarks can provide estimates for the size of different overheads and hint at where optimisations might be possible and necessary. Together with algorithmic advances, the optimisation for new architectures has been a driving force of lattice QCD and this will likely become much more important as new machines with increasingly hybridised compute units emerge on the path to exa-scale performance. In the future, therefore, it would be desirable to have more formal performance models of the tmLQCD computational kernels in order to enable optimisations to be planned.

11 INTEGRATOR TUNING VIA FOURIER TRANSFORMS

In the last chapter of this part, a potentially interesting observation regarding the behaviour of the molecular dynamics forces of quark determinants and determinant ratios will be explored. Using the notation of Section 9.2, the gauge derivative of a given term in the effective action is given by $\delta P_\mu^a(x,\tau)$. The magnitude of the forces has been studied empirically as a function of the mass preconditioning parameters in Chapter 5. As seen there, light dynamical fermions are known to produce strong fluctuations, including deviations which exceed the average forces by orders of magnitude. In addition to these exceptional spikes, it seems from practical experience that the frequency of the fluctuations in the forces depends inversely on some unknown power of the quark mass in the operator in the determinant and on some complicated combination of the parameters in the determinant ratio.

The argument can be made that naïvely there are thus two sources of finite step size integration errors that occur in molecular dynamics integration of these force contributions. On the one hand, when $\delta\tau\delta P$ is too large, the momentum update T_P (Equation (1.100)) will not be a small increment, thus inducing a source of error. This reasoning historically led to improvements in integrators and the determinant splitting techniques that have been discussed in Chapter 5.

Another potential source of finite step size errors may be seen from a different perspective. Clearly, $\delta P_\mu^a(x,\tau)$ are smooth functions of τ, a fact not reflected by the molecular dynamics histories shown in Chapter 5 because the measurements are only taken once every trajectory of length $\tau = 1.0$. As a result, although these histories indicate average and maximal forces rather well and show signs of the frequency of exceptionally large fluctuations, they do not provide a measure of the characteristic frequencies which contribute to $\delta P_\mu^a(x,\tau)$. The integration process can also be interpreted as a sampling problem of the "signal" δP, sampled at intervals $\delta\tau$, the integrator step size. When this has contributions from fluctuations with characteristic frequencies f which exceed

$$f > \frac{1}{2\delta\tau}, \tag{11.1}$$

they cannot be sampled correctly, leading to integration errors of a different kind. It is clear that the Nyquist-Shannon sampling theorem does not really apply here because δP may not be sufficiently band-limited, but it may still provide some level of guidance for choosing a minimum step size to avoid these integration errors. It will be shown in the following that in the cases that were studied, δP does appear to be largely band-limited, with the limiting frequency increasing as the quark mass is lowered. In addition, mass preconditioning seems to have a substantial effect also on the dominant characteristic frequencies present in $\delta P_\mu^a(x,\tau)$.

BACKGROUND

Using the developments of Section 9.2, the molecular dynamics histories of the $\delta P_\mu^a(x,\tau)$ for the different monomials in the action can be analysed directly. The idea is to use the fast Fourier transform (FFT) of $\delta P_\mu^a(x,\tau)$ to

determine the minimum number of integration steps required to resolve the finest fluctuations. To this end, after a simulation is thermalised, one long trajectory ($\Delta\tau \geqslant 10$ or so) would be simulated with many integrations steps. During this trajectory, $\delta P_\mu^a(x,\tau)$ would be sampled for a randomly chosen subset of a, μ and x for all τ. The FFT of this trajectory can then be taken to provide estimates of the dominant frequencies in the signal

$$\mathcal{F}\left[\delta P_\mu^a(x,\tau)\right](f_m) = \sum_{k=1}^{n}\left[\delta P_\mu^a(x,k) - \langle \delta P_\mu^a(x)\rangle_{\Delta\tau}\right]\cdot \exp\left\{-2\pi i\frac{(k-1)(m-1)}{n}\right\}, \quad (11.2)$$

where the mean value of δP has been subtracted to remove the constant term[1], k is the integration step and n is the total number of integration steps on the considered time scale. Further, f_m is the discrete frequency component in units of the trajectory length which goes from $f_1 = 0, f_2 = 1/\Delta\tau$ to $f_n = (n-1)/\Delta\tau$. Since δP is real, only $n/2-1$ amplitudes are unique and because the mean has been subtracted, $\mathcal{F}\left[\delta P_\mu^a(x,\tau)\right](f_0) = 0$. The number of steps on each time scale t is specified by N_t, resulting in $n_t = 2N_t + 1$ effective steps for the 2MN integrator.

As will be seen below, the FFT for a single $\delta P_\mu^a(x,\tau)$ is noisy, making the interpretation difficult. Increasing the trajectory length would help, but $\Delta\tau$ would need to be taken impractically large to avoid this. However, it seems that because the forces everywhere on the lattice are similar, the dominant discrete frequencies can be isolated by averaging the amplitudes of the FFT over the combinations of a, μ and x that were sampled. Labelling these combinations with the super-index $A_i = (x,\mu,a)$, the frequency spectrum for each of the monomials is then given by the moduli of the amplitudes in Equation (11.2), averaged over these N_c time histories

$$\|\mathcal{F}[\delta P](f_m)\|_{av} = \frac{1}{N_c}\sum_{i=1}^{N_c}\|\mathcal{F}[\delta P(A_i)](f_m)\|. \quad (11.3)$$

It is at present unclear to what extent this averaging may introduce systematic biases into the observed average amplitudes.

NUMERICAL SETUP

The idea in its current form is demonstrated here on the basis of a simulation using $N_f = 2$ Wilson twisted mass clover fermions and Iwasaki gauge action with simulation parameters:

- $T/2a = L/a = 16$, $\beta = 2.1$, $\kappa = 0.1373$, $c_{sw} = 1.57551$, $a\mu_\ell = 0.01$($\tilde{\mu}_\ell = 0.002746$), corresponding to $M_{\pi^\pm} \sim 450$ MeV
- four time-scales t_i with the gauge monomial on the finest one (t_0)
- mass preconditioning with two determinant ratios and parameters $\tilde{\rho}_3 = 0.3$, $\tilde{\rho}_2 = 0.03$, $\tilde{\rho}_1 = 0.0$.
- second order minimal norm (2MN) integrator [74] on all time scales
- trajectory length $\Delta\tau = 1.0$ with N_t integration steps on each time scale t, $N_t = \{1,1,1,13\}$.

[1] for $\Delta\tau$ sufficiently large this is zero anyway

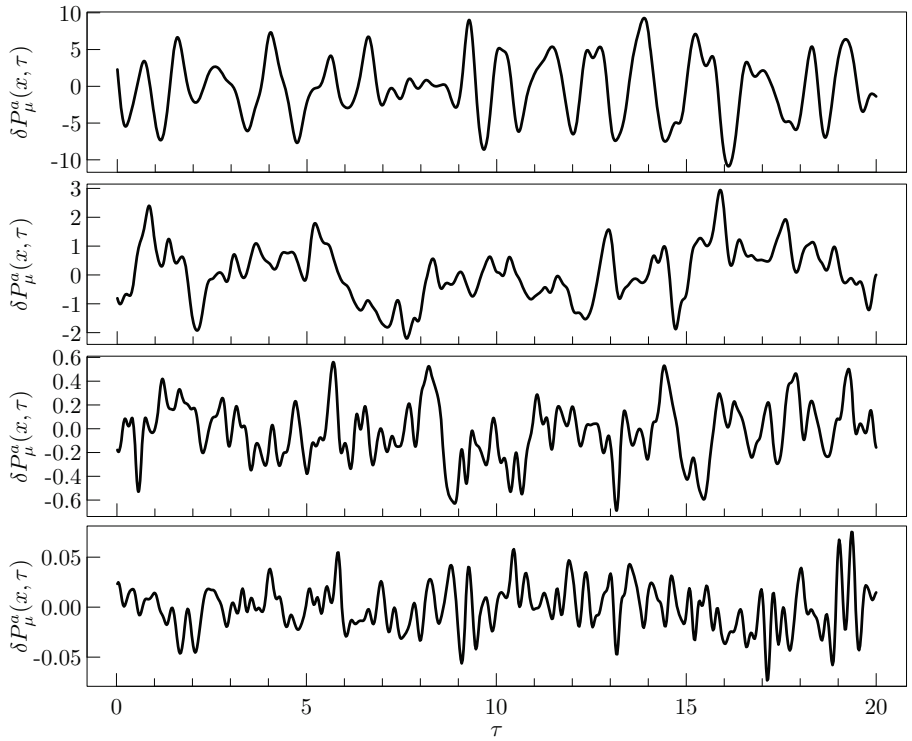

Figure 11.1: Contributions to the derivative for one μ, a, x of the four monomials along a trajectory of length $\Delta\tau = 20.0$. The gauge contribution is in the top-most panel. The next panel gives the contribution of the determinant (with $\tilde{\rho} = 0.3$) and the lower-most panel of the determinant ratio with the target mass $\tilde{\mu} = 0.002437$ in the numerator.

This is the same ensemble that was used for Chapter 5 and after thermalisation, the acceptance rate was found to be around 93% with these parameters.

RESULTS

To test the method, a trajectory of length $\Delta\tau = 20.0$ was integrated using the 2MN integrator with the number of steps on the coarsest time scale set to $N_3 = 400$, corresponding to an effective number of integration steps $n_3 = 801$. The time history of $\delta P^a_\mu(x,\tau)$ for each monomial in the MD Hamiltonian is recorded for 1024 combinations x, μ and a. The total number of collected samples is thus

- $N_c = 1024$ combinations of x, μ and a (hence, effectively, 1024 trajectories),
- $n_t = 2N_t + 1$ effective integration steps on a given time scale,

which corresponds to 820224 samples for $N_3 = 400$ and increases by a factor of two for each of N_2, N_1, N_0. Typical molecular dynamics histories of the derivative contributions for the gauge, determinant and the two determinant ratio monomials are shown in Figure 11.1 from top to bottom. The FFT of a single $\delta P^a_\mu(x,\tau)$ is shown in Figure 11.2. Evaluation of Equation (11.3) for the trajectory described above leads to Figure 11.3, which shows the averaged FFT moduli of the trajectories ordered from top left to bottom right. Because the contributions from different time scales are sampled at different rates, the discrete frequencies f_m in this figure have been rescaled to be comparable and are expressed in units of unit length trajectories (so $f = 2.0$ corresponds to a mode with period $T = 0.5$). The data is further normalised such that the maximum amplitude has unit modulus and the dotted vertical and horizontal lines indicate at which frequency this relative amplitude drops below 0.01, chosen as an arbitrary point of reference.

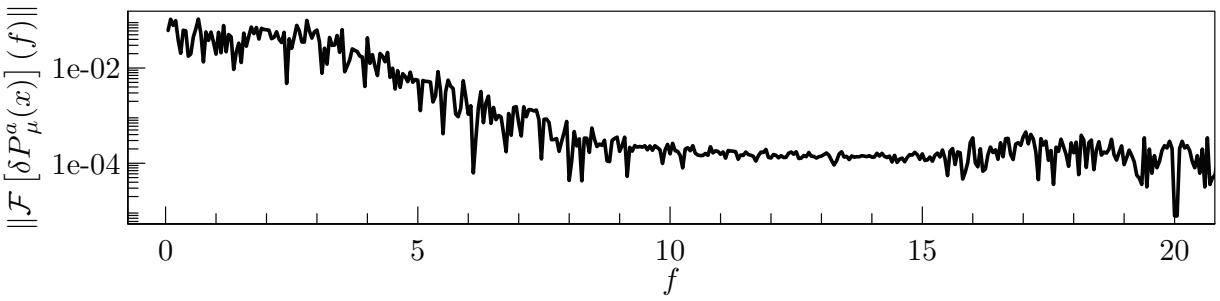

Figure 11.2: The modulus of Equation (11.2) evaluated for a single trajectory of the force contribution due to the determinant ratio on the outermost time scale with $\tilde{\rho}_1 = 0.0, \tilde{\rho}_2 = 0.03$. The discrete frequencies have been rescaled to be in terms of unit length trajectories.

It should be noted at this point that it was checked whether the trajectories shown in Figure 11.1 are stable with respect to changes in the number of integration steps. For this purpose, the sampling was repeated and N_3 was reduced in steps of 50. Until $N_3 < 200$, the resulting trajectories are very similar to Figure 11.1, with first differences starting to appear at around $\tau = 15$. In any case, up to this point the frequency spectra can be completely resolved including the point where the relative amplitude reaches 0.01. For trajectories with even fewer integration steps, differences appear much earlier in the course of trajectory. Especially for monomials on the outer time scales, it is visually obvious that the fluctuations are not being sampled correctly. In the frequency spectra one observes that the frequency where the relative amplitude drops below 0.01 is not sampled. MD histories and frequency spectra with $N_3 = 300$ and $N_3 = 150$ are shown in Appendix F. Increasing the number of integration steps above $N_3 = 400$ does not lead to visually discernible differences, nor does it affect the conclusions regarding the drop of the relative amplitude below 0.01.

The first notable feature is that the derivative from the gauge monomial has a clear dominant frequency around $f = 1$, which can be confirmed (approximately) by looking at top-most panel of Figure 11.1. Further, the point at which the normalised modulus drops below 0.01 is highest for the determinant ratio with the target mass in the numerator on the outermost time scale and seems to decrease monotonically as the mass in the numerator is increased. In addition, as the mass parameter in the numerator is decreased, a larger range of frequencies contribute significantly to the dynamics.

The idea now is to use the point where the relative amplitude of the signal on the outermost time scale drops below 0.01 as a basis for selecting the number of integration steps. For the determinant ratio with the target mass in the numerator shown here, this frequency is seen to be $f \sim 10$ and so it follows that around 20 effective integration steps per $\tau = 1.0$ should be required on the outermost time scale to capture the dominant fluctuations in the force. This would correspond to $N_t = \{1, 1, 1, 10\}$ in the notation that has been used above. Testing this integration scheme in practice shows that indeed the acceptance rate remains around 90% going from $N_3 = 13, 12, 11, 10$ and drops significantly when the number of integration steps is reduced further. Of course, given the limited scope of this analysis, this serves as no more than an indication.

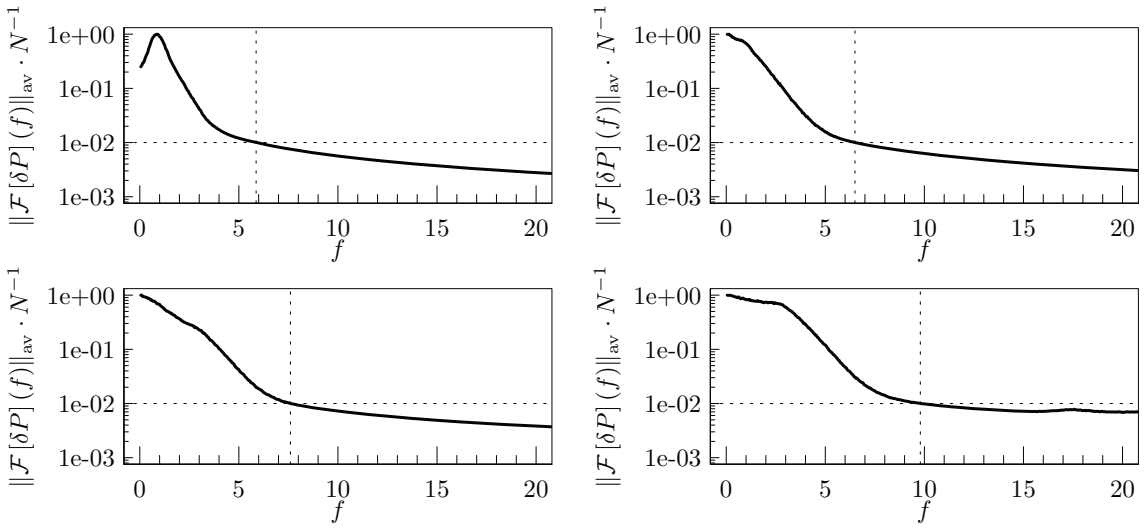

Figure 11.3: Evaluation of Equation (11.3) on the test trajectory introduced in the text. From top left to bottom right: gauge monomial, determinant, determinant ratio with $\tilde{\rho}_3 = 0.3$ in the denominator, determinant ratio with $\tilde{\rho}_2 = 0.03$ in the denominator. The data has been normalised with respect to the maximum average modulus. The discrete frequencies have been rescaled to be comparable.

Trajectory without Mass Preconditioning

It is interesting to use the technique to explore the role that mass preconditioning plays in suppressing large amplitude, high frequency oscillations in the fermionic contribution to the force. For this purpose, a trajectory was integrated with just a quark determinant with the target mass parameter on time scale t_1 and the gauge monomial on time scale t_0. On account of the substantial computational cost, it had to be limited to a length of $\Delta\tau = 1.0$ with $N_t = \{1, 400\}$. Typical trajectories of the gauge monomial and the fermion determinant forces are shown in Figure 11.4, together with the averaged amplitudes of the frequency modes.

It is clear that the fermion force is of substantial magnitude, comparable to that of the determinant term in Figure 11.1. The fluctuations are more regular, but the FFT average shows that many more frequency modes contribute strongly with a prominent *shoulder* of frequencies extending to about $f = 30$. It is interesting to note that the determinant ratios in Figure 11.3 show a similar feature to a much lesser extent and that the width seems to be depend on the mass parameter in the numerator. The quark determinant can be thought of as a determinant ratio with an infinite mass in the denominator, which suggests that this could be taken as the limiting case in an attempt to theoretically understand the observed patterns.

The gauge force, on the other hand, is comparable to the one shown in Figure 11.1, considering the different trajectory lengths. The FFT of the gauge force trajectory of course has no meaning because the trajectory was too short for this purpose. It is tempting to conclude that mass preconditioning seems to suppress not only the amplitude of the average forces, it also significantly affects the dominant frequencies which contribute to the forces of the fermion monomials in the MD Hamiltonian. However, a more comprehensive study of the force histories for determinants with different quark masses is necessary to form a complete picture.

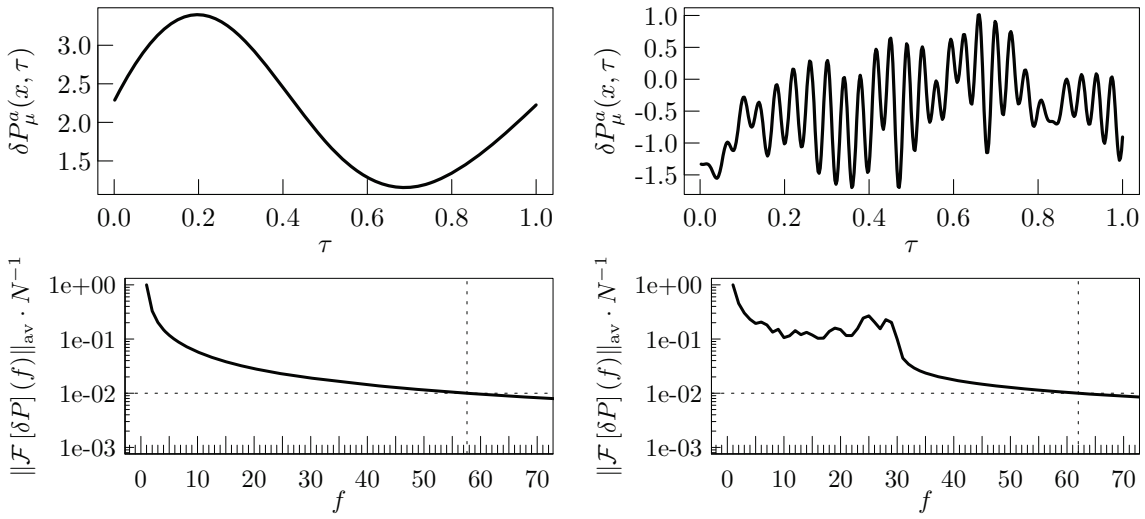

Figure 11.4: (top): Representative trajectories of the gauge (left) and determinant ratio (right) monomials over a trajectory without mass preconditioning. **(bottom)**: Equation (11.3) evaluated on the collection of these trajectories.

Trajectory with Optimized Mass Preconditioning

Using a mass preconditioning scheme along the lines of Section 5.4 with five time scales and $x = 4$ results in MD histories for the gauge, determinant and three determinant ratio monomials as shown in Figure 11.5. Compared to the original scheme, high frequency modes seem to be even further suppressed with the fermionic monomials not showing signs of the kind of shouldering visible in Figure 11.3. Although the cut-off frequency for the determinant ratio on the outermost time scale is somewhat higher, the frequencies just below this cut-off are much more strongly suppressed than in the original scheme. This seems to confirm that mass preconditioning indeed suppresses higher frequency modes, although the results should be carefully checked before being used to guide the optimisation of simulations.

Conclusions

In this chapter a tentative attempt was made at using long molecular dynamics histories of forces to explore the effect of mass preconditioning on the frequency of large amplitude, high frequency oscillations in the fermionic forces. The fast Fourier transform of long trajectories with many time steps of these forces was taken and the frequency spectrum was averaged over many lattice sites, directions and SU(3) generators. The resulting frequency distributions are seen to drop off significantly for high frequencies when mass preconditioning is used. In particular, the frequency spectra of the fermionic contributions seem to approach each other as the extent of mass preconditioning is increased. One should note that the resulting spectra are reminiscent of Brownian noise, but are not seen to drop off as strongly with the frequency for low frequencies.

For the gauge contribution, a peak is identified at a frequency $f \sim 1$ and it would be insightful to understand where this originates and whether it depends on relevant parameters such as the inverse gauge coupling β or the type of gauge action used. It is interesting that the monomials that are usually integrated on finer time scales have large magnitudes, but seemingly show smaller contributions from higher frequencies, compared to

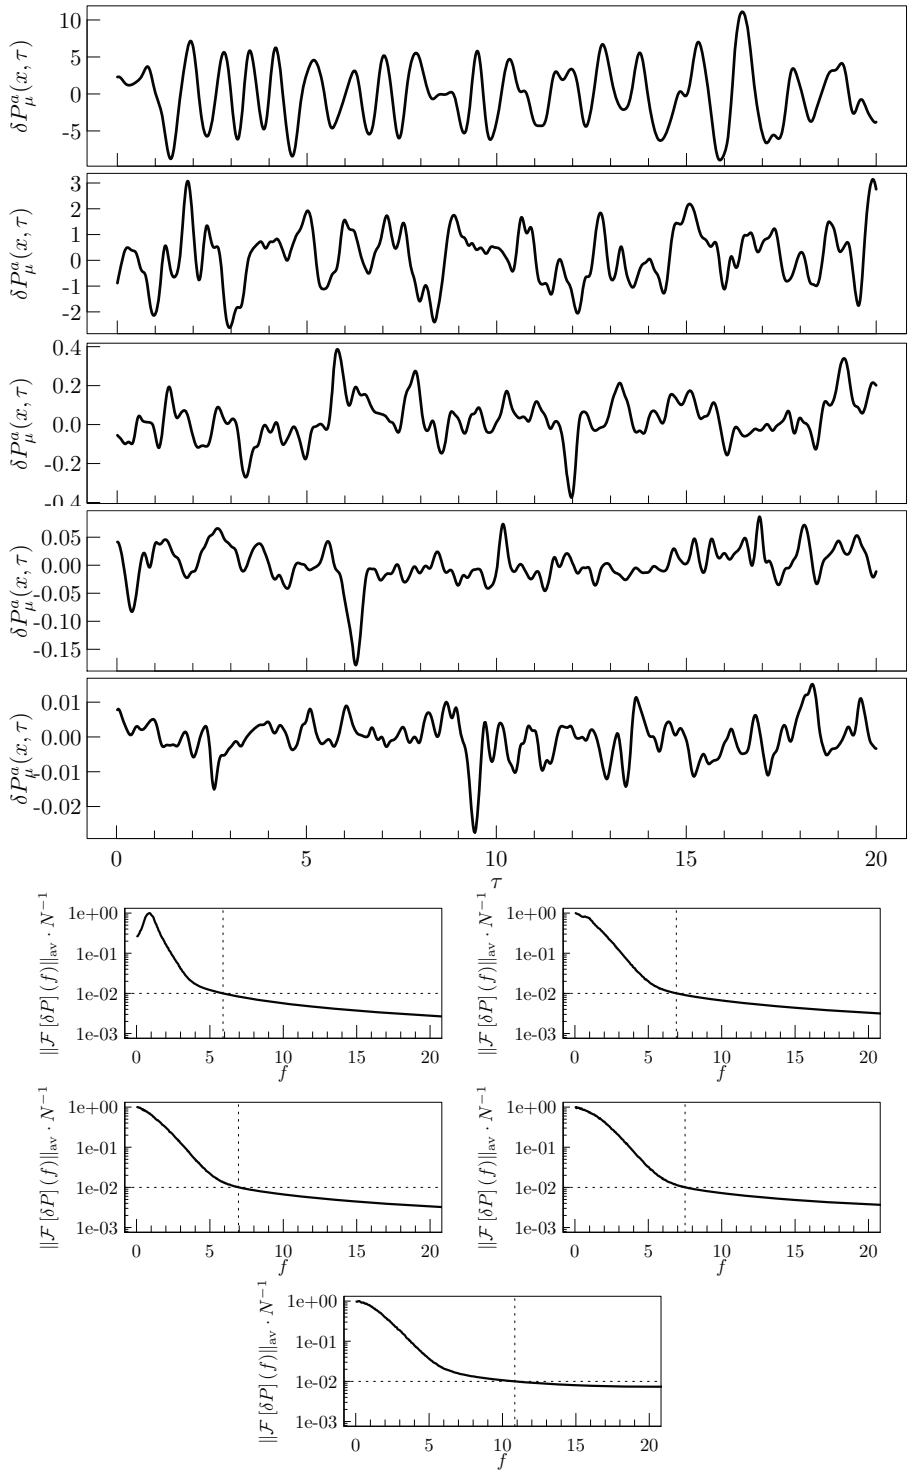

Figure 11.5: Representative MD trajectories and their normalised average frequency amplitudes of the forces in a simulation with mass preconditioning following Section 5.4 on five time scales and $x = 4$. The gauge monomial is in the topmost row, followed by the determinant and three determinant ratios. For the FFTs, the same order is given top-left to bottom-right and then bottom-centre.

the monomials usually integrated on the coarser time scales. This would suggest that achieving good acceptance may require balancing the high frequency oscillations on the coarsest time scales (and the required number of integration steps) against the number of integration steps required on the finer time scales due to the magnitude of the forces.

The frequency at which the amplitude of the frequency spectrum drops by a factor of 100 relative to its maximum, is seen to be affected by the preconditioning masses. Using this criterion to select the number of effective integration steps on the outermost time scale was shown to provide a good estimate of the minimum number of integration steps required to retain a high acceptance rate, at least for the single simulation tested here.

It would therefore be interesting to extend the present investigation to a wider range of mass preconditioning schemes, in particular in large volume and at, or close to, the physical pion mass. Studies with a range of inverse gauge couplings and perhaps different gauge actions would likely also provide interesting data. Further, if a more general pattern were found – perhaps a functional dependence of f on the preconditioning masses – this could be used to supplement the proposal of Section 5.4 with a suggested number of integration steps on the outermost time scale. Finally, lacking theoretical guidance, the analysis could be applied to ongoing simulations to check if, for example, the kind of shouldering of the frequency spectrum is observed. This could be taken as a signal for possible optimisations to improve acceptance rates by adjusting the mass preconditioning parameters to the number of integration steps.

CONCLUSIONS AND OUTLOOK

The aim of this thesis was to investigate the feasibility and practicality of simulations with physically light quarks in the Wilson twisted mass discretisation of Lattice QCD at maximal twist. With the addition of the Sheikholeslami-Wohlert term to the twisted mass action, it was found that in two flavour simulations, the physical charged pion mass can be successfully reached and the theory tuned to maximal twist, ensuring $\mathcal{O}(a)$-improvement. It was shown further that the bare quark mass parameter at which the charged pion mass takes its physical value can be extrapolated successfully from only a few simulations with heavier than physical quarks. No signs of meta-stabilities were observed , indicating that any phase transition points in the space of bare parameters must be sufficiently far away from the physically interesting region. In addition, no unexpectedly long auto-correlations were observed which could have hinted at the presence of very light stable particles in the spectrum. Large fluctuations in the molecular dynamics forces as well as spikes in the molecular dynamics histories of the energy violation were investigated. It was found that integration schemes with more time scales and a determinant splitting over more terms seem to reduce their occurrence, suggesting that they do not constitute a fundamental problem.

The involved parameter tuning for simulations with light, strange and charm quarks was discussed and parameter sets were proposed for first simulations at a lattice spacing of around 0.095 fm. A two-pronged tuning strategy was investigated which separates the tuning of strange and charm quark masses from the tuning of β and c_{sw}, resulting in seemingly good parameter estimates at a reasonable computational cost. It was found that if β and c_{sw} are chosen appropriately, the values of μ_σ and μ_δ tuned in the valence sector are good estimates and do not seem to require retuning once they are used as sea quark mass parameters in $N_f = 2+1+1$ simulations.

The dependence on c_{sw} and β of the lattice spacing, critical hopping parameter and plaquette expectation value was determined from simple empirical fits to simulation data spanning a wide range of parameter values. Somewhat surprisingly, it is seen that this procedure is sufficiently accurate for the purpose of choosing simulation points, likely because it appears to be rather insensitive to the values of μ_σ and μ_δ or the light quark mass. As a result, values of β and c_{sw} can be found which correspond to a given lattice spacing within a few per-cent and to the tadpole-improved value of c_{sw} within the available perturbative precision. In addition, the estimate of the critical hopping parameter κ_c at these values of β and c_{sw} is seen to be accurate to within around half a per-mille for light quarks corresponding to a pion mass of around 250 MeV.

Simulations with light, strange and charm quarks in the sea were carried out with these parameter sets with the aim of studying the impact of the value of c_{sw} on the simulation algorithm and the pion mass splitting. It was found that the rational hybrid Monte Carlo algorithm works well for the simulation of the non-degenerate strange and charm quark doublet and that no issues appear with regards to the positivity of the heavy quark determinant. Significant fluctuations are seen in the maximal eigenvalues of the non-degenerate twisted mass clover operator and these

make simulations somewhat more expensive because a generous eigenvalue interval must be chosen for the rational approximation. No issues with stability were observed, although simulations with approximately matched physical situations at different values of c_{sw} indicated that larger values of c_{sw} induce larger molecular dynamics forces and correspondingly lower acceptance and higher cost.

A set of tuning runs with different light quark masses was used to provide an estimate of the bare quark mass corresponding to the physical charged pion mass in four flavour simulations. Runs at this light quark mass were started, showing promising thermalisation properties and no problems with tuning to maximal twist. First mass degenerate four flavour simulations – historically the most problematic – were successfully tuned to maximal twist at several values of the bare light quark mass, providing compelling evidence for a substantial reduction of the pion mass splitting due to the clover term. The accumulated evidence seems to indicate that the tuning of $N_f = 2+1+1$ simulations at the physical charged pion mass is feasible and the simulations done so far do not seem to point to any practical problems.

The pion mass splitting was studied on $N_f = 2$ and $N_f = 2+1+1$ gauge configurations and it was found that compared to previous $N_f = 2$ simulations, at a charged pion mass of 250 MeV the splitting is reduced somewhat, although an exact comparison is difficult because of differing lattice spacings. In the $N_f = 2+1+1$ case, it was seen that the improvement amounts to almost a factor of two reduction at a charged pion mass of around 250 MeV, where the comparison is more straightforward on account of the size of the improvement. The dependence of the pion mass splitting on the value of c_{sw} was studied in the range $c_{sw} = 1.45$ to $c_{sw} = 2.05$ and approximately constant physical situation, showing no effect within the sizeable uncertainties. This suggests that if simulations with small to moderate values of c_{sw} are stable and show no indications of a large pion mass splitting, it might be beneficial to simulate at smaller values of c_{sw} than would be required for on-shell $\mathcal{O}(a)$-improvement for standard Wilson quarks. Smaller values of c_{sw} show smaller fluctuations in the forces due to the quark determinant and smaller fluctuations in the maximal eigenvalues of the non-degenerate twisted mass operator, thus reducing simulation cost.

Benchmark calculations of pseudoscalar meson masses and decay constants were carried out at the physical light quark mass, leading also to estimates of the light, strange and charm quark masses and their ratios. It was shown that in a simple computation, ignoring all finite size and lattice artefact corrections, compatibility with phenomenological determinations could be obtained. Only in the charm sector, where sizeable lattice artefacts are to be expected, were significant deviations from experiment observed. Their magnitude, however, does not preclude a well-behaved continuum extrapolation, although of course this will have to be confirmed. The statistical and systematic uncertainties so far obtained on phenomenologically highly relevant quantities like f_D, f_{D_s} and their ratio are still somewhat large. Obtaining lattice results with phenomenologically interesting uncertainties will likely require an analysis involving multiple light quark masses and interpolations using heavy meson χPT. When this is done, however, special care must be taken to consider the relative size of the light quark masses and lattice artefacts as sources of chiral symmetry breaking.

To tune the valence strange and charm quark masses, linear interpolations of all observables in the analysis were carried out and the ratios $M_K/M_{\pi^\pm}$ and $M_D/M_{\pi^\pm}$ were matched to their phenomenological values.

From the data it is seen that this simple approach is well-justified, especially since the interpolations are quite short. The resulting strange and charm quark masses and their ratios are shown to be compatible with various determinations in the literature after being multiplied by the appropriate RI'-MOM matching factor.

A particular source of systematic uncertainty which may become significant on very large lattices is the ambiguity of the fit range for correlation functions. Excited state contributions, round-off errors of heavy quark propagators as well as correlations in the lattice data can result in a situation where identifying a plateau in the effective mass of a given correlation function is non-trivial. In addition, even if a plateau is identified and a fit range chosen, analysis shows that there exist many fit ranges with similar reduced χ^2 values. It was argued, therefore, that all of these fit ranges should be considered and a method was proposed for estimating the resulting systematic uncertainty in meson masses, decay constants and their ratios. In most cases on a $L/a = 48$ lattice, the uncertainty was found to be small, but in the case of pion correlation functions, an effect at the order of the statistical error was seen, such that it should be taken into account. It is likely that the technique will become even more relevant for lattices of spatial extent $L/a = 64$ and above.

Significant importance was paid to exploring possible optimisations of the simulation algorithm. The parameter dependence of mass preconditioning was studied and simple empirical formulae with only two fit parameters were proposed which allow the forces of determinant and determinant ratio terms in the effective action to be expressed in terms of the mass preconditioning parameters. Although lacking a thorough theoretical motivation, when fitted to to data spanning many orders of magnitude, the resulting curves were seen to represent the data rather well. It is hoped that the formulae can be used in practice by fitting them to a few data points for a given simulation and using the resulting parameter dependence to tune the mass preconditioning scheme.

Based on these formulae, a mass preconditioning and integration scheme was suggested which was shown to be effective at reducing simulation cost by providing a more predictable hierarchy of molecular dynamics forces, even without knowledge of the values of the fit parameters. This work was complemented by a tentative analysis of fast Fourier transforms of long molecular dynamics histories of the forces due to various terms in the effective action. It was shown that the spectrum of frequencies which contribute substantially to the dynamics seems to depend inversely on the preconditioning masses in some, as yet unspecified way. Finally, it was suggested that these frequency spectra might be used to tune the number of integration steps in Hybrid Monte Carlo simulations with multiple time scales.

Despite the potential for efficiency gains from tuning the mass preconditioning scheme, the simulations are rather expensive and algorithmic advances will have to be made in order for a complete simulation programme to be viable. On lattices with spatial extent up to $L = 48$, the CG algorithm is rather competitive with twisted mass quarks, even at the physical charged pion mass. When the volume is increased, however, the usage of better solvers becomes a necessity. Inexact deflation [210] has been investigated by the ETMC, but the behaviour of the inner solver is inconsistent with the situation of standard Wilson quarks. This situation, likely caused because the twisted mass Dirac operator exhibits a shifted zero quark mass Wilson spec-

trum at maximal twist, is in the final stages of being resolved. Hierarchically deflated CG [211] might be ideally suited for twisted mass quarks because it operates directly on the normal equations, but its implementation is highly non-trivial, especially given the relative rigidity of the tmLQCD code-base. Similarly, adaptive aggregation based domain decomposition [212] would likely perform well, but the implementation is also quite involved. It should be noted that for the purpose of analysis, exact deflation and EigCG [213] are seen to work very well, although these are of course unsuitable for ensemble generation since the number of systems to be solved for each gauge configuration must be very large.

A number of essential additions were made to the tmLQCD software suite. Multi-threading was introduced into the entire code-base via OpenMP, following a strategy which aimed at preserving existing single-thread optimisations. At the same time, care was taken to make the resulting code easy to read and to maintain, especially by exploiting variable scoping rules for automatic management of thread-level data locality. The multi-threaded routines were carefully profiled using SCALASCA and a graphical representation was developed for the identification of bottlenecks. Using this information, a number of optimisations were introduced and their effectiveness confirmed through benchmarks. In addition, orphaned directives were explored as a means of reducing or eliminating the substantial thread management overhead. Benchmarks of a test implementation of the hopping matrix kernel using orphaned directives showed that the SCALASCSA profiles indeed provide relevant information for the identification of overheads. Unfortunately, the usage of orphaned directives makes the management of data locality significantly more complicated and it is currently unclear how they can be deployed more generally.

Optimisations for the IBM BlueGene/Q supercomputer architecture were introduced which, combined, led to a factor of six improvement in the hopping matrix kernel. The performance characteristics on this machine were carefully mapped and recommendations formulated for the efficient simulation on lattices of different volumes. In this context, particular effort was invested into the study of hopping matrix kernels which attempt to overlap communication and computation with the help of threads. An optimised overlapping kernel was devised for BlueGene/Q using SPI communicators and a naïve thread workload assignment.

The standard MPI kernel with overlapping communication and computation, hybrid parallelisation and optimised thread scheduling is unfortunately generally not as fast as pure MPI parallelisation. However, multi-threading helps significantly when scaling to very large machine partitions and threads allow to efficiently absorb prime factors in the number of cores, a circumstance which is becoming increasingly common.

Several techniques for testing the simulation algorithm were presented and demonstrated using specific examples from the tmLQCD software suite. A combined approach was explored for testing the molecular dynamics implementation. First, numerical derivatives are computed and then compared to their analytical counterparts to check if the action and the derivatives are implemented consistently. In addition, the expectation value of some observable is expanded around the vanishing limit of a relevant parameter in the action and an expression is derived for the derivative of this expectation value with respect to the parameter. This derivative can then be computed in a partially quenched manner given a sufficiently strong signal and checked against the value obtained directly from simulations at non-zero

values of the relevant parameter. With the growing complexity of LQCD software packages in mind, high statistics runs were proposed as reliable integration testing frameworks, the use of which should become part of the development process of LQCD codes. Their effectiveness at discovering bugs which are otherwise very hard or impossible to identify was demonstrated with three examples encountered in tmLQCD, one of which had even gone unnoticed in production simulations, altough its effects were á posteriori seen to be quite substantial.

From the numerical side, continuum and infinite volume limits at the physical pion mass will be subject to two major challenges. The complexity of simulation software has reached a point at which software packages for LQCD cannot be written specifically for a given architecture any more. Coupled with the many architectures that are in use, this requires many different optimisations to be provided in a modular fashion in one software package. These different optimisations are often at odds with each other, resulting in a significant engineering challenge which will have to be overcome. A sign of the times is the emergence of dedicated libraries of extremely highly optimised, architecture-specific computational kernels such as BAGEL [202], QUDA [214] and QPhiX [203]. Except for BAGEL, these projects have many developers working on them practically full-time. Although it is possible and worthwhile to rely on them, providing the necessary interfaces is often non-trivial. When that is achieved, it is unclear to what extent they will remain supported by the original developers, potentially causing significant future maintenance overheads. These issues must be overcome if phenomenologically relevant simulations of tmLQCD at the physical pion mass are to be carried out.

The second and related complication which has yet to be dealt with properly is the increasing relevance of hybrid architectures as supercomputers are pushed towards the *exa-scale* using accelerators. A simple example where Amdahl's law is already having an impact is in the generation of random numbers. On Intel Xeon Phi, for instance, the low performance of individual execution units leads to significant idling as random numbers are generated by a single core. This may be alleviated by having more copies of the pseudorandom number generator, but this may induce problems with correlations. The alternative is to implement multi-threaded generators, but this is complicated significantly by the non-deterministic execution order of multi-threaded routines. Hybrid architectures pose a particular challenge for the tmLQCD simulation software because of its rigid implementation: using these new systems will require structural changes which amount to rewriting many of its basic components. Given these challenges, the testing frameworks that have been investigated will certainly prove useful.

In conclusion it can be said that $N_f = 2$ simulations with twisted mass clover quarks at the physical pion mass are certainly possible and they seem to produce promising physics results, although of course their relevance in light of $N_f = 2+1$ and $N_f = 2+1+1$ data is limited to the exploration of unquenching effects. The tuning of $N_f = 2+1+1$ simulations at the physical pion mass was shown to be unproblematic and apparently unaffected by complications from the interaction of multiple scales. Simulations at relatively coarse lattice spacings can proceed and are practical on the currently modest lattice volumes. Continuum and infinite volume limits, however, hinge on the successful implementation of better algorithms for the inversion of the twisted mass clover Dirac operator at maximal twist, especially for the calculation of molecular dynamics forces.

Without these additions, simulations on large lattices of size $L = 64$ and up are prohibitively expensive and not presently viable.

ACKNOWLEDGEMENTS

This thesis would not have been possible without the efforts and personal impact of a number of people whom I would like to acknowledge here. I hope that I can do them justice with my feeble words.

I would like to start by expressing my sincerest gratitude to Karl Jansen and Michael Müller-Preußker for giving me the opportunity to work under their supervision. It pains me greatly that I will not be able to discuss with Michael the result of his trust and patience. Whenever I was feeling overwhelmed, he had an anecdote to share that would cheer me up and I am deeply grateful for that. One of the last things he said to me before he passed away was: "Erstens kommt es anders, und zweitens als man denkt!". Even though I knew him only for a short time, I feel that he lived his life according to this rule and that he did so to the fullest. His jovial attitude, warmth of character, ability to make everyone feel welcome and of course his stupendous talent for formulating eloquent toasts at any imaginable occasion will always be missed. I'm sorry I can't thank you personally, Michael.

Four years ago I contacted him by accident, inquiring about PhD positions in phenomenology, and what a serendipitous accident it turned out to be. He handed me a copy of his trusty Rothe, which I read in leaps and bounds on the night before my interview, realising that this field could be my future. When I presented my master thesis to Karl and Michael, I had no idea that lattice QCD would shape my understanding of physics in the deep way that it has.

During my thesis work, Karl gave me the freedom to find my interests mostly on my own terms but was always available when I had questions. He was especially patient listening to my tirades when I was once again incensed at some circumstance or other. His critical reading of this thesis resulted in many important alterations for which I am grateful. Thank you, Karl, for allowing me to grow at my own pace.

I would like to express very special thanks to Carsten Urbach, whose impact on this thesis was tremendous. His many explanations and his guidance with the technical aspects of the analyses presented in this work were invaluable. Much of what was done here builds directly on his work: high quality, extensible and maintainable scientific software with good documentation. I believe that what Carsten does, is what it means to do reproducible frontier science based on numerical simulations, which can actually be peer-reviewed down to the last detail. Thank you, Carsten, I hope to follow your example.

I am also very grateful to Rainer Sommer for taking responsibility for my thesis with Humboldt-University when Michael passed away.

To all my colleagues at Humboldt University and DESY Zeuthen: thank you for the pleasant welcome that I was offered from the first day. In particular, I thank Grit, Otti and Florian for the many lunch discussions and our communal upset over social injustice. You made me feel at home.

To all the members of the European Twisted Mass Collaboration I would like to express my thanks for the many insightful discussions and for allowing me to rant during our meetings.

Outside of academia, a number of people have contributed to my personal and professional growth and I would like to acknowledge them here. My parents have played a very special role in nurturing my approach to everything in life and this thesis was no exception. You always cherished my opinion and included me in decisions, building my sense of worth and self-confidence. Never did you let me feel that I could not count on your support or guidance, even when my failings as a person made it very hard on you. For this, I am and will always be deeply grateful, I love you both.

My high-school chemistry teacher, John Hall, gifted me "The Pleasure of Finding Things Out" by Richard Feynman. His encouragement and Feynman's wit convinced me to follow my desire for understanding and to enrol in a physics degree. Thank you, Mr. Hall, wherever you are now. I don't have the understanding, but at least I now know how to increase the amount of stuff that I don't know.

My former partner Sofia will always have a very special place in my life for many reasons. A particularly important one is the impact that our time together had on how I approach and think about fellow humans. Science is a human enterprise and thrives most when we collaborate deeply with others, your lessons in empathy and understanding have surely made me a better human being and thus a better scientist. Thank you, darling.

I would also like to thank Ulrike, whose friendship, infectious laugh and amazing sense of humour gave me strength in my darkest moments. Thank you, Uljaana.

Finally, I would like to express my gratitude to all the people who had an indirect impact on this work. I thank the Luxembourgish tax payer for giving me the opportunity of working on my thesis in the pleasant environments at Humboldt-University and DESY Zeuthen. Without the *Aide Formation Recherche* grant through the Fonds National de la Recherche, finding a PhD place would have been much harder. I would also like thank the staff at the FNR for always handling all issues with utmost care and professionalism.

The calculations for this thesis were carried out on supercomputers and clusters managed by JSC Jülich, HLRN Berlin and Hannover, LRZ München, CINECA Bologna and DESY Zeuthen. I am indebted to the staff of these computing centres for their support, especially at times when the technical challenges of running these machines were at their peak. I used the R language for statistical computing for most analyses in this thesis and would like to acknowledge the many contributors to this project as well as the authors of the many add-on packages that I have used.

A

EVEN-ODD PRECONDITIONED TWISTED MASS CLOVER DIRAC OPERATOR

MASS DEGENERATE TWISTED MASS CLOVER DOUBLET

The implementation of the even-odd preconditioned Wilson Dirac operator in the presence of the clover term, the related even-odd preconditioned effective action and the integration of the equations of motion have been discussed in Ref. [64]. The tmLQCD software suite largely follows the strategy outlined there, but a subtlety arises when the twisted mass term is added. Starting from the fermion action in the hopping parameter representation in the twisted basis

$$\begin{aligned} S_{F,2} = \sum_x \Big\{ & \bar{\chi}(x)[1 + 2\kappa c_{sw} T(x) + 2ia\kappa\mu\gamma_5\tau^3]\chi(x) \\ & - \kappa\bar{\chi}(x) \sum_{\mu=1}^{4} \Big[U_\mu(x)(r+\gamma_\mu)\chi(x+a\hat{\mu}) \\ & + U_\mu^\dagger(x-a\hat{\mu})(r-\gamma_\mu)\chi(x-a\hat{\mu}) \Big] \Big\} \\ \equiv & \sum_{x,y} \bar{\chi}(x) M_{xy} \chi(y) \,, \end{aligned} \tag{A.1}$$

with the clover term $T(x)$. For convenience, the definitions

$$\tilde{\mu} \equiv 2a\kappa\mu \,, \qquad \tilde{c}_{sw} \equiv 2\kappa c_{sw} \,,$$

will be used and the SW coefficient will be absorbed

$$T = \tilde{c}_{sw} T(x) \,.$$

The hermitian, two flavour twisted mass clover operator, Q, is given by

$$Q \equiv \gamma_5 M = \begin{pmatrix} Q^+ & \\ & Q^- \end{pmatrix} , \tag{A.2}$$

where the subscripts refer to the sign of the twisted mass term from the τ^3 in flavour space. Decomposing the lattice into even and odd sites, the Schur decomposition of the sub-matrices $Q^\pm$ follows

$$\begin{aligned} Q^\pm &= \gamma_5 \begin{pmatrix} 1 + T_{ee} \pm i\tilde{\mu}\gamma_5 & M_{eo} \\ M_{oe} & 1 + T_{oo} \pm i\tilde{\mu}\gamma_5 \end{pmatrix} = \gamma_5 \begin{pmatrix} M^\pm_{ee} & M_{eo} \\ M_{oe} & M^\pm_{oo} \end{pmatrix} \\ &= \begin{pmatrix} \gamma_5 M^\pm_{ee} & 0 \\ \gamma_5 M_{oe} & 1 \end{pmatrix} \begin{pmatrix} 1 & (M^\pm_{ee})^{-1} M_{eo} \\ 0 & \gamma_5 \left(M^\pm_{oo} - M_{oe} (M^\pm_{ee})^{-1} M_{eo} \right) \end{pmatrix} , \end{aligned} \tag{A.3}$$

where the even and odd subscripts, e and o, have their obvious meanings. When the clover term is not present, the same decomposition applies and $(M^\pm_{ee})^{-1}$ is simply a $\tilde{\mu}$-dependent constant. With it present, it depends on the gauge field and must be inverted numerically, which will be discussed below. Unlike in the standard Wilson case, this inverse now also depends on $\pm\tilde{\mu}$.

The quark determinant splits into the factors

$$\det(Q) = \det(Q^+)\det(Q^-)\,,$$

and the $\det(Q^\pm)$ are proportional to

$$\det(Q^\pm) \propto \det(1 + T_{ee} \pm i\tilde{\mu}\gamma_5)\det(\hat{Q}^\pm)\,, \tag{A.4}$$

where $\hat{Q}^\pm$ is referred to as the even-odd preconditioned Dirac operator and is defined only on odd lattice sites

$$\hat{Q}^\pm = \gamma_5\left[(1 + T_{oo} \pm i\tilde{\mu}\gamma_5) - M_{oe}(1 + T_{ee} \pm i\tilde{\mu}\gamma_5)^{-1}M_{eo}\right]\,. \tag{A.5}$$

Hence,

$$\det(Q) = \|\det(1 + T_{ee} + i\tilde{\mu}\gamma_5)\|^2\det(\hat{Q}^+\hat{Q}^-)\,,$$

and the contribution to the path integral weight is given by

$$e^{-S_{\text{eff}}} = \exp\left[-\sum_x \log\det M_{ee}(x)\right]\cdot\int \mathcal{D}\phi_o^\dagger\mathcal{D}\phi_o\,\exp\left[-\phi_o^\dagger\frac{1}{\hat{Q}^+\hat{Q}^-}\phi_o\right]\,, \tag{A.6}$$

where $M_{ee}(x)$ is a 24×24 matrix in colour, spin and flavour which splits into four 6×6 blocks with $A(x)$ and $B(x)$ individually hermitian

$$M_{ee}(x) = \begin{pmatrix} A(x) + i\tilde{\mu} & 0 & 0 & 0 \\ 0 & B(x) - i\tilde{\mu} & 0 & 0 \\ 0 & 0 & A(x) - i\tilde{\mu} & 0 \\ 0 & 0 & 0 & B(x) + i\tilde{\mu} \end{pmatrix}\,, \tag{A.7}$$

such that

$$\log\det M_{ee}(x) = \log\|\det[A(x) + i\tilde{\mu}]\|^2 + \log\|\det[B(x) + i\tilde{\mu}]\|^2\,. \tag{A.8}$$

The second term in Equation (A.6) is expressed in terms of pseudofermion fields defined only on odd sites. As a result, the inverses of $(1 + T_{ee}(x) + i\tilde{\mu}\gamma_5)$ and $(1 + T_{oo}(x) - i\tilde{\mu}\gamma_5)$ are needed only on even and odd sites respectively which is exploited to save memory and reduce the amount of computation required.

Mass Preconditioning

Mass preconditioning in the presence of the clover term has been discussed in Chapter 5. Adding the preconditioning parameter $\tilde{\rho} = 2a\kappa\rho$ simply as a shift in the twisted quark mass $\tilde{\mu}_\ell$ is complicated by the $\tilde{\mu}$-dependence of the inverse of M_{ee}, as this would now have to be computed for every $\tilde{\mu}$, $\tilde{\rho}$ combination in the action. For efficiency reasons, mass preconditioning in the presence of the clover term is thus implemented as follows. The even-odd preconditioned Dirac operator with mass preconditioning takes the form

$$\begin{aligned}\hat{W}^\pm(\tilde{\mu}_\ell, \tilde{\rho}) =& \gamma_5\left[1 + T_{oo} \pm i\gamma_5(\tilde{\mu}_\ell + \tilde{\rho}) - M_{oe}(1 + T_{ee} \pm i\gamma_5\tilde{\mu}_\ell)^{-1}M_{eo}\right] \\ =& \hat{Q}^\pm(\tilde{\mu}_\ell) \pm i\tilde{\rho}\,,\end{aligned} \tag{A.9}$$

such that M_{ee} and M_{oo} depend only on the target mass $\tilde{\mu}_\ell$. The even-odd preconditioned quark determinant with mass preconditioning is thus

$$\|\det(1 + T_{ee} + i\tilde{\mu}_\ell\gamma_5)\|^2\cdot\int\mathcal{D}\phi_o^\dagger\mathcal{D}\phi_o\exp\left[-\phi_o^\dagger\frac{1}{\hat{W}^+\hat{W}^-}\phi_o\right]\,. \tag{A.10}$$

MASS NON-DEGENERATE TWISTED MASS CLOVER DOUBLET

Even-odd preconditioning for the mass non-degenerate twisted mass clover doublet proceeds in the same fashion, but the two flavour operator does not split into blocks. Starting again with the hopping representation of the fermionic action

$$\begin{aligned} S_{F,ND} = \sum_x \Big\{ & \bar{\chi}_h(x)[1 + 2\kappa c_{sw} T(x) + 2i a\kappa\mu_\sigma\gamma_5\tau^3 - 2\kappa\mu_\delta\tau^1]\chi(x) \\ & - \kappa\bar{\chi}(x) \sum_{\mu=1}^{4} \Big[U_\mu(x)(r+\gamma_\mu)\chi(x+a\hat{\mu}) \\ & + U_\mu^\dagger(x-a\hat{\mu})(r-\gamma_\mu)\chi_h(x-a\hat{\mu}) \Big] \Big\} \\ \equiv & \sum_{x,y} \bar{\chi}(x) M^h_{xy} \chi(y)\,, \end{aligned} \tag{A.11}$$

and proceeding in exact analogy with Appendix A.1, defines the operator

$$Q^h = \gamma_5 M^h\,. \tag{A.12}$$

The even-odd decomposition is then given by

$$\begin{aligned} Q^h &= \begin{pmatrix} \gamma_5(1 + T_{ee} + i\tilde{\mu}_\sigma\gamma_5\tau^3 - \tilde{\mu}_\delta\tau^1) & Q^h_{eo} \\ Q^h_{oe} & \gamma_5(1 + T_{oo} + i\tilde{\mu}_\sigma\gamma_5\tau^3 - \tilde{\mu}_\delta\tau^1) \end{pmatrix} \\ &= \begin{pmatrix} Q^h_{ee} & 0 \\ Q^h_{oe} & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & (Q^h_{ee})^{-1} Q_{eo} \\ 0 & \hat{Q}^h_{oo} \end{pmatrix} \end{aligned} \tag{A.13}$$

where $\hat{Q}^h_{oo}$ is given by

$$\hat{Q}^h_{oo} = \gamma_5 (M^h_{oo} - (M^h_{oe}\,(M^h_{ee})^{-1}\,M^h_{eo})\,, \tag{A.14}$$

with

$$M^h_{oo|ee} = 1 + T_{oo|ee} + i\tilde{\mu}\gamma_5\tau^3 - \tilde{\mu}_\delta\tau^1\,, \tag{A.15}$$

and $\hat{Q}^h$ now satisfies τ^1 hermiticity

$$(\hat{Q}^h)^\dagger = \tau^1\ \hat{Q}^h\ \tau^1\,.$$

Because $1 + T_{ee}$ is unchanged and hermitian, the inverse of M^h_{ee} can be computed via

$$(1 + T_{ee} + i\tilde{\mu}_\sigma\gamma_5\tau^3 - \tilde{\mu}_\delta\tau^1)^{-1} = \frac{(1 + T_{ee} - i\tilde{\mu}_\sigma\gamma_5\tau^3 + \tilde{\mu}_\delta\tau^1)}{(1 + T_{ee})^2 + \tilde{\mu}_\sigma^2 - \tilde{\mu}_\delta^2}\,. \tag{A.16}$$

In practice, $((1 + T_{ee})^2 + \tilde{\mu}_\sigma^2 - \tilde{\mu}_\delta^2)^{-1}$ is computed and stored and $(M^h_{ee})^{-1}$ is applied in two stages in the current implementation.

The heavy quark determinant factorises into

$$\det(Q^h) = \det[\gamma_5(1 + T_{ee} + i\tilde{\mu}_\sigma\gamma_5\tau^3 - \tilde{\mu}_\delta\tau^1)] \cdot \det[\hat{Q}^h_{oo}]\,, \tag{A.17}$$

where the first factor is computed as above, via the 24×24 colour, spin and flavour matrix

$$M^h_{ee}(x) = \begin{pmatrix} A(x) + i\tilde{\mu}_\sigma & 0 & -\tilde{\mu}_\delta & 0 \\ 0 & B(x) - i\tilde{\mu}_\sigma & 0 & -\tilde{\mu}_\delta \\ -\tilde{\mu}_\delta & 0 & A(x) - i\tilde{\mu}_\sigma & 0 \\ 0 & -\tilde{\mu}_\delta & 0 & B(x) + i\tilde{\mu}_\sigma \end{pmatrix}, \tag{A.18}$$

where A and B are the same 6×6 matrices as in the mass degenerate case and are individually hermitian. This can be simplified by a Schur decomposition in flavour space:

$$\begin{aligned}
\det(M^h_{ee}) &= \det \begin{pmatrix} K & D \\ D & K^\dagger \end{pmatrix} = \det \left[\begin{pmatrix} K & D - KD^{-1}K^\dagger \\ D & 0 \end{pmatrix} \cdot \begin{pmatrix} 1 & D^{-1}K^\dagger \\ 0 & 1 \end{pmatrix} \right] \\
&= -\det(D) \cdot \det(D - KD^{-1}K^\dagger) \\
&= \det(KK^\dagger - D^2) \\
&= \det(A^2 + \tilde{\mu}_\sigma^2 - \tilde{\mu}_\delta^2) \cdot \det(B^2 + \tilde{\mu}_\sigma^2 - \tilde{\mu}_\delta^2)\,,
\end{aligned}$$

where the sign in the second line comes from the first term and in the third line the proportionality of D to the identity matrix was used. The second term, $\det(\hat{Q}^h_{oo})$ is computed through a rational approximation using standard methods [137, 138].

B

PSEUDOSCALAR INTERPOLATIONS

In this appendix, the interpolations in the valence strange and charm quark masses of all observables in Chapter 7 are shown, as computed on ensemble *cA2.09.48* at the physical charged pion mass. For quantities involving decay constants, the figure on the left hand side shows data with decay constants using the continuum meson dispersion relation while the right hand side shows the lattice dispersion relation. For those quantities which do not involve a decay constants, the estimate of $a\mu_s$ from matching M_K/f_K is computed with f_K in the continuum definition on the left and the lattice definition on the right.

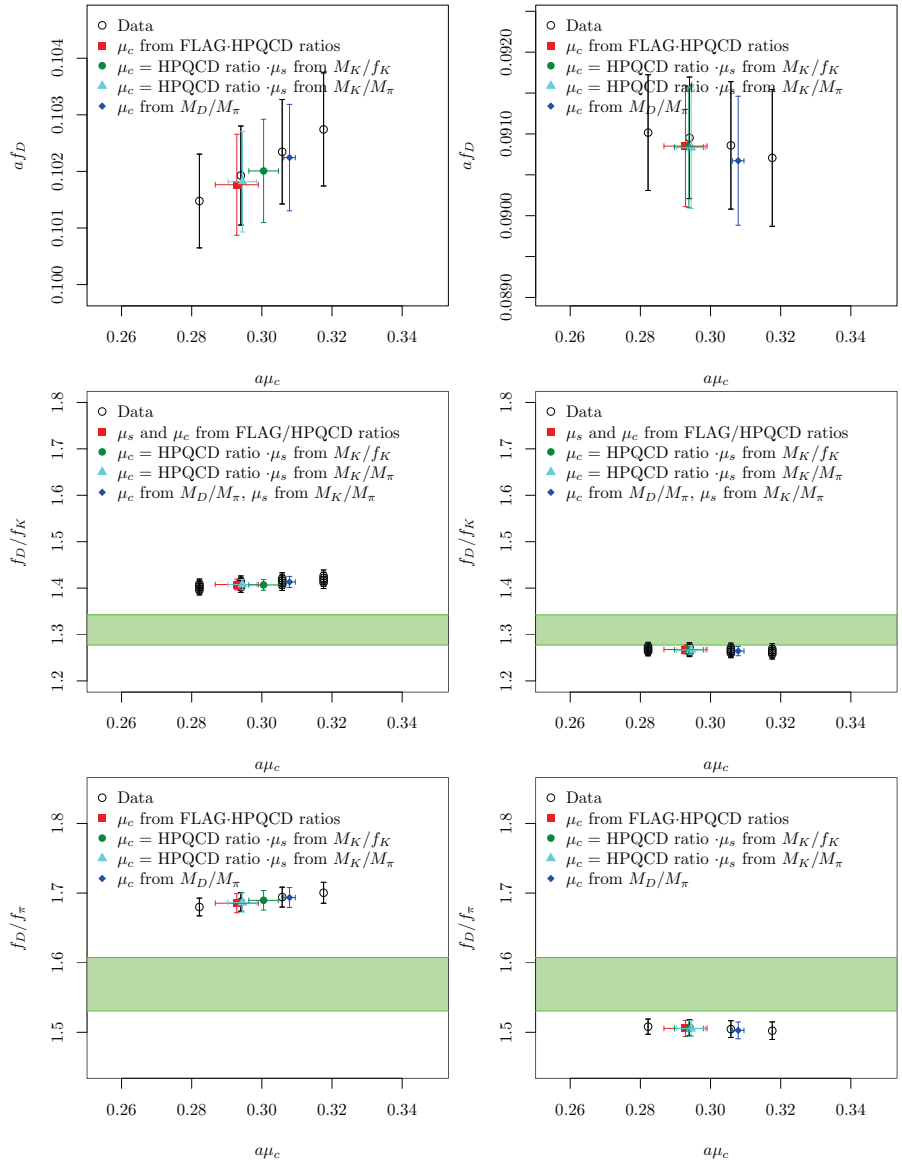

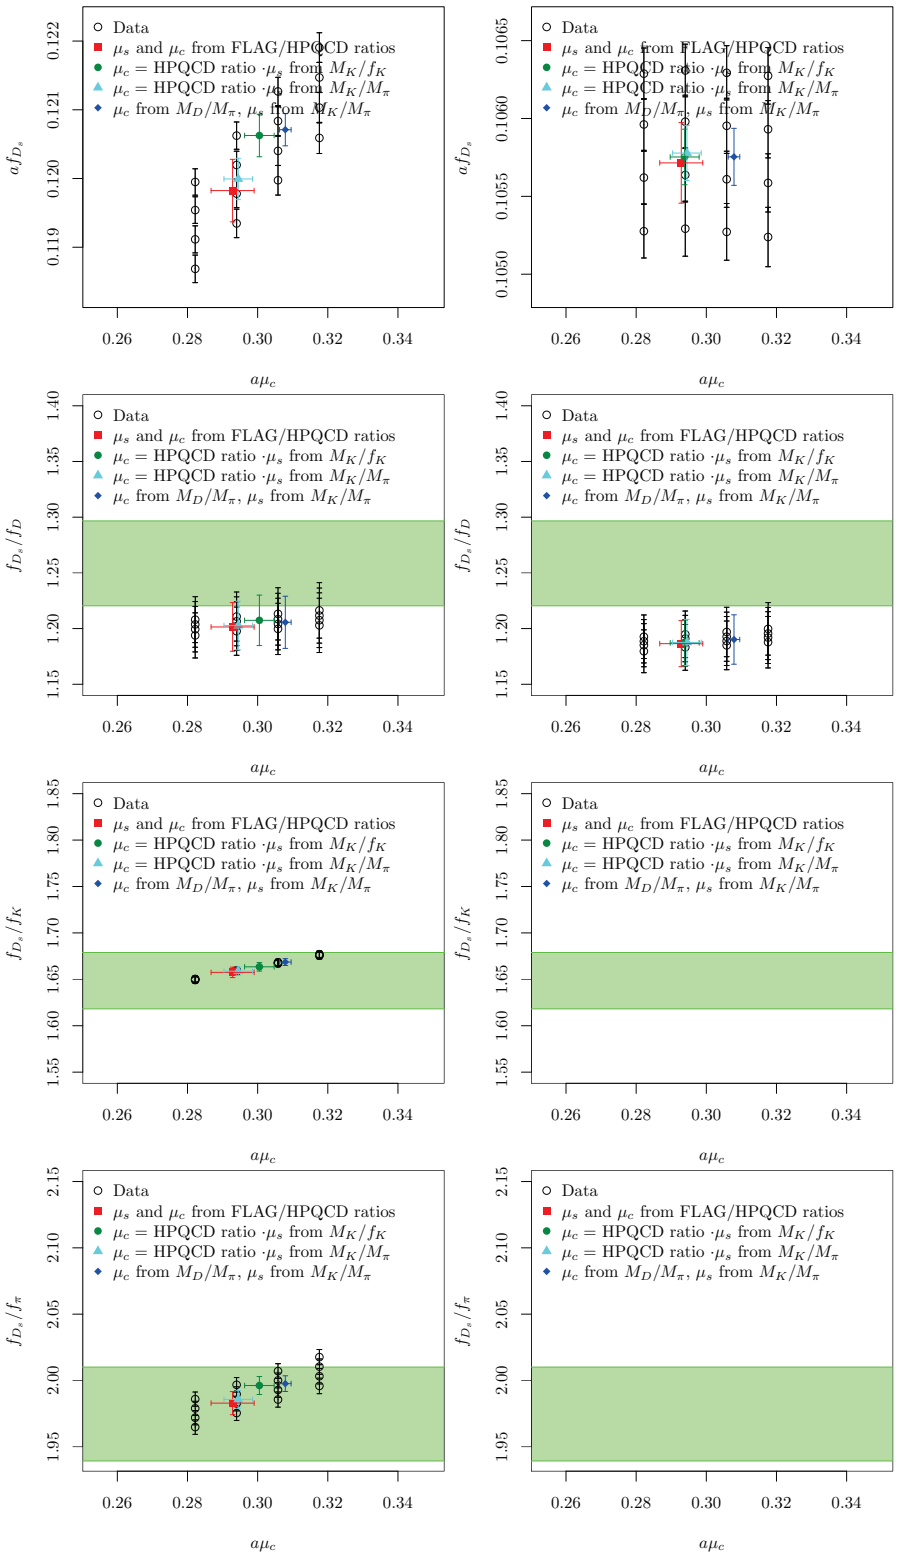
Data
μ_s and μ_c from FLAG/HPQCD ratios
μ_c = HPQCD ratio $\cdot\mu_s$ from M_K/f_K
μ_c = HPQCD ratio $\cdot\mu_s$ from M_K/M_π
μ_c from M_D/M_π, μ_s from M_K/M_π
af_{D_s}
f_{D_s}/f_D
f_{D_s}/f_K
f_{D_s}/f_π
$a\mu_c$

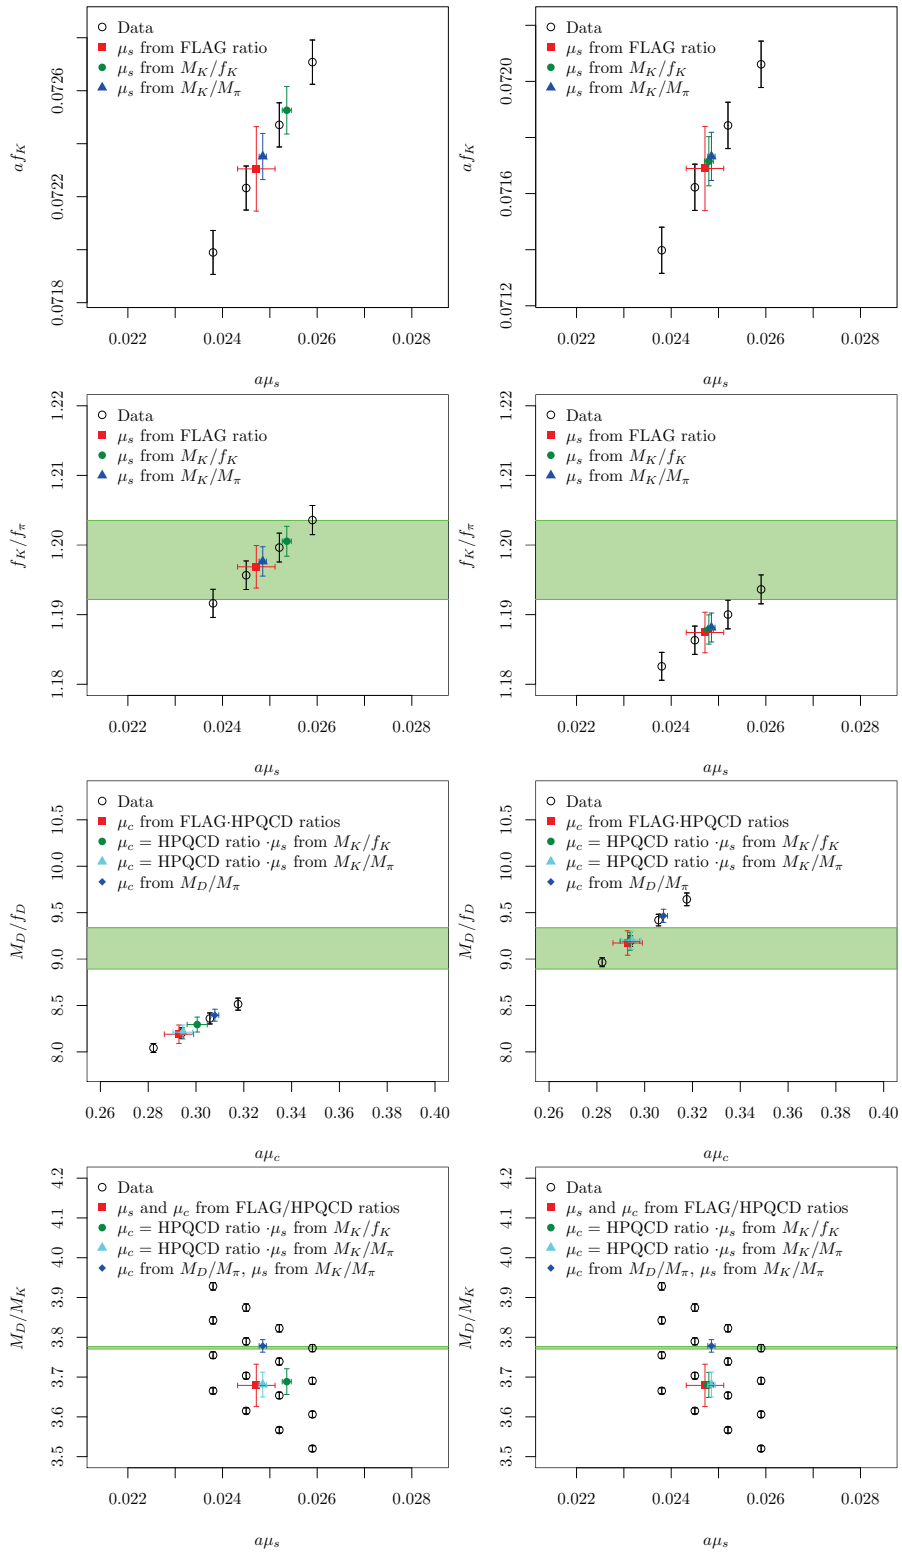
Data
μ_s from FLAG ratio
μ_s from M_K/f_K
μ_s from M_K/M_π
af_K
$a\mu_s$
f_K/f_π
μ_c from FLAG-HPQCD ratios
μ_c = HPQCD ratio $\cdot\mu_s$ from M_K/f_K
μ_c = HPQCD ratio $\cdot\mu_s$ from M_K/M_π
μ_c from M_D/M_π
M_D/f_D
$a\mu_c$
μ_s and μ_c from FLAG/HPQCD ratios
μ_c from M_D/M_π, μ_s from M_K/M_π
M_D/M_K

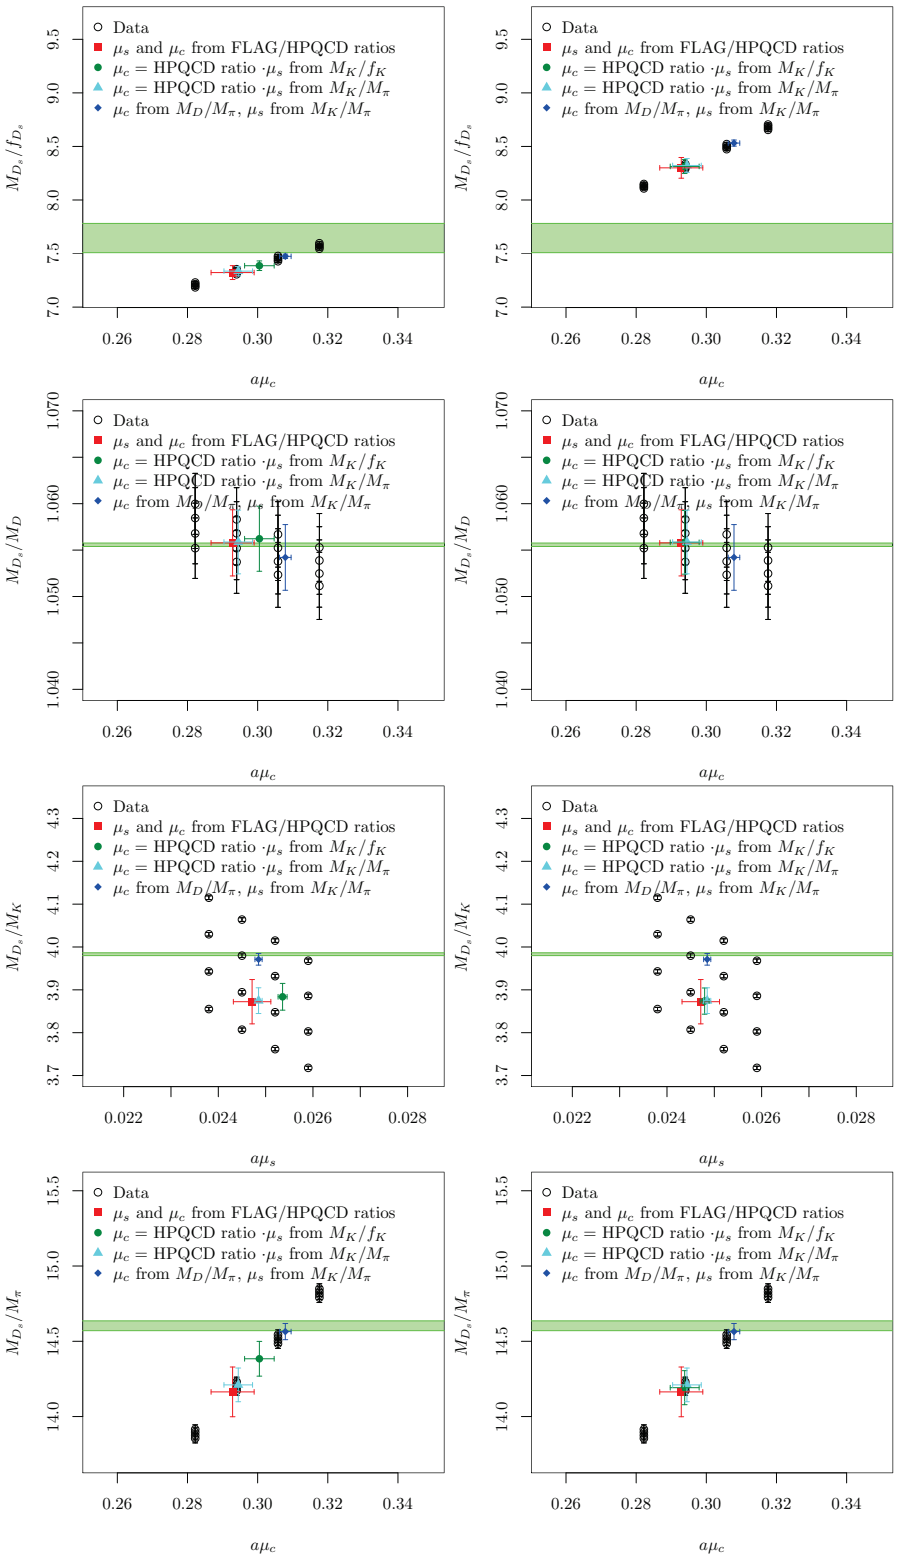
Data
μ_s and μ_c from FLAG/HPQCD ratios
μ_c = HPQCD ratio $\cdot\mu_s$ from M_K/f_K
μ_c = HPQCD ratio $\cdot\mu_s$ from M_K/M_π
μ_c from M_D/M_π, μ_s from M_K/M_π
M_{D_s}/f_{D_s}
M_{D_s}/M_D
M_{D_s}/M_K
M_{D_s}/M_π
$a\mu_c$
$a\mu_s$

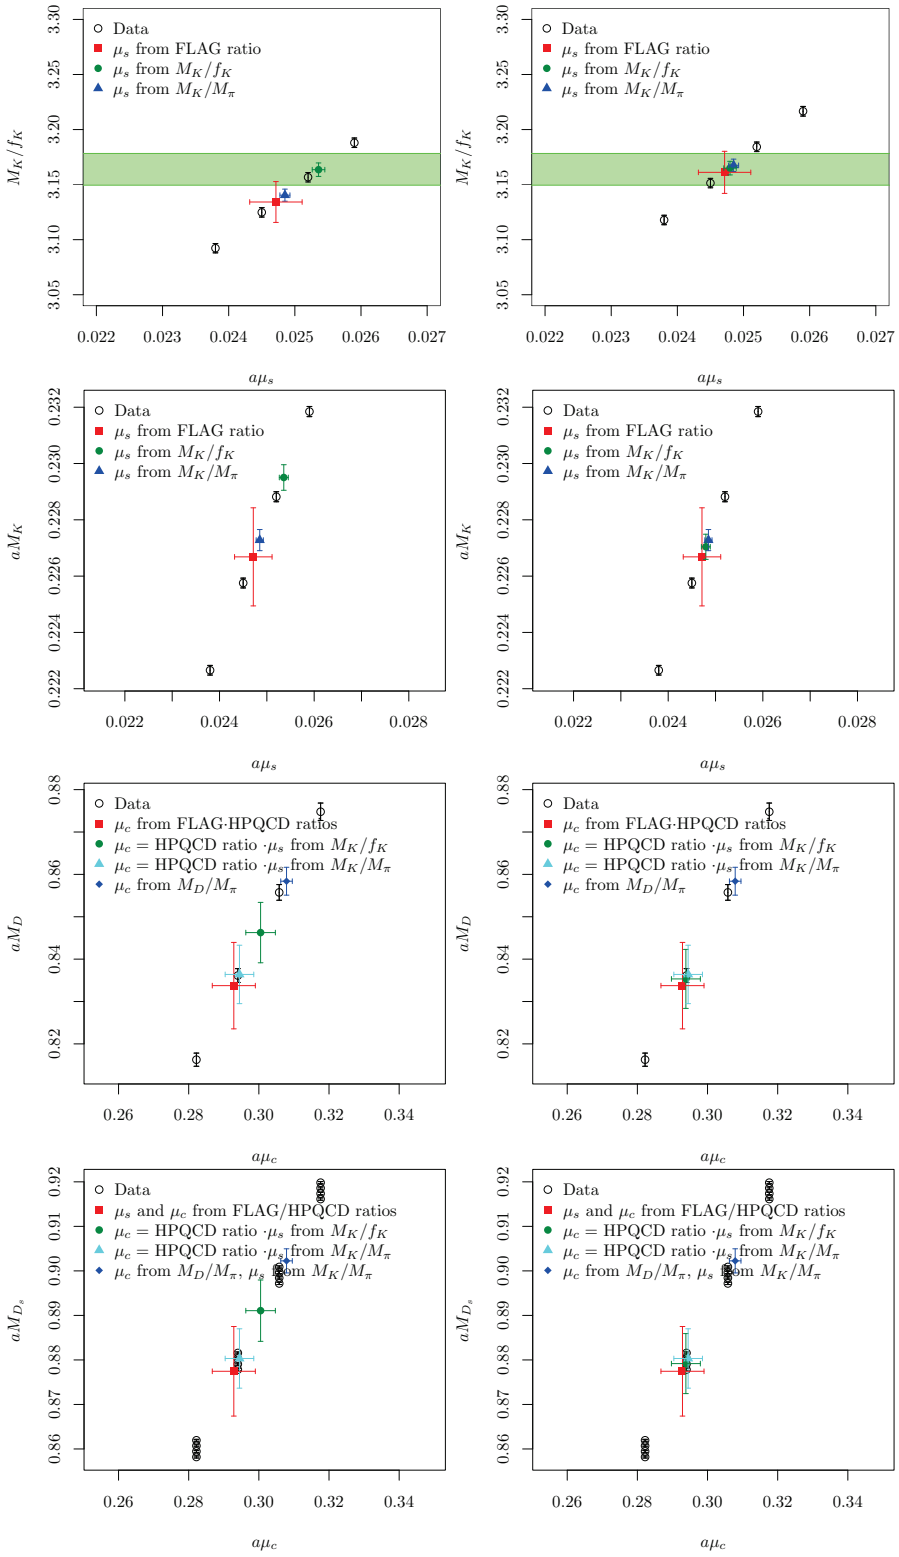
Data
μ_s from FLAG ratio
μ_s from M_K/f_K
μ_s from M_K/M_π
M_K/f_K
$a\mu_s$
aM_K
μ_c from FLAG·HPQCD ratios
μ_c = HPQCD ratio ·μ_s from M_K/f_K
μ_c = HPQCD ratio ·μ_s from M_K/M_π
μ_c from M_D/M_π
aM_D
$a\mu_c$
μ_s and μ_c from FLAG/HPQCD ratios
μ_c from M_D/M_π, μ_s from M_K/M_π
aM_{D_s}

C

MD FORCES AS A FUNCTION OF THE SW COEFFICIENT

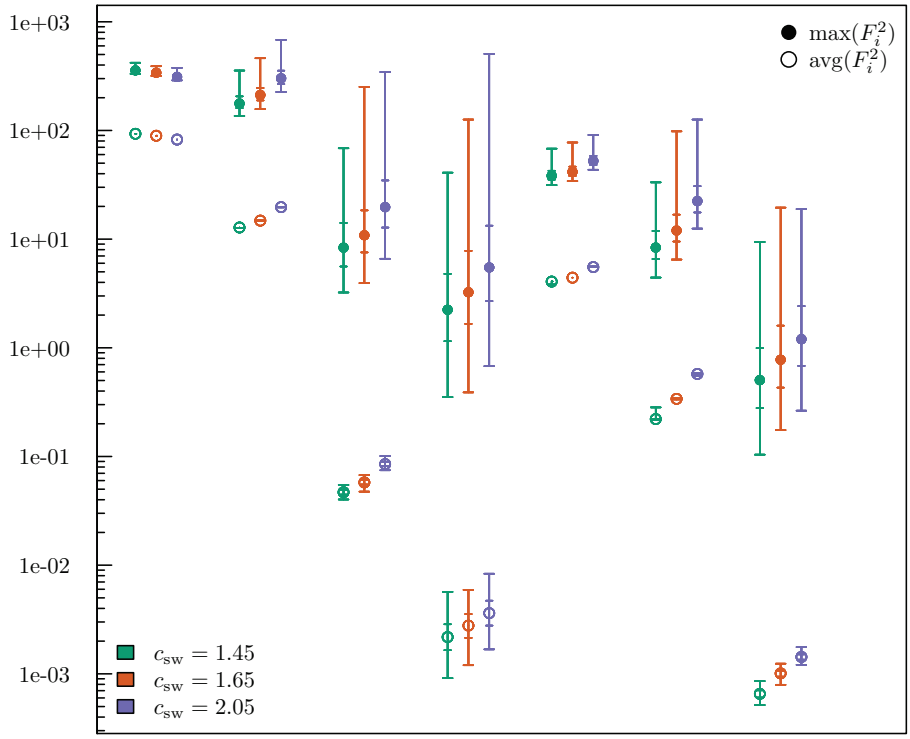

Figure C.1: Box-and-whiskers type plot of the distributions of average and maximal forces in $N_f = 2+1+1$ simulations at roughly constant physical situation as a function of c_{sw}. The quantiles are $0.0, 0.1573, 0.5, 0.8427, 1.0$. From left to right, the groups represent: **(1)** gauge, **(2)** light degenerate doublet determinant with largest preconditioning mass $2a\kappa\rho = 0.1$, **(3)** determinant ratio with $2a\kappa\rho_i = \{0.01, 0.1\}$, **(4)** determinant ratio with target quark mass in numerator, **(5)** partial fractions $k = 0, 1, 2$ of non-degenerate strange-charm doublet, **(6)** partial fractions $k = 3, 4, 5$, **(7)** partial fractions $k = 6, 7, 8$.

This appendix collects a number of force measurements from three of the four ensembles used for the measurement pion mass splitting as a function of c_{sw} in Section 6.2. Since the physical situation is tuned to be approximately constant, the volumes and lattice spacings are comparable and so is the dynamical quark content. The simulations were almost of the same length such that a sufficient number of force measurements can usefully be compared via histograms. In addition, all the algorithmic settings were exactly the same between the three ensembles, except for the eigenvalue approximation interval for the rational approximation at $c_{sw} = 2.05$. The upper end of the approximation interval had to be increased substantially compared to the ensembles at $c_{sw} = 1.45, 1.65$ because the maximum eigenvalues seem to scale with c_{sw}, likely stronger than linearly.

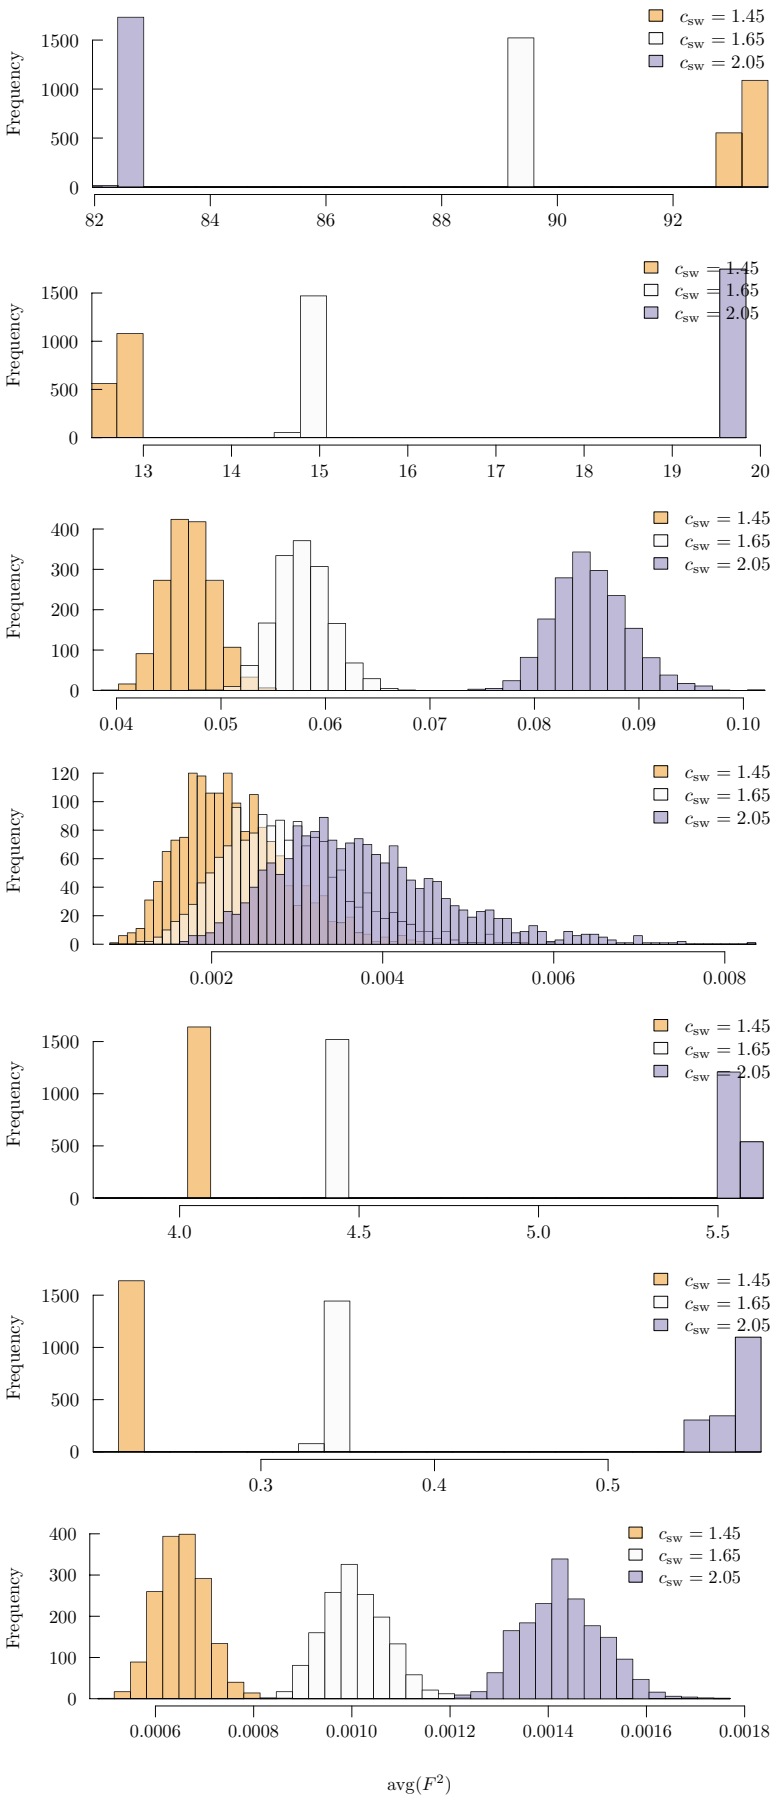

Figure C.2: Distribution of average forces as a function of c_{sw} in $N_f = 2+1+1$ simulations. From top to bottom: **(1)** gauge, **(2)** light degenerate doublet determinant with largest preconditioning mass $2a\kappa\rho = 0.1$, **(3)** determinant ratio with $2a\kappa\rho_i = \{0.01, 0.1\}$, **(4)** determinant ratio with target quark mass in numerator, **(5)** partial fractions $k = 0, 1, 2$ of non-degenerate strange-charm doublet, **(6)** partial fractions $k = 3, 4, 5$, **(7)** partial fractions $k = 6, 7, 8$.

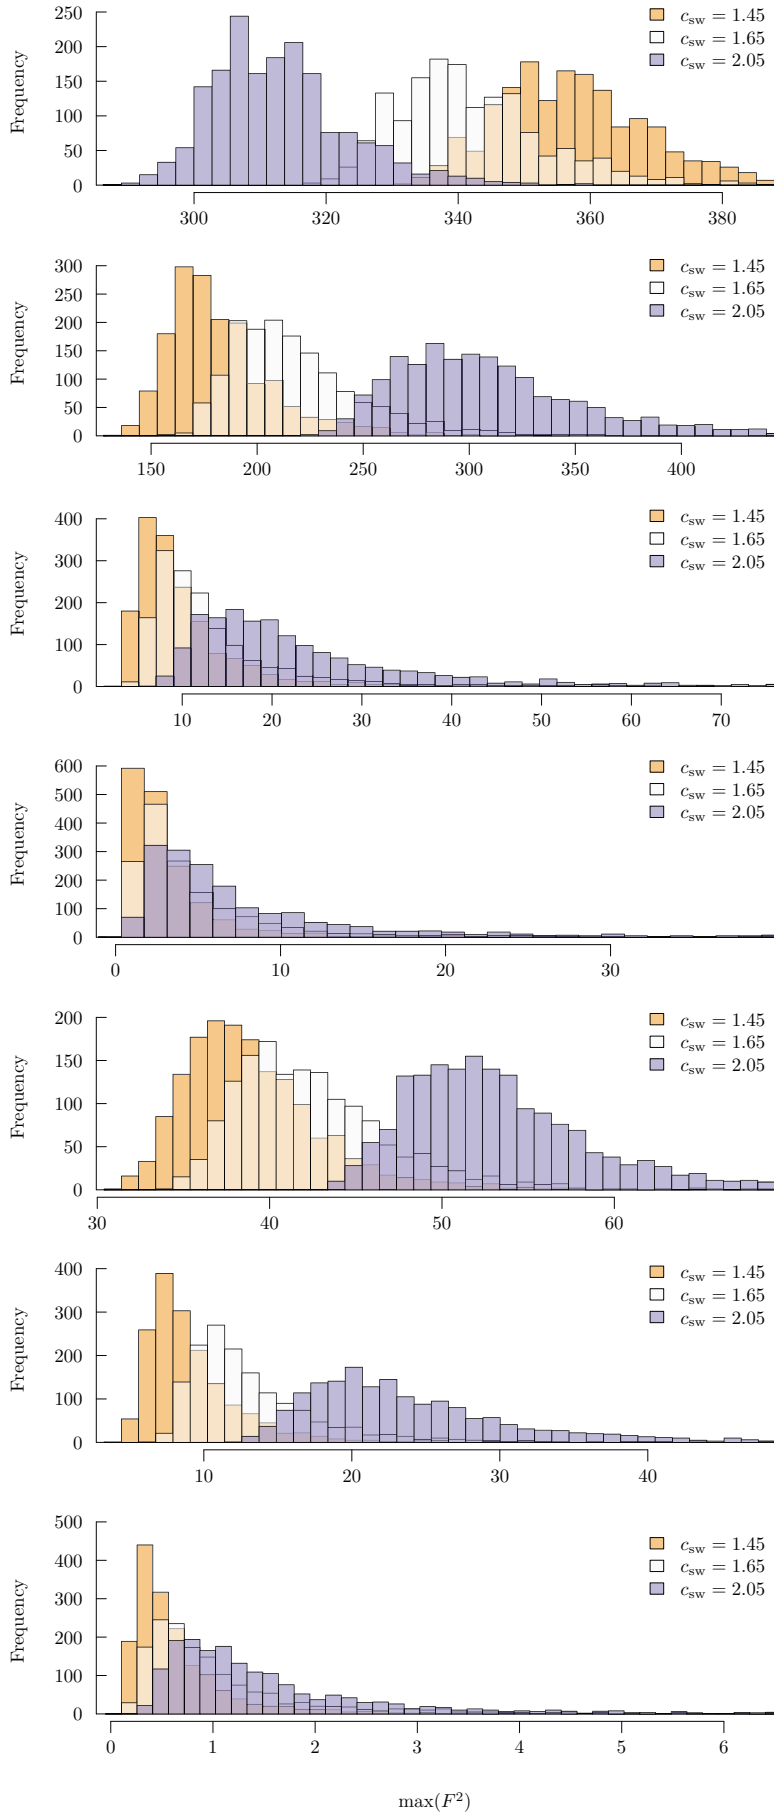

Figure C.3: Distribution of maximum forces as a function of c_{sw} in $N_f = 2+1+1$ simulations. From top to bottom: **(1)** gauge, **(2)** light degenerate doublet determinant with largest preconditioning mass $2a\kappa\rho = 0.1$, **(3)** determinant ratio with $2a\kappa\rho_i = \{0.01, 0.1\}$, **(4)** determinant ratio with target quark mass in numerator, **(5)** partial fractions $k = 0, 1, 2$ of non-degenerate strange-charm doublet, **(6)** partial fractions $k = 3, 4, 5$, **(7)** partial fractions $k = 6, 7, 8$.

D

ADDITIONAL FORCE MEASUREMENTS

DETERMINANT FORCES

In this section, further measurements of the square norm forces of a mass-degenerate even-odd preconditioned twisted mass clover determinant on the test ensemble presented in Chapter 5. Complementary to the results presented in Section 5.1.1, the measurements were repeated varying $\tilde{\rho}$ while keeping $\tilde{\mu}_\ell$ constant, which is actually the situation encountered in practice.

Figure D.1 shows the analysis resutl for $\tilde{\mu}_\ell = 0.0010984$, resulting in fits of the form

$$\|F\|^2_{\text{av}}(\tilde{\mu}) = \frac{4.59(16)}{\tilde{\mu} + 0.29(2)} \tag{D.1}$$

$$\|F\|^2_{\text{max}}(\tilde{\mu}) = \frac{39(5)}{\tilde{\mu} + 0.19(4)} \,. \tag{D.2}$$

Figure D.2 shows the analysis result for $\tilde{\mu}_\ell = 0.0002746$, resulting in fits of the form

$$\|F\|^2_{\text{av}}(\tilde{\mu}) = \frac{4.59(16)}{\tilde{\mu} + 0.29(2)} \tag{D.3}$$

$$\|F\|^2_{\text{max}}(\tilde{\mu}) = \frac{39(5)}{\tilde{\mu} + 0.19(4)} \,. \tag{D.4}$$

Finally, Figure D.3 shows the analysis results for $\tilde{\mu}_\ell = 0.0$, resulting in fits of the form

$$\|F\|^2_{\text{av}}(\tilde{\mu}) = \frac{4.59(17)}{\tilde{\mu} + 0.29(2)} \tag{D.5}$$

$$\|F\|^2_{\text{max}}(\tilde{\mu}) = \frac{39(5)}{\tilde{\mu} + 0.19(4)} \,. \tag{D.6}$$

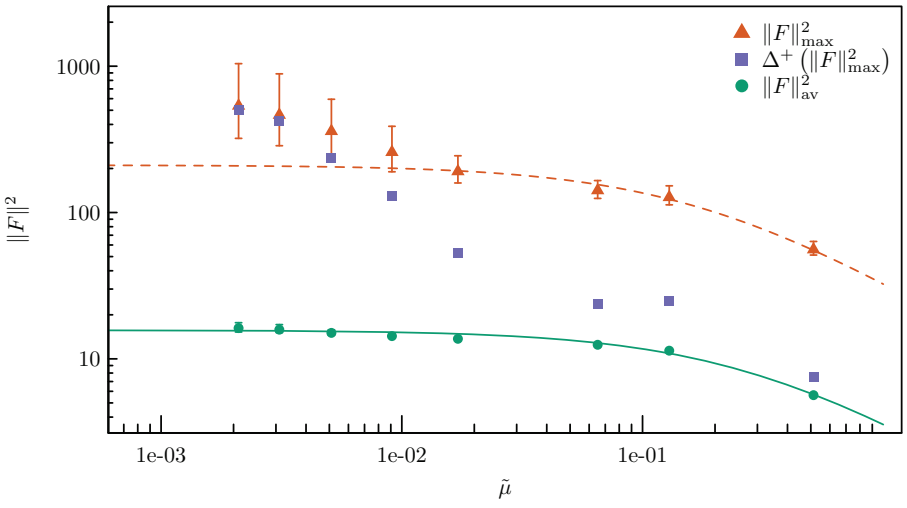

Figure D.1: Square norm of the forces from an even-odd preconditioned determinant for two flavours of mass degenerate twisted mass quarks at maximal twist as a function of the mass parameter $\tilde{\mu} = \tilde{\mu}_\ell + \tilde{\rho} = 2a\kappa(\mu_\ell + \rho)$, where $2\kappa\mu_\ell = 0.0010984$ and $\tilde{\rho}$ is varied. The notation Δ^+ indicates the difference between the median and 84.27 percentile of the maximal square norm force.

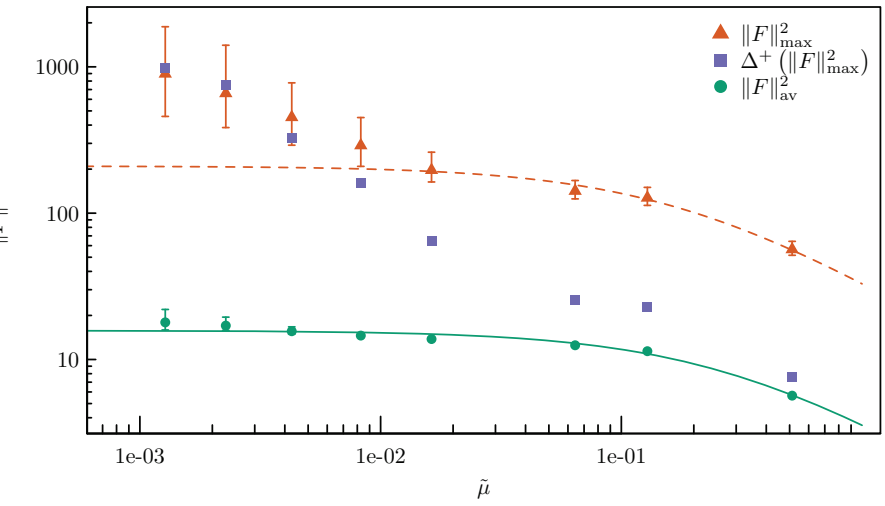

Figure D.2: Square norm of the forces from an even-odd preconditioned determinant for two flavours of mass degenerate twisted mass quarks at maximal twist as a function of the mass parameter $\tilde{\mu} = \tilde{\mu}_\ell + \tilde{\rho} = 2a\kappa(\mu_\ell + \rho)$, where $2\kappa\mu_\ell = 0.0002746$ and $\tilde{\rho}$ is varied. The notation Δ^+ indicates the difference between the median and 84.27 percentile of the maximal square norm force.

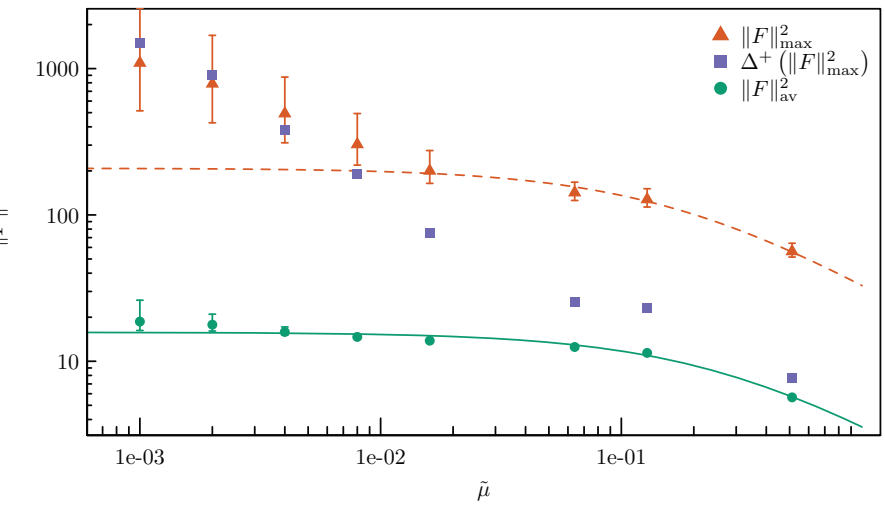

Figure D.3: Square norm of the forces from an even-odd preconditioned determinant for two flavours of mass degenerate twisted mass quarks at maximal twist as a function of the mass parameter $\tilde{\mu} = \tilde{\mu}_\ell + \tilde{\rho} = 2a\kappa(\mu_\ell + \rho)$, where $2\kappa\mu_\ell = 0.0$ and $\tilde{\rho}$ is varied. The notation Δ^+ indicates the difference between the median and 84.27 percentile of the maximal square norm force.

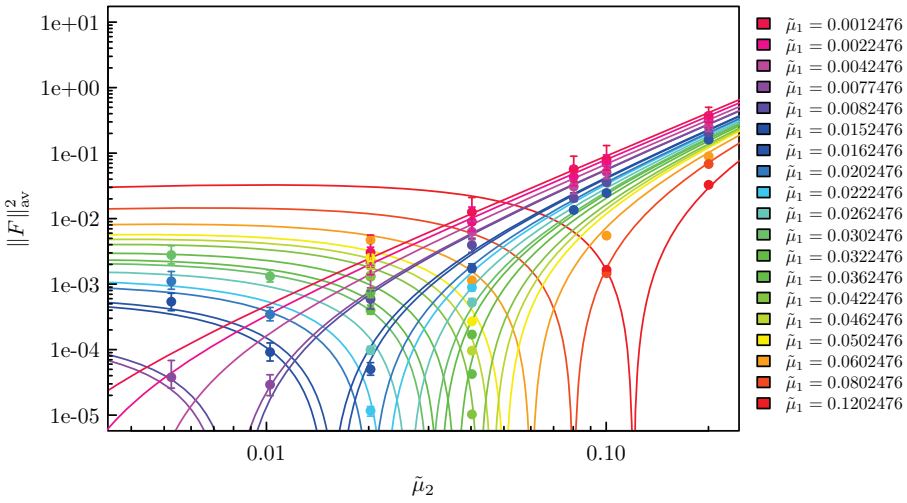

Figure D.4: Average forces coming from a determinant ratio as a function of the mass preconditioning parameters $\tilde{\mu}_1$ and $\tilde{\mu}_2$ on gauge configurations with the $N_f = 2$ twisted mass clover action. The lines correspond to the fit of the form Equation (5.12).

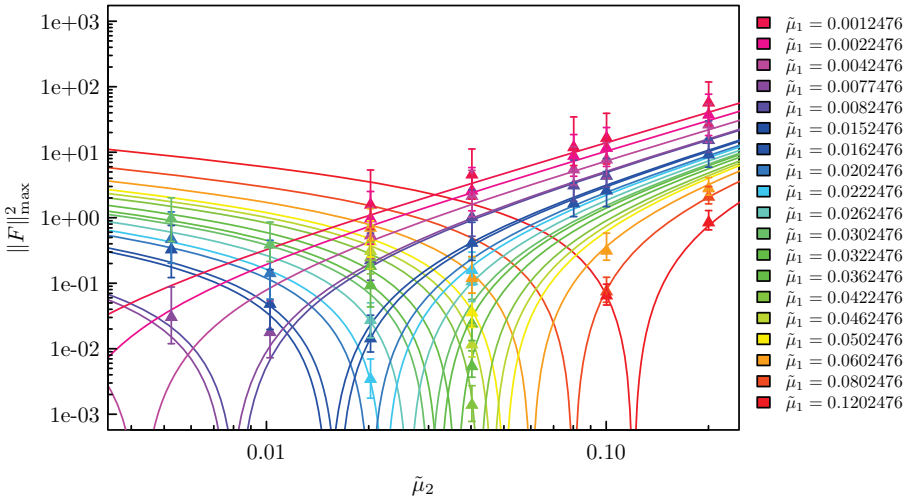

Figure D.5: Maximal forces coming from an even-odd preconditioned determinant ratio as a function of the mass preconditioning parameters $\tilde{\mu}_1$ and $\tilde{\mu}_2$ on gauge configurations with the $N_f = 2$ twisted mass clover action. The lines correspond to the fit in Equation (5.13).

DETERMINANT RATIO FORCES

For completeness, the results of Section 5.1.2 are shown in Figures D.4 and D.5 with $\tilde{\mu}_2$ on the abscissa.

Forces with a Higher Target Mass

As discussed in Section 5.1.2, when even-odd preconditioning is used with the twisted mass clover action, the target mass $\tilde{\mu}_\ell$ appears in the inverse of the clover term without a preconditioning mass $\tilde{\rho}$. The results in Section 5.1.2 were obtained with $\tilde{\mu}_\ell = 0.0002746$, relevant for simulations at the physical pion mass. For completeness, the analysis was repeated with a tenfold higher $\tilde{\mu}_\ell = 0.002746$ for a smaller range of $\tilde{\rho}_1$ and $\tilde{\rho}_2$. The results are given in Figures D.6 to D.11 and Equations (D.7) to (D.9).

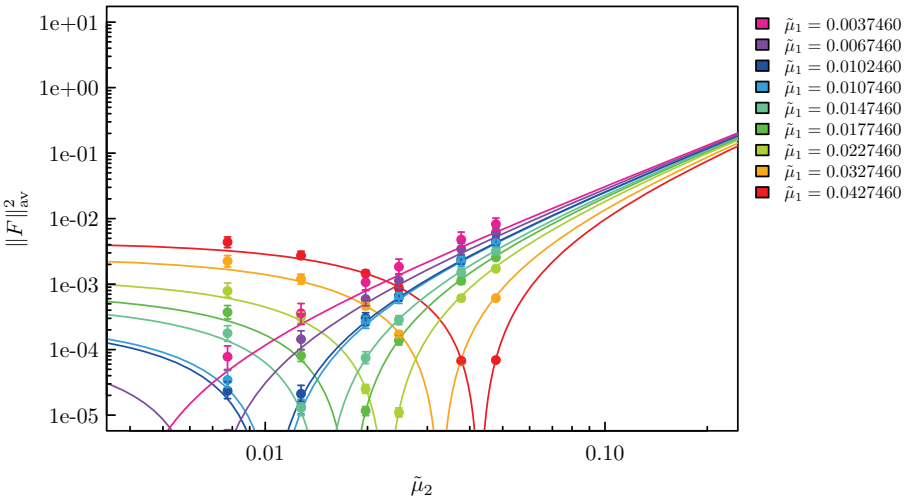

Figure D.6: Average forces coming from a determinant ratio as a function of the mass preconditioning parameters $\tilde{\mu}_1$ and $\tilde{\mu}_2$ on gauge configurations with the $N_f = 2$ twisted mass clover action. The lines correspond to the fit of the form Equation (D.7).

$$\|F\|^2_{\mathrm{av}}(\tilde{\mu}_1, \tilde{\mu}_2) = 3.21(9)\,(\tilde{\mu}_2 - \tilde{\mu}_1)^2 \left|\frac{\tilde{\mu}_2}{\tilde{\mu}_1}\right|^{0.07(1)} \tag{D.7}$$

$$\|F\|^2_{\mathrm{max}}(\tilde{\mu}_1, \tilde{\mu}_2) = 17(2)\frac{(\tilde{\mu}_2 - \tilde{\mu}_1)^2}{|\tilde{\mu}_1\tilde{\mu}_2|^{0.47(2)}} \tag{D.8}$$

$$\Delta^+\left(\|F\|^2_{\mathrm{max}}(\tilde{\mu}_1, \tilde{\mu}_2)\right) = 22(9)\frac{(\tilde{\mu}_2 - \tilde{\mu}_1)^2}{|\tilde{\mu}_1\tilde{\mu}_2|^{0.48(5)}} \tag{D.9}$$

Although more measurements would be necessary to draw firm conclusions, it appears that

- In $\|F\|^2_{\mathrm{av}}$, the phenomenological factor $|\tilde{\mu}_2/\tilde{\mu}_1|$ becomes less significant as $\tilde{\mu}_\ell$ is increased. The scale of the forces is somewhat reduced as expected.
- In $\|F\|^2_{\mathrm{max}}$, the overall scale is a little bit lower while the phenomenological factor is the same within errors, suggesting that this form may be universal.
- In $\Delta^+(\|F\|^2_{\mathrm{max}})$, both parameters take on very different values, but the suggested form does describe the data well.

One is thus tempted to conclude that the proposed fit forms are indeed appropriate and that as argued in Section 5.1.2, a few measurements should be sufficient to determine the parameters. In addition, the suggested simplified forms remain valid.

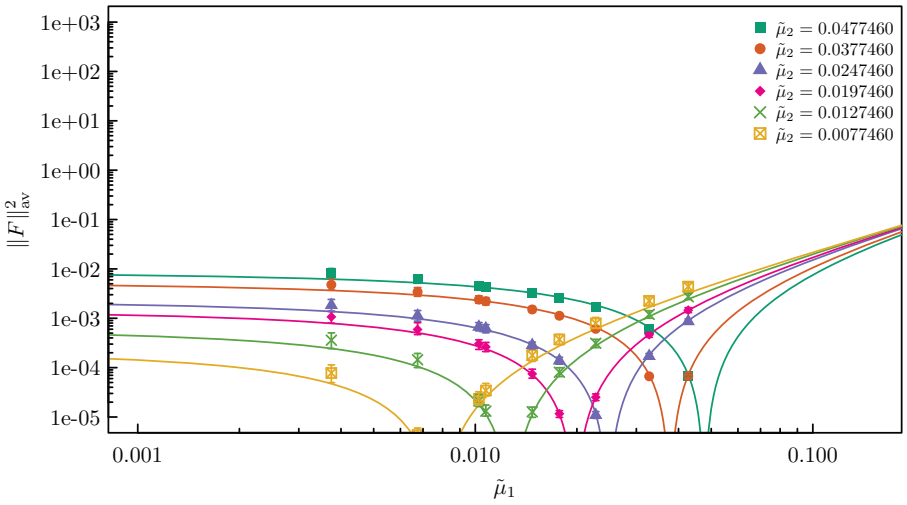

Figure D.7: Average forces coming from a determinant ratio as a function of the mass preconditioning parameters $\tilde{\mu}_1$ and $\tilde{\mu}_2$ on gauge configurations with the $N_f = 2$ twisted mass clover action. The lines correspond to the fit of the form Equation (D.7).

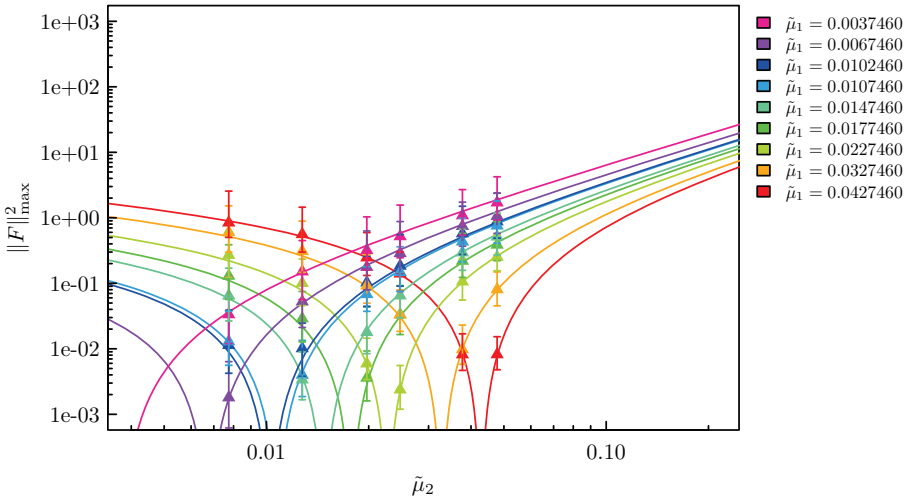

Figure D.8: Maximal forces coming from an even-odd preconditioned determinant ratio as a function of the mass preconditioning parameters $\tilde{\mu}_1$ and $\tilde{\mu}_2$ on gauge configurations with the $N_f = 2$ twisted mass clover action. The lines correspond to the fit in Equation (D.8).

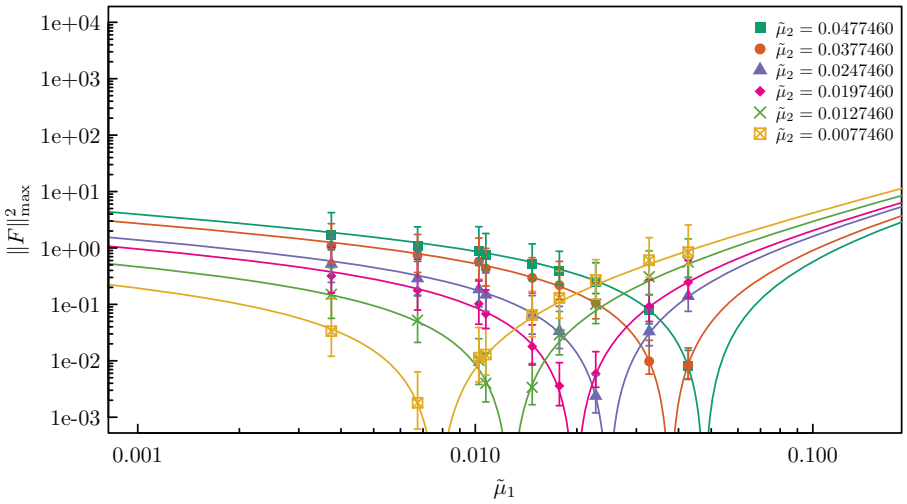

Figure D.9: Maximal forces coming from an even-odd preconditioned determinant ratio as a function of the mass preconditioning parameters $\tilde{\mu}_1$ and $\tilde{\mu}_2$ on gauge configurations with the $N_f = 2$ twisted mass clover action. The lines correspond to the fit in Equation (D.8).

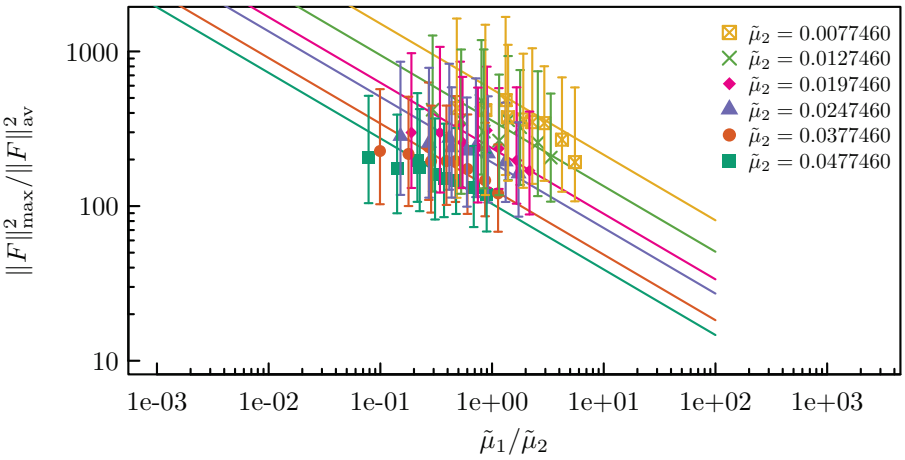

Figure D.10: Ratio of the maximal and average forces from even-odd preconditioned determinant ratios with different parameters. The lines give the ratio of the model in Equations (D.7) and (D.8) for constant $\tilde{\mu}_2$ as a function of $\tilde{\mu}_1/\tilde{\mu}_2$, as indicated by the colours.

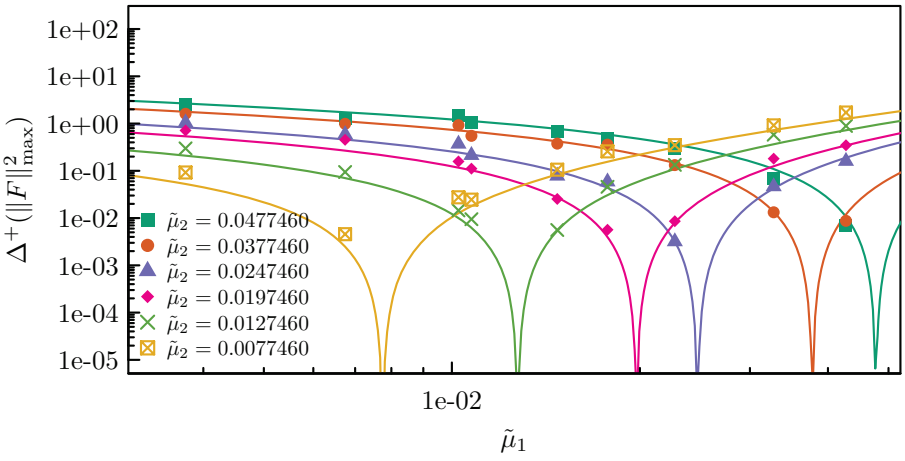

Figure D.11: Difference between 84.27 percentiles and the median of the maximal forces of even-odd preconditioned determinant ratios with different parameter combinations, the lines correspond to the functional form in Equation (D.9)

E

DEBUGGING CASE STUDIES

This appendix collects the result of a number of high statistics runs which uncovered three particularly subtle issues with the tmLQCD code-base. The issue described in Appendix E.3 even affected $N_f = 2+1+1$ production simulations which had to be repeated as a consequence. It is clear that these exact issues will not be encountered in different code-bases, but they do serve as very good case studies of the kinds of subtle faults which can be uncovered and which are otherwise very difficult or impossible to identify.

NON-DETERMINISTIC BUG IN CLOVER FORCE COMPUTATION

Parallelisation using threads is complicated chiefly by data locking issues as discussed in detail in Section 8.4. The computation of the derivative of the clover term required for molecular dynamics is the most non-trivial calculation in the code-base when it comes to the potential for data races. Even though the routine `sw_all` was reviewed and tested thoroughly during development and no apparent problems were found with 2, 4 or 8 threads on an Intel® Xeon™ machine, once simulations were started on IBM® BlueGene/Q™ with 64 threads, there was a puzzling loss of acceptance in the HMC. At the time it was of course not clear where this issue originated and the problem was only identified after the procedure discussed in the following.

It turns out that the issue was caused by a lack of proper memory locking in just a few lines of the routine mentioned above. The bug was not discovered because even though the conflicts in this function are not so rare even with just 2 threads, the non-determinism of the instruction order in a multi-threaded application makes it very difficult to predict by how much δH, for example, should differ between a multi-threaded and a serial execution. The very small deviations in δH that were seen on a trajectory to trajectory basis in initial testing, were attributed to this non-determinism. It should be noted that after the issue has been resolved, there continue to be small deviations in δH between serial and multi-threaded executions. However, they grow at a much smaller rate with the volume, and à posteriori, it is clear that these are now purely the result of round-off differences due to differences in execution order.

These kinds of failures are extremely dependent on the execution environment on which the probability of conflicts strongly depends and will therefore often appear with clear consequences only in production situations. They are difficult to find by inspection because multi-threading makes it much harder to reason about program execution order. Further, as explained above, under typical testing conditions it is often difficult to differentiate their consequences from benign differences in round-off. However, as will be shown below, high statistics runs can help in detecting this kind of issue early on and in a methodical manner, if they are conducted across a sufficiently wide set of parameter combinations (in this case this is the number of threads).

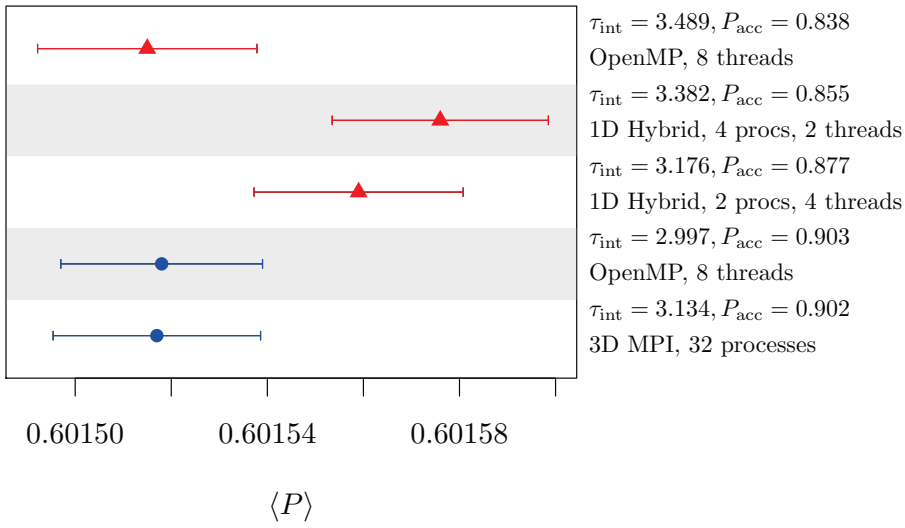

Figure E.1: Plaquette expectation values from different parallelisations of high-statistics runs with the $N_f = 2$ tmclover action indicating the effect of a data race in the derivative of the clover term. The red ▲ show the situation with the bug while the blue ● show the situation with the bug fixed. All the points are compatible within errors, but the different acceptance rates reveal that there must be an issue.

For the purpose of identifying the origin of the loss of acceptance, a set of runs with 10^5 trajectories each was set up with parameters mirroring the problematic situation but in smaller volume:

- Iwasaki gauge action, $V = 8^4$, $\beta = 2.1$, $\kappa = 0.13729$, $a\mu_\ell = 0.1$, $c_{\text{sw}} = 1.57551$
- Parallelisations: OpenMP with 8 threads, Hybrid with 2 processes à 4 threads, Hybrid with 4 processes à 2 threads and, as a reference, MPI using 32 processes and no multi-threading
- Random numbers chosen in such a way as to be completely identical between the different runs.

The plaquette expectation values from these runs are shown in Figure E.1 together with the respective integrated autocorrelation times and acceptance rates. As can be seen, all the points are compatible within errors and, most importantly, if there were just one simulation, one would not be able to see that there is something wrong. However, the acceptance rate decreases with as the number of concurrent threads is increased (red triangles in the figure). Once the bug was fixed, there is only a mild difference in autocorrelation times which can be understood because OpenMP effectively increases the amount of randomness in the simulation (and thus lowers the autocorrelation time).

This result indicates that the accept-reject step successfully rejected those trajectories where thread collisions led to deviations in the value of the derivative. It can also be interpreted further. The fact that the acceptance rate is reduced combined with the apparent lack of any bias points to the molecular dynamics part of the simulation and in particular, to the computation of derivatives. If there were a thread-related issue with the heatbath or acceptance steps, the simulated probability distributions would be different and bias would be bound to appear in the central values. Once one has established that the heatbath and acceptance steps are implemented correctly, one is led to conclude that the increase in energy violation responsible for the decrease in acceptance rate, must originate in the computation of the derivative.

A further conclusion one can draw from this test is that multi-threaded code can indeed fail in very subtle ways. Without a comprehensive testing

programme in place, wrong conclusions may be reached about the stability of simulations, autocorrelation times, acceptance rates or even about physical results. In this particular case, one might have come to the conclusion that the twisted mass clover action simply requires a somewhat higher number of integration steps to give good acceptance. This may even have allowed successful simulations (possibly significantly overestimating their cost), until some bias might have possibly surfaced in some other observable at a much later point in time. Clearly, high statistics runs are a valuable tool for identifying issues early on and with the added non-determinism of threads, they have become essential. Their central role in the identification of various problems will be further underlined below.

RANDOM NUMBERS IN PARALLEL APPLICATIONS

A very important consideration in any kind of simulation which makes use of stochastic properties is the quality and reliability of the pseudo-random number generators (PRNGs) used in the simulation. There exist numerous types of PRNG and the criteria for labelling them "good" or "bad" depend on the application. A comprehensive review of PRNG types, quality and testing methodology is given in Ref. [215] and references therein. The fact that PRNGs are crucial is best exemplified by the well-known PRNG-dependent systematic errors in Monte Carlo simulations of the Ising model, presented in Ref. [216]. Particular problems can be expected if the random number chains contain correlations between sequences or individual elements in a sequence, but also when chains generated with different starting values (seeds) show correlations of some type.

As an aside: the advent of hybrid architectures with high levels of shared memory parallelism presents a new challenge in this respect. Random number generation may become a bottleneck as the power of single execution units on these architectures is reduced while their number increases. As a result, unless a PRNG with hybrid parallelisation is used, many execution units will idle while one is generating random numbers. In fact, this has been observed when running tmLQCD on Intel® Xeon Phi™ in native mode. Overcoming this bottleneck is a non-trivial problem because the non-determinism inherent in the usage of threads presents a formidable challenge in the implementation of multi-threaded random number generators which need to be both reliable and fast. This should be considered an unsolved problem in the case of tmLQCD.

The need for massive parallelism in LQCD calculations further complicates the quality criteria that a PRNG has to satisfy because using PRNGs in parallel is not straightforward. In principle, three modes of operation can be considered:

1. A *sequential* mode which generates the same random numbers as in the sequential execution of a program by transferring the PRNG state from one process to the next in the right order, as depicted in the left panel of Figure E.2.

2. A *parallel* mode in which each process contains a random number generator which has been seeded independently and chains of random numbers are generated without any inter-process communication, as shown in the right panel of Figure E.2.

3. A generalisation of the parallel mode in which each sub-lattice of some size is associated with a PRNG. The most extreme such implementation

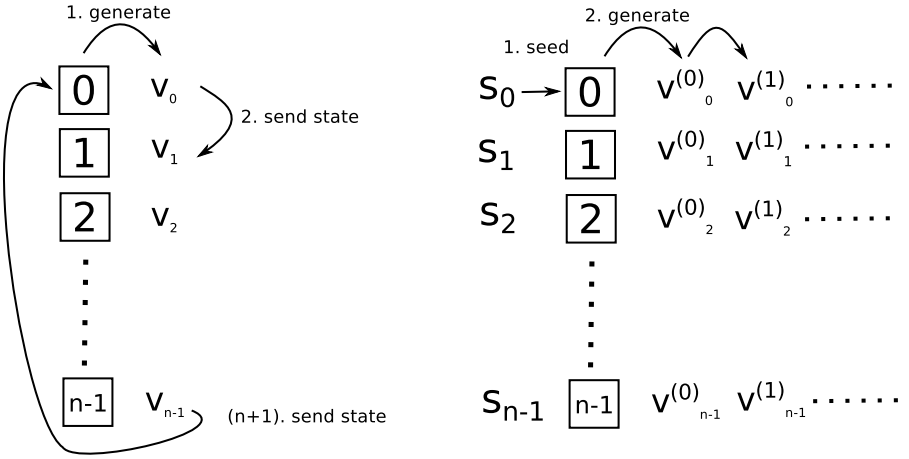

Figure E.2: Depiction of two modes of operation of pseudo-random number generators (PRNG) in parallel applications with n processes. v_i are vectors of random numbers of some length l. In the left panel, the PRNG of process 0 is seeded upon program initialisation. It then generates a vector v_0 of random numbers and afterwards sends its state to process 1, which generates its vector v_1 and so on until all processes have generated random numbers and process $n-1$ sends its state to process 0. In the right panel, the PRNG of all processes are initialised with different seeds s_i and generate independent chains of pseudo-random numbers without any further inter-process communication.

would have one PRNG for each lattice site. Each one is seeded independently and random numbers are generated without inter-process communication. This last mode is not implemented in tmLQCD.

While the sequential mode has the benefit of reproducing the random number chains used in sequential execution in a parallel program, the necessity of inter-process communication can make it substantially slower than the parallel mode. This mode can be implemented in an alternative way by having each process generate random numbers for the entire lattice but using only those numbers that are necessary for the process-local domain while discarding all others. Although it may appear that this could be substantially more expensive, when many processes are used, the serialisation and communication overheads of the implementation depicted in Figure E.2 are much greater. In the parallel mode, it must be ensured that any possible correlations in the seeds do not produce correlations between the different random number chains. In addition, the latter mode makes debugging more complicated because the random number chains depend on the number of processes. From these compromises one can conclude that the sequential mode is more suitable and necessary for testing and debugging, while the parallel mode is advisable for production runs, as long as the absence of correlations has been demonstrated.

The random number generator used in the tmLQCD code-base is RANLUX [217], operated in the parallel mode by default with some exceptions. The sequential mode is available for testing purposes and has been implemented in the more efficient way. The high statistics runs presented in this section uncovered two bugs in the PRNG framework implemented in the tmLQCD code-base one of which affected two $N_f = 2+1+1$ ETMC production ensembles.

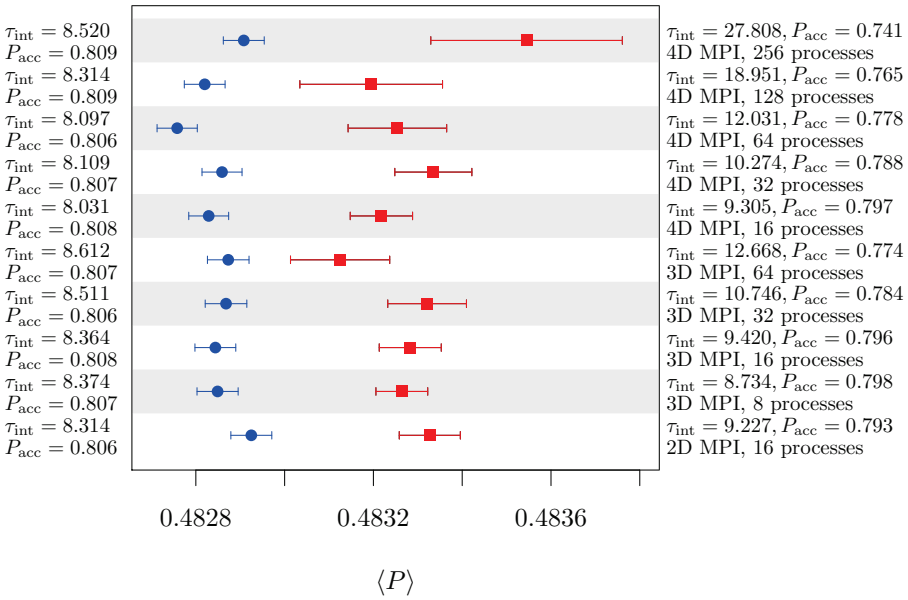

Figure E.3: Plaquette expectation values from $N_f = 2 + 1 + 1$ high statistics runs with different parallelisations showing the effect of incorrectly switching between sequential and parallel modes in the PRNG. The red ■ correspond to the situation with the bug present while the blue ● show the situation after the bug has been fixed. The autocorrelation times and acceptance probabilities correspond to the respective side and the parallelisation is listed on the right.

INCORRECTLY SWITCHING BETWEEN PRNG MODES

The default behaviour of the tmLQCD software suite is to operate the PRNG such that each process has its own independently seeded generator, relying on the good quality criteria of RANLUX to ensure that no correlations are produced between the random number chains of different processes. The various routines which provide random numbers for gauge, pseudo-fermion and conjugate momentum fields for different purposes. They are called with an argument which selects between the sequential and parallel modes and in theory this can be switched back and forth as program execution progresses. In addition, a global flag is provided in the input file which is supposed to make the mode consistent across all functions. However, when the routines for the Polynomial Hybrid Monte Carlo algorithm [218] were added to version 5 of the code-base, a particular routine requested random numbers in the sequential mode, whatever the aforementioned global flag was set to. In principle this should not be a problem, but unfortunately the particular function that it called assumed that the entire program was running in the sequential mode if it was called with the respective flag.

As a result, the PRNGs would behave as follows for N processes:

1. The program would be started in the parallel mode and a number of routines would request random numbers normally.
2. The problematic routine would be called, and processes 0 to $N-1$ would each generate m random numbers in sequence, correctly.
3. At the end of the routine, process 0 and $N-1$ would be left in the same state while all other processes would be in intermediate states.
4. The remaining routines of the program would now operate in the parallel mode.
    - Process 0 and $N-1$ would generate the same random numbers until the next call of the problematic routine.
    - For any other process $n \mid 1 < n < N-1$, the generation of m random numbers would result in the PRNG state of process $n+1$,

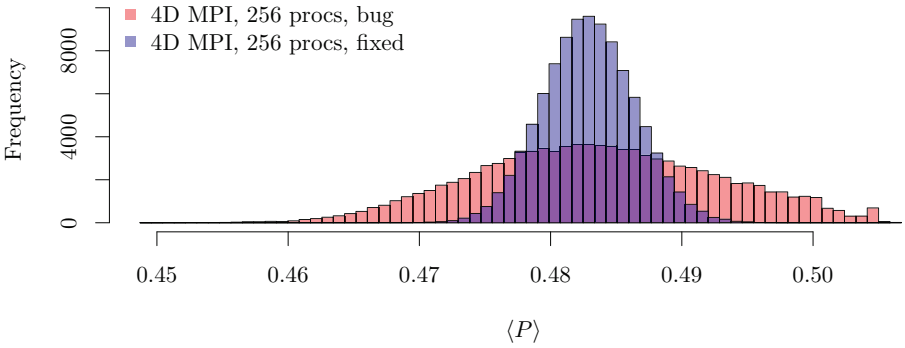

Figure E.4: Histograms showing the distribution of the plaquette observable in the molecular dynamics history of identically parallelised $N_f = 2+1+1$ high statistics runs with and without the bug described in the text.

the same random numbers would thus be recycled for different purposes and different parts of the lattice.

The effect of this bug can best be understood from Figure E.3 which shows the plaquette expectation values, their autocorrelation times and the acceptance rates for $N_f = 2+1+1$ high statistics runs with the non-degenerate doublet simulated using the PHMC. All the simulations with the bug present show a clear bias which increases the plaquette expectation value. In addition, a marked increase in the variance can be seen in Figure E.4. At first sight this seems to be counter-intuitive because the amount of randomness is reduced by the recycling of random numbers, but it is likely caused by the fact that cancellations between negative and positive deviations from the mean value are suppressed. The width of the distribution increases autocorrelation times as expected, but there is only a mild effect on the acceptance rate. This explains why the bug was not noticed in the production of the $N_f = 2+1+1$ ensembles D15.48 and D45.32, which had to be regenerated as a consequence and once again reinforces that high statistics runs with different parametrisations should be an essential part of the software development process of LQCD applications.

After the bug had been fixed the set of high statistics runs was repeated and a situation was seen where all parallelisations were compatible within errors with regards to their central value, integrated autocorrelation time and acceptance rate. Upon closer inspection of Figure E.3, however, there seems to be a pattern in the data points with four-dimensional parallelisation, indicating some bias which pushes the expectation value up mildly and which affects the autocorrelation time slightly as the number of processes is varied. In the next section it will be shown that this almost undetectably mild bias is in fact yet another bug which was introduced in the process of fixing the present one. It has only a minimal effect on the plaquette expectation value in the $N_f = 2+1+1$ simulations done here, but will be shown to emerge clearly with specially chosen simulation parameters.

INCORRECT PRNG INITIALISATION

As a result of the issue described in the previous section, changes were made to the entire PRNG framework in the tmLQCD software suite. In addition to a safer and faster implementation of the sequential mode, the seeding of the PRNGs in the parallel mode was modified to ensure that in two identical runs with the same number of processes, the distribution of PRNG seeds in the parallel mode was guaranteed to be the same. Unfortunately, a bug was introduced which would have been quite undetectable in the early stages of a production run, unless the parallelisation was changed or two

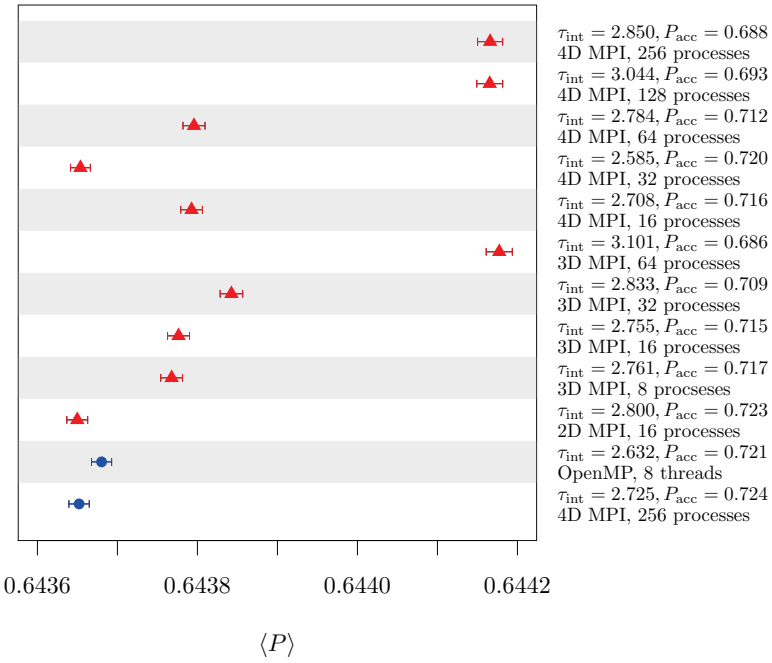

Figure E.5: Plaquette expectation values from $N_f = 8$ high statistics runs with different parallelisations hinting at a bug in the computation of the process-local seed for the PRNG. Red ▲ correspond to the situation with the bug present while blue ● show the situation after the bug has been fixed.

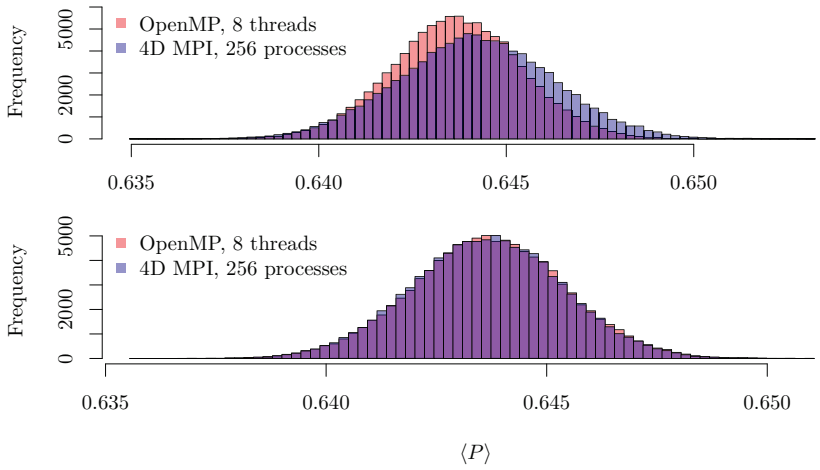

Figure E.6: Histograms of the plaquette expectation value from two $N_f = 8$ high statistics runs with different parallelisations. **(top)**: Some processes have been started with the same seed. **(bottom)**: The bug has been fixed.

different volumes would be run concurrently. This bug neither affected the autocorrelation time nor the acceptance rate, leading instead to statistically significant biases given sufficient statistical accuracy. Luckily, at this point in time the framework for high-statistics tests was well established and more tests were carried out which uncovered this issue.

As described in the previous section, it was seen that the bug had a minor but visible effect on the plaquette expectation value but it was initially unclear what the problem was, where it originated from and whether there was some discernible pattern in the deviations. Based on the assumption that it was related to random numbers, a special set of $N_f = 8$ high statistics runs was carried out with the hope that the increased number of fermions – and hence random numbers for pseudo-fermion fields – would make some pattern emerge which would help in identifying the issue. As can be seen in Figure E.5 after 10^5 trajectories, this assumption was well justified. The significant bias was already visible after a fraction of the total simulation time, but the pattern which eventually hinted at the origin of the problem in

the calculation of the seed for each process, only became apparent after the statistical error was sufficiently suppressed.

This set of data has the quite remarkable feature that even though several processes were given the same seed, the autocorrelation time and error of the observable was only very mildly affected. In a production situation, there would therefore have been no hint whatsoever of the bug, unless the parallelisation was changed midway or some other observable showed unexpected behaviour. This becomes quite clear looking at the histograms in the top panel of Figure E.6 which shows the distribution of the plaquette observable in the molecular dynamics histories with the most extreme deviation between two parallelisations (OpenMP vs. MPI with 256 processes). As can be imagined, in this case *individual* observations of the plaquette or in fact the energy violation show no hint of there being something significantly wrong. The situation after the bug had been fixed is shown in the lower panel of the same figure from which it can be concluded that the calculation of the process-local seed is now correctly implemented.

F

ADDITIONAL FAST FOURIER TRANSFORMS

Figures F.1 and F.2 are complementary material to Chapter 11 and show trajectories of gauge, quark determinant and determinant ratio forces and averaged FFTs thereof. The number of integration steps is halved between Figure F.1 and Figure F.2.

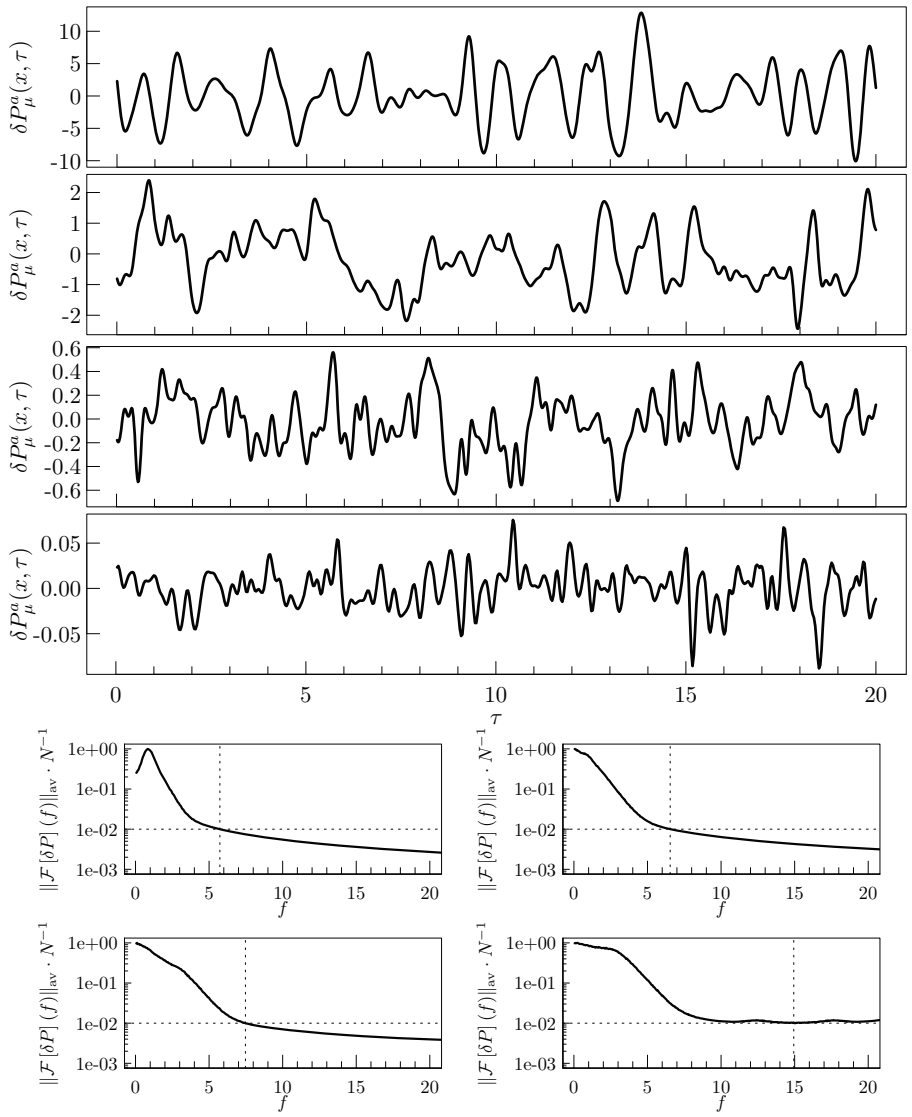

Figure F.1: Representative MD trajectories and the normalised average frequency amplitudes of the forces in a simulation with mass preconditioning as in Chapter 11 with $N_4 = 300$ integration steps on the outermost time scale. Top to bottom, top-left to bottom right: gauge, determinant and two determinant ratios.

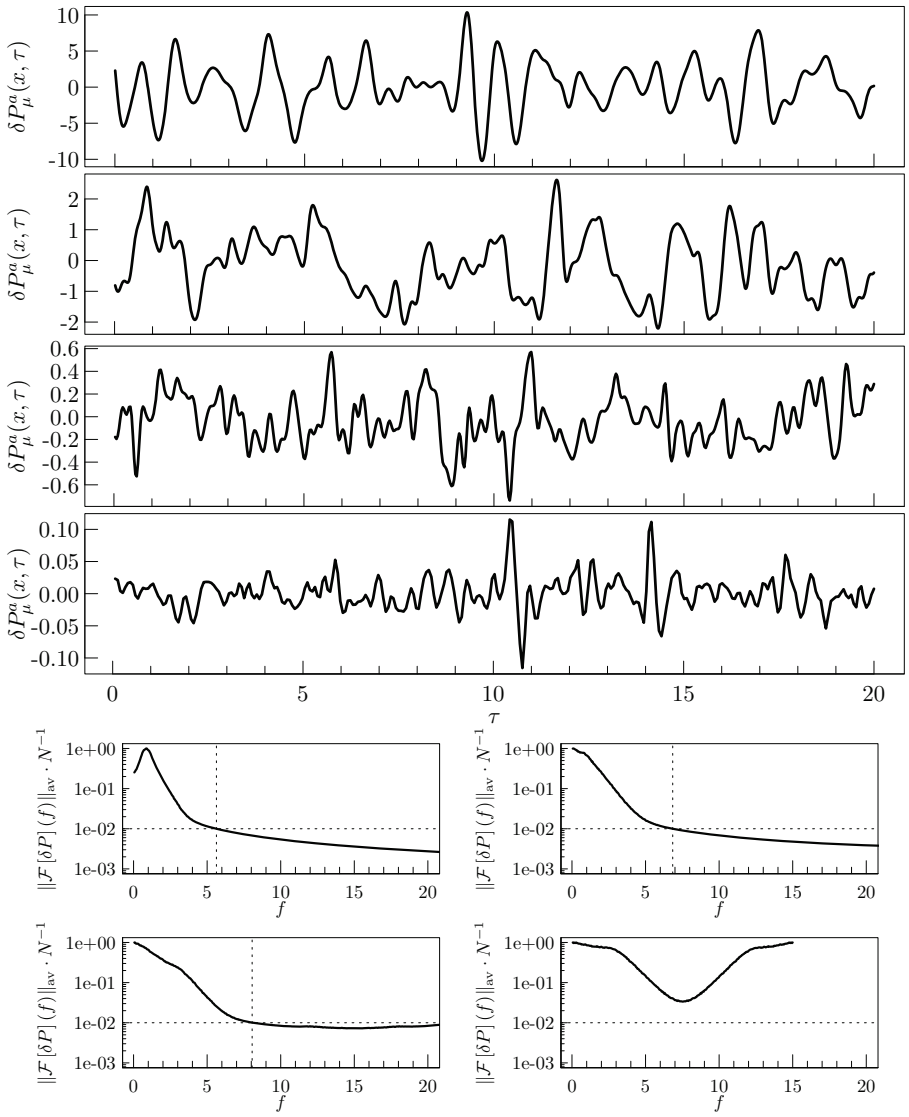

Figure F.2: Representative MD trajectories and the normalised average frequency amplitudes of the forces in a simulation with mass preconditioning as in Chapter 11 with $N_4 = 150$ integration steps on the outermost time scale. Top to bottom, top-left to bottom right: gauge, determinant and two determinant ratios.

PUBLICATIONS AND CONFERENCE CONTRIBUTIONS

[P1] A. Abdel-Rehim, P. Boucaud, N. Carrasco, A. Deuzeman, P. Dimopoulos, R. Frezzotti, G. Herdoiza, K. Jansen, B. Kostrzewa, M. Mangin-Brinet, I. Montvay, D. Palao, G. Rossi, F. Sanfilippo, L. Scorzato, A. Shindler, C. Urbach, and U. Wenger. A first look at maximally twisted mass lattice QCD calculations at the physical point. *PoS*, LATTICE2013:264, 2013. arXiv:1311.4522 [hep-lat].

[P2] A. Deuzeman, K. Jansen, B. Kostrzewa, and C. Urbach. Experiences with OpenMP in tmLQCD. In *31st International Symposium on Lattice Field Theory*, number 416 in PoS(LATTICE 2013). arXiv:1311.4521 [hep-lat].

[P3] A. Abdel-Rehim, F. Burger, B. Kostrzewa, A. Deuzeman, L. Scorzato, and C. Urbach. Recent developments in the tmLQCD software suite. In *31st International Symposium on Lattice Field Theory*, number 414 in PoS(LATTICE 2013). arXiv:1311.5495 [hep-lat].

[P4] C. Alexandrou, V. Drach, K. Hadjiyiannakou, K. Jansen, B. Kostrzewa, and C. Wiese. Looking at the gluon moment of the nucleon with dynamical twisted mass fermions. *PoS*, LATTICE2013:289, 2014. arXiv:1311.3174 [hep-lat].

[P5] A.Abdel-Rehim, C. Alexandrou, P. Dimopoulos, R. Frezzotti, K. Jansen, C. Kallidonis, B. Kostrzewa, M. Mangin-Brinet, G. Rossi, C. Urbach, and U. Wenger. Progress in Simulations with Twisted Mass Fermions at the Physical Point. *PoS*, LATTICE2014:119, 2015. arXiv:1411.6842 [hep-lat].

[P6] A. Abdel-Rehim, C. Alexandrou, F. Burger, M. Constantinou, P. Dimopoulos, R. Frezzotti, K. Hadjiyiannakou, K. Jansen, C. Kallidonis, B. Kostrzewa, G. Koutsou, M. Mangin-Brinet, M. Petschlies, G. Pientka, G. Rossi, C. Urbach, and U. Wenger. Simulating QCD at the Physical Point with $N_f = 2$ Wilson Twisted Mass Fermions at Maximal Twist. 2015. arXiv:1507.05068 [hep-lat].

[P7] A. Abdel-Rehim, Alexandrou, M. C. Constantinou, P. Dimopoulos, R. Frezzotti, K. Hadjiyiannakou, K. Jansen, C. Kallidonis, B. Kostrzewa, G. Koutsou, M. Mangin-Brinet, M. Oehm, G. Rossi, C. Urbach, and U. Wenger. Nucleon and pion structure with lattice QCD simulations at physical value of the pion mass. *Phys. Rev.*, D92(11):114,513, 2015. arXiv:1507.04936 [hep-lat].

SOFTWARE DEVELOPMENT

[C1] `tmLQCD` – Versatile simulation and inversion code for various types of Wilson quarks. `https://github.com/etmc/tmLQCD`

[C2] `hadron` – Comprehensive library of R routines for the analysis of lattice QCD data such as correlation functions. `https://github.com/etmc/hadron`

[C3] `Rscripts` – Large collection of R routines for the analysis of various lattice QCD data. Large array of specialised and generalised plotting routines for various types of data. Integrates with and relies heavily on routines from [C2]. Suitable for automated analysis of data produced with [C5, C6, C4]. `https://github.com/kostrzewa/misc_R_scripts`

[C4] `numderiv` – Extension of [C1] with support for the computation of numerical molecular dynamics derivatives for most of the fermionic monomials supported by tmLQCD. Based on advanced memory management framework written by A. Deuzeman. `https://github.com/kostrzewa/tmLQCD` branch `ndcloverrat_numderiv`

[C5] `libvcvpp` – C++ interface and wrapper for a generalized contraction code with a flexible input file format which allows the definition of correlation functions in terms of quark lines. The current implementation has support for meson correlation functions with smearing and fuzzing. Point, timeslice and volume sources are supported and various levels of dilution can be used. Propagators should be provided in LIME format. Particularly well suited for automating workflows with many valence quark masses. Also includes support for propagators computed using the tmLQCD multiple mass solver. Designed to be straight-forward to extend to more complicated observables. Based on MPI-parallel contraction routines from CVC by M. Petschlies, it provides a class library defining the basic objects required for the contraction of quark propagators. Output format is fully backwards-compatible to ETMC contraction code written by C. Urbach and M. Wagner. `https://github.com/kostrzewa/cvc`, branch `libcvcpp`

[C6] `bgqmeson_2pt` – Jobscript and input file generator for the automatic computation of meson two-point functions on BG/Q, designed to seamlessly integrate [C1] and [C5]. Trivially extensible to other machines with the LoadLeveler (LL) workload scheduler. Makes extensive use of advanced LL scheduling features. Modular design should allow extension to other architectures. `https://github.com/kostrzewa/jobscripts` subproject `generators/bgqmeson_2pt`

[C7] `highstat-gen` – Jobscript and input file generator for the automated running of high statistics runs using [C1]. In its current form limited for usage on a particular cluster. Somewhat modular design should allow generalisation. Integrates with [C8]. `https://github.com/kostrzewa/jobscripts` subproject `generators/bgqmeson_2pt`

[C8] `highstat-analysis` – Collection of R routines for the analysis of data generated in high statistics runs using [C1]. Integrates with [C7].

BIBLIOGRAPHY

[1] V. Fock. On the invariant form of the wave equation and the equations of motion for a charged point mass. (In German and English). *Z. Phys.*, 39:226–232, 1926. [Surveys High Energ. Phys.5,245(1986)].

[2] H. Weyl. Electron and Gravitation. 1. (In German). *Z. Phys.*, 56:330–352, 1929. [Surveys High Energ. Phys.5,261(1986)].

[3] C.-N. Yang and R. L. Mills. Conservation of Isotopic Spin and Isotopic Gauge Invariance. *Phys. Rev.*, 96:191–195, 1954.

[4] S. Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19:1264–1266, 1967.

[5] S. L. Glashow. Partial Symmetries of Weak Interactions. *Nucl. Phys.*, 22:579–588, 1961.

[6] A. Salam and J. C. Ward. Electromagnetic and weak interactions. *Phys. Lett.*, 13:168–171, 1964.

[7] P. W. Higgs. Broken symmetries, massless particles and gauge fields. *Phys. Lett.*, 12:132–133, 1964.

[8] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Phys. Rev. Lett.*, 13:321–323, 1964.

[9] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. Global Conservation Laws and Massless Particles. *Phys. Rev. Lett.*, 13:585–587, 1964.

[10] G. Aad. Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments. *Phys. Rev. Lett.*, 114:191,803, 2015. `arXiv:1503.07589 [hep-ex]`.

[11] H. Fritzsch, M. Gell-Mann, and H. Leutwyler. Advantages of the Color Octet Gluon Picture. *Phys. Lett.*, B47:365–368, 1973.

[12] D. J. Gross and F. Wilczek. Asymptotically Free Gauge Theories. 1. *Phys. Rev.*, D8:3633–3652, 1973.

[13] S. Weinberg. Nonabelian Gauge Theories of the Strong Interactions. *Phys. Rev. Lett.*, 31:494–497, 1973.

[14] D. J. Gross and F. Wilczek. Ultraviolet Behavior of Nonabelian Gauge Theories. *Phys. Rev. Lett.*, 30:1343–1346, 1973.

[15] Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\tilde{\Lambda}_b^0 \to J/\psi K^- p$ Decays. *Phys. Rev. Lett.*, 115:072,001, Aug 2015. URL `http://link.aps.org/doi/10.1103/PhysRevLett.115.072001`.

[16] K. G. Wilson and W. Zimmermann. Operator product expansions and composite field operators in the general framework of quantum field theory. *Commun. Math. Phys.*, 24:87–106, 1972.

[17] M. Luscher. Signatures of unstable particles in finite volume. *Nucl. Phys.*, B364:237–254, 1991.

[18] M. Luscher. Volume Dependence of the Energy Spectrum in Massive Quantum Field Theories. 2. Scattering States. *Commun. Math. Phys.*, 105:153–188, 1986.

[19] M. Lüscher and U. Wolff. HOW TO CALCULATE THE ELASTIC SCATTERING MATRIX IN TWO-DIMENSIONAL QUANTUM FIELD THEORIES BY NUMERICAL SIMULATION. *Nucl.Phys.*, B339:222–252, 1990.

[20] X. Ji. Parton Physics on a Euclidean Lattice. *Phys. Rev. Lett.*, 110:262,002, 2013. arXiv:1305.1539 [hep-ph].

[21] J. Rojo et al. The PDF4LHC report on PDFs and LHC data: Results from Run I and preparation for Run II. *J. Phys.*, G42:103,103, 2015. arXiv:1507.00556 [hep-ph].

[22] T. Blum, M. Hayakawa, and T. Izubuchi. Hadronic corrections to the muon anomalous magnetic moment from lattice QCD. *PoS*, LATTICE2012:022, 2012. arXiv:1301.2607 [hep-lat].

[23] A. Vladikas. FLAG: Lattice QCD Tests of the Standard Model and Foretaste for Beyond. In *13th Conference on Flavor Physics and CP Violation (FPCP 2015) Nagoya, Japan, May 25-29, 2015*. 2015. URL http://inspirehep.net/record/1391651/files/arXiv:1509.01155.pdf. arXiv:1509.01155 [hep-lat].

[24] E. Shintani, T. Blum, and T. Izubuchi. Electric Dipole Moment of the Neutron. *PoS*, ConfinementX:348, 2012.

[25] K. Symanzik. Continuum Limit and Improved Action in Lattice Theories. 1. Principles and phi**4 Theory. *Nucl. Phys.*, B226:187, 1983.

[26] K. Symanzik. Continuum Limit and Improved Action in Lattice Theories. 2. O(N) Nonlinear Sigma Model in Perturbation Theory. *Nucl. Phys.*, B226:205, 1983.

[27] M. Constantinou. Hadron Structure. *PoS*, LATTICE2014:001, 2015. arXiv:1411.0078 [hep-lat].

[28] R. P. Feynman. Space-time approach to nonrelativistic quantum mechanics. *Rev. Mod. Phys.*, 20:367–387, 1948.

[29] R. P. Feynman. Space - time approach to quantum electrodynamics. *Phys. Rev.*, 76:769–789, 1949.

[30] K. Symanzik. Small distance behavior in field theory and power counting. *Commun. Math. Phys.*, 18:227–246, 1970.

[31] J. C. G. Callan. Broken scale invariance in scalar field theory. *Phys. Rev.*, D2:1541–1547, 1970.

[32] S. Weinberg. Dynamical approach to current algebra. *Phys. Rev. Lett.*, 18:188–191, 1967.

[33] M. Bochicchio, L. Maiani, G. Martinelli, G. C. Rossi, and M. Testa. Chiral Symmetry on the Lattice with Wilson Fermions. *Nucl. Phys.*, B262:331, 1985.

[34] G. Martinelli, C. Pittori, C. T. Sachrajda, M. Testa, and A. Vladikas. A General method for nonperturbative renormalization of lattice operators. *Nucl. Phys.*, B445:81–108, 1995. arXiv:hep-lat/9411010 [hep-lat].

[35] K. Jansen, C. Liu, M. Luscher, H. Simma, S. Sint, R. Sommer, P. Weisz, and U. Wolff. Nonperturbative renormalization of lattice QCD at all scales. *Phys. Lett.*, B372:275–282, 1996. arXiv:hep-lat/9512009 [hep-lat].

[36] M. Luscher, R. Sommer, P. Weisz, and U. Wolff. A Precise determination of the running coupling in the SU(3) Yang-Mills theory. *Nucl. Phys.*, B413:481–502, 1994. arXiv:hep-lat/9309005 [hep-lat].

[37] K. G. Wilson. Confinement of Quarks. *Phys. Rev.*, D10:2445–2459, 1974.

[38] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar. Sign problem in the numerical simulation of many-electron systems. *Phys. Rev. B*, 41:9301–9307, May 1990. URL http://link.aps.org/doi/10.1103/PhysRevB.41.9301.

[39] K. Osterwalder and R. Schrader. AXIOMS FOR EUCLIDEAN GREEN'S FUNCTIONS. *Commun. Math. Phys.*, 31:83–112, 1973.

[40] K. Osterwalder and R. Schrader. Axioms for Euclidean Green's Functions. 2. *Commun. Math. Phys.*, 42:281, 1975.

[41] I. Montvay and G. Münster. *Quantum Fields on a Lattice*. Cambridge University Press, 1994. ISBN 9780511470783. URL http://dx.doi.org/10.1017/CBO9780511470783.

[42] M. Luscher. Construction of a Selfadjoint, Strictly Positive Transfer Matrix for Euclidean Lattice Gauge Theories. *Commun. Math. Phys.*, 54:283, 1977.

[43] K. Osterwalder and E. Seiler. Gauge Field Theories on the Lattice. *Annals Phys.*, 110:440, 1978.

[44] Y. Iwasaki. Renormalization Group Analysis of Lattice Theories and Improved Lattice Action: Two-Dimensional Nonlinear O(N) Sigma Model. *Nucl. Phys.*, B258:141–156, 1985.

[45] A. Haar. Der Massbegriff in der Theorie der Kontinuierlichen Gruppen. *Annals of Mathematics*, 34(1):pp. 147–169, 1933. ISSN 0003486X. URL http://www.jstor.org/stable/1968346.

[46] S. Schaefer, R. Sommer, and F. Virotta. Critical slowing down and error analysis in lattice QCD simulations. *Nucl. Phys.*, B845:93–119, 2011. arXiv:1009.5228 [hep-lat].

[47] M. Luscher and S. Schaefer. Lattice QCD without topology barriers. *JHEP*, 07:036, 2011. arXiv:1105.4749 [hep-lat].

[48] S. Aoki and A. Gocksch. Spontaneous breaking of flavor symmetry and parity in lattice QCD with Wilson fermions. *Phys. Rev.*, D45:3845–3853, 1992.

[49] R. Frezzotti, P. A. Grassi, S. Sint, and P. Weisz. A Local formulation of lattice QCD without unphysical fermion zero modes. *Nucl. Phys. Proc. Suppl.*, 83:941–946, 2000. arXiv:hep-lat/9909003 [hep-lat].

[50] R. Frezzotti, P. A. Grassi, S. Sint, and P. Weisz. Lattice QCD with a chirally twisted mass term. *JHEP*, 08:058, 2001. hep-lat/0101001.

[51] R. Frezzotti, S. Sint, and P. Weisz. O(a) improved twisted mass lattice QCD. *JHEP*, 07:048, 2001. arXiv:hep-lat/0104014 [hep-lat].

[52] M. Della Morte, R. Frezzotti, J. Heitger, and S. Sint. Cutoff effects in twisted mass lattice QCD. *JHEP*, 10:041, 2001. arXiv:hep-lat/0108019 [hep-lat].

[53] R. Frezzotti and G. C. Rossi. Chirally improving Wilson fermions. I: O(a) improvement. *JHEP*, 08:007, 2004. hep-lat/0306014.

[54] R. Frezzotti and G. C. Rossi. Chirally improving Wilson fermions. II: Four-quark operators. *JHEP*, 10:070, 2004. arXiv:hep-lat/0407002.

[55] S. Sint. Lattice QCD with a chiral twist. In *Workshop on Perspectives in Lattice QCD Nara, Japan, October 31-November 11, 2005*. 2007. arXiv:hep-lat/0702008 [HEP-LAT].

[56] A. Shindler. Twisted mass lattice QCD. *Phys.Rept.*, 461:37–110, 2008. arXiv:0707.4093 [hep-lat].

[57] M. Luscher and P. Weisz. On-Shell Improved Lattice Gauge Theories. *Commun. Math. Phys.*, 97:59, 1985. [Erratum: Commun. Math. Phys.98,433(1985)].

[58] B. Sheikholeslami and R. Wohlert. Improved Continuum Limit Lattice Action for QCD with Wilson Fermions. *Nucl. Phys.*, B259:572, 1985.

[59] S. Aoki and O. Bar. Automatic O(a) improvement for twisted-mass QCD in the presence of spontaneous symmetry breaking. *Phys. Rev.*, D74:034,511, 2006. arXiv:hep-lat/0604018 [hep-lat].

[60] K. Cichy, E. Garcia-Ramos, and K. Jansen. Short distance singularities and automatic O(a) improvement: the cases of the chiral condensate and the topological susceptibility. *JHEP*, 04:048, 2015. arXiv:1412.0456 [hep-lat].

[61] F. Burger, G. Hotzel, K. Jansen, and M. Petschlies. The hadronic vacuum polarization and automatic $\mathcal{O}(a)$ improvement for twisted mass fermions. *JHEP*, 03:073, 2015. arXiv:1412.0546 [hep-lat].

[62] S. R. Sharpe. Observations on discretization errors in twisted-mass lattice QCD. *Phys. Rev.*, D72:074,510, 2005. arXiv:hep-lat/0509009 [hep-lat].

[63] R. Frezzotti, G. Martinelli, M. Papinutto, and G. C. Rossi. Reducing cutoff effects in maximally twisted lattice QCD close to the chiral limit. *JHEP*, 04:038, 2006. arXiv:hep-lat/0503034 [hep-lat].

[64] K. Jansen and C. Liu. Implementation of Symanzik's improvement program for simulations of dynamical Wilson fermions in lattice QCD. *Comput.Phys.Commun.*, 99:221–234, 1997. arXiv:hep-lat/9603008 [hep-lat].

[65] T. Chiarappa et al. Numerical simulation of QCD with u, d, s and c quarks in the twisted-mass Wilson formulation. *Eur. Phys. J.*, C50:373–383, 2007. arXiv:hep-lat/0606011.

[66] R. Frezzotti and G. C. Rossi. Twisted-mass lattice QCD with mass non-degenerate quarks. *Nucl. Phys. Proc. Suppl.*, 128:193–202, 2004. hep-lat/0311008.

[67] R. Baron et al. Computing K and D meson masses with N_f = 2+1+1 twisted mass lattice QCD. *Comput.Phys.Commun.*, 182:299–316, 2011. arXiv:1005.2042 [hep-lat].

[68] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953. URL http://scitation.aip.org/content/aip/journal/jcp/21/6/10.1063/1.1699114.

[69] W. K. HASTINGS. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970. URL http://biomet.oxfordjournals.org/content/57/1/97.abstract. http://biomet.oxfordjournals.org/content/57/1/97.full.pdf+html.

[70] M. Luscher. A New approach to the problem of dynamical quarks in numerical simulations of lattice QCD. *Nucl. Phys.*, B418:637–648, 1994. arXiv:hep-lat/9311007 [hep-lat].

[71] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Phys. Lett.*, B195:216–222, 1987.

[72] A. D. Kennedy and P. Rossi. CLASSICAL MECHANICS ON GROUP MANIFOLDS AND APPLICATIONS TO HYBRID MONTE CARLO. *Nucl. Phys.*, B327:782, 1989.

[73] K. Bitar, A. D. Kennedy, R. Horsley, S. Meyer, and P. Rossi. Hybrid Monte Carlo and Quantum Chromodynamics. *Nucl. Phys.*, B313:377–392, 1989.

[74] I. Omelyan, I. Mryglod, and R. Folk. Symplectic analytically integrable decomposition algorithms: classification, derivation, and application to molecular dynamics, quantum and celestial mechanics simulations. *Computer Physics Communications*, 151(3):272–314, 2003.

[75] J. C. Sexton and D. H. Weingarten. Hamiltonian evolution for the hybrid Monte Carlo algorithm. *Nucl. Phys.*, B380:665–678, 1992.

[76] M. Hasenbusch. Speeding up the hybrid Monte Carlo algorithm for dynamical fermions. *Phys. Lett.*, B519:177–182, 2001. arXiv:hep-lat/0107019 [hep-lat].

[77] M. Hasenbusch and K. Jansen. Speeding up lattice QCD simulations with clover improved Wilson fermions. *Nucl. Phys.*, B659:299–320, 2003. arXiv:hep-lat/0211042 [hep-lat].

[78] A. Ali Khan, T. Bakeyev, M. Gockeler, R. Horsley, D. Pleiter, P. E. L. Rakow, A. Schafer, G. Schierholz, and H. Stuben. Accelerating the hybrid Monte Carlo algorithm. *Phys. Lett.*, B564:235–240, 2003. arXiv:hep-lat/0303026 [hep-lat].

[79] C. Urbach, K. Jansen, A. Shindler, and U. Wenger. HMC algorithm with multiple time scale integration and mass preconditioning. *Comput.Phys.Commun.*, 174:87–98, 2006. arXiv:hep-lat/0506011 [hep-lat].

[80] M. Luscher and S. Schaefer. Lattice QCD with open boundary conditions and twisted-mass reweighting. *Comput. Phys. Commun.*, 184:519–528, 2013. arXiv:1206.2809 [hep-lat].

[81] S. Schaefer. Status and challenges of simulations with dynamical fermions. *PoS*, LATTICE2012:001, 2012. arXiv:1211.5069 [hep-lat].

[82] E. Follana, Q. Mason, C. Davies, K. Hornbostel, G. P. Lepage, J. Shigemitsu, H. Trottier, and K. Wong. Highly improved staggered quarks on the lattice, with applications to charm physics. *Phys. Rev.*, D75:054,502, 2007. arXiv:hep-lat/0610092 [hep-lat].

[83] A. Bazavov et al. Lattice QCD ensembles with four flavors of highly improved staggered quarks. *Phys. Rev.*, D87(5):054,505, 2013. arXiv:1212.4768 [hep-lat].

[84] M. Bruno et al. Simulation of QCD with N_f $= 2 + 1$ flavors of non-perturbatively improved Wilson fermions. *JHEP*, 02:043, 2015. arXiv:1411.3982 [hep-lat].

[85] W. Bietenholz et al. Tuning the strange quark mass in lattice simulations. *Phys. Lett.*, B690:436–441, 2010. arXiv:1003.1114 [hep-lat].

[86] S. Durr, Z. Fodor, C. Hoelbling, S. D. Katz, S. Krieg, T. Kurth, L. Lellouch, T. Lippert, K. K. Szabo, and G. Vulvert. Lattice QCD at the physical point: Simulation and analysis details. *JHEP*, 08:148, 2011. arXiv:1011.2711 [hep-lat].

[87] S. Durr, Z. Fodor, C. Hoelbling, S. D. Katz, S. Krieg, T. Kurth, L. Lellouch, T. Lippert, K. K. Szabo, and G. Vulvert. Lattice QCD at the physical point: light quark masses. *Phys. Lett.*, B701:265–268, 2011. arXiv:1011.2403 [hep-lat].

[88] S. Borsanyi et al. Ab initio calculation of the neutron-proton mass difference. *Science*, 347:1452–1455, 2015. arXiv:1406.4088 [hep-lat].

[89] P. H. Ginsparg and K. G. Wilson. A Remnant of Chiral Symmetry on the Lattice. *Phys. Rev.*, D25:2649, 1982.

[90] H. Neuberger. Exactly massless quarks on the lattice. *Phys. Lett.*, B417:141–144, 1998. arXiv:hep-lat/9707022 [hep-lat].

[91] W. Bietenholz et al. Going chiral: Overlap versus twisted mass fermions. *JHEP*, 12:044, 2004. arXiv:hep-lat/0411001 [hep-lat].

[92] H. Matsufuru et al. Simulation with 2+1 flavors of dynamical overlap fermions. *PoS*, LATTICE2008:077, 2008.

[93] D. B. Kaplan. A Method for simulating chiral fermions on the lattice. *Phys. Lett.*, B288:342–347, 1992. arXiv:hep-lat/9206013 [hep-lat].

[94] T. Blum et al. Domain wall QCD with physical quark masses. 2014. arXiv:1411.7017 [hep-lat].

[95] J. Noaki, S. Aoki, G. Cossu, H. Fukaya, S. Hashimoto, and T. Kaneko. Fine lattice simulations with the Ginsparg-Wilson fermions. *PoS*, LATTICE2014:069, 2014.

[96] R. Baron et al. Light Meson Physics from Maximally Twisted Mass Lattice QCD. *JHEP*, 1008:097, 2010. arXiv:0911.5061 [hep-lat].

[97] R. Baron et al. Light hadrons from lattice QCD with light (u,d), strange and charm dynamical quarks. *JHEP*, 06:111, 2010. arXiv:1004.5284 [hep-lat].

[98] M. G. Beckett, B. Joo, C. M. Maynard, D. Pleiter, O. Tatebe, and T. Yoshie. Building the International Lattice Data Grid. *Comput. Phys. Commun.*, 182:1208–1214, 2011. arXiv:0910.1692 [hep-lat].

[99] C. Alexandrou, V. Drach, K. Jansen, C. Kallidonis, and G. Koutsou. Baryon spectrum with $N_f = 2 + 1 + 1$ twisted mass fermions. *Phys. Rev.*, D90(7):074,501, 2014. arXiv:1406.4310 [hep-lat].

[100] S. R. Sharpe and R. L. J. Singleton. Spontaneous flavor and parity breaking with Wilson fermions. *Phys.Rev.*, D58:074,501, 1998. arXiv:hep-lat/9804028 [hep-lat].

[101] G. Rupak and N. Shoresh. Chiral perturbation theory for the Wilson lattice action. *Phys. Rev.*, D66:054,503, 2002. arXiv:hep-lat/0201019 [hep-lat].

[102] S. Aoki. New Phase Structure for Lattice QCD with Wilson Fermions. *Phys. Rev.*, D30:2653, 1984.

[103] S. Aoki and A. Gocksch. Spontaneous Breaking of Parity in Quenched Lattice QCD With Wilson Fermions. *Phys. Lett.*, B231:449, 1989.

[104] S. Aoki. On the phase structure of QCD with Wilson fermions. *Prog. Theor. Phys. Suppl.*, 122:179–186, 1996. arXiv:hep-lat/9509008 [hep-lat].

[105] G. Munster and C. Schmidt. Chiral perturbation theory for lattice QCD with a twisted mass term. *Europhys.Lett.*, 66:652–656, 2004. arXiv:hep-lat/0311032 [hep-lat].

[106] G. Munster, C. Schmidt, and E. E. Scholz. Chiral perturbation theory for twisted mass QCD. *Nucl. Phys. Proc. Suppl.*, 140:320–322, 2005. [,320(2004)], arXiv:hep-lat/0409066 [hep-lat].

[107] S. R. Sharpe and J. M. S. Wu. Applying chiral perturbation to twisted mass lattice QCD. *Nucl. Phys. Proc. Suppl.*, 140:323–325, 2005. [,323(2004)], arXiv:hep-lat/0407035 [hep-lat].

[108] G. Munster. On the phase structure of twisted mass lattice QCD. *JHEP*, 09:035, 2004. arXiv:hep-lat/0407006 [hep-lat].

[109] L. Scorzato. Pion mass splitting and phase structure in twisted mass QCD. *Eur.Phys.J.*, C37:445–455, 2004. arXiv:hep-lat/0407023 [hep-lat].

[110] S. R. Sharpe and J. M. S. Wu. The Phase diagram of twisted mass lattice QCD. *Phys. Rev.*, D70:094,029, 2004. arXiv:hep-lat/0407025 [hep-lat].

[111] S. Aoki and O. Bar. Twisted-mass QCD, O(a) improvement and Wilson chiral perturbation theory. *Phys. Rev.*, D70:116,011, 2004. arXiv:hep-lat/0409006 [hep-lat].

[112] O. Bar. Chiral logs in twisted mass lattice QCD with large isospin breaking. *Phys. Rev.*, D82:094,505, 2010. arXiv:1008.0784 [hep-lat].

[113] F. Bernardoni, J. Bulava, and R. Sommer. Determination of the Wilson ChPT low energy constant c_2. *PoS*, LATTICE2011:095, 2011. arXiv:1111.4351 [hep-lat].

[114] F. Farchioni, R. Frezzotti, K. Jansen, I. Montvay, G. Rossi et al. Twisted mass quarks and the phase structure of lattice QCD. *Eur.Phys.J.*, C39:421–433, 2005. arXiv:hep-lat/0406039 [hep-lat].

[115] F. Farchioni, K. Jansen, I. Montvay, E. Scholz, L. Scorzato et al. Lattice spacing dependence of the first order phase transition for dynamical twisted mass fermions. *Phys.Lett.*, B624:324–333, 2005. arXiv:hep-lat/0506025 [hep-lat].

[116] E.-M. Ilgenfritz, W. Kerler, M. Müller-Preussker, A. Sternbeck, and H. Stüben. A Numerical reinvestigation of the Aoki phase with N(f) = 2 Wilson fermions at zero temperature. *Phys. Rev.*, D69:074,511, 2004. arXiv:hep-lat/0309057 [hep-lat].

[117] C. Urbach. *Wilson Twisted Mass Fermions: Towards Realistic Simulations of Lattice QCD*. Ph.D. thesis, Free University of Berlin, September 2005. URL http://www.itkp.uni-bonn.de/~urbach/thesis.pdf.

[118] D. Becirevic, P. Boucaud, V. Lubicz, G. Martinelli, F. Mescia et al. Exploring twisted mass lattice QCD with the Clover term. *Phys.Rev.*, D74:034,501, 2006. arXiv:hep-lat/0605006 [hep-lat].

[119] S. Aoki et al. Bulk first-order phase transition in three-flavor lattice QCD with O(a)-improved Wilson fermion action at zero temperature. *Phys. Rev.*, D72:054,510, 2005. arXiv:hep-lat/0409016 [hep-lat].

[120] F. Farchioni, K. Jansen, I. Montvay, E. Scholz, L. Scorzato et al. The Phase structure of lattice QCD with Wilson quarks and renormalization group improved gluons. *Eur.Phys.J.*, C42:73–87, 2005. arXiv:hep-lat/0410031 [hep-lat].

[121] P. Weisz. Continuum Limit Improved Lattice Action for Pure Yang-Mills Theory. 1. *Nucl. Phys.*, B212:1, 1983.

[122] P. Weisz and R. Wohlert. Continuum Limit Improved Lattice Action for Pure Yang-Mills Theory. 2. *Nucl. Phys.*, B236:397, 1984. [Erratum: Nucl. Phys.B247,544(1984)].

[123] K. Jansen, C. McNeile, I. Montvay, C. M. Richards, E. E. Scholz, C. Urbach, and U. Wenger. Stout Smearing for Twisted Mass Fermions. *PoS*, LAT2007:036, 2007. arXiv:0709.4434 [hep-lat].

[124] G. Herdoiza, K. Jansen, C. Michael, K. Ottnad, and C. Urbach. Determination of Low-Energy Constants of Wilson Chiral Perturbation Theory. *JHEP*, 1305:038, 2013. arXiv:1303.3516 [hep-lat].

[125] M. T. Hansen and S. R. Sharpe. Determining low-energy constants in partially quenched Wilson chiral perturbation theory. *Phys. Rev.*, D85:054,504, 2012. arXiv:1112.3998 [hep-lat].

[126] M. T. Hansen and S. R. Sharpe. Constraint on the Low Energy Constants of Wilson Chiral Perturbation Theory. *Phys. Rev.*, D85:014,503, 2012. arXiv:1111.2404 [hep-lat].

[127] G. Colangelo, U. Wenger, and J. M. S. Wu. Twisted Mass Finite Volume Effects. *Phys. Rev.*, D82:034,502, 2010. arXiv:1003.0847 [hep-lat].

[128] N. Carrasco et al. Up, down, strange and charm quark masses with N_f = 2+1+1 twisted mass lattice QCD. *Nucl.Phys.*, B887:19–68, 2014. arXiv:1403.4504 [hep-lat].

[129] S. R. Sharpe and J. M. Wu. Twisted mass chiral perturbation theory at next-to-leading order. *Phys.Rev.*, D71:074,501, 2005. arXiv:hep-lat/0411021 [hep-lat].

[130] Y. Namekawa et al. Light hadron spectroscopy in two-flavor QCD with small sea quark masses. *Phys. Rev.*, D70:074,503, 2004. arXiv:hep-lat/0404014 [hep-lat].

[131] S. Aoki et al. Lattice quantum chromodynamics at the physical point and beyond. *PTEP*, 2012:01A102, 2012.

[132] S. Aoki et al. Nonperturbative O(a) improvement of the Wilson quark action with the RG-improved gauge action using the Schrodinger functional method. *Phys.Rev.*, D73:034,501, 2006. arXiv:hep-lat/0508031 [hep-lat].

[133] G. Colangelo, S. Durr, and C. Haefeli. Finite volume effects for meson masses and decay constants. *Nucl. Phys.*, B721:136–174, 2005. arXiv:hep-lat/0503014 [hep-lat].

[134] P. Boucaud et al. Dynamical Twisted Mass Fermions with Light Quarks: Simulation and Analysis Details. *Comput.Phys.Commun.*, 179:695–715, 2008. arXiv:0803.0224 [hep-lat].

[135] M. Lüscher. Properties and uses of the Wilson flow in lattice QCD. *JHEP*, 08:071, 2010. [Erratum: JHEP03,092(2014)], arXiv:1006.4518 [hep-lat].

[136] N. Ukita et al. 2+1 flavor lattice QCD simulation with O(a)-improved Wilson quarks. *PoS*, LATTICE2008:097, 2008. arXiv:0810.0563 [hep-lat].

[137] M. A. Clark and A. D. Kennedy. Accelerating Staggered Fermion Dynamics with the Rational Hybrid Monte Carlo (RHMC) Algorithm. *Phys. Rev.*, D75:011,502, 2007. arXiv:hep-lat/0610047 [hep-lat].

[138] M. A. Clark and A. D. Kennedy. Accelerating dynamical fermion computations using the rational hybrid Monte Carlo (RHMC) algorithm with multiple pseudofermion fields. *Phys. Rev. Lett.*, 98:051,601, 2007. arXiv:hep-lat/0608015 [hep-lat].

[139] G. Parisi. Recent progresses in gauge theories. *AIP Conference Proceedings*, 68(1):1531–1568, 1980. ICHEP.

[140] S. Aoki, R. Frezzotti, and P. Weisz. Computation of the improvement coefficient c(SW) to one loop with improved gluon actions. *Nucl. Phys.*, B540:501–519, 1999. arXiv:hep-lat/9808007 [hep-lat].

[141] C. Davies, C. McNeile, K. Wong, E. Follana, R. Horgan et al. Precise Charm to Strange Mass Ratio and Light Quark Masses from Full Lattice QCD. *Phys.Rev.Lett.*, 104:132,003, 2010. arXiv:0910.3102 [hep-ph].

[142] S. Borsanyi, S. Durr, Z. Fodor, C. Hoelbling, S. D. Katz et al. High-precision scale setting in lattice QCD. *JHEP*, 1209:010, 2012. arXiv:1203.4469 [hep-lat].

[143] K. Jansen and C. Urbach. tmLQCD: A Program suite to simulate Wilson Twisted mass Lattice QCD. *Comput. Phys. Commun.*, 180:2717–2738, 2009. arXiv:0905.3331 [hep-lat].

[144] U. Wolff. Monte Carlo errors with less errors. *Comput. Phys. Commun.*, 156:143–153, 2004. hep-lat/0306017.

[145] F. Tekin, R. Sommer, and U. Wolff. Symanzik improvement of lattice QCD with four flavors of Wilson quarks. *Phys. Lett.*, B683:75–79, 2010. arXiv:0911.4043 [hep-lat].

[146] F. Stollenwerk, P. Fritzsch, R. Sommer, and U. Wolff. Determination of c_{sw} in $N_f = 3 + 1$ Lattice QCD with massive Wilson fermions. *PoS*, LATTICE2014:293, 2015. arXiv:1501.06753 [hep-lat].

[147] B. Joó, B. Pendleton, A. D. Kennedy, A. C. Irving, J. C. Sexton, S. M. Pickles, and S. P. Booth. Instability in the molecular dynamics step of a hybrid Monte Carlo algorithm in dynamical fermion lattice QCD simulations. *Phys. Rev. D*, 62:114,501, Oct 2000. URL http://link.aps.org/doi/10.1103/PhysRevD.62.114501.

[148] R. G. Edwards, I. Horvath, and A. D. Kennedy. Instabilities and nonreversibility of molecular dynamics trajectories. *Nucl. Phys.*, B484:375–402, 1997. arXiv:hep-lat/9606004 [hep-lat].

[149] M. A. Clark, A. D. Kennedy, and P. J. Silva. Tuning HMC using Poisson brackets. *PoS*, LATTICE2008:041, 2008. arXiv:0810.1315 [hep-lat].

[150] M. A. Clark, B. Joo, A. D. Kennedy, and P. J. Silva. Better HMC integrators for dynamical simulations. *PoS*, LATTICE2010:323, 2010. arXiv:1011.0230 [hep-lat].

[151] M. A. Clark, B. Joo, A. D. Kennedy, and P. J. Silva. Improving dynamical lattice QCD simulations through integrator tuning using Poisson brackets and a force-gradient integrator. *Phys. Rev.*, D84:071,502, 2011. arXiv:1108.1828 [hep-lat].

[152] A. D. Kennedy, P. J. Silva, and M. A. Clark. Shadow Hamiltonians, Poisson Brackets, and Gauge Theories. *Phys. Rev.*, D87(3):034,511, 2013. arXiv:1210.6600 [hep-lat].

[153] V. Gülpers, G. von Hippel, and H. Wittig. Scalar pion form factor in two-flavor lattice QCD. *Phys. Rev.*, D89(9):094,503, 2014. arXiv:1309.2104 [hep-lat].

[154] A. Bazavov et al. Gradient flow and scale setting on MILC HISQ ensembles. 2015. arXiv:1503.02769 [hep-lat].

[155] F. Farchioni et al. Numerical simulations with two flavors of twisted-mass Wilson quarks and DBW2 gauge action. *Eur. Phys. J.*, C47:453–472, 2006. arXiv:hep-lat/0512017 [hep-lat].

[156] K. A. Olive et al. Review of Particle Physics. *Chin. Phys.*, C38:090,001, 2014.

[157] N. Cabibbo. Unitary Symmetry and Leptonic Decays. *Phys. Rev. Lett.*, 10:531–533, 1963.

[158] M. Kobayashi and T. Maskawa. CP Violation in the Renormalizable Theory of Weak Interaction. *Prog. Theor. Phys.*, 49:652–657, 1973.

[159] M. Artuso, E. Barberio, and S. Stone. B Meson Decays. *PMC Phys.*, A3:3, 2009. arXiv:0902.3743 [hep-ph].

[160] P. Lacock, A. McKerrell, C. Michael, I. M. Stopher, and P. W. Stephenson. Efficient hadronic operators in lattice gauge theory. *Phys. Rev.*, D51:6403–6410, 1995. arXiv:hep-lat/9412079 [hep-lat].

[161] C. Michael and I. Teasdale. EXTRACTING GLUEBALL MASSES FROM LATTICE QCD. *Nucl.Phys.*, B215:433, 1983.

[162] B. Blossier, M. Della Morte, G. von Hippel, T. Mendes, and R. Sommer. On the generalized eigenvalue method for energies and matrix elements in lattice field theory. *JHEP*, 0904:094, 2009. arXiv:0902.1265 [hep-lat].

[163] J. Gasser and H. Leutwyler. Light Quarks at Low Temperatures. *Phys. Lett.*, B184:83, 1987.

[164] S. Weinberg. Phenomenological Lagrangians. *Physica*, A96:327, 1979.

[165] J. Gasser and H. Leutwyler. Chiral Perturbation Theory to One Loop. *Ann. Phys.*, 158:142, 1984.

[166] S. Dürr et al. Lattice QCD at the physical point meets SU(2) chiral perturbation theory. *Phys. Rev.*, D90(11):114,504, 2014. arXiv:1310.3626 [hep-lat].

[167] C. Alexandrou, R. Baron, J. Carbonell, V. Drach, P. Guichon, K. Jansen, T. Korzec, and O. Pene. Low-lying baryon spectrum with two dynamical twisted mass fermions. *Phys. Rev.*, D80:114,503, 2009. arXiv:0910.2419 [hep-lat].

[168] B. Blossier et al. Average up/down, strange and charm quark masses with Nf=2 twisted mass lattice QCD. *Phys.Rev.*, D82:114,513, 2010. arXiv:1010.3659 [hep-lat].

[169] B. Blossier et al. Pseudoscalar decay constants of kaon and D-mesons from $N_f = 2$ twisted mass Lattice QCD. *JHEP*, 0907:043, 2009. arXiv:0904.0954 [hep-lat].

[170] S. Aoki et al. Review of lattice results concerning low-energy particle physics. *Eur. Phys. J.*, C74:2890, 2014. arXiv:1310.8555 [hep-lat].

[171] M. Schmelling. Averaging correlated data. *Phys.Scripta*, 51:676–679, 1995.

[172] R. Sommer. A New way to set the energy scale in lattice gauge theories and its applications to the static force and alpha-s in SU(2) Yang-Mills theory. *Nucl.Phys.*, B411:839–854, 1994. arXiv:hep-lat/9310022 [hep-lat].

[173] N. Carrasco, P. Dimopoulos, R. Frezzotti, P. Lami, V. Lubicz et al. Leptonic decay constants f_K, f_D, and f_{D_s} with $N_f = 2+1+1$ twisted-mass lattice QCD. *Phys.Rev.*, D91(5):054,507, 2015. arXiv:1411.7908 [hep-lat].

[174] C. Alexandrou, M. Constantinou, and H. Panagopoulos. RI-MOM Renormalisation constants. *in preparation*, 2015.

[175] M. Göckeler, R. Horsley, H. Oelrich, H. Perlt, D. Petters et al. Nonperturbative renormalization of composite operators in lattice QCD. *Nucl.Phys.*, B544:699–733, 1999. arXiv:hep-lat/9807044 [hep-lat].

[176] M. Constantinou, V. Lubicz, H. Panagopoulos, and F. Stylianou. O(a**2) corrections to the one-loop propagator and bilinears of clover fermions with Symanzik improved gluons. *JHEP*, 0910:064, 2009. arXiv:0907.0381 [hep-lat].

[177] C. Alexandrou, M. Constantinou, T. Korzec, H. Panagopoulos, and F. Stylianou. Renormalization constants for 2-twist operators in twisted mass QCD. *Phys.Rev.*, D83:014,503, 2011. arXiv:1006.1920 [hep-lat].

[178] B. Blossier et al. Light quark masses and pseudoscalar decay constants from Nf=2 Lattice QCD with twisted mass fermions. *JHEP*, 04:020, 2008. arXiv:0709.4574 [hep-lat].

[179] A. Bazavov, C. Bernard, C. DeTar, X. Du, W. Freeman et al. Staggered chiral perturbation theory in the two-flavor case and SU(2) analysis of the MILC data. *PoS*, LATTICE2010:083, 2010. arXiv:1011.1792 [hep-lat].

[180] R. Arthur et al. Domain Wall QCD with Near-Physical Pions. *Phys.Rev.*, D87:094,514, 2013. arXiv:1208.4412 [hep-lat].

[181] A. Bazavov et al. Charmed and light pseudoscalar meson decay constants from four-flavor lattice QCD with physical light quarks. *Phys.Rev.*, D90(7):074,509, 2014. arXiv:1407.3772 [hep-lat].

[182] D. N. Politis and J. P. Romano. The Stationary Bootstrap. *J. of the American Stat. Assoc.*, 89(428):1303–1313, 1994.

[183] S. Aoki, M. Fukugita, S. Hashimoto, Y. Iwasaki, K. Kanaya, Y. Kuramashi, H. Mino, M. Okawa, A. Ukawa, and T. Yoshié. Analysis of hadron propagators with one thousand configurations on a $24^3 \times 64$ lattice at $\beta = 6.0$. *Nuclear Physics B - Proceedings Supplements*, 47(1–3):354–357, 1996. ISSN 0920-5632. URL http://www.sciencedirect.com/science/article/pii/0920563296000722.

[184] C. Helmes, C. Jost, B. Knippschild, C. Liu, J. Liu et al. Hadron-Hadron Interactions from $N_f = 2+1+1$ Lattice QCD: isospin-2 $\pi\pi$ scattering length. 2015. arXiv:1506.00408 [hep-lat].

[185] A. Juttner and M. Della Morte. Heavy quark propagators with improved precision using domain decomposition. *PoS*, LAT2005:204, 2006. arXiv:hep-lat/0508023 [hep-lat].

[186] M. P. Forum. MPI: A Message-Passing Interface Standard. Technical report, Knoxville, TN, USA, 1994.

[187] G. E. Moore. Cramming More Components Onto Integrated Circuits. *Electronics Magazine*, 38(8), April 1965.

[188] J. M. Tendler, J. S. Dodson, J. Fields, H. Le, and B. Sinharoy. POWER4 system microarchitecture. *IBM Journal of Research and Development*, 46(1):5–25, 2002.

[189] B. Sinharoy, R. N. Kalla, J. M. Tendler, R. J. Eickemeyer, and J. B. Joyner. POWER5 system microarchitecture. *IBM journal of research and development*, 49(4.5):505–521, 2005.

[190] E. A. Lee. The problem with threads. *COMPUTER*, 39:2006, 2006.

[191] L. Dagum and R. Menon. OpenMP: An Industry-Standard API for Shared-Memory Programming. *IEEE Comput. Sci. Eng.*, 5(1):46–55, January 1998. ISSN 1070-9924. URL http://dx.doi.org/10.1109/99.660313.

[192] K. Jansen and C. Liu. Kramers equation algorithm for simulations of {QCD} with two flavors of Wilson fermions and gauge group {SU} (2). *Nuclear Physics B*, 453(1–2):375–392, 1995. ISSN 0550-3213. URL http://www.sciencedirect.com/science/article/pii/055032139500427T.

[193] C. Liu, A. Jaster, and K. Jansen. Liapunov exponents and the reversibility of molecular dynamics algorithms. *Nucl.Phys.*, B524:603–617, 1998. arXiv:hep-lat/9708017 [hep-lat].

[194] W. Kahan. Pracniques: Further Remarks on Reducing Truncation Errors. *Commun. ACM*, 8(1):40–, January 1965. ISSN 0001-0782. URL http://doi.acm.org/10.1145/363707.363723.

[195] J. M. Bull. Measuring synchronisation and scheduling overheads in OpenMP. In *Proceedings of First European Workshop on OpenMP*, volume 8, page 49. 1999.

[196] N. R. Fredrickson, A. Afsahi, and Y. Qian. Performance Characteristics of openMP Constructs, and Application Benchmarks on a Large Symmetric Multiprocessor. In *Proceedings of the 17th Annual International Conference on Supercomputing*, ICS '03, pages 140–149. ACM, New York, NY, USA, 2003. ISBN 1-58113-733-8. URL http://doi.acm.org/10.1145/782814.782835.

[197] M. Geimer, F. Wolf, B. J. N. Wylie, E. Ábrahám, D. Becker, and B. Mohr. The Scalasca performance toolset architecture. *Concurrency Computat.: Pract. Exper.*, 22:702–719, 2010.

[198] O. Tatebe, M. Sato, and S. Sekiguchi. Impact of OpenMP Optimizations for the MGCG Method. In M. Valero, K. Joe, M. Kitsuregawa, and H. Tanaka, editors, *High Performance Computing*, volume 1940 of *Lecture Notes in Computer Science*, pages 471–481. Springer Berlin Heidelberg, 2000. ISBN 978-3-540-41128-4. URL http://dx.doi.org/10.1007/3-540-39999-2_44.

[199] G. M. Amdahl. Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, spring joint computer conference*, pages 483–485. ACM, 1967.

[200] J. Shalf, S. Dosanjh, and J. Morrison. Exascale computing technology challenges. In *High Performance Computing for Computational Science–VECPAR 2010*, pages 1–25. Springer, 2011.

[201] W. J. Bolosky and M. L. Scott. False Sharing and Its Effect on Shared Memory Performance. In *Proc. of the USENIX Symposium on Experiences with Distributed and Multiprocessor Systems (SEDMS IV)*, pages 57–71. San Diego, CA, 1993.

[202] P. A. Boyle. The BAGEL assembler generation library. *Computer Physics Communications*, 180(12):2739–2748, 2009. ISSN 0010-4655. URL http://www.sciencedirect.com/science/article/pii/S0010465509002616.

[203] S. Heybrock, B. Joó, D. D. Kalamkar, M. Smelyanskiy, K. Vaidyanathan, T. Wettig, and P. Dubey. Lattice QCD with Domain Decomposition on Intel®Xeon Phi™Co-processors. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '14, pages 69–80. IEEE Press, Piscataway, NJ, USA, 2014. ISBN 978-1-4799-5500-8. URL http://dx.doi.org/10.1109/SC.2014.11.

[204] G. Bronevetsky and B. Supinski. Complete Formal Specification of the OpenMP Memory Model. *International Journal of Parallel Programming*, 35(4):335–392, 2007. ISSN 0885-7458. URL http://dx.doi.org/10.1007/s10766-007-0051-4.

[205] A. Abdel-Rehim, C. Alexandrou, N. Anastopoulos, G. Koutsou, I. Liabotis et al. PLQCD library for Lattice QCD on multi-core machines. *PoS*, LATTICE2013:419, 2014. arXiv:1405.0700 [hep-lat].

[206] C. Jung. Overview of Columbia Physics System. *PoS*, LATTICE2013:417, 2014.

[207] S. A. Gottlieb and S. Tamhankar. Benchmarking MILC code with OpenMP and MPI. *Nucl.Phys.Proc.Suppl.*, 94:841–845, 2001. arXiv:hep-lat/0011037 [hep-lat].

[208] B. Joó, M. Smelyanskiy, D. D. Kalamkar, and K. Vaidyanathan. Chapter 9 - Wilson Dslash Kernel From Lattice QCD Optimization. In J. Reinders and J. Jeffers, editors, *High Performance Parallelism Pearls Volume Two: Multicore and Many-core Programming Approaches*, volume 2, pages 139–170. Morgan Kaufmann, Boston, MA, USA, 2015. ISBN 978-0-12-803819-2. URL http://www.sciencedirect.com/science/article/pii/B9780128038192000239.

[209] R. Haring, M. Ohmacht, T. Fox, M. Gschwind, D. Satterfield, K. Sugavanam, P. Coteus, P. Heidelberger, M. Blumrich, R. Wisniewski, A. Gara, G.-T. Chiu, P. Boyle, N. Chist, and C. Kim. The IBM Blue Gene/Q Compute Chip. *Micro, IEEE*, 32(2):48–60, March 2012. ISSN 0272-1732.

[210] M. Luscher. Local coherence and deflation of the low quark modes in lattice QCD. *JHEP*, 07:081, 2007. arXiv:0706.2298 [hep-lat].

[211] P. A. Boyle. Hierarchically deflated conjugate gradient. 2014. arXiv:1402.2585 [hep-lat].

[212] A. Frommer, K. Kahl, S. Krieg, B. Leder, and M. Rottmann. Adaptive Aggregation Based Domain Decomposition Multigrid for the Lattice Wilson Dirac Operator. *SIAM J. Sci. Comput.*, 36:A1581–A1608, 2014. arXiv:1303.1377 [hep-lat].

[213] A. Stathopoulos and K. Orginos. Computing and deflating eigenvalues while solving multiple right hand side linear systems in quantum chromodynamics. *SIAM J. Sci. Comput.*, 32:439–462, 2010. arXiv:0707.0131 [hep-lat].

[214] M. A. Clark, R. Babich, K. Barros, R. C. Brower, and C. Rebbi. Solving Lattice QCD systems of equations using mixed precision solvers on GPUs. *Comput. Phys. Commun.*, 181:1517–1528, 2010. arXiv:0911.3191 [hep-lat].

[215] P. L'Ecuyer. Random Number Generation. In J. E. Gentle, W. K. Härdle, and Y. Mori, editors, *Handbook of Computational Statistics*, Springer Handbooks of Computational Statistics, pages 35–71. Springer Berlin Heidelberg, 2004. ISBN 978-3-642-21550-6. URL http://dx.doi.org/10.1007/978-3-642-21551-3_3.

[216] A. M. Ferrenberg, D. P. Landau, and Y. J. Wong. Monte Carlo simulations: Hidden errors from "good" random number generators. *Phys. Rev. Lett.*, 69:3382–3384, Dec 1992. URL http://link.aps.org/doi/10.1103/PhysRevLett.69.3382.

[217] M. Lüscher. A portable high-quality random number generator for lattice field theory simulations. *Computer Physics Communications*, 79(1):100–110, 1994. ISSN 0010-4655. URL http://www.sciencedirect.com/science/article/pii/0010465594902321.

[218] R. Frezzotti and K. Jansen. A polynomial hybrid Monte Carlo algorithm. *Physics Letters B*, 402(3–4):328–334, 1997. ISSN 0370-2693. URL http://www.sciencedirect.com/science/article/pii/S0370269397004759.

LIST OF FIGURES

SELBSTÄNDIGKEITSERKLÄRUNG

Ich erkläre, dass ich die Dissertation selbständig und nur unter Verwendung der von mir gemäß § 7 Abs. 3 der Promotionsordnung der Mathematisch-Naturwissenschaftlichen Fakultät, veröffentlicht im Amtlichen Mitteilungsblatt der Humboldt-Universität zu Berlin Nr. 126/2014 am 18.11.2014 angegebenen Hilfsmittel angefertigt habe.

Bartosz Kostrzewa Berlin, 28.12.2015