A study of QCD at finite density using complex Langevin dynamics

Swansea University
Prifysgol Abertawe

Felipe Attanasio

Department of Physics

Swansea University

Submitted to Swansea University in fulfilment
of the requirements for the degree of

Doctor of Philosophy

2017

Abstract

Numerical simulations are a standard tool to investigate field theories in non-perturbative regimes. Typical algorithms used to evaluate path integrals in Euclidean space rely on importance sampling methods; i.e., a probabilistic interpretation of the Boltzmann weight e^{-S}. However, many theories of interest suffer from the infamous sign problem: the action is complex and the Boltzmann weight cannot be used as a probability distribution. Complex Langevin simulations allow numerical studies of theories that exhibit the sign problem, such as QCD at finite density.

In this thesis, we study methods to investigate the phase diagram of QCD in the temperature–chemical potential plane, using the complex Langevin method. We provide results on the phase diagram for the heavy-dense approximation of QCD (HDQCD) for three spatial volumes, using complex Langevin and the gauge cooling technique. We also present polynomial fits of the critical temperature as function of the chemical potential for each volume. Subsequently, we discuss instabilities encountered during this study, which motivated a novel technique, named Dynamic Stabilisation, which will be introduced and the theoretical ideas behind it, explained.

Dynamic stabilisation was, then, used in an investigation of the dependency of the critical chemical potential on the hopping parameter. The two previous studies were used to guide a second examination of the HDQCD phase diagram, focussed around the phase boundary.

Lastly, we present preliminary results on the phase diagram of QCD with fully dynamical quarks at high temperatures. This shows that complex Langevin, augmented with gauge cooling and dynamic stabilisation, is suited for investigating QCD at finite chemical potential.

Declaration

This work has not previously been accepted in substance for any degree and is not being concurrently submitted for any degree.

Signed: ..... (candidate)

Date: .....

Statement 1

This thesis is the result of my own investigations, except where otherwise stated. Where correction services have been used, the extent and nature of the correction is clearly marked in a footnote(s). Other sources are acknowledged by footnotes giving explicit references. A bibliography is appended.

Signed: ..... (candidate)

Date: .....

Statement 2

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed: ..... (candidate)

Date: .....

Contents

Acknowledgements

Thanks to my supervisors Gert Aarts and Chris Allton, for the support and guidance during my PhD studies.

I would also like to thank my collaborators, Benjamin Jäger, Dénes Sexty, Erhard Seiler and Ion-Olimpiu Stamatescu for invaluable discussions throughout my work. Special thanks to Dénes Sexty for providing the BGQCD code, upon which our current code was based, and to Ben Jäger for the (un)countably many discussions and digressions.

I am also thankful to all my friends and colleagues in Swansea, who made possible for me to have a great time here.

I thank the Brazilian government, who funded my work through the Science without Borders programme, scholarship number Bex 9463/13-5. This work has been carried out at the DiRAC2 BlueGene/Q supercomputing facility at the University of Edinburgh's Advanced Computing Facility, and the HPC-Wales clusters, funded by the European Regional Development Fund through the Welsh European Funding Office, part of the Welsh Government.

I am immensely grateful to my mother, Celia, my grandfather, Nivaldo, and my wife, Mayra, for the uncountably infinite amount of support and help.

あなたの隣(となり)を歩(ある)いていきたい
あなたの隣(となり)を明日(あした)も明後日(あさって)も
相合傘(あいあいがさ)ー고윤하

List of Lists

List of Figures

List of Tables

Chapter 1

Introduction

Since being proposed in the 1970s, the lattice formulation of Quantum Chromodynamics (QCD) has stabilished itself as the main tool to investigate non-perturbative phenomena in the interactions of hadrons. The technique enables the study of strongly interacting matter, for which traditional perturbation theory is not applicable. Moreover, with ever increasing computational power available, both precision measurements and investigations into questions concerning non-perturbative behaviour are becoming feasible.

An important milestone of modern particle physics is the understanding of the QCD phase diagram—i.e., how strongly interacting matter behaves under different thermodynamical conditions. The thermodynamics of QCD will be reviewed in the next section, where we explain what the current knowledge of the phase diagram is, and how analytical and numerical techniques have been used to make further conjectures. We present a sketch of the diagram, mention some applications, and discuss modern computational methods currently used to study the phase diagram. Good reviews on this topic include [1–4].

In the subsequent section we proceed to examine the motivation for the lattice approach to strongly interacting theories. We then discuss how the quark and gluon actions are translated into the lattice formulation and explain how the infamous *sign problem*, which burdens computer simulations of QCD at finite density, arises. We finish the section explaining the heavy-dense approximation of QCD, which shares many interesting features with QCD at nonzero baryon density, but is computationally

cheaper, allowing rather inexpensive testing of numerical methods before applying them to the full theory.

1.1 Thermodynamics of QCD

The strongly interacting sector of the Standard Model, composed of quarks and gluons, is described by QCD. This non-abelian gauge theory exhibits three striking features, namely chiral symmetry breaking, asymptotic freedom and confinement. These features will be reviewed below and are largely responsible for the existence of hadrons at small temperatures and low net baryon densities, compared to the intrinsic QCD scale $\Lambda_{\rm QCD} \sim 200$ MeV.

An early understanding of how different thermodynamical environments affect the strongly interacting matter came from the study of the hadron resonance gas model [5], where it was realised that the observed resonance spectrum leads to a limiting temperature, the Hagedorn temperature, in a gas of hadrons. In other words, above this temperature ordinary nuclear matter is not expected to exist.

Since the formulation of QCD, it has been shown via renormalisation group analysis [6, 7] that its running coupling vanishes in the limit of high energies. This is known as *asymptotic freedom*: the dominant degrees of freedom in high temperature environments are no longer hadrons but their constituents, quarks and gluons, which form what is known as the quark–gluon plasma (QGP). Conversely, at low energies the coupling constant is very large and quarks are permanently bound inside hadrons. However, at large baryonic densities, n_B, hadrons start to overlap and screen this potential. Quarks are then liberated, and may move over macroscopic distances. In both high energy regimes (high temperature T or large n_B) QCD can be studied perturbatively.

Observed values for hadronic masses, such as the proton's ($\sim$ 1GeV), suggest that their three constituent quarks should have masses in the order of hundreds of MeV, assuming quarks carry all the mass. This mass is generated by QCD interactions which, because of the large coupling, cannot be studied perturbatively. However, in the QGP phase quarks are deconfined, and propagate with their bare massees. This change suggests a phase transition from a state of heavy constituent quarks to another, with light current quarks. In these energetic media the light quarks are essentially massless and thus possess approximate chiral symmetry. This symmetry, realised for

massless quarks, implies that left- and right-handed fermions behave in the same way. A mass term, which mixes both helicity states, explicitly breaks this symmetry. Its order parameter, the chiral condensate $\langle\overline{\psi}\psi\rangle$, can be used to identify the broken and unbroken phases. The chiral condensate is related to the pion mass via the Gell-Mann–Oakes–Renner relation [8], and chiral perturbation theory and lattice results indicate that it decreases with the temperature and chemical potential.

QCD at thermodynamical equilibrium has, as mentioned above, two external control parameters: the temperature, T, and the baryon number density, n_B. We study QCD in the grand canonical ensemble, where the quark chemical potential, $\mu \equiv \mu_q = \mu_B/3$, is introduced as a conjugate variable to the quark number density $n \equiv n_q = 3n_B$. Considering the topics discussed above, some conjectures have been made about the QCD phase diagram in the T–μ plane. Either at finite temperature, above the critical one, and small baryon density ($T \gg T_c$ and $\mu_B \ll T$) or at asymptotically high densities ($\mu_B \gg \Lambda_{\mathrm{QCD}}$), firm statements can be made from perturbation theory. Other parts of the phase diagram have been investigated using numerical methods, some of which will be briefly discussed below. Some noteworthy sections of the diagram are:

- *Quark deconfinement at* $\mu_B = 0$: The gauge part of QCD (with $N_f = 0$ quark flavours) has a first-order deconfinement transition at the critical temperature $T_c \sim 270$ MeV [9]. When light quark flavours are present, chiral symmetry analyses indicate a crossover between the hadronic phase and the QGP for realistic u, d and s quark masses [10, 11] at $T \sim 155$ MeV [12, 13].

- *Liquid–gas transition of nuclear matter*: When the baryon chemical potential is around the nucleon mass, the density is expected to vary from zero to the normal nuclear density $n_0 = 0.17$ fm^{-3}. Below n_0 the nuclear matter is fragmented into droplets with $n_B = n_0$, such that $n_B < n_0$ is achieved on spatial average. This behaviour, typical of first order phase transitions of the liquid-gas type, weakens as the temperature grows and is expected to end at a second order critical point.

- *QCD critical points*: Chiral models suggest that QCD has a critical point in the transition between hadronic matter and QGP. This transition is of first order, above a certain critical chemical potential, while below it, it is a crossover that connects to the $\mu_B = 0$ region [14–17].

- *Colour superconductivity*: For asymptotically large baryon chemical potentials, QCD becomes weakly coupled, and some knowledge can be borrowed from condensed matter physics, with quarks playing the part of electrons. With this analogy in mind, it is possible to expect the formation of Cooper pairs, leading to a superconducting phase carrying colour charge [18–21]. This state, known as colour-flavour locked phase, is a superfluid and breaks chiral symmetry.

Figure 1.1 summarises this discussion in a sketch of the QCD phase diagram. In it, the different phases and their expected locations can be seen. The dashed lines represent first order phase transitions, with the conjectured critical point being of second order.

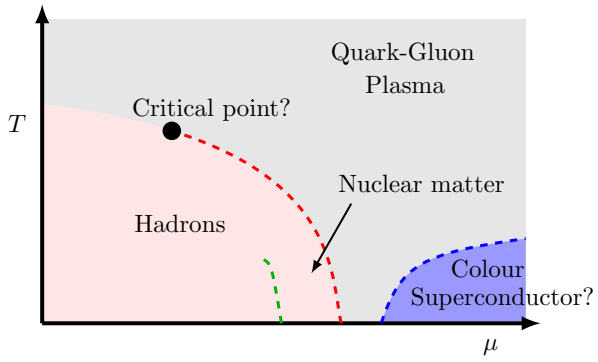

Figure 1.1: A sketch of the phase diagram of QCD, based on the statements above, originating from perturbation theory and numerical calculations.

An important application of understanding the QCD phase diagram is the determination of the QCD equation of state. This is of particular phenomenological significance for the study of hot or dense systems. Heavy-ion collision experiments, which have a hydrodynamical description, as well as investigations of cosmological objects and the physics of the early universe, can use the equation of state as input. One such example is the study of the stability of neutron stars. The Tolman–Oppenheimer–Volkoff equation [22], when supplied with an equation of state relating density and pressure, completely determines the structure of a spherically symmetric body of isotropic material in static gravitational equilibrium.

Experiments involving heavy-ion collisions have been very successful in creating quark-gluon plasmas. In the high temperature phase there is evidence from the Relativistic Heavy Ion Collider (RHIC), at Brookhaven National Laboratory (BNL), USA,

for an 'almost perfect liquid' picture for the strongly interacting medium. Hydrodynamical quantities, such as transport coefficients, shear and bulk viscosities and conductivities, can then be analysed numerically and compared to these experiments.

Current experiments at RHIC, at the Large Hadron Collider (LHC) in Switzerland, and planned facilities and experiments at the Facility for Antiproton and Ion Research (FAIR), in Darmstadt, Germany, and the Nuclotron-based Ion Collider Facility (NICA), in Dubna, Russia, will further explore the QCD phase diagram. This allows for the study of how quarks and gluons interact in extreme conditions, where the high baryonic density dominates the energy scale, or the early universe where very high temperatures were present. Future experiments are being proposed and constructed seeking deeper understanding of cold/dense systems.

1.1.1 Simulation methods

The study of QCD outside limiting cases of very small chemical potentials and high temperature, or very large chemical potentials outlined above, has to be done via numerical simulations on Euclidean spacetimes. However, a problem hinders these investigations: the addition of chemical potential to Euclidean field theories leads to the so-called sign problem. These topics will be discussed in more detail in the next section. For now we just quote the main effects related to the sign problem and briefly discuss the main methods to deal with it.

At a finite chemical potential the effective quark action obtained by integrating out the fermionic fields in the Euclidean path integral is complex-valued. The evaluation of the functional integral requires dealing with very precise cancelations of positive and negative numbers. Some of the most common methods to circumvent the sign problem are [1, 23, 24]

- *Reweighting*: Given an operator $\mathcal{O}$, its expectation value at finite μ can be written in terms of an ensemble average at $\mu = 0$ as

$$\langle \mathcal{O} \rangle_\mu = \frac{\langle \mathcal{O} R(\mu) \rangle_{\mu=0}}{\langle R(\mu) \rangle_{\mu=0}} , \tag{1.1}$$

  where $R(\mu)$ is referred to as the reweighting factor. The right hand side of the equation is evaluated at zero chemical potential and is therefore sign problem-free. However, the overlap between configurations at $\mu = 0$ and finite μ is known

to decrease exponentially with the volume, in a situation known as the *overlap problem*, and the evaluation of $\langle \mathcal{O} \rangle_\mu$ becomes difficult when $\langle R(\mu) \rangle_{\mu=0}$ is close to zero or its error becomes large. Nevertheless, reweighting is a powerful tool to guide, or when combined with, other methods, and when the overlap problem is under control. A short introduction to reweighting can be found in appendix A.

- *Taylor expansion*: For sufficiently small μ it is possible to expand observables in powers of μ/T. Examples include expansions around $\mu = 0$ of the reweighting factor mentioned previously [25], pressure and quark number susceptibility [26] and baryon number susceptibility [27]. This method can be used to estimate the location of critical point via the radius of convergence, but extrapolations for larger μ/T are difficult.

- *Imaginary chemical potential*: If μ is purely imaginary, the effective quark action becomes real and there is no sign problem [28, 29]. Analytical continuation can then be used to get information about QCD at real chemical potentials from lattice simulations at imaginary ones,

$$\langle \mathcal{O} \rangle_{\mu_i} = \sum_{n=0}^{\infty} c_n \left(\frac{\mu_i}{T} \right)^n \rightarrow \langle \mathcal{O} \rangle_\mu = \sum_{n=0}^{\infty} c_n \left(\frac{-i\mu}{T} \right)^n , \tag{1.2}$$

  where we used μ_i to identify imaginary chemical potentials. This region of the phase diagram has interesting phases of its own, with periodicity in μ_i known as the Roberge–Weiss periodicity [30].

- *Canonical ensemble*: In the thermodynamic limit, all statistical mechanics ensembles are equivalent. In particular, the grand canonical ensemble with chemical potential μ and average particle number N is equivalent to the canonical ensemble with particle number N. The conversion between them requires a Fourier transform in terms of the imaginary chemical potential [28, 31–33],

$$\langle \mathcal{O} \rangle_N = \int_0^{2\pi} \frac{d\phi}{2\pi} e^{-iN\phi} \langle \mathcal{O} \rangle_{\mu_i = \phi T} . \tag{1.3}$$

  This conversion is done with a Laplace transform for real chemical potentials. For large volumes the difficulties caused by the sign problem manifest in the

integration with respect to $\phi = \mu_i/T$. Moreover, the order of the phase transition is sensitive to how the thermodynamic limit is approached [34].

- *Density of states*: The density of states method involves writing the partition function in terms of a density function that enumerates the number of states for a given hyper-surface in configuration space. Some examples of surfaces include the phase of the complex-valued effective fermion action [35], the plaquette [36], or the action [37]. This density is given, in the latter case, by

$$\rho(E) = \langle \delta \left(S[\phi] - E \right) \rangle \,, \tag{1.4}$$

with expectation values now reading

$$\langle \mathcal{O} \rangle = \int dE \, \rho(E) \mathcal{O}(E) e^{-\beta E} \,, \tag{1.5}$$

where ϕ generically represents the relevant fields and $S[\phi]$ is the classical action of the theory. This method's success relies on determining ρ effectively and to high precision. Two recent techniques are the Logarithmic Linear Relaxation (LLR) [37, 38] and the Functional Fit Approach (FFA) [39] methods. Recent works include [40, 41].

- *Lefschetz thimbles*: In this approach, the integration is done over the paths known as Lefschetz thimbles [42]. These are paths that pass through fixed (or critical) points of the complex action, which live in a complexified space, and follow contours where the imaginary part of the action is constant. This does not eliminate the sign problem, but replaces it by two: one related to the curvature of the integration contour (known as residual); and the other appearing in cases where more than one thimble contributes (known as global). These can be dealt with independently, with the first one being easily tractable when one thimble has a dominant contribution, and reweighting being applicable to the latter, at least in simple cases [43]. The first use of Lefschetz thimbles on the lattice was carried out in [44], and a comparison with complex Langevin can be found in [45].

- *Complex Langevin*: This method is based on stochastic quantisation [46], where the fields evolve in a fictitious time dimension following a Langevin equation.

Using the generic representation shown above,

$$\frac{\partial\phi(\theta)}{\partial\theta} = -\frac{\delta S[\phi]}{\delta\phi(\theta)} + \eta(\theta)\,, \tag{1.6}$$

with θ being the fictitious (Langevin) time and $\eta(\theta)$ a white noise field. Quantum expectation values are recovered in the $\theta \to \infty$ limit. Allowing the fields to be complex-valued and complexifying their Langevin evolution can circumvent the sign problem by deforming the path integration into a larger (complexified) phase space. Some early studies include [47–49], while more recent results are found in [50–52]. This method is the central tool used in this thesis and will be reviewed in detail in chapter 2.

1.2 The lattice approach

We start by considering the QCD action and its fermionic and bosonic constituents in Euclidean spacetime. From those, we can justify the need for the lattice approach as a technique that allows for computations not restricted to a particular interval in the energy spectrum. Then the lattice formulation of the quark and gluon actions are discussed, as well as how the sign problem manifests itself in this context.

The action that describes the strong interactions, $S_{\mathrm{QCD}} = S_{\mathrm{F}} + S_{\mathrm{G}}$, consists of fermionic, S_{F}, and bosonic, S_{G}, parts. Throughout this thesis we work in natural units,

$$\hbar = c = 1\,. \tag{1.7}$$

The fermionic portion is formed by Dirac fermions minimally coupled to SU(3) gauge fields,

$$S_{\mathrm{F}} = \int d^4x\, \overline{\psi}\,(\gamma_\mu(\partial_\mu + iA_\mu) + m)\,\psi\,, \tag{1.8}$$

where the Dirac spinors $\psi = \psi(x)$ and $\overline{\psi} = \overline{\psi}(x)$ carry Dirac, colour and flavour indices, $A_\mu = \lambda^a A^a_\mu$ is the gluon field, and m is the mass for each quark flavour. The group's generators, λ^a, are reviewed in appendix B. The Euclidean γ matrices, with $\mu = 1, 2, 3, 4$, satisify the anti-commutation relations

$$\{\gamma_\mu, \gamma_\nu\} = 2\delta_{\mu\nu}\mathbb{1}\,. \tag{1.9}$$

In QCD, the quark fields have 4 spinor, 3 colour and 6 flavour (up, down, charm, strange, top and bottom) components. Quarks also carry weak and electric charges.

The bosonic part of the QCD action, the SU(3) Yang–Mills term, is given by

$$S_{\mathrm{G}} = \frac{1}{2g^2} \int d^4x \, \mathrm{Tr}\left[F^{\mu\nu} F_{\mu\nu}\right] , \quad F_{\mu\nu} = \left[\partial_\mu A^c_\nu - \partial_\nu A^c_\mu - f^{cde} A^d_\mu A^e_\nu\right] \lambda^c , \tag{1.10}$$

with g representing the gauge coupling constant, $A^c_\mu = A^c_\mu(x)$ the colour and spacetime components of the gauge field, $F_{\mu\nu}$ the field strength tensor and λ^c and f^{cde} the generators and structure constants of SU(3), respectively. The above action is similar to that of electrodynamics, with one copy for each generator, together with the self-interacting term proportional to the structure constants. This latter term gives rise to cubic and quartic self-interactions in the action, making Yang–Mills theories highly complex.

Renormalisation group calculations made in the 1970s [6, 7] have shown that Yang–Mills actions are asymptotically free—i.e., they become weakly coupled at high energies. Because of this result, conventional perturbation theory methods could be applied to QCD leading to, e.g., calculations of cross sections in deep inelastic processes.

On the other hand, for low energy processes QCD was still untractable. After the seminal work by Kenneth Wilson in 1974 [53], which provided a way of writing the QCD action in a format suitable for computer simulations, the study of low energy hadronic interactions became possible.

Typically, the Euclidean spacetime is discretised in a 4-dimensional lattice, with neighbouring points separated by the lattice spacing, a. Spacetime points are then identified by integer multiples of a: $x_\mu = a n_\mu$, with n_μ a 4-dimensional vector with integer components. The simulation box then has spatial and temporal extents

$$L = a N_\sigma \quad \text{and} \quad L_t = a N_\tau = 1/T , \tag{1.11}$$

respectively, with N_σ and N_τ integers and T being the temperature. QCD at finite temperature will be discussed in sec. 1.2.3. In this context, integrals become sums over all lattice points

$$\int d^4x \to a^4 \sum_x , \tag{1.12}$$

with x also being used to identify lattice points. For the discrete derivatives we choose

the symmetric form

$$\partial_\mu \phi(x) \to \frac{\phi(x+\hat{\mu}) - \phi(x-\hat{\mu})}{2a}\,, \tag{1.13}$$

where $\hat{\mu}$ denotes a unit vector in the μ direction.

One characteristic of the lattice approach is that it naturally provides a UV momentum regulator. In each spacetime direction the momentum is constrained in the interval $[-\pi/a, \pi/a]$, known as the Brillouin zone. In the continuum limit, as $a \to 0$, the usual integrations in momentum space are recovered.

1.2.1 Quarks on the lattice

We begin with the free fermion action in continuum Euclidean spacetime, given by eq. (1.8) with $A_\mu = 0$, and apply the prescriptions of eqs. (1.12) and (1.13):

$$S_{\mathrm{F}} \to a^4 \sum_x \overline{\psi}_x \left(\gamma_\mu \frac{\psi_{x+\hat{\mu}} - \psi_{x-\hat{\mu}}}{2a} + m\psi_x \right)\,, \tag{1.14}$$

where we have written the spacetime position as a subscript. Usually periodic or anti-periodic boundary conditions are imposed to mimic an infinite geometry,

$$\psi_{x+L\hat{\mu}} = \pm \psi_x\,. \tag{1.15}$$

Anti-periodic boundary conditions are used in the temporal direction to account for Fermi–Dirac statistics.

The terms in the discretised action must be made dimensionless in order to make it suitable for computer simulations. In four spacetime dimensions this is accomplished via the replacements

$$a^{3/2}\psi \to \hat{\psi}\,, \tag{1.16}$$

$$a^{3/2}\overline{\psi} \to \hat{\overline{\psi}}\,, \tag{1.17}$$

$$am \to \hat{m}\,. \tag{1.18}$$

This way S_{F} is written in terms of dimensionless variables only as

$$S_{\mathrm{F}} = \sum_x \overline{\psi}_x \left(\gamma_\mu \frac{\psi_{x+\hat{\mu}} - \psi_{x-\hat{\mu}}}{2} + m\psi_x \right)\,, \tag{1.19}$$

where, for simplicity, we have dropped the hats. Unless otherwise stated, we will refer exclusively to dimensionless quantities from this point on.

This formulation is known as the naïve discretisation of the fermion action. The reason behind this name is that it contains lattice artefacts that do not vanish in the continuum limit. For convenience, we rewrite it as

$$S_{\mathrm{F}} = \sum_{x,y} \overline{\psi}_x D(x,y) \psi_y \,, \tag{1.20}$$

with the Dirac operator $D(x,y)$ given by

$$D(x,y) = \frac{\gamma_\mu}{2} \left(\delta_{x+\hat{\mu},y} - \delta_{x-\hat{\mu},y}\right) + m\delta_{x,y} \,. \tag{1.21}$$

The lattice quark inverse propagator, obtained by the Fourier transform of D, reads

$$\tilde{D}(q,p) = \delta(p-q) \left(\gamma_\mu \sin(p_\mu) + m\right) \,. \tag{1.22}$$

The propagator has the usual pole at $p = 0$, but also additional ones at $p_\mu = \pi$. This issue is known as the *fermion doubling problem* in d spacetime dimensions: the continuum limit of the free quark propagator contains contributions from 2^d fermion-like excitations.

1.2.1.1 Removing fermion doublers

There are many ways of dealing with doublers, each with their own advantages and disadvantages. Here we will focus on two of the most well established: the Wilson and Staggered (or Kogut–Susskind) formulations.

Wilson proposed in [53] the addition of a second-derivative term that removes the additional poles and vanishes in the contiuum limit. The Wilson fermion action is defined as

$$S_{\mathrm{W}} = \sum_x \overline{\psi}_x \left[\kappa \left(\gamma_\mu - r\right) \psi_{x+\hat{\mu}} + \kappa \left(\gamma_\mu + r\right) \psi_{x-\hat{\mu}} + \psi_x\right] \,, \tag{1.23}$$

where the hopping parameter,

$$\kappa = \frac{1}{2m + 8r} \,, \tag{1.24}$$

and the coefficient r have been introduced. Typically r is set to one. This change requires the quark fields to be re-scaled as $\psi \to (m+4r)^{-1/2}\psi$ and $\overline{\psi} \to (m+4r)^{-1/2}\overline{\psi}$.

The addition included by Wilson to remove the unphysical fermion poles, however, causes an explicit breaking of the chiral symmetry even for massless quarks,

$$\{D_{\mathrm{W}}, \gamma_5\} \neq 0. \tag{1.25}$$

In the continuum limit, with the lattice spacing restored, the Wilson action for massless fermions reads

$$S_{\mathrm{W}}^{\mathrm{cont}} = \int d^4x\, \overline{\psi} \left(\gamma_\mu \partial_\mu + ra\partial^2\right) \psi \,. \tag{1.26}$$

The second term on the right-hand side is irrelevant for low momenta (near $p = 0$), but is relevant at the unphysical ones. When $a \to 0$ the doublers are then removed from the theory, at the expense of the chiral symmetry.

Staggered fermions were proposed in [54], and reduce the degeneracy of the poles described above from 16 to 4[1], while maintaining a remnant chiral symmetry. It makes use of the symmetry made explicit by applying the *staggered transformation*

$$\psi_x \to \gamma_1^{x_1}\gamma_2^{x_2}\gamma_3^{x_3}\gamma_4^{x_4}\psi_x \,, \quad \overline{\psi}_x \to \overline{\psi}_x\gamma_4^{x_4}\gamma_3^{x_3}\gamma_2^{x_2}\gamma_1^{x_1} \,. \tag{1.27}$$

This transformation mixes Dirac and spacetime indices via products of gamma matrices raised to the power of the corresponding x_μ of site $x = (x_1, x_2, x_3, x_4)$.

The mass term in the action is invariant, since $\gamma_\mu^2 = \mathbb{1}$. In the kinetic term, ψ fields are shifted with respect to $\overline{\psi}$. Therefore, a factor of γ_μ will remain after the transformation. This extra γ_μ cancels the one already present in the action. Accounting for the necessary reordering of gamma matrices to achieve this cancellation, a generic term in the kinetic part of the action becomes

$$\overline{\psi}_x\gamma_\mu\psi_{x+\hat{\mu}} \to \eta_{x,\mu}\overline{\psi}_x\mathbb{1}\psi_{x+\hat{\mu}} \,, \tag{1.28}$$

where we have introduced the staggered sign function, also known as Kawamoto–Smit phases

$$\eta_{x,\mu} = (-1)^{\sum_{i=1}^{\mu-1} x_\mu} \,, \quad \eta_{x,1} = 1 \,. \tag{1.29}$$

The new action, which is diagonal in Dirac space and has the same form for all four

[1] In general, from 2^d to d.

Dirac components, reads

$$S_{\mathrm{KS}} = \sum_x \overline{\psi}_x \mathbb{1} \left(\eta_{x,\mu} \frac{\psi_{x+\hat{\mu}} - \psi_{x-\hat{\mu}}}{2} + m\psi_x \right) . \tag{1.30}$$

Keeping only one of the four identical components gives us the staggered fermion action, where the fields possess spacetime and colour indices only. By discarding 3 of the 4 components it is expected that only 4 degrees of freedom are left. More details on staggered fermions can be found in [55, 56].

Modern implementations of lattice quark actions include Ginsparg-Wilson [57], overlap [58], and domain wall [59]. These fermionic actions obey the Ginsparg-Wilson relation,

$$\gamma_5 D(x,y) + D(x,y)\gamma_5 = a \sum_z D(x,z)\gamma_5 D(z,x) . \tag{1.31}$$

This equation was based on renormalisation group transformations and replaces the continuum expression for the chiral symmetry. The factor of a on the right-hand side appears for dimensional reasons, and ensures the continuum result is recovered.

1.2.1.2 Gauging the colour symmetry

For now, we shall stick to the naïve lattice fermion implementation due to its simplicity. The concepts discussed here translate immediately to the other formulations. We shall see how the local SU(3) gauge symmetry of QCD is realised on the lattice, thus paving the way for a discrete version of the Yang–Mills action to be reviewed in the following section.

We start by considering the mass term $m\overline{\psi}_x\psi_x$. This term is trivially invariant under global and local SU(3) transformations, since

$$\psi_x \to \Omega\psi_x , \quad \overline{\psi}_x \to \overline{\psi}_x \Omega^\dagger , \tag{1.32}$$

with $\Omega = e^{i\lambda^a\omega^a}$ or $\Omega = \Omega_x = e^{i\lambda^a\omega^a_x}$ for global and local transformations, respectively. On the other hand, the kinetic term's elements change under a local transformation as

$$\overline{\psi}_x \psi_{x+\hat{\mu}} \to \overline{\psi}_x \Omega_x^\dagger \Omega_{x+\hat{\mu}} \psi_{x+\hat{\mu}} ; \tag{1.33}$$

i.e., they are not generally invariant.

As is done in the continuum, we augment the action by coupling the fermions to fields that behave under gauge transformations in such a way as to keep the action invariant. Elements of the SU(3) group, $U_{x,\mu} = \exp\left[i\lambda^a A^a_{x,\mu}\right]$, transform as

$$U_{x,\mu} \to \Omega_x U_{x,\mu} \Omega^\dagger_{x+\hat{\mu}}\,, \tag{1.34}$$

so that the combination

$$\overline{\psi}_x U_{x,\mu} \psi_{x+\hat{\mu}} \tag{1.35}$$

is invariant under local gauge transformations. The elements $U_{x,\mu}$ are known as gauge links, since they are the parallel transporters of colour charge between neighbouring lattice sites. The gauged fermion action then reads

$$S_{\mathrm{F}} = a^4 \sum_x \overline{\psi}_x \left(\gamma_\mu \frac{U_{x,\mu}\psi_{x+\hat{\mu}} - U^\dagger_{x-\hat{\mu},\mu}\psi_{x-\hat{\mu}}}{2} + m\psi_x \right). \tag{1.36}$$

This action is invariant under local gauge transformations at finite lattice spacing, and reduces to the continuum action in the naïve continuum limit.

1.2.2 Gluons on the lattice

With the gauge links introduced in the previous section, we now proceed to make them dynamical. The lattice action has to be gauge invariant, and reduce to the continuum one when $a \to 0$. The simplest non-trivial objects that can be constructed from the gauge links, $U_{x,\mu}$, are elementary closed loops, known as *plaquettes*, $U_{x,\mu\nu}$, defined as

$$U_{x,\mu\nu} = U_{x,\mu} U_{x+\hat{\mu},\nu} U^\dagger_{x+\hat{\nu},\mu} U^\dagger_{x,\nu}\,. \tag{1.37}$$

It is easy to see that upon a gauge transformation the plaquette changes as

$$U_{x,\mu\nu} \to \Omega_x U_{x,\mu\nu} \Omega^\dagger_x\,, \tag{1.38}$$

such that its trace is invariant.

Since the plaquettes transport the gauge fields along an elementary closed path it is expected that they are related to the field strenght tensor $F_{\mu\nu}(x)$ in some fashion. The

relation between them can be seen by using the Baker–Campbell–Hausdorff formula

$$e^A e^B = \exp\left(A + B + \frac{1}{2}[A,B] + \cdots\right) \tag{1.39}$$

on the gauge links in eq. (1.37) and the fact that

$$U_{x+\hat{\mu},\nu} = \exp\left[iA_{x+\hat{\mu},\nu}\right] = \exp\left[iA_{x,\nu} + a\partial_\mu A_{x,\nu} + O(a^2)\right] , \tag{1.40}$$

where we have used $A_{x,\mu} \equiv \lambda^a A^a_{x,\mu}$. These expansions combined lead to

$$\begin{aligned} U_{x,\mu\nu} &= \exp\left[ia^2\left(\partial_\mu A_{x,\nu} - \partial_\nu A_{x,\mu} + i[A_{x,\mu}, A_{x,\nu}]\right) + O(a^3)\right] \\ &= \exp\left[ia^2 F_{x,\mu\nu} + O(a^3)\right] . \end{aligned} \tag{1.41}$$

The Wilson gauge action is then given by

$$S_{\mathrm{G}}[U] = \frac{2}{g^2}\sum_x \sum_{\mu<\nu} \mathrm{ReTr}\left[\mathbb{1} - U_{x,\mu\nu}\right] = \frac{a^4}{2g^2}\sum_x \sum_{\mu,\nu} \mathrm{Tr}\left[F^2_{x,\mu\nu}\right] + O(a^2)\,. \tag{1.42}$$

Traditionally the coefficient in front of the Wilson action is written using the *inverse coupling*

$$\beta = \frac{6}{g^2}\,. \tag{1.43}$$

With that, and using the definition $\mathrm{Re}[z] = \frac{1}{2}(z + \overline{z})$, we finally write

$$S_{\mathrm{G}}[U] = \frac{\beta}{3}\sum_x \sum_{\mu<\nu} \mathrm{Tr}\left[\mathbb{1} - \frac{1}{2}\left(U_{x,\mu\nu} + U^\dagger_{x,\mu\nu}\right)\right] . \tag{1.44}$$

1.2.3 QCD at finite temperature

In statistical mechanics, the canonical partition function is given by

$$Z = \mathrm{Tr}\left[e^{-H/k_B T}\right] = \int d\phi\, \langle\phi|e^{-H/k_B T}|\phi\rangle\,, \tag{1.45}$$

where H is the system's Hamiltonian, k_B is Boltzmann's constant and T the temperature. From here on, we shall take $k_B = 1$. The trace indicates a sum over all possible states; this is generically written as the integral on the right-hand side in terms of a field ϕ. Similarly, the quantum mechanical evolution between states ϕ_i, at time 0, and

ϕ_f, at time t_f can be written in terms of a path integral as

$$\langle \phi_f | e^{iHt_f} | \phi_i \rangle = \int_{\phi_i}^{\phi_f} \mathcal{D}\phi \, \exp\left[i \int_0^{t_f} dt \mathcal{L} \right] , \tag{1.46}$$

with $\mathcal{L}$ being the Lagrangian density of the system. By doing the associations

$$it_f \to 1/T \, , \tag{1.47}$$

$$\phi_i = \phi_f = \phi \tag{1.48}$$

$$t \to \tau \tag{1.49}$$

it is possible to write the partition function in eq. (1.45) using eq. (1.46) as

$$\begin{aligned} Z &= \int d\phi \int_{\phi_i=\phi}^{\phi_f=\phi} \mathcal{D}\phi \, \exp\left[- \int_0^{1/T} d\tau \mathcal{L} \right] \\ &= \int_{\phi(0)=\phi(1/T)} \mathcal{D}\phi \, e^{-S[1/T]} \, , \end{aligned} \tag{1.50}$$

where S is the classical action of the system, and the integration is done over fields satisfying the periodic boundary condition $\phi(0) = \phi(1/T)$.

Hence, in the language of path integrals, where the classical action is the central object, finite temperature is introduced by restricting the integration in the Euclidean time direction in eqs. (1.8) and (1.10) to the interval $[0, 1/T]$, where we have taken Boltzmann's constant to be unity. On the lattice, the temperature can be controlled by changing the number of points in the temporal direction or by making the lattice anisotropic, with the spatial and temporal lattice spacings being different. This makes the Euclidean lattice formulation the natural framework to study field theories at non-zero temperature. Throughout this thesis we work with isotropic lattices.

At finite temperature, a commonly studied observable is the trace of the Polyakov loop $\mathcal{P}$ (also known as thermal Wilson line). On the lattice it is given by

$$P_{\vec{x}} = \frac{1}{3} \mathrm{Tr} \left[\mathcal{P}_{\vec{x}} \right] = \frac{1}{3} \mathrm{Tr} \left[\prod_{x_4=0}^{N_\tau - 1} U_{(\vec{x}, x_4), \hat{4}} \right] . \tag{1.51}$$

It represents the world-line of a static quark evolving in Euclidean time and wrapping around the (anti-)periodic box, while $P_{\vec{x}}^\dagger$ represents that of an anti-quark. The expec-

tation value of the product $P_{\vec{x}}P_{\vec{y}}^{\dagger}$ gives the free energy difference between an ensemble with a quark–anti-quark pair at positions $\vec{x}$ and $\vec{y}$, respectively, and that of the vacuum, via [55, 56]

$$\langle P_{\vec{x}}P_{\vec{y}}^{\dagger}\rangle = \exp\left[-\left(F_{q\bar{q}} - F_0\right)/T\right] . \tag{1.52}$$

It is also interesting to study the expectation value of the spatial average of the Polyakov loop

$$P = \frac{1}{N^3}\sum_{\vec{x}} P_{\vec{x}} , \tag{1.53}$$

since we assume translation invariance. By the reasoning outlined above it can be seen to be related to $e^{-(F_q - F_0)/T}$. In a confining regime, the energy necessary to create a single quark would be infinite. Therefore it is expected that $\langle P\rangle = 0$ for such a case. Conversely, if a non-zero expectation value is observed it means that $F_q - F_0$ is finite and deconfined quarks are possible. This indicates that $\langle P\rangle$ is an order parameter for the deconfinement transition for pure gauge theories.

As is known from statistical mechanics, phase transitions are accompanied by the breaking of a global symmetry. In the case of pure SU(3) Yang–Mills the action is not only invariant under gauge transformations, but posseses an additional symmetry. Transformations of all gauge links in the temporal direction in a given time slice x_4 of the form

$$U_{(\vec{x},x_4)\hat{4}} \to zU_{(\vec{x},x_4)\hat{4}} , \tag{1.54}$$

with $z = e^{2i\pi n/3}\mathbb{1}$, $n = 0, 1, 2$, belonging to the centre of the group, do not change the action. This happens because centre elements commute with all group elements and trivially cancel in closed loops that do not wrap around the time direction. The Polyakov loop, consequently, is not invariant under this transformation,

$$P \to zP . \tag{1.55}$$

If the ground state of the quantum system respects the symmetries of the classical action, configurations related to one another by centre transformations should have equal statistical weight making ensemble averages of the Polyakov loop vanish, since

$$\langle P\rangle \sim \langle P + zP + z^2P\rangle = \left(1 + e^{2i\pi/3} + e^{4i\pi/3}\right)\langle P\rangle = 0 \tag{1.56}$$

Thus, an unbroken centre symmetry is expected to be a signal of confinement. In situations where the ground state breaks the centre symmetry, a non-zero average Polyakov loop is observed and the theory is deconfined.

Strictly speaking, the Polyakov loop is only an order parameter of the pure Yang–Mills theory, since quarks do explicitly break the centre symmetry. One way to see this is by considering the hopping expansion of the fermion determinant [55, 56, 60, 61], where it is written as a sum over all possible closed loops on the lattice, including Polyakov loops. In particular, this sum will include terms of the form

$$\kappa^{N_\tau} \left(P + P^\dagger \right) , \tag{1.57}$$

where the hopping parameter $\kappa \sim 1/ma$. A centre transformation will modify such a term non-trivially, with the full action being no longer invariant. Nevertheless, despite $\langle P \rangle$ not being an order parameter when quarks are present, some information can still be gained by studying it. An especially good example of this is the heavy-dense approximation, to be discussed later, where the quarks are very heavy and remnants of the centre symmetry are present in the theory. Alternatively, the explicit breaking of the centre symmetry can be seen by applying a centre transformation to eq. (1.36), which will not leave the action invariant.

We note here that finite systems cannot exhibit phase transitions. They are associated with non-analytic behaviour in the partition function, which can only be achieved in the thermodynamic limit. In principle, assuming ergodicity and an infinitely long simulation time for a finite system, all allowed configurations should contribute and the expectation value of the Polyakov loop should always vanish. In this case, analyses of scatter plots of the Polyakov loop's real and imaginary parts, or of its absolute value, indicate whether only the trivial vacuum or the three vacua allowed by centre symmetry contribute.

1.2.4 Finite density and the sign problem in lattice QCD

In order to achieve a dynamical asymmetry between the numbers of quarks and anti-quarks, and therefore a non-zero density, a chemical potential is introduced. The grand canonical ensemble of statistical mechanics is, in a nutshell, constructed by adding

$\mu N/T$ to the statistical weight in the partition function,

$$Z = \mathrm{Tr}\left[e^{-(H-\mu N)/T}\right] , \tag{1.58}$$

with μ being the chemical potential and N the number of particles.

Naïvely adding the new term to a lattice field theory leads to a theory that does not possess a continuum limit. A proper discretisation of the chemical potential is obtained by realising that it couples to the Noether charge of the U(1) global symmetry of the quarks. Again, this explanation will use the naïve free fermion discretisation; all steps can be easily translated to the appropriate formulations, if necessary. The lattice version of the Noether current density reads [62]

$$j_{x,\mu} = \frac{1}{2}\left[\overline{\psi}_{x-\hat{\mu}}\gamma_\mu\psi_x + \overline{\psi}_x\gamma_\mu\psi_{x-\hat{\mu}}\right] . \tag{1.59}$$

Adding the corresponding charge density with suitable μ-dependent coefficients to the lattice fermion Lagrangian density leads to the correct way of including a chemical potential on the lattice. The charge densities can be combined with the temporal part of the kinetic term of the naïve fermion lattice action as

$$\sum_x \overline{\psi}_x\gamma_4\frac{f(\mu)\psi_{x+\hat{4}} - f^{-1}(\mu)\psi_{x-\hat{4}}}{2} \tag{1.60}$$

with $f(\mu)$ a real function. The anti-quark term is modified by the inverse of f so that closed quark loops with trivial topology; i.e., those that do not wrap around the periodic boundaries, are not affected. At $\mu = 0$ the original action should be recovered, therefore $f(0) = 1$. In order to recover the correct continuum behaviour f has to be linear in the chemical potential, $f(\mu) = 1 + \mu + O(\mu^2)$. Introducing the parallel transporter of the quark U(1) global symmetry in the same way as for the colour symmetry we arrive at $f(\mu) = e^\mu$ [63, 64]. Hence, the lattice action of fermions at finite density coupled to a gauge field can be written as

$$S_\mathrm{F} = \sum_x \overline{\psi}_x\left(\gamma_\nu\frac{e^{\mu\delta_{\nu,4}}U_{x,\nu}\psi_{x+\hat{\nu}} - e^{-\mu\delta_{\nu,4}}U^\dagger_{x-\hat{\nu},\nu}\psi_{x-\hat{\nu}}}{2} + m\psi_x\right) . \tag{1.61}$$

We remark that different prescriptions are possible, provided that they agree in the continuum limit [62].

A side-effect of including the chemical potential to an Euclidean field theory is the so-called *sign problem*. At $\mu = 0$, the Dirac operator $D(U)$ in Euclidean spacetime is γ_5-Hermitian,

$$D^\dagger(U) = \gamma_5 D(U) \gamma_5 \,. \tag{1.62}$$

This implies that its eigenvalues are either real or come in complex conjugate pairs; and hence the determinant is real. Since the quadratic fermion action can be integrated analyticaly,

$$Z = \int \mathcal{D}\overline{\psi}\mathcal{D}\psi\mathcal{D}U \, e^{-S_\mathrm{F}-S_\mathrm{G}} = \int \mathcal{D}U \, \det\left[D(U)\right] e^{-S_\mathrm{G}} = \int \mathcal{D}U \, e^{-S_\mathrm{G}+\ln\det[D(U)]} \,, \tag{1.63}$$

this leaves $\ln\det\left[D(U)\right]$ as an effective fermion action in terms of the gauge fields. For $\mu \neq 0$, however, $D(U,\mu)$ is not γ_5-Hermitian,

$$D^\dagger(U,\mu) = \gamma_5 D(U,-\mu)\gamma_5 \,, \tag{1.64}$$

which means that its determinant is a complex number, and consequently $e^{\ln\det[D(U,\mu)]}$ can no longer be interpreted as a probability distribution. This is the origin of the sign problem: the weight of the field configuration in the path integral ceases to be positive definite, in general. An accurate computer simulation would require precise cancellations of these oscillating factors among different configurations.

If a complex chemical potential is considered, then

$$D^\dagger(U,\mu) = \gamma_5 D(U,-\mu^*)\gamma_5 \,, \tag{1.65}$$

and hence for purely imaginary μ, the Dirac operator is again γ_5-Hermitian, with a real determinant. Then, the study of the phase structure of QCD becomes possible, as mentioned in section 1.1.

An interesting special case is the study of isospin chemical potential, μ_I, where two quark species with degenerate mass and opposite isospin charges are considered. Their chemical potentials are related via $\mu_a = \mu_I$ and $\mu_b = -\mu_I$. This causes the Dirac operator to have a block diagonal form, with each entry being a single flavour operator

$$\begin{pmatrix} D(U,\mu_I) & 0 \\ 0 & D(U,-\mu_I) \end{pmatrix} = \begin{pmatrix} D(U,\mu_I) & 0 \\ 0 & \gamma_5 D^\dagger(U,\mu_I)\gamma_5 \end{pmatrix} , \tag{1.66}$$

making use of the relation in eq. (1.64). Its determinant is real and positive, since

$$\det\left[D(U,\mu_I)\gamma_5 D^\dagger(U,\mu_I)\gamma_5\right] = \det\left[D(U,\mu_I)D^\dagger(U,\mu_I)\right] = \left|\det\left[D(U,\mu_I)\right]\right|^2 . \quad (1.67)$$

It is worth noting that at finite isospin density, the quark determinant corresponds to that of the phase quenched theory. Works on the thermodynamics of QCD at finite isospin density can be found in [65–67].

An important note is that the sign problem is not exclusive to fermionic fields. Bosonic fields in Euclidean spacetime also display a sign problem at finite real chemical potential. Some examples include the Bose gas [68, 69], the XY model [70] and the SU(3) spin model [71, 72].

1.2.5 The heavy-dense approximation

Because the evaluation of the fermion determinant in QCD is very demanding, it is useful to have a simpler model which shares some of QCD's features, such as the sign problem, where new methods and techniques can be tested first. One such model is QCD in the limit of heavy and dense quarks (HDQCD) [47, 73], where quarks are considered static in space but the full chemical potential dependency, which resides in the temporal hopping terms, is retained. In this approximation, the gluonic terms are not changed; i.e., the standard Wilson gauge action is used.

We start with the Dirac operator for Wilson fermions,

$$D_{xy} = \delta_{xy} - 2\kappa\left[e^{\mu\delta_{\nu,4}}\Gamma_{-\nu}U_{x,\nu}\delta_{x+\hat{\nu},y} + e^{-\mu\delta_{\nu,4}}\Gamma_{+\nu}U^\dagger_{x-\hat{\nu},\nu}\delta_{x-\hat{\nu},y}\right] , \quad (1.68)$$

where we have used $\Gamma_{\pm\nu} = (\mathbb{1}\pm\gamma_\nu)/2$. By dropping the spatial hopping terms we arrive at

$$\left[D_{\mathrm{HD}}(U,\mu)\right]_{xy} \approx \delta_{xy} - 2\kappa\left[e^{\mu}\Gamma_{-4}U_{x,4}\delta_{x+\hat{4},y} + e^{-\mu}\Gamma_{+4}U^\dagger_{x-\hat{4},4}\delta_{x-\hat{4},y}\right] , \quad (1.69)$$

whose determinant in spacetime and Dirac indices takes the form

$$\det D_{\mathrm{HD}}(U,\mu) = \prod_{N_f}\prod_{\vec{x}}\left\{\det\left[1+he^{\mu N_\tau}\mathcal{P}_{\vec{x}}\right]^2\det\left[1+he^{-\mu N_\tau}\mathcal{P}^{-1}_{\vec{x}}\right]^2\right\} . \quad (1.70)$$

The + sign originates from the anti-periodic boundary conditions in the temporal direction, and the powers of 2 from the γ-matrix structure. The Polyakov loop $\mathcal{P}$ and its inverse were defined in eq. (1.51); the exponent μN_τ is the lattice version of the ratio μ/T, $h = (2\kappa)^{N_\tau}$, and N_f is the number of quark flavours. We have used the inverse rather than the Hermitian adjoint of the Polyakov loop because they coincide for SU(N) fields, but using the inverse will be necessary in later chapters. The term related to the contribution from anti-quarks has been kept, despite being exponentially suppressed compared to that of quarks, so that the fermion matrix obeys

$$D^\dagger_{\rm HD}(U,\mu) = \gamma_5 D_{\rm HD}(U,-\mu^*)\gamma_5\,. \tag{1.71}$$

The fermion determinant in the HDQCD approximation then exhibits the sign problem for real chemical potential, as in QCD,

$$\det\left[D_{\rm HD}(U,\mu)\right]^* = \det D_{\rm HD}(U,-\mu^*)\,, \tag{1.72}$$

and is real when μ is purely imaginary.

The heavy dense approximation also allows for a closed expression for the expectation value of the quark number density in terms of the Polyakov loop,

$$\langle n\rangle = \frac{T}{V}\frac{\partial \ln Z}{\partial\mu} = \frac{1}{V}\sum_{\vec{x}}\langle n_{\vec{x}}\rangle\,, \tag{1.73}$$

with V the spatial volume and $n_{\vec{x}}$ is given by [47]

$$n_{\vec{x}} = 6N_f\left[\frac{zP_{\vec{x}} + 2z^2P^{-1}_{\vec{x}} + z^3}{1 + 3zP_{\vec{x}} + 3z^2P^{-1}_{\vec{x}} + z^3} - \frac{\overline{z}P^{-1}_{\vec{x}} + 2\overline{z}^2P_{\vec{x}} + \overline{z}^3}{1 + 3\overline{z}P^{-1}_{\vec{x}} + 3\overline{z}^2P_{\vec{x}} + \overline{z}^3}\right]\,, \tag{1.74}$$

where $z = he^{\mu/T}$ and $\overline{z} = he^{-\mu/T}$, and we have used the identity for SL$(3,\mathbb{C})$ matrices

$$\det(1 + cU) = 1 + c{\rm Tr}U + c^2{\rm Tr}U^{-1} + c^3\,. \tag{1.75}$$

Additionally, in the zero temperature limit, $N_\tau \to \infty$ at fixed lattice spacing, it is possible to determine analyticaly the critical chemical potential at which the quarks condense. The contribution from the anti-quarks to $n_{\vec{x}}$ is exponentially suppressed for

$\mu > 0$. For the quark contribution we can write

$$z = he^{\mu/T} = (2\kappa e^{\mu})^{N_\tau} \equiv \exp\left[\left(\mu - \mu_c^0\right) N_\tau\right] , \quad \mu_c^0 \equiv -\ln(2\kappa) , \tag{1.76}$$

where we have introduced the critical chemical potential at zero temperature, μ_c^0. It follows that, at zero temperature, $z \to 0$ when $\mu < \mu_c^0$ (the Silver Blaze region [74, 75]), and the density reaches its saturation value[1], $n_{\text{sat}} = 2 \times N_f \times N_c$ for $\mu > \mu_c^0$, irrespective of the value of the Polyakov loop. Thus μ_c^0 is the critical chemical potential at $T = 0$. Additionally, the expectation value of the Polyakov loop must vanish at zero temperature: in the $z \to 0$ sector the HDQCD fermion determinant becomes 1, which means the quarks have no influence and the theory reduces to pure Yang–Mills. On the other hand, above the critical chemical potential $z \to \infty$ for $T = 0$, and $\langle P \rangle = 0$ in order to keep the quark determinant finite.

The region beyond the critical chemical potential at zero temperature is a lattice artefact. At zero temperature, as soon as $\mu > \mu_c^0$ every lattice site is filled with the maximum number of quarks allowed by the Pauli exclusion principle, n_{sat}. For higher temperatures this transition is expected to be smoother, as thermal effects lower the threshold for the creation of deconfined quarks, for $\mu < \mu_c^0$.

The thermal transition at $\mu \approx 0$ can be understood in terms of that of pure Yang–Mills, due to the heavy quarks of this model. In this situation, the centre symmetry is approximate; a critical temperature, T_c, at which the quarks deconfine is expected. For temperatures above T_c, quarks are deconfined for all chemical potentials. A sketch of this behaviour for the HDQCD phase diagram is shown in figure 1.2.

It is important to stress that the heavy-dense approximation is based on the hopping parameter expansion, where the fermion determinant is written in powers of $\kappa \sim 1/ma$. In other words, the approximation assumes very heavy quarks, compared to the energy scale set by the lattice spacing. In turn, this implies that a continuum extrapolation is not possible, as the quark masses would have to become infinite as $a \to 0$, in order to keep κ fixed, leading to a quenched approximation of QCD with completely static quarks.

[1] The 2 comes from spin orientations (up and down), N_f is the number of flavours and N_c the number of colours.

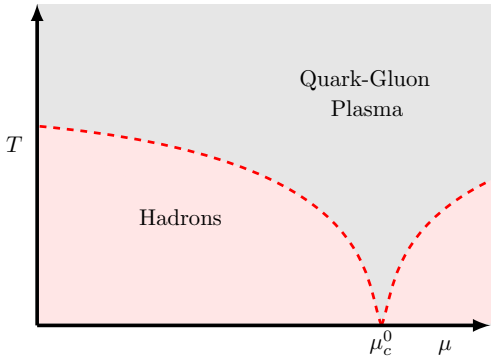

Figure 1.2: Sketch of the phase diagram of HDQCD. Pink areas represent regions in the parameter space where quarks are confined, whereas in the grey region quarks are free.

Chapter 2

Theoretical framework

Traditionally, Monte Carlo techniques have been used to simulate QFTs on the lattice at zero density, but the sign problem makes their results unreliable when one tries to apply it to finite density, especially for $\mu/T \gtrsim 1$. Therefore, it is instructive to briefly review these standard procedures, in order to understand where they differ from stochastic quantisation, which is the basis for complex Langevin. We, then, introduce the complex Langevin method as a possible way to perform lattice simulations of quantum field theories in Euclidean spacetime at finite chemical potential. Subsequently, stochastic quantisation applied to gauge fields will be discussed in more detail and with emphasis on a discrete evolution time, more suitable for computer simulations. We then review the complex Langevin method, including a discussion on its correctness and on methods to improve the technique—namely gauge cooling and dynamic stabilisation. This chapter concludes with a discussion on using the complex Langevin method to simulate dynamical quarks.

2.1 Monte Carlo formalism

Path integrals on the lattice typically have a very large number of degrees of freedom, making them practically impossible to calculate directly. A good example is a 2-state system on a lattice of N_s points (e.g., the Ising model). In this case, the number of states in the phase space would be 2^{N_s}, a number prohibitively large if one hopes to get close to the thermodynamical limit. In a quantum field theory the situation is even

more problematic as there are infinitely many values a field could have at each lattice point.

One way to estimate a path integral for a gauge theory is to randomly generate a sequence of configurations of gauge links $\{U_n\}_{n=1}^{N}$ (where spacetime and direction indices have been suppressed for clarity) such that the entire phase space may be visited, and average the observables over its elements. In other words,

$$\langle \mathcal{O} \rangle = \frac{1}{Z} \int \mathcal{D}U \, \mathcal{O}(U) \, e^{-S} \approx \frac{1}{N} \sum_{n=1}^{N} \mathcal{O}(U_n) \, e^{-S[U_n]} \,, \tag{2.1}$$

where $\mathcal{O}(U)$ represents any observable that is a function of the gauge links, and the partition function Z is given by

$$Z = \int \mathcal{D}U \, e^{-S} \,. \tag{2.2}$$

In the above equations $\mathcal{D}U$ is the path integration measure on the lattice, which can be expressed as

$$\mathcal{D}U = \prod_{x,\mu} dU_{x,\mu} \,, \tag{2.3}$$

with x indicating spacetime points, $\mu = 1, \ldots, 4$ are the Euclidean directions and $dU_{x,\mu}$ is the Haar measure of the group.

The fact that the action in Euclidean spacetime is positive semi-definite can be used to improve the generation of the U_n. If the U_n are generated with, e.g., an uniform probability distribution, many configurations will have their contributions exponentially suppressed by the weight e^{-S}, while configurations close to a minimum of S will have a larger contribution to the sum in eq. (2.1). It is useful, then, to produce field configurations using e^{-S}/Z as probability distribution and estimate $\langle \mathcal{O} \rangle$ simply as

$$\langle \mathcal{O} \rangle \approx \frac{1}{N} \sum_{n=1}^{N} \mathcal{O}(U_n) \,. \tag{2.4}$$

There is a variety of different Monte Carlo algorithms: Metropolis, Heat-Bath and Hybrid Monte Carlo, to name a few. The underlying principle in all of them is that described above. Their differences reside in how the sequence $\{U_n\}$ is generated. Typically an update procedure is applied to a field configuration to generate the next one,

in a Markov chain. Each algorithm has a different updating process, making some of them more efficient in certain situations. Further information can be found in [55, 56] and references therein.

2.2 Stochastic quantisation

Stochastic quantisation is an alternative formulation of quantum mechanics and quantum field theory [46] in Euclidean spacetime. Its workings are similar to stochastic processes like the Brownian motion, where the long time distribution for the velocity of the Brownian particle is $e^{-mv^2/2kT}$. The idea is that evolving the fields in an extra "time" dimension according to a Langevin equation can lead to an equilibrium distribution that is proportional to e^{-S}. Then, ensemble averages (i.e., path integrals) can be calculated as "time" averages over this extra dimension. The connection between the equilibrium distributions is given by the Fokker–Planck equation, to be discussed in this section.

Calling the extra temporal dimension θ—also known as Langevin time—the evolution of a gauge link $U_{x,\mu}(\theta)$ with dynamics given by the action S is

$$\frac{\partial U_{x,\mu}}{\partial \theta}\, U_{x,\mu}^{-1} = i\left[-\nabla_{x,\mu} S + \eta_{x,\mu}\right] . \tag{2.5}$$

Here $\nabla_{x,\mu} = \lambda^a \nabla^a_{x,\mu}$, where ∇^a is defined in appendix B and $\eta_{x,\mu}(\theta) = \lambda^a \eta^a_{x,\mu}(\theta)$ is a white noise field that satisfies

$$\left\langle \eta^a_{x,\mu}(\theta) \right\rangle_\eta = 0 \quad \text{and} \quad \left\langle \eta^a_{x,\mu}(\theta)\eta^b_{y,\nu}(\theta') \right\rangle_\eta = 2\,\delta^{ab}\delta_{xy}\delta_{\mu\nu}\delta(\theta-\theta') . \tag{2.6}$$

Here $\langle\cdots\rangle_\eta$ represents an average over the noise fields,

$$\langle\cdots\rangle_\eta = \frac{\int \mathcal{D}\eta\,(\cdots)\,\exp\left[-\frac{1}{4}\int d\theta\,\eta^2(\theta)\right]}{\int \mathcal{D}\eta\,\exp\left[-\frac{1}{4}\int d\theta\,\eta^2(\theta)\right]} . \tag{2.7}$$

Quantum expectation values are then calculated as

$$\langle \mathcal{O}(U)\rangle = \lim_{\theta\to\infty} \langle \mathcal{O}(U(\theta))\rangle_\eta . \tag{2.8}$$

Alternatively, the noise average in eq. (2.8) can be formulated in terms of a proba-

bility distribution $P(U,\theta)$,

$$\langle \mathcal{O}(U(\theta))\rangle_\eta = \int \mathcal{D}U\, \mathcal{O}(U)\, P(U,\theta)\,, \tag{2.9}$$

provided that

$$\lim_{\theta\to\infty} P(U,\theta) \equiv P_{eq}(U) = \frac{e^{-S}}{Z}\,. \tag{2.10}$$

There are many ways of showing that eq. (2.8) correctly reproduces the quantum correlation functions. The most common ones use eqs. (2.5) and (2.9) to derive a Fokker–Planck equation describing the Langevin time evolution of P (see, e.g., [76]), and show that eq. (2.10) holds. We shall focus on one that involves a discrete Langevin time, where it is incremented by finite steps, as it is the most natural one for computer simulations. Results for continuous fictitious time can be obtained by taking the limit of the Langevin step size going to zero.

2.2.1 Stochastic quantisation for SU(N) gauge links

Upon discretisation of the Langevin time, eq. (2.5) takes the form [76]

$$U(\theta+\varepsilon) = \exp\left[i\left(-\varepsilon\nabla S + \sqrt{\varepsilon}\,\eta\right)\right] U(\theta) \equiv \exp\left[iX(\varepsilon)\right] U(\theta)\,, \tag{2.11}$$

where ε is the Langevin step size and, from here on, the indices on U will be suppressed when no ambiguity arises. In order to define the probability distribution $P(U,\theta)$, introduced in eq. (2.9), we must define a delta function on the group manifold that works similarly to the usual one on real or complex space. The equivalent of a subtraction in group space is multiplication by the inverse, while the neutral element is the identity:

$$\begin{cases} x-y=0 \iff UV^{-1}=\mathbb{1} \\ \delta(x-y) \iff \delta(UV^{-1}) \end{cases}\,. \tag{2.12}$$

This way we have

$$\int dU\, \delta(UV^{-1}) f(U) = \int dU\, \delta(VU^{-1}) f(U) = f(V)\,. \tag{2.13}$$

Formally, the probability distribution $P(W,\theta)$ is connected to the noise averages via

$$P(W,\theta) = \left\langle \delta(U(\theta)W^{-1})\right\rangle_\eta\,. \tag{2.14}$$

The transition probability between two configurations W at time θ and V at $\theta + \varepsilon$ reads [77]

$$P(V, \theta + \varepsilon \,|W, \theta) = \left\langle \delta \left(e^{iX(\varepsilon)} W \, V^{-1} \right) \right\rangle_\eta \,, \tag{2.15}$$

where the delta function enforces the equation of motion, eq. (2.11), and the noise average guarantees that all realisations compatible with the initial and final states in the ensemble are taken into account. It is important to notice that, despite similar functional forms, equations (2.14) and (2.15) represent different objects: the former is the probability of the gauge link U having value W after time θ, while the latter is the conditional probability of the link being equal to V at time $\theta + \varepsilon$ given that it equated W at time θ. From the above equation one obtains

$$\begin{aligned} P(V, \theta + \varepsilon) &= \int dW \, P(V, \theta + \varepsilon \,|W, \theta) P(W, \theta) \\ &= \left\langle \delta \left(e^{iX} U(\theta) V^{-1} \right) \right\rangle_\eta \,, \end{aligned} \tag{2.16}$$

as expected from the equation of motion.

Following [78] we expand the delta function of eq. (2.15) as a power series in $\overline{\varepsilon} = \sqrt{\varepsilon}$

$$\begin{aligned} \delta \left(e^{iX} U V^{-1} \right) = \delta \left(U V^{-1} \right) \\ + \overline{\varepsilon} \left[X'^a \, \nabla^a \right]_{\overline{\varepsilon}=0} \delta \left(U V^{-1} \right) \\ + \overline{\varepsilon}^2 / 2 \left[X'^a X'^b \, \nabla^a \nabla^b + X''^a \, \nabla^a \right]_{\overline{\varepsilon}=0} \delta \left(U V^{-1} \right) \\ + O(\overline{\varepsilon}^3) \,, \end{aligned} \tag{2.17}$$

with $X'^a = dX^a / d\overline{\varepsilon}$, X^a representing the components of the Langevin drift, X, in the Lie algebra. We then use the transition probability to write $P(V, \theta + \varepsilon)$ as

$$P(V, \theta + \varepsilon) = P(V, \theta) + \overline{\varepsilon} \, P_1(V, \theta) + \overline{\varepsilon}^2 / 2 \, P_2(V, \theta) + O(\overline{\varepsilon}^3) \,, \tag{2.18}$$

where, after integrating by parts,

$$P_1(V, \theta) = -\nabla^a \left\langle \overline{X}'^a P(V, \theta) \right\rangle_\eta \,, \tag{2.19}$$

$$P_2(V, \theta) = \nabla^a \nabla^b \left\langle \overline{X}'^a \overline{X}'^b P(V, \theta) \right\rangle_\eta - \nabla^a \left\langle \overline{X}''^a P(V, \theta) \right\rangle_\eta \,, \tag{2.20}$$

with, c.f. appendix C,

$$\overline{X}^{(n)} = \left.\frac{\partial^n X}{\partial \bar{\varepsilon}^n}\right|_{\bar{\varepsilon}=0} . \tag{2.21}$$

Using the definition of $X(\varepsilon)$ from eq. (2.11) and the properties of the white noise η, we can see that

$$P_1(V,\theta) = -\nabla^a \left\langle \eta^a P(V,\theta) \right\rangle_\eta = -\left\langle \eta^a \right\rangle_\eta \nabla^a P = 0 \,, \tag{2.22}$$

$$\begin{aligned} P_2(V,\theta) &= \nabla^a\nabla^b \left\langle \eta^a\eta^b P(V,\theta) \right\rangle_\eta + \nabla^a \left\langle 2\nabla^a S\, P(V,\theta) \right\rangle_\eta \\ &= 2\,\nabla^a \left[(\nabla^a S + \nabla^a) P(V,\theta) \right] . \end{aligned} \tag{2.23}$$

Combining the above with eq. (2.18) we arrive at the (discrete time) Fokker–Planck equation for P

$$\frac{P(V,\theta+\varepsilon) - P(V,\theta)}{\varepsilon} = \nabla^a \left[(\nabla^a + \nabla^a S) P(V,\theta) \right] + O(\varepsilon) \,. \tag{2.24}$$

It is easy to see that $P_{eq}(V) = Z^{-1} e^{-S[V]}$ is an equilibrium solution of the Fokker–Planck equation. Using the transformation

$$P(V,\theta) = \psi(V,\theta)\, e^{-\frac{1}{2}S} \,, \tag{2.25}$$

one can write the Fokker–Planck equation in a Schrödinger-like form

$$\frac{\psi(V,\theta+\varepsilon) - \psi(V,\theta)}{\varepsilon} = -H_{\mathrm{FP}}\, \psi(V,\theta) \,, \tag{2.26}$$

where the operator

$$H_{\mathrm{FP}} = \left(-\nabla^a + \frac{1}{2}\nabla^a S \right) \left(\nabla^a + \frac{1}{2}\nabla^a S \right) \tag{2.27}$$

is known as the Fokker–Planck Hamiltonian. It is self-adjoint for real actions, and a discrete and positive spectrum is strongly suggested by numerical evidence, depending on the model and parameter values, in many interesting models [47, 69].

Assuming the discrete spectrum mentioned above, a generic solution ψ can be expanded as

$$\psi(V,\theta) = \sum_{n=0}^{\infty} a_n \psi_n(V) \exp\left(-E_n\, \theta\right) , \tag{2.28}$$

with E_n being the eigenvalues of H_{FP} and ψ_n the respective eigenfunctions. From the factorisation of H_{FP} it is clear that the eigenfunction of zero energy is

$$\psi_0(V) = a_0 \exp\left(-\frac{1}{2}S\right) . \tag{2.29}$$

Since $E_n > 0$ for $n > 1$, it follows that the probability distribution for $\theta \to \infty$ is indeed given by $Z^{-1}e^{-S}$. This shows the equivalence between stochastic quantisation and conventional path integrals for large Langevin time.

2.2.2 Finite Langevin step size corrections

For a finite Langevin step size the equilibrium distribution exhibits an ε-dependence: $P \sim e^{-S[V]-\varepsilon S_2[V]}$. The terms that give rise to S_2 are obtained by continuing the Taylor expansion of eq. (2.17) up to order $\overline{\varepsilon}^4 = \varepsilon^2$, and modifying the expansion for $P(V, \theta+\varepsilon)$ accordingly, by adding the coefficients of $\overline{\varepsilon}^3$ and $\overline{\varepsilon}^4$,

$$P_3(V,\theta) = \left\langle \eta^a\eta^b\eta^c \, \nabla^a\nabla^b\nabla^c P - 3\left(\eta^a\nabla^b S + \eta^b\nabla^a S\right)\nabla^a\nabla^b P \right\rangle_\eta , \tag{2.30}$$

$$\begin{aligned} P_4(V,\theta) = {} & 12\left\langle \nabla^a S \nabla^b S \, \nabla^a\nabla^b P \right\rangle_\eta \\ & - 4\left\langle \left(\eta^a\eta^b \, \nabla^c S + \eta^b\eta^c \, \nabla^a S + \eta^a\eta^c \, \nabla^b S\right)\nabla^a\nabla^b\nabla^c P \right\rangle_\eta \\ & + \left\langle \eta^a\eta^b\eta^c\eta^d \, \nabla^a\nabla^b\nabla^c\nabla^d P \right\rangle_\eta , \end{aligned} \tag{2.31}$$

respectively. Details of this expansion can be found in appendix C.

Due to the odd numbers of noise fields, P_3 will average out to zero. For P_4 we repeatedly use the ansatz proposed in [78], $\nabla^a P = -\left(\nabla^a S\right) P + O(\overline{\varepsilon}^2)$, except for the last derivative:

$$P_4(V,\theta) = 12\,\nabla^a\left[\left(\nabla^2\nabla^a S - \nabla^b\nabla^a S \, \nabla^b S + \frac{\alpha^2}{6} C_A \nabla^a S\right) P\right] , \tag{2.32}$$

where $C_A = N$ is the Casimir invariant for SU(N) in the adjoint representation and α is related to the normalisation of the group generators (c.f. appendix B) ,

$$\mathrm{Tr}\left[\lambda^a\lambda^b\right] = \frac{\alpha^2}{2}\delta^{ab} . \tag{2.33}$$

The discrete-time Fokker–Planck equation including corrections up to ε^2 then reads

$$\frac{P(V,\theta+\varepsilon)-P(V,\theta)}{\varepsilon} = \nabla^a\left[\left(\nabla^a + \nabla^a S + \varepsilon\nabla^a S_2\right)P(V,\theta)\right] + O(\varepsilon^2)\,, \tag{2.34}$$

with

$$S_2 = \frac{1}{2}\nabla^2 S - \frac{1}{4}\nabla^a S\,\nabla^a S + \frac{\alpha^2}{12}C_A\,S\,. \tag{2.35}$$

The equilibrium distribution with order ε corrections is

$$P_{eq}(V) = \frac{1}{Z}\exp\left[-S-\varepsilon S_2\right]\,. \tag{2.36}$$

2.2.3 Second order algorithms

In the numerical analysis of ordinary differential equations, Runge–Kutta methods are used to decrease the dependence on the discretisation step size, at the expense of more complicated algorithms. This is done by generating "intermediate" solutions between consecutive steps and using those to cancel out the dependence on the lowest powers of the step size. Here we show how a similar procedure can be applied to stochastic differential equations on a gauge group. This was first done in [77, 79]

It is necessary for the second order Langevin equation to have a Boltzmann-like equilibrium probability distribution. Therefore, the new drift, $\mathcal{X}$, must have a similar form to X, as defined in eq. (2.11). Additionally, information from the first order solution must be used in order to cancel S_2.

We use a "generalised" first-order drift [80],

$$X^a = -\overline{\varepsilon}^2 k_1 \nabla^a S[U] + \overline{\varepsilon}\,k_2\,\eta^a\,, \tag{2.37}$$

where k_1 and k_2 will be set later, to generate intermediate gauge links $U' = e^{iX}U$. These links are used in the second-order drift $\mathcal{X}^a$, defined as

$$\begin{aligned}\mathcal{X}^a &= -\overline{\varepsilon}^2 k_3 \nabla^a S[U] - \overline{\varepsilon}^2 k_4 \nabla^a S[U'] - \overline{\varepsilon}^4 k_5 \frac{\alpha^2}{2} C_A \nabla^a S[U] + \overline{\varepsilon}\,k_6\,\eta^a \qquad (2.38)\\ &= -\overline{\varepsilon}^2(k_3+k_4)\nabla^a S[U] - \overline{\varepsilon}^3 k_2 k_4\,\eta^b\nabla^b\nabla^a S[U]\\ &\quad - \overline{\varepsilon}^4\left(\frac{1}{2}\,k_2^2 k_4\,\eta^b\eta^c\nabla^b\nabla^c\nabla^a S[U] - k_1 k_4 \nabla^b S[U]\nabla^c\nabla^a S[U]\right.\\ &\quad \left. + k_5 C_A \frac{\alpha^2}{2}\nabla^a S[U]\right) + \overline{\varepsilon}\,k_6\,\eta^a\,, \qquad (2.39)\end{aligned}$$

where $\nabla^a S[U']$ has been expanded to order $\bar{\varepsilon}^2$. Ignoring terms of $O(\bar{\varepsilon}^3)$ and following the same steps that lead to eq. (2.24), the first order Fokker–Planck equation now reads

$$\frac{P(V,\theta+\varepsilon)-P(V,\theta)}{\varepsilon} = \nabla^a \left\{ \left[(k_3+k_4)\,\nabla^a + k_6^2 \nabla^a S \right] P(V,\theta) \right\} + O(\varepsilon)\,. \tag{2.40}$$

In this form, it is clear that for an equilibrium distribution of the form e^{-S} it is necessary to have $k_3 + k_4 = 1$ and $k_6^2 = 1$.

The expansion of $P(V,\theta+\varepsilon)$, again up to order $\bar{\varepsilon}^4$, and with $\mathcal{X}$ for the drift, leads once more to $P_3 = 0$ while P_4 is modified to

$$P_4^{\mathcal{X}} = P_4 - \nabla^a \left\langle \overline{\mathcal{X}}^{(4)a} P \right\rangle_\eta + 2\nabla^a \nabla^b \left\langle \overline{\mathcal{X}}'^a \overline{\mathcal{X}}^{(3)b} + \overline{\mathcal{X}}^{(3)a} \overline{\mathcal{X}}'^b P \right\rangle_\eta . \tag{2.41}$$

Following the previously used procedure, we insert the derivatives of the drift into $P_4^{\mathcal{X}}$, and use the ansatz $\nabla^a P = -\left(\nabla^a S\right) P$, except for the last derivative of P[1]

$$\begin{aligned} P_4^{\mathcal{X}} = 4!\,\nabla^a \Bigg\{ \Bigg[& \left(k_2^2 k_4 - 2k_2 k_4 k_6 + \frac{1}{2} \right) \nabla^2 \nabla^a S \\ & + \left(2k_2 k_4 k_6 - k_1 k_4 - \frac{1}{2} \right) \nabla^b S \nabla^b \nabla^a S \\ & + \left(k_5 - k_2 k_4 k_6 + \frac{1}{6} \right) \frac{\alpha^2}{2} C_A \nabla^a S \Bigg] P \Bigg\} . \end{aligned} \tag{2.42}$$

Therefore, in order to eliminate all $\bar{\varepsilon}^4 = \varepsilon^2$ dependence from the expansion of $P(V,\theta+\varepsilon)$, it is required that

$$k_2 = \frac{2k_4 k_6 \pm \sqrt{2k_4\,(2k_4-1)}}{2k_4}\,, \tag{2.43}$$

$$k_1 = \frac{4k_2 k_4 k_6 - 1}{2k_4}\,, \tag{2.44}$$

$$k_5 = k_2 k_4 k_6 - \frac{1}{6}\,. \tag{2.45}$$

[1] The terms coming from P_4 are not changed, as they depend only on $k_3 + k_4$ or k_6^2.

Two possible choices for the k's are [79]

$$k_3 = k_4 = \frac{1}{2}\,, \tag{2.46}$$

$$k_1 = k_2 = k_6 = 1\,, \tag{2.47}$$

$$k_5 = \frac{1}{3}\,, \tag{2.48}$$

in which the intermediate and final evolution steps resemble Heun's method for solving ODE's, and [80]

$$k_4 = k_6 = 1 \tag{2.49}$$

$$k_3 = 0 \tag{2.50}$$

$$k_1 = \frac{3 \pm 2\sqrt{2}}{2} \tag{2.51}$$

$$k_2 = \frac{2 \pm \sqrt{2}}{2} \tag{2.52}$$

$$k_5 = \frac{10 \pm 6\sqrt{2}}{12} \tag{2.53}$$

where X and $\mathcal{X}$ have similar forms, with $\mathcal{X}$ depending only on $\nabla S[U']$. From a computational point of view, this second method is preferred as it reduces the memory needs, since $\nabla S[U]$ does not need to be stored for the second step.

2.3 Complex Langevin method

For theories with a complex action, such as QCD at finite chemical potential, the force on the right-hand side of (2.5) will cause the gauge links to become non-unitary. Allowing the physical fields to explore a larger (complexified) phase space is central to the complex Langevin method, proposed in the early 1980s by Klauder [81] and Parisi [82]. Because field configurations generated by stochastic processes do not rely on e^{-S} as a probability distribution, but have it as a consequence of the evolution over Langevin time, the complex Langevin method is a very good candidate for simulating theories with complex actions, in particular those that exhibit the sign problem.

After a certain initial popularity [76, 83], some problems were found with the complex Langevin method. Those can be categorised in two classes: runaway trajectories, where the simulation would not converge, and convergence to a wrong limit [84, 85].

These difficulties have recently begun to be overcome, allowing the method to be successfully revived [47, 68, 69, 86–90]. Runaway trajectories have been eliminated with the use of adaptive Langevin step size [48], while the convergence to wrong limits is subject of active research and methods to deal with it will be reviewed in the following subsections.

In essence, this method requires adding an imaginary part to each component of the fields and re-writing the action and observables in terms of them. For instance, in a 0-dimensional model one has $x \to z = x + iy$ as the complexified field, whose components evolve as

$$\begin{aligned} \dot{x} &= -\mathrm{Re}\nabla S(x+iy) + \sqrt{N_R}\eta_R\,, \\ \dot{y} &= -\mathrm{Im}\nabla S(x+iy) + \sqrt{N_I}\eta_I\,, \end{aligned} \tag{2.54}$$

with η_R and η_I being the real and imaginary noise fields, respectively, and $N_R - N_I = 1$, with $N_I \geq 0$. For example, an observable x^2 becomes $z^2 = x^2 - y^2 + 2ixy$. In the language of gauge fields on the lattice, this complexification amounts to allowing for non-unitary gauge links, i.e. $U^\dagger \neq U^{-1}$.

A key observation is that the new phase space has twice the dimension of the original one. Because of the extra degrees of freedom, the sign problem can be circumvented, as the path integration is deformed in the enlarged space. However, compact manifolds such as U(1) and SU(N) are no longer compact after complexification, allowing for potential runaway trajectories. Additionally, the Langevin evolution of the probability distribution and the observables might cause them to break the proof of convergence, to be explained in the next subsection, culminating in convergence to wrong limits.

2.3.1 Proof of convergence

The usual proofs of convergence to the probability weight e^{-S} are problematic for complex actions. We shall discuss here the necessary criteria for those fields' Langevin evolution to converge to the correct distributions. More details can be found in [49], while an analysis focussing on issues related to finite Langevin step sizes is done at [91].

When simulating theories with complex actions, one is faced with a complex distribution $\rho = e^{-S}$, with S a holomorphic function on a real manifold $\mathscr{M}$. The complex Langevin approach concerns the evolution of real and imaginary parts of complexified

fields, defined on the complexification $\mathscr{M}_c$ of $\mathscr{M}$. The idea is to simulate the evolution of *entire holomorphic* observables, $\mathcal{O}$, on $\mathscr{M}_c$, such that their expectation values converge to those with complex distribution ρ,

$$\langle \mathcal{O} \rangle = \int_{\mathscr{M}_c} dU dU^\dagger P(U, U^\dagger) \mathcal{O}(U) = \int_{\mathscr{M}} dV \rho(V) \mathcal{O}(V) \,. \tag{2.55}$$

The Langevin evolution of a complex density $\rho(V, \theta)$, with $V \in \mathrm{SU}(N)$ and $S(V)$ a holomorphic action, follows the same steps outlined above for real PDFs, whose Taylor expansion can be written generically as

$$\rho(V, \theta + \varepsilon) = \sum_{n=0}^{\infty} \frac{\varepsilon^n}{(2n)!} L^T_{\mathbb{1},2n} \left[\rho(V, \theta)\right] \,, \tag{2.56}$$

with $L^T_{\mathbb{1},2n}$ being an operator that generates the terms from the expansion of P shown in section 2.2. A slight generalisation of this operator, $L^T_{P_0,2n}$, involves replacing V as argument in the action by $P_0 V$, where P_0 is a positive semi-definite Hermitian matrix. This change is akin to adding a constant imaginary part to a real number: $x \to x + iy_0$, and leads to a stationary solution $\rho_{P_0} \sim \exp\left[-S(P_0 V)\right]$.

One must be careful when considering these processes on non-unitary manifolds, such as $\mathrm{SL}(N, \mathbb{C})$, since U and $U^\dagger$ are independent group elements. A suitable redefinition of the gauge group derivative is found in appendix C. It suffices here to say that terms of the form $\overline{X}^{(n)a} \nabla^a$, in Taylor expansions in powers of the Langevin step size, are replaced by $(\overline{X}^{(n)a}_R \nabla^a_R + \overline{X}^{(n)a}_I \nabla^a_I)$, where X_R and X_I are the real and imaginary parts, respectively, of the complex Langevin drift

$$X^a_{\mathbb{C}}(\bar{\varepsilon}) = \left(\sqrt{N_R}\eta^a_R + \sqrt{N_I}\eta^a_I\right) \bar{\varepsilon} - \nabla^a_R S(U) \bar{\varepsilon}^2 \,, \tag{2.57}$$

with the assumption of a holomorphic action.

Analogous to the Langevin evolution over $\mathrm{SU}(N)$, evolving a PDF $P(U, U^\dagger; \theta)$ with $U, U^\dagger \in \mathrm{SL}(N, \mathbb{C})$ results in a Taylor expansion similar to eq. (2.56)

$$P(U, U^\dagger; \theta + \varepsilon) = \sum_{n=0}^{\infty} \frac{\varepsilon^n}{(2n)!} \tilde{L}^T_{2n} \left[P(U, U^\dagger; \theta)\right] \,, \tag{2.58}$$

where $\tilde{L}^T_{2n}$ also generates the terms shown in section 2.2, using the redefined derivative and eq. (2.57).

The holomorphic observables, $\mathcal{O}(U,\theta)$, can be evolved using the operators obtained from the above equation as

$$\begin{aligned}\langle \mathcal{O}(U,\theta+\varepsilon)\rangle_\eta &= \int \mathcal{D}U\mathcal{D}U^\dagger\, \mathcal{O}(U)\, P(U,U^\dagger;\theta+\varepsilon) \\ &= \sum_{n=0}^{\infty} \frac{\varepsilon^n}{(2n)!} \int \mathcal{D}U\mathcal{D}U^\dagger\, \mathcal{O}(U)\, \tilde{L}_{2n}^T \Big[P(U,U^\dagger;\theta)\Big] \\ &= \sum_{n=0}^{\infty} \frac{\varepsilon^n}{(2n)!} \int \mathcal{D}U\mathcal{D}U^\dagger\, P(U,U^\dagger;\theta)\, \tilde{L}_{2n}\left[\mathcal{O}(U)\right] \\ &= \left\langle \sum_{n=0}^{\infty} \frac{\varepsilon^n}{(2n)!} \tilde{L}_{2n}\left[\mathcal{O}(U)\right] \right\rangle_\eta \equiv \left\langle \tilde{\Theta}_\varepsilon\left[\mathcal{O}(U,\theta)\right]\right\rangle_\eta , \end{aligned} \tag{2.59}$$

where $\tilde{\Theta}$ has been introduced as a shorthand notation for the power series in ε and $\tilde{L}_{2n}$ represents the formal adjoint of $\tilde{L}_{2n}^T$ with respect to the above integral. It is important to notice that $\mathcal{O}(U,\theta+\varepsilon)$ will be holomorphic if $\mathcal{O}(U,\theta)$ is. This can be seen from the fact that, by hypothesis $\overline{\partial}_{ij}\mathcal{O}(U,\theta)=0$, where $\overline{\partial}_{ij}$ represents the derivative with respect to the (i,j)-th component of $U^\dagger$. Since the operators $\tilde{L}_{2n}$ only contain ∂_{ij}, $\overline{\partial}_{ij}$ and S, which is assumed to be holomorphic, we have that

$$\overline{\partial}_{ij}\mathcal{O}(U,\theta+\varepsilon) = \overline{\partial}_{ij}\tilde{\Theta}_\varepsilon\left[\mathcal{O}(U,\theta)\right] = \tilde{\Theta}_\varepsilon\left[\overline{\partial}_{ij}\mathcal{O}(U,\theta)\right] = 0\,. \tag{2.60}$$

We now proceed to show that

$$\langle \mathcal{O}\rangle_{P(\theta)} = \langle \mathcal{O}\rangle_{\rho(\theta)}\,, \tag{2.61}$$

where $\langle\cdot\rangle_{P(\theta)}$ and $\langle\cdot\rangle_{\rho(\theta)}$ represent expectation values with respect to P and ρ, respectively, provided that the initial conditions agree,

$$\langle \mathcal{O}\rangle_{P(0)} = \langle \mathcal{O}\rangle_{\rho(0)}\,. \tag{2.62}$$

This is achieved by choosing

$$P(U,U^\dagger;0) = \rho(P_0V,0)\delta(UU^\dagger(P_0^2)^{-1})\,, \tag{2.63}$$

with $P_0V=U$, V being a unitary matrix and P_0 a constant, self-adjoint and positive semi-definite matrix. The delta function enforces $UU^\dagger = P_0^2$, or $U^\dagger = U^{-1}P_0^2$. Start-

ing from the expectation value $\langle \mathcal{O} \rangle_{P(\theta)}$ of a holomorphic observable $\mathcal{O}$ and using the properties of the time evolution operators we see that

$$\begin{aligned}
\langle \mathcal{O} \rangle_{P(\theta)} &= \int \mathcal{D}U\mathcal{D}U^\dagger P(U,U^\dagger;\theta)\,\mathcal{O}(U) \\
&= \int \mathcal{D}U\mathcal{D}U^\dagger \tilde{\Theta}_\theta^T \left[P(U,U^\dagger;0)\right] \mathcal{O}(U) \\
&= \int \mathcal{D}U \rho(U,0) \Theta_\theta \left[\mathcal{O}(U)\right] \\
&= \int \mathcal{D}U \Theta_\theta^T \left[\rho(U,0)\right] \mathcal{O}(U) \\
&= \int \mathcal{D}U \rho(U,\theta) \mathcal{O}(U) = \langle \mathcal{O} \rangle_{\rho(\theta)} \,,
\end{aligned} \tag{2.64}$$

where we have introduced $\Theta_\theta[\cdot]$ as a shorthand for $\Theta_\varepsilon[\Theta_\varepsilon[\cdots]]$ with $n = \theta/\varepsilon$ operators chained, and used the fact that $\mathcal{O}$ is holomorphic on going from $\tilde{\Theta}[\mathcal{O}]$ to $\Theta[\mathcal{O}]$, since $\nabla_I^a \mathcal{O} = i\nabla_R^a \mathcal{O}$ and terms of the form $(\overline{X}_R^{(n)a} \nabla_R^a + \overline{X}_I^{(n)a} \nabla_I^a)$ reduce to $\overline{X}_{\mathbb{C}}^{(n)a} \nabla_R^a$. In doing the above manipulations it has been assumed that boundary terms do not arise.

It is expected that if the distribution of the observables, $\mathcal{O}(U,U^\dagger)P(U,U^\dagger,\theta)$, decays fast enough in the non-compact directions, the Langevin process should converge to the right result [92–94]. So far we have assumed that the Langevin drift X^a is a regular function. If, however, the derivative of the action, $\nabla^a S$, has poles, it is necessary to study in more detail how they influence the dynamics and the correctness of the results [95–97].

2.3.2 Gauge cooling

The addition of new degrees of freedom via the complexification of the original manifold allows for the circumvention of the sign problem, as already discussed, but the non-compactness of groups like $\mathrm{SL}(N,\mathbb{C})$ poses extra problems, as the Langevin process could run away to infinity by following an unstable classical trajectory. In some situations, the noise can knock the system off such trajectories; in others, the adaptive step size can help in avoiding them. However, it is not always possible to keep the system away from them. Therefore, it is instructive to consider the gauge links' average distance from $\mathrm{SU}(N)$, as a function of the Langevin time.

Given that $\mathrm{SL}(N,\mathbb{C})$ matrices can be written, as used in the previous section, according to the polar decomposition as $U = PV$, with P being Hermitian and positive

semi-definite and V unitary, one may consider the *unitarity norm* [98, 99]

$$d = \frac{1}{N}\mathrm{Tr}\left[UU^\dagger - \mathbb{1}\right] \geq 0\,. \tag{2.65}$$

This distance is only null if U is unitary and is otherwise strictly positive[1]. Similar measures of this distance can be defined, but this is the simplest one and our choice for this work. It is worth noting that d is not invariant under general $\mathrm{SL}(N,\mathbb{C})$ gauge transformations

$$U_{x,\mu} \to \Omega_x U_{x,\mu}\Omega^{-1}_{x+\mu}\,, \quad U^\dagger_{x,\mu} \to \left(\Omega^{-1}_{x+\mu}\right)^\dagger U^\dagger_{x,\mu}\Omega^\dagger_x\,, \tag{2.66}$$

but is unchanged if the Ω_x are unitary.

Deep explorations of the non-compact directions of the complexified manifolds of gauge theories can lead to wide probability distributions that violate the criteria for correctness mentioned above, as has been shown in previous studies [88, 98, 99]. The $\mathrm{SU}(N)$ gauge invariance of the unitarity norm can be used to reduce these explorations without changing the physics, in a procedure known as *gauge cooling* [98].

Taking $\Omega_x = \exp\left[\varepsilon\alpha\lambda^a f^a_x\right]$ as the transformation matrix, with ε the Langevin time, α a positive control parameter and $f^a_x \in \mathbb{C}$, we consider the changes in d to first order in ε

$$\delta d \approx 2\varepsilon\alpha f^a_x \mathrm{Tr}\left[\lambda^a\left(U_{x\mu}U^\dagger_{x\mu} - U^\dagger_{x-\mu,\mu}U_{x-\mu,\mu}\right)\right]\,. \tag{2.67}$$

Since we want d to decrease after the transformation, choosing

$$f^a_x = -\mathrm{Tr}\left[\lambda^a\left(U_{x\mu}U^\dagger_{x\mu} - U^\dagger_{x-\mu,\mu}U_{x-\mu,\mu}\right)\right] \tag{2.68}$$

ensures that $\delta d \leq 0$. This choice also guarantees that the f^a_x are real numbers, and therefore that Ω_x acts orthogonally to $\mathrm{SU}(N)$, and directed toward it. Gauge fixing on the lattice works similarly, with f^a_x as imaginary numbers and the "distance" function $d = \mathrm{Tr}\left[U + U^\dagger\right]$, for the Landau gauge.

[1] In the case of U(1) gauge theories one must also include the inverted links as

$$d = UU^\dagger - U^{-1}U^{-1\dagger}$$

In simulations, gauge cooling steps are applied between consecutive Langevin evolution steps. Any potential non-unitarity of the gauge links, originating from the dynamics, is reduced in a gauge invariant way. The links are transformed into those equivalent to them that are closer to the unitary manifold in a steepest descent fashion. Many studies of this technique have been undertaken in different models, such as one $\mathrm{SL}(N,\mathbb{C})$ link and SU(2) Polyakov chain models [99], where an adaptive control parameter has also been employed, $0+1$ and $1+1$ dimensional QCD [100], and also different norms have been investigated [101]. Other variations of the method, where different norms have been studied for, e.g., random matrix theories, can be found in [102, 103].

2.3.3 Dynamic stabilisation

The aforementioned studies indicate that gauge cooling seems to be a necessary condition to avoid convergence to wrong limits, but it is not sufficient, as has been pointed out in [50, 99]. In particular, it is known that gauge cooling becomes less effective on coarser lattices—i.e., for smaller values of the gauge coupling β. This limits the applicability of complex Langevin on lower temperatures, which would require potentially very large lattices.

Since the simulations cannot be trusted when the unitarity norm is large (i.e., of $O(1)$) one is faced with the need to remove all data points after the norm becomes large from the analysis. This can be problematic, as in certain cases the unitarity norm grows very quickly, leaving only a small statistical sample available.

We have proposed a new method, called dynamic stabilisation [52, 104], which modifies the Langevin drift, X^a, for each lattice site and spacetime direction by adding a $\mathrm{SU}(N)$—but not $\mathrm{SL}(N,\mathbb{C})$—gauge invariant term that is trivial in the continuum limit, and grows with the distance from the unitary manifold. This lattice artefact has the effect of a restoring force pointing toward the $\mathrm{SU}(N)$ submanifold, and is not holomorphic, since it must depend on $UU^\dagger$. As such, it cannot be derived from an action.

Applying this method to the simple 0-dimensional model of eq. (2.54) modifies the drift of the imaginary component y by a force of the form $-\alpha y^n$, such that

$$\begin{aligned} \dot{x} &= -\mathrm{Re}\nabla S(x+iy) + \eta\,, \\ \dot{y} &= -\mathrm{Im}\nabla S(x+iy) - \alpha y^n\,, \end{aligned} \tag{2.69}$$

since y is the equivalent of the unitarity norm for this model and we have specialised for real noises. In this case, large excursions into the imaginary direction can be controlled by changing α and n, while still allowing for a certain exploration of the complexified phase space.

One possible implementation of dynamic stabilisation in gauge theories is given by

$$X^a_{x,\mu} \to X^a_{x,\mu} + i\varepsilon\alpha_{\mathrm{DS}} M^a_x \,, \tag{2.70}$$

with α_{DS} a real control coefficient and

$$M^a_x = i b^a_x \left(\sum_c b^c_x b^c_x \right)^3 \,, \quad b^a_x = \sum_\nu \mathrm{Tr}\left[\lambda^a U_{x,\nu} U^\dagger_{x,\nu} \right] \,. \tag{2.71}$$

This choice of the dynamic stabilisation drift straightforwardly implements the criteria outlined above: M^a_x depends only on the non-unitary parts of the gauge links, and the power of 3 ensures scaling with a high power of the lattice spacing. Moreover, since M^a_x is purely imaginary, it can only affect the non-unitary parts of the gauge links.

To check for the triviality of M^a_x in the continuum limit, we use the definition of the gauge links $U_{x,\nu}$ in terms of continuum gauge fields,

$$U_{x,\nu} = \exp\left[i a \lambda^c \left(A^c_{x,\nu} + i B^c_{x,\nu} \right) \right] \,, \tag{2.72}$$

where the complexified field has been split as $A + iB$, and a is the lattice spacing. An expansion in powers of a shows that

$$U_{x,\nu} U^\dagger_{x,\nu} = \mathbb{1} - 2a\lambda^c B^c_{x,\nu} + a^2 \left(2\lambda^c \lambda^d B^c_{x,\nu} B^d_{x,\nu} + i B^c_{x,\nu} A^d_{x,\nu} \left[\lambda^c, \lambda^d \right] \right) + O(a^3) \,, \tag{2.73}$$

which leads to

$$b^a_x \sim -a \sum_\nu B^a_{x,\nu} \equiv -a \overline{B}^a_x \,, \tag{2.74}$$

and

$$M^a_x \sim -i a^7 \left(\sum_c \overline{B}^c_x \overline{B}^c_x \right)^3 \overline{B}^a_x \,. \tag{2.75}$$

It is clear that M^a_x should become irrelevant compared to, for example, $F^2_{\mu\nu}$ in the continuum limit. With this modification to the drift, the Fokker–Planck equation

becomes, neglecting $O(\varepsilon)$ effects,

$$\frac{P(U,\theta+\varepsilon)-P(U,\theta)}{\varepsilon} = \nabla^a \left[(\nabla^a + \nabla^a_R S + i\alpha_{\mathrm{DS}} M^a) P(U,\theta)\right] . \tag{2.76}$$

We point out that while our choice for M^a_x was isotropic, this is not necessary. Other choices can be proposed, but we do not have any method to know their efficacy *a priori*. Furthermore, it is worth emphasizing this procedure is different from gauge cooling, where gauge transformations are used between Langevin steps, with the drift unchanged. Dynamic stabilisation is a direct modification of the Langevin drift by a non-holomorphic term. This cannot be derived from an action, and as such the proof of convergence outlined in section 2.3.1 is not directly applicable. It is not yet known how it can be extended to include such terms. However, promising results can be found in the references given and in section 3.4, where studies of the dependence of both the Langevin drift and of observables on α_{DS} are reviewed, together with comparisons with results from the literature.

2.4 Simulating dynamical quarks

The simulation of dynamical fermions using Langevin dynamics is not as straightfoward as that of gluons. The effective quark action—basically the determinant of the Dirac operator, obtained after integrating out the fermionic fields—is highly non-local. The Langevin method requires calculating a derivative of the action with respect to the gauge links, which, in turn, requires inverting the Dirac matrix. We will briefly discuss here how the necessary steps can be carried out and potential issues related to them. Our approach follows that of [105].

The effective fermion action at finite chemical potential μ generically reads (c.f. section 1.2.4)

$$S_{\mathrm{F}} = -N_f \ln\det D(U,\mu) , \tag{2.77}$$

where N_f is the number of quark flavours, assuming degenerate masses, and U are the gauge links. Using the identity $\ln\det M = \mathrm{Tr}\ln M$, we can rewrite the drift term originating from S_{F} as

$$\nabla^a S_{\mathrm{F}} = -N_f \nabla^a \mathrm{Tr}\ln D(U,\mu) = -N_f \mathrm{Tr}\left[D^{-1}(U,\mu)\nabla^a D(U,\mu)\right] . \tag{2.78}$$

Given that the size of D scales as the number of lattice sites squared, exact evaluations of its inverse are possible only in very few cases.

One cost-effective alternative, proposed in [77, 106], is known as the bilinear noise scheme. This method, which is related to pseudofermionic variables, involves first writing the trace in eq. (2.78) as an average over vectors η^i of Gaussian random numbers

$$\mathrm{Tr}\left[D^{-1}(U,\mu)\nabla^a D(U,\mu)\right] \approx \frac{1}{N_{\mathrm{vec}}}\sum_{i=1}^{N_{\mathrm{vec}}} \eta^{i\dagger} D^{-1}(U,\mu)\nabla^a D(U,\mu)\eta^i\,, \tag{2.79}$$

with

$$\frac{1}{N_{\mathrm{vec}}}\sum_{i=0}^{N_{\mathrm{vec}}} {\eta^i}^{\dagger}_x \eta^i_y = \delta_{xy}\,. \tag{2.80}$$

In the limit of $N_{\mathrm{vec}} \to \infty$ this approximation becomes exact. Then a new set of vectors, ψ^i, is introduced as the solution to the following linear system

$$D^{\dagger}(U,\mu)\psi^i = \eta^i\,. \tag{2.81}$$

This system can be solved using, e.g., the Conjugate Gradient (CG) algorithm. In terms of η^i and ψ^i the fermion drift reads

$$\nabla^a S_F \approx -N_f \frac{1}{N_{\mathrm{vec}}}\sum_{i=1}^{N_{\mathrm{vec}}} \psi^{i\dagger}\nabla^a D(U,\mu)\eta^i\,. \tag{2.82}$$

Using this method, it is in principle necessary to compute the inverse of $D(U,\mu)$ for N_{vec} random vectors at each Langevin step.

If N_{vec} is sufficiently large, the actual drift is expected to be well approximated. However, the linear system in eq. (2.81) is typically very expensive to solve, thus forcing $N_{\mathrm{vec}} \sim O(1)$. Unfortunately, since this approximation is only exact for a large number of vectors, the drift can develop an imaginary component even at zero chemical potential, where the determinant of the Dirac matrix is real. Therefore care must be taken in controlling potential systematic uncertainties stemming from this. An extrapolation to zero Langevin step size can be used to obtain correct results for finer lattices [105].

A source of difficulties when inverting the Dirac matrix is the non-zero chemical potential. This can be seen more easily with the continuum version of the Dirac operator in momentum space,

$$D = \gamma_\nu p_\nu + m + \gamma_4 \mu\,. \tag{2.83}$$

When $\mu = 0$ the smallest eigenvalue of D is proportional to m. However, the combination $m + \gamma_4\mu$ can make the Dirac matrix singular. In such situations the matrix does not have a formal inverse, while when $m + \gamma_4\mu \approx 0$ the numerical inversion procedure becomes very expensive and prone to errors. This issue makes the study of nearly massless quarks and certain regions of the phase diagram very difficult.

A subject of ongoing research is the presence of poles in the drift of eq. (2.78), due to branch cuts of the complex logarithm on the negative real axis, as suggested in [95]. This has been studied in [97], where the effect of these poles in the formal justification of complex Langevin has been investigated and the findings applied to HDQCD and full QCD.

Chapter 3

Tests of complex Langevin

We consider a wide variety of tests of the complex Langevin method, applied to pure SU(3) Yang–Mills theory, and then to QCD with heavy-dense quarks (HDQCD). QCD with staggered quarks will be mentioned briefly. Our main goal here is to test the ideas described in the previous chapter against known results generated using Monte Carlo and analytical expectations.

The first test considered is the application of complex Langevin to the pure Yang–Mills theory, where we study the deconfinement transition as function of the gauge coupling. In this theory the sign problem is not present, and therefore this is typically done with Monte Carlo or real Langevin simulations. However, here, we do not apply any reunitarisation procedure to counter the numerical round-off errors that push the evolution off the SU(3) manifold. Instead, gauge gooling is applied in combination with the Langevin evolution to contain the potential non-unitarity of the gauge links. This provides a good consistency check for complex Langevin, as complexification is not necessary, but the known results for the critical coupling should be reproduced.

Secondly, we consider HDQCD in the vicinity of $\mu^2 = 0$. For purely imaginary chemical potentials, $\mu^2 < 0$, the sign problem is absent and simulations using real Langevin are possible. Upon considering observables that are continuous functions of μ^2 we can check if and when complex Langevin is applicable by verifying if this continuity[1] is observed around $\mu^2 = 0$.

Then, we study the dependence of the plaquette on the Langevin step size. We compare complex Langevin simulations of HDQCD using the first and second order

[1]Provided that there is no transition to a different state of matter.

algorithms, described in sections 2.2.1 and 2.2.3, at a variety of average step sizes with known results from Metropolis-based Monte Carlo methods, which do not have step size dependences, due to their characteristic accept/reject step. The latter cannot be applied to Langevin simulations because they follow trajectories given by a (stochastic) differential equation. The Monte Carlo results were possible in the cases studied because they were carried out in regions of the parameter space where the sign problem is mild and the reweighting technique is applicable.

Lastly, we perform a number of tests of our dynamic stabilisation (DS) method, detailed in section 2.3.3. To motivate this method, we begin by showing a situation where gauge cooling is not enough to keep the dynamics stable. Then we study two situations, with mild and severe sign problems, for a wide range of the DS control parameter α_{DS}. We study histograms of the DS drift to see whether it introduces boundary terms to the partial integrations necessary in the proof of convergence, and also of the Langevin drift to see how it responds to the added force. Studies of how the DS drift behaves as function of the gauge coupling β are also performed, in order to check for the expected scaling with the lattice spacing. Then we move to a study of the deconfinement transition in HDQCD, comparing results from complex Langevin with DS and results from Monte Carlo. Again, this comparison is possible because the sign problem for the studied case is mild and reweighting can be used. Lastly, we show how dynamic stabilisation can help when simulating dynamical quarks on the lattice.

Before going further, it is useful to review the actions discussed in section 1.2 for Yang–Mills, HDQCD and QCD with staggered quarks, with the relevant fields being rescaled by the appropriate power of the lattice spacing in order to make them dimensionless. Throughout this chapter we will be using the Wilson plaquette action for the $\mathrm{SL}(N,\mathbb{C})$ gauge links:

$$S_{\mathrm{G}} = \frac{\beta}{N} \sum_{x} \sum_{\mu<\nu} \mathrm{Tr} \left[\mathbb{1} - \frac{1}{2} \left(U_{x,\mu\nu} + U_{x,\mu\nu}^{-1} \right) \right] , \tag{3.1}$$

with β being the gauge coupling constant and the plaquette

$$U_{x,\mu\nu} = U_{x,\hat{\mu}} U_{x+\mu,\hat{\nu}} U_{x+\nu,\hat{\mu}}^{-1} U_{x,\hat{\nu}}^{-1} . \tag{3.2}$$

The quarks' actions, being bilinear in the fermion fields, can be integrated analytically,

and lead to the quark determinants. We will consider that of staggered quarks,

$$\det\left[D_{\mathrm{KS}}\right] = \det_{x,c,s}\left\{m\delta_{xy} + \sum_{\nu}\frac{\eta_{x,\nu}}{2a}\left[e^{\mu\delta_{\nu,4}}U_{x,\nu}\delta_{x+\hat{\nu},y} - e^{-\mu\delta_{\nu,4}}U^{-1}_{x-\hat{\nu},\nu}\delta_{x-\hat{\nu},y}\right]\right\}, \quad (3.3)$$

where $\det_{x,c,s}$ means that the determinant is taken over spacetime, spinor and colour indices, μ is the chemical potential, and $\eta_{x,\nu}$ is the staggered fermion sign function. We also consider the limit of heavy quarks (HDQCD) that only propagate in the Euclidean time direction,

$$\det\left[D_{\mathrm{HD}}\right] = \prod_{N_f}\prod_{\vec{x}}\left\{\det\left[1 + \left(2\kappa e^{\mu}\right)^{N_\tau}\mathcal{P}_{\vec{x}}\right]^2 \det\left[1 + \left(2\kappa e^{-\mu}\right)^{N_\tau}\mathcal{P}^{-1}_{\vec{x}}\right]^2\right\}, \quad (3.4)$$

with N_f the number of (mass degenerate) flavours and the Polyakov loop defined as

$$\mathcal{P}_{\vec{x}} = \prod_{\tau=0}^{N_\tau-1} U_{(\vec{x},\tau),\hat{4}}\,. \quad (3.5)$$

For HDQCD, we have studied the expectation value of the traced Polyakov loop and its traced inverse

$$\langle P\rangle = \frac{1}{V}\sum_{\vec{x}}\langle P_{\vec{x}}\rangle = \frac{1}{3V}\sum_{\vec{x}}\langle \mathrm{Tr}\left[\mathcal{P}_{\vec{x}}\right]\rangle\,, \quad (3.6)$$

$$\langle P^{-1}\rangle = \frac{1}{V}\sum_{\vec{x}}\langle P^{-1}_{\vec{x}}\rangle = \frac{1}{3V}\sum_{\vec{x}}\langle \mathrm{Tr}\left[\mathcal{P}^{-1}_{\vec{x}}\right]\rangle\,. \quad (3.7)$$

Both the Polyakov loop and its inverse are complex-valued for a given gauge configuration, but with real expectation values, since they are related to the free energy of a single (anti-)quark. Therefore we consider the symmetrised combination

$$\mathcal{P}^{s}_{\vec{x}} = \frac{1}{2}\left(\mathcal{P}_{\vec{x}} + \mathcal{P}^{-1}_{\vec{x}}\right), \quad (3.8)$$

which matches the definition of real part of a complex number for SU(3) configurations.

Additionally, when searching for the critical parameters of the models where the Polyakov loop is a (quasi-)order parameter, we have looked at other statistical quantities, namely the susceptibility,

$$\chi = N^3\left[\langle\mathcal{O}^2\rangle - \langle\mathcal{O}\rangle^2\right], \quad (3.9)$$

and Binder cumulant [107]

$$B = 1 - \frac{\langle \mathcal{O}^4 \rangle}{3\langle \mathcal{O}^2 \rangle^2}\,, \tag{3.10}$$

shown here written in terms of a generic observable, $\mathcal{O}$. The susceptibility is defined as the (pure or mixed) second derivative of a free energy, measuring the response of an observable with respect to changes in an external parameter. Analyses of the system's susceptibility with respect to different paramaters, at various volumes, provides information on its critical behaviour.

The Binder cumulant has limiting values 0 and 2/3, which can be understood as follows: Let $\langle \mathcal{O} \rangle$ be zero in one phase and non-zero in another. Assuming the higher moments are governed by Gaussian fluctuations one can see that

$$\langle \mathcal{O} \rangle = 0 \iff B = 0\,, \qquad \langle \mathcal{O} \rangle \neq 0 \iff B = \frac{2}{3}\,, \tag{3.11}$$

where it is assumed that $\langle \mathcal{O}^2 \rangle - \langle \mathcal{O} \rangle^2 \ll \langle \mathcal{O} \rangle^2$ in the latter case. These values have corrections that scale with the inverse of the lattice volume. Additionally, for large enough volumes, curves for the Binder cumulant as a function of an external parameter cross at a 'fixed point' value, whose location is a critical point.

In our tests involving staggered quarks we have looked at the chiral condensate,

$$\Sigma = \frac{1}{N^3 N_\tau} \frac{\partial \log Z}{\partial m} = \langle \overline{\psi}\psi \rangle\,, \tag{3.12}$$

where Z is the partition function, m is the quark mass, and N and N_τ are the lattice extents in the spatial and temporal directions, respectively. It is the order parameter for the chiral symmetry of the massless fermion action. When light quarks are present it is no longer an order parameter, but similar to the Polyakov loop when the centre symmetry is not exact, it provides information on the different phases of matter.

3.1 Deconfinement transition in pure Yang–Mills

In order to validate the complex Langevin and gauge cooling methods we have investigated their range of validity as function of the gauge coupling β in the SU(3) pure Yang–Mills theory. This theory has a real action, and so does not exhibit a sign problem; consequently, traditional methods like Monte-Carlo and real Langevin dynamics

would suffice. However, the updates could lead to non-unitary gauge links due to numerical roundoff errors, which are typically dealt with by periodically reunitarising the links. Simulations of pure Yang–Mills theories are good consistency checks for complex Langevin and gauge cooling, as they should agree with known results generated with other techniques, but with gauge cooling providing an alternative way to control the non-unitarity of the gauge links. These results have first appeared in [108].

This agreement was tested by searching for the critical gauge coupling, β_c, at which the SU(3) centre symmetry of the Yang–Mills theory is broken. This breaking is marked by a first order phase transition to non-zero expectation values for the Polyakov loop, which also signifies the transition from the confined to the deconfined phase.

Simulations were performed for two lattice volumes, $12^3 \times 3$ and $16^3 \times 4$ with the gauge coupling in the range $0.5 \leq \beta \leq 6.5$ using complex Langevin with adaptive step size and gauge cooling. Volumes were chosen such that a comparison with the literature [109] was possible. Our typical simulations have run until a Langevin time of 500, with the first 100 discarded to account for the thermalisation time, and average step sizes ranging from 10^{-3} to 10^{-4}. We have compared our results for β_c, determined using the Binder cumulant, with those of [109]. Figures 3.1 and 3.2 show the average Polyakov loop, while figs. 3.3 and 3.4 show the Binder cumulant for these lattices as functions of β.

Figures 3.1 and 3.2 show the expected behaviour, with the average Polyakov loop vanishing in the confined phase, with $\beta < \beta_c$, and jumping to a non-zero value after the critical gauge coupling. The region where the jump happens is consistent with the value obtained in [109]. A clearer signal of the transition can be seen in figures 3.3 and 3.4, where Binder cumulants of 0 and 2/3 identify each phase of the theory.

It is noticeable that the critical β for the second lattice, $16^3 \times 4$, is slightly larger than that of the first one. If we consider the Yang–Mills theory as the description of some physical phenomenon, the transition happens at well-defined temperature. In terms of lattice parameters, the physical temperature is given by $T = 1/aN_\tau$, with a the lattice spacing and N_τ the lattice extent in the temporal direction. By increasing the number of points in the Euclidean time direction from 3 to 4, a decrease in a is to be expected for a fixed temperature. From renormalisation group arguments [55], $\beta \sim 1/g^2 \sim -\log(a)$, hence the observed larger inverse coupling is to be expected.

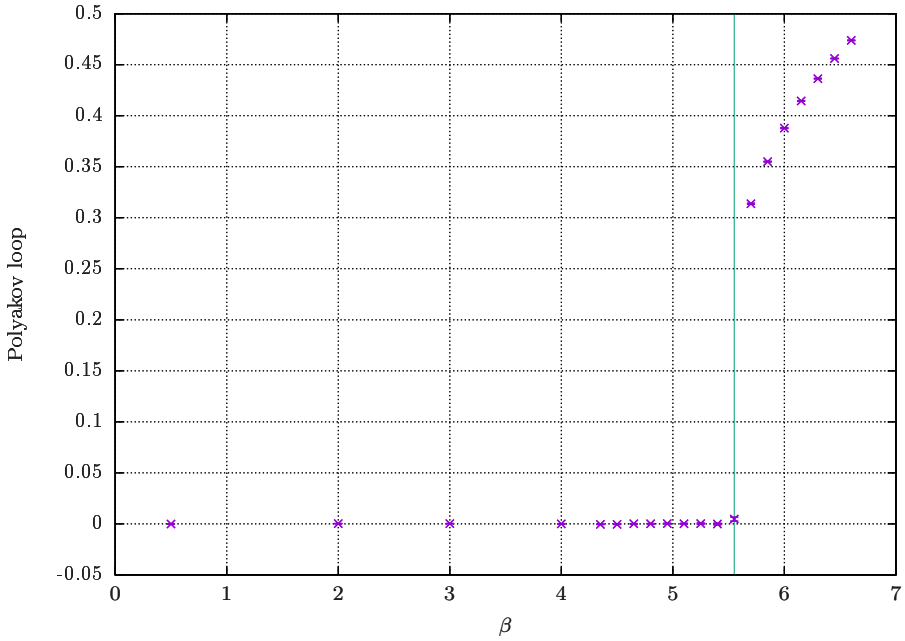

Figure 3.1: Average value of the Polyakov loop in a $12^3 \times 3$ lattice as a function of β for a pure SU(3) Yang–Mills theory. The vertical line indicates β_c from [109].

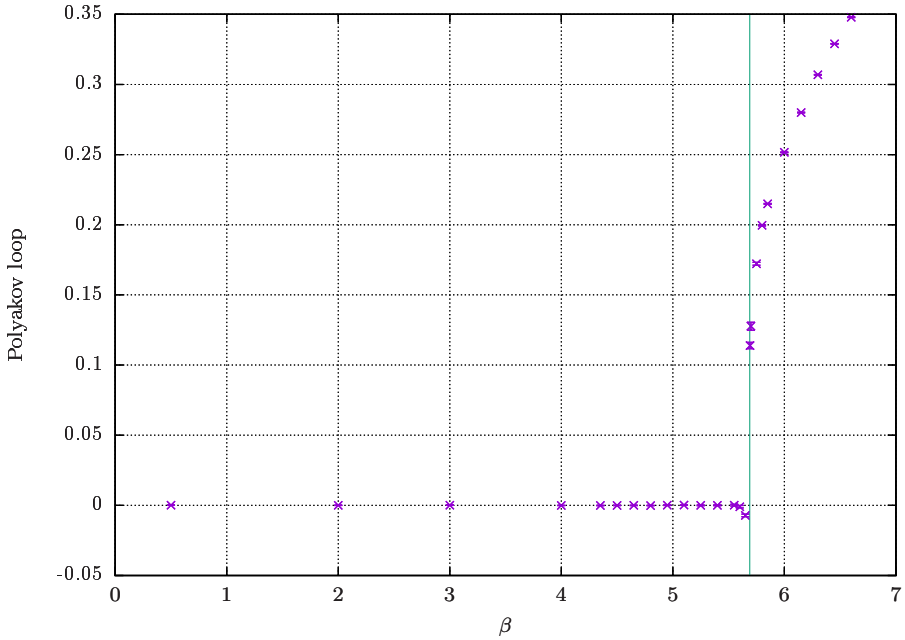

Figure 3.2: Average value of the Polyakov loop in a $16^3 \times 4$ lattice as a function of β for a pure SU(3) Yang–Mills theory. The vertical line indicates β_c from [109].

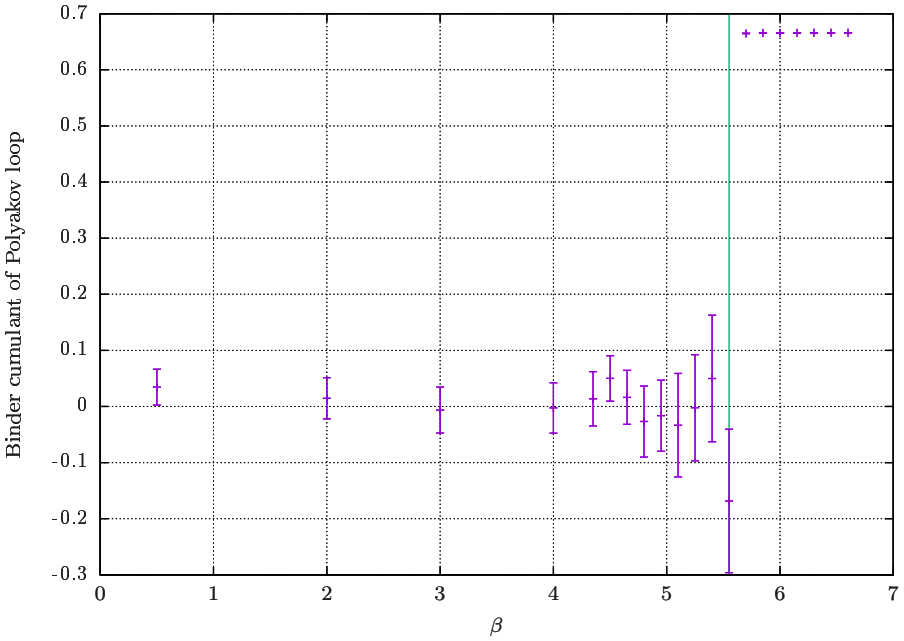

Figure 3.3: Binder cumulant of the Polyakov loop in a $12^3 \times 3$ lattice as a function of β for a pure SU(3) Yang–Mills theory. The vertical line indicates β_c from [109].

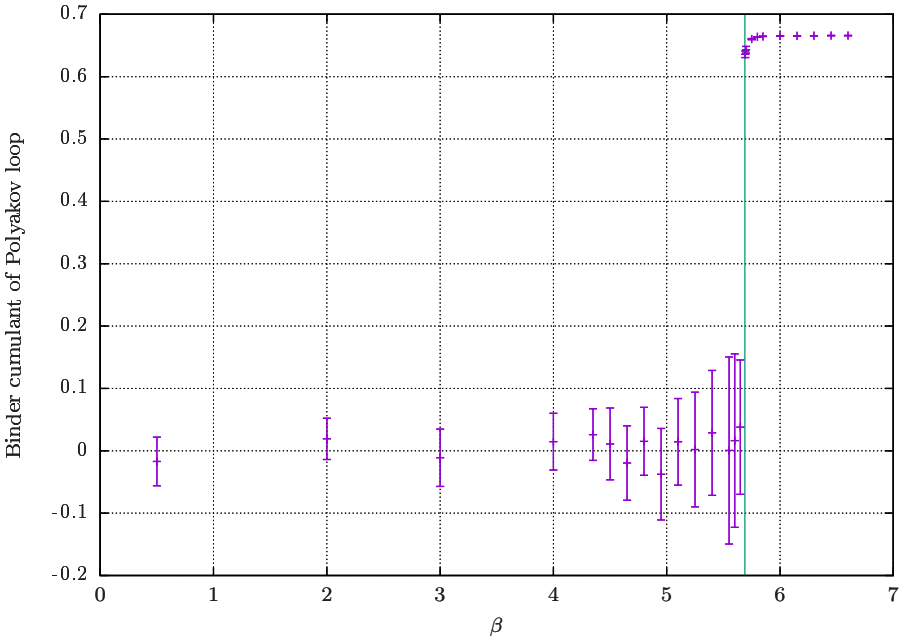

Figure 3.4: Binder cumulant of the Polyakov loop in a $16^3 \times 4$ lattice as a function of β for a pure SU(3) Yang–Mills theory. The vertical line indicates β_c from [109].

The critical gauge coupling extracted from our simulations via the Binder cumulant is in accordance with the expected value from the literature. Figures 3.1 to 3.4 also support that, despite the enlarged phase space, the complex Langevin method is able to produce the correct results. Excursions into the non-compact directions of $\mathrm{SL}(3,\mathbb{C})$ have been prevented without reunitarising the gauge links, as the typical average unitarity norm of our simulations was of $O(10^{-27})$.

We stress that, strictly speaking, the behaviour shown in figs. 3.1 to 3.4 arises from the apparent absence of tunnelling between the three vacua allowed by the $\mathbb{Z}_3$ symmetry. This is further demonstrated in fig. 3.5, where results from three sets of simulations are shown, each of them starting close to a different vacuum. This shows that, for the system size and simulation time used, tunnelling is absent and, moreover, the $\mathbb{Z}_3$ symmetry was not affected by the complexification procedure.

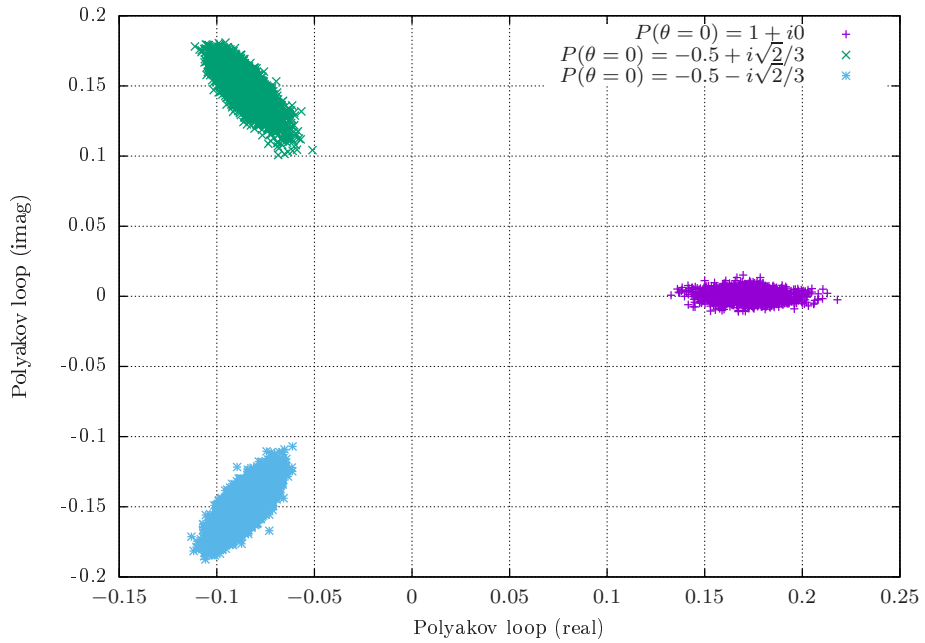

Figure 3.5: Scatter plot of the real and imaginary parts of the Polyakov loop of three simulations with different initial conditions. The simulations were carried out in a volume of $16^3 \times 4$ and gauge coupling $\beta = 5.75$.

3.2 Continuity of the plaquette around $\mu = 0$

As previously discussed, QCD in the limit of heavy quarks (HDQCD) suffers from the sign problem at real chemical potentials. On the other hand, when μ is purely imaginary, the quark determinant is real and the model can be simulated using, e.g., the real Langevin equation method, where the gauge links are periodically reunitarised. Therefore, observables that are continuous functions of μ^2 should have compatible average values for both purely imaginary and purely real chemical potentials. These results have first appeared in [108].

One such observable is the plaquette, which we have studied as a function of μ^2, for different values of the gauge coupling, using the first order algorithm. We have used $\kappa = 0.04$ as the hopping parameter, $N_f = 2$ quark flavours in a $8^3 \times 16$ lattice and $-0.3 < \mu^2 < 0.3$, well below the critical chemical potential at zero temperature of $\mu_c^0 \approx 2.53$. At this temperature the system is within the Silver Blaze region; very little dependence on the chemical potential is expected, as it is exponentially suppressed. As with the previous tests, the points before Langevin time of 100 were discarded, with the simulation running until $\theta = 500$, employing gauge cooling and adaptive Langevin step size averaging around $O(10^{-4})$. Our results are shown in figures 3.6, 3.7 and 3.8.

For $\beta = 6.2$ and $\beta = 5.8$ we provide linear fits for the real part of the plaquette, together with their error bands. The figures show clearly that the points fluctuate randomly around the fitted line, thus confirming the expected continuity. It is worth noting that the statistical fluctuations are in the fifth decimal place. Reducing them at this precision is considerably expensive and beyond the scope of this test.

At stronger couplings (lower β) the complex Langevin method has been known to give inaccurate results. Some can be seen in [99], where it fails to match results generated using reweighting, in a region where the latter is reliable. The graph in fig. 3.8 verifies this statement, showing a discontinuity at $\beta = 5.4$. Besides the discontinuity, we note that the large error bars on the right-hand side of figure 3.8 are due to the difficulty in determining the plaquette accurately. Another contributing factor is that the unitarity norm grew to $O(1)$ very quickly; as will be discussed in more details in section 4.2, data generated after the unitarity norm becomes too large is unreliable and should not be included in the statistical analysis.

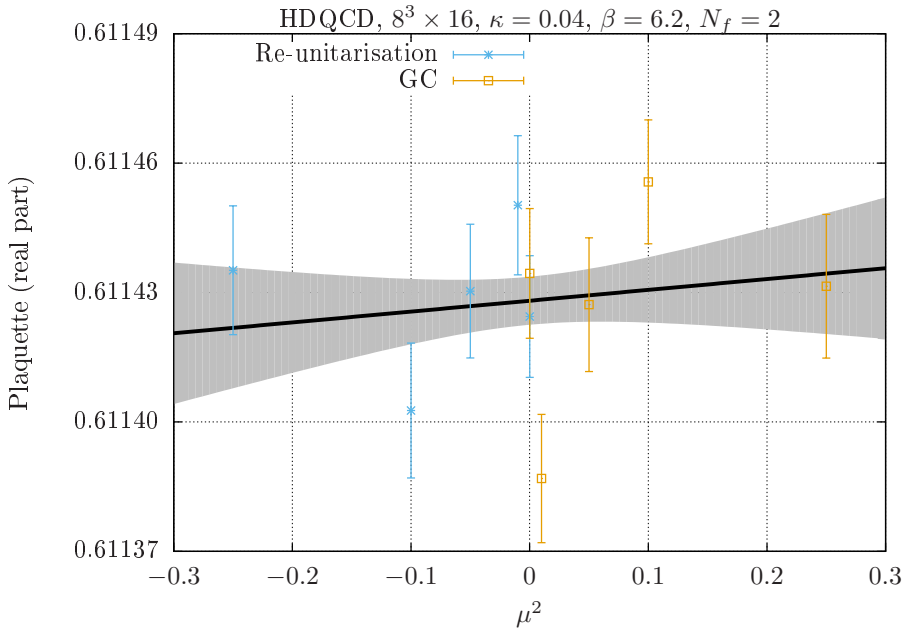

Figure 3.6: The average plaquette as function of μ^2 at $\beta = 6.2$. Points at $\mu^2 < 0$, generated with real Langevin, and at $\mu^2 > 0$, using complex Langevin, are compatible with the linear fit (black line). An error band for the fit is also shown.

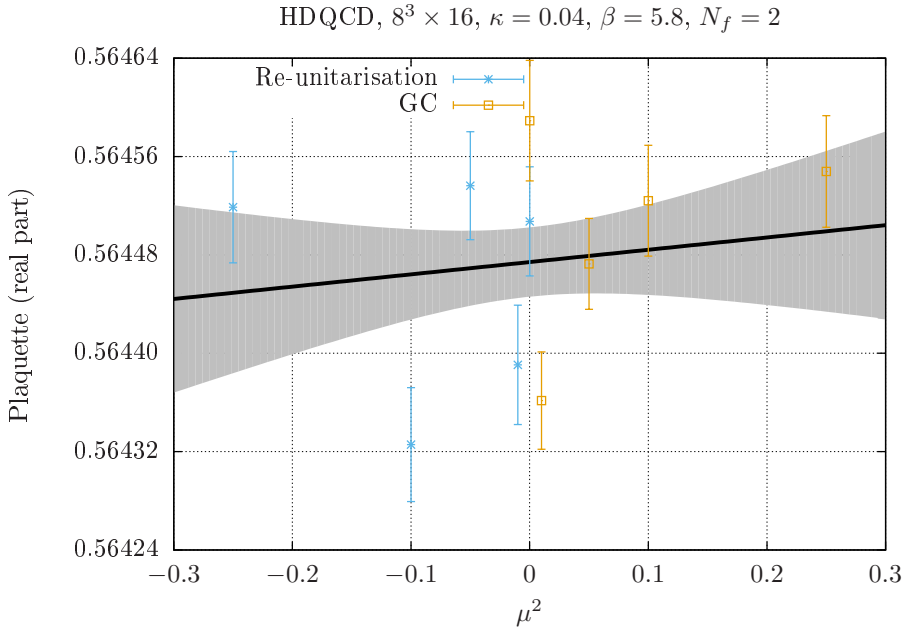

Figure 3.7: The average plaquette as function of μ^2 at $\beta = 5.8$. Points at $\mu^2 < 0$, generated with real Langevin, and at $\mu^2 > 0$, using complex Langevin, are compatible with the linear fit (black line). An error band for the fit is also shown.

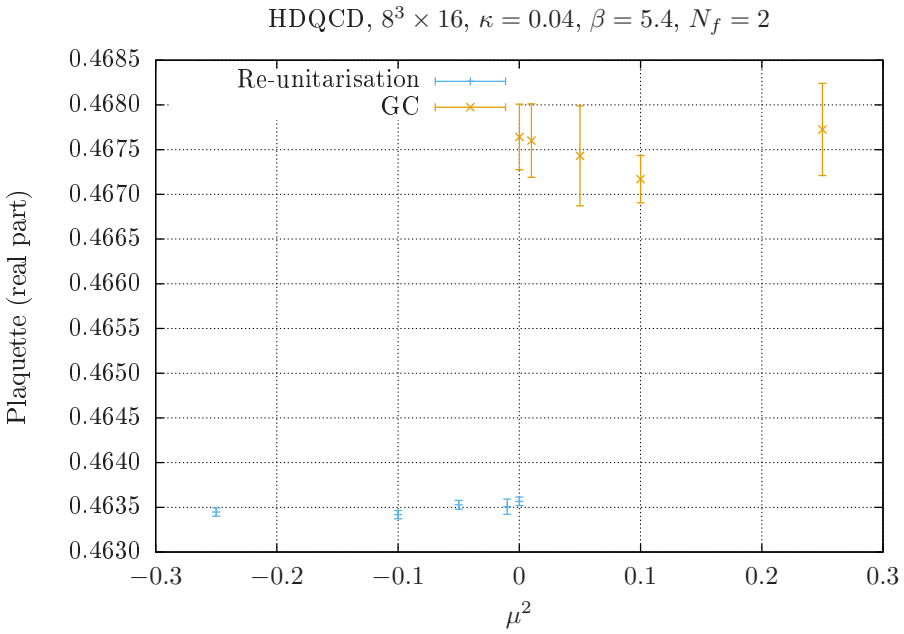

Figure 3.8: The average plaquette as function of μ^2 at $\beta = 5.4$. Points at $\mu^2 < 0$, have been generated with real Langevin, and those at $\mu^2 > 0$, with complex Langevin.

It is important to note that this study has been carried out using a first order discretisation scheme, with average Langevin step size of $\langle\varepsilon\rangle \sim 1.5 \times 10^{-4}$, and therefore systematic errors are to be expected. A simulation at $\langle\varepsilon\rangle \sim 1.5 \times 10^{-5}$, $\beta = 5.8$ and $\mu = 0$ using real Langevin and the same discretisation method gives an average plaquette of 0.567458(46), which differs from the value of figure 3.7 in the third decimal place. We stress that despite the high statistical precision of the results quoted in this section, our aim was to verify the plaquette's continuity around $\mu^2 = 0$ and not make precision measurements.

3.3 Dependence of observables on Langevin step size

We have tested the effectiveness of the $O(\varepsilon^2)$ improvement to the Langevin process discussed in section 2.2.3 in the HDQCD model. This improvement should cause the finite Langevin step size corrections to be smaller, making an eventual extrapolation $\varepsilon \to 0$ more reliable and closer to the correct answer for finite ε. With these systematic uncertainties reduced, better comparison with results from Monte Carlo and other methods

are possible, when the sign problem is milder and techniques such as reweighting can be applied.

Our first test was done using HDQCD in a lattice of size $6^3 \times 6$, coupling $\beta = 5.9$, hopping parameter $\kappa = 0.12$ and chemical potential $\mu = 0.85$, and used gauge cooling and dynamic stabilisation, with $\alpha_{\mathrm{DS}} = 100$, to guarantee that the unitarity norm was under control. These parameters were chosen to match those used in [99], from which the value generated using reweighting was used to compare with our simulations. As before, a thermalisation time of 100 has been used, but now with longer runs, going up to 2100. Figure 3.9 shows, in a logarithmic scale for the horizontal axis, our results for the usual first order algorithm, both choices of parameters (eqs. (2.46)-(2.48) and (2.49)-(2.53)) for the second order one, and the result from reweighting.

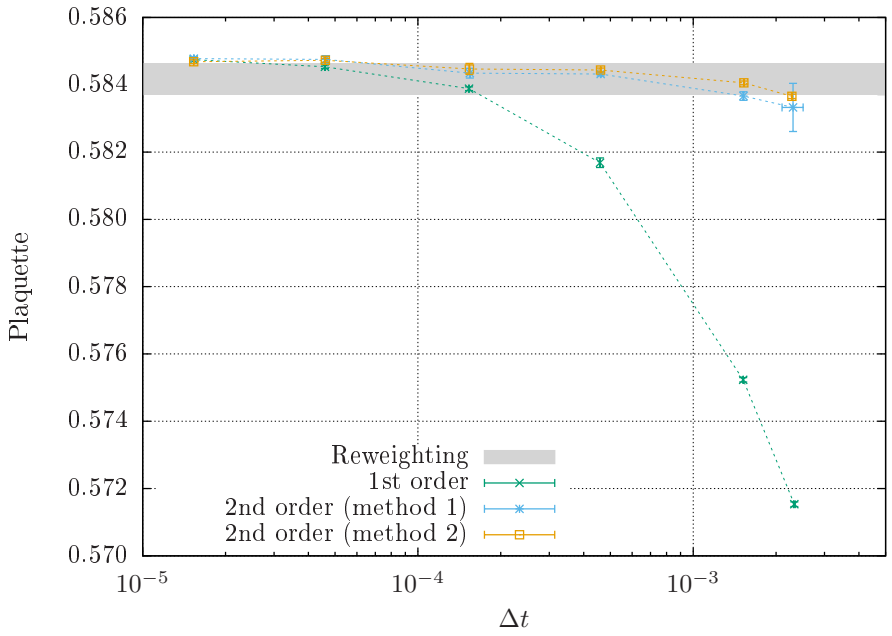

Figure 3.9: A comparison of the step size dependence of the plaquette in HDQCD between 1st and 2nd order algorithms and reweighting, for $6^3 \times 6$, $\beta = 5.9$, $\kappa = 0.12$ and $\mu = 0.85$ using dynamic stabilisation with $\alpha_{\mathrm{DS}} = 100$. Dashed lines are to guide the eye.

The graph shows that both second order methods agree with one another and with the grey band representing the reweighting result already at step sizes of $O(10^{-3})$, while the first order scheme required $\langle \varepsilon \rangle \lesssim O(10^{-4})$. One of the main reasons for using an improved discretisation scheme is that, for given finite ε corrections, simulations are

faster, since an improvement in this systematic uncertainty by a factor of 10 compared to 1st order outweighs their greater computational complexity.

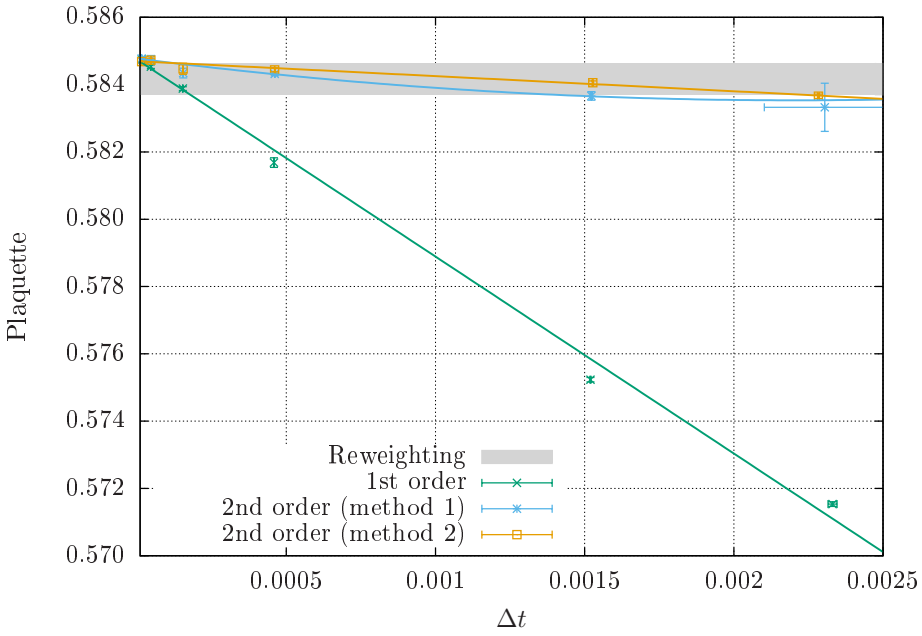

Figure 3.10: A comparison of step size dependence of the plaquette in HDQCD between 1st and 2nd order algorithms and reweighting for $6^3 \times 6$, $\beta = 5.9$, $\kappa = 0.12$ and $\mu = 0.85$ using dynamic stabilisation with $\alpha_{\mathrm{DS}} = 100$. Linear and quadratic fits are provided.

Figure 3.10 shows the same data as figure 3.9 with a linear scale for the Langevin step size, together with linear and quadratic fits, for the first and second order points, respectively. The linear fit shows an evident trend for the first order points, with clear indication that the improved algorithm removes this linear dependence on the step size.

We have also performed tests including staggered fermions. However, it is known that dynamical fermions cannot be fully improved to $O(\varepsilon^2)$ due to the non-integrability of certain terms in the Langevin drift [106]. Our results using the second order scheme were incompatible with those of first order, further confirming that the Langevin evolution of dynamical fermions cannot be improved in the same way as that of gluons, as shown in fig. 3.11. These simulations have been performed in a lattice of volume $8^3 \times 2$, inverse coupling $\beta = 5.4$, chemical potential $\mu = 0$ and quark mass $m = 0.05$. The average Langevin step sizes have been of $O(10^{-5})$; i.e., in the region where step size corrections should be small enough that both algorithms should agree.

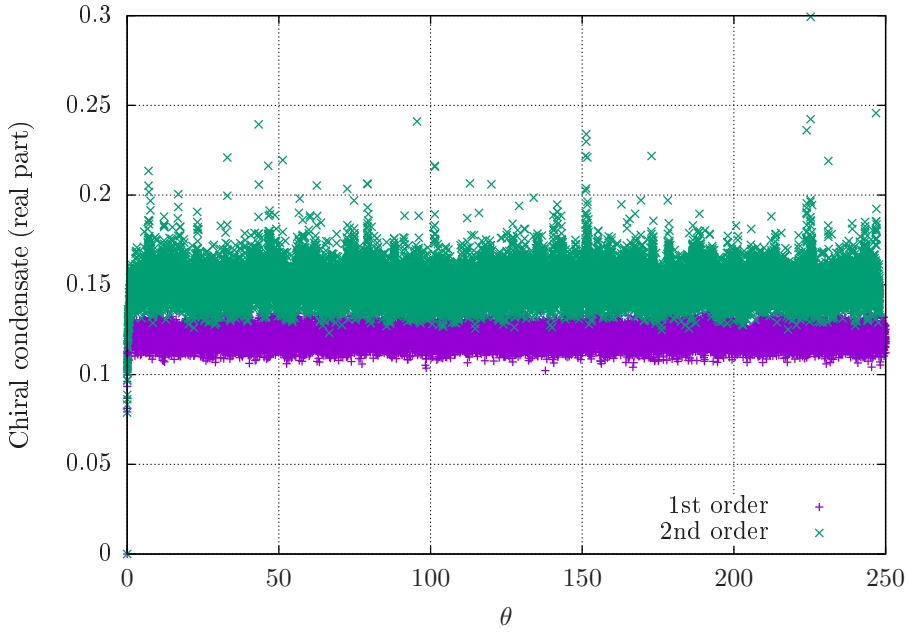

Figure 3.11: A comparison of first and second order algorithms applied to a Langevin simulation of staggered fermions. The simulations have been performed in a lattice of volume $8^3 \times 2$, with $\beta = 5.4$, $\mu = 0$ and $m = 0.05$.

3.4 Tests of dynamic stabilisation

3.4.1 Dependence of observables on α_{DS}

Our proposed method to constrain the exploration of the $\mathrm{SL}(N, \mathbb{C})$ phase space, dynamic stabilisation (DS), has one free parameter labelled α_{DS} that controls how strong the restoring force is. Its effect on the Langevin dynamics is difficult to predict, however, due to the high complexity of the gauge theories and of the complex Langevin method itself. In the low α_{DS} limit the DS drift becomes very small and the dynamics is essentially unaffected. For large α_{DS}, on the other hand, excursions into the non-unitary directions of $\mathrm{SL}(N, \mathbb{C})$ become heavily suppressed. In practice, this can be viewed as a "soft" reunitarisation of the gauge links.

The question of whether there is an interval of values for α_{DS} that keeps the unitarity norm bounded while still allowing for exploration of the enlarged group manifold is the central point of this section. This has been tested using the HDQCD model. In some cases, results from reweighting were available, so a comparison was possible. However, this means that the average phase of the quark determinant is not close to

zero, indicating the absence of (or a very mild) sign problem. More interesting tests have been done closer to the deconfinement transition, where reweighting cannot be applied, but simulations with gauge cooling suffered from instabilities that reduced the size of the statistical sample available.

Our initial tests have been done in a situation where the average phase of the quark determinant is close to unity, and therefore a comparison with reweighting was possible. We have chosen a volume of $10^3 \times 4$, $\beta = 5.8$, $\kappa = 0.04$ and $\mu = 0.7$. In this situation, the average Polyakov loop generated using complex Langevin and gauge cooling does not agree with that from reweighting. Further investigation has shown that their values did match for Langevin time below ~ 50, while the unitarity norm was below $O(0.1)$, as shown in figure 3.12. Further discussion about the instabilities

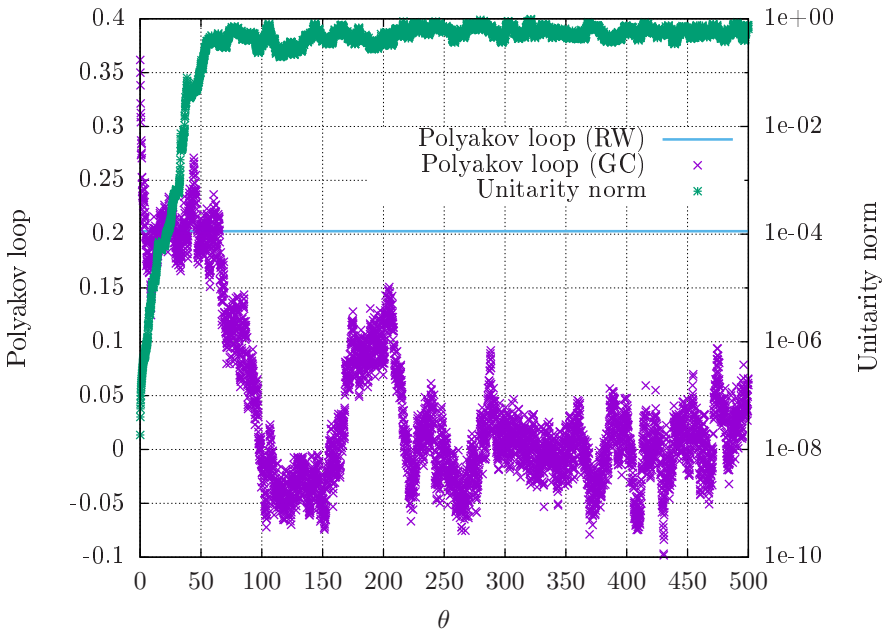

Figure 3.12: The average Polyakov loop and unitarity norm as functions of the Langevin time for HDQCD in a $10^3 \times 4$ lattice, with $\kappa = 0.04$, $\beta = 5.8$ and $\mu = 0.7$. Values generated using gauge cooling and reweighting agree when the unitarity norm is lower than $O(0.1)$.

displayed by the average Polyakov loop can be found in section 4.2. For now we consider the effect of adding the DS, term $i\varepsilon\alpha_{\mathrm{DS}}M_x^a$, defined in section 2.3.3, to the Langevin drift for a variety of values for α_{DS}. The results of this change applied to the same simulation parameters used above are shown in figs. 3.13 and 3.14, where the

simulations included one step of gauge cooling between successive Langevin updates. The former shows the Langevin time evolution of the average Polyakov loop, comparing gauge cooling, dynamic stabilisation for $\alpha_{\mathrm{DS}} = 100$, and reweighting, while the latter presents the same observable for different values of α_{DS}, with a panel zooming around the reweighting result.

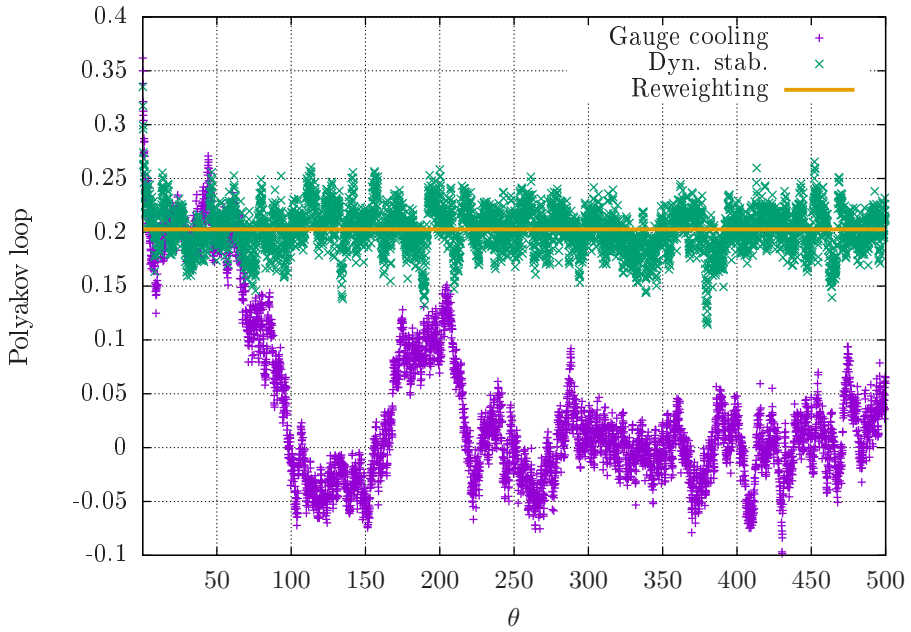

Figure 3.13: The average Polyakov loop as function of the Langevin time for HDQCD in a $10^3 \times 4$ lattice, with $\kappa = 0.04$, $\beta = 5.8$, $\mu = 0.7$ and $\alpha_{\mathrm{DS}} = 100$. The plot shows results using only gauge cooling, gauge cooling with dynamic stabilisation, and reweighting.

We find agreement between reweighting and the simulation using DS for all Langevin times after thermalisation, in fig. 3.13. Its average unitarity norm was of $O(10^{-3})$, below the region known to cause convergence to wrong limits. In fig. 3.14, we see that complex Langevin results agree with those from reweighting for sufficiently large α_{DS}. When the contribution from dynamic stabilisation is small, the instabilities shown in fig. 3.12 are still present.

Figures 3.13 and 3.14 seem to indicate that α_{DS} could be chosen arbitrarily large without consequences to the dynamics—i.e., without changing the physics. As mentioned previously, for the parameters used in those simulations the average phase of the quark determinant is close to unity, which means that the sign problem is very mild.

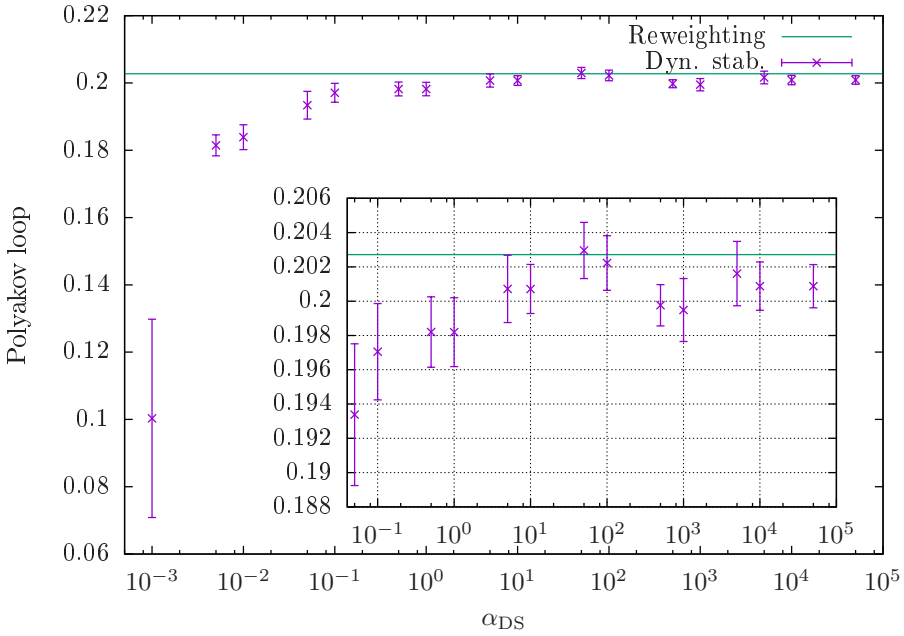

Figure 3.14: The average Polyakov loop as a function of α_{DS} compared with the result generated with reweighting for HDQCD in a $10^3 \times 4$ lattice, with $\kappa = 0.04$, $\beta = 5.8$ and $\mu = 0.7$. Agreement is reached once α_{DS} is sufficiently large.

In situations where the average phase of the fermion determinant is (close to) zero, reweighting is unreliable, and restricting the simulation to the unitary manifold would mean ignoring the oscillating phase of the quark determinant. Complex Langevin combined with gauge cooling is, in principle, a viable option, but the instabilities shown in fig. 3.12 have also been observed in this context. Dynamic stabilisation, given the previous results, is the ideal method to be applied in this case. Moreover, in such situations an excessive suppression of explorations of $\mathrm{SL}(3,\mathbb{C})$ could have effects similar to those of projecting the gauge links onto the unitary submanifold. This can give indications of upper bounds on α_{DS}.

Our choice of parameters for these tests are a lattice of size $8^3 \times 20$, with coupling $\beta = 5.8$, hopping parameter $\kappa = 0.04$ and chemical potential $\mu = 2.45$, which is very close to the critical one at zero temperature, $\mu_c^0 \approx 2.53$. The average phase of the fermion determinant has been compatible with zero for all of our choices of α_{DS}, being of $O(10^{-20})$ in some cases, indicating a severe sign problem. As before, we have used a Langevin time of 100 for thermalisation and have ran the simulations until $\theta = 500$.

The effect of DS on the unitarity norm is shown in fig. 3.15, where we analyse it as

function of $\alpha_{\rm DS}$ for the previously discussed set of parameters ($10^3 \times 4$, $\beta = 5.8$, $\kappa = 0.04$ and $\mu = 0.7$), which has an average quark determinant close to unity, and that of the previous paragraph, with $\det M \approx 0$. Essentially, the DS term was able to restrict the excursions in $\mathrm{SL}(3,\mathbb{C})$ to a submanifold whose distance from SU(3) decreases with $\alpha_{\rm DS}$. When the sign problem is more severe, a plateau seems to emerge for large $\alpha_{\rm DS}$, whereas for milder sign problems it follows an almost straight line in the log-log plot, as indicated by the fit of form $y = ax^b$.

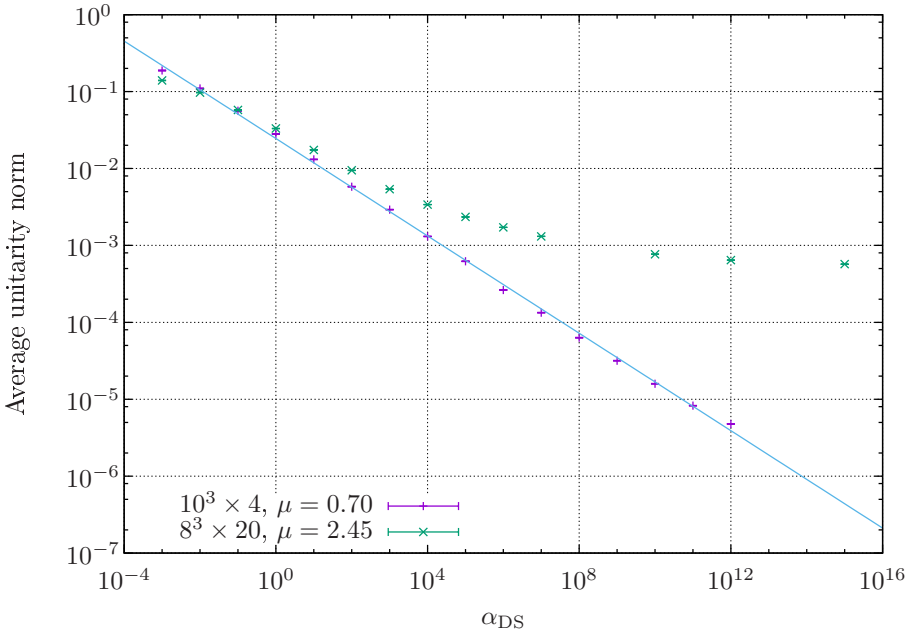

Figure 3.15: The average unitarity norm as a function of $\alpha_{\rm DS}$ for HDQCD, with $\kappa = 0.04$, $\beta = 5.8$, and volume and chemical potential indicated on the figure. A power-law fit is provided for the data with average fermion determinant close to unity.

Due to the severity of the sign problem, comparisons with reweighting are not possible. Fig. 3.16 shows results for the average Polyakov loop using only gauge cooling, and also with dynamic stabilisation, with $\alpha_{\rm DS} = 10^3$ and Langevin time $0 < \theta < 100$. The gauge cooled simulation exhibits the aforementioned instabilities when the unitarity norm becomes too large, which happens sooner than in the previous scenario. Nevertheless, the system seems to thermalise quickly, before converging to the wrong limit, and DS results are compatible with those of gauge cooling while $d < O(0.1)$, which corresponds to $\theta \lesssim 22$.

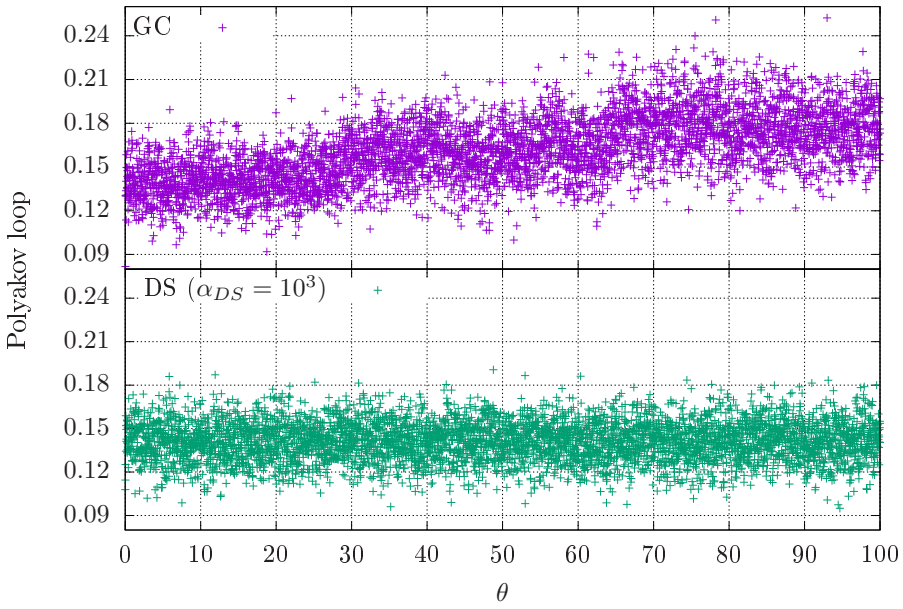

Figure 3.16: The Polyakov loop as function of Langevin time using gauge cooling and dynamic stabilisation for HDQCD in a $8^3 \times 20$ lattice, with $\kappa = 0.04$, $\beta = 5.8$ and $\mu = 2.45$. Convergence to wrong limiting value is seen in the gauge cooling data after $\theta \approx 22$.

A comparison amongst simulations using real Langevin, gauge cooling with unitarity norms above and below our threshold of 0.03 and dynamic stabilisation at various values of α_{DS} can be found in fig. 3.17. We have also added a line representing the value from the phase quenched version of HDQCD, where the fermion determinant is replaced by its absolute value. As in fig. 3.14 there is a region where DS agrees with gauge cooling results prior to the instability. However, for larger values of α_{DS} we observe a departure from this agreement, possibly towards the value simulated using real Langevin.

As will be discussed in 4.2, gauge cooling simulations provide robust results when the unitarity norm is small. The agreement between DS and gauge cooling with low unitarity norm, shown in fig. 3.17, is further evidence that dynamic stabilisation succeeds in keeping explorations of $\mathrm{SL}(3, \mathbb{C})$ under control. Thus, convergence to the right limit can be achieved with reasonable values for α_{DS}. For the moment, there is no known *a priori* way of determining them, which must be done in a case-by-case basis. Fortunately, α_{DS} does not require fine-tuning, since the range of values that results in correct convergence spans a few orders of magnitude, at least in the studied cases.

There is evidence in random matrix models [110] that complex Langevin simulations

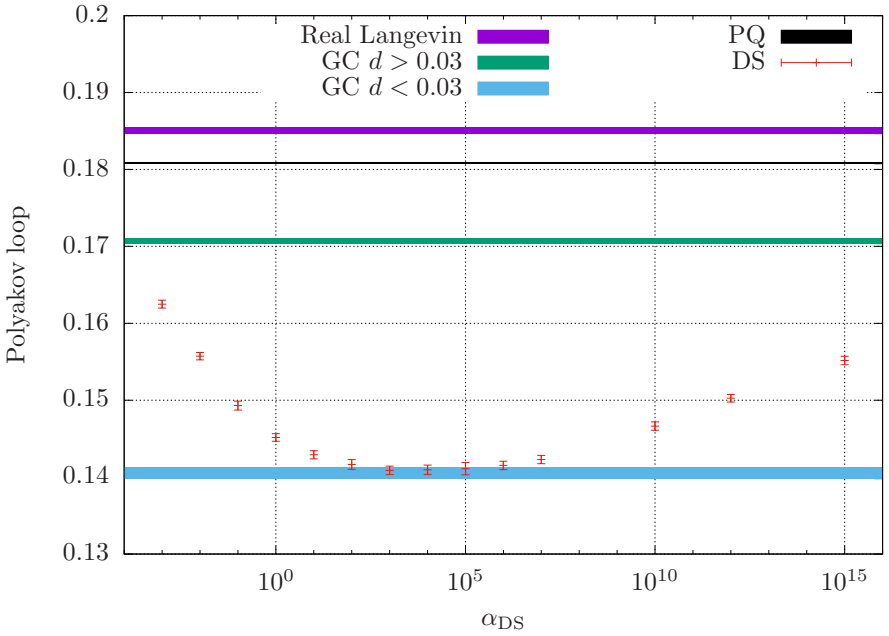

Figure 3.17: The average Polyakov loop as function of α_{DS}. Also shown are results from simulations with gauge cooling (GC) at different unitarity norm cutoffs, phase quenched (PQ) HDQCD, and with real Langevin.

can converge to phase quenched results in certain situations. When that is the case, a reweighting procedure can assure convergence to the right limit. However, we have found no indication of this behaviour in HDQCD, as the phase quenched result for the average Polyakov loop in the setup outlined above is 0.180844(91), shown in fig. 3.17.

For some of the simulations discussed above, we have plotted histograms of the absolute value of the dynamic stabilisation drift, $\alpha_{\mathrm{DS}}\varepsilon|M_x^a|$, as shown in fig. 3.18 on a log-log scale. In the context of the criteria for correctness and convergence to the correct limit, discussed in section 2.3.1, these histograms show how the count of the new addition to the Langevin drift decays. A slow decay, perhaps with heavy "tails", would spoil the arguments used in the proof. On the other hand, a localised distribution, with fast decay, should not prevent the convergence to the right limit, despite the term's non-holomorphy.

At $\alpha_{\mathrm{DS}} = 1.0$, the dynamic stabilisation force has little effect and the unitarity norm is still high, causing larger values of $|M_x^a|$. As α_{DS} increases we see an initial decrease in the number of larger values of $|M_x^a|$, due to the smaller unitarity norms, followed by an increase. This is due to α_{DS} increasing faster than the decrease in M_x^a.

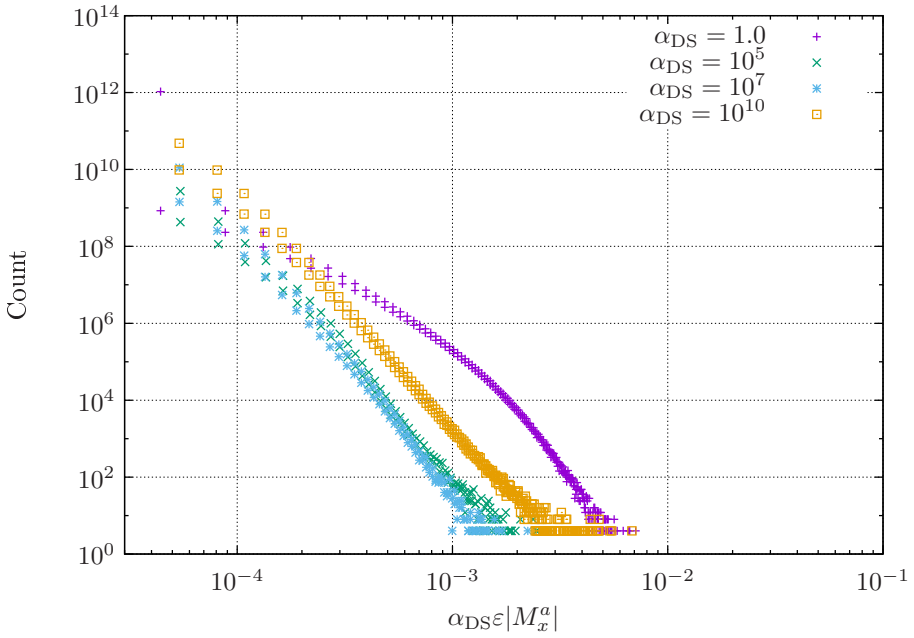

Figure 3.18: Histogram of the drift added by dynamic stabilisation for different values of the control parameter α_{DS} in a log-log scale.

When this happens, the DS drift risks overshadowing the Langevin drift coming from the physical action if α_{DS} is made larger. It is also clear that the drifts are localised for the intermediate range of α_{DS} where the results match those of gauge cooling prior to instabilities: there are no countings of $\alpha_{\mathrm{DS}}\varepsilon|M_x^a| > 10^{-2}$. This means that dynamic stabilisation does not add boundary terms to the partial integrations needed in the proof of convergence discussed in section 2.3.1. The values of the DS control parameter shown in the figure were chosen to display the DS drift's behaviour at extreme ($\alpha_{\mathrm{DS}} = 1.0$ and $\alpha_{\mathrm{DS}} = 10^{10}$) and more sensible cases ($\alpha_{\mathrm{DS}} = 10^5$ and 10^7). The former two can be seen to give results away from that of gauge cooling with low unitarity norm, while the latter two are much closer to it.

We have also prepared histograms of the real and imaginary parts of the Langevin drift, $K_{x,\mu}^a = -\nabla_{x,\mu}^a S$. The imaginary part, that which is orthogonal to SU(3), is suppressed by dynamic stabilisation. This clearly is visible in fig. 3.19, which also indicates a decreasing unitarity norm. Additionally, the real part is essentially unchanged by the DS term once α_{DS} is large enough, as seen in fig. 3.20. The reason $\mathrm{Re}[K_{x,\mu}^a]$ is different for lower DS parameters is that the system is still exploring a wide region of $\mathrm{SL}(3,\mathbb{C})$

and converging to a wrong limit. The kink in the graph is due to a change in the Langevin step size, from the adaptive algorithm.

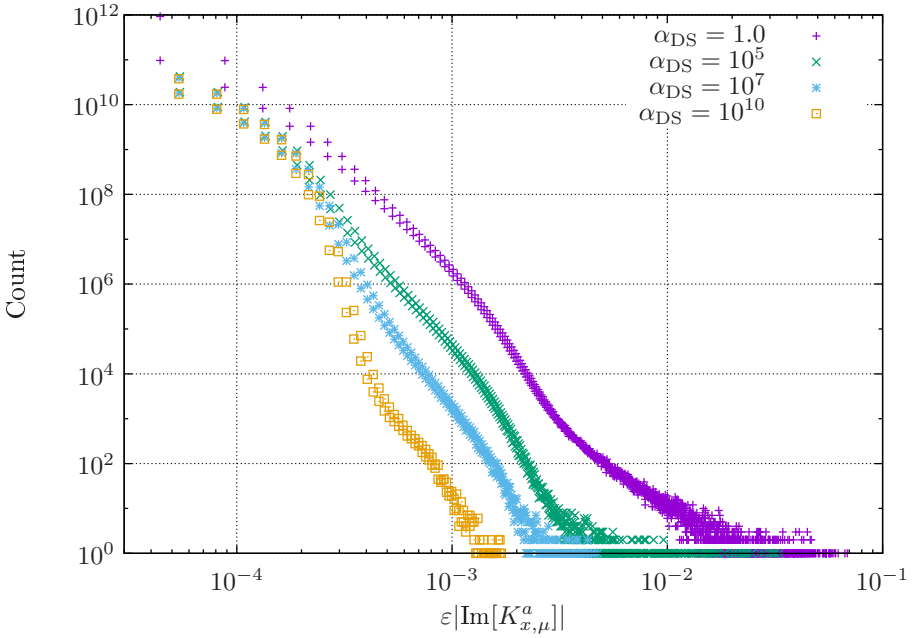

Figure 3.19: Histogram of the imaginary part of the Langevin drift $K^a_{x,\mu} = -\nabla^a_{x,\mu} S$, multiplied by the Langevin step size, for different values of the DS control parameter α_{DS} in a log-log scale.

3.4.2 Dependence of dynamic stabilisation on the gauge coupling

As mentioned in section 2.3.3, the dynamic stabilisation has been constructed to be irrelevant in the continuum limit. One way to verify this irrelevance is to analyse the distribution of the values of the DS drift at different lattice spacings. We have again used $\kappa = 0.04$ and $\mu = 2.45$ in a $8^3 \times 20$ lattice in a further three simulations at different values of the gauge coupling, specifically $\beta = 5.8$, 6.0 and 6.2, with $\alpha_{\mathrm{DS}} = 10^3$. From the results we have generated histograms of the DS drift, shown in fig. 3.21 in a log-log scale, together with power-law fits of the form $y = bx^a$, and also of the real part of the Langevin drift, shown in fig. 3.22. The fitted parameters, a and b, are shown in table 3.1.

The occurence of larger values of $|M^a_x|$ in fig. 3.21 decreases at finer lattices—i.e., larger β. In addition, comparing figs. 3.21 and 3.22 confirms that $|M^a_x|$ is more sensitive

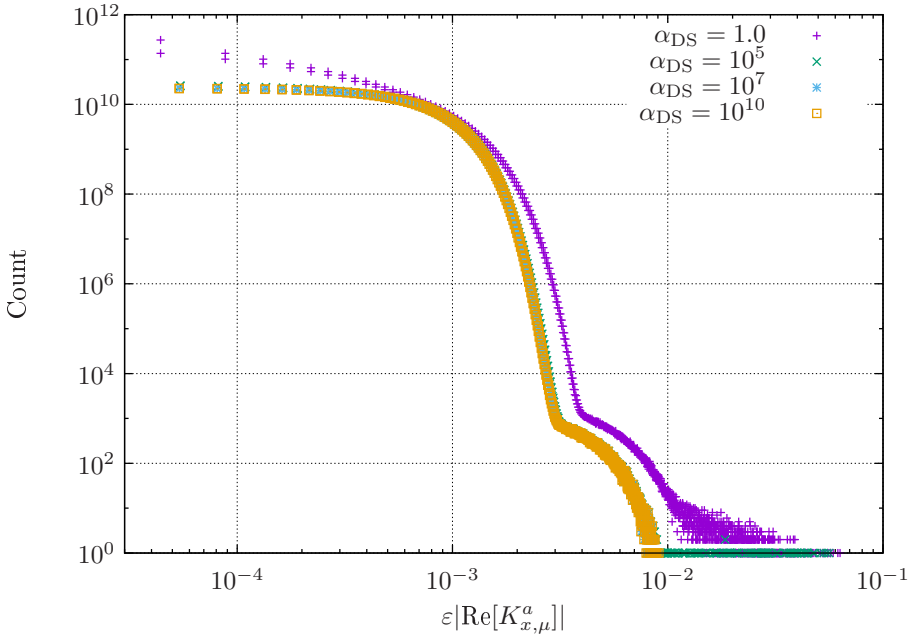

Figure 3.20: Histogram in a log-log scale of the real part of the Langevin drift $K^a_{x,\mu} = -\nabla^a_{x,\mu} S$, multiplied by the Langevin step size, for different values of the DS control parameter α_{DS}.

β	a	$\log_{10}(b)$
5.8	−4.372(44)	−9.48(12)
6.0	−4.637(58)	−11.50(17)
6.2	−5.64(10)	−15.61(34)

Table 3.1: Values and uncertainties for the parameters of the power law $y = bx^a$ fitted to the data shown in fig. 3.21.

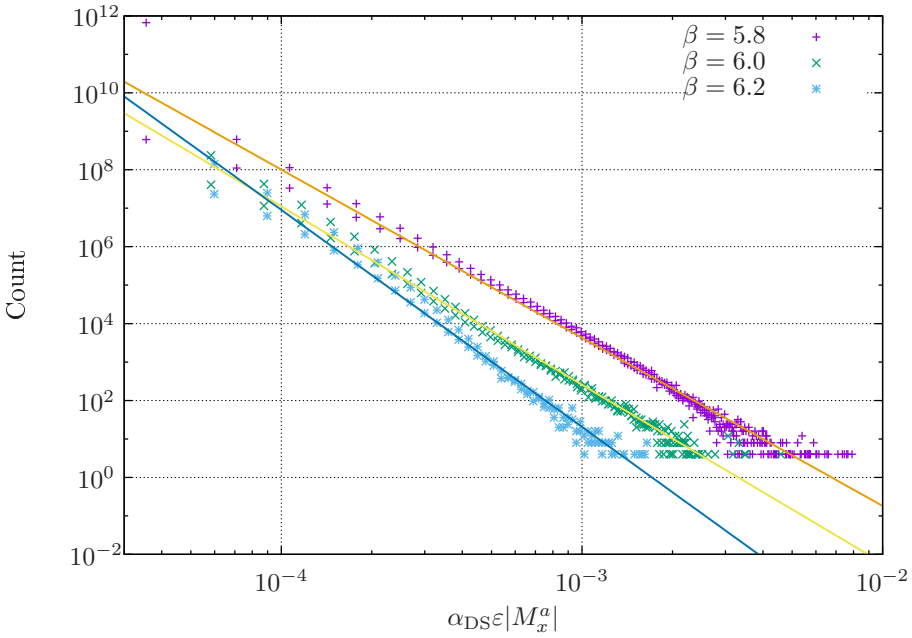

Figure 3.21: Histogram of the drift added by dynamic stabilisation for different values of the gauge coupling β in a log-log scale. Power law fits are also provided.

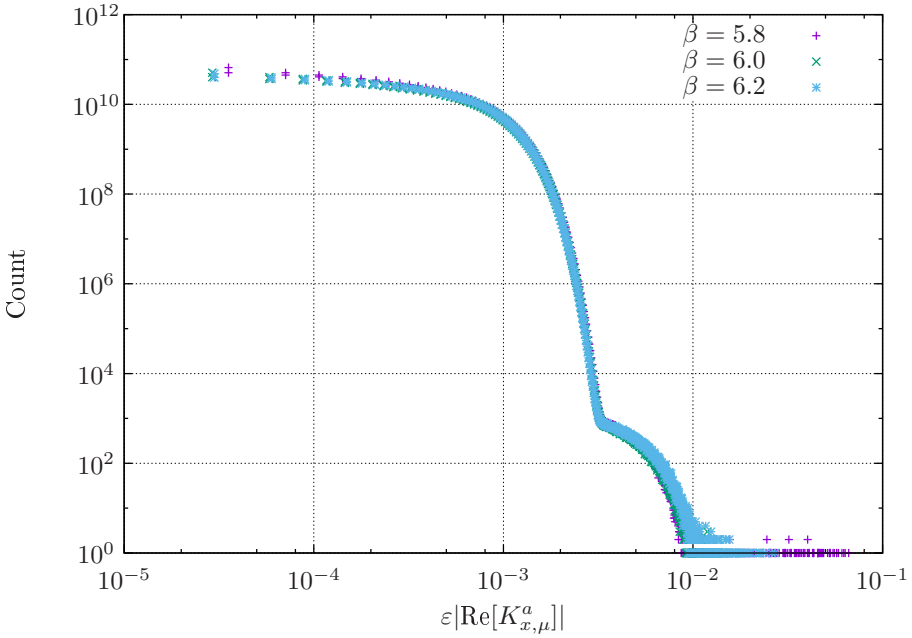

Figure 3.22: Histogram of the real part of the Langevin drift, for different values of the gauge coupling β, in a log-log scale. The kink in the graph is caused by a change in step size, introduced by the adaptive algorithm.

to changes in the gauge coupling than the real part of the Langevin drift. This indicates that $|M_x^a|$ is decreasing with the coupling faster than $K_{x,\mu}^a$, as it is supposed to. The small differences between the different data sets in fig. 3.22 are caused by the different physics being simulated because of the smaller coupling.

This is a crucial test for the dynamic stabilisation method, since it has been constructed to

- curb large exploration of $\mathrm{SL}(3,\mathbb{C})$;
- not change the SU(3) part of the Langevin drift;
- be irrelevant in the continuum limit.

The first point above has been demonstrated in the previous section, while the second is true in our construction shown section 2.3.3. It is worth noting that, as with gauge cooling, DS does not directly change the real part of the drift. However, due to the non-linearity of the update process, both can affect the observables, as expected. The last point is seen to be true from fig. 3.21 and table 3.1, where larger values for the inverse coupling resulted in larger slopes for the fitted lines.

3.4.3 Deconfinement transition for HDQCD

In this section we apply dynamic stabilisation to QCD with heavy quarks in order to reproduce the deconfinement transition by varying the gauge coupling. Our results are compared to those of [99], which have been generated with reweighting and with complex Langevin with adaptive step size and gauge cooling. As noted before, complex Langevin has been known be inaccurate at lower values of the coupling, namely $\beta \lesssim 5.5$. Even though DS has been conceived to deal with large explorations of the enlarged phase space necessary for the complex Langevin method, its effect on the dynamics at different couplings is not manifest.

We have employed dynamic stabilisation with $\alpha_{\mathrm{DS}} = 10^3$ on complex Langevin simulations of HDQCD with parameters chosen to match those used in [99], namely a lattice of size $6^3 \times 6$, $\kappa = 0.12$ and $\mu = 0.85$, with one gauge cooling step between successive Langevin updates. The same parameters have been used in simulations using complex Langevin and gauge cooling only, with a variable number of gauge cooling steps, depending on its effectiveness. Our average Langevin step size is of $O(10^{-4})$,

with thermalisation time of 100 and total Langevin time ranging from 600 to 5000 for gauge cooling simulations, and 5000 for all DS runs. The gauge coupling has been varied from $\beta = 5.4$ to 6.2. Fig. 3.23 shows the Polyakov loop and its inverse as functions of β, comparing reweighting results to our DS results. As already known, it is possible

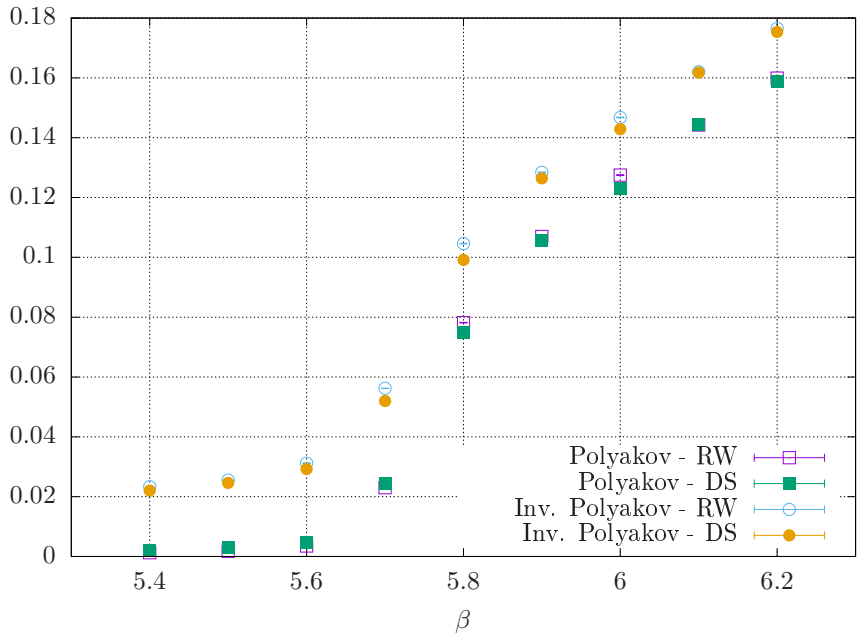

Figure 3.23: The average of Polyakov loop and of its inverse as functions of the gauge coupling. Data has been generated using dynamic stabilisation and compared with reweighting.

to go from a deconfined to a confined phase using complex Langevin. There is good agreement between our Polyakov loop averages and reweighting results for all β, and the same is true for the inverse Polyakov loop. Both methods see a zero Polyakov loop average at the same coupling, meaning that the critical gauge coupling should be compatible between them.

We compare the spatial plaquette's expectation value simulated using dynamic stabilisation, gauge cooling and reweighting in fig. 3.24 and table 3.2. As already noticed in [99], there is a clear discrepancy between gauge cooling and reweighting results for lower values of the gauge coupling. Including all points generated with gauge cooling in the analysis makes the disagreement appear much sooner than in the reference. When simulation data after the unitarity norm reaches our threshold of 0.03 is excluded from

the analysis, also shown in the figure, we are able to reproduce their results. Unfortunately, this also makes our statistical sample small, since the unitarity norms rise very quickly, making the error bars much larger than those from reweighting.

Using dynamic stabilisation, however, largely seems to resolve this tension. Both the figure and the table show excellent agreement between reweighting and DS for all values of β used in the simulations: at $\beta = 5.4$ the difference in the spatial plaquette between these methods is approximately of 1.08σ, while for $\beta \geq 5.5$ it is less than one standard deviation.

This indicates that going through the deconfinement transition is not the cause of the disagreement between complex Langevin with gauge cooling and reweighting. Large explorations of $\mathrm{SL}(3, \mathbb{C})$ could be related to the observed discrepancy, since cutting off data points after the unitarity norm became too large improved the results. Nevertheless, the non-trivial effects of DS on the dynamics seem to be more beneficial than just keeping the unitarity norm under control.

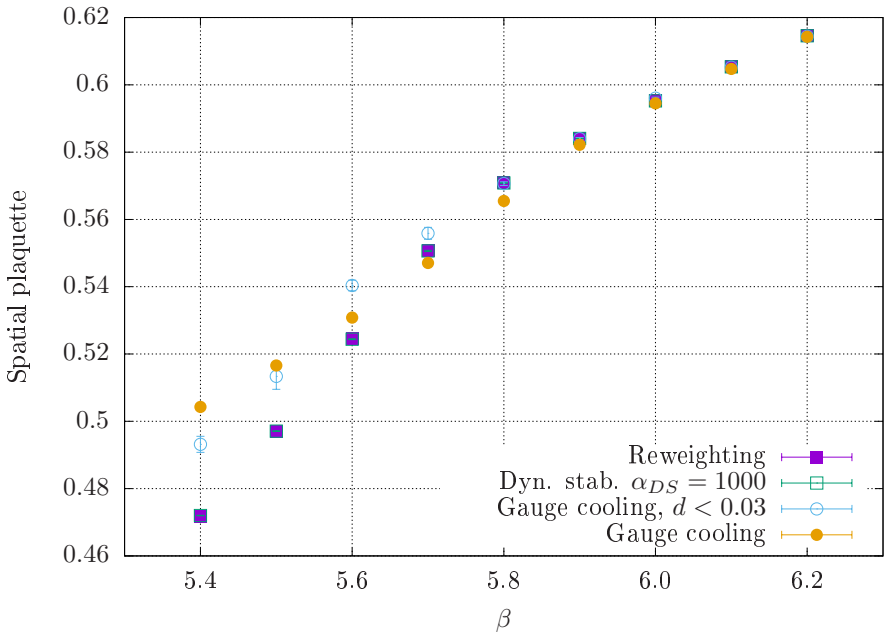

Figure 3.24: The expectation value of the spatial plaquette as function of the gauge coupling β from HDQCD simulations at $6^3 \times 6$, $\kappa = 0.12$ and $\mu = 0.85$. Results generated with dynamic stabilisation, gauge cooling and reweighting, from [99], are shown.

β	RW	DS	GC $d < 0.03$	GC
5.4	0.47164(33)	0.472007(86)	0.4931(24)	0.504292(75)
5.5	0.49687(38)	0.49708(11)	0.5133(38)	0.516607(56)
5.6	0.52461(47)	0.52441(12)	0.5404(16)	0.530817(72)
5.7	0.55086(63)	0.55064(19)	0.5559(18)	0.54705(97)
5.8	0.57097(58)	0.570849(69)	0.57062(55)	0.56547(22)
5.9	0.58417(47)	0.584086(37)	0.584(34)	0.5822(16)
6.0	0.59533(42)	0.59522(28)	0.59604(31)	0.59449(67)
6.1	0.60533(38)	0.605332(24)	0.60524(30)	0.604713(50)
6.2	0.6146(36)	0.614567(22)	0.61441(15)	0.614275(33)

Table 3.2: The average value for the spatial plaquette, from HDQCD simulations at $6^3 \times 6$, $\kappa = 0.12$ and $\mu = 0.85$, using reweighting, dynamic stabilisation and gauge cooling. Reweighting data have been taken from [99]. In the fourth column data after the unitarity norm became too large was excluded from analysis.

3.4.4 Staggered fermions with dynamical stablisation

As a final test of dynamic stabilisation we have investigated whether agreement between QCD simulations with lighter quarks using gauge cooling and dynamical stabilisation is observed, as it is with HDQCD, when the unitarity norm is below $O(0.1)$. These tests have been performed at zero chemical potential, since the number of random vectors used in the bilinear noise scheme, discussed in section 2.4, has to be small to keep the simulation costs low. Therefore a non-zero unitarity norm is to be expected, which can lead to instabilities if uncontrolled.

We have simulated QCD with staggered fermions in a lattice of volume 8^3, inverse coupling $\beta = 5.4$, quark mass $m = 0.03$ and $\mu = 0$, and two different temporal extents, $N_\tau = 2$ and $N_\tau = 16$. These simulations used one random vector in the evaluation of the fermionic drift. Small tests with 2 and 4 random vectors have not resulted in a cost-effective improvement, especially in the unitarity norm. As is common practice also in Monte Carlo simulations, we have not inverted the Dirac matrix at every single Langevin step, but at every four steps.

Our simulations were set to run for 1100 Langevin time units. In the cases where only gauge cooling was applied, the unitarity norm rose very quickly and instabilities set in. This has severely limited the statistics available, as evident in figs. 3.25 and 3.26. They show the Langevin time evolution of the real part of the chiral condensate for the studies with dynamical stabilisation combined with gauge cooling and just with gauge cooling for $N_\tau = 2$ and $N_\tau = 16$, respectively. Agreement between both cases

can be seen for small Langevin time, with the DS simulations maintaining their trends for the full stipulated simulation time.

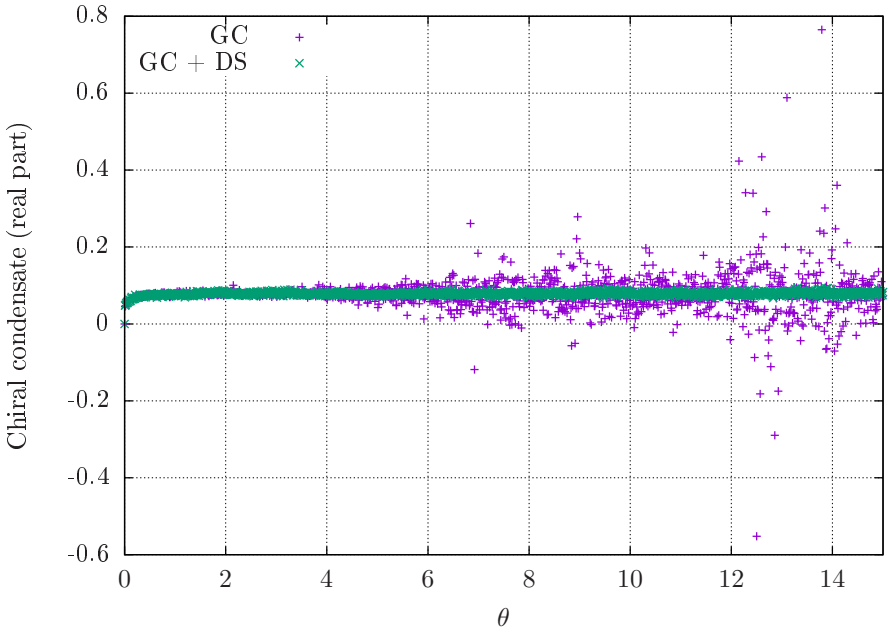

Figure 3.25: Langevin time evolution of the chiral condensate, in a lattice of $8^3 \times 2$, $\beta = 5.4$, $m = 0.03$ and $\mu = 0$, of simulations using only gauge cooling and gauge cooling with dynamic stabilisation. Large fluctions develop on the former.

Despite the window where gauge cooling tests gave reliable results being small, the agreement with dynamic stabilisation is clear. We conclude that DS can be applied to QCD with fully dynamical quarks, and is able to control the rising unitarity norm caused by the small number of random vectors used to calculate the fermionic drift.

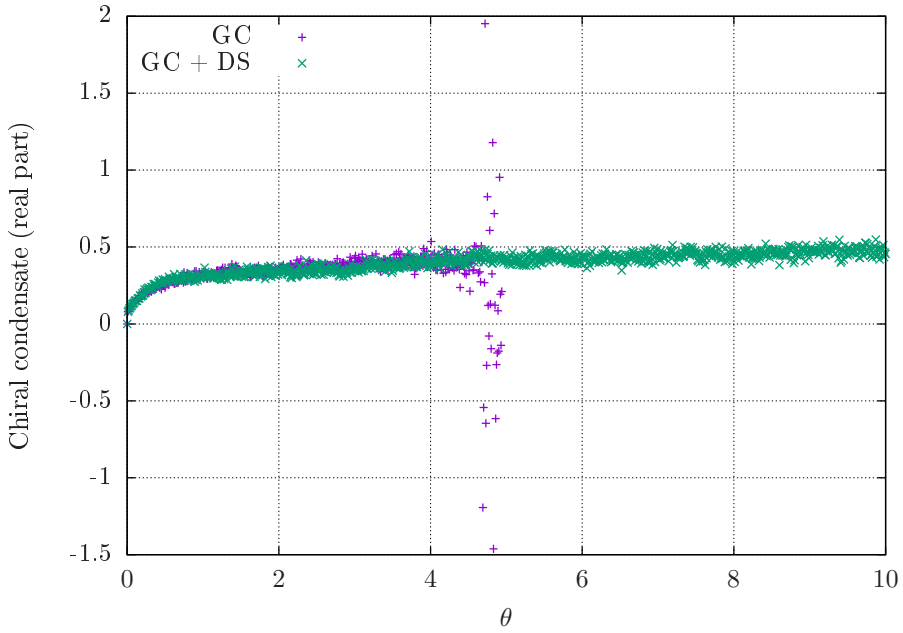

Figure 3.26: Langevin time evolution of the chiral condensate, in a lattice of $8^3 \times 16$, $\beta = 5.4$, $m = 0.03$ and $\mu = 0$, of simulations using only gauge cooling and gauge cooling with dynamic stabilisation. The former has stopped around $\theta \sim 5$.

Chapter 4

Lattice results using complex Langevin

We have applied the techniques described in chapter 2 and the tests of chapter 3 to different studies of lattice QCD and its heavy-dense approximation at finite temperature and chemical potential. We present our results in this chapter. Preliminary results from the studies discussed in sections 4.1 and 4.2 have appeared in [111–114], with the final results in [50]. Results from the subsequent sections will appear in [115].

The first section concerns the exploration of the phase diagram of HDQCD using the first order discretisation algorithm of complex Langevin combined with gauge cooling. We have done multiple simulations of this theory at a variety of temperatures and chemical potentials for three different lattice volumes. By analysing statistical quantities of the Polyakov loop, such as average, susceptibility and Binder cumulant we have been able to determine the boundary between the different phases of the heavy-quark approximation of QCD, as sketched in fig. 1.2. A polynomial fit to the critical temperature as a function of the chemical potential is also provided.

Afterwards, we discuss instabilities in the Langevin evolution encountered during the study of the HDQCD phase diagram. These have been discovered to be related to large unitarity norms—i.e., large explorations of the $\mathrm{SL}(3,\mathbb{C})$ manifold. We have studied how cutting off data after the norm becomes sufficiently large from the analyses can improve them and allow for reliable results.

Section 4.3 discusses the application of the second order discretisation scheme, combined with gauge cooling and dynamic stabilisation, to the determination of the critical chemical potential of HDQCD as a function of the hopping parameter, κ, for low temperatures ($T < 100$MeV). This analysis again uses the Binder cumulant of the Polyakov loop. We have concluded that, for these low temperatures, the critical chemical potential can be described by linear fits in κ and T.

Using the information obtained in the studies of sections 4.1 to 4.3, we revisit the phase diagram of HDQCD in sec. 4.4. In this second study, we have employed the second order algorithm and dynamic stabilisation in a scan of the $T-\mu$ plane focussing on points around the boundary line. For the thermal transition, a qualitative analysis was used, due to the coarseness in the temperature axis. A quantitative result was possible when studying the transition in the μ direction for fixed temperatures. We fitted the data with a function that shows the quartic dependency of T on μ, for higher temperatures, and the linear behaviour, close to $T = 0$.

We close this section with our preliminary results on the use of Dynamic Stabilisation to simulate QCD with dynamical quarks, at finite chemical potential, on the lattice. We have restricted our studies to the high temperature region, where errors stemming from the inversion of the Dirac operator are more easily under control. At our lowest temperature we see large fluctuations developing after a certain chemical potential, where the fermion matrix becomes harder to invert. Our studies concerned varying the chemical potential until the saturation region, for four different temperatures, and analysing the effects on chiral symmetry, via the chiral condensate, and on the centre symmetry, via the Polyakov loop.

4.1 HDQCD phase diagram

Our objective in this study was an exploration of the phase diagram of HDQCD, discussed in section 1.2.5. As said there, this model shares interesting features with QCD with fully dynamical quarks, such as the sign problem. On the other hand, because its fermion determinant is simpler, it is much cheaper to simulate, and offers a good testing ground for complex Langevin in a more realistic case.

We have investigated the phase diagram of QCD in the limit of heavy quarks via direct simulation of points in the temperature–chemical potential plane. We have em-

ployed complex Langevin simulations with the first order algorithm, a variable number of gauge cooling steps applied between consecutive Langevin updates and adaptive Langevin step size. The reliability of this study has been tracked by monitoring the unitarity norm and the distribution of the observables, and by making comparisons with reweighting, when applicable. When analysing the data, we have imposed our unitarity norm cutoff of 0.03, meaning that all data points after this threshold was reached were omitted from the analysis. Larger unitarity norms will be discussed in the next section.

We have taken an approach of fixed lattice spacing, a, in this study, with the various numerical parameters, such as gauge coupling β, lattice volume V, hopping parameter κ, number of flavours N_f and critical chemical potential μ_c^0, described in table 4.1.

$\beta = 5.8$					$V = 6^3, 8^3, 10^3$					$a \sim 0.15$ fm				
$\kappa = 0.04$					$N_f = 2$					$\mu_c^0 = 2.53$				
N_τ	28	24	20	16	14	12	10	8	7	6	5	4	3	2
T [MeV]	48	56	67	84	96	112	134	168	192	224	268	336	447	671

Table 4.1: Parameters used in our study of the HDQCD phase diagram. The chemical potential μ is varied from 0 to $1.3\mu_c^0$, with $\mu_c^0 = -\ln(2\kappa)$. The lattice spacing was set in [116] using the gradient flow [117–119].

A wide range of chemical potentials and temperatures has been covered, with a total of 880 ensembles with different combinations of N_τ and μ, for each of our three volumes. The lattice spacing has been determined in [116] using the gradient flow [117–119], and has been used to convert the temperature into physical units. The fixed lattice spacing approach provides good coverage of the phase diagram for low temperatures, with fixed artefacts, but gives poor coverage at high temperatures.

Figures 4.1 and 4.2 show the average quark number density $\langle n \rangle$ and symmetrised Polyakov loop $\langle P^s \rangle = \frac{1}{2}\langle P + P^{-1} \rangle$, respectively, as functions of the chemical potential and temperature for a spatial volume of 10^3. We show interpolated surfaces to guide the eye, and black points, representing the individual simulations, on each figure. We have expressed the chemical potential in units of μ_c^0 and used the lattice spacing to write the temperature in physical units.

Figure 4.1 shows a sharp rise in the quark density to the saturation value[1] of $n_{\text{sat}} = 12$ around $\mu = \mu_c^0$ for low temperatures. At higher temperatures, this be-

[1] $n_{\text{sat}} = N_{\text{spins}} \times N_c \times N_f = 2 \times 3 \times 2 = 12$.

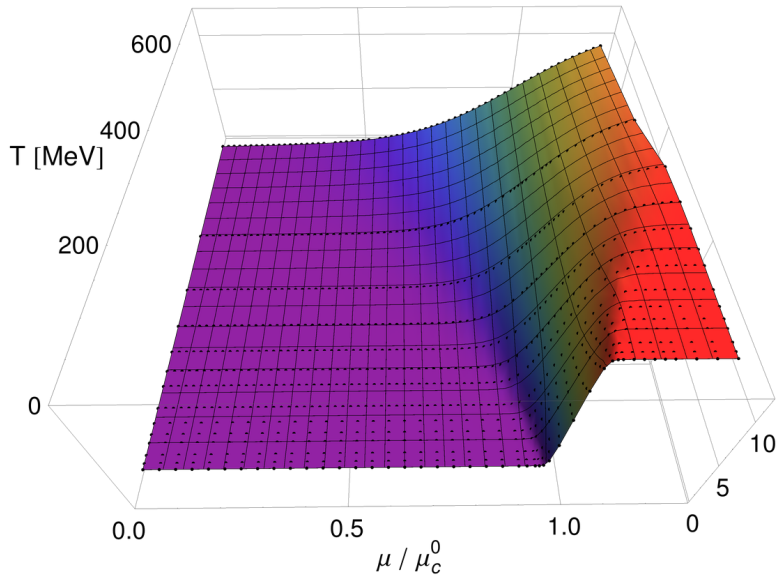

Figure 4.1: Average quark number density $\langle n \rangle$ as function of the temperature and chemical potential for a volume of 10^3. In addition to the black points, representing simulations' results, we show an interpolated surface to guide the eye.

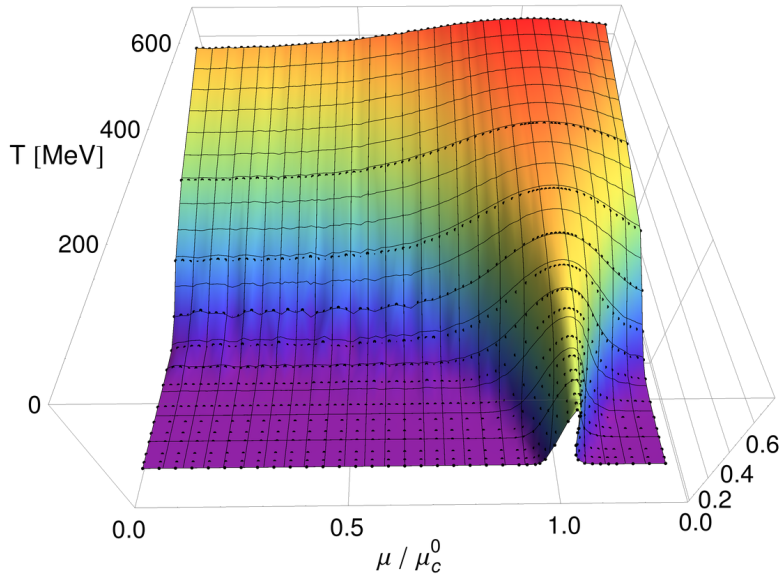

Figure 4.2: Average symmetrised Polyakov loop $\langle P^s \rangle$ as function of the temperature and chemical potential for a volume of 10^3. In addition to the black points, representing simulations' results, we show an interpolated surface to guide the eye.

haviour becomes smoother. Due to the quark number susceptibility at $\mu = 0$ being exponentially suppressed for heavy quarks, the quark number density only rises very slowly as μ increases from zero. The region where $\mu \gg \mu_c^0$ is a lattice artefact, with the lattice entirely filled with quarks, and the Pauli exclusion principle forbidding further addition of quarks.

The Polyakov loop of fig. 4.1 exhibits both the transition to high densities, driven by quark dynamics, and the deconfinement transition, drive by gluonic dynamics. In the saturation region we see the Polyakov loop dropping to zero at low temperatures. This can be understood in terms for the particle-hole symmetry of the HD approximation, described in section 1.2.5. We then expect the Polyakov loop to be symmetric around μ_c^0. At high temperatures, the thermal energy is sufficient to deconfine the quarks and we see a non-zero Polyakov loop for all chemical potentials.

In figs. 4.3 and 4.4, we show the quark density and Polyakov loop susceptibilities, where outlines of the phase transitions can be seen. Note that, to improve visibility, in both figures the dominant peaks are not shown. One of the peaks not displayed in the figure has a negative value for the susceptibility. This happens because of the complex nature of the observables, whose real and imaginary parts cannot be studied separately, as that would break holomorphy. In principle, the susceptibilities can be used to determine the phase boundary. However, a better signal has been obtained

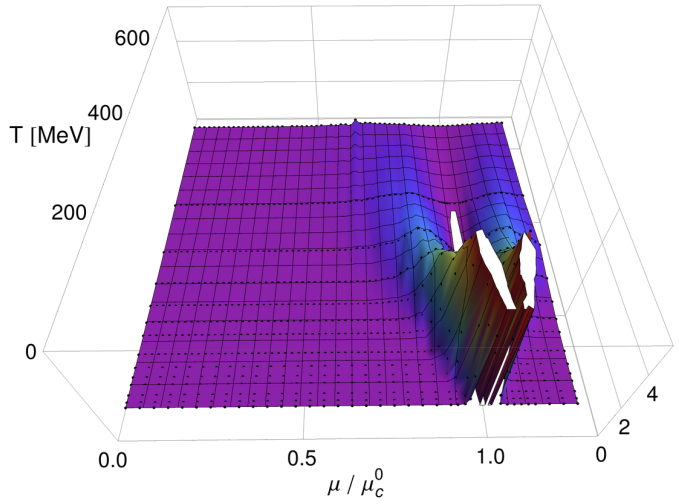

Figure 4.3: Quark density susceptibility as function of the temperature and chemical potential for a volume of 10^3. Peak heights have been cut, resulting in white plateaus.

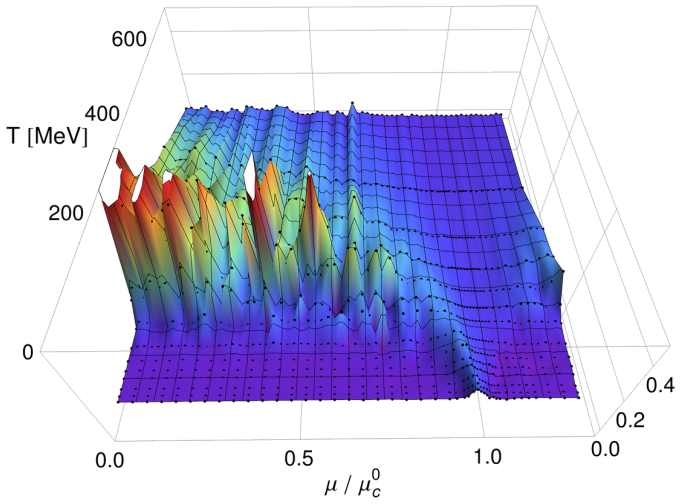

Figure 4.4: Quark number susceptibility as function of the temperature and chemical potential for a volume of 10^3. Peak heights have been cut, resulting in white plateaus.

by employing the Binder cumulant, introduced in chapter 3,

$$B = 1 - \frac{\langle \mathcal{O}^4 \rangle}{3\langle \mathcal{O}^2 \rangle^2}\,, \tag{4.1}$$

where $\mathcal{O}$ represents a generic observable.

We have studied the Binder cumulant of the symmetrised Polyakov loop P^s, shown in fig. 4.5. The two phases, confined, with $\langle P^s \rangle = 0$, and deconfined, with $\langle P^s \rangle \neq 0$, are clearly separated. Here it shows that our fixed lattice spacing approach does not have sufficient resolution to determine the thermal transition at high temperature with precision. On the other hand, the coverage of the parameter space at low T is good, and the transition can be easily identified, with a clear phase boundary seen to emerge.

In order to identify the parameters for which the phase transition occurs, we looked for the points in the $T - \mu$ plane where the Binder cumulant equalled the mid-point between the characteristic values for each phase—$B = 1/3$. The results are shown in fig. 4.6. We have used half of the distance between the neighbouring points, in both the T and μ directions, as estimates for the uncertainties. As previously mentioned, the number of points in the temporal direction being an integer limits the resolution in temperature, causing the thermal transition to suffer from larger discretisation effects.

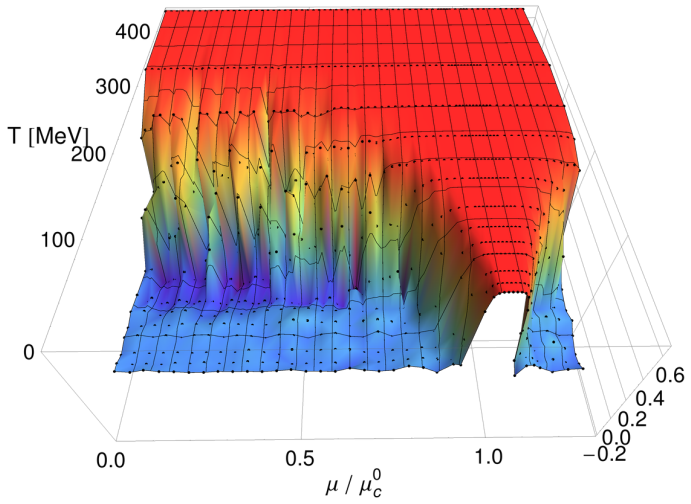

Figure 4.5: Binder cumulant for the symmetrised Polyakov loop as function of T and μ.

Conversely, no such limitations apply to the chemical potential; the transition to higher densities can be mapped with much more precision.

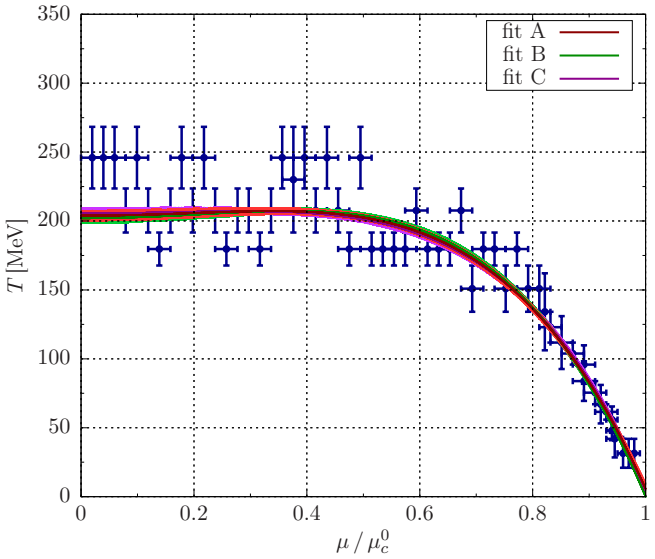

Figure 4.6: Estimates of the phase boundary of HDQCD, with three fitted functions being compared.

A description of the transition (critical) temperature as a function of the chemical potential can be obtained by fitting estimates for $T_c(\mu)$ to a number of fitting functions.

We have used the notation

$$x = \left(\frac{\mu}{\mu_c^0}\right)^2 , \tag{4.2}$$

and considered an expansion around $x = 0$, i.e.,

$$\text{Fit A:} \quad T_c(\mu) = \sum_{k=0}^{n} a_k x^k , \tag{4.3}$$

where we have taken into account that $T_c(\mu)$ is an even function of μ [29]. Alternatively, due to our lattice setup, the transition can be better determined with an expansion around $x = 1$, instead of $x = 0$. Given that, in HDQCD, $T_c(\mu_c^0) = 0$, we have also considered a power series

$$\text{Fit B:} \quad T_c(\mu) = \sum_{k=1}^{n} b_k (1-x)^k . \tag{4.4}$$

The coefficients $\{a_k\}$ and $\{b_k\}$ can be trivially related, provided that the fitted a_k obey $\sum_k a_k = 0$. Finally, to account for the non-analytic behaviour described by the Clausius-Clapeyron relation [120], $\partial T_c(\mu)/\partial\mu \to \infty$ at $\mu = \mu_c^0$, we have fitted

$$\text{Fit C:} \quad T_c(\mu) = c_0 (1-x)^\alpha + \sum_{k=1}^{n} b_k (1-x)^k , \tag{4.5}$$

with $0 < \alpha < 1$.

Figure 4.6 shows fits A, B and C for $n = 2$ applied to the data from our largest volume, $V = 10^3$. Because our lowest temperature is not close enough to zero, the non-analytic behaviour around $\mu \approx \mu_c^0$ is not evident. Forcing α to be between zero and one resulted in it being as close to 1 as possible with single-precision numbers; it did not yield any new information and we did not study fit C any further. The fit coefficients for A and B, together with the corresponding reduced χ^2 value are shown in tables 4.2 and 4.3 for $n = 2$. Fits A and B are seen to be compatible. Furthermore, this indicates that $T_c(\mu_c^0) = 0$ emerges naturally from the data. It is worth noting that a rough estimate for $T_c(\mu = 0)$, in MeV, is given by $b_1 + b_2 \sim a_0$, which also sets the scale of the coefficients.

We compare three different polynomials for fit B in fig. 4.7, with $n = 2$, 3 and 4, for $V = 10^3$. Using higher-order polynomials resulted in almost identical curves as

	fit A, $n = 2$			
V	a_0	a_1	a_2	χ^2_{red}
6^3	276.9 (7.2)	7.4 (33.7)	-283.4 (31.8)	0.85
8^3	216.4 (5.0)	86.0(25.5)	-305.8 (24.8)	1.51
10^3	203.9 (4.3)	58.9 (23.1)	-257.1 (23.2)	1.62

Table 4.2: Fit parameters and reduced χ^2 for fit A, as defined in eq. (4.3), used to describe $T_c(\mu)$ for three spatial volumes.

	fit B, $n = 2$		
V	b_1	b_2	χ^2_{red}
6^3	564.3 (15.2)	-287.8 (19.2)	0.83
8^3	507.8 (12.8)	-289.9 (15.7)	1.49
10^3	481.4 (12.4)	-279.3 (15.0)	1.62

Table 4.3: Fit parameters and reduced χ^2 for fit B, as defined in eqs. (4.4), used to describe $T_c(\mu)$ for three spatial volumes.

that of $n = 2$, and a χ^2 analysis showed that the latter best represented the data. Therefore, adding more parameters does not improve the fitting. A study of fit B with

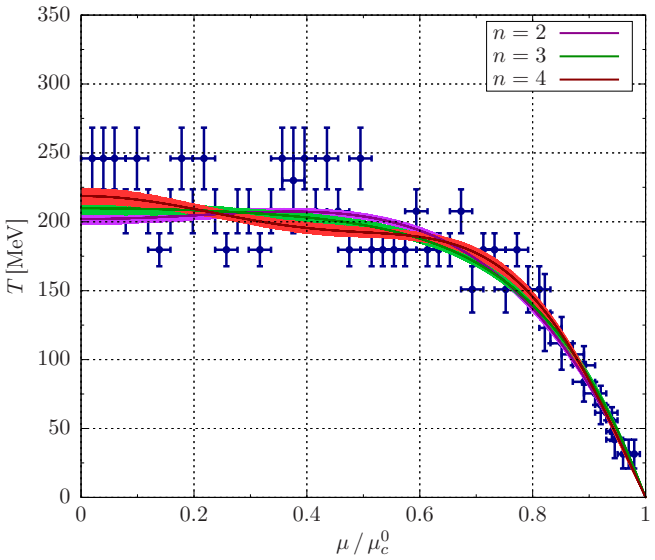

Figure 4.7: Estimates of the phase boundary of HDQCD, with three different orders for fit B being compared.

$n = 2$ for all three volumes studied here is shown in fig. 4.8. Evident finite-size effects can be seen, in particular for the smallest simulation volume of 6^3. A considerably smaller trend is visible for the two larger volumes. However, as previously discussed, our

dominant limitation comes from discretisation effects at high temperature. In principle,

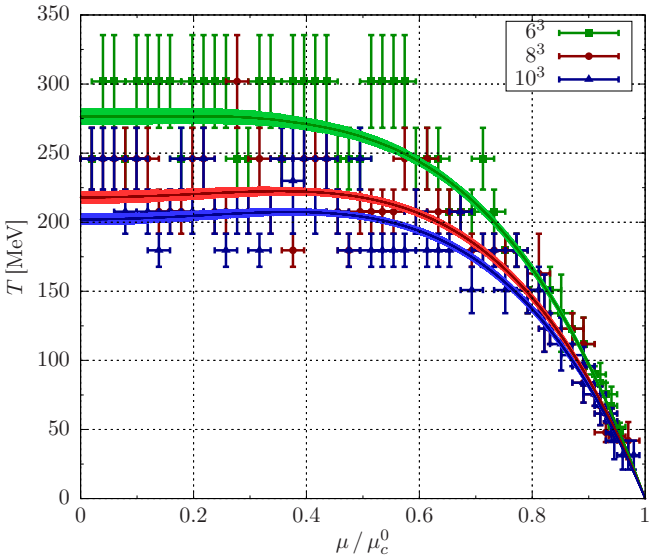

Figure 4.8: Volume dependence of the phase boundary of HDQCD, using fit B and $n = 2$.

the Binder cumulant is suitable for determining the order of the phase transition, since its value at the transition point depends only on the universality class [107]. This would require, however, a deeper volume dependence analysis, with more precise determination of $T_c(\mu)$ throughout the $T - \mu$ plane, and smaller uncertainties.

4.2 Instabilities

As briefly mentioned in section 3.4, and exemplified in fig. 3.12, complex Langevin simulations can suffer from instabilities. In this section we focus on those encountered when simulating the phase diagram of HDQCD, described in the previous section. These are generic problems, however, not exclusive to the parameters we have used. Our goal is to show they are related to convergence to wrong limits, when the unitarity norm becomes too large, and that keeping the latter under control improves the analysis of the results.

Typically, the instabilities seen in complex Langevin result in a widening of the distribution of observables during the simulations, significantly affecting susceptibilities and other quantities. One can conclude that these wider distributions do not reflect

the original theory, based on comparisons with reweighting [121] and on the formal justification discussed in section 2.3.1 and references therein.

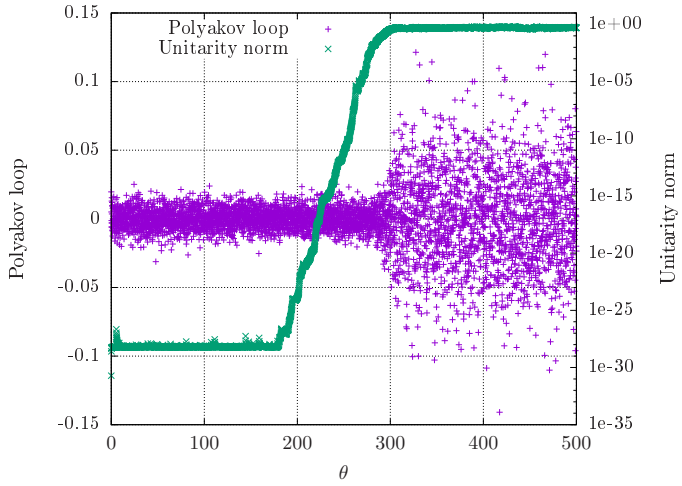

Figure 4.9: The average Polyakov loop and unitarity norm as functions of the Langevin time at a low temperature setup ($N_\tau = 20$ and $\mu = 0.5$).

Figure 4.9 shows the real part of the Polyakov loop and the unitarity norm as functions of the Langevin time in a low temperature ($N_\tau = 20$) setup. Similar to fig. 3.12, we observe two distinct regions, the first characterised by small unitarity norm and controlled fluctuations, followed by large unitarity norm and fluctations afterwards. Unlike in fig. 3.12, though, in fig. 4.9 we do not see a change in the average value for the Polyakov loop. Rather, in this case, the instability affects its fluctuations.

Table 4.4 has a further analysis of the data shown in figs. 3.12 and 4.9. We have studied the expectation value of the Polyakov loop, its susceptibility, χ_P, and its Binder cumulant, for each of the intervals where the Polyakov loop fluctuates consistently around a certain value. We also present results obtained using reweighting.

The table shows that, within statistical errors, in the first interval the observables are compatible with reweighting, while in the second one they are not. At low temperature the susceptibility is a more sensitive measure of accuracy, since the apparent agreement around $B \sim 0$ for the entire interval is a reflection of $\langle P \rangle \sim 0$. In figs. 4.10 and 4.11, we compare histograms for the low and high temperature scenarios, respectively, with Gaussian fits added to guide the eye. The distributions are broader, with larger tails, for the regions with higher unitarity norm. Additionally, at high tempera-

$N_\tau = 20,\ \mu = 0.5$

	$100 < \theta < 250$	$330 < \theta < 500$	Reweighting
$\langle P \rangle$	0.00009(65)	−0.0002(44)	0.000032(22)
χ_P	0.0542(68)	0.0510(1796)	0.055473(68)
B	0.01(17)	−22(207)	0.0013(19)

$N_\tau = 4,\ \mu = 0.7$

	$20 < \theta < 60$	$100 < \theta < 500$	Reweighting
$\langle P \rangle$	0.2043(53)	0.0069(115)	0.202717(66)
χ_P	0.37(17)	1.44(73)	0.37993(17)
B	0.6544(57)	−0.6332(8105)	0.65487(18)

Table 4.4: Analysis of the real part of the Polyakov loop, its susceptibility, and its Binder cumulant, for the data presented in figs. 3.12 and 4.9. In each case, the two intervals correspond to the regions where the Polyakov loop fluctuations are consistent around a given value. Reweighting results are added for comparison.

ture, there is a shift of the mean value. Our conclusion is that the regions with small unitarity norm lead to adequate results, whilst those with larger norms do not.

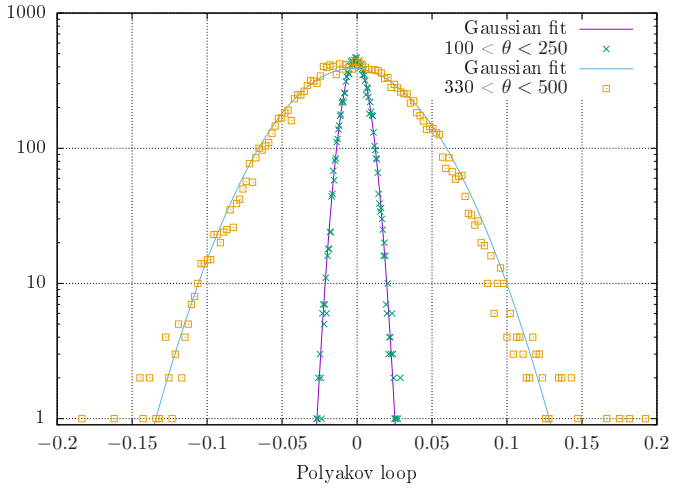

Figure 4.10: Histograms of the real part of the Polyakov loop before and after the unitarity norm becoming too large, for the simulation presented in fig. 4.9.

In all situations where the behaviour described above has been observed, the widening of distributions has coincided with severe changes in the unitarity norm. Reducing the average Langevin step size used in the adaptive algorithm has not helped to prevent this transition from occurring. The inability to control the unitarity norm of complex

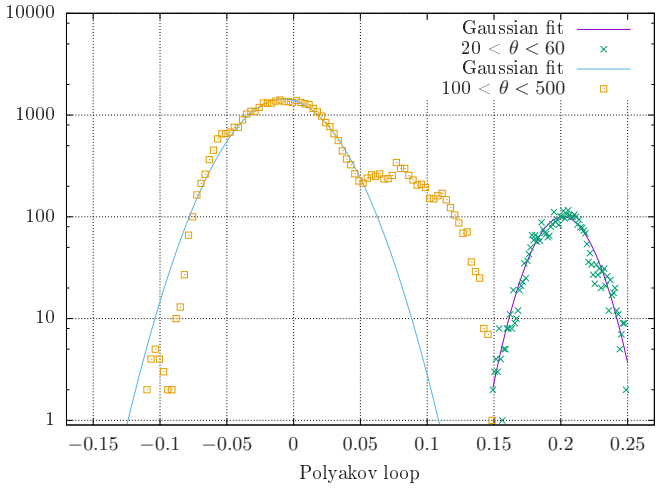

Figure 4.11: Histograms of the real part of the Polyakov loop before and after the unitarity norm becoming too large, for the simulation presented in fig. 3.12.

Langevin simulations on coarser lattices using only gauge cooling has been already noted in [98], and mentioned in section 3.4.

We have carried out additional simulations with larger gauge couplings, $\beta = 6.0$ and 6.2, to check the behaviour close to the continuum limit. We show the real part of the Polyakov loop for setups identical to those described in table 4.4 and in fig. 4.12, with results for low temperature ($N_\tau = 20$) on the left hand side and high temperature ($N_\tau = 4$) on the right hand side. The combination of finer lattices and low temperature produced unitarity norms that remain practically at 0 for the entire simulation. There is, however, an exponential increase in the norm for higher temperatures, albeit with smaller exponent than at $\beta = 5.8$. After the unitarity norm becomes too large, substantial fluctuations develop and skirts emerge. This behaviour is visible in fig. 4.13, where we present histograms for the high-temperature simulations for the two larger β values on a 10^3 lattice. Consequently, we conclude that instabilities can still develop in finer lattices, except that they set in later, for high temperature, or only appear beyond the length of the Langevin trajectory, for low temperature.

The Langevin time when the unitarity norm begins rising varied considerably among the simulations shown in the preceeding section. Fortunately, that happened sufficiently long after the thermalisation stage, in most cases, to leave enough data points to allow for statistical analyses. However, we found different uncertainties, including auto-

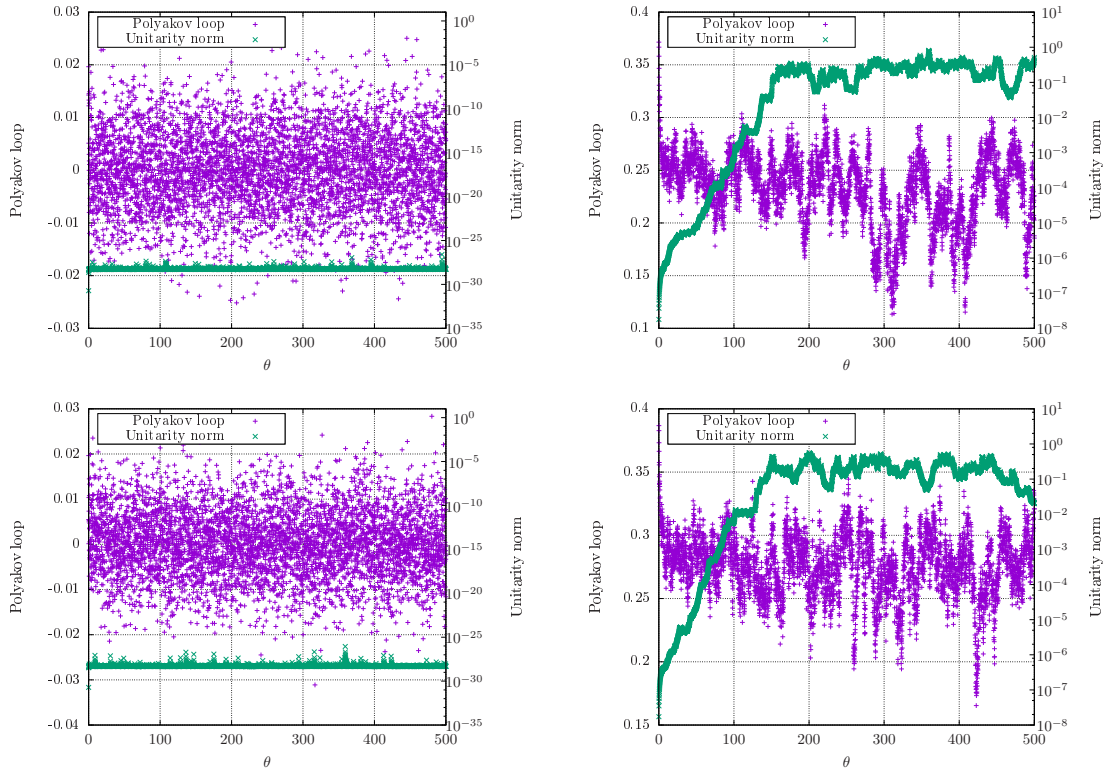

Figure 4.12: Real part of the Polyakov loop P and unitarity norm d_2 for a larger gauge coupling of $\beta = 6.0$ (top) and $\beta = 6.2$ (bottom) with low temperature ($N_\tau = 20, \mu = 0.5$, left) and high temperature ($N_\tau = 4, \mu = 0.7$, right) on a 10^3 lattice.

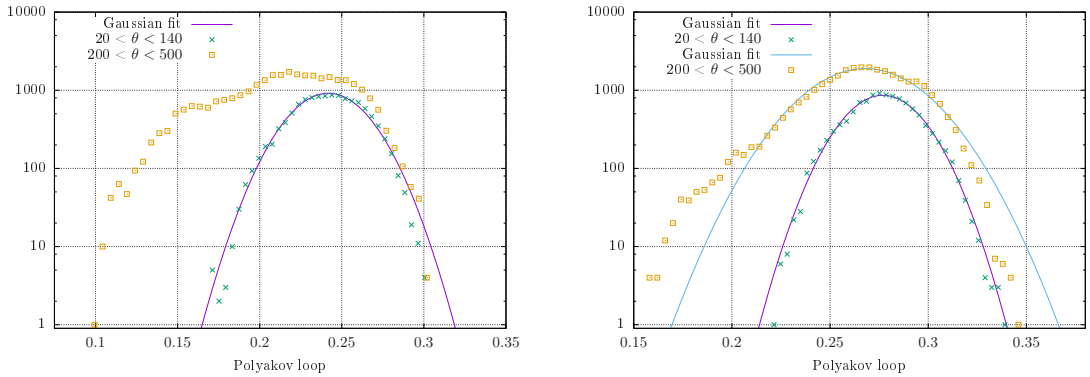

Figure 4.13: Histograms of the real part of the Polyakov loop before and after the rise of the unitarity norm, for the larger gauge couplings of $\beta = 6.0$ (left) and 6.2 (right), at high temperature ($N_\tau = 4, \mu = 0.7$).

correlation time [122], in each setup, due to the variable amount of available data among ensembles. We have taken these findings into account by making sure that the results presented for the phase diagram contained only data for which the unitarity norm was smaller than 0.03.

Because of the smaller lattice spacings, simulations with larger gauge couplings require larger lattice volumes to keep the physical volume constant. We have found that a gauge coupling of $\beta = 5.8$ offers a good compromise between the ability to extract reliable information about the phase diagram and the simulation cost.

A study of the robustness and sensitivity of the physical observables, with respect to changes in the cutoff imposed to the unitarity norm, is shown in fig. 4.14, with graphs of the average quark density and Polyakov loop as functions of the maximum unitarity norm allowed, for the high temperature setup discussed previously ($N_\tau = 4$, $\mu = 0.7$ and volume 10^3). It is clear that the observables are stable and independent of the cutoff over a broad range, up to ~ 0.5. The widening of distributions happens at the same Langevin time as the transition seen for larger values of the cutoff. We remark that the change in the statistical uncertainties, shown in the insets, stems from the decrease in the number of available data points.

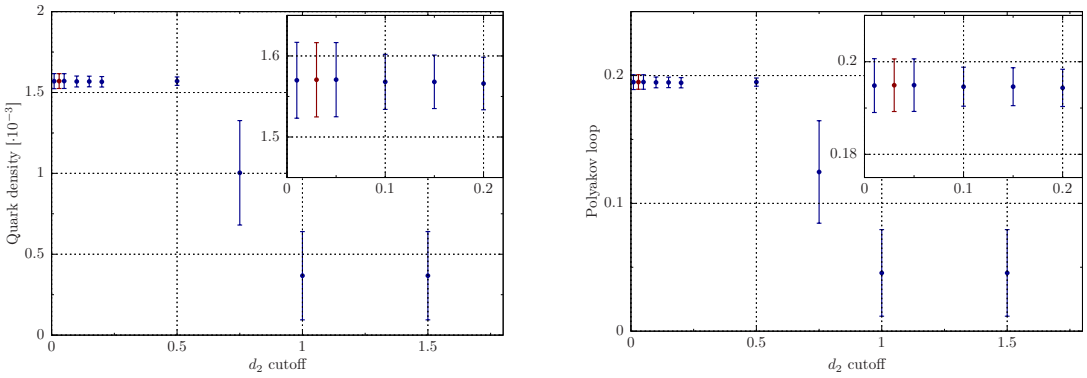

Figure 4.14: The quark density (left) and the Polyakov loop (right) as a function of the cutoff imposed on the unitarity norm d_2 for $N_\tau = 4, \mu = 0.7$ on a 10^3 lattice. For small unitarity norms, $d_2 < 0.5$, the observables are independent of the cutoff. The red point indicates the value chosen in this study. The insets focus on the region of smaller cutoffs.

Our conclusion is that observables are independent of the cutoff imposed on the unitarity norm, provided that it is below $O(0.1)$. Throughout this thesis, unless otherwise stated, we have conservatively chosen a cutoff of 0.03, to guarantee a safe distance

from the regions affected by instabilities.

4.3 Critical chemical potential at different hopping parameters for HDQCD

In this section, we apply the second order algorithm described in section 2.2.3, in combination with dynamic stabilisation, discussed in 2.3.3, to a series of simulations of HDQCD. Our objective is to employ these two methods, in conjunction with the known fact that at zero temperature the critical chemical potential for heavy quarks is analyticaly known, to extract further information from HDQCD.

At $T = 0$ we expect to see the transition to higher densities at

$$\mu = \mu_c^0 \equiv -\ln(2\kappa)\,, \tag{4.6}$$

with κ being the hopping parameter. By studying the Binder cumulant of the Polyakov loop as a function of μ for various values of κ and N_τ, we expect to find how the critical chemical potential at finite temperature relates to μ_c^0.

Our simulations have been done at a volume of 8^3 and gauge coupling $\beta = 5.8$, at temporal extents of $N_\tau = 16$ and 32. The chemical potentials used were fractions of the relevant critical value at zero temperature for each value of the hopping parameter. We have used $\alpha_{\rm DS} = 10^3$ and one gauge cooling step between Langevin steps to control excursions into $\mathrm{SL}(3,\mathbb{C})$, and an average Langevin step size of $O(10^{-4})$.

Figures 4.15 and 4.16 show the Binder cumulant of the symmetrised Polyakov loop as a function of the chemical potential for several values of the hopping parameter and $N_\tau = 16$ and 32, respectively. In both plots, the horizontal axis has been re-scaled by the corresponding μ_c^0 for each value of κ, and dashed lines have been added to guide the eye. Both phases, confined and deconfined, are easily seen for all simulation setups, as well as the region where the systems transition between them.

As the quark masses increase ($\kappa \to 0$), more energy is required to allow their creation from the vacuum, as clearly shown in the figures. At low, but non-zero, temperatures, the energy available in the environment allows for critical chemical potentials smaller than that of zero temperature for all quarks masses. This results in the curves of fig. 4.16 being shifter towards $\mu/\mu_c^0 = 1$, in comparison to those of fig. 4.15.

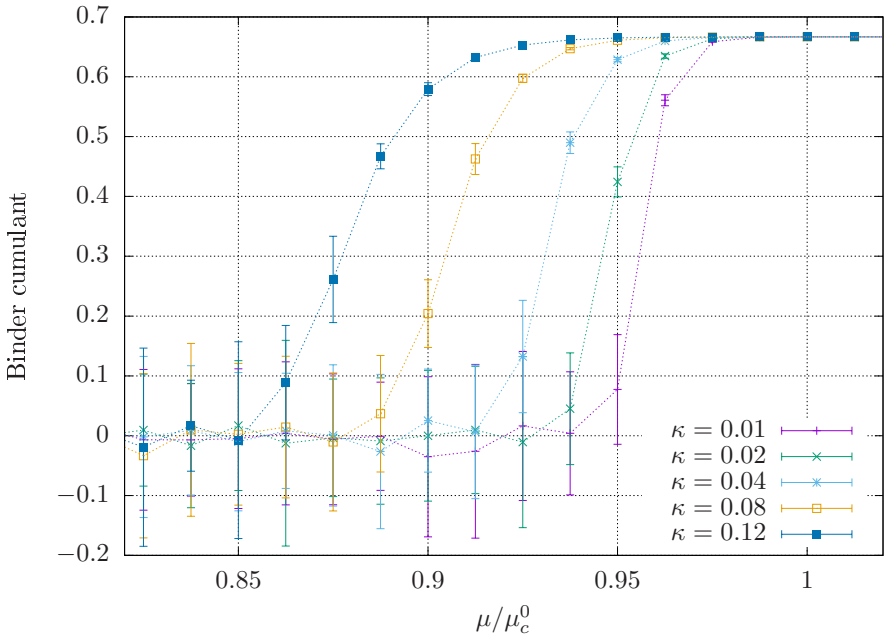

Figure 4.15: Binder cumulant of the Polyakov loop as a function of the chemical potential for various values of the hopping parameter at $N_\tau = 16$. The x-axis has been re-scaled by the relevant μ_c^0 for each κ.

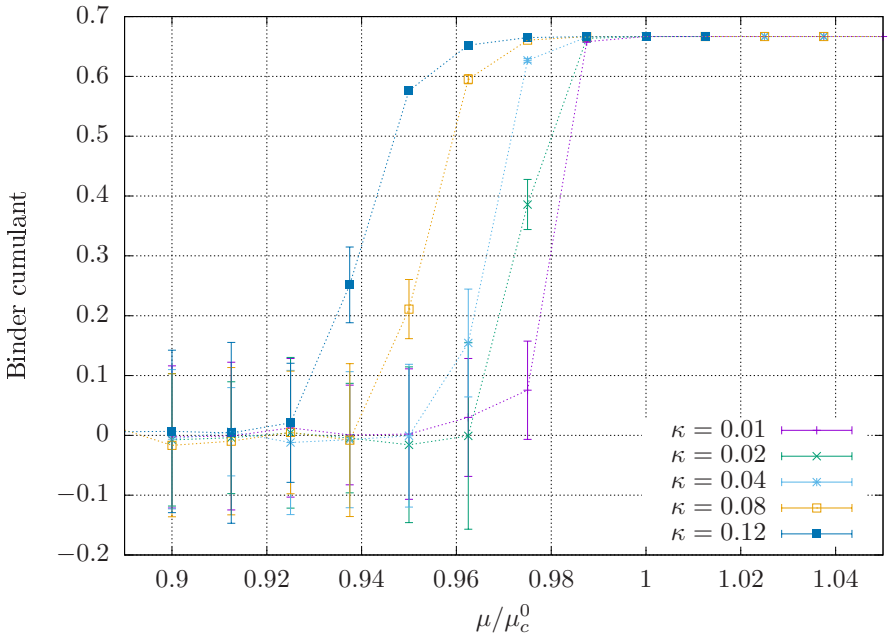

Figure 4.16: Binder cumulant of the Polyakov loop as a function of the chemical potential for various values of the hopping parameter at $N_\tau = 32$. The x-axis has been re-scaled by the relevant μ_c^0 for each κ.

The critical chemical potential has been studied in a manner similar to that used in section 4.1—namely, by considering the points where the Binder cumulant was equal to the mid-point between the characteristic values for each phase. Again, we use half of the distance between neighbouring points in the μ direction as estimate for the uncertainties. We show the results in fig. 4.17, together with linear fits of the form $\mu_c/\mu_c^0 = a + b\kappa$, the corresponding error bands, and a line representing $\mu_c(T=0) = \mu_c^0$. We have used the lattice spacing, estimated using the gradient flow [117–119], to convert the temperatures into physical units, also shown in the graph. As the temperature gets closer to zero the fitted lines become increasingly more horizontal; i.e., μ_c tends to μ_c^0 as expected.

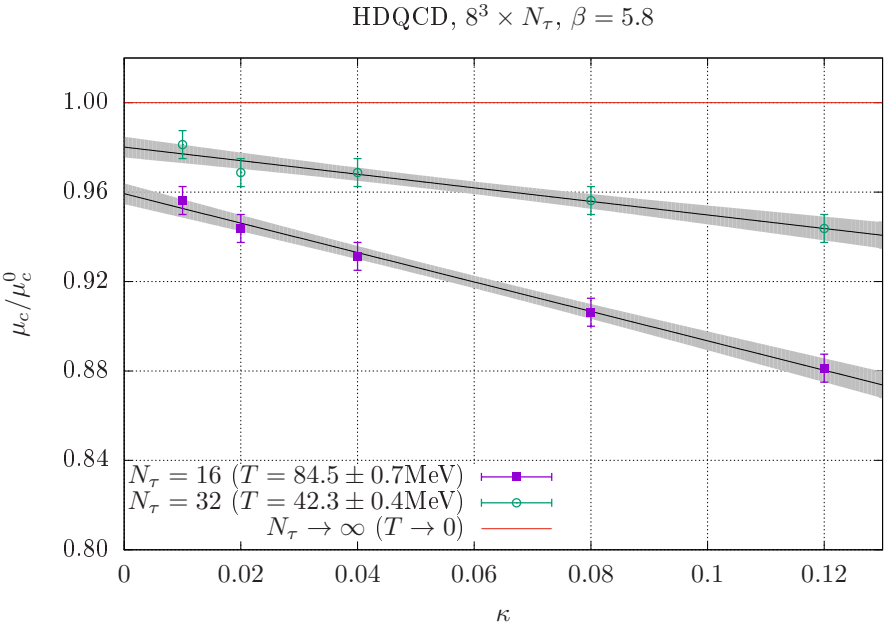

Figure 4.17: The critical chemical potential, in units of that at zero temperature, as a function of the hopping parameter. Linear fits and error bands are also shown.

The graph shown in fig. 4.17, together with the fact that at $T = 0$ we have $\mu_c = \mu_c^0 \equiv -\ln(2\kappa)$, suggests that $\mu_c(\kappa, T)$ can be parametrised as

$$\frac{\mu_c(\kappa, T)}{\mu_c^0} = a(T) + b(T)\kappa\,, \tag{4.7}$$

with $a(T=0) = 1$ and $b(T=0) = 0$. We have fitted $a(T)$ and $b(T)$ from eq. (4.7) with linear functions to the data shown in the figure. Because of the above constraints, only

the functions' slopes had to be fitted. The results are $-0.000582(49)$ and $-0.01192(73)$, for the functions a and b, respectively.

Our studies have covered most of the range of validity of κ, since the hopping parameter is restricted to be smaller than 1/8. It is, then, possible to state that μ_c/μ_c^0 has a linear dependecy on κ, at low temperatures, for all allowed values. The linear dependency on the temperature will be used, in the next section, to choose values of the chemical potential around the phase boundary.

We stress that, strictly speaking, the critical chemical potential is undefined at $\kappa = 0$, as at that point the theory reduces to pure Yang–Mills. At zero temperature, μ_c^0 diverges logarithmically with the hopping parameter, and our results show that the same behaviour is seen at finite T, such that μ_c/μ_c^0 converges to a finite value.

4.4 HDQCD phase diagram revisited

With the information shown in the previous three sections—i.e., the exploratory scan of the HDQCD phase diagram, study of instabilities in complex Langevin simulations, and the search for the critical chemical potential for different hopping parameters—we have performed a second study of HDQCD in the $T - \mu$ plane. We have used the acquired knowledge on the location of the phase transition to do a finer scan around the phase boundary. This included lower temperatures, up to $N_\tau = 88$, and increased density of points in the chemical potential axis using the second order algorithm and dynamic stabilisation.

This study was done on a lattice of spatial volume 8^3, fixed inverse coupling $\beta = 5.8$, quark hopping parameter $\kappa = 0.04$ and DS control parameter $\alpha_{\mathrm{DS}} = 10^3$. All simulations have run for 1500 Langevin time units, with the first 100 excluded for thermalisation. As described before, the fixed β scans led to rather coarse steps in the temperature axis for higher temperatures. Thus, our focus has been on improving the precision of our phase boundary with better simulation techniques and a finer sweep of the chemical potential axis for the density transition and a qualitative study of the thermal one.

Our use of dynamic stabilisation has allowed for faster simulations, since only one step of gauge cooling was necessary, and increased statistics, due to the unitarity norm being under control. The longer simulation times allowed the Polyakov loop to visit

the other minima, when its value was non-zero and the ground state still exhibited a remnant of the $\mathbb{Z}_3$ centre symmetry, which is explicitly broken by the quarks. These situations typically occur at low chemical potential and higher temperature. Owing to the coarseness of our temperature axis, these cases also benefit from a qualitative analysis in terms of histograms, as a crossover phase between broken and symmetric Polyakov loops emerges.

We have divided our analysis in two parts. In the lower μ region we have carried out simulations at fixed chemical potential, varying the temperature. Since the phase boundary line is expected to be almost flat, within our statistical precision, for $0 \leq \mu/\mu_c^0 \lesssim 0.5$, as seen in fig. 4.8. For the intermediate to large chemical potential section of the phase diagram we have employed the same analysis as before, and scanned in the μ directions for fixed temperatures.

For temperatures in the range of 224MeV$\leq T \leq$ 336MeV ($4 \leq N_\tau \leq 6$) some qualitative statements can be made using histograms of the symmetrised Polyakov loop. At $N_\tau = 4$ we expect the system to be in a deconfined phase, whereas at $N_\tau = 6$ the average Polyakov loop should be close to zero. Between them, at $N_\tau = 5$, we can see from the histograms in fig. 4.18 that the Polyakov loop exhibits influence from both trivial and non-trivial vacua at $\mu = 0$. It has almost equal probability of being found on each of them, with the higher peak stemming from two of the non-trivial vacua being on the negative x-axis. This crossover behaviour becomes less pronounced as the chemical potential is increased, shown in fig. 4.19, and completely vanishes, as seen in fig. 4.20, indicating that μ_c has been reached.

Concerning the lower temperature section of the phase diagram ($7 \leq N_\tau \leq 88$, or 15MeV $\leq T \leq$ 192MeV), figure 4.21 shows typical cases, where the Polyakov loop has a single peak that drifts from zero to a non-vanishing value as μ increases. In the figure, with data collected from simulations at $N_\tau = 88$, the top-most and right-most plots show histograms of the real and imaginary parts of the symmetrised Polyakov loop, respectively. The bottom left part displays a scatter plot. This shows similarities and differences between each situation: in both cases the Polyakov loop has a localised distribution, but around very different, μ-dependent, average values. Intermediate values for the chemical potential, between those shown in the figure, exhibit the same behaviour, with the distributions shifting to the right as μ increases. Such a sharp

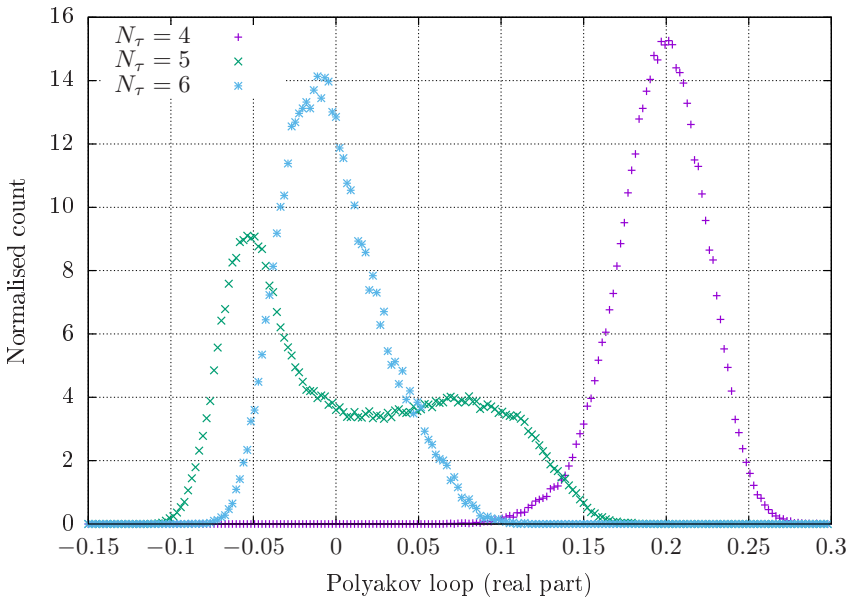

Figure 4.18: Histograms of the Polyakov loop for HDQCD in a lattice with $N_\tau = 4$, 5, and 6, volume 8^3, $\beta = 5.8$, $\kappa = 0.04$ and $\mu = 0$.

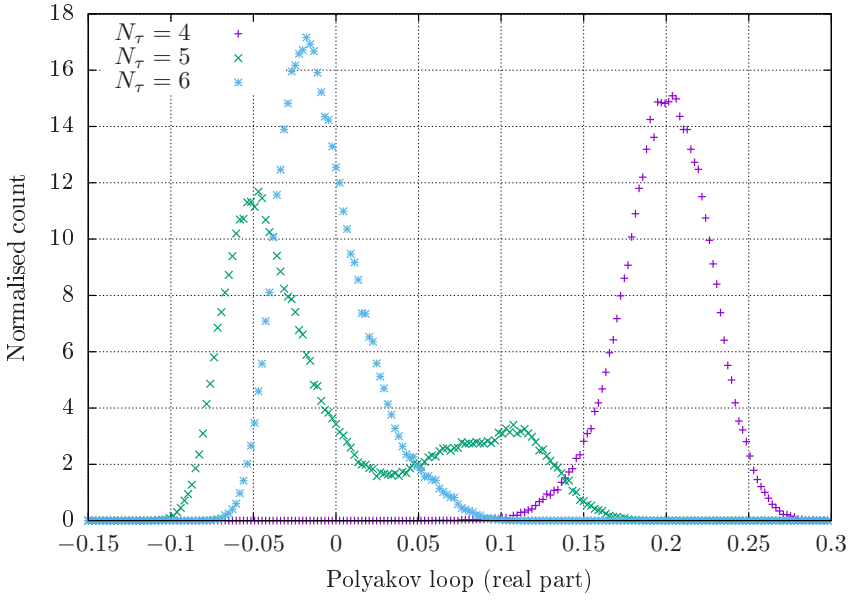

Figure 4.19: Histograms of the Polyakov loop for HDQCD in a lattice with $N_\tau = 4$, 5, and 6, volume 8^3, $\beta = 5.8$, $\kappa = 0.04$ and $\mu = 0.79$ ($\mu/\mu_c^0 = 0.31$).

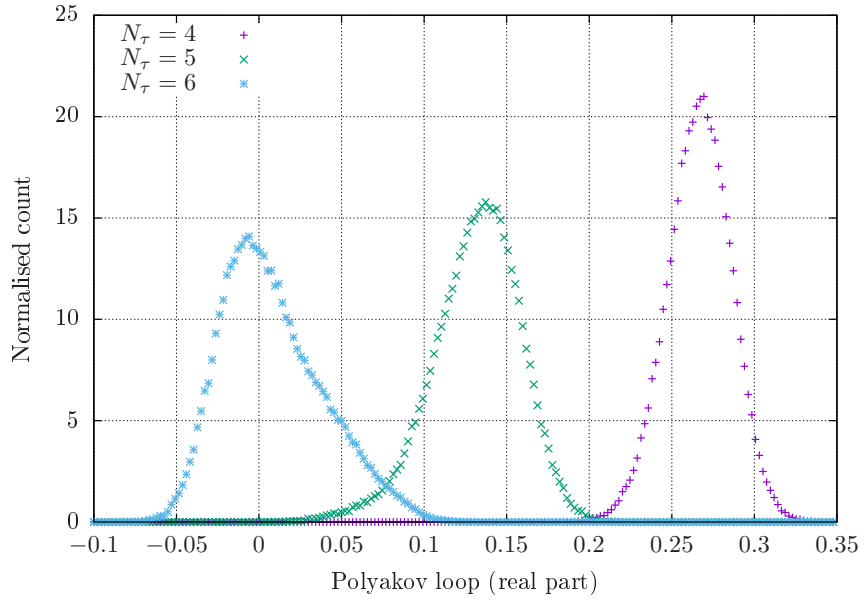

Figure 4.20: Histograms of the Polyakov loop for HDQCD in a lattice with $N_\tau = 4$, 5, and 6, volume 8^3, $\beta = 5.8$, $\kappa = 0.04$ and $\mu = 1.58$ ($\mu/\mu_c^0 = 0.63$).

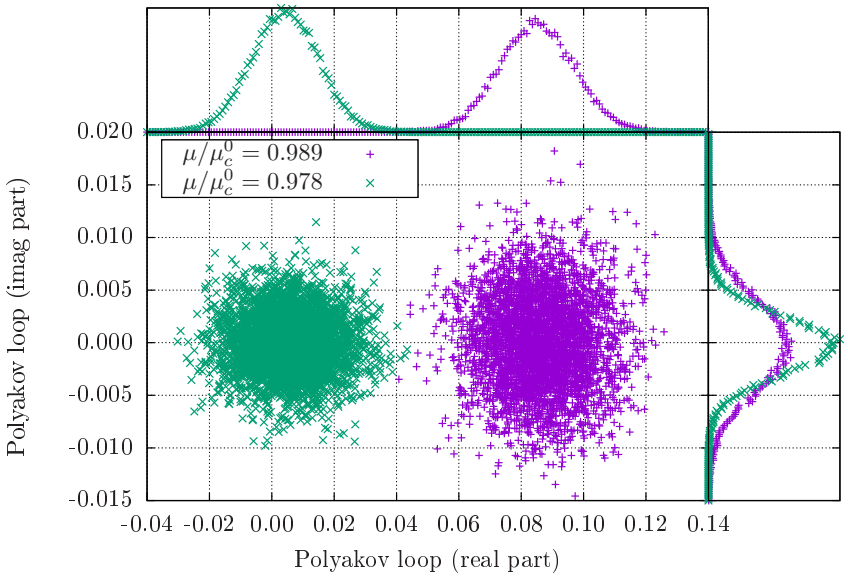

Figure 4.21: Analysis of the symmetrised Polyakov loop for HDQCD at $8^3 \times 88$, $\kappa = 0.04$ and $\beta = 5.8$. The scatter plot and histograms show similar behaviour for both chemical potentials, with different average values.

change from essentially zero $\langle P \rangle$ to non-zero ones is expected due to the low temperature of this simulation and fig. 4.2. Studies at higher temperatures revealed similar behaviour, only with a wider range of chemical potentials needed for the same effect on the Polyakov loop.

Our choices of chemical potentials for this region of the phase diagram were influenced by the linear behaviour observed in section 4.3. We have used the critical chemical potential, μ_c^{fit}, predicted by the linear fit at $\kappa = 0.04$ as function of the temperature, as base point, and used its uncertainty, $\sigma_{\mu_c^{\text{fit}}}$, to choose neighbouring points. They are summarised in table 4.5. For each temperature we have simulated 15 points, from $-7\sigma_{\mu_c^{\text{fit}}}$ to $+7\sigma_{\mu_c^{\text{fit}}}$. These ranges have been chosen to accommodate for potential deviations from the linear behaviour, especially at higher temperatures. We have also included two higher temperatures in this study, where we have used $14\sigma_{\mu_c^{\text{fit}}}$ instead of 7. The extra points, at $N_\tau = 6$ and 5, have the highest temperatures we can reach at our fixed inverse coupling of $\beta = 5.8$ that exhibit chemical potential-dependent phases. At $N_\tau = 5$ the system transitions between the crossover and broken phases, as shown in figs. 4.18, 4.19 and 4.20, whereas $N_\tau = 4$ has deconfined quarks for all chemical potentials and does not provide any new information for this study.

N_τ	T/MeV	μ_c^{fit}	$\mu_c^{\text{fit}}/\mu_c^0$	μ_c^{data}	$\mu_c^{\text{data}}/\mu_c^0$
5	268.44	1.808(39)	0.716(15)	1.347(33)	0.533(13)
6	223.70	1.927(32)	0.763(13)	1.675(33)	0.663(13)
7	191.74	2.013(28)	0.797(11)	1.911(14)	0.7566(55)
8	167.78	2.077(24)	0.8224(96)	2.037(12)	0.8065(48)
10	134.22	2.167(19)	0.8579(77)	2.154(10)	0.8528(40)
12	111.85	2.227(16)	0.8816(64)	2.2160(80)	0.8774(32)
14	95.87	2.269(14)	0.8985(55)	2.2600(70)	0.8948(28)
16	83.89	2.301(12)	0.9112(48)	2.2930(60)	0.9079(24)
18	74.57	2.326(11)	0.9210(43)	2.3195(55)	0.9183(22)
20	67.11	2.3463(97)	0.9289(38)	2.3398(49)	0.9264(19)
32	41.94	2.4136(60)	0.9556(24)	2.4096(30)	0.9540(12)
48	27.96	2.4509(40)	0.9704(16)	2.4483(20)	0.96932(81)
64	20.97	2.4696(30)	0.9778(12)	2.4677(15)	0.97702(59)
88	15.25	2.4849(22)	0.98385(87)	2.4835(11)	0.98328(44)

Table 4.5: Temperatures used in our simulations with the respective projected critical chemical potentials and uncertainties, calculated using the linear fits from section 4.3, and results from simulations.

Since we have done fixed temperature scans, it is reasonable now to parametrise

a	b
$-0.000225(46)$	$0.003154(12)$

Table 4.6: Values for the parameters in eq. (4.8), used to fit the data shown in fig. 4.22.

the phase boundary using the critical chemical potential as function of temperature, $\mu_c(T)$, as opposed to our previous mapping of $T_c(\mu)$. As before, we have used as μ_c the midpoint between the first points whose Binder cumulants were below and above $1/3$, and the error bars came from half of the distance between them. The resulting values for the critical chemical potential, μ_c^{data}, are shown in table 4.5 and fig. 4.22. Based on the quartic polynomial used to fit $T_c(\mu)$ in section 4.1, we have fitted the current data with

$$\frac{\mu_c(T)}{\mu_c^0} = aT + (1 - bT)^{1/4} \ , \tag{4.8}$$

which enforces $\mu_c(T = 0) = \mu_c^0$, and whose parameters are shown in table 4.6. The fourth root is based on the previously observed relation between μ and T, while the linear term has been added to account for the low-temperature behaviour.

The fitted function and curves representing a 1σ confidence interval are also shown in the figure. As expected from the discussion in the previous section, there is a clear linear trend for $\mu_c(T)$ for lower temperatures. Despite the lowest temperature being a third of that used the study of section 4.1, the non-analytic behaviour of the critical temperature at $\mu = \mu_c$ still could not be observed. On the other side of the chemical potential spectrum, a departure from the linear trend is visible for temperatures above 150MeV.

4.5 Finite density QCD with staggered quarks

We have conducted studies of QCD at high temperatures, spanning a wide range of chemical potentials. Our objective was to test dynamic stabilisation applied to lattice QCD with dynamical quarks at finite chemical potential, in a region where the inversion of the fermion matrix is under control. First studies of staggered quarks at finite chemical potentials were carried out in [105], while recent works include [123, 124].

In high temperature scenarios, quarks are deconfined—no Silver Blaze region is expected, with the quark number density increasing for any $\mu > 0$. Additionally, the absence of a Silver Blaze region makes the inversion of the Dirac operator considerably

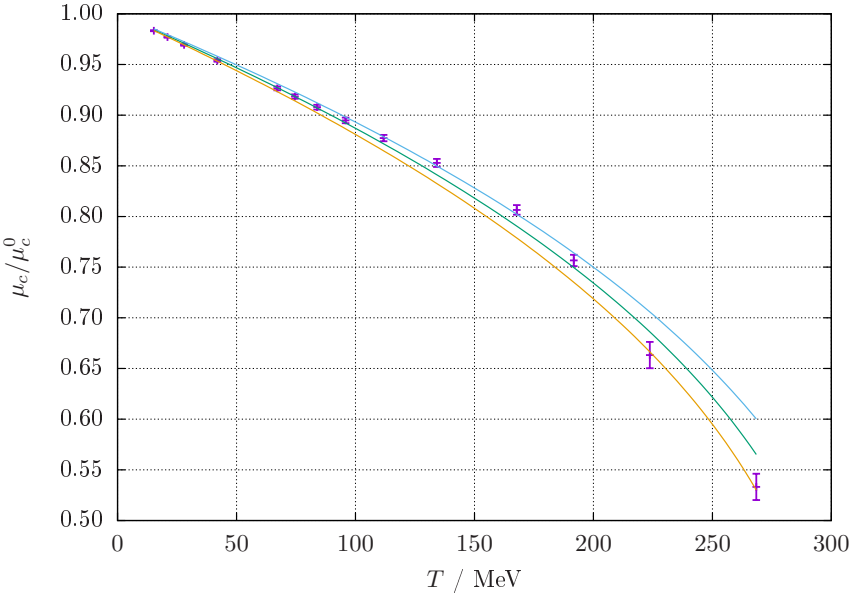

Figure 4.22: The critical chemical potential as a function of the temperature for HDQCD at a volume of 8^3, $\beta = 5.8$ and $\kappa = 0.04$. Linear behaviour can be seen for $T \lesssim 150$MeV. A 1σ confidence interval is provided.

easier, since there is no threshold above which particles will be created, leading to a proliferation of eigenvalues close to zero. Nevertheless, the combination of mass and chemical potential in the Dirac matrix can make it ill-conditioned and its inversion prohibitively expensive.

Our simulations have been performed in a lattice of volume $V = 12^3$, with four different temperatures ($N_\tau = 2, 4, 6$ and 8), inverse coupling $\beta = 5.6$, and two quark flavours of degenerate mass $m = 0.025$. These input parameters result in a pion mass of $m_\pi \approx 0.42$, and the nucleon mass $m_N \approx 0.93$, in lattice units [125].

Our studies used chemical potentials in the range of $0 \leq \mu/T \leq 8$. The simulations have had a total run time of 300 Langevin time units, and we have taken the thermalisation time to be 100. For $N_\tau = 6$ and 8 and $\mu/T \gtrsim 3.5$, the number of conjugate gradient steps needed to invert the fermion matrix increased by an order of magnitude compared to other regions. Therefore, large fluctuations developed in the observables, making their measurements unreliable. These cases were removed from subsequent analyses.

In a low temperature regime, pion condensation is expected to happen at $\mu =$

$m_\pi/2 \approx 0.21$, when there is enough energy available for both quark species to start forming pions. At $\mu = m_N/3 \approx 0.31$, nucleons can form, leading to a non-zero quark number density. For high temperatures, quarks are deconfined for all values of the chemical potential. The Silver Blaze region is absent and the quark number density is expected to be non-zero at any finite μ.

Figures 4.23 and 4.24 show the average quark number density, normalised by the saturation value[1], n_{sat}. The latter, similar to the one observed in HDQCD simulations, is reached for large values of the chemical potential. In the figures, for $N_\tau = 6$ and

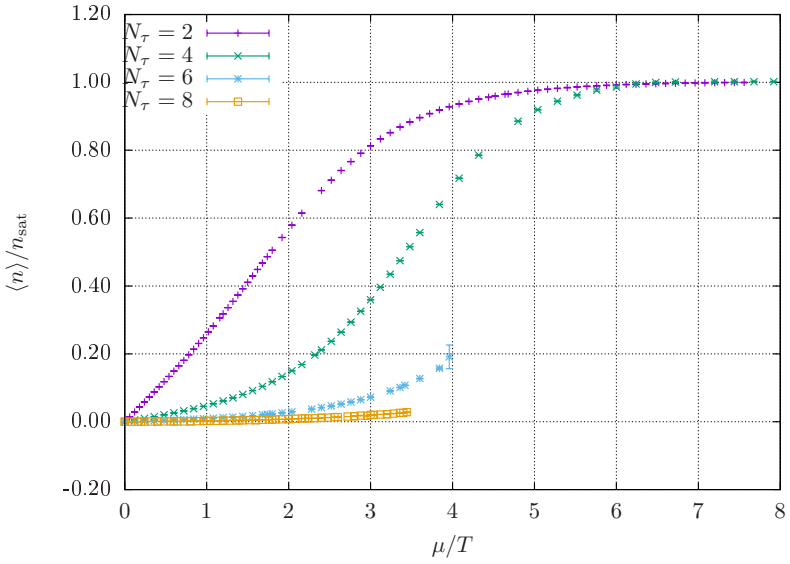

Figure 4.23: The average quark number density $\langle n \rangle$ as function of the chemical potential, in units of the temperature, for four temperatures, at a volume of 12^3. The x-axis extends into the saturation region, while the y-axis has been divided by the saturation density.

8 only the results with $\mu/T \lesssim 4$ are shown, since for larger chemical potentials the inversion of the Dirac operator became prohibitively expensive. It is possible to see that as the temperature decreases, the density's dependency on the chemical potential reduces, for low values of μ. For temperatures below the critical one, the Silver Blaze phenomenon ensures that the average density exponentially suppressed for $\mu < m_N/3$.

The average Polyakov loop, shown in fig. 4.25, is non-zero before the saturation region, due to the centre symmetry being explicitly broken by the quarks. After satu-

[1]For staggered fermions, there is no spin structure. Therefore, $n_{\text{sat}} = 3 \times N_f/4$, with the divisor being due to staggered fermions keeping only one of the four identical components of the action.

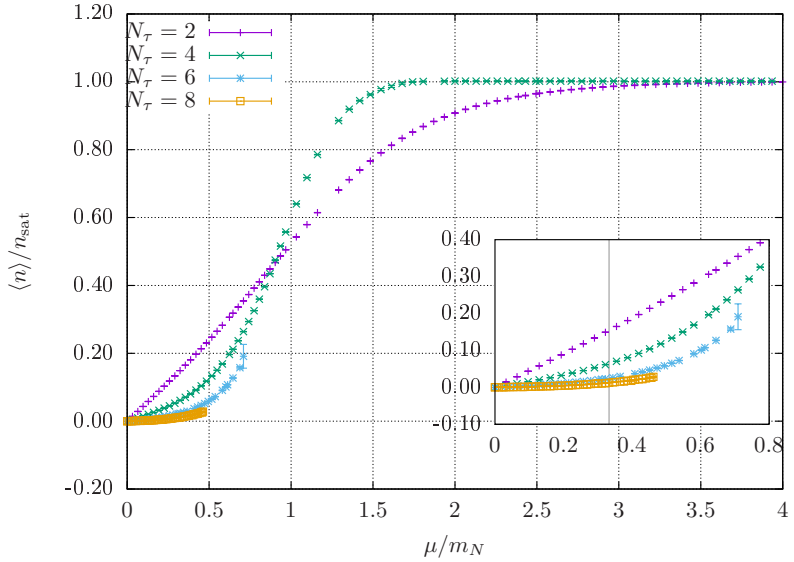

Figure 4.24: The average quark number density $\langle n \rangle$ as function of the chemical potential, in units of the nucleon mass, for four temperatures, at a volume of 12^3. The x-axis extends into the saturation region, while the y-axis has been divided by the saturation density. The vertical line indicates $\mu = m_N/3$, with the inset zooming around that region.

ration, quarks have no influence and the dynamics reduces to that of pure a Yang-Mills theory; $N_\tau = 2$ corresponds to a temperature above the critical one, with $\langle P \rangle \neq 0$, while at $N_\tau = 4$ the centre symmetry is restored. This is corroborated by the agreement with the bands displayed in the plot, which have been generated with quenched simulations. The same behaviour is expected to occur for $N_\tau \geq 6$, as the temperatures are even lower.

At our highest temperature ($N_\tau = 2$) we observe a small but non-zero (unrenormalised) chiral condensate, shown in fig. 4.26. In this high energy environment, the quark masses are almost negligible, leading to a partial restoration of the chiral symmetry. Owing to the finite quark mass, the chiral symmetry is not completely restored at $\mu = 0$. With higher chemical potentials, however, chiral symmetry is seen to be restored, as the medium becomes more energetic and the small fermion masses become negligible. For lower temperatures, the explicit breaking of the chiral symmetry is evident, with the chiral condensate growing with N_τ. Figure 4.27 displays the chiral condensate as a function of μ/m_N. The vertical line in the figure indicates the position of the phase transition at $\mu = m_N/3$. A stronger dependence on the chemical potential

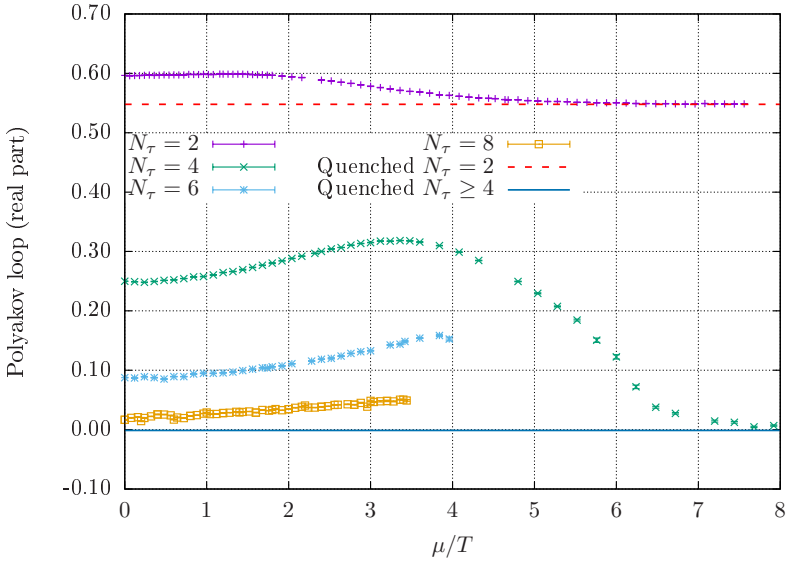

Figure 4.25: The average Polyakov loop as function of the chemical potential for four temperatures, at a volume of 12^3. The x-axis extends into the saturation region. Values for quenched simulations at different temperatures are also shown.

can be seen for the simulations at $N_\tau \gtrsim 4$, whose saturation stage is confining.

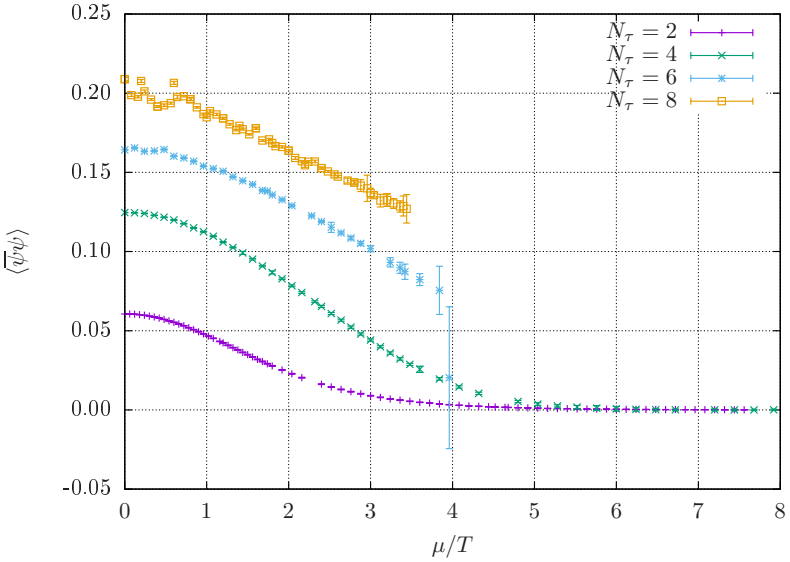

Figure 4.26: The average chiral condensate as function of the chemical potential, in units of the temperature, for four temperatures, at a volume of 12^3. The x-axis extends into the saturation region.

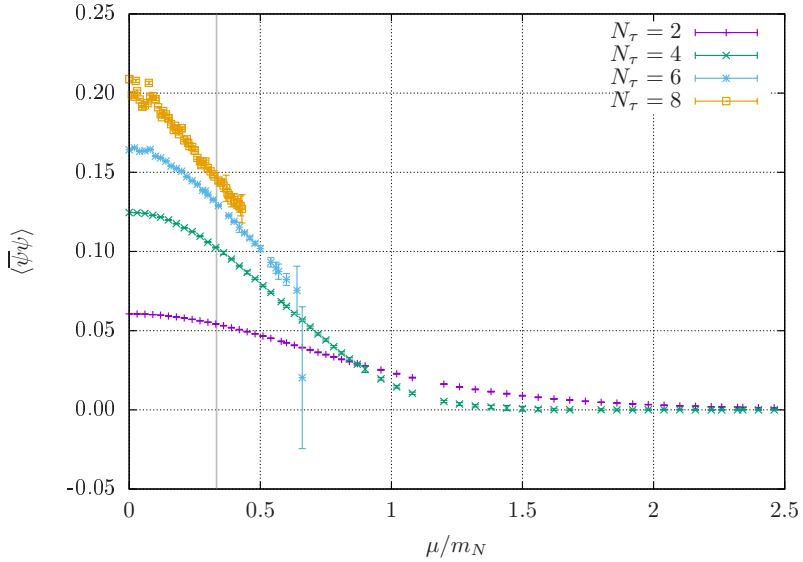

Figure 4.27: The average chiral condensate as function of the chemical potential, in units of the nucleon mass, for four temperatures, at a volume of 12^3. The x-axis extends into the saturation region, and the vertical line indicates $\mu = m_N/3$.

The main issues in our studies have been related to the inversion of the fermion matrix. As it becomes more difficult with lower temperatures and chemical potentials close to the critical one, the number of iterations required in the conjugate gradient procedure grows to unsustainable levels. Moreover, a proper calculation of the quark contribution to the Langevin drift, using the bilinear noise scheme, requires as many random vectors as possible, with a new inversion performed for each of them.

It is clear that a different approach is required here. One example is the addition of a source to the Dirac operator that explicitly "lifts" the lowest eigenvalue, and therefore improves the infrared behaviour of the operator. The reliability of this procedure, then, depends on how the source is removed. The simple extrapolation to zero can be ameliorated using, e.g., the valence quark improvement method [67], which uses the singular value decomposition of the observables that depend on the added source.

Chapter 5

Summary and outlook

In chapter 1, we reviewed the main features of QCD and how they shape its thermodynamics. We discussed firmly stabilished and conjectured phases of hadronic matter, based on analytical and numerical arguments and presented a sketch of the QCD phase diagram. Then we examined the main techniques used to simulate QCD at finite chemical potential and mentioned how current and future experiments will investigate the phases of strongly interacting matter. Afterwards, we showed how the lattice fomulation of QCD is obtained from the continuum one, and how to include finite temperature and chemical potential on the simulations. Emphasis was given on the sign problem and how it appears in such scenarios. The heavy-dense approximation of QCD (HDQCD) was introduced as a numerically cheaper alternative that shares interesting features with QCD, such as the sign problem

Chapter 2 is devoted to stochastic quantisation at zero and finite chemical potentials. We showed how, at zero chemical potential, fields evolving in an fictitious (Langevin) time, according to a Langevin equation, can be used to calculate quantum expectation values for continuous Langevin time. Then, we show how this argument is translated to a discrete Langevin time, suitable for computer simulations, with a discussion on algorithms to reduce the results' dependence on the discretisation step size. For finite chemical potentials, the sign problem causes the theory to explore a complexified phase space. We reviewed the complex Langevin method, and how it allows one to obtain meaningful results in such situations. A proof of convergence (to the correct limit for long Langevin times) is discussed, as well as methods to improve such convergence, such as gauge cooling and Dynamic Stabilisation (DS). The chapter ends

with notes on the Langevin simulation of dynamical quark fields and the associated difficulties.

A wide variety of tests of the complex Langevin method is presented in chapter 3. We first show how complex Langevin can be used for systems without a sign problem, with gauge cooling providing an alternative to reunitarisation. Then we assess the continuity of the plaquette in HDQCD, when transitioning from purely imaginary to purely real chemical potentials. We show that the combination of complex Langevin and gauge cooling results in discontinuous plaquettes only for $\beta \lesssim 5.4$. A study of the observables' dependence on the Langevin step size follows. There, we test the second order algorithms discussed in the previous chapter and show how they improve the results, making extrapolations to zero step size easier. We also present results verifying that the second order improvements cannot be applied to dynamical fermions. Lastly, we present investigations related to our method of Dynamic Stabilisation. We analysed how the observables depend on the parameter α_{DS}, and verified that the forces added to the Langevin drift decrease with the inverse gauge coupling. We show how DS allows results compatible with Monte Carlo methods (e.g., reweighting) around the deconfinement transition of HDQCD, and present data showing that DS also works for simulations using staggered fermions.

Our main results are shown and discussed in chapter 4. We used complex Langevin and gauge cooling to map the phase diagram of HDQCD, with simulations in a variety of temperatures and chemical potentials, for three different lattice volumes. We determined, via polynomial fits, the boundary between the different phases of QCD in this approximation, and were able to see the influence of the finite simulation box size. During these studies, we have encountered instabilities in the Langevin evolution, which have been shown to be related to large unitarity norms. We, then, discuss how restricting the analyses to the region with small unitarity norms can lead to reliable results. These instabilities were the original motivation for our method of Dynamic Stabilisation. Afterwards, we applied the second order discretisation scheme and DS to the determination of the critical chemical potential as a function of the hopping parameter for HDQCD. For low temperatures, linear fits in κ and T were seen to describe well the critical chemical potential. Combining the information from these previous studies, we revisited the HDQCD phase diagram, focussing around the phase boundary. The use of second order algorithms and DS allowed for more data points to be collected,

thus improving the analyses. We provided qualitative results for the thermal transition at higher temperatures, since our fixed lattice spacing approach only allows for a coarse exploration of this axis, with a quantitative analysis being employed for the transition along the chemical potential axis. Finally, we show preliminary results of complex Langevin and DS being applied to simulations of QCD with staggered quarks at high temperature. In this regime, there is no Silver Blaze region, making the inversion of the Dirac operator less expensive. We interpret the results based on the behaviour of the chiral condensate and the Polyakov loop. The latter was compared to results in quenched simulations when the system was in the saturation stage. These results indicate that DS is able to prevent issues related to large unitarity norm for staggered quarks, as our results agreed with expectations and known cases.

We intend to expand our studies of QCD with staggered quarks at finite chemical potential even further. Initially, a more thorough scan of the high temperature phase should be feasible, as inverting the fermion matrix in such situations is not as difficult as for low temperatures. This can potentially provide some information on the crossover region between quark–gluon plasma and the hadronic phase of matter. Additionally, it is a good testing ground for novel techniques to deal with the inversion of the Dirac operator when its eigenvalues become too small.

With the fermionic matrix under control in the infrared, our plan is to carry on with simulations at lower temperatures. A fine scan in the chemical potential direction of the phase diagram is expected to give information about the transition to nuclear matter and also to quark–gluon plasma. Furthermore, the existence of the QCD critical point can be assessed. Given the current success of complex Langevin, combined with gauge cooling and dynamic stabilisation, such studies should be possible.

Appendix A

The reweighting method

We consider a system whose dynamics is given by a Hamiltonian H, at fixed volume V and temperature T, and assume that expectation values can be calculated in a satisfactory manner. By that we mean that the average value of some observable $\mathcal{O}$ can be calculated analytically or, e.g., using Monte Carlo techniques. The reweighting method consists of using the system subject the Hamiltonian H to extract information about a different one, with Hamiltonian H'.

The expectation value of an observable $\mathcal{O}$ in the system with statistical weight $\rho' = e^{-H'/T}$ can be written in terms of $\rho = e^{-H/T}$ as

$$\langle \mathcal{O} \rangle_{\rho'} = \frac{\mathrm{Tr}\left[\mathcal{O}\rho'\right]}{\mathrm{Tr}\left[\rho'\right]} = \frac{\mathrm{Tr}\left[\mathcal{O}\frac{\rho'}{\rho}\rho\right]}{\mathrm{Tr}\left[\frac{\rho'}{\rho}\rho\right]} = \frac{\mathrm{Tr}\left[\mathcal{O}\frac{\rho'}{\rho}\rho\right]}{\mathrm{Tr}\left[\rho\right]} \frac{\mathrm{Tr}\left[\rho\right]}{\mathrm{Tr}\left[\frac{\rho'}{\rho}\rho\right]} = \frac{\left\langle \mathcal{O}\frac{\rho'}{\rho} \right\rangle_{\rho}}{\left\langle \frac{\rho'}{\rho} \right\rangle_{\rho}}, \tag{A.1}$$

where the ratio ρ'/ρ is known as the reweighting factor. Essentially, this method allows one to write $\langle \mathcal{O} \rangle_{\rho'}$ as the average of the reweighted observable, divided by that of the reweighting factor.

In lattice QCD, the expectation values calculated using path integrals weighted by $\rho_{\mathrm{QCD}} = e^{-S}$. At finite chemical potential μ, the action becomes complex and ρ_{QCD} loses its probabilistic interpretation. Reweighting can be used to calculate quantum expectation values of quantities at finite μ using, e.g., simulations at $\mu = 0$. This is but one particular choice of the reweighting factor, with others being available [126].

The parameter range in which reweighting is still reliable is usually estimated as follows. We require that the reweighted average of an observable $\mathcal{O}$, $\langle \mathcal{O} \rangle_{\rho'}$, should be

less than $\sigma_{\mathcal{O}}$ away from $\langle\mathcal{O}\rangle_\rho$, with $\sigma_{\mathcal{O}}$ being the statistical uncertainty in $\mathcal{O}$. We assume that H and H' differ by a change in a certain parameter λ, to $\lambda + \Delta\lambda$, and that the average of $\mathcal{O}$ can be obtained by differentiating the free energy, F, with respect to some parameter η. Then, we can write our requirement as

$$\langle\mathcal{O}\rangle_{\rho'} - \langle\mathcal{O}\rangle_\rho \lesssim \left[\langle\mathcal{O}^2\rangle_\rho - \langle\mathcal{O}\rangle_\rho^2\right]^{1/2} = \sqrt{\frac{\partial^2 F}{\partial\eta\partial\lambda}}\,. \tag{A.2}$$

A simplification can be made by Taylor expanding $\langle\mathcal{O}\rangle_{\rho'}$ around $\langle\mathcal{O}\rangle_\rho$,

$$\langle\mathcal{O}\rangle_{\rho'} \approx \langle\mathcal{O}\rangle_\rho + \Delta\lambda\frac{\partial^2 F}{\partial\eta\partial\lambda}\,, \tag{A.3}$$

from which it follows that

$$\Delta\lambda \lesssim \left[\frac{\partial^2 F}{\partial\eta\partial\lambda}\right]^{-1/2} = \left[V\frac{\partial^2 f}{\partial\eta\partial\lambda}\right]^{-1/2}\,, \tag{A.4}$$

where we have used $F = fV$, with f being the free energy density. It is clear, then, that reweighting has a range proportional to $V^{-1/2}$.

Other limitations on the applicability of reweighting, in particular to systems subject to a complex probability distribution with oscillating phase factor, stem from the choice of reweighting factor. The denominator of eq. (A.1), being the ratio of two partition functions, can be rewritten as

$$\left\langle\frac{\rho}{\rho'}\right\rangle_{\rho'} = \frac{Z_\rho}{Z_{\rho'}} = \exp\left[-\frac{V}{T}\Delta f\right]\,, \tag{A.5}$$

with ρ' assumed to be real and non-negative, and $\Delta f = f_\rho - f_{\rho'}$ being the free energy density difference between the two ensembles. This difference is non-negative, since the complex phase in ρ can cause field configurations to have opposite contributions to the partition function, thus making $Z_\rho \leq Z_{\rho'}$. This implies that the reweighting factor vanishes exponentially fast with the volume, being ill-defined in the thermodynamic limit, and is known as the *overlap problem*. Suitable choices of ρ' can lead to smaller free energy differences, allowing for larger volumes to be used.

Appendix B

Conventions on su(N) and SU(N)

In this appendix we explain our conventions regarding the elements of the Lie group of special unitary matrices, SU(N), and its corresponding Lie algebra, su(N). We also define the gauge group derivative in terms of conventional partial derivatives.

B.1 Conventions on Lie algebra elements

We denote the normalised generators of SU(N), which form a basis in the vector space of su(N), in the fundamental representation by g_F^a. They are $N \times N$ traceless Hermitian matrices that can be uniquely defined by the trace of their product and commutation relations,

$$\mathrm{Tr}\left[g_F^a g_F^b\right] = \frac{1}{2}\delta^{ab}\,, \tag{B.1}$$

$$\left[g_F^a, g_F^b\right] = i f^{abc} g_F^c\,, \tag{B.2}$$

where δ^{ab} is a Kronecker delta and f^{abc} are the completely anti-symmetric structure constants. The superscript indices run from 1 to $N^2 - 1$ and repeated indices are assumed to be summed over. The normalised generators in the adjoint representation are defined from eq. (B.2), with the matrix elements

$$(g_A^a)_{bc} = -i f^{abc}\,. \tag{B.3}$$

The Casimir invariants of su(N) in the fundamental and adjoint representations are constructed, respectively, as

$$C_F \delta^{bd} \equiv (g_F^a g_F^a)_{bd} = \frac{N^2-1}{2N}\delta^{bd}\,, \tag{B.4}$$

$$C_A \delta^{bd} \equiv (g_A^a g_F^a)_{bd} = -f^{abc}f^{acd} = N\delta^{bd}\,. \tag{B.5}$$

Other choices for the normalisation are, obviously, possible. To account for that and allow for easier translation of group theoretical expressions amongst various publications we re-write the above expressions for the generators $\lambda_F^a = \alpha g_F^a$. The generators' defining properties then become

$$\mathrm{Tr}\left[\lambda_F^a g_F^b\right] = \frac{\alpha^2}{2}\delta^{ab}\,, \tag{B.6}$$

$$\left[\lambda_F^a, g_F^b\right] = i\alpha f^{abc} g_F^c\,, \tag{B.7}$$

with their adjoint representation now being

$$(\lambda_A^a)_{bc} = -i\alpha f^{abc}\,. \tag{B.8}$$

The respective Casimir operators can easily be seen to be re-scaled by α^2 in comparison to the ones defined above:

$$C_F \to \alpha^2 C_F = \alpha^2 \frac{N^2-1}{2N}\,, \tag{B.9}$$

$$C_A \to \alpha^2 C_A = \alpha^2 N\,. \tag{B.10}$$

Using the generators for su(N), a generic $N \times N$ Hermitian matrix can be written as $X = X^0 \mathbb{1} + X^a \lambda^a$, from which the appropriate coefficients can be extracted using the properties outlined above:

$$X^0 = \frac{\mathrm{Tr}\,[X]}{N}\,, \tag{B.11}$$

$$X^b = \frac{2}{\alpha^2}\mathrm{Tr}\left[\lambda^b X\right]\,. \tag{B.12}$$

From the above relations one has that

$$X = \frac{1}{N}\mathrm{Tr}\,[X]\,\mathbb{1} + \frac{2}{\alpha^2}\mathrm{Tr}\,[\lambda^a X]\,\lambda^a\,, \tag{B.13}$$

which, in turn, leads to the Fierz identity

$$\frac{2}{\alpha^2}\left(\lambda^a\right)_{ij}\left(\lambda^a\right)_{kl}+\frac{1}{N}\delta_{ij}\delta_{kl}=\delta_{il}\delta_{kj}\,. \tag{B.14}$$

Some of the references cited in this thesis use different conventions for their generators: [76] uses $\alpha=1$, [79,106] take $\alpha=\sqrt{2}$ while [47–50,52,68,69,93,94,97,99,104] have chosen $\alpha=2$. The choice of anti-Hermitian generators can be achieved by taking α as a purely imaginary number: [78] used $\alpha=i$. While this list is certainly not exhaustive, it should help the interested reader.

B.2 Gauge group derivative

Following the derivations in [61] and [127] we proceed to define the gauge group derivative used in chapter 2. First we invoke the definition of the adjoint action in a Lie algebra

$$\mathrm{ad}_\omega(\xi)\equiv[\omega,\xi]\;, \tag{B.15}$$

where $\omega,\xi\in\mathrm{su}(N)$ and define the function

$$E(x)=\frac{1-\exp(-x)}{x}=\sum_{n=0}^{\infty}\frac{1}{(n+1)!}(-x)^n\,. \tag{B.16}$$

Then we use the fact that a SU(3) element can be parametrised in terms of the generators as $U=\exp(i\lambda^a A^a)$, with A^a real numbers, to write

$$U^{-1}\frac{\partial}{\partial A^b}U=iE(\mathrm{ad}_A)\lambda^b\,. \tag{B.17}$$

It is possible to represent the adjoint action defined above as an anti-symmetric matrix, $\overline{\omega}^{cb}=\alpha\omega^a f^{cba}$, through

$$\mathrm{ad}_\omega(\xi)=i\alpha f^{abc}\omega^a\xi^b\lambda^c\equiv-i\lambda^c\overline{\omega}^{cb}\xi^b\,. \tag{B.18}$$

Therefore, we can write

$$U^{-1}\frac{\partial}{\partial A^b}U=iE(\overline{A})^{bc}\lambda^c\,. \tag{B.19}$$

Letting

$$\nabla^a=\left[E(\overline{A})^{-1}\right]^{ab}\frac{\partial}{\partial A^b} \tag{B.20}$$

we have $\nabla^a U = i\lambda^a U$, analogous to what happens in the simpler case of U(1) gauge fields. An alternative way to represent this derivative is obtained by applying it to a function of U and using the usual chain rule

$$\begin{aligned}\nabla^a f(U) &= \left[E(\overline{A})^{-1}\right]^{ab} \frac{\partial f}{\partial U_{ij}} \frac{\partial U_{ij}}{\partial A^b} \\ &= \left[E(\overline{A})^{-1}\right]^{ab} E(\overline{A})^{bc} \frac{\partial f}{\partial U_{ij}} \left(i\lambda^c U\right)_{ij} \\ &= \frac{\partial f}{\partial U_{ij}} (i\lambda^a U)_{ij} \,, \end{aligned} \tag{B.21}$$

from which we get

$$\nabla^a = (i\lambda^a U)_{ij} \frac{\partial}{\partial U_{ij}} \,. \tag{B.22}$$

Another equivalent way, which is operationally simpler, is found in [47]

$$\nabla^a f(U) = \frac{\partial}{\partial \alpha} f(e^{i\alpha\lambda^a} U)\bigg|_{\alpha=0} \,. \tag{B.23}$$

Gauge group derivatives, in general, do not commute. They follow a commutation relation similar to that of the generators,

$$\left[\nabla^a, \nabla^b\right] = \alpha f^{abc} \nabla^c \,. \tag{B.24}$$

Additionally, the contraction $\nabla^2 \equiv \nabla^a \nabla^a$ commutes with all derivatives,

$$\left[\nabla^2, \nabla^b\right] = 0 \,. \tag{B.25}$$

Two useful identities that arise from the commutation relation are

$$f^{abc} \nabla^a \nabla^b = \frac{\alpha}{2} C_A \nabla^c \,, \tag{B.26}$$

and

$$f^{abc} \nabla^a \nabla^b \nabla^c = \frac{\alpha}{2} C_A \nabla^2 \,, \tag{B.27}$$

where we have used the complete anti-symmetry of the structure constants, and eqs. (B.5) and (B.10) for the Casimir invariant of the adjoint representation.

Appendix C

Taylor expansion of $f\left(e^{iX(\overline{\varepsilon})}U\right)$

In this appendix we present the steps necessary to write the Taylor expansion of a holomorphic scalar function of a gauge link after one Langevin evolution step. This can be particularly involved, since the terms in the expansion will contain non-commuting elements. Writing the final result in terms of the gauge group derivative of appendix B allows for easier use of the expansion.

We consider a scalar function f that takes gauge links into complex numbers, $f : \mathrm{SL}(N,\mathbb{C}) \to \mathbb{C}$. We also assume f to be holomorphic; i.e., $f(U,U^\dagger) = f(U)$, $U, U^\dagger \in \mathrm{SL}(N,\mathbb{C})$. After a Langevin step, a generic gauge link can be written as $U' = \exp[iX(\overline{\varepsilon})]U$, with U being the link from the previous time step, c.f. eq. (2.11), and $X(0) = 0$. In order to facilitate subsequent calculations we introduce the following notations for the derivatives of $f(U)$,

$$\frac{\partial^n f}{\partial U_{ij}\,\partial U_{i'j'}\,\cdots\,\partial U_{i^{(n)}j^{(n)}}} \equiv \partial^{(n)}_{ji,j'i',\cdots,j^{(n)}i^{(n)}} f \equiv f^{(n)}_{ji,j'i',\cdots,j^{(n)}i^{(n)}}\,, \tag{C.1}$$

and for those of X,

$$\overline{X}^{(n)} = \left.\frac{\partial^n X}{\partial \overline{\varepsilon}^n}\right|_{\overline{\varepsilon}=0}. \tag{C.2}$$

We show here how to obtain the first two terms of the Taylor expansion of $f(U')$ around $\overline{\varepsilon} = 0$. The first term is given by

$$\left.\frac{\partial f(U')}{\partial \overline{\varepsilon}}\right|_{\overline{\varepsilon}=0} = \left.\frac{\partial f(U')}{\partial (U')_{ij}}\frac{\partial (U')_{ij}}{\partial \overline{\varepsilon}}\right|_{\overline{\varepsilon}=0}. \tag{C.3}$$

The second term in the right-hand side can be calculated using a further Taylor expan-

sion, evaluated at $\overline{\varepsilon}=0$,

$$\left.\frac{\partial(U')_{ij}}{\partial\overline{\varepsilon}}\right|_{\overline{\varepsilon}=0} = \left.\frac{\partial(e^{iX(\overline{\varepsilon})}U)_{ij}}{\partial\overline{\varepsilon}}\right|_{\overline{\varepsilon}=0} = \left.\frac{\partial}{\partial\overline{\varepsilon}}\left(\delta_{ij}+iX_{ia}U_{aj}+O(\overline{\varepsilon}^2)\right)\right|_{\overline{\varepsilon}=0} = i(\overline{X}'U)_{ij}\,. \quad \text{(C.4)}$$

Then, using eq. (C.1) we may write eq. (C.3) as

$$\left.\frac{\partial f(U')}{\partial\overline{\varepsilon}}\right|_{\overline{\varepsilon}=0} = i(\overline{X}'U)_{ij}f'_{ji}(U) = i\overline{X}'^{a}(\lambda^a U)_{ij}f'_{ji}(U) = \overline{X}'^{a}\nabla^a f(U)\,, \quad \text{(C.5)}$$

using the notation of eq. B.22.

The second order term requires the use of the identity

$$\frac{\partial U_{ij}}{\partial U_{kl}} = \delta_{ik}\delta_{jl}\,, \quad \text{(C.6)}$$

in order to be written in a closed form in terms of eq. (B.22). We start by applying the chain rule

$$\left.\frac{\partial^2 f(U')}{\partial\overline{\varepsilon}^2}\right|_{\overline{\varepsilon}=0} = \left[\frac{\partial^2 f(U')}{\partial U'_{ij}\partial U'_{i'j'}}\frac{\partial U'_{ij}}{\partial\overline{\varepsilon}}\frac{\partial U'_{i'j'}}{\partial\overline{\varepsilon}} + \frac{\partial f(U')}{\partial U'_{ij}}\frac{\partial^2 U'_{ij}}{\partial\overline{\varepsilon}^2}\right]_{\overline{\varepsilon}=0}. \quad \text{(C.7)}$$

The first term in the right-hand side is just $f''_{ij,i'j'}$, using eq. (C.1). The first derivative of U' with respect to $\overline{\varepsilon}$ has been calculated in eq. (C.4), while its second derivative can be manipulated into a convenient expression

$$\begin{aligned}
\left.\frac{\partial^2 U'_{ij}}{\partial\overline{\varepsilon}^2}\right|_{\overline{\varepsilon}=0} &= \frac{\partial^2}{\partial\overline{\varepsilon}^2}\left[\delta_{ij}+i(XU)_{ij}+\frac{i^2}{2}(XXU)_{ij}+O(\overline{\varepsilon}^3)\right]\Big|_{\overline{\varepsilon}=0} \\
&= i(\overline{X}''U)_{ij} + i^2\overline{X}'_{ip}\overline{X}'_{pq}U_{qj} \\
&= i(\overline{X}''U)_{ij} + i^2\overline{X}'_{ip}\overline{X}'_{i'q}\underbrace{\delta_{i'p}U_{qj'}\delta_{j'j}} \\
&= i(\overline{X}''U)_{ij} + i^2\overline{X}'_{i'q}U_{qj'}\overline{X}'_{ip}\frac{\partial U_{pj}}{\partial U_{i'j'}} \\
&= i(\overline{X}''U)_{ij} + i^2\overline{X}'^{a}(\overline{X}'U)_{i'j'}\frac{\partial(\lambda^a U)_{ij}}{\partial U_{i'j'}} \\
&= i(\overline{X}''U)_{ij} + i^2\overline{X}'^{a}(\overline{X}'U)_{i'j'}\partial_{j'i'}(\lambda^a U)_{ij}\,,
\end{aligned} \quad \text{(C.8)}$$

where we inserted Kronecker δ's on the third and used the identity of eq. (C.6) on going

to the fourth line. Combining eqs. (C.7) and (C.8) and integrating by parts we have

$$\begin{aligned}
\left.\frac{\partial^2 f(U')}{\partial\overline{\varepsilon}^2}\right|_{\overline{\varepsilon}=0} &= i^2(\overline{X}'U)_{ij}(\overline{X}'U)_{i'j'}f''_{ji,j'i'} + i(\overline{X}''U)_{ij}f'_{ji} + i^2\overline{X}'^a(\overline{X}'U)_{i'j'}\partial_{j'i'}(\lambda^a U)_{ij}f'_{ji} \\
&= i^2\overline{X}'^a(\overline{X}'U)_{i'j'}\left[(\lambda^a U)_{ij}\partial_{j'i'}f'_{ji} + \partial_{j'i'}(\lambda^a U)_{ij}f'_{ji}\right] + i(\overline{X}''U)_{ij}f'_{ji} \\
&= i^2\overline{X}'^a(\overline{X}'U)_{i'j'}\partial_{j'i'}\left[(\lambda^a U)_{ij}f'_{ji}\right] + i(\overline{X}''U)_{ij}f'_{ji} \\
&= \overline{X}'^a\overline{X}'^b\nabla^b\nabla^a f(U) + \overline{X}''^a\nabla^a f(U)\,. \qquad \text{(C.9)}
\end{aligned}$$

Subsequent terms in the expansion can be dealt with in a similar manner: inserting Kronecker deltas and using the identity in eq. (C.6) allows them to be written in terms of the gauge group derivatives, ∇.

A suitable rederivation of the first order term for non-holomorphic functions $f(U', U'^\dagger) = f(e^{iX(\overline{\varepsilon})}U, U^\dagger e^{-iX^\dagger(\overline{\varepsilon})})$ proceeds as follows:

$$\begin{aligned}
\frac{d}{d\overline{\varepsilon}}f(e^{iX(\overline{\varepsilon})}U, U^\dagger e^{-iX^\dagger(\overline{\varepsilon})})\Big|_{\overline{\varepsilon}=0} &= \left(i\overline{X}'^a\lambda^a U\right)_{ij}\partial_{ij}f + \left(-i\overline{X}'^{*a}U^\dagger\lambda^a\right)_{ij}\overline{\partial}_{ij}f \\
&= \left[\left(i\overline{X}'^a_R - \overline{X}'^a_I\right)\lambda^a U\right]_{ij}\partial_{ij}f + \left[\left(-i\overline{X}'^a_R - \overline{X}'^a_I\right)U^\dagger\lambda^a\right]_{ij}\overline{\partial}_{ij}f \\
&= \overline{X}'^a_R\left[(i\lambda^a U)_{ij}\,\partial_{ij} - \left(iU^\dagger\lambda^a\right)_{ij}\overline{\partial}_{ij}\right]f + \overline{X}'^a_I\left[(-\lambda^a U)_{ij}\,\partial_{ij} + \left(-U^\dagger\lambda^a\right)_{ij}\overline{\partial}_{ij}\right]f \\
&\equiv \left(\overline{X}'^a_R\nabla^a_R + \overline{X}'^a_I\nabla^a_I\right)f\,. \qquad \text{(C.10)}
\end{aligned}$$

This motivates the prescription

$$X^a\nabla^a \to X^a_R\nabla^a_R + X^a_I\nabla_I\,, \qquad \text{(C.11)}$$

where X^a_R and X^a_I are the real and imaginary parts of X^a, respectively, for writing the Taylor expansion above for a non-holomorphic function. It is important to stress that when applied to holomorphic functions, this redefined derivative obeys the Cauchy-Riemann equation

$$i\nabla^a_R f = \nabla^a_I f\,, \qquad \text{(C.12)}$$

resulting in

$$X^a\nabla^a f \to \left(X^a_R\nabla^a_R + X^a_I\nabla^a_I\right)f = X^a\nabla^a_R f\,, \qquad \text{(C.13)}$$

for complex X^a.

References

[1] K. Fukushima and T. Hatsuda, "The phase diagram of dense QCD," *Rept. Prog. Phys.* **74** (2011) 014001, arXiv:1005.4814 [hep-ph]. 1, 5

[2] O. Philipsen, "The QCD equation of state from the lattice," *Prog. Part. Nucl. Phys.* **70** (2013) 55–107, arXiv:1207.5999 [hep-lat]. 1

[3] H.-T. Ding, F. Karsch, and S. Mukherjee, "Thermodynamics of Strong-Interaction Matter from Lattice QCD," in *Quark-Gluon Plasma 5*, X.-N. Wang, ed., pp. 1–65. 2016. 1

[4] C. Schmidt and S. Sharma, "The phase structure of QCD," arXiv:1701.04707 [hep-lat]. 1

[5] R. Hagedorn, "Statistical thermodynamics of strong interactions at high-energies," *Nuovo Cim. Suppl.* **3** (1965) 147–186. 2

[6] H. D. Politzer, "Reliable Perturbative Results for Strong Interactions?," *Phys. Rev. Lett.* **30** (1973) 1346–1349. 2, 9

[7] D. J. Gross and F. Wilczek, "Asymptotically Free Gauge Theories. 1," *Phys. Rev.* **D8** (1973) 3633–3652. 2, 9

D. J. Gross and F. Wilczek, "Ultraviolet Behavior of Nonabelian Gauge Theories," *Phys. Rev. Lett.* **30** (1973) 1343–1346.

[8] M. Gell-Mann, R. J. Oakes, and B. Renner, "Behavior of current divergences under SU(3) x SU(3)," *Phys. Rev.* **175** (1968) 2195–2199. 3

[9] M. Fukugita, M. Okawa, and A. Ukawa, "Finite Size Scaling Study of the Deconfining Phase Transition in Pure SU(3) Lattice Gauge Theory," *Nucl. Phys.* **B337** (1990) 181–232. 3

[10] Y. Aoki, G. Endrodi, Z. Fodor, S. D. Katz, and K. K. Szabo, "The Order of the quantum chromodynamics transition predicted by the standard model of particle physics," *Nature* **443** (2006) 675–678, `arXiv:hep-lat/0611014 [hep-lat]`. 3

[11] C. DeTar and U. M. Heller, "QCD Thermodynamics from the Lattice," *Eur. Phys. J.* **A41** (2009) 405–437, `arXiv:0905.2949 [hep-lat]`. 3

[12] **Wuppertal-Budapest** Collaboration, S. Borsanyi, Z. Fodor, C. Hoelbling, S. D. Katz, S. Krieg, C. Ratti, and K. K. Szabo, "Is there still any T_c mystery in lattice QCD? Results with physical masses in the continuum limit III," *JHEP* **09** (2010) 073, `arXiv:1005.3508 [hep-lat]`. 3

[13] A. Bazavov *et al.*, "The chiral and deconfinement aspects of the QCD transition," *Phys. Rev.* **D85** (2012) 054503, `arXiv:1111.1710 [hep-lat]`. 3

[14] M. Asakawa and K. Yazaki, "Chiral Restoration at Finite Density and Temperature," *Nucl. Phys.* **A504** (1989) 668–684. 3

[15] F. Wilczek, "Application of the renormalization group to a second order QCD phase transition," *Int. J. Mod. Phys.* **A7** (1992) 3911–3925. [Erratum: Int. J. Mod. Phys.A7,6951(1992)]. 3

[16] J. Berges and K. Rajagopal, "Color superconductivity and chiral symmetry restoration at nonzero baryon density and temperature," *Nucl. Phys.* **B538** (1999) 215–232, `arXiv:hep-ph/9804233 [hep-ph]`. 3

[17] M. A. Stephanov, K. Rajagopal, and E. V. Shuryak, "Signatures of the tricritical point in QCD," *Phys. Rev. Lett.* **81** (1998) 4816–4819, `arXiv:hep-ph/9806219 [hep-ph]`. 3

[18] B. C. Barrois, "Superconducting Quark Matter," *Nucl. Phys.* **B129** (1977) 390–396. 4

[19] D. Bailin and A. Love, "Superfluidity and Superconductivity in Relativistic Fermion Systems," *Phys. Rept.* **107** (1984) 325. 4

[20] K. Rajagopal and F. Wilczek, "The Condensed matter physics of QCD," in *At the frontier of particle physics. Handbook of QCD. Vol. 1-3*, M. Shifman and B. Ioffe, eds., pp. 2061–2151. 2000. `arXiv:hep-ph/0011333 [hep-ph]`. 4

[21] M. G. Alford, A. Schmitt, K. Rajagopal, and T. Schäfer, "Color superconductivity in dense quark matter," *Rev. Mod. Phys.* **80** (2008) 1455–1515, `arXiv:0709.4635 [hep-ph]`. 4

[22] J. R. Oppenheimer and G. M. Volkoff, "On Massive neutron cores," *Phys. Rev.* **55** (1939) 374–381. 4

[23] P. de Forcrand, "Simulating QCD at finite density," *PoS* **LAT2009** (2009) 010, `arXiv:1005.0539 [hep-lat]`. 5

[24] C. Gattringer and K. Langfeld, "Approaches to the sign problem in lattice field theory," *Int. J. Mod. Phys.* **A31** no. 22, (2016) 1643007, `arXiv:1603.09517 [hep-lat]`. 5

[25] C. R. Allton, S. Ejiri, S. J. Hands, O. Kaczmarek, F. Karsch, E. Laermann, C. Schmidt, and L. Scorzato, "The QCD thermal phase transition in the presence of a small chemical potential," *Phys. Rev.* **D66** (2002) 074507, `arXiv:hep-lat/0204010 [hep-lat]`. 6

[26] C. R. Allton, M. Doring, S. Ejiri, S. J. Hands, O. Kaczmarek, F. Karsch, E. Laermann, and K. Redlich, "Thermodynamics of two flavor QCD to sixth order in quark chemical potential," *Phys. Rev.* **D71** (2005) 054508, `arXiv:hep-lat/0501030 [hep-lat]`. 6

[27] R. V. Gavai and S. Gupta, "QCD at finite chemical potential with six time slices," *Phys. Rev.* **D78** (2008) 114503, `arXiv:0806.2233 [hep-lat]`. 6

[28] M. G. Alford, A. Kapustin, and F. Wilczek, "Imaginary chemical potential and finite fermion density on the lattice," *Phys. Rev.* **D59** (1999) 054502, `arXiv:hep-lat/9807039 [hep-lat]`. 6

[29] P. de Forcrand and O. Philipsen, "The QCD phase diagram for small densities from imaginary chemical potential," *Nucl. Phys.* **B642** (2002) 290–306, arXiv:hep-lat/0205016 [hep-lat]. 6, 82

[30] A. Roberge and N. Weiss, "Gauge Theories With Imaginary Chemical Potential and the Phases of QCD," *Nucl. Phys.* **B275** (1986) 734–745. 6

[31] D. E. Miller and K. Redlich, "Exact Implementation of Baryon Number Conservation in Lattice Gauge Theory," *Phys. Rev.* **D35** (1987) 2524. 6

[32] A. Hasenfratz and D. Toussaint, "Canonical ensembles and nonzero density quantum chromodynamics," *Nucl. Phys.* **B371** (1992) 539–549. 6

[33] S. Kratochvila and P. de Forcrand, "The Canonical approach to finite density QCD," *PoS* **LAT2005** (2006) 167, arXiv:hep-lat/0509143 [hep-lat]. 6

[34] K. Fukushima, "Thermodynamic limit of the canonical partition function with respect to the quark number in QCD," *Annals Phys.* **304** (2003) 72–88, arXiv:hep-ph/0204302 [hep-ph]. 7

[35] A. Gocksch, "Simulating lattice QCD at finite density," *Phys. Rev. Lett.* **61** (1988) 2054. 7

[36] Z. Fodor, S. D. Katz, and C. Schmidt, "The Density of states method at non-zero chemical potential," *JHEP* **03** (2007) 121, arXiv:hep-lat/0701022 [hep-lat]. 7

[37] K. Langfeld, B. Lucini, and A. Rago, "The density of states in gauge theories," *Phys. Rev. Lett.* **109** (2012) 111601, arXiv:1204.3243 [hep-lat]. 7

[38] K. Langfeld, B. Lucini, R. Pellegrini, and A. Rago, "An efficient algorithm for numerical computations of continuous densities of states," *Eur. Phys. J.* **C76** no. 6, (2016) 306, arXiv:1509.08391 [hep-lat]. 7

[39] C. Gattringer and P. Törek, "Density of states method for the $\mathbb{Z}_3$ spin model," *Phys. Lett.* **B747** (2015) 545–550, arXiv:1503.04947 [hep-lat]. 7

[40] N. Garron and K. Langfeld, "Anatomy of the sign-problem in heavy-dense QCD," arXiv:1605.02709 [hep-lat]. 7

[41] M. Giuliani, C. Gattringer, and P. Törek, "Developing and testing the density of states FFA method in the SU(3) spin model," *Nucl. Phys.* **B913** (2016) 627–642, arXiv:1607.07340 [hep-lat]. 7

[42] E. Witten, "Analytic Continuation Of Chern-Simons Theory," *AMS/IP Stud. Adv. Math.* **50** (2011) 347–446, arXiv:1001.2933 [hep-th]. 7

E. Witten, "A New Look At The Path Integral Of Quantum Mechanics," arXiv:1009.6032 [hep-th].

[43] H. Fujii, D. Honda, M. Kato, Y. Kikukawa, S. Komatsu, and T. Sano, "Hybrid Monte Carlo on Lefschetz thimbles - A study of the residual sign problem," *JHEP* **10** (2013) 147, arXiv:1309.4371 [hep-lat]. 7

[44] **AuroraScience** Collaboration, M. Cristoforetti, F. Di Renzo, and L. Scorzato, "New approach to the sign problem in quantum field theories: High density QCD on a Lefschetz thimble," *Phys. Rev.* **D86** (2012) 074506, arXiv:1205.3996 [hep-lat]. 7

[45] G. Aarts, L. Bongiovanni, E. Seiler, and D. Sexty, "Some remarks on Lefschetz thimbles and complex Langevin dynamics," *JHEP* **10** (2014) 159, arXiv:1407.2090 [hep-lat]. 7

[46] G. Parisi and Y.-s. Wu, "Perturbation Theory Without Gauge Fixing," *Sci. Sin.* **24** (1981) 483. 7, 27

[47] G. Aarts and I.-O. Stamatescu, "Stochastic quantization at finite chemical potential," *JHEP* **09** (2008) 018, arXiv:0807.1597 [hep-lat]. 8, 21, 22, 30, 35, 113, 114

[48] G. Aarts, F. A. James, E. Seiler, and I.-O. Stamatescu, "Adaptive stepsize and instabilities in complex Langevin dynamics," *Phys. Lett.* **B687** (2010) 154–159, arXiv:0912.0617 [hep-lat]. 8, 35, 113

[49] G. Aarts, E. Seiler, and I.-O. Stamatescu, "The Complex Langevin method: When can it be trusted?," *Phys. Rev.* **D81** (2010) 054508, arXiv:0912.3360 [hep-lat]. 8, 35, 113

[50] G. Aarts, F. Attanasio, B. Jäger, and D. Sexty, "The QCD phase diagram in the limit of heavy quarks using complex Langevin dynamics," *JHEP* **09** (2016) 087, arXiv:1606.05561 [hep-lat]. 8, 40, 75, 113

[51] G. Aarts, F. Attanasio, B. Jäger, and D. Sexty, "Results on the heavy-dense QCD phase diagram using complex Langevin," *PoS* **LATTICE2016** (2016) 074, arXiv:1610.04401 [hep-lat]. 8

[52] F. Attanasio and B. Jäger, "Testing dynamic stabilisation in complex Langevin simulations," *PoS* **LATTICE2016** (2016) 053, arXiv:1610.09298 [hep-lat]. 8, 40, 113

[53] K. G. Wilson, "Confinement of Quarks," *Phys. Rev.* **D10** (1974) 2445–2459. [,45(1974)]. 9, 11

[54] J. B. Kogut and L. Susskind, "Hamiltonian Formulation of Wilson's Lattice Gauge Theories," *Phys. Rev.* **D11** (1975) 395–408. 12

[55] H. J. Rothe, "Lattice gauge theories: An Introduction," *World Sci. Lect. Notes Phys.* **43** (1992) 1–381. [World Sci. Lect. Notes Phys.82,1(2012)]. 13, 17, 18, 27, 49

[56] C. Gattringer and C. B. Lang, "Quantum chromodynamics on the lattice," *Lect. Notes Phys.* **788** (2010) 1–343. 13, 17, 18, 27

[57] P. H. Ginsparg and K. G. Wilson, "A Remnant of Chiral Symmetry on the Lattice," *Phys. Rev.* **D25** (1982) 2649. 13

[58] R. Narayanan and H. Neuberger, "A Construction of lattice chiral gauge theories," *Nucl. Phys.* **B443** (1995) 305–385, arXiv:hep-th/9411108 [hep-th]. 13

[59] D. B. Kaplan, "A Method for simulating chiral fermions on the lattice," *Phys. Lett.* **B288** (1992) 342–347, arXiv:hep-lat/9206013 [hep-lat]. 13

[60] I. O. Stamatescu, "A Note on the Lattice Fermionic Determinant," *Phys. Rev.* **D25** (1982) 1130. 18

[61] I. Montvay and G. Munster, *Quantum fields on a lattice*. Cambridge University Press, 1997. 18, 113

[62] N. Bilic and R. V. Gavai, "On the Thermodynamics of an Ideal Fermi Gas on the Lattice at Finite Density," *Z. Phys.* **C23** (1984) 77. 19

[63] P. Hasenfratz and F. Karsch, "Chemical Potential on the Lattice," *Phys. Lett.* **B125** (1983) 308–310. 19

[64] J. B. Kogut, H. Matsuoka, M. Stone, H. W. Wyld, S. H. Shenker, J. Shigemitsu, and D. K. Sinclair, "Chiral Symmetry Restoration in Baryon Rich Environments," *Nucl. Phys.* **B225** (1983) 93–122. 19

[65] D. T. Son and M. A. Stephanov, "QCD at finite isospin density," *Phys. Rev. Lett.* **86** (2001) 592–595, `arXiv:hep-ph/0005225 [hep-ph]`. 21

[66] P. de Forcrand, M. A. Stephanov, and U. Wenger, "On the phase diagram of QCD at finite isospin density," *PoS* **LAT2007** (2007) 237, `arXiv:0711.0023 [hep-lat]`. 21

[67] B. B. Brandt and G. Endrodi, "QCD phase diagram with isospin chemical potential," *PoS* **LATTICE2016** (2016) 039, `arXiv:1611.06758 [hep-lat]`. 21, 103

[68] G. Aarts, "Can stochastic quantization evade the sign problem? The relativistic Bose gas at finite chemical potential," *Phys. Rev. Lett.* **102** (2009) 131601, `arXiv:0810.2089 [hep-lat]`. 21, 35, 113

[69] G. Aarts, "Complex Langevin dynamics at finite chemical potential: Mean field analysis in the relativistic Bose gas," *JHEP* **05** (2009) 052, `arXiv:0902.4686 [hep-lat]`. 21, 30, 35, 113

[70] G. Aarts and F. A. James, "On the convergence of complex Langevin dynamics: The Three-dimensional XY model at finite chemical potential," *JHEP* **08** (2010) 020, `arXiv:1005.3468 [hep-lat]`. 21

[71] G. Aarts and F. A. James, "Complex Langevin dynamics in the SU(3) spin model at nonzero chemical potential revisited," *JHEP* **01** (2012) 118, `arXiv:1112.4655 [hep-lat]`. 21

[72] Y. Delgado Mercado and C. Gattringer, "Monte Carlo simulation of the SU(3) spin model with chemical potential in a flux representation," *Nucl. Phys.* **B862** (2012) 737–750, arXiv:1204.6074 [hep-lat]. 21

[73] I. Bender, T. Hashimoto, F. Karsch, V. Linke, A. Nakamura, M. Plewnia, I. O. Stamatescu, and W. Wetzel, "Full QCD and QED at finite temperature and chemical potential," *Nucl. Phys. Proc. Suppl.* **26** (1992) 323–325. 21

[74] T. D. Cohen, "Functional integrals for QCD at nonzero chemical potential and zero density," *Phys. Rev. Lett.* **91** (2003) 222001, arXiv:hep-ph/0307089 [hep-ph]. 23

[75] G. Aarts, "Introductory lectures on lattice QCD at nonzero baryon number," *J. Phys. Conf. Ser.* **706** no. 2, (2016) 022004, arXiv:1512.05145 [hep-lat]. 23

[76] P. H. Damgaard and H. Huffel, "Stochastic Quantization," *Phys. Rept.* **152** (1987) 227. 28, 34, 113

[77] G. G. Batrouni, G. R. Katz, A. S. Kronfeld, G. P. Lepage, B. Svetitsky, and K. G. Wilson, "Langevin Simulations of Lattice Field Theories," *Phys. Rev.* **D32** (1985) 2736. 29, 32, 43

[78] A. S. Kronfeld, "Dynamics of Langevin simulations," *Prog. Theor. Phys. Suppl.* **111** (1993) 293–312, arXiv:hep-lat/9205008 [hep-lat]. 29, 31, 113

[79] A. Ukawa and M. Fukugita, "Langevin simulation including dynamical quark loops," *Phys. Rev. Lett.* **55** (1985) 1854. 32, 34, 113

[80] G. S. Bali, C. Bauer, A. Pineda, and C. Torrero, "Perturbative expansion of the energy of static sources at large orders in four-dimensional SU(3) gauge theory," *Phys. Rev.* **D87** (2013) 094517, arXiv:1303.3279 [hep-lat]. 32, 34

[81] J. R. Klauder, "Stochastic quantization," *Acta Phys. Austriaca Suppl.* **25** (1983) 251–281. 34

J. R. Klauder, "A Langevin Approach to Fermion and Quantum Spin Correlation Functions," *J. Phys.* **A16** (1983) L317.

J. R. Klauder, "Coherent State Langevin Equations for Canonical Quantum Systems With Applications to the Quantized Hall Effect," *Phys. Rev.* **A29** (1984) 2036–2047.

[82] G. Parisi, "On complex probabilities," *Phys. Lett.* **B131** (1983) 393–395. 34

[83] F. Karsch and H. W. Wyld, "Complex Langevin Simulation of the SU(3) Spin Model With Nonzero Chemical Potential," *Phys. Rev. Lett.* **55** (1985) 2242. 34

[84] J. Ambjorn and S. K. Yang, "Numerical Problems in Applying the Langevin Equation to Complex Effective Actions," *Phys. Lett.* **B165** (1985) 140. 34

[85] J. Ambjorn, M. Flensburg, and C. Peterson, "The Complex Langevin Equation and Monte Carlo Simulations of Actions With Static Charges," *Nucl. Phys.* **B275** (1986) 375. 34

[86] J. Berges and I. O. Stamatescu, "Simulating nonequilibrium quantum fields with stochastic quantization techniques," *Phys. Rev. Lett.* **95** (2005) 202003, `arXiv:hep-lat/0508030 [hep-lat]`. 35

[87] J. Berges, S. Borsanyi, D. Sexty, and I. O. Stamatescu, "Lattice simulations of real-time quantum fields," *Phys. Rev.* **D75** (2007) 045007, `arXiv:hep-lat/0609058 [hep-lat]`. 35

[88] J. Berges and D. Sexty, "Real-time gauge theory simulations from stochastic quantization with optimized updating," *Nucl. Phys.* **B799** (2008) 306–329, `arXiv:0708.0779 [hep-lat]`. 35, 39

[89] C. Pehlevan and G. Guralnik, "Complex Langevin Equations and Schwinger-Dyson Equations," *Nucl. Phys.* **B811** (2009) 519–536, `arXiv:0710.3756 [hep-th]`. 35

[90] G. Guralnik and C. Pehlevan, "Effective Potential for Complex Langevin Equations," *Nucl. Phys.* **B822** (2009) 349–366, `arXiv:0902.1503 [hep-lat]`. 35

[91] K. Nagata, J. Nishimura, and S. Shimasaki, "Argument for justification of the complex Langevin method and the condition for correct convergence," *Phys. Rev.* **D94** no. 11, (2016) 114515, `arXiv:1606.07627 [hep-lat]`. 35

[92] G. Aarts, F. A. James, E. Seiler, and I.-O. Stamatescu, "Complex Langevin: Etiology and Diagnostics of its Main Problem," *Eur. Phys. J.* **C71** (2011) 1756, arXiv:1101.3270 [hep-lat]. 38

[93] G. Aarts, F. A. James, J. M. Pawlowski, E. Seiler, D. Sexty, and I.-O. Stamatescu, "Stability of complex Langevin dynamics in effective models," *JHEP* **03** (2013) 073, arXiv:1212.5231 [hep-lat]. 38, 113

[94] G. Aarts, P. Giudice, and E. Seiler, "Localised distributions and criteria for correctness in complex Langevin dynamics," *Annals Phys.* **337** (2013) 238–260, arXiv:1306.3075 [hep-lat]. 38, 113

[95] A. Mollgaard and K. Splittorff, "Complex Langevin Dynamics for chiral Random Matrix Theory," *Phys. Rev.* **D88** no. 11, (2013) 116007, arXiv:1309.4335 [hep-lat]. 38, 44

[96] J. Greensite, "Comparison of complex Langevin and mean field methods applied to effective Polyakov line models," *Phys. Rev.* **D90** no. 11, (2014) 114507, arXiv:1406.4558 [hep-lat]. 38

[97] G. Aarts, E. Seiler, D. Sexty, and I.-O. Stamatescu, "Complex Langevin dynamics and zeroes of the fermion determinant," *JHEP* **05** (2017) 044, arXiv:1701.02322 [hep-lat]. 38, 44, 113

[98] E. Seiler, D. Sexty, and I.-O. Stamatescu, "Gauge cooling in complex Langevin for QCD with heavy quarks," *Phys. Lett.* **B723** (2013) 213–216, arXiv:1211.3709 [hep-lat]. 39, 87

[99] G. Aarts, L. Bongiovanni, E. Seiler, D. Sexty, and I.-O. Stamatescu, "Controlling complex Langevin dynamics at finite density," *Eur. Phys. J.* **A49** (2013) 89, arXiv:1303.6425 [hep-lat]. x, xiv, 39, 40, 53, 56, 69, 70, 71, 72, 113

[100] J. Bloch, J. Mahr, and S. Schmalzbauer, "Complex Langevin in low-dimensional QCD: the good and the not-so-good," *PoS* **LATTICE2015** (2016) 158, arXiv:1508.05252 [hep-lat]. 40

[101] S. Schmalzbauer and J. Bloch, "Complex Langevin Dynamics in 1+1d QCD at Non-Zero Densities," *PoS* **LATTICE2016** (2016) 362, arXiv:1611.00702 [hep-lat]. 40

[102] K. Nagata, J. Nishimura, and S. Shimasaki, "Testing a generalized cooling procedure in the complex Langevin simulation of chiral Random Matrix Theory," *PoS* **LATTICE2015** (2016) 156, arXiv:1511.08580 [hep-lat]. 40

[103] K. Nagata, J. Nishimura, and S. Shimasaki, "Gauge cooling for the singular-drift problem in the complex Langevin method - a test in Random Matrix Theory for finite density QCD," *JHEP* **07** (2016) 073, arXiv:1604.07717 [hep-lat]. 40

[104] G. Aarts, F. Attanasio, B. Jäger, and D. Sexty, "Complex Langevin in Lattice QCD: dynamic stabilisation and the phase diagram," *Acta Phys. Polon. Supp.* **9** (2016) 621, arXiv:1607.05642 [hep-lat]. 40, 113

[105] D. Sexty, "Simulating full QCD at nonzero density using the complex Langevin equation," *Phys. Lett.* **B729** (2014) 108–111, arXiv:1307.7748 [hep-lat]. 42, 43, 98

[106] M. Fukugita, Y. Oyanagi, and A. Ukawa, "Langevin Simulation of the Full QCD Hadron Mass Spectrum on a Lattice," *Phys. Rev.* **D36** (1987) 824. 43, 57, 113

[107] K. Binder, "Finite size scaling analysis of Ising model block distribution functions," *Z. Phys.* **B43** (1981) 119–140. 48, 84

[108] G. Aarts, F. Attanasio, B. Jäger, E. Seiler, D. Sexty, and I.-O. Stamatescu, "Towards the heavy dense QCD phase diagram using Complex Langevin simulations," *PoS* **LATTICE2015** (2016) 154, arXiv:1510.09098 [hep-lat]. 49, 53

[109] G. Cella, G. Curci, A. Vicere, and B. Vigna, "The SU(3) deconfining phase transition with Symanzik action," *Phys. Lett.* **B333** (1994) 457–460, arXiv:hep-lat/9405018 [hep-lat]. viii, 49, 50, 51

[110] J. Bloch, J. Glesaaen, O. Philipsen, J. Verbaarschot, and S. Zafeiropoulos, "Complex Langevin simulations of a finite density matrix model for QCD," *EPJ Web Conf.* **137** (2017) 07030, arXiv:1612.04621 [hep-lat]. 63

[111] G. Aarts, F. Attanasio, B. Jäger, E. Seiler, D. Sexty, and I.-O. Stamatescu, "Exploring the phase diagram of QCD with complex Langevin simulations," *PoS* **LATTICE2014** (2014) 200, arXiv:1411.2632 [hep-lat]. 75

[112] G. Aarts, F. Attanasio, B. Jäger, E. Seiler, D. Sexty, and I.-O. Stamatescu, "QCD at nonzero chemical potential: recent progress on the lattice," *AIP Conf. Proc.* **1701** (2016) 020001, arXiv:1412.0847 [hep-lat]. 75

[113] G. Aarts, F. Attanasio, B. Jäger, E. Seiler, D. Sexty, and I.-O. Stamatescu, "The phase diagram of heavy dense QCD with complex Langevin simulations," *Acta Phys. Polon. Supp.* **8** no. 2, (2015) 405, arXiv:1506.02547 [hep-lat]. 75

[114] G. Aarts, F. Attanasio, B. Jäger, E. Seiler, D. Sexty, and I.-O. Stamatescu, "Insights into the heavy dense QCD phase diagram using Complex Langevin simulations," *PoS* **LATTICE2015** (2016) 155, arXiv:1510.09100 [hep-lat]. 75

[115] F. Attanasio and B. Jäger, "Dynamic stabilisation of complex Langevin simulations of QCD, *in preparation*,". 75

[116] D. Sexty. Private communication. xiv, 77

[117] M. Lüscher, "Properties and uses of the Wilson flow in lattice QCD," *JHEP* **08** (2010) 071, arXiv:1006.4518 [hep-lat]. [Erratum: JHEP03,092(2014)]. xiv, 77, 92

[118] M. Luscher and P. Weisz, "Perturbative analysis of the gradient flow in non-abelian gauge theories," *JHEP* **02** (2011) 051, arXiv:1101.0963 [hep-th]. xiv, 77, 92

[119] S. Borsanyi *et al.*, "High-precision scale setting in lattice QCD," *JHEP* **09** (2012) 010, arXiv:1203.4469 [hep-lat]. xiv, 77, 92

[120] A. Barducci, R. Casalbuoni, G. Pettini, and R. Gatto, "Discontinuity theorem for first order phase transitions. Implications for QCD," *Phys. Lett.* **B301** (1993) 95–100, arXiv:hep-ph/9212276 [hep-ph]. 82

[121] R. De Pietri, A. Feo, E. Seiler, and I.-O. Stamatescu, "A Model for QCD at high density and large quark mass," *Phys. Rev.* **D76** (2007) 114501, arXiv:0705.3420 [hep-lat]. 85

[122] **ALPHA** Collaboration, U. Wolff, "Monte Carlo errors with less errors," *Comput. Phys. Commun.* **156** (2004) 143–153, arXiv:hep-lat/0306017 [hep-lat]. [Erratum: Comput. Phys. Commun.176,383(2007)]. 89

[123] D. K. Sinclair and J. B. Kogut, "Exploring Complex-Langevin Methods for Finite-Density QCD," *PoS* **LATTICE2015** (2016) 153, arXiv:1510.06367 [hep-lat]. 98

[124] D. K. Sinclair and J. B. Kogut, "Complex Langevin for Lattice QCD at $T = 0$ and $\mu \geq 0$," *PoS* **LATTICE2016** (2016) 026, arXiv:1611.02312 [hep-lat]. 98

[125] K. M. Bitar *et al.*, "Effects of spatial size, lattice doubling and source operators on the hadron spectrum with dynamical staggered quarks at 6/g**2 = 5.6," *Phys. Rev.* **D49** (1994) 6026–6038, arXiv:hep-lat/9311027 [hep-lat]. 99

[126] P. de Forcrand, S. Kim, and T. Takaishi, "QCD simulations at small chemical potential," *Nucl. Phys. Proc. Suppl.* **119** (2003) 541–543, arXiv:hep-lat/0209126 [hep-lat]. [,541(2002)]. 109

[127] B. Sakita, "Stochastic quantization," in *22nd International Conference on High Energy Physics. Vol. 1: Leipzig, Germany, July 19-25, 1984*, p. I.107. 1984. http://inspirehep.net/record/191267/files/c84-07-19-p107.pdf. 113