

ABSTRACT

Title of dissertation: **SIGN PROBLEMS IN QUANTUM
FIELD THEORY: CLASSICAL
AND QUANTUM APPROACHES**

Scott Lawrence
Doctor of Philosophy, 2020

Dissertation directed by: **Professor Paulo F. Bedaque
Department of Physics**

Monte Carlo calculations in the framework of lattice field theory provide non-perturbative access to the equilibrium physics of quantum fields. When applied to certain fermionic systems, or to the calculation of out-of-equilibrium physics, these methods encounter the so-called sign problem, and computational resource requirements become impractically large. These difficulties prevent the calculation from first principles of the equation of state of quantum chromodynamics, as well as the computation of transport coefficients in quantum field theories, among other things.

This thesis details two methods for mitigating or avoiding the sign problem. First, via the complexification of the field variables and the application of Cauchy's integral theorem, the difficulty of the sign problem can be changed. This requires searching for a suitable contour of integration. Several methods of finding such a contour are discussed, as well as the procedure for integrating on it. Two notable examples are highlighted: in one case, a contour exists which entirely removes the sign problem, and in another, there is provably no contour available to improve the

sign problem by more than a (parametrically) small amount.

As an alternative, physical simulations can be performed with the aid of a quantum computer. The formal elements underlying a quantum computation — that is, a Hilbert space, unitary operators acting on it, and Hermitian observables to be measured — can be matched to those of a quantum field theory. In this way an error-corrected quantum computer may be made to serve as a well controlled laboratory. Precise algorithms for this task are presented, specifically in the context of quantum chromodynamics.

SIGN PROBLEMS IN QUANTUM FIELD THEORY:
CLASSICAL AND QUANTUM APPROACHES

by

Scott Lawrence

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2020

Advisory Committee:
Professor Andrew Baden
Professor Paulo F. Bedaque, Chair/Advisor
Professor Zackaria Chacko
Professor Massimo Ricotti
Professor Raman Sundrum

Acknowledgments

The work described in this thesis was performed in collaboration with others at the University of Maryland and the George Washington University: Andrei Alexandru, Paulo Bedaque, Siddhartha Harmalkar, Neill Warrington, and Yukari Yamauchi. Most especially, credit is due to Henry Lamm, who aside from being an excellent collaborator, is unfailing generous with his time.

That this thesis exists at all is down to the aid and encouragement of a long series of mentors, who helped me along throughout undergraduate and graduate school. In chronological order, I am deeply indebted to: Drew Baden, Raman Sundrum, Victor Yakovenko, Cole Miller, Tom Cohen, and of course my Ph.D. advisor Paulo Bedaque.

Lastly, for their ceaseless encouragement throughout this decade, I am grateful to Brian McPeak, Ophir Lifshitz, and Yukari Yamauchi.

Table of Contents

Acknowledgements	ii
Table of Contents	iii
List of Figures	v
1 Introduction	1
2 Two Views of Lattice Field Theory	4
2.1 Hamiltonian Lattice Field Theory	4
2.1.1 Fermionic Theories	7
2.2 Lattice Actions	8
2.2.1 Fermionic Path Integral	10
2.3 Gauge Theories	11
2.3.1 General Gauge Groups	13
2.3.2 Path Integral	16
3 Computational Difficulties	20
3.1 Monte Carlo Methods	21
3.1.1 Markov-Chain Monte Carlo Methods	22
3.2 Finite Fermion Density	23
3.3 Real-Time Linear Response	26
3.4 Noisy Correlators	29
4 Complexification	31
4.1 A One-Dimensional Example	32
4.2 Cauchy's Integral Theorem	34
4.2.1 Holomorphic Boltzmann Factors and Observables	38
4.2.1.1 Complex Scalar Fields	38
4.2.1.2 Fermionic Determinant and Correlators	39
4.3 Lefschetz Thimbles	41
4.4 Holomorphic Gradient Flow	43
4.4.1 Algorithmic Costs	44
4.5 Machine Learning Manifolds	45
4.5.1 Feed-Forward Networks	46
4.5.2 Training	47
4.6 Manifold Optimization	48

4.6.1	Another View of the Flow	50
4.7	Application to the Thirring Model	51
4.7.1	Field Complexification	53
4.7.2	An Ansatz	54
4.7.3	Phase Diagram	55
4.8	Optimal Manifolds	56
4.8.1	Lefschetz Thimbles Are Not Optimal	57
4.8.2	A Complexification-Immune Sign Problem	59
5	Quantum Simulations	62
5.1	Digital Quantum Computers	63
5.1.1	A Single Qubit	63
5.1.2	Coupling Many Qubits	65
5.1.3	Some Simple Algorithms	66
5.1.3.1	Controlled Nots	66
5.1.3.2	Circuits From Circuits	67
5.1.3.3	Quantum Fourier Transform	69
5.2	Quantum Simulations in General	70
5.2.1	Mapping Hilbert Spaces	70
5.2.2	Suzuki-Trotter Decomposition	71
5.2.3	Aside: Disordered Potentials	72
5.2.3.1	Simulation	73
5.2.3.2	A Pseudo-Random Number Generator	76
5.2.3.3	Demonstration	78
5.3	Simulating a Field Theory	79
5.3.1	Linear Response, Two Ways	81
5.3.2	Measuring Masses	83
5.3.3	Scalar Fields	85
5.3.4	Fermions	86
5.4	Simulating a Gauge Theory	87
5.4.1	Time Evolution	90
5.4.2	Gauge Invariance	90
5.4.3	Adding Fermions	92
5.4.4	Demonstration: D_4 Gauge Theory	93
5.5	Simulations of QCD	95
5.5.1	The Valentiner Group	96
5.5.2	State Preparation	100
5.5.3	Hadronic Tensor	102
5.5.3.1	Parton Distribution Function	104
5.6	Avoiding State Preparation	105
A	Grassmann Numbers	108
A.1	Integration	108
A.2	Coherent States	109
	Bibliography	111

List of Figures

2.1	The lattice associated to a gauge theory in two dimensions. Each link is an independent (commuting) degree of freedom. The lower-leftmost plaquette is given by $\text{Re Tr } U_3^\dagger U_2 U_1 U_0^\dagger$.	11
4.1	The sign problem, and its removal, of a gaussian integral. On the left is the Boltzmann factor e^{-S} without any contour deformation performed (black) and with the optimal contour (red). The original manifold of integration (black) and the deformed contour (red) are on the right.	33
4.2	The asymptotically safe regions for the Gaussian integral.	37
4.3	Figure from [52] depicting a feed-forward network, used to define a manifold in the complexified field space.	46
4.4	Figure from [64] of the chiral symmetry breaking phase transition of the Thirring model. On the left, the condensate $\langle \bar{\psi}\psi \rangle$, as a function of chemical potential μ on $\beta \times 6^2$ lattices. On the right, the full T - μ plane for the same spatial volume. The central band indicates the location of $\langle \bar{\psi}\psi \rangle_{\mu,T} = 0.5 \langle \bar{\psi}\psi \rangle_0$; the thinner lines on either side indicate the width of the crossover	56
4.5	Thimbles and an optimized manifold from the ansatz 4.35 for the heavy-dense limit of the Thirring model, or equivalently the one-site Thirring model.	57
5.1	A two-qubit circuit, implementing time-evolution under the Hamiltonian $H = \sigma_x \otimes \sigma_x$.	65
5.2	Implementation of the Toffoli gate from 1- and 2-qubit gates.	67
5.3	Construction of a 4-controlled not gate from the Toffoli gate, with the aid of ancilla. The ancillary qubits are on the top; the target is the bottom qubit.	68
5.4	Construction of the controlled Hadamard (top) and controlled $\pi/8$ (bottom) gates, from the primitive gateset of H , $\mathcal{R}_Z(\theta)$, and CX . For brevity, an additional gate is defined as $S = T^2$, and we have used $X = HT^4H$.	69

5.5	Quantum circuits for the simulation of the Anderson model. The time-evolution is trotterized: on top is the kinetic piece, and on bottom is the evolution under a random Hamiltonian. For brevity, $R_X(\theta)$ and $R_Z(\theta)$ denote $e^{i\theta\sigma_x}$ and $e^{i\theta\sigma_z}$, respectively. The construction of a pseudo-random permutation operator P is discussed in the text. . . .	76
5.6	Figure from [79], showing average distance as defined by (5.13), as a function of the disorder parameter W , for a 16^3 lattice, in the limit of long time evolution. The blue triangles show results obtained with a conventional PRNG, and the red squares give results obtained with the circuit-based PRNG described in Section 5.2.3.2. The dashed line shows the delocalized limit for D	78
5.7	Figure from [79], showing average distance as measured by (5.13), as a function of evolution time for a 4-site lattice with disorder parameter $W = 5$ and Trotterization step size $\Delta t = 0.2$	79
5.8	Quantum circuit for a single Trotter step of time-evolution for the Heisenberg spin chain with Hamiltonian (5.14), on four lattice sites. .	81
5.9	Quantum circuit for a single Trotter step of the time-evolution of \mathbb{Z}_2 gauge theory with one plaquette.	88
5.10	Circuits implementing the time-evolution of a pure-gauge lattice field theory. The first circuit implements the quantum-mechanical kinetic term, and the second the quantum-mechanical potential term associated to a single plaquette $\text{Re Tr } U_{13}^\dagger U_{34}^\dagger U_{24} U_{12}$. Note that in these circuits, the primitive object is a G -register, denoted with a doubled line.	90
5.11	The lattice geometry of the D_4 gauge theory simulated here. Dashed lines indicate repeated links due to the periodic boundary conditions.	94
5.12	Figure from [20]: simulation of the two-plaquette D_4 gauge theory. The expectation value of one plaquette as a function of time. The exact result is shown in black, with sampled data shown in red (blue) for a Trotterization time-step of $\Delta t = 0.2$ ($\Delta t = 0.5$).	95
5.13	Figure from [100] of the average energy per plaquette vs β_0 for $SU(3)$, \mathcal{V} with $\beta_1 = 0$, and \mathcal{V} with $\beta_1 = -0.6$	97
5.14	Continuum extrapolations of $T_c\sqrt{t_0}$ for $SU(3)$ and $S(1080)$ gauge theories. The $S(1080)$ result is shown in blue, with the extrapolated value marked by the solid shape; the $SU(3)$ extrapolated results are also indicated by solid shapes.	99

Chapter 1: Introduction

The majority of observed non-gravitational phenomena in laboratories and the universe are believed to be described by the quantum field theory of the standard model. The standard model can be crudely divided into two sectors, one governed by the electroweak force and the other by the strong force. The electroweak force possesses a small expansion parameter. As a result, the physics of electroweak phenomena can be computed in perturbation theory. In the other sector, governed by the strong force, there is no small expansion parameter available.

The physics of the strong force governs phenomena ranging from heavy ion collisions to the structure of neutron stars. Furthermore, various extensions of the standard model (of relevance, for instance, in the search for dark matter [1]) also involve no small expansion parameter; indeed there is no *a priori* reason why we should expand physics beyond the standard model to be perturbative. Computational methods capable of probing non-perturbative physics are therefore critical not only for understanding the physics of the standard model, but also for designing tests for beyond-standard-model physics.

The main nonperturbative tool for field theory is the lattice. Whereas a quantum field theory, as usually understood, involves arbitrarily small distance scales and

therefore arbitrarily high momenta (and with them, a host of technical complications), the lattice regularization removes momenta above a certain cutoff, reducing the physics of field theory to the task of evaluating a finite-dimensional (albeit the number of dimensions is large) integral. The lattice regularization, in addition to being a central subject of the formal mathematical study of field theories [2], is therefore a prime candidate for *numerical* study, the subject of this thesis.

The problems of lattice field theory are problems of evaluating high-dimensional integrals. Monte Carlo simulations of lattice field theory provide the main tool by which nonperturbative physics may be accessed in a practical way [3,4]. In certain regimes, known algorithms for simulating lattice field theory encounter severe computation obstacles. The most prominent of these is the *fermion sign problem*: when simulating field theories with a finite density of fermions, evaluating the integral by usual methods requires resolving fine cancellations of large terms. This prevents calculations, for example, of the equation of state of a neutron star. Another sign problem occurs when applying lattice methods to determine transport coefficients of quantum fluids, restricting the ability to compare the experimentally observed behavior of heavy ion collisions to theoretical calculations.

This thesis concerns the task of performing nonperturbative computations in regimes where Monte Carlo methods encounter a sign problem. One approach is to re-structure the integral in such a way as to alleviate the sign problem. We will see in Chapter 4 that the methods of complex analysis — in particular, an application of a multidimensional generalization of Cauchy’s integral theorem — may be used to construct such algorithms. The rapid progress being made on the construction of

practical quantum computers [5,6] suggests another promising avenue of attack. As a quantum computer is, in fact, a quantum-mechanical system, it is an ideal device to simulate other quantum-mechanical systems [7,8]. This suggests the possibility of using a quantum computer as a tool for simulating lattice quantum field theory [9].

This thesis is structured in three broad parts. Chapters 2 and 3 introduce lattice field theory, the standard context in which nonperturbative quantum field theory computations are performed, and describe the ways in which sign problems obstruct calculations performed at finite fermion density or in real time. Chapter 4 describes several closely related methods for thwarting the sign problem on classical computers, based on the complexification of the path integral. These methods are applicable to a wide range of theories, but the Thirring model (in both $1 + 1$ and $2 + 1$ dimensions) is investigated in detail. Finally in Chapter 5, quantum computing is introduced as a tool for the simulation of field theories, which entirely circumvents the sign problem. The necessary algorithms for simulating QCD as a lattice gauge theory are detailed, and the costs (in terms of quantum gates and qubits) are estimated.

Chapter 2: Two Views of Lattice Field Theory

In this chapter we describe lattice-regularized field theory from two perspectives: first the Hamiltonian formulation, in which time is continuous and the field theory is a regular quantum mechanical system, and second from a lattice action, where space and time are both discretized, as is appropriate for a relativistic theory.

2.1 Hamiltonian Lattice Field Theory

We begin our overview of the Hamiltonian formulation of lattice field theory with the example of a noninteracting real scalar field [10]. The Hamiltonian of a noninteracting scalar field in the continuum may be written in momentum space as

$$H = \sum_k \frac{1}{2} \Pi_k^2 + \frac{\sqrt{m^2 + k^2}}{2} \Phi_k^2, \quad (2.1)$$

where (Φ_k, Π_k) are the position and conjugate momentum coordinates of harmonic oscillators indexed by momentum k .

This form makes clear the noninteracting structure of the theory. Each momentum mode can contain any non-negative number of particles, and in a mode with momentum k , each particle contributes an energy of $\sqrt{m^2 + k^2}$, where m is the

mass of the field. The possible momenta in a box of side length L are $k = \frac{2\pi n}{L}$ where $n = 0, 1, 2, \dots$. The same theory may also be rewritten in position space, yielding the Hamiltonian

$$H = \int d^d x \frac{1}{2} \pi(x)^2 + \frac{1}{2} (\partial_i \phi(x)) (\partial_i \phi(x)) + \frac{m^2}{2} \phi(x)^2, \quad (2.2)$$

where the spatial index is summed over $i = 1, 2, 3$.

This field theory is ‘free’ in the sense that the partition function factorizes into a product of factors, each involving only one momentum mode. Introducing an interaction directly in the continuum field theory is technically difficult. In the example discussed here of a real scalar field, it is believed that no interacting theory can be constructed in three spatial dimensions [11]. Moreover, the continuum formulation, with an infinite number of degrees of freedom even in a finite-sized box, isn’t amenable to numerical simulation.

For these reasons, we introduce the lattice regularization of the field theory. In the case of scalar field theory, this regularization is obtained by placing a cutoff on the momenta k included in the sum of (2.1). Equivalently, and more conveniently, the integral in (2.2) is changed to a discrete sum over lattice sites, and the derivative is changed to a finite difference. The resulting lattice hamiltonian is

$$H = \sum_x \left[\frac{1}{2} \pi_x^2 + \frac{m^2}{2} \phi_x^2 \right] + \sum_{\langle xy \rangle} \frac{1}{2} (\phi_x - \phi_y)^2, \quad (2.3)$$

where the second sum denotes the sum over all pairs of neighboring lattice sites. It

is this object (and others like it) that will serve as the starting point for numerical work. At this point, an interaction is introduced by adding a term proportional to $\phi(x)^4$ to the Hamiltonian¹.

The original field theory consisted of non-interacting harmonic oscillators, each associated to a separate momentum mode. Before introducing an interaction, the same decomposition could be done for the Hamiltonian (2.3), now with a finite number of momentum modes. However, it will be useful to take an alternate view of the system, where each *position* is associated to a harmonic oscillator. Even in the absence of an interaction, these oscillators are coupled by the finite difference term. When the interaction is introduced, the oscillators become anharmonic. Allowing for an arbitrary potential $V(\phi)$, the final lattice hamiltonian is

$$H = \sum_x \left[\frac{1}{2} \pi_x^2 + V(\phi_x) \right] + \sum_{\langle xy \rangle} \frac{1}{2} (\phi_x - \phi_y)^2. \quad (2.4)$$

On a lattice with N sites, the Hilbert space is the tensor product of N copies of that of the harmonic oscillator. The operators ϕ_x and π_x act on the portion of the Hilbert space associated to site x , and commute with ϕ_y and π_y at any site $y \neq x$. Moreover, these operators can be written in terms of the raising and lowering operators associated to a site:

$$\phi_x = \sqrt{\frac{1}{2m}} (a_x^\dagger + a_x) \quad \text{and} \quad \pi_x = i\sqrt{2m} (a_x^\dagger - a_x). \quad (2.5)$$

¹Any polynomial of even degree greater than 2 will do.

2.1.1 Fermionic Theories

For a theory of lattice fermions, due to Pauli exclusion, at most one fermion will be permitted per degree of freedom. (There may be multiple degrees of freedom per lattice site; for example, a lattice site may be occupied by both a spin-up and a spin-down fermion.) The Hilbert space of one degree of freedom, therefore, is two-dimensional, and the Hilbert space of the full lattice is again the tensor product of N copies of that local Hilbert space.

The two-dimensional fermionic Hilbert space associated to degree of freedom i is acted on by creation and annihilation operators a_i^\dagger and a_i , respectively. These operators obey the anticommutation relations

$$\{a_i, a_j\} = 0 \text{ and } \{a_i^\dagger, a_i\} = \delta_{ij}. \quad (2.6)$$

Most theories of physical interest have multiple fermionic degrees of freedom per lattice site. Typically these degrees of freedom are related by some global symmetry.

A typical theory of lattice fermions is given by the Hamiltonian

$$H = -t \sum_{\langle xy \rangle, s} (a_{xs}^\dagger a_{ys} + a_{ys}^\dagger a_{xs}) - U \sum_x a_{x\downarrow}^\dagger a_{x\downarrow} a_{x\uparrow}^\dagger a_{x\uparrow}, \quad (2.7)$$

where each lattice site (indexed by x or y) contains a spin-up and spin-down degree of freedom, and the first sum is taken over all pairs of neighboring lattice sites, and

$s = \downarrow, \uparrow$. This Hamiltonian describes the nonrelativistic fermions of the Hubbard model [12], a common target of numerical work [13].

Due to the nature of the dispersion relation of lattice fermions, the number of low-energy modes of such a Hamiltonian will generically be larger than naive counting would suggest [14]. Various methods for removing these modes exist [3, 4]; alternatively, one may simply accept that the theory being simulated has more particles than originally intended.

2.2 Lattice Actions

Quantum field theories can be described by a path integral; in this section we derive the lattice path integral. We begin by considering the thermal properties of a lattice field theory at temperature $T \equiv 1/\beta$, defined by the hermitian operator $\rho = e^{-\beta H}$, termed the *density matrix*. The thermal partition function is given by $Z = \text{Tr } \rho$, and thermal expectation values are given by various derivatives of $\log Z$.

To obtain the path integral, we expand the trace by summing over all possible intermediate states:

$$Z = \text{Tr} \left(e^{-\Delta t H} \right)^{N_t} = \sum_{\Psi_1, \Psi_2, \dots} \langle \Psi_1 | e^{-\Delta t H} | \Psi_2 \rangle \dots \langle \Psi_N | e^{-\Delta t H} | \Psi_1 \rangle, \quad (2.8)$$

where $\Delta_t N_t = \beta$, and use has been made of the completeness relation $I = \sum_{\Psi} |\Psi\rangle\langle\Psi|$. The operator $e^{-\Delta t H}$ is termed the *transfer matrix*.

In the case of the bosonic theory defined by Hamiltonian (2.4), an appropriate set of states is given by the simultaneous eigenstates of the field operators $\psi(x)$,

denotes $|\psi\rangle$. The resulting resolution of the identity is

$$I = \int d^V \phi |\phi\rangle \langle \phi|. \quad (2.9)$$

It is natural to take Δ_t to be equal to 1, and after approximation by the Suzuki-Trotter decomposition [15, 16], the resulting lattice path integral is

$$Z = \int d^{\beta V} \phi \exp \left\{ - \sum_{\langle xy \rangle} \frac{(\phi_x - \phi_y)^2}{2} - \sum_x V(\phi_x) \right\}. \quad (2.10)$$

The continuum limit of (2.10) is approached by tuning the potential $V(\phi)$ such that correlation functions (e.g. $\langle \phi(x)\phi(0) \rangle$) decay slowly with $|x|$. As mentioned previously, the process of taking a continuum limit will not be of much interest here; however it must be noted that the continuum limit obtained in this way need not be the same as that obtained by working with the Hamiltonian (2.4) directly. In order for the two methods to be equivalent, we must first take $\Delta t \rightarrow 0$ in the path integral (termed the Hamiltonian limit) and only then performing the tuning of the potential.

There is always some ambiguity in constructing the path integral: for instance, the Suzuki-Trotter decomposition is not unique, and there are many possible resolutions of the identity to be used. If we are interested in the Hamiltonian limit of the lattice theory, then this is not an issue at all. In fact, in that limit, the multiple lattice path integrals simply correspond to distinct classical algorithms for computing thermodynamic quantities. Difficulties can potentially arise when taking

a continuum limit directly from the lattice theory, as there is generally no guarantee that all of the possible lattice path integrals will yield the same continuum limit.

2.2.1 Fermionic Path Integral

The fermionic path integral is derived in the same way as (2.10) above, but with a particular choice of completeness relation. Consider first a single fermionic degree of freedom. Define the state $|0\rangle$ to be the eigenstate of $a^\dagger a$ with eigenvalue 0 — this is the occupation number basis. We introduce the coherent state $|\psi\rangle$ and its dual $\langle\psi|$, defined by

$$|\psi\rangle \equiv e^{-\psi a^\dagger} |0\rangle \quad \text{and} \quad \langle\psi| \equiv \langle 0| e^{-a\psi^\dagger}, \quad (2.11)$$

where ψ and ψ^\dagger are independent Grassmann variables (see Appendix A). With these definitions, a new completeness relation is available:

$$I = \int d\psi^\dagger d\psi e^{-\psi^\dagger\psi} |\psi\rangle \langle\psi|. \quad (2.12)$$

For a lattice with multiple fermionic degrees of freedom, we introduce a pair of grassman variables ψ_i and ψ_i^\dagger for each. The completeness relation for the full lattice system is

$$I = \int \prod_i \left(d\psi_i^\dagger d\psi_i \right) e^{-\sum_i \psi_i^\dagger \psi_i} |\psi_1 \cdots \psi_N\rangle \langle\psi_1 \cdots \psi_N| \quad (2.13)$$

where the products and sums run over all N fermionic degrees of freedom i . At this point the derivation of the path integral proceeds just as for a scalar field theory,

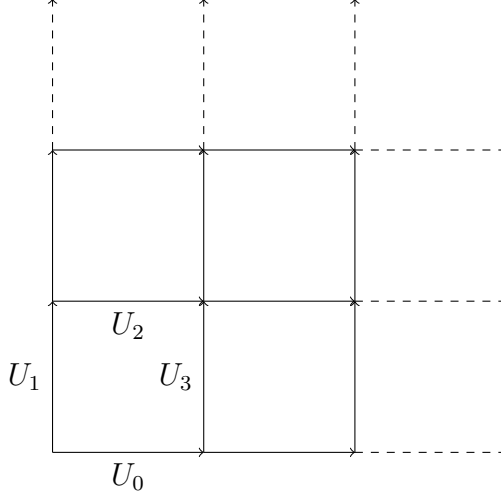


Figure 2.1: The lattice associated to a gauge theory in two dimensions. Each link is an independent (commuting) degree of freedom. The lower-leftmost plaquette is given by $\text{Re Tr } U_3^\dagger U_2 U_1 U_0^\dagger$.

using the completeness relation of (2.13).

2.3 Gauge Theories

We now return to the Hamiltonian picture. Some Hamiltonians possess locally conserved charges. Perhaps the simplest example is \mathbb{Z}_2 gauge theory introduced by Wegner [17]. This theory (like other lattice gauge theories we will discuss) lives on the lattice shown in Figure 2.1; to each link ℓ is associated a two-dimensional Hilbert space acted on by local Pauli operators $\sigma_x(\ell), \sigma_y(\ell), \sigma_z(\ell)$. The operators at separate links are mutually commuting. The Hamiltonian of the theory is

$$H = \sum_{\ell} \sigma_x(\ell) + \sum_P \sigma_z(P_1) \sigma_z(P_2) \sigma_z(P_3) \sigma_z(P_4), \quad (2.14)$$

where the first sum is taken over all links in the lattice, and the second is taken over all ‘plaquettes’ consisting of four links arranged in a square.

For each site r of the lattice, this Hamiltonian commutes with the operator

$$G(r) = \prod_{\ell \text{ at } r} \sigma_x(r), \quad (2.15)$$

where the product is taken over all links which have an endpoint at r . The action of this operator is called a gauge transformation; in the Z basis, it has the effect of flipping any link in contact with r and leaving all other links invariant.

The group of symmetries associated to gauge transformations, for this theory, is $(\mathbb{Z}_2)^V$, where V is the number of sites on the lattice.

A local conservation law is in some sense more restrictive than a global symmetry. Consider a scattering experiment: we begin with a vacuum, introduce particles ‘at infinity’ (at the boundary of a large box), and then observe what particles are measured at later times, again ‘at infinity’. This is typical of physical experiments, in which the experimentalist can only act on the boundary of the laboratory. Such experiments can introduce charged particles into the theory, and thus explore the sectors of Hilbert space labelled by different global charges; however, locally conserved charges cannot be introduced this way. Thus, only one sector of Hilbert space can be considered physically relevant.

The ‘physical sector’ of Hilbert space is taken to be the space of vectors invariant under gauge transformations; that is, the set of vectors that are simultaneous eigenvectors of all $G(r)$, with eigenvalue 1. The remainder of the full Hilbert space

is physically irrelevant. We consider it to exist only for convenience in writing down the Hamiltonian (2.14).

2.3.1 General Gauge Groups

In general, a gauge theory can be defined for any group G chosen to be the gauge group. Lattice gauge theories for continuous groups were initially formulated from the path integral; here we begin with the Hamiltonian formulation [18] and obtain the path integral as a consequence.

As with the \mathbb{Z}_2 gauge theory, we define the gauge theory for gauge group G on a rectangular (or cubic) lattice, such as shown in Figure 2.1. To each link ℓ is associated a local Hilbert space $\mathcal{H}_\ell = \mathbb{C}G$, the space of square-integrable complex-valued functions on G . A large Hilbert space $\mathcal{H} = \mathcal{H}_\ell^{\otimes N}$ is constructed from the tensor product of N copies of \mathcal{H}_ℓ , where there are N total links on the lattice.

A local Hilbert space $\mathbb{C}G$ has a basis² consisting of one state $|U\rangle$ for every element $U \in G$. The basis is the eigenbasis of a G -valued operator \hat{U} defined by $\hat{U}|U\rangle = U|U\rangle$. On a lattice with many links, the generalization of this basis is one state $|U(\ell)\rangle$ for every function U from the set of links to the gauge group G . This set is a basis for \mathcal{H} . To every link is associated an operator \hat{U}_ℓ , defined as \hat{U} above acting on the Hilbert space of that link.

The Hilbert space of the physical theory is in fact a subspace of \mathcal{H} . To con-

²Strictly speaking, for a continuous group G , this set of vectors is overcomplete and does not lie in the Hilbert space, as they are not normalizable. Much as for ordinary quantum mechanics, where eigenstates of the position operator are referred to as $|x\rangle$, we can disregard this issue without affecting any of the results. It is most visible in the fact that the dimension of the Hilbert space of a lattice is countably infinite, whereas the set of basis vectors we use is uncountably infinite.

struct the physical Hilbert space, we need to consider the group of gauge symmetries. For any site x and $V \in G$, we define an operator $\phi_x(V)$ on \mathcal{H} which performs a gauge transformation by V at x . The operator acts on each link independently. Links going out of site x transform in the right regular representation; links going into site x transform in the adjoint of the left regular representation (other links do not transform, or rather, transform in the trivial representation):

$$\phi_x(V) |U_{xy}\rangle = |UV\rangle \quad \text{and} \quad \phi_x(V) |U_{yx}\rangle = |V^\dagger U\rangle. \quad (2.16)$$

Here we have introduced the subscript notation xy to denote a link from y to x . A general gauge transformation is obtained by specifying an element V at every site of the lattice³. Under such a gauge transformation, the transformation law for the Hilbert space is

$$\phi(V) |\cdots U_{xy} \cdots\rangle = |\cdots (V_x^\dagger U_{xy} V_y) \cdots\rangle. \quad (2.17)$$

The physical Hilbert space \mathcal{H}_P is the subspace of \mathcal{H} consisting only of states invariant under gauge transformations. This can be defined constructively with the aid of the gauge projection operator

$$P = \int dV \phi(V) \quad (2.18)$$

which has the effect of integrating over all possible gauge transformations.

³Performing a ‘constant’ gauge transformation, where the group element V is the same at every site, has no effect if g is in the center Z of G . Therefore, the full group of gauge symmetries is not G^V , but rather G^V/Z .

The lattice Hamiltonian for gauge group G is

$$H = \sum_{\ell} \nabla_{\ell}^2 + \sum_P \text{Re Tr } U_P. \quad (2.19)$$

where, just as for the \mathbb{Z}_2 gauge theory, the first sum is over all links and the second sum goes over all plaquettes P . The operator U_P is defined as the product of the U_{ℓ} operators for links going around the plaquette: for the plaquette shown in Figure 2.1 we have $U_P = U_0 U_1^{\dagger} U_2^{\dagger} U_3$. (Under the trace, the choice of starting link does not matter.)

For a continuous gauge group G , the operator ∇_{ℓ}^2 is the Laplace-Beltrami operator; i.e. the kinetic energy of the wavefunction on the curved surface G . This is the generalization of the Laplacian for a curved manifold. For discrete gauge groups, an appropriate generalization of this is obtained by noting that the Laplace-Beltrami operator is diagonal in Fourier space, and is proportional to the identity when restricted to any irreducible representation. It remains to pick one real number for each irreducible representation of G , and *any* such choice will yield a gauge-invariant Hamiltonian.

Viewing the lattice theory as a regular quantum mechanical system, the first term is a kinetic term (momentum squared on a curved manifold), and the second is a potential term, defining how the degrees of freedom are coupled. In the language of field theory, and specifically by analogy with the $U(1)$ gauge theory of electromagnetism, the first term is the electric term and the second the magnetic term, so that the Hamiltonian can be re-written as $H = \frac{1}{2}E^2 + \frac{1}{2}B^2$.

Each term in the lattice Hamiltonian (2.19) is individually gauge-invariant. The kinetic term at link ℓ is invariant under any rotation of the group G , and therefore under all gauge transformations. The plaquette term in the Hamiltonian is invariant under gauge transformation because each closed path (termed a Wilson loop) must go into and out of every link, and therefore the gauge transformation is immediately undone.

The gauge theories of greatest physical interest are those with continuous Lie groups. The gauge group of quantum chromodynamics is $SU(3)$; that of the standard model is $U(1) \times SU(2) \times SU(3)$. Much of the focus of numerical simulation work is on the calculation of quantities in the $SU(3)$ gauge theory.

2.3.2 Path Integral

We now derive lattice path integral for a gauge theory (following [19,20]), and discuss the matter of gauge fixing.

The construction of the transfer matrix T presented here, differs slightly from the usual one [21–25] in that T is defined on the entire space \mathcal{H} . The usual transfer matrix is defined only on the physical Hilbert space, and may be obtained by projection with P .

Fixing a timestep Δt , the transfer matrix is an approximation to imaginary-time evolution: $T \approx e^{-\Delta t H}$. We will work in the fiducial basis of \mathcal{H} of eigenstates $|U\rangle$ of the field operators. We construct the transfer matrix in this basis from separate

kinetic T_K and potential T_V contributions via Suzuki-Trotter approximation [15,16]

$$T = T_V^{1/2} T_K T_V^{1/2} \quad (2.20)$$

where the potential evolution resembles the product of spatial plaquettes that appears in the Wilson action

$$\langle \tilde{U}_{12} \dots | T_V | U_{12} \dots \rangle = \delta_{U_{12} \dots}^{\tilde{U}_{12} \dots} \exp \left(\frac{\Delta t}{ag^2} \sum_x W_{\mu\nu}(x) \right). \quad (2.21)$$

Here a is the spatial lattice spacing, not necessarily equal to Δt . We have borrowed from lattice field theory the Wilson plaquette

$$W_{\mu\nu}(x) = \text{Re} \text{Tr} [U_{x,x+\hat{\nu}}^\dagger U_{x+\hat{\nu},x+\hat{\mu}+\hat{\nu}}^\dagger U_{x+\hat{\mu},x+\hat{\mu}+\hat{\nu}} U_{x,x+\hat{\mu}}], \quad (2.22)$$

and μ, ν are restricted to space-like directions. The kinetic evolution acts on each link independently.

$$\langle \tilde{U}_{12} \dots | T_K | U_{12} \dots \rangle = \prod_{\langle ij \rangle} e^{\frac{a}{g^2 a_0} \text{Re} \text{Tr} [\tilde{U}_{ij}^\dagger U_{ij}]} \quad (2.23)$$

Note that $[T, P] = 0$ due to the fact that T_V and T_K individually commute with P .

At this point a path integral may be obtained from the approximate partition function $\text{Tr} T^{-\beta/\Delta t}$; however, this partition function includes contributions from non-physical gauge-variant states. The correct partition function is projected onto the ground state. There is some freedom in how we do this. The most straight-forward

approach is to insert a single projection operator, writing $Z = \text{Tr} e^{-\beta H} P$. This yields the correct physics, but the resulting lattice path integral is awkward to work with, because the coupling between the first and last time-slices will be different than the coupling between any other pair. (A calculation performed with this path integral is partly gauge-fixed to the $A_0 = 0$ gauge.) It is instead conventional to insert multiple projection operators, writing

$$Z = \text{Tr}_P T^\beta = \text{Tr}(PTP)^\beta \quad (2.24)$$

This is possible because the projection operator commutes with the transfer matrix T , and in fact with T_K and T_V individually. Projecting the kinetic part of the transfer matrix yields

$$\langle \tilde{U}_{12} \dots | PT_K P | U_{12} \dots \rangle = \int dV \prod_{\langle ij \rangle} e^{\frac{\alpha}{g^2 a_0} \text{Re} \text{Tr} [\tilde{U}_{ij}^\dagger V_i^\dagger U_{ij} V_j]}. \quad (2.25)$$

The V_i are group elements that live at a single site in the Hamiltonian picture. Visualizing a Euclidean lattice (with a separate time-like dimension), the V_i connect a lattice site on one time-slice to the same site on the next time-slice. They constitute a gauge-transformation performed in going from one time-slice to the next. This is usually visualized in the form of timelike links connecting one spatial slice to the next.

The resulting partition function is

$$Z = \left(\int_G dU_{12} \cdots \right) \exp \left[\frac{1}{g^2} \sum \text{Re Tr} \prod U_{ij} \right]. \quad (2.26)$$

Here Tr_P denotes the trace over only the physical subspace \mathcal{H}_P , and the sum is taken over both spatial and temporal plaquettes on a $d + 1$ lattice.

It remains to show that, for vanishing temporal lattice spacing $\Delta t \rightarrow 0$, the transfer matrix corresponds exactly to imaginary time under the gauge Hamiltonian. This is done in detail by Creutz [25] for the gauge-fixed transfer matrix. The result for the gauge-free transfer matrix used here is the same. The potential transfer matrix T_V is exactly $e^{-\Delta t H_V}$. The correspondence between H_K and T_K is not exact, but indeed $\lim_{\Delta t \rightarrow 0} T_K = e^{-\Delta t H_K}$.

Chapter 3: Computational Difficulties

The most straightforward method for studying the physics of a lattice Hamiltonian is to construct the Hamiltonian as an explicit matrix, and perform computational linear algebra on that matrix. For instance, the matrix may be diagonalized to reveal the masses of particles and bound states. Such methods are in practice useless for three-dimensional field theories: linear algebra algorithms scale polynomially with the dimension of the vector space in question, and the Hilbert space of a lattice theory is exponential in the volume of the lattice¹.

The most widely-used nonperturbative tool for studying lattice QCD in practice is the Markov chain Monte Carlo (MCMC) method. Much of the time, this algorithm scales polynomially with the volume of the system being studied, rather than exponentially. At finite density of relativistic fermions, or when studying real-time evolution, the MCMC method reverts to exponential scaling. In these cases, we are left without a general nonperturbative method.

¹Or, for theories with continuous degrees of freedom, the Hilbert space is infinite.

3.1 Monte Carlo Methods

Monte Carlo methods for lattice field theory are based on the path integral representation of the partition function. For concreteness, we consider here a scalar field theory, but the same ideas generalize to gauge theories and theories of interacting fermions (discussed in more detail in Section 3.2 below).

The lattice partition function for a theory of one real scalar field, as described by the Hamiltonian (2.4), is

$$Z[J] = \int d^{\beta V} \phi e^{-S[\phi]} e^{\sum_x J_x \phi_x}$$

where $S[\phi] = \sum_{\langle xy \rangle} \frac{(\phi_x - \phi_y)^2}{2} + \sum_x \frac{m^2}{2} \phi^2 + \sum_x \lambda \phi_x^4.$ (3.1)

The functional S is referred to as the action. This expression is only an approximation (à la Suzuki-Trotter) to the true partition function of the Hamiltonian, which is itself only an approximation to the continuum theory. After performing a calculation, one must extrapolate to the continuum and infinite volume limits to obtain physically meaningful results. The process of extrapolation is largely independent from the rest of the calculation, and we will ignore it for the remainder of this chapter.

The partition function (3.1) couples the fields linearly to a spacetime-dependent source field J . Expectation values are obtained by differentiating $\log Z$ with respect to J . For an arbitrary observable \mathcal{O} (typically some polynomial of the fields), the

expectation value is given by

$$\langle \mathcal{O} \rangle = \frac{\int d^{\beta V} \phi e^{-S[\phi]} \mathcal{O}[\phi]}{\int d^{\beta V} \phi e^{-S[\phi]}}. \quad (3.2)$$

When the action $S[\phi]$ is guaranteed to be real, the “Boltzmann factor” e^{-S} is non-negative, and this expectation value may be viewed as an expectation value over the probability distribution proportional to e^{-S} . It follows that, to calculate arbitrary expectation values, one need only sample from the distribution given by the Boltzmann factor.

3.1.1 Markov-Chain Monte Carlo Methods

A Markov chain is a discrete-time stochastic process in which the state at time $t+1$ depends only on the state at time t . The chain is defined by a matrix P_{ij} giving the probability of transitioning to state i at step $t+1$, given that the state was j at step t . The matrix P should be thought of as a linear operation on the space of probability distributions.

The long-time behavior of a Markov chain is determined by the eigenvector of P with largest eigenvalue. Markov-Chain Monte Carlo (MCMC) algorithms sample from a distribution by setting up a Markov chain with the target distribution as this eigenvalue. For many distributions encountered in practice, including many lattice field theories, the Markov chain mixes in time polynomial in the volume (and all other physical parameters), making these methods viable for simulations.

The particular method most widely used in lattice field theory is the Metropolis

algorithm [26].

3.2 Finite Fermion Density

When simulating a field theory with fermions, the fermions are typically integrated out analytically before performing the numerical integral [3, 4]. An example of this is a simulation of $SU(3)$ gauge theory with fermions, i.e. QCD. The lattice action is

$$S = \sum_P \text{Re Tr } P + \sum_i m \bar{\psi}_i \psi_i + \frac{1}{2} \sum_{i,\mu} [\bar{\psi}_i U_{i,i+\hat{\mu}} \gamma^\mu \psi_{i+\hat{\mu}} + \text{h.c.}] . \quad (3.3)$$

Here the first sum is over all plaquettes, the second over all sites, and the third over all sites and spacetime directions μ . Because this action is only quadratic in the fermion fields, they can be integrated out analytically for any fixed configuration of the gauge fields. This yields the lattice partition function

$$Z = \int dU e^{-S[U]} \det D[U] \equiv \int dU e^{-S_{\text{eff}}[U]} \quad (3.4)$$

where S_G is the pure-gauge piece of the original action, and $D[U]$ is the gauge-field-dependent matrix that gave the quadratic part of the action.

Although this saves us from having to work directly with anticommuting numbers, it introduces a new complication: the determinant of D may not be a non-negative real number, in which case $e^{-S} \det D$ cannot be interpreted as a probability distribution. This situation is termed a *fermion sign problem*. In the context of rel-

ativistic field theories, $\det D$ is usually guaranteed to be positive and real at zero density, but at finite density picks up a complex phase.

A standard approach at this point is to define the *quenched* Boltzmann factor as the absolute value $e^{-S_G} |\det D|$ of the original Boltzmann factor. Expectation values of the physical system can be rewritten as ratios of expectation values of the quenched system:

$$\langle \mathcal{O} \rangle = \frac{\langle \mathcal{O} e^{-i\text{Im} S_{\text{eff}}} \rangle_Q}{\langle e^{-i\text{Im} S_{\text{eff}}} \rangle_Q} \quad (3.5)$$

where $\langle \cdot \rangle_Q$ denotes a quenched expectation value.

The denominator of (3.5) is difficult to evaluate. Each sample will be a number with magnitude 1; once averaged, these number must cancel out to yield a value of considerably smaller magnitude. In fact, a simple argument shows that the denominator $\langle e^{-i\text{Im} S_{\text{eff}}} \rangle_Q$, termed the ‘‘average sign’’ $\langle \sigma \rangle$, is characteristically exponentially small in the volume. The partition function of a field theory in a large volume should be approximately equal to a product of two partition functions with half the volume: the contribution of the boundary to the free energy is negligible in this limit. The same statement is true of the quenched partition function $Z_Q = \int |e^{-S_{\text{eff}}}|$. The average sign is just the ratio of the physical partition function to the quenched partition function, $\langle \sigma \rangle = Z/Z_Q$. It therefore follows that

$$\langle \sigma \rangle (V) = \frac{Z(V)}{Z_Q(V)} \approx \frac{Z(V/2)^2}{Z_Q(V/2)^2} = [\langle \sigma \rangle (V/2)]^2 \quad (3.6)$$

and therefore the average sign shrinks exponentially with the volume. It follows that

an exponentially large number of samples are needed to even resolve the denominator of (3.5) from 0. This exponential scaling is characteristic of methods that encounter a fermion sign problem.

The fermion sign problem is a major obstacle to nonperturbative calculations in several physical regimes. Prominent in nuclear physics is the problem of determining the low-temperature limit nuclear equation of state [27]. To a reasonable approximation, the interior of a neutron star is at zero temperature, and so the energy density as a function of number density (or equivalently, pressure as a function of energy density) of zero-temperature nuclear matter determines the mass-radius curve of neutron stars. This function is beginning to be constrained through astronomical observations [28], but remains largely out of the realm of first-principles calculations.

It has been shown that the most general case of a fermion sign problem is NP-hard [29]. Under standard assumptions of computational complexity, this implies that classical (or even quantum) simulations of such systems cannot be achieved in polynomial time [30]. It is important to bear in mind, however, that this result does not exclude (even heuristically) the polynomial-time simulation of specific instances of systems that suffer from a fermion sign problem. In particular, the system used to prove the NP-hardness of the general case has an inhomogeneous Hamiltonian, and indeed the proof relies heavily on that fact by encoding particular combinatorial problems into the Hamiltonian. The Hamiltonians associated to field theory are homogenous and have little information content.

3.3 Real-Time Linear Response

So far we have discussed the difficulties encountered when trying to determine the equilibrium properties of quantum matter. Out-of-equilibrium physics is also difficult to access with nonperturbative techniques.

A general class of experiments we might perform involve preparing a thermal state of some Hamiltonian H_0 , and then changing the Hamiltonian to some (possibly time-dependent) $H(t)$, and measuring an expectation value $\langle \mathcal{O}(T) \rangle$ at some later time. The Schwinger-Keldysh formalism [31] presents us with the possibility of performing such calculations with lattice methods [32, 33]. The time-dependent expectation value, at inverse temperature β , is given by

$$\langle \mathcal{O}(T) \rangle = \frac{\text{Tr} e^{-\beta H_0} e^{iHT} \mathcal{O} e^{-iHT}}{\text{Tr} e^{-\beta H_0} e^{iHT} e^{-iHT}}. \quad (3.7)$$

Note that the denominator is just the partition function. This expression can be transformed into a path integral in the same manner as the pure-imaginary time partition function. The result, for scalar field theory, is a path integral in both imaginary and real time, with lattice action

$$S = \sum_{x, \langle tt' \rangle} \frac{(\phi_{x,t} - \phi_{x,t'})^2}{2a_0} + \sum_{\langle xy \rangle, t} a_0 \frac{(\phi_{x,t} - \phi_{y,t})^2}{2} + \sum_{x,t} a_0 \left(\frac{m^2}{2} \phi_{x,t}^2 + \lambda \phi_{x,t}^4 \right), \quad (3.8)$$

where $a_0(t)$ depends on the time-slice being considered, being 1 for the first β slices (corresponding to the thermodynamic part of the lattice), i for the next T slices

(yielding forward time evolution), and $-i$ for the rest (backward time evolution).

Here we have assumed that the lattice spacing is 1.

A less ambitious version of this task comes from considering the case where $H(t)$ is equal to H_0 , except for a small, delta-like term added at $t = 0$:

$$H_\epsilon(t) = H_0 + \epsilon\delta(t)H'. \quad (3.9)$$

The response of \mathcal{O} at some later time, to leading order in ϵ , is termed linear response, and is given by a time-separated correlator evaluated in equilibrium.

$$\langle \mathcal{O}(T) \rangle = \langle \mathcal{O}(0) \rangle + \epsilon \langle [H', \mathcal{O}(t)] \rangle \quad (3.10)$$

Transport coefficients, such as bulk and shear viscosity, fall into the category of linear response. The calculation of these time-separated correlators still suffers from a sign problem, and will be our main focus.

An efficient algorithm for classically computing the real-time *non*-linear response of a quantum system, with an arbitrary time-varying Hamiltonian, would imply the ability to efficiently simulate a quantum computer with a classical computer [34]. Thus, under common computational complexity assumptions, it is expected that no such algorithm exists. However, just as Troyer and Wiese's result [29] on the hardness of the inhomogeneous fermion sign problem does not forbid a solution to the homogeneous problem, the hardness of nonlinear simulation with a time-varying Hamiltonian does not seem to forbid the efficient computation of two-

point correlators.

Unlike in the fermion case, however, more directly relevant results have been recently developed. Two developments are worth highlighting here. First, under stronger (but still widely believed) assumptions about computational complexity², the simulation of a sequence of *commuting* quantum gates is inaccessible by any polynomial-time classical algorithm [35]. This problem corresponds to the physical task of computing the nonlinear response of an arbitrary homogeneous state under a time-constant, but spatially inhomogeneous, Hamiltonian. Separately, again under standard assumptions, it was shown in [36] that the task of simulating quantum scattering, beginning from an arbitrary initial state, is also inaccessible by efficient classical algorithms.

These barriers do not provide evidence that time-separated two-point functions are inaccessible classically. The lattice Schwinger-Keldysh method discussed above is not the only approach to computing these functions on the lattice. A common approach, applied for example to the shear viscosity of lattice Yang-Mills [37] is to compute the two-point function at Euclidean separation, and attempt to analytically continue to timelike Minkowski separation. This approach ultimately suffers from the fact that the analytic continuation is ill-posed, and some modeling assumptions are needed.

²Specifically, the fact that the polynomial hierarchy does not collapse.

3.4 Noisy Correlators

Particle masses may be measured in a lattice calculation by considering the long-time behavior of a correlator separated in imaginary time. Let \mathcal{O} be an operator that, when applied to the vacuum state $|\Omega\rangle$, has some overlap with the ground state $|P\rangle$ of a single particle whose mass we would like to know: $\langle P|\mathcal{O}|\Omega\rangle \neq 0$. The Euclidean time-separated correlator has an exponential decay characterized by the mass of the particle:

$$C(\tau) = \langle \Omega | e^{\tau H} \mathcal{O} e^{-\tau H} \mathcal{O} | \Omega \rangle = \sum_i e^{\tau(E_\Omega - E_i)} |\langle i | \mathcal{O} | \Omega \rangle|^2 \quad (3.11)$$

where the sum is taken over all eigenstates of the Hamiltonian, of which $|P\rangle$ is one. When $|P\rangle$ is the lowest-lying eigenstate with nonvanishing overlap, the asymptotic behavior of $C(\tau)$ reveals the mass.

The measurement of $C(\tau)$ on the lattice has some noise, characterized by the variance $\langle (\mathcal{O}(\tau)\mathcal{O}(0))^2 \rangle - \langle \mathcal{O}(\tau)\mathcal{O}(0) \rangle^2$. The difficulty of obtaining an accurate measurement of $C(\tau)$ is measured by the signal-to-noise ratio; that is, the ratio the expectation value to the standard deviation of the estimator. An argument due to Parisi and Lepage [38,39] shows that, for the proton, the signal-to-noise ratio falls off exponentially with τ . The correlator that yields the proton mass is $\langle \bar{q}\bar{q}\bar{q}(\tau)qqq(0) \rangle$, and asymptotically decays with $e^{-\tau m_p}$, where m_p is the proton mass. The variance is given by $\langle \bar{q}\bar{q}\bar{q}(\tau)qqq(\tau)\bar{q}\bar{q}\bar{q}(0)qqq(0) \rangle$. The operator in that correlator at $\tau = 0$ has overlap with a three-pion state, and so the asymptotic behavior of the variance is

$e^{-3\tau m_\pi}$. The noise thus decays less quickly than the signal. As a result, the signal $C(\tau)$ is exponentially difficult to measure at large separations τ .

This signal-to-noise problem is not as severe in practice as the sign problems associated to finite fermion density and real-time correlators. In particular, it has not prevented the accurate measurements of hadronic masses on the lattice [40]. Although it does not outright prevent these calculations, it does make them more expensive. The Parisi-Lepage signal-to-noise problem can be reduced by complexification [41] and, as we will see in Section 5.3.2, can be evaded entirely on a quantum computer.

Chapter 4: Complexification

Motivated by the previous chapters, we would like to compute via Monte Carlo sampling with reweighting, the expectation value of a function $\mathcal{O}(A)$, defined as

$$\langle \mathcal{O} \rangle = \frac{\int \mathcal{D}A e^{-S(A)} \mathcal{O}(A)}{\int \mathcal{D}A e^{-S(A)}}, \quad (4.1)$$

where $S(A)$ denotes the action, $\mathcal{O}(A)$ comes from some Hermitian observable, and the integral is taken over all Euclidean lattice field configurations. For the moment, we will abstract the problem somewhat, allowing S and \mathcal{O} to be arbitrary functions, with the sole restriction that they be holomorphic¹.

To evaluate the expectation value via reweighting more efficiently, we will explore a method for alleviating the sign problem present in the denominator. A sign problem is present whenever S fails to be real, and is characterized by the average sign

$$\langle \sigma \rangle = \frac{\int \mathcal{D}A e^{-S(A)}}{\int \mathcal{D}A e^{-\text{Re} S(A)}} \equiv \frac{Z}{Z_Q}; \quad (4.2)$$

smaller $\langle \sigma \rangle$ correspond to worse sign problems. Note that despite the notation $\langle \sigma \rangle$ is not an expectation value with respect to e^{-S} . In fact, it is an expectation value with

¹We will see in Section 4.2.1 that even for fermionic theories, functions \mathcal{O} coming from arbitrary correlators are in fact holomorphic.

respect to the *phase-quenched* action $\text{Re } S$, and we have introduced the quenched partition function Z_Q , defined as the integral of the quenched Boltzmann factor $|e^{-S}|$.

In this chapter we will consider the integral over fields A as a contour integral in the sense of complex analysis. As written, the integral is taken over $\mathbb{R}^N \subset \mathbb{C}^N$, but we will deform this contour to a different N -manifold $\mathcal{M} \subset \mathbb{C}^N$, and integrate over \mathcal{M} instead. This procedure is motivated by two observations: first, that the action and observables are holomorphic functions of A , and therefore the expectation value $\langle \mathcal{O} \rangle$ will have no dependence on \mathcal{M} ; second, that the quenched partition function is the integral of a non-holomorphic function, and therefore the average sign will generically depend upon the choice of \mathcal{M} .

This chapter proceeds as follows. After a one-dimensional motivating example, the general procedure is rigorously described in the N -dimensional case, with a proof of the theorem that prevents expectation values from depending on the choice of manifold. Next we discuss two methods for selecting a manifold of integration, and apply each to the previously-discussed Thirring model. Finally, we discuss one case in which these methods completely remove the sign problem, and one case in which these methods provably have no effect.

4.1 A One-Dimensional Example

Our motivating example is the sign problem that comes from considering an action of one variable, $S(x) = x^2 + 2i\alpha x$. In this case, the partition function and




Figure 4.1: The sign problem, and its removal, of a gaussian integral. On the left is the Boltzmann factor e^{-S} without any contour deformation performed (black) and with the optimal contour (red). The original manifold of integration (black) and the deformed contour (red) are on the right.

quenched partition function can both be evaluated exactly, and the average sign is

$$\langle \sigma \rangle (\alpha) = \frac{\int e^{-x^2 - 2i\alpha x}}{\int e^{-x^2}} = e^{-\alpha^2}. \quad (4.3)$$

As an aside, note that one can increase the ‘volume’ of the system by adding more dimensions to the integral. The partition function of the volume- V system is then $Z_V = Z_1^V$, where Z_1 is the partition function given by the single-dimensional integral. The resulting sign problem is

$$\langle \sigma \rangle (\alpha, V) = \langle \sigma \rangle (\alpha, 1)^V = e^{-\alpha^2 V}, \quad (4.4)$$

a simple demonstration of the general fact that sign problems scale exponentially in the volume. (By contrast, the parameter α doesn’t correspond to any parameter in a physical system, and so the scaling with α shouldn’t be taken seriously.)

This sign problem can be removed entirely by deforming the contour of inte-

gration from the real plane to the line defined by $\text{Im } x = -\alpha$. The situation is shown in Figure 4.1: after the contour deformation, the partition function is written

$$Z = \int_{\infty}^{\infty} dx e^{-(x+i\alpha)^2} = \int_{\infty-i\alpha}^{\infty-i\alpha} dx e^{-x^2} = \int_{\infty}^{\infty} dx e^{-x^2} \quad (4.5)$$

which is of course a sign-problem free integral. The first step is merely a change-of-variable; in the second step, Cauchy's theorem must be invoked to show that the two contour integrals are the same.

This motivates the broad strategy of 'complexification' for attacking sign problems. In general, the partition function Z (from which physical quantities are obtained) is invariant under contour deformations, as it is an integral of the holomorphic function $e^{-S(A)}$. The quenched partition function Z_Q , by contrast, is the integral of the non-holomorphic $e^{-\text{Re } S(A)}$, so the value of Z_Q , and therefore the average sign, are not invariant under contour deformations. Thus one may improve the sign problem by deforming the contour of integration, without changing the observables measured. The precise requirements on the contour deformation are imposed by Cauchy's theorem, which we discuss next.

4.2 Cauchy's Integral Theorem

Cauchy's integral theorem is the tool that allows us to deform a contour integral without changing the value of the integral. A function $f : \mathbb{C} \rightarrow \mathbb{C}$ is termed

holomorphic where it obeys the Cauchy-Riemann equations:

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y} \quad \text{and} \quad \frac{\partial v}{\partial x} = -\frac{\partial u}{\partial y}, \quad (4.6)$$

where we have decomposed $f(x + iy) = u(x, y) + iv(x, y)$ into its real and imaginary parts. By introducing the holomorphic (Wirtinger [42]) and antiholomorphic derivatives

$$\partial = \frac{\partial}{\partial z} \equiv \frac{1}{2} \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right) \quad \text{and} \quad \bar{\partial} = \frac{\partial}{\partial \bar{z}} \equiv \frac{1}{2} \left(\frac{\partial}{\partial x} + i \frac{\partial}{\partial y} \right), \quad (4.7)$$

the Cauchy-Riemann equations can be rewritten as $\bar{\partial}f = 0$. Note that the holomorphic and antiholomorphic derivatives, just like the ordinary derivatives ∂_x and ∂_y , are derivatives taken along orthogonal vectors.

Cauchy's integral theorem states that for holomorphic functions, integrals around closed contours vanish.

Theorem 1 (Cauchy's Integral Theorem). *For a closed region $\Omega \subset \mathbb{C}$ and a function f holomorphic on Ω , the integral of f around the boundary of Ω must vanish:*

$$\oint_{\partial\Omega} f(z) \, dz = 0 \quad (4.8)$$

Proof. By Stokes' theorem,

$$\oint_{\partial\Omega} f(z) \, dz = \int_{\Omega} d(f(z)dz) = \int_{\Omega} df \wedge dz. \quad (4.9)$$

The differential of f is given by $df = \partial f dz + \bar{\partial} f d\bar{z}$. As the anti-holomorphic derivative of f vanishes, the differential of f is simply $df = \partial f dz$. However, the wedge product $dz \wedge dz$ vanishes, and so must the integral. \square

The usual form of Cauchy's theorem — the one just given — applies to functions of one complex variable. The theorem has a natural generalization to functions of many complex variables. A $f : \mathbb{C}^N \rightarrow \mathbb{C}$ of N complex variables is termed holomorphic where it obeys the Cauchy-Riemann equations in each complex dimension independently; that is, where

$$\frac{\partial}{\partial \bar{z}_i} f(z_1, \dots, z_i, \dots, z_N) = 0 \quad (4.10)$$

for all i . A multidimensional generalization of Cauchy's theorem states that the integral around the boundary of any $(N + 1)$ -real-dimensional region, in which f is holomorphic, must vanish.

Theorem 2 (Multidimensional Cauchy's Theorem). *For a closed region $\Omega \subset \mathbb{C}^N$ of real dimension $N + 1$, and a function f holomorphic on Ω , the integral of f around the boundary of Ω vanishes:*

$$\oint_{\partial\Omega} f(z) \bigwedge_{i=1}^N dz_i = 0 \quad (4.11)$$

Proof. As before, Stokes' theorem yields

$$\oint_{\partial\Omega} f(z) \wedge \bigwedge_{i=1}^N dz_i = \int_{\Omega} df \wedge \bigwedge_{i=1}^N dz_i. \quad (4.12)$$




Figure 4.2: The asymptotically safe regions for the Gaussian integral.

The differential df now has $2N$ terms, of which N vanish by the Cauchy-Riemann equations $\bar{\partial}_i f = 0$. Each of the remaining terms has the form $\partial_j f dz_j$ for some j , and therefore is annihilated when the wedge product is taken with $\bigwedge_i dz_i$. \square

A complication arises when deforming integration contours which extend to infinity. Application of Cauchy's integral theorem alone does not allow the asymptotic behavior of such a contour to be changed. However, in cases where the function being integrated decays rapidly² at infinity, the contribution of the part of the contour at infinity vanishes, and the asymptotic behavior can be changed without affecting the result. Figure 4.2 shows the asymptotically safe regions for the Gaussian model. The shaded regions mark the regions that decay exponentially at infinity, and as long as a contour's asymptotics remain in a shaded region, the integral will be unchanged by any deformation.

Tracking the asymptotically safe regions and ensuring the manifold deformation never leaves them can be technically challenging. It is often possible to arrange the physical model so that the domain of integration is compact³. This ensures

²An exponential decay is always sufficient.

³The standard formulation of lattice gauge theories accomplishes this.

that the original integration manifold does not go near any asymptotic region, and any finite deformation will be permissible. This is the approach we will take when applying the method to fermionic models.

4.2.1 Holomorphic Boltzmann Factors and Observables

The utility of a Cauchy's-theorem-based procedure comes from the observation that the integrands of interest are holomorphic functions of the integration variables. Here we discuss two major cases in which the holomorphicity of the integrands is not obvious: a theory of complex scalar fields, and a theory of fermions in which the fermions have been integrated out.

4.2.1.1 Complex Scalar Fields

The (Euclidean) lagrangian of a complex scalar field is given by

$$L = (\partial_\mu \phi^\dagger) (\partial_\mu \phi) + V(|\phi|), \quad (4.13)$$

and it is immediately obvious that the lagrangian is not, and therefore the lattice action is not, a holomorphic function of the field variable ϕ . The resolution to this dilemma comes upon writing out the full lattice path integral:

$$Z = \int_{\mathbb{C}^V} d\phi \exp \left\{ - \sum_{\langle ij \rangle} \frac{|\phi_i - \phi_j|^2}{2} - \sum_i V(\phi_i) \right\} \quad (4.14)$$

Here we see that the domain of integration (for a theory with V sites) has V *complex* dimensions, rather than V real dimensions. Expressing the complex scalar field at site i as a sum $\phi_i = u_i + iv_i$ of two real scalar fields at the same site, we note that both ϕ and ϕ^\dagger are in fact analytic functions of the new field variables u and v . The partition function may now be written

$$Z = \int d^V u d^V v \exp \left\{ - \sum_{\langle ij \rangle} \frac{(u_i - u_j)^2}{2} - \sum_{\langle ij \rangle} \frac{(v_i - v_j)^2}{2} - \sum_i V(u, v) \right\}, \quad (4.15)$$

in which the integrand is manifestly a holomorphic function of u and v . The methods of the previous sections may now be applied, and the contour integral deformed into the space of imaginary u and v .

4.2.1.2 Fermionic Determinant and Correlators

Thus far we have been concerned with ensuring the lattice action is a holomorphic function. This guarantees that the integrand of the partition function is holomorphic, and that the partition function is unchanged by the deformation. Holomorphicity of the action, however, is in general neither necessary nor sufficient. In fact it is required that $e^{-S[\phi]}$ and $\mathcal{O}[\phi]e^{-S[\phi]}$ both be holomorphic. We will now see that this does not imply that either S or \mathcal{O} must themselves be holomorphic [43].

Consider a theory of fermions ψ , with interactions mediated by a gauge field A_μ . After the fermions have been integrated out, the lattice partition function is

$$Z = \int d^d A e^{-S_{\text{gauge}}[A]} \det D[A], \quad (4.16)$$

where the lattice has d dimensions (making A_μ a d -vector), S_{gauge} gives the terms in the lattice action involving only the gauge field, and $D[A]$ is the fermion inverse propagator in the presence of a fixed background field A .

A typical observable of interest is a meson propagator

$$\langle \bar{\psi}_i \bar{\psi}_j \psi_j \psi_i \rangle = \frac{1}{Z} \int \mathcal{D}\phi e^{-S[\phi]} [D_{ij}^{-1} D_{ji}^{-1} - D_{ii}^{-1} D_{jj}^{-1}], \quad (4.17)$$

so we require this integrand to be holomorphic as well.

Note that the effective action on the gauge fields — i.e. the logarithm of the integrand in the partition function — contains a term $\log \det D$. This term has logarithmic singularities where $\det D = 0$. Also at these points, D^{-1} is not well-defined, so the meson propagator given above (along with many others) involves a singular $\mathcal{O}[A]$.

Despite this, the integrands $e^{-S[\phi]}$ and $\mathcal{O}[\phi]e^{-S[\phi]}$ are always holomorphic (with lattice regularized actions). The holomorphicity of e^{-S} is easiest to understand. The gauge action is of course holomorphic in the fields A , and for typical gauge-fermion interactions each element $D[A]_{ij}$ of the fermion matrix is also a holomorphic function of A . As the determinant is merely a polynomial of the elements of the matrix, it follows that $\det D$ itself is holomorphic. This establishes that the partition function (4.16) will remain unchanged under appropriate deformations.

We now discuss fermionic observables. In the case of a fermion propagator $\langle \bar{\psi}_i \psi_j \rangle$, there is only a single D^{-1} in the integrand, and it is easy to see that the singularity of this factor is cancelled by the zero of the fermion determinant coming

from the action. To see that integrands involving fermionic observables are holomorphic in general, we write an expectation value in terms of the original, fermionic path integral.

$$\langle \bar{\psi}_a \psi_b \rangle = \frac{1}{Z} \int \mathcal{D}A e^{-S_{\text{gauge}}[A]} \int \mathcal{D}\bar{\psi} \mathcal{D}\psi \bar{\psi}_a \psi_b e^{-\bar{\psi}_i D_{ij}[A] \psi_j}$$

With V sites, the fermionic exponential $e^{-\bar{\psi} D \psi}$ may be expanded in 2^{2V} terms, identified by what subset of the $2V$ Grassmann variables is included in each term. The \mathbb{C} -number part of each term is a product of finitely many components of $D[A]$, and therefore is a holomorphic function of the gauge field A . Multiplying by any combination of $\bar{\psi} \psi$ and integrating over $\mathcal{D}\bar{\psi} \mathcal{D}\psi$ has the effect of selecting one of these coefficients. Therefore, the integral over fermionic fields yields a holomorphic function of A .

The story remains the same no matter how many fermionic fields are inserted in the expectation value, as long as no Grassmanns are repeated. If a Grassmann is repeated (if the expectation value $\langle \bar{\psi}_i \psi_i \rangle$ is requested), then the expectation value is simply zero.

4.3 Lefschetz Thimbles

At this point we have motivated the use of deformed contour integrals, and shown that physical observables will remain unchanged, but we have no general principles for selecting integration manifolds on which the average sign is likely to be improved. The Lefschetz thimbles [44] provide an attractive choice of manifold

for improving the sign problem [45, 46]. Each thimble is an N -dimensional manifold extending from a critical point z_c of the action obeying $\partial S|_{z=z_c} = 0$. The thimble extends from the critical point along the paths of steepest descent of the real part of the action $\text{Re } S$. The thimble terminates either at infinite, or at a point where the real part of the action diverges (zeros of the fermion determinant have this effect).

The union of all thimbles is not necessarily obtainable as a smooth deformation of the real plane; however, some linear combination of the thimbles always is [44]. In other words, there is some linear combination of the thimbles that, when integrated, gives a result equal to the integral along the real plane (for any holomorphic integrand). Determining what linear combination is needed may be computationally difficult, as indeed may be the task of enumerating all critical points of the action. Section 4.4 below provides a closely related algorithm which circumvents this problem.

The usefulness of the thimbles is related to the fact that $\text{Im } S$ is constant on each thimble. Note first that the path of steepest descent, defined by

$$\frac{d\text{Re } A}{dt} = \frac{\partial S}{\partial \text{Re } A} \quad \text{and} \quad \frac{d\text{Im } A}{dt} = \frac{\partial S}{\partial \text{Im } A}, \quad (4.18)$$

can also be written in the form $\dot{A} = \bar{\partial} S$. It follows that the change of S with flow time is given by

$$\frac{dS}{dt} = \frac{dA}{dt} \frac{\partial S}{\partial A} = \left| \frac{\partial S}{\partial A} \right|^2. \quad (4.19)$$

As the change in the action is real, the imaginary part is constant along each path of

steepest descent, and therefore all over the thimble. The fact that $\text{Im } S$ is constant along a thimble means that, within one thimble and neglecting the Jacobian, there can be no phase cancellations in the integral.

It is not the case, however, that Lefschetz thimbles completely remove the sign problem. Lattice theories often have multiple thimbles contributing to the integral [47], with different phases, creating a sign problem. Additionally, although $\text{Im } S$ is constant, the Jacobian (i.e. the curvature of the thimble) introduces its own contribution to the phase [48].

4.4 Holomorphic Gradient Flow

A thimble may be defined as the union of all solutions $A(t)$ to the holomorphic gradient flow equation

$$\frac{dA}{dt} = \frac{\partial S}{\partial \bar{A}} \tag{4.20}$$

that approach a certain critical point A_c in the early-time limit: $\lim_{t \rightarrow -\infty} A(t) = A_c$. This definition makes apparent the fact that any thimble is left invariant under the action of the gradient flow. In fact, thimbles are not only fixed points of the flow, but attractive fixed points: any nearby manifold will evolve to become closer to the thimble, under the flow.

It follows that the holomorphic gradient flow can be used to construct an algorithm for integrating along the thimbles [49]. Define a function $\tilde{A}_T(A)$ to be the result of evolving the point A under (4.20) for time T . Under mild conditions on the action (holomorphicity of e^{-S} is sufficient), this is a continuous function

of A and T , and therefore defines a contour of integration homotopic to the real plane. Moreover, in the limit of large T , this integration contour approaches some linear combination of the Lefschetz thimbles, and is therefore expected to have an improved sign problem. The Monte Carlo integration is performed by sampling A according to a modified action

$$S_{\text{eff}}(A) = S[\tilde{A}_T(A)] - \log \det J, \quad (4.21)$$

where J is the Jacobian of \tilde{A} , i.e. the matrix of complex first derivatives.

Although the thimbles are only obtained in the long-time limit of the gradient flow, any manifold created by flowing for a finite amount of time can be used as an integration contour. In practice, it is found that flowing for a short amount of time can dramatically improve the average sign at a relatively small computational cost [50]. This choice of integration manifold, defined by flowing the real plane for some fixed time T , is both useful in practice and provides a guide for the search for other manifolds.

4.4.1 Algorithmic Costs

Methods based directly on the holomorphic flow have a substantial drawback: the computation of $\det J$ is expensive. The evolution of the Jacobian induced by (4.20) is given by

$$\frac{dJ}{dt} = \bar{H} \bar{J}, \quad (4.22)$$

where H , the Hessian, is the matrix of holomorphic second derivatives of the action. The matrix multiplication is unavoidable when computing $\det J$, and requires about $O(n^3)$ steps in practice for an $n \times n$ matrix. The asymptotic time complexity of one step of the flow, then, is approximately cubic in the volume of the lattice. Much of the technical effort around flow is motivated by the desire to avoid this cost, including by computing an approximation to the determinant and reweighting [51], or modifying the Monte Carlo sampling to automatically include the Jacobian [33].

Although flow-based methods can improve a wide variety of sign problems without much need for model-specific tweaks, the expense of the procedure restricts the method to small lattices in practice. The parameterization of the integration manifold by the real plane is also not particularly convenient: in the limit of long flow times, an entire thimble is mapped to by a single point, creating large potential barriers (in parameterization space) between different thimbles. Finally, we will see in Section 4.8.1 that the Lefschetz thimbles, the “ultimate goal” of flow-based methods, are not the best possible manifold, and that for large lattices it may be possible to dramatically improve the sign problem with a different integration contour. This motivates us to continue the search for other manifolds.

4.5 Machine Learning Manifolds

Instead of directly integrating on the flowed manifold $\tilde{A}_T(A)$, we can use machine learning methods to create a computationally efficient approximation, and perform the integration on the approximated manifold instead [52].




Figure 4.3: Figure from [52] depicting a feed-forward network, used to define a manifold in the complexified field space.

4.5.1 Feed-Forward Networks

Feed-forward networks are a particular class of nonlinear functions $f(\vec{x})$ of many variables which are particularly fast to compute. A feed forward network, depicted in Figure 4.3, consists of several layers, each containing a set of nodes. The initial layer (shown on the left) is the input layer, and to each node is associated one input variable x_i . Values are propagated through the network from left to right. The values y_j in the second layer are determined from the values of the first layer by first performing some linear transformation on the x_i (usually thought of as a set of weights associated to the edges between the first and second layers), and then performing a nonlinear transformation on each node in the second layer separately. Letting w be the matrix of weights, b_j be a linear bias associated to node j of the second layer, and $\sigma(\cdot)$ be the nonlinear transformation, we have $y_j = \sigma(b_j + w_{ji}x_i)$. This process is then repeated for every subsequent layer. The output

$f(\vec{x})$ is extracted from the output node, or output nodes in the case of a vector-valued function.

A feed-forward network $f(\vec{x})$ can be made to define a manifold by taking it to yield the imaginary part of a point on the manifold when given the real part as input:

$$\tilde{A}(A) = A + if(A). \tag{4.23}$$

Here, because A has many components, there may in general be many functions $f(A)$, which must be trained separately. The number of separate functions required can be reduced by imposing translational invariance and other symmetries found in the target action. This construction of a manifold is not completely general, because it requires that there be only one point on the manifold with any given set of real coordinates. Nevertheless, this suffices to describe any manifold obtained by a sufficiently small amount of flow, and no example of a flowed manifold that folds back on itself (that is, for which the imaginary coordinate isn't a single-valued function of the real coordinate) has been found in practice.

4.5.2 Training

In order to find a suitable function $f(\vec{x})$ for our manifold, a cost function must be defined, which we will seek to minimize. We may use the holomorphic gradient flow as a guide. After flowing some set of points from the real manifold, we obtain

a training set of \mathcal{N} points (x, y) located on the flowed manifold. The cost function

$$C(w, b) = \frac{1}{\mathcal{N}} \sum_{h=1}^{\mathcal{N}} \left| \vec{f}_{w,b}(\vec{x}^{(h)}) - \vec{y}^{(h)} \right| \quad (4.24)$$

attempts to compute the imaginary part of each of these points by looking at the real part, and takes the average error.

It remains only to perform the minimization of the cost function; this minimization is usually done with some form of gradient descent algorithm. The space of biases and weights is of quite large dimension, and the cost function has many local minima, so some experimentation with different algorithms is advisable. An extensive review of gradient descent algorithms is given in [53]; the Adaptive Moment Estimate algorithm [54] (dubbed ADAM) was used in [52] for the purposes of training the manifold.

4.6 Manifold Optimization

The manifold learning method used a cost function defined by taking the distance between an ansatz manifold (e.g. a feed-forward network) and a set of training data constructed using the holomorphic flow. We can, however, use a more directly relevant cost function, constructed from the sign problem itself [55, 56]. An immediate obstacle is that, if we chose the cost function to be the average sign evaluated on the manifold \mathcal{M} , we find that the task of evaluating the cost function on a given manifold is as difficult as measuring any other observable on that manifold. The average sign is a noisy observable, and it is difficult to measure precisely when

there is a bad sign problem.

It turns out, though, that an inability to evaluate the cost function is no obstacle to its optimization. We select as our cost function the log of the average sign:

$$C(\mathcal{M}) = -\log\langle\sigma\rangle_{\mathcal{M}}. \tag{4.25}$$

Here we have written the cost function directly as a function of the manifold \mathcal{M} , and the average sign on that manifold is denoted $\langle\sigma\rangle_{\mathcal{M}}$. Where a family of manifolds parameterized by some λ is used as an ansatz, this induces a cost function of the space of λ .

This cost function is no easier to evaluate. However, the derivative with respect to some manifold parameter λ is quite simple, as a result of the fact that the physical partition function Z cannot depend on the choice of manifold.

$$\frac{\partial}{\partial\lambda}C(\mathcal{M}_\lambda) = -\frac{\partial}{\partial\lambda}\log Z_Q(\mathcal{M}_\lambda) \tag{4.26}$$

We see that $\partial_\lambda C$ is a derivative of the log of the *quenched* partition function. This is an expectation value of the quenched system, which can be evaluated without encountering a sign problem. At this point, we may again apply any minimization method to optimize the cost function, just as was done for the manifold learning procedure above.

This method has two substantial advantages over manifold learning. First, it does not require the potentially expensive step of preparing a library of flowed

points to use as training data. Second, while the manifold learning procedure can at best be expected to perform (as measured by $\langle\sigma\rangle$) as well as the holomorphic flow, manifold optimization makes no reference to the flow and can in principle find manifolds with milder sign problems than any reached by flowing. We will see later that this is in fact the case for the Thirring model, even with a relatively simple ansatz.

4.6.1 Another View of the Flow

The flow was originally motivated by the observation that, in the limit of long flow times, the manifold would approach the Lefschetz thimbles, which generally have a substantially milder sign problem than the real plane. However, the flow has been found in practice to greatly improve the sign problem even for quite small flow times, without coming particularly close to the thimbles. This should be surprising: why does the flow perform so well, away from the regime where it is a well-motivated procedure?

The picture of manifold optimization above provides us with an answer. Take as an ansatz the family of manifolds $\tilde{z}(x)$ defined by interpolating from a fine mesh. The real plane itself is in this ansatz: the value of y associated each x in the mesh is $y_i(x) = 0$. Taking this as our starting point, we perform gradient descent on the cost function $C[y] = -\log\langle\sigma\rangle$. The gradient is

$$-\frac{\partial}{\partial y(x)} \log Z_Q = \frac{1}{Z_Q} \int e^{-\text{Re } S} \frac{\partial \text{Re } S}{\partial y(x)}. \quad (4.27)$$

We see that, starting from the real plane, the holomorphic gradient flow is in fact moving (in manifold space) in the direction which most quickly improves the average sign. Unfortunately, after the first infinitesimal step of flow has been performed, there is no longer a simple expression for the behavior of the manifold optimizing flow.

4.7 Application to the Thirring Model

The Thirring model [57] is a common target for methods designed to alleviate or remove a fermionic sign problem. In $1 + 1$ dimensions, it is defined in the continuum by the Euclidean action

$$S = \int d^2x \left[\bar{\psi}^\alpha (\not{\partial} + \mu\gamma_0 + m)\psi^\alpha + \frac{g^2}{2N_F} \bar{\psi}^\alpha \gamma_\mu \psi^\alpha \bar{\psi}^\beta \gamma_\mu \psi^\beta \right], \quad (4.28)$$

where the flavor indices take values $\alpha, \beta = 1, \dots, N_F$, μ is the chemical potential, and ψ is a two-component spinor.

The four-fermi interaction is removed by introducing an auxilliary field A_μ , which we take to be periodic with period 2π . The resulting lattice action is

$$S = \frac{N_F}{g^2} \sum_{x,\nu} (1 - \cos A_\nu(x)) + \sum_{x,y} \bar{\psi}^a(x) D_{xy}(A) \psi^a(y) \quad (4.29)$$

where the spin index a is implicitly summed over $a = 1, 2$. For Kogut-Susskind

staggered fermions [18], the matrix D is defined by

$$D_{xy} = m\delta_{xy} + \frac{1}{2} \sum_{\nu=0}^2 \left[\eta_{\nu}(x) e^{iA_{\nu}(x) + \mu\delta_{\nu 0}} \delta_{x+\hat{\nu},y} - \eta_{\nu}^{\dagger}(y) e^{-iA_{\nu}(y) - \mu\delta_{\nu 0}} \delta_{x,y+\hat{\nu}} \right]. \quad (4.30)$$

This is of course not the only discretization possible. Another, with Wilson fermions [58], yields the fermion matrix

$$D_{xy}^W = \delta_{xy} - \kappa \sum_{\nu=0,1} \left[(1 - \gamma_{\nu}) e^{iA_{\nu}(x) + \mu\delta_{\nu 0}} \delta_{x+\nu,y} + (1 + \gamma_{\nu}) e^{-iA_{\nu}(x) - \mu\delta_{\nu 0}} \delta_{x,y+\nu} \right]. \quad (4.31)$$

Except where otherwise noted, statements in this chapter are applicable to both discretizations. As usual, because the lattice action is quadratic in the fermion fields, they can be integrated out of (4.29), yielding

$$S = \frac{N_F}{g^2} \sum_{x,\nu} (1 - \cos A_{\nu}(x)) - \frac{N_F}{2} \log \det D(A). \quad (4.32)$$

We will work in the case $N_F = 2$.

The Thirring model has no sign problem at vanishing chemical potential. At finite chemical potential there is, as usual, a sign problem exponentially bad in the volume. The sign problem is made worse at larger couplings and larger chemical potentials. The sign problem of the Thirring model has been extensively investigated with flow-based methods, in both 0 + 1 dimensions [49, 50, 59] and 1 + 1 dimensions [60]. Attempts have also been made to approximate the Thirring model by integrating on a single thimble in isolation [61–63].

4.7.1 Field Complexification

The path integral for the lattice Thirring model defined in this way is an integral over the manifold S^{dV} , that is, one copy of the unit circle for each variable A , of which there are dV , where V is the volume and d the dimension of the lattice. In order to apply the methods of complexification, we need to construct a space with complex structure which includes S^{dV} .

The complexification of S^1 is a cylinder: the set of points (x, y) such that $x \in [0, 2\pi)$ and y is an unbounded real number. Under the exponential map, the original domain of A_ν maps to the unit circle in the complex plane. The full complexified space maps to the complex plane with one point removed, $\mathbb{C} \setminus \{0\}$. As the integration space is just the product of many copies of S^1 , the complexification is the product of many cylinders. Topologically the space is $(S^1 \times \mathbb{R})^{dV}$.

The real plane, which we will refer to as \mathcal{M}_0 , may be deformed to another manifold \mathcal{M}_1 without changing the value of the path integral as long as the two manifolds together form the boundary of a closed region in $(S^1 \times \mathbb{R})^{dV}$. A sufficient condition is that there exists a homotopy between \mathcal{M}_0 and \mathcal{M}_1 , that is, a continuous family of functions $g_t : \mathcal{M}_0 \rightarrow (S^1 \times \mathbb{R})^{dV}$ from the real plane to the complex plane, such that g is continuous in both t and its argument, g_0 is the identity function, and the range of g_1 is \mathcal{M}_1 . This sort of construction is also convenient from a computational point of view, as the function g_1 already provides a parameterization of the manifold to be integrated on.

For the purposes of the Thirring model, we can be even less general. Any

manifold of the form

$$\tilde{A}_\nu(x) = A_\nu(x) + i f_\nu^{(x)}(\vec{A}) \quad (4.33)$$

such that $f_\nu^{(x)}$ is a continuous function in its dV arguments, is homotopically connected to the real plane. The homotopy is constructed by scaling the function $f_\nu^{(x)}$ by t .

4.7.2 An Ansatz

This is the manifold ansatz we will consider [55]:

$$\tilde{A}_0(x) = A_0(x) + i f(A_0(x)) \quad \text{and} \quad \tilde{A}_i(x) = A_i(x). \quad (4.34)$$

This is an enormously constrained ansatz: we have left all dimensions other than $\nu = 0$ undeformed, and the deformation of the integral over $A_0(x)$ does not depend on the value of A at any other link. Additionally, we have used the fact that the action is translationally invariant to infer that the ansatz should be as well. Nevertheless, the ansatz still has an infinite number of parameters. The requirement that f be a continuous function suggests that we expand it in a fourier series:

$$f(z) = a_0 + a_1 \cos(x) + a_2 \cos(2x) + \dots + b_1 \sin(x) + \dots \quad (4.35)$$

The action is symmetric under $A_0 \rightarrow -A_0$, which suggests that the chosen manifold ought to be as well, so we can set $b_i = 0$. Finally, to have a finite number of manifold parameters, we truncate the fourier series to the first 3 even terms, with coefficients

a_0 , a_1 , and a_2 .

4.7.3 Phase Diagram

Now that a plausible ansatz is constructed, it can be optimized with the methods of Section 4.6. The manifold parameters depend on the model parameters, so the optimization must be performed separately for every set of model parameters.

We simulate the 2 + 1-dimensional Thirring model [64] with bare lattice parameters $g = 1.08$ and $m = 0.01$. All results are quoted in lattice units; physical quantities may be recovered by multiplication with the appropriate power of the lattice spacing. This choice of g and m puts the lattice model in the strong coupling regime: in a box of size 10^2 , we measure a fermion mass $m_f = 0.46(1)$ and a boson mass $m_b = 0.21(1)$.

We focus on the chiral condensate, defined by the expectation value $\langle \bar{\psi}\psi \rangle$. At low temperatures and low chemical potentials, the chiral condensate has a non-zero value, indicating the breaking of chiral symmetry. At either high temperature or larger chemical potential, chiral symmetry is nearly (because $m > 0$) restored, and the chiral condensate drops to near zero.

Figure 4.4 shows measurements of the chiral condensate on a 6^2 spatial lattice; the size of the time dimension is temperature-dependent. At low temperatures, a relatively sharp transition between the broken and unbroken phases is seen near $\mu \sim 0.4$. The crossover broadens at higher temperatures, and moves to lower chemical potentials.




Figure 4.4: Figure from [64] of the chiral symmetry breaking phase transition of the Thirring model. On the left, the condensate $\langle \bar{\psi}\psi \rangle$, as a function of chemical potential μ on $\beta \times 6^2$ lattices. On the right, the full T - μ plane for the same spatial volume. The central band indicates the location of $\langle \bar{\psi}\psi \rangle_{\mu,T} = 0.5 \langle \bar{\psi}\psi \rangle_0$; the thinner lines on either side indicate the width of the crossover

4.8 Optimal Manifolds

The general method of complexification may fail for two different reasons. For any given model, the complexification method may fail because no manifold that removes the sign problem exists, or it may fail because the manifold is just too computationally expensive to integrate on (for instance, it may be difficult to find in the first place!). In practice it is difficult to distinguish these two failure modes unless we can prove no satisfactory manifold exists. In this section we consider general questions about the “best possible” manifolds, that is, those that minimize the quenched partition function, and look in particular at one case where a “perfect” manifold can be found, and at another where we can prove none exists.




Figure 4.5: Thimbles and an optimized manifold from the ansatz 4.35 for the heavy-dense limit of the Thirring model, or equivalently the one-site Thirring model.

4.8.1 Lefschetz Thimbles Are Not Optimal

The lattice Thirring model (4.29) becomes trivial in the heavy-dense limit of large chemical potential μ . Physically, the large chemical potential pushes the Fermi momentum up to the lattice cutoff, so that every site of the lattice is filled by a fermion. To see this effect algebraically, we can expand the fermion determinant in powers of $e^{-\mu}$. The leading-order term is the one in which only time-like links are included. The physical ‘saturation’ effect of the lattice manifests in the partition function factorizing to leading order $e^{-\mu}$, so that each link is now independent and uncoupled from all other links:

$$Z = \left[\int dA_0 e^{\frac{N_F}{g^2} \cos A_0 + \mu + iA_0} \right]^{\beta L} \left[\int dA_1 e^{\frac{N_F}{g^2} \cos A_1} \right]^{\beta L}. \quad (4.36)$$

With exponents appropriately modified, this factorization holds independent of the number of spacetime dimensions. This observation provides a convenient *post-hoc*

rationalization for the ansatz of Section 4.7: that ansatz defines the most general manifold which maintains all the symmetries of the heavy-dense limit of the Thirring model⁴.

This trivial limit also allows us to study how the Lefschetz thimbles compare to a “best-possible” manifold [48]. Because the partition function factorizes, the average sign does as well:

$$\langle \sigma \rangle = \left(\frac{Z_1}{Z_{1,Q}} \right)^{\beta L} \quad (4.37)$$

where Z_1 and $Z_{1,Q}$ denote the partition function and quenched partition function, respectively, of the one-site model. This allows us to accurately compute the average sign at large volumes, where a direct calculation of $\langle \sigma \rangle$ would be impractical. The thimbles and the ansatz of that model are shown in Figure 4.5.

Numerical evidence indicates that the ansatz (with all Fourier coefficients maintained) achieves an average phase of exactly 1, regardless of coupling. The average phase on the thimbles, meanwhile, is less than 1 for any nonvanishing coupling. In the case of $g = 1.08$, the average phase obtained on the thimbles is 0.985 for the one-link model, and therefore $(0.985)^{\beta L}$ for the full lattice.

This heavy-dense limit of the Thirring model is one of the few physically-inspired models in which both the Lefschetz thimbles and a provably optimal manifold can be understood. In this case, a manifold that completely solves the sign problem does exist, and the thimbles do not.

The fact that the sign problem of this model can be solved exactly may be

⁴Strictly speaking, we have also imposed that $\text{Im } A_0$ be a single-valued function of $\text{Re } A_0$.

very special — in fact, in the next section, we will see a (less physical) model in which no manifold can solve the sign problem. However, the fact that the Lefschetz thimbles are non-optimal is probably less special. Looking at Figure 4.5, note that the thimbles contain a sharp ‘cusp’, where the contour folds back along itself. At stronger couplings, the cusp becomes sharper, and the thimbles come closer to each other. The action on one side of the cusp does not differ much from the action on the other, but the sign of the integration element dz flips. Therefore, the contributions to the integral from the two sides of the cusp almost exactly cancel. The ansatz manifold cuts off the cusp, removing the considerable residual sign problem in this region.

4.8.2 A Complexification-Immune Sign Problem

Not every sign problem can be removed with complexification. A simple example serves to prove the point:

$$Z(\epsilon) = \int_0^{2\pi} d\theta (\epsilon + \cos \theta). \quad (4.38)$$

This partition function should be thought of as a ‘lattice’ with a single degree of freedom θ on a single site, and an action of $S_\epsilon(\theta) = -\log(\epsilon + \cos \theta)$. We are interested in the regime of small ϵ , where we will be able to establish upper bounds on the best possible average phase $\langle \sigma \rangle$.

We begin by taking $\epsilon = 0$, where the partition function vanishes. Here there are two thimbles, constituting the two halves of the real line, and the antisymmetry

of the Boltzmann factor under $\theta \rightarrow \theta + \pi$ causes them to exactly cancel. The partition function will vanish no matter what manifold is chosen. The quenched partition function depends strongly on the manifold, but can be rigorously bounded from below. For any $x \in [0, 2\pi)$, the chosen manifold must have a point θ with $\text{Re } \theta = x$. Because $\cos \theta$ is minimized, for any fixed $\text{Re } \theta$, by $\text{Im } \theta = 0$, the minimum possible value of $|e^{-S(x)}|$ is achieved at $\theta = x$. Therefore we can do no better in minimizing Z_Q than to integrate over the real line, resulting in the bound $Z_Q \geq 4$. Because the partition function itself vanishes, the average sign will always be zero. This is a pathological example.

The pathology is lifted by introducing ϵ . At small $\epsilon > 0$, the thimbles remain on the real line, but the cancellation is no longer exact, and the partition function no longer vanishes, but is instead given by $Z = 2\pi\epsilon + O(\epsilon^2)$. The average sign, then, is forced to be of order ϵ as well:

$$\langle \sigma \rangle \leq \epsilon \frac{\pi}{2} + O(\epsilon^2). \quad (4.39)$$

We conclude that, at small but nonvanishing ϵ , there is no manifold that can improve the sign problem beyond what is achieved on the real line. Moreover, even with only one degree of freedom, the sign problem on the real line can be made arbitrarily bad.

A key feature of this example is the presence of multiple cancelling thimbles. As mentioned earlier, the other way in which the Lefschetz thimbles fail to completely remove the sign problem is via the residual phase introduced by the Jacobian.

Whether this residual phase can always be counteracted by deforming away from the thimbles (at in the heavy-dense limit of the Thirring model) remains an open question.

Chapter 5: Quantum Simulations

In this chapter we discuss the use of a quantum computer in studying the time-evolution of physical quantum systems. Sections 5.1 and 5.2 provide introductions to quantum computing and quantum simulations, respectively; however, a cursory overview of quantum computers suffices to show that they are a powerful tool for studying quantum systems.

After abstracting away implementation details¹, a quantum computer consists of a set of qubits, and the ability to apply arbitrary unitary operations on pairs of qubits. The state of a single qubit is described by a two-dimensional Hilbert space $\mathcal{H}_2 = \text{span}\{|0\rangle, |1\rangle\}$. In a computer with N qubits, the full Hilbert space is given by the tensor product of N copies of the single-qubit Hilbert space, $\mathcal{H}_{QC} = \mathcal{H}_1^{\otimes N}$. This Hilbert space describes the set of possible states of the quantum computer; in addition, there is a set of unitary operations (termed ‘gates’) on this Hilbert space, which may be performed in any order in order to manipulate the qubits.

We are interested in simulating some physical system, described by a Hilbert space \mathcal{H} , and a time-evolution operator $U(t) = e^{-iHt}$. Here it becomes clear how a quantum computer might be useful. Two Hilbert spaces of equal dimension are

¹Crucially, this also requires abstracting away the fact that, as of this writing, no quantum computer exists at the scale necessary to perform any field theory simulation discussed in this chapter.

necessarily isomorphic, so it is possible to establish a mapping between the physical Hilbert space \mathcal{H} and (some linear subspace of) the quantum computer's \mathcal{H}_{QC} . If, after this mapping is established, the time-evolution operator $U(t)$ can be efficiently implemented in terms of the available quantum gates, then it will be possible to simulate time-evolution of the physical system with the quantum computer. In Section 5.2 we will see that, as shown in [65], this is true for a large class of physically relevant systems.

5.1 Digital Quantum Computers

In this section we give an expedited overview of quantum computation, tailored to those aspects which will be important in designing quantum simulations.

5.1.1 A Single Qubit

For physical intuition, one may think of a single qubit as being implemented by a quantum spin-1/2 system, although any two-state system will suffice and many are used in practice. The state of a single qubit is a vector in the Hilbert space $\mathcal{H}_1 = \text{span}\{|0\rangle, |1\rangle\} \approx \mathbb{C}^2$. There are two types of manipulations we perform on a qubit: quantum gates acting on 1 or 2 qubits, which correspond to unitary 2×2 or 4×4 matrices, and measurements, which yield classical information while collapsing the state of the qubit.

The set of unitary operators on \mathbb{C}^2 is denoted $U(2)$. An overall phase on a quantum state cannot be measured and is treated as physically irrelevant. For this

reason, the set of physically distinct quantum operations on one qubit is actually $U(2)/U(1) \approx SU(2)/\mathbb{Z}_2$.

Implementing the uncountable set of operations directly is often inconvenient, particularly when constructing an error-correct quantum computer. Instead, one implements a small discrete subset of these operations (fundamental gates), such that any unitary operator can be arbitrarily well approximated by a sequence of fundamental gates. A common set of fundamental gates are

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad \text{and} \quad T = \begin{pmatrix} e^{i\pi/8} & 0 \\ 0 & e^{-i\pi/8} \end{pmatrix}, \quad (5.1)$$

and as shown by Solovay and Kitaev, any operation in $U(2)/U(1)$ can be approximated to within ϵ with $O(1/\epsilon)$ gates chosen from this set [66–68].

The gate T , often referred to as the $\frac{\pi}{8}$ -gate, is the exponential of the Pauli matrix σ_z . Often T^\dagger is included in the set of fundamental gates, but it can of course be obtained as $T^\dagger = T^3$ (note that this equality is true in $U(2)/U(1)$, and not in $U(2)$). The Hadamard gate H , which is its own inverse, corresponds to a change-of-basis between the x and z bases: $H\sigma_zH = \sigma_x$.

In addition to gates, we may also perform measurements. For our purposes, we will consider all measurements to be performed in the z -basis of $|0\rangle$ and $|1\rangle$. If a qubit is in state $\alpha|0\rangle + \beta|1\rangle$, measurement changes the state to $|0\rangle$ (resp. $|1\rangle$) with probability $|\alpha|^2$ (resp. $|\beta|^2$), and yields the classical bit 0 (resp. 1).




Figure 5.1: A two-qubit circuit, implementing time-evolution under the Hamiltonian $H = \sigma_x \otimes \sigma_x$.

5.1.2 Coupling Many Qubits

A quantum computer with only one qubit is of no use — after all, it can be efficiently simulated by a classical computer by multiplying $SU(2)$ matrices. Quantum computers become interesting when we add the ability to operate on multiple qubits simultaneously. The resulting operations, with N qubits, are unitaries on the Hilbert space $\mathcal{H}_1^{\otimes N}$, modulo an irrelevant overall phase. This is the group $U(2^N)/U(1)$.

It is sufficient to add to our set of primitive gates only a single extra gate², which couples two qubits [70]. A common choice is the controlled-not gate, defined by

$$CX = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (5.2)$$

This gate may be thought of as acting on a ‘control’ and a ‘target’ qubit: the target qubit is flipped when the control qubit is 1. The matrix above is written in the basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$, so that the first (“high-order”) qubit is the control.

A quantum circuit is a composition of primitive gates. A simple example is shown in Figure 5.1, in which a change of basis is achieved by Hadamard gates,

²In fact, universal quantum computation can be achieved with almost any multi-qubit gate [69].

followed by a phase rotation of the $|11\rangle$ state, followed by a NOT gate on each qubit and another phase rotation. The result is equivalent to time-evolution under a Hamiltonian $H = \sigma_x \otimes \sigma_x$.

5.1.3 Some Simple Algorithms

A small number of quantum algorithms will be directly relevant to the task of creating a simulation, and are introduced here.

The first algorithm to discuss is really a meta-algorithm — that is, a procedure for producing a quantum algorithm: any reversible classical circuit may be re-interpreted as a quantum circuit. If a reversible quantum circuit starts with some bitstring x and yields $x + f(x)$ (that is, the concatenation of x and $f(x)$), then the corresponding quantum circuit yields the transformation $|x\rangle|0\rangle \rightarrow |x\rangle|f(x)\rangle$. (As $|x\rangle$ and $|f(x)\rangle$ are both computational basis states, this is a unitary transformation.)

As any classical circuit can be easily transformed to be reversible [71], this implies that any classical algorithm yields a quantum algorithm. This procedure requires polynomially many (in the memory size of the original, non-reversible classical circuit) ancillary qubits.

5.1.3.1 Controlled Nots

The controlled-not operation is usually considered to be a primitive gate, implemented directly by the quantum computing hardware. It is usually convenient, however, to make use of many-controlled nots, such as the Toffoli gate (here called




Figure 5.2: Implementation of the Toffoli gate from 1- and 2-qubit gates.

CCX) defined by

$$U_{\text{CCX}} |110\rangle = |111\rangle \quad \text{and} \quad U_{\text{CCX}} |111\rangle = |110\rangle, \quad (5.3)$$

and acting as the identity on all other basis states. These can be implemented from the single-controlled not CX and general one-qubit gates. Furthermore, arbitrarily-controlled nots $C^n X$ may be efficiently constructed (with $\text{poly}(n)$ gates) by introducing ancillary qubits. In fact this can be improved to remove the need for ancilla [72, 73], but these constructions will not be discussed here.

A common construction [70] of the twice-controlled not gate is shown in Figure 5.2. This is in fact the construction that minimizes the number of CX gates required [74]. From this building block, n -controlled not gates may be constructed with the introduction of $n - 2$ ancillary qubits, as demonstrated in Figure 5.3. Intuitively, each Toffoli may be thought of as adding two bits (base 2), and storing the result in the target bit. Eventually, the sum of all bits is accumulated.

5.1.3.2 Circuits From Circuits

Given a quantum circuit implementing a unitary U , certain related unitaries can be easily obtained by modification of the quantum circuit. For concreteness, we




Figure 5.3: Construction of a 4-controlled not gate from the Toffoli gate, with the aid of ancilla. The ancillary qubits are on the top; the target is the bottom qubit.

will assume that the circuit is implemented from the Hadamard gate, the $\frac{\pi}{8}$ gate, and the controlled-not. This constraint can be substantially relaxed without materially changing the methods discussed in this section.

Our first task, given a circuit implementing U , is to obtain a circuit implementing the inverse U^\dagger . Note first that, for every gate in our gateset, the inverse gate is already known: the Hadamard and **CX** gates are their own inverses, and $T^\dagger = T^3$. This allows us to construct an inverse circuit simply by inverting each gate and reversing the order of application.

$$U^\dagger = \left[\prod_{i=1}^K V_i \right]^\dagger = \prod_{i=K}^1 V_i^\dagger \quad (5.4)$$

Our second task, again given a quantum circuit implementing U , is to implement the controlled- U operation U_C defined by

$$U_C (|0\rangle \otimes |\Psi\rangle) = |0\rangle \otimes |\Psi\rangle \quad \text{and} \quad U_C (|1\rangle \otimes |\Psi\rangle) = |0\rangle \otimes U |\Psi\rangle \quad (5.5)$$

General constructions of controlled circuits are given in [73]; for our purposes, we

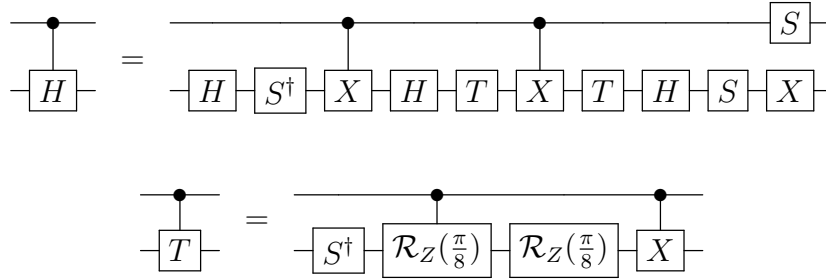


Figure 5.4: Construction of the controlled Hadamard (top) and controlled $\pi/8$ (bottom) gates, from the primitive gateset of H , $\mathcal{R}_Z(\theta)$, and CX . For brevity, an additional gate is defined as $S = T^2$, and we have used $X = HT^4H$.

will specialize to the case where the available gates are H , T , T^\dagger , and CX . The technique stems from the observation that if $U = VV'$ can be decomposed as a product of (potentially simpler) unitaries, then the controlled unitary is given by the product $U_C = V_C V'_C$ of the controlled versions of the simpler unitaries. By assumption we have U expressed as a product of the fundamental gates, and so we need only construct controlled versions of those four gates.

The Toffoli gate CCX has been constructed above. Constructions of the controlled Hadamard and controlled $\frac{\pi}{8}$ gates are shown in Figure 5.4.

5.1.3.3 Quantum Fourier Transform

The Fourier transform, suitably generalized to nonabelian groups is a change-of-basis operation between the regular representation of a group and the Fourier basis. This operation is unitary, and therefore can be implemented by quantum circuits. An efficient (polylogarithmic in the size of the group) implementation for arbitrary groups is not known; nevertheless, circuits are known for large classes of groups [75], including abelian groups [76].

The quantum Fourier transform is central to major quantum algorithms, most notably Shor's factoring algorithm [77] and Grover's search algorithm [78]. For quantum simulations, the relevance of the quantum Fourier transform comes from the fact that it diagonalizes the quantum-mechanical kinetic term ∇^2 of the Hamiltonian. We will see that the asymptotic scaling of the quantum Fourier transform with group size is not very relevant to the efficiency of lattice simulations.

5.2 Quantum Simulations in General

5.2.1 Mapping Hilbert Spaces

As discussed at the beginning of this chapter, the first step to setting up a simulation of a physical system on a quantum computer is to establish a mapping between the Hilbert space of the physical system and that of the quantum computer. The Hilbert space of the quantum computer is necessarily finite-dimensional, having 2^Q dimensions for a quantum computer with Q qubits; therefore, we are constrained to consider a physical system with a similarly finite Hilbert space. In particular, we will be working at finite volume and lattice spacing.

Any two Hilbert spaces of equal finite dimension are isomorphic. However, it is helpful to have the mapping between the states of the quantum computer and the physical states be a natural one. In particular, in simulating a field theory, it is often helpful to preserve the notion of locality. Many lattice field theories have a Hilbert space which is naturally expressed as a tensor product of many local Hilbert spaces, each associated to some lattice site (or link, for gauge theories). By grouping qubits,

and allowing each group to correspond to a single site, we can write the state space of the quantum computer similarly as a tensor product of simpler, local Hilbert spaces.

The nature of the mapping of local Hilbert spaces depends strongly on the system being simulated. The simplest case is when the local Hilbert space is two-dimensional, therefore mapping cleanly to a single qubit. This is true for a spin chain or the \mathbb{Z}_2 gauge theory, and we will use these systems as illustrative examples below.

5.2.2 Suzuki-Trotter Decomposition

The basic gates discussed above can be viewed as time-evolution under simple Hamiltonians affecting one or two qubits at a time. The one-qubit gates, applied to site i , yield evolution under $\vec{n} \cdot \vec{\sigma}_i$, while the two qubit gates applied to sites i and j yield evolution under $\sigma_i^\mu \otimes \sigma_j^\nu$ for any $\mu, \nu \in \{x, y, z\}$. A general Hamiltonian of interest, however, is not so simple. The Heisenberg spin chain, for example, has the Hamiltonian

$$H = -J_x \sum_{\langle ij \rangle} \sigma_i^x \otimes \sigma_j^x - J_y \sum_{\langle ij \rangle} \sigma_i^y \otimes \sigma_j^y - J_z \sum_{\langle ij \rangle} \sigma_i^z \otimes \sigma_j^z. \quad (5.6)$$

Note that because the terms of the Hamiltonian do not commute, the unitary time-evolution operator does not factorize. More complicated systems of physical interest will also have Hamiltonians which can be expressed as sums of few-qubit Hermitian operators which fail to mutually commute.

The Suzuki-Trotter decomposition [15, 16] provides an approximate factorization of the time-evolution operator in the case where the terms of the Hamiltonian do not commute. With two terms in the Hamiltonian $H = A + B$, we have

$$e^{-i\delta(A+B)} \approx e^{-i\delta A} e^{-i\delta B} + O(\delta^2). \quad (5.7)$$

If the decomposition is such that A and B have well-understood diagonal bases, the operators e^{-iAt} and e^{-iBt} are readily implemented, and therefore we have an easy implementation of approximate time-evolution, becoming exact in the limit $\delta \rightarrow 0$.

In general, this procedure can be generalized to any sparse efficiently computable sparse Hamiltonian [65]. In the case of field theories, the Hamiltonian can typically be split into two terms, one diagonal in field space and the other diagonal in conjugate momentum space; thus the general theorem is not needed.

5.2.3 Aside: Disordered Potentials

Here we describe the simulation [79] of a particular quantum-mechanical system: the Anderson tight-binding model [80] of a single particle living on V lattice sites. This is a model of a particle in a random potential. The Hamiltonian is

$$H = - \sum_{\langle ij \rangle} (c_i^\dagger c_j + c_j^\dagger c_i) + W \sum_i u_i c_i^\dagger c_i. \quad (5.8)$$

Here the u_i are random variables, taken to be independently and identically distributed on the interval $[0, 1]$. The first sum is taken over all pairs of neighboring

sites, and the parameter W gives the strength of the disorder.

Depending on W and the dimension of the lattice, this model may exhibit *Anderson localization*. When $W = 0$, this is a model of a free particle in a box. A wavefunction that begins concentrated at one site, spreads out throughout the lattice over time. For sufficiently strong disorder, however, an initially concentrated wavefunction will remain concentrated at all later times [80–82]. This localization effect is not due to any potential well trapping, but rather interference effects between the different paths a particle could take to propagate.

In one or two dimensions, any amount of disorder yields Anderson localization. In three dimensions, Anderson localization only sets in above the critical disorder of $W_c \approx 16.5$ [83], with a second-order transition at that point.

5.2.3.1 Simulation

The Hilbert space of this model is of dimension V . At large volumes, where the second-order phase transition is most visible, it becomes numerically difficult to simulate: naive algorithms run in time at least $O(V^2)$. When performing a simulation on a quantum computer, we expect to require only $O(\log V)$ qubits to represent the Hilbert space and similarly $O(\log V)$ operations per time step, indicating that far larger volumes can be obtained at relatively little cost.

We first discuss the simulation of a particle in a one-dimensional random potential [79]. Each of the V sites is labelled by an integer $0 \dots V - 1$, and the state $|i\rangle$ is the position eigenstate of an electron located at site i . This will be the com-

putational basis. The mapping of the computation basis of the quantum computer is achieved by representing the integer i in binary; thus the physical state $|6\rangle$ is mapped to the state $|110\rangle$ on the computer.

The time-evolution is simulated according to the Suzuki-Trotter decomposition, splitting the Hamiltonian into three pieces $H = H_{K,e} + H_{K,o} + H_V$, where:

$$H_{K,e} = \sum_{i=0,2,\dots} \left(c_i^\dagger c_j + c_j^\dagger c_i \right) \quad (5.9)$$

$$H_{K,o} = \sum_{i=1,3,\dots} \left(c_i^\dagger c_j + c_j^\dagger c_i \right) \quad (5.10)$$

$$H_V = W \sum_i u_i c_i^\dagger c_i. \quad (5.11)$$

Splitting the kinetic term into even and odd links in this fashion allows it to be simulated without a change of basis to momentum-space, and allows the algorithm to generalize to the case where the disorder in the potential lives on the links instead of the sites. When simulating in d dimensions, this Trotterization scheme requires $2d + 1$ steps.

The even links couple states that differ only in the last qubit

$$H_{K,e} = |0\rangle \langle 1| + |2\rangle \langle 3| + \dots + \text{h.c.} = I^{\otimes(V-1)} \otimes \sigma_x, \quad (5.12)$$

and $e^{-iH_{K,e}\Delta t}$ is therefore obtained by a rotation about the X axis of the least-significant qubit. The evolution of the odd links can be put into a similar form by first shifting the whole lattice by 1. This is a change of basis that maps $H_{K,o}$ to

$H_{K,e}$. The shift corresponds to the addition of 1 modulo V , for which a classical circuit (and therefore quantum circuit) is readily constructed.

The evolution under the disordered potential is, on its face, more difficult. This evolution requires the phase of the state to be changed by the same random number each time the electron finds itself at a particular site. In a classical simulation, this effect is accomplished by generating a list of V random numbers at the beginning of the computation. This step already exponentially exceeds our $O(\log V)$ budget.

To avoid this, we note that the u_i are typically not truly random variables, but instead are defined to be the output of a pseudo-random number generator (PRNG) with a seed chosen in advance. A PRNG is a circuit sufficiently complicated that the u_i look random to any practical statistical test. Given a PRNG $f(i)$ returning a Q -bit number, we may take $u_i = 2^{-Q}f(i)$. Critically, there exist *seekable* PRNGs, from which the i th element $f(i)$ can be obtained in fixed time for any i . Two appropriate constructions of PRNGs are discussed below.

Given a classical circuit for a suitable PRNG, we can construct a quantum circuit U_f defined by $U_f |i\rangle |0\rangle = |i\rangle |f(i)\rangle$. The evolution under H_V is then implemented by applying U_f to compute the PRNG, and a diagonal phase rotation by the value specified in the ancillary register.

The resulting circuits corresponding to kinetic and potential evolution are shown in Figure 5.5.

Whether a system is localized or not can be detected by placing a particle on the lattice (at the origin, say), allowing it to diffuse for a long time, and then observing how close to the origin it remains on average. In the non-localized phase,

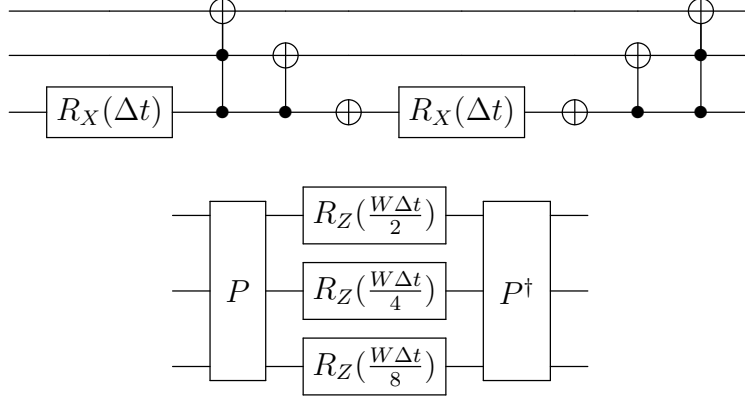


Figure 5.5: Quantum circuits for the simulation of the Anderson model. The time-evolution is trotterized: on top is the kinetic piece, and on bottom is the evolution under a random Hamiltonian. For brevity, $R_X(\theta)$ and $R_Z(\theta)$ denote $e^{i\theta\sigma_x}$ and $e^{i\theta\sigma_z}$, respectively. The construction of a pseudo-random permutation operator P is discussed in the text.

the particle will diffuse out to infinity; in the localized phase, it will remain within some finite distance. When simulating the Anderson transition at finite volume on a periodic lattice, the degree of localization of a wavefunction can be measured by

$$D = \frac{L}{\pi\sqrt{2}} \sqrt{1 - \left\langle \cos \frac{2\pi\hat{x}}{L} \right\rangle}, \quad (5.13)$$

which, in the large volume limit, yields the average distance. This serves as an order parameter for the Anderson transition. On a three-dimensional lattice, this order parameter as a function of disorder W is seen in Figure 5.6, with the expected transition being visible even at such small volumes, near $W \sim 15$.

5.2.3.2 A Pseudo-Random Number Generator

We now discuss the construction of a suitable (seekable) pseudo-random number generator. A particularly straightforward construction of a seekable PRNG uses

a cryptographic hash function like SHA256 [84]. A sequence of K PRNGs indexed by seeds k are constructed via $f_k(i) = \text{SHA256}(i * K + k)$. This construction is validated and known to perform well in practice [85]. Unfortunately, near-term quantum computers do not have enough qubits available to compute modern cryptographic hash functions, which operate on fixed-size registers of hundreds of bits.

An alternative approach is to construct the function $f(i)$ from a random reversible classical circuit (which therefore implements a random permutation matrix).

A random permutation matrix P defines a seekable PRNG via $P|i\rangle = |f(i)\rangle$. Although the reversibility of the operation implies weak correlations between the different values of $f(i)$, these correlations are unmeasurable in the large volume limit and can be neglected.

We construct a random circuit by appending a fixed-length sequence of not Toffoli gates, acting on random arguments. With the aid of $Q - 3$ ancillary qubits, any permutation matrix can be obtained in this way [71]. In the limit of a large number of gates, this samples uniformly from the distribution of permutation matrices.

This construction of a PRNG was validated against the `dieharder` [86, 87] battery of statistical tests in [79]. For circuits acting on 30 bits, a sequence of 600 random gates was sufficient to consistently pass all statistical tests in the battery, and the number of gates required was found to scale polynomially with the number of bits (and therefore polylogarithmically with the volume).

This construction is further validated in Figure 5.6, where it is shown that the same physical results are obtained in a simulation using this construction as in a simulation using a conventional (sequential) PRNG.




Figure 5.6: Figure from [79], showing average distance as defined by (5.13), as a function of the disorder parameter W , for a 16^3 lattice, in the limit of long time evolution. The blue triangles show results obtained with a conventional PRNG, and the red squares give results obtained with the circuit-based PRNG described in Section 5.2.3.2. The dashed line shows the delocalized limit for D .

5.2.3.3 Demonstration

The method described above has the nice property that even very small numbers of qubits can be used in a sensible simulation, albeit on a small lattice. The field theory simulations we will look at later don't have this property, as representing even one link can require many qubits (11, in the main example). Furthermore, because the size of the lattice is exponential in the number of qubits (in contrast to the scaling for simulating a field theory), even near-term quantum computers can simulate lattice sizes at which localization can be seen.

Figure 5.7 demonstrates the algorithm on two physical quantum processors, one provided by IBM (programmed with `qiskit` [88]) and the other by Rigetti (programmed with `quil` [89]). The simulation is done on two qubits, and therefore involves four lattice sites. There is no sensible notion of a random potential on four




Figure 5.7: Figure from [79], showing average distance as measured by (5.13), as a function of evolution time for a 4-site lattice with disorder parameter $W = 5$ and Trotterization step size $\Delta t = 0.2$.

sites; we fix the potential to be $V(0, 1, 2, 3) = 0, 1, 3, 2$, which is computed with a single CX gate. On each processor, for each point, we perform 300 quantum measurements to estimate $D(t)$.

5.3 Simulating a Field Theory

Now we turn to the problem of simulating a field theory on a quantum compute. For concreteness, let us focus on the Heisenberg spin chain, with Hamiltonian

$$H = -J \sum_{\langle ij \rangle} \sigma_z(i) \sigma_z(j) - \mu \sum_i \sigma_x(i) \quad (5.14)$$

where the first sum runs over all pairs of adjacent sites on a one-dimensional lattice, J is the ferromagnetic coupling, and μ is a magnetic field. Because the local degrees of freedom possess a two-dimensional Hilbert space, this model is particularly amenable to qubit-based quantum simulation. Nevertheless, the features of the simulation

of this model are essentially the same as those of other field theories; moreover, many models can be rewritten (at least approximately) as spin chains, such as σ models [90]. The generalizations to scalar ϕ^4 field theory and fermionic fields are discussed in Sections 5.3.3 and 5.3.4, respectively.

The first step of preparing a quantum simulation is mapping the physical Hilbert space to that of the quantum computer. For the spin chain, the mapping is trivial: each site of the spin chain corresponds to a single qubit.

The next step is to select a Suzuki-Trotter decomposition of the Hamiltonian. A natural choice for (5.14) is

$$e^{-iH} \approx \left(e^{i(\Delta t)\mu \sum \sigma_x} e^{i(\Delta t)J \sum \sigma_x \sigma_x} \right)^{t/\Delta t}. \quad (5.15)$$

A more general Heisenberg Hamiltonian is possible, with arbitrary couplings and magnetic fields along all three axes:

$$H = - \sum_{\langle ij \rangle} [J_x \sigma_x(i) \sigma_x(j) + J_y \sigma_y(i) \sigma_y(j) + J_z \sigma_z(i) \sigma_z(j)] - \mu \cdot \sum_i \sigma(i). \quad (5.16)$$

In this case a natural Trotterization has three factors, diagonal in the x , y , and z -bases individually.

Because the mapping between the physical Hilbert space and that of the quantum processor is so clean, the resulting time-evolution circuits are particularly simple. Evolution for a single Trotterization step is shown in Figure 5.8. Note that the order in which the coupling terms are applied is irrelevant, as they all commute.




Figure 5.8: Quantum circuit for a single Trotter step of time-evolution for the Heisenberg spin chain with Hamiltonian (5.14), on four lattice sites.

This completes the implementation of time-evolution on the quantum computer. Two important steps remain: the preparation of a physically relevant initial state, and measurement at the end of some time evolution. Typically the initial state of greatest interest is the ground state of the Hamiltonian, or the lowest-lying state constrained to have some quantum numbers. Preparation of such states is a major area of study, and some strategies are discussed in detail in Sections 5.5.2 and 5.6.

Assuming that the ground state $|\Omega\rangle$ has been prepared, there are a wide selection of physically interesting observables readily accessible. By allowing the magnetic field to be time-dependent (and even space-dependent, at no cost to circuit complexity), the response of the system to time-varying magnetic fields can be measured.

5.3.1 Linear Response, Two Ways

A frequently relevant subset of these observables are those obtained by considering the limit of a weak perturbation. In the case of a delta-function perturbation

described by the Hamiltonian

$$H_\epsilon(t) = H_0 + \epsilon\delta(t)H', \quad (5.17)$$

the expectation value of an operator \mathcal{O} at time T after the perturbation hits is given, to leading order, by

$$\langle \mathcal{O}(t) \rangle = \langle \mathcal{O}(0) \rangle + \epsilon \langle [H', \mathcal{O}(t)] \rangle + O(\epsilon^2). \quad (5.18)$$

Thus we see that *linear* response is governed by time-dependent correlation functions evaluated in the original Hamiltonian.

This discussion can be re-cast as a method for evaluating time-dependent correlation functions. Evolution under a time-dependent Hamiltonian is easily achieved on a quantum computer. We can evaluate, therefore, the expectation value $\langle \mathcal{O}(t) \rangle$ under no perturbation and under a small perturbation, and finite-differencing yields an approximation to the expectation value $\langle [H', \mathcal{O}(t)] \rangle$.

An alternative procedure [91], not limited to the commutator of Hermitian operators, is available with the use of an ancillary qubit. To begin with, we describe a procedure for measuring the expectation value of an arbitrary unitary operator [92]. Let U be a unitary operator, for which we would like to measure the expectation value $\langle \Psi | U | \Psi \rangle$ in some state $|\Psi\rangle$. Define U_C to be the controlled- U unitary acting on the combination of the original system with the one ancillary qubit. Thus U_C is

defined as

$$U_C |\Psi\rangle |0\rangle = |\Psi\rangle |0\rangle \quad \text{and} \quad U_C |\Psi\rangle |1\rangle = (U |\Psi\rangle) |1\rangle. \quad (5.19)$$

A circuit implementing this unitary can be obtained using the technique described in Section 5.1.3.2. We now begin by applying a hadamard gate to the ancillary qubit, then apply U_C , and finally measure the controlled qubit. Depending on the basis chosen to measure the controlled qubit, obtain either the real or imaginary part of the desired expectation value.

An obvious application is to the Heisenberg spin chain. The time-separated correlator $\sigma_x(t)\sigma_x(0) = e^{iHt}\sigma_x e^{-iHt}\sigma_x$ is unitary, and therefore can be measured directly. However, the method is slightly more general: the ability to measure an arbitrary unitary allows us also to measure any operator which can be decomposed as a sum of unitaries.

5.3.2 Measuring Masses

Measuring energies on a quantum computer may be accomplished via the algorithm of quantum phase estimation [93] (QPE). Given a unitary operator U (implemented via quantum circuits) and a prepared eigenstate $|\Psi\rangle$, QPE is a procedure for estimating the phase θ of the eigenvalue $e^{i\theta}$ of the prepared state. When the unitary operator is time-evolution, this phase is of course the energy of the state. The simplest form of quantum phase estimation proceeds by introducing an ancillary qubit in the state $|0\rangle + |1\rangle$, and performing controlled evolution under U . After this evolution, the state of the system is $(|0\rangle + e^{i\theta}|1\rangle)|\Psi\rangle$, and θ may be estimated

modulo π via repeated measurements in the Z -basis. With a little more sophistication, the binary representation of θ can be determined with a single measurement, to precision $1/\epsilon$, with $\log(1/\epsilon)$ ancilla.

We would like to measure the mass of a hadron — that is, the difference in energies between the vacuum and the lowest-lying state with quantum numbers of that hadron. Assume we have the ability to prepare both the ground state of the lattice theory and the ground state of the sector with quantum numbers of some hadron. The most straightforward procedure to obtain the mass of that hadron is to first prepare the ground state $|\Omega\rangle$, and then measure via QPE the energy E_Ω of this state. On a lattice of the same parameters, we may prepare the ground state $|P\rangle$ of the hadron, and similarly measure the energy E_P of that state. The mass is then given by $E_P - E_\Omega$. This method is simple, but suffers from a significant flaw, related to the fact that E_Ω and E_P are not sensible physical quantities. The vacuum energy is divergent in both the continuum and infinite-volume limits. Therefore, as these limits are approached, both energies must be measured with increasing precision to resolve the cancellation, before any information about the mass is obtained. This is another signal-to-noise problem.

This signal-to-noise problem can be done away with by preparing two lattices at once on the same quantum processor. These lattices are uncoupled, and we prepare in the first the ground state and in the second the hadron state, so that the quantum processor is in the state $|\Omega\rangle \otimes |P\rangle$. We now consider the unitary operator $U(t) = e^{-iHt} \otimes e^{iHt}$. The prepared state is an eigenstate of this operator; moreover, the divergent part of the energies cancel. Thus, QPE applied to U directly yields

the hadron mass, with no need to resolve fine cancellations.

5.3.3 Scalar Fields

Lattice scalar field theory was examined as a target for quantum simulation in [9]. For a lattice scalar field theory described by the Hamiltonian (2.4), each lattice site is associated to an anharmonic oscillator, and the oscillators are coupled by the term $(\phi_x - \phi_y)^2$. This system presents a new difficulty for quantum simulations, which is in fact characteristic of most bosonic field theories: the local Hilbert space is of infinite dimension, and there is therefore no isomorphic Hilbert space that can be created with a finite number of qubits.

The only solution is to truncate the physical Hilbert space. To maintain the locality of the theory, it is convenient to truncate the local Hilbert spaces independently, so that the full Hilbert space remains a tensor product, but now of finite-dimensional systems.

A reasonable truncation for this system becomes apparent once the Hamiltonian is rewritten in terms of creation and annihilation operators:

$$H = \sum_{\langle xy \rangle} (a_x^\dagger a_y + \text{h.c.}) + \sum_x \left[a_x^\dagger a_x + \lambda (a_x^\dagger + a_x)^4 \right]. \quad (5.20)$$

Truncating each local Hilbert space to the lowest K eigenstates of the harmonic oscillator Hamiltonian $a^\dagger a$, the number of qubits required for the full simulation scales with $V \log K$. When performing computations, there are now three extrapolations that must be performed: removing the truncation ($K \rightarrow \infty$), removing the lattice

cutoff ($a \rightarrow 0$), and the infinite volume limit ($V \rightarrow \infty$), in that order.

5.3.4 Fermions

Fermionic lattice theories are limited to a finite Hilbert space by the Pauli exclusion principle. Each fermionic degree of freedom (there may be many per lattice site, due to spin, flavor, and internal symmetries) is associated to a two-dimensional Hilbert space, preparing a convenient mapping between the physical space and that of the processor.

A difficulty arises, however, when attempting to map the operators used to define the physical Hamiltonian, to operators defined on the qubits of the quantum computer. The fundamental operators of a fermionic theory are the raising and lowering operators a^\dagger and a , defined to anticommute:

$$\{a_i, a_j\} = 0 \quad \text{and} \quad \{a_i^\dagger, a_j\} = \delta_{ij}. \quad (5.21)$$

The natural raising and lowering operators defined on qubits, given by $\sigma^\pm(i) = \sigma_x(i) \pm i\sigma_y(i)$, commute at different qubits. Anticommuting operators must be constructed on the quantum computer in order to map the physical Hamiltonian over to the qubits.

A suitable appropriate mapping of operators is provided by the Jordan-Wigner transformation [92, 94], in which a_i maps to

$$\left[\bigotimes_{k=1}^{i-1} \sigma_z \right] \otimes \sigma^- \otimes I. \quad (5.22)$$

Two such operators are readily seen to anticommute, and so a fermionic Hamiltonian can be rewritten in terms of them.

The Jordan-Wigner transformation does violence to locality³. The string of σ_z operators imposes an extra cost, typically polynomial in the volume ($V^{2/3}$ for a three-dimensional lattice). The transformation can be improved to alleviate or remove this asymptotic cost. The first such improvement was due to Bravyi and Kitaev, and reduced this to a logarithm of the volume of the lattice [95]. More recent improvements result in a constant overhead [96, 97], removing all asymptotic penalty.

5.4 Simulating a Gauge Theory

The simplest gauge theory to simulate is the \mathbb{Z}_2 gauge theory of Section 2.3, with Hamiltonian (2.14). This is only a small modification from the spin system considered above. Each degree of freedom is now associated to a link on the lattice, instead of a site, and the coupling term diagonal in the Z basis couples four degrees of freedom (one plaquette) rather than just two.

The procedure for simulating time-evolution of such a theory is consequentially simple. Figure 5.9 shows a single step of time evolution of a one-plaquette model. This major new feature of this procedure is the use of an ancillary qubit to compute the value of the plaquette. The qubit begins in state $|0\rangle$, and each of the four links in the plaquette are combined via controlled not operations (corresponding to the

³With the exception of one-dimensional lattices without periodic boundary conditions, where an Jordan-Wigner transformation can be constructed such that local fermion bilinears map to local spin operators.

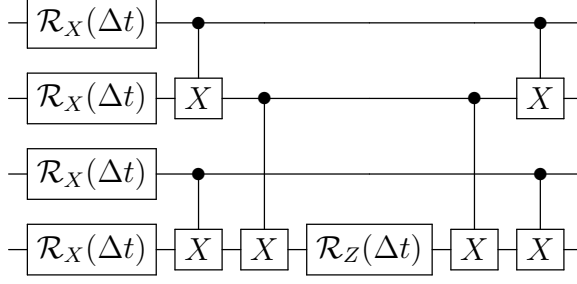


Figure 5.9: Quantum circuit for a single Trotter step of the time-evolution of \mathbb{Z}_2 gauge theory with one plaquette.

\mathbb{Z}_2 group operation) until the ancillary qubit contains the value of the plaquette.

Generalizing to an arbitrary gauge group G , we introduce the concept of a G -register [20]: a collection of qubits whose Hilbert space is mapped to $\mathbb{C}G$. For continuous gauge groups, where the space $\mathbb{C}G$ is infinite-dimensional, this is necessarily some approximation, however we will ignore this difficulty for the time being. The Hilbert space of a G -register is spanned by states $|g\rangle$, which we will take to be the computational basis. On this space a set of primitive operations are needed:

- An inversion gate, which takes a G -register and transforms, in the computational basis, by taking the inverse of the group element. The gate is defined by $\mathfrak{U}_{-1} |U\rangle = |U^{-1}\rangle$. This gate is self-adjoint.
- A multiplication gate, which acts on two G -registers and transforms the second. This gate is defined by $\mathfrak{U}_{\times} |g\rangle |h\rangle = |g\rangle |gh\rangle$.
- A trace gate, which gives each computational basis state a phase proportional to the trace of the stored group element. This gate is defined by $\mathfrak{U}_{\text{Tr}}(\theta) |U\rangle = e^{i\theta \text{Re Tr } U} |U\rangle$. Note that the definition of this gate depends on the representation, and indeed, some lattice Hamiltonians may involve traces

taken in multiple representations.

- The (nonabelian) Fourier transform \mathfrak{U}_F , which transforms a G -register into Fourier space (a \hat{G} -register). It is defined by

$$\mathfrak{U}_F \sum_{U \in G} f(U) |U\rangle = \sum_{\rho \in \hat{G}} \hat{f}(\rho)_{ij} |\rho, i, j\rangle \quad (5.23)$$

where the second sum is taken over all representations ρ of G , and \hat{f} denotes the Fourier transform of f . This gate is the operation described in Section 5.1.3.3, and diagonalizes the kinetic part of the Hamiltonian.

- The Laplace-Beltrami gate \mathfrak{U}_{LB} , which acts on a \hat{G} -register and gives each state a diagonal phase, which is a function of the representation alone (not the indices i, j).

A few note about the generality of these operations are in order. First, although the multiplication gate \mathfrak{U}_\times is defined here to perform left multiplication, a gate for right multiplication is obtained from the combination of \mathfrak{U}_\times and \mathfrak{U}_{-1} via the identity $\mathfrak{U}_{\times,R}(1, 2) = \mathfrak{U}_{-1}(1)\mathfrak{U}_{-1}(2)\mathfrak{U}_\times(2, 1)\mathfrak{U}_{-1}(2)\mathfrak{U}_\times(1, 2)$.

From these operations we can construct time-evolution for a pure gauge theory. In the presence of matter fields, a few additional ones will be needed, discussed in Section 5.4.3 below.

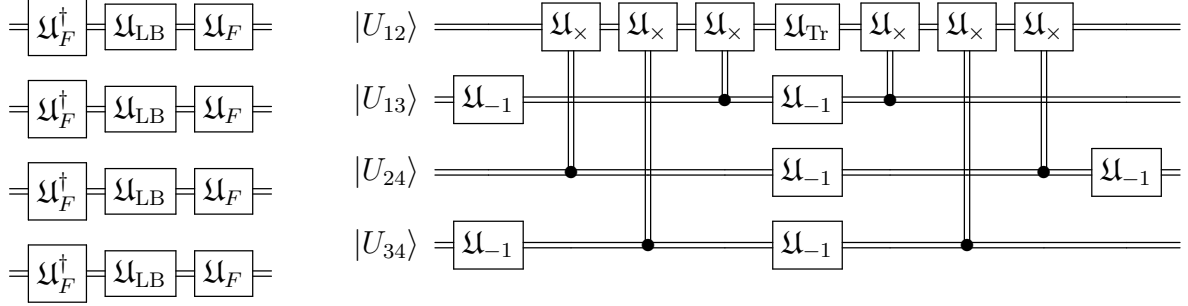


Figure 5.10: Circuits implementing the time-evolution of a pure-gauge lattice field theory. The first circuit implements the quantum-mechanical kinetic term, and the second the quantum-mechanical potential term associated to a single plaquette $\text{Re Tr } U_{13}^\dagger U_{34}^\dagger U_{24} U_{12}$. Note that in these circuits, the primitive object is a G -register, denoted with a doubled line.

5.4.1 Time Evolution

Time-evolution of the pure gauge theory for a general group G is performed, as usual, with a Trotter-Suzuki decomposition. The kinetic piece of the evolution is diagonalized by the Fourier transform. The potential piece of the evolution requires, as for the \mathbb{Z}_2 gauge theory, the accumulation of a plaquette in a single G -register. This is accomplished by picking one link in the plaquette and repeatedly multiplying with all other links in the plaquette.

The resulting circuits for the propagation of a single-plaquette, nonabelian gauge theory are shown in Figure 5.10.

5.4.2 Gauge Invariance

The time-evolution circuits presented above can only be said to represent a gauge theory when the initial state lies in the physical subspace; i.e., is gauge-invariant. It is critical, therefore, that we are able to prepare gauge-invariant states

in general.

Two gauge-invariant states are particularly easy to prepare: the strong-coupling ground state and the weak-coupling ground state. In the strong-coupling limit, the kinetic term of the Hamiltonian dominates. This term does not couple links, and the resulting ground state is a product state, particularly easy to prepare:

$$|\Omega_{\text{strong}}\rangle = \bigotimes_L \left(\sum_{U \in G} |U\rangle \right). \quad (5.24)$$

Here each link is in an equal superposition of all group elements. The state described here is already gauge-invariant, so no further symmetrization is needed.

In the weak-coupling limit, each plaquette is forced to be the identity in the ground state⁴. Naively, then, the ground state is the product state where each link is set to $|I\rangle$; however, this state is not gauge-invariant. We can prepare a gauge-symmetric version of this state with the aid of one ancillary G -register per lattice site. These registers represent a gauge transformation, and we will denote them V . Begin by initializing each link to $|I\rangle$, and each ancillary register to an equal superposition of all group elements. We may now perform the gauge transformation: the link U_{ji} from site i to site j is multiplied on the right by V_i^\dagger and on the left by V_j . This yields the state

$$\sum_V \left| (V_2 V_1^\dagger) \cdots \right\rangle |V_1 V_2 \cdots\rangle. \quad (5.25)$$

⁴Assuming that such a configuration is permitted by boundary conditions — otherwise, we would have a frustrated system, and no general efficient algorithm for preparing the ground state.

At this point, the physical link registers are entangled with the (unphysical) gauge registers. The entanglement is removed by again repeatedly applying \mathfrak{U}_x , but now with the link registers as the control and the gauge registers as the target. A privileged site i is selected (the choice will not affect the final state), and for each other site j , a particular path from j to i is selected. After one multiplication for each link in this path, the ancillary register associated to site j is transformed to the state $|V_i\rangle$. Repeating for all ancilla, we obtain the state

$$\sum_V |(V_2 V_1^\dagger) \cdots\rangle |V_i V_i \cdots\rangle. \quad (5.26)$$

That this state is in fact a product state may be seen by noting that a global gauge transformation by V_i^\dagger leaves the physical registers invariant, while rotating each ancilla into $|I\rangle$. The disentangled ancilla may now be discarded, and the resulting state is the gauge-projected $P|I\rangle$.

This suffices to show that the gauge-invariant sector is efficiently accessible. Further discussion of state preparation, in particular the preparation of physical ground states and thermal states, is in Sections 5.5.2 and 5.6 below.

5.4.3 Adding Fermions

The most interesting gauge theories have matter fields coupled to the gauge degrees of freedom. We will consider here the case of fermionic matter fields, as in QCD. In order to introduce matter fields, one must first pick a representation of the

group G under which the matter fields are to transform⁵. For matrix groups (e.g. $SU(N)$), it is most common to select the fundamental representation.

Labeling the dimension of the selected representation by N , and ignoring for simplicity spin and flavor, we will have N independent fermionic degrees of freedom at each lattice site, creating a local Hilbert space of dimension 2^N . Denote the creation and annihilation operators a_i^\dagger and a_i , where $i \in 1 \dots N$ is the ‘color’ index of the chosen representation ρ . These operators transform into each other under the action of the group G :

$$U^\dagger a_i U = \sum_j \rho(U)_{ij} a_j. \quad (5.27)$$

Gauge transformations, and the projection operator P , now affect the fermionic modes as well.

Crucially, the VN fermionic degrees of freedom can be created via the Jordan-Wigner transformation (or any other method) before making reference to the fact that they transform into each other under various symmetries. In other words, the presence of gauge-invariance does not complicate the task of creating anticommuting operators from the fundamental commuting operators of a quantum computer.

5.4.4 Demonstration: D_4 Gauge Theory

The methods of the previous sections can be demonstrated on a classically simulated quantum computer, with a small finite gauge group. The smallest two nonabelian groups are D_3 and D_4 , defined as the group of isometries of the triangle

⁵The gauge fields themselves have no choice but to transform in the adjoint representation.




Figure 5.11: The lattice geometry of the D_4 gauge theory simulated here. Dashed lines indicate repeated links due to the periodic boundary conditions.

and the square, respectively. Because D_3 has 6 elements and D_4 has 8, they each require three qubits per link to simulate, and therefore we may as well simulate D_4 .

The first order of business is to construct a G -register; that is, to choose a particular isomorphism between the space of complex-valued functions on D_4 , and the Hilbert space of three qubits. The group D_4 can be defined as the subgroup of $U(2)$ generated by the matrices

$$\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \text{ and } \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \quad (5.28)$$

The state $|abc\rangle$ is defined to correspond to the matrix

$$\left[\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \right]^a \left[\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \right]^{2b+c}. \quad (5.29)$$

We next construct the inversion, multiplication, trace, and Fourier transform circuits. The inversion and multiplication circuits are classical circuits, easily constructed. As the only element of D_4 with a nonvanishing trace in this representation is the identity, the trace circuit is a three-qubit-controlled phase gate. The cir-




Figure 5.12: Figure from [20]: simulation of the two-plaquette D_4 gauge theory. The expectation value of one plaquette as a function of time. The exact result is shown in black, with sampled data shown in red (blue) for a Trotterization time-step of $\Delta t = 0.2$ ($\Delta t = 0.5$).

cuits for these three operations and the Fourier transform are given in the appendix of [20].

Each link requires three qubits, so a classical simulation can simulate a lattice containing four links without much trouble. (Substantially larger lattices — up to ~ 30 qubits — could be obtained with more sophisticated algorithms and reasonable computer time; however, this is pointless for the purposes of demonstration.) A two-plaquette geometry based on a four links is shown in Figure 5.11. On this geometry, we initialize the gauge-invariant ground state in the weak-coupling limit, and measure the expectation value of a plaquette as a function of time. The result is shown in Figure 5.12.

5.5 Simulations of QCD

We now apply the general method of [20] to the simulation of QCD with a quantum computer. The largest difficulty encountered is the one mentioned previously: the dimension of the Hilbert space of $SU(3)$ gauge theory, even on a finite lattice, is infinite, and therefore any quantum system consisting of qubits must be only an approximation. A particular subgroup of $SU(3)$, the Valentiner group, is shown to provide an adequate approximation, and from there we can discuss how to prepare interesting states and extract partonic physics.

5.5.1 The Valentiner Group

The Valentiner group [98] \mathcal{V} , also referred to as $\tilde{S}(1080)$ [99], is a finite subgroup of $SU(3)$ with 1080 elements. It is not the largest finite subgroup: \mathbb{Z}_n , for instance, is a subgroup for any n , due to the presence of $U(1) < SU(3)$. It is, however, the largest *exceptional* subgroup, that is, the largest subgroup that does not fall into one of a small number of infinite families.

The Valentiner group is particularly suitable as an approximation to $SU(3)$ because it tiles the surface of $SU(3)$ evenly. In particular, the Voronoi diagram of $\mathcal{V} \subset SU(3)$ has two special properties:

- Each region is isomorphic to every other region. This follows from the fact that \mathcal{V} , being a subgroup of $SU(3)$, is also a symmetry group of the Voronoi diagram itself.




Figure 5.13: Figure from [100] of the average energy per plaquette vs β_0 for $SU(3)$, \mathcal{V} with $\beta_1 = 0$, and \mathcal{V} with $\beta_1 = -0.6$.

- Each face — the boundary between two regions — is isomorphic to every other face, and in fact every face is the same distance from the center of the neighboring regions.

The group \mathcal{V} thus serves as a good approximation to $SU(3)$ in much the same sense as a dodecahedron might serve as an approximation to the sphere⁶.

The degree to which \mathcal{V} is a good approximation for $SU(3)$ is a matter for empirical study, first investigated in [101]. Critically, this question can be addressed by performing classical Monte Carlo calculations (in imaginary time), without the need for a quantum processor [100]. We will work with the lattice action

$$S = - \sum_P \left(\frac{\beta_0}{3} \text{Re Tr } P + \beta_1 \text{Re Tr } P^2 \right), \quad (5.30)$$

where the partition function is defined over all gauge configurations with links taking values in \mathcal{V} , rather than the $SU(3)$. With $\beta_1 = 0$, this is the standard Wilson action.

This lattice theory is compared to the Wilson gauge theory of $SU(3)$ in Fig-

⁶In fact the platonic solids correspond directly to the nice approximations of $SU(2)$.

ure 5.13, using the average energy per plaquette as a probe. The two are qualitatively different. The $SU(3)$ theory possesses a crossover from strong coupling (at small β_0) to weak coupling around $\beta_0 \sim 5$. For the $\beta_0 = 0$ version of the Valentiner gauge theory, a first-order transition is encountered before this crossover. Smaller average energy is a proxy for a closer approach to continuum physics; failing to reach the so-called scaling regime indicates that the Valentiner theory is far from the $SU(3)$ continuum theory.

The breakdown of the Valentiner gauge theory as we approach the continuum limit can be understood by noting that, in $SU(3)$ gauge theory, the fluctuations of the gauge fields about the identity become smaller as β_0 is increased. In the theory of $S(1080)$, there is limit past which the fluctuations cannot get any smaller; there are no group elements arbitrarily close to the identity. Beyond a critical value $\beta_0 \sim 4$, the fields become fixed.

The situation can be improved by noting that for the $SU(3)$ theory, the same continuum limit is expected to be reached as $\beta_0 \rightarrow \infty$ regardless of the value of β_1 . We see in Figure 5.13 that, increasing β_0 along this trajectory, lower values of the average energy are achieved, suggesting a closer approach to the continuum physics of $SU(3)$.

So far this is just a heuristic argument that the gauge theory of \mathcal{V} can approximate the $SU(3)$ theory, if an appropriate trajectory in (β_0, β_1) is selected. We can make this more rigorous by performing, on each theory separately, a measurement at multiple points along the continuum trajectory. Then, we extrapolate (again in each theory separately) to the continuum, and compare. Note that we perform a




Figure 5.14: Continuum extrapolations of $T_c\sqrt{t_0}$ for $SU(3)$ and $S(1080)$ gauge theories. The $S(1080)$ result is shown in blue, with the extrapolated value marked by the solid shape; the $SU(3)$ extrapolated results are also indicated by solid shapes.

continuum extrapolation on the \mathcal{V} theory despite the fact that, due to the first-order transition, we know that the \mathcal{V} theory cannot reach the continuum at all. The process of extrapolation, however, cannot see the (nonanalytic) first-order transition, and ends up extrapolating to the continuum $SU(3)$ result.

Both the $SU(3)$ and $S(1080)$ gauge theories possess a phase transition in the temperature (that is, the temporal extent of the Euclidean lattice). The chosen measurement is the dimensionless ratio $T_c\sqrt{t_0}$, where T_c is the temperature of this phase transition, and t_0 is a scale set by the Wilson flow [102]. Figure 5.14 shows the continuum extrapolation of $T_c\sqrt{t_0}$ for the Valentiner theory with a trajectory defined by

$$\beta_1 = -0.1267\beta_0 + 0.253 \quad (5.31)$$

against the measurements of [103, 104]. The Valentiner group serves as a good approximation to the $SU(3)$ theory for this low-energy observable.

The task of comparing other low-energy observables, particularly spectroscopy, is not complete as this is written. It must be remembered that the use of \mathcal{V} as an

approximation to $SU(3)$ carries a serious drawback: because \mathcal{V} is the largest ‘nice’ subgroup, this is not a systematically improvable approximation. The only hope for improvement is to add more terms to the action, and seek a trajectory that provides a better approximation to continuum physics.

5.5.2 State Preparation

The easiest states to prepare on a quantum computer are the eigenstates of field operators. The infinite-temperature thermal state is also readily accessible. Unfortunately, these states are not of much physical interest. We would like to prepare the ground state of our Hamiltonian, perhaps restricted to some sector. For example, the ground state of the baryon-number-1 sector of QCD contains a single proton.

A large body of literature exists with many proposed methods for this task; see [9, 34, 105–107] for a very incomplete sample. Quantum computers have not yet been constructed that are large enough to investigate the practical performance of these methods, and formal analysis of most preparation methods is not available. One exception is adiabatic state preparation [9, 34, 107]. This method is backed by the adiabatic theorem [108], allowing us to make crude estimates of the costs of this method as applied to physical theories. Here we will discuss in detail the cost of preparing the ground state of the baryon-number-1 sector, i.e. the proton. This does not imply that adiabatic state preparation is the most efficient or most practical method, merely that it is the only one for which *a priori* analysis is currently

possible.

The adiabatic theorem [108] constrains the behavior of a quantum system with a time-dependent Hamiltonian. Let $H(t)$ be our Hamiltonian, and let $|\Omega\rangle$ be an eigenstate of $H(0)$, prepared at time $t = 0$. When the time-dependence of the Hamiltonian is mild compared to the size Δ of the gap $\Delta(t)$ between $|\Omega\rangle$ and the nearest eigenstate (when $\dot{H}/\Delta(t)^2 \ll 1$ for all t), the adiabatic theorem states that the system will never transition to a different eigenstate⁷. In particular, if $|\Omega(0)\rangle$ is the ground state of $H(0)$, then $|\Omega(t)\rangle$ will be the ground state of $H(t)$ for all t .

Adiabatic state preparation exploits the adiabatic theorem by beginning with a well-understood Hamiltonian $H(0)$, and slowly deforming along some trajectory to the desired Hamiltonian $H(T)$. As long as the evolution is performed slowly enough, the ground state is prepared with high probability. The computational cost of this method is dominated by the smallest gap along the trajectory: we must spend time proportional to Δ^{-2} in preparing the state.

It is useful to note that, although phrased as a method for finding the ground state of a Hamiltonian, the adiabatic state preparation method allows access to ground states of restricted sectors of the Hamiltonian as well. As an example, as long as all Hamiltonians along the adiabatic trajectory are translation-invariant, the total momentum will commute with each time-evolution operator e^{-iHt} . If the initially-prepared state has some non-vanishing momentum, then the adiabatically prepared state will preserve that momentum. The same holds for any conserved quantum

⁷This notion is well-defined as long as $\Delta > 0$, which is of course necessary for the theorem to make any statement at all.

numbers, as long as the conservation law holds at each point on the trajectory.

We now consider the application of adiabatic state preparation [109] to the proton state $|P\rangle$. We chose the adiabatic trajectory to begin with the zero-coupling theory, and the coupling is slowly increased until the desired lattice spacing is reached. The ground state of the baryon-number-1 sector of the free theory is three zero-momentum fermions in a box, with the gauge fields in a Gaussian configuration (for the $SU(3)$ theory) or pegged to the identity (for the Valentiner gauge theory); these states are easily prepared. At the end of the trajectory, in the physical regime, the gap is equal to the pion mass m_π and is relatively large. The other end is more problematic. At vanishing coupling, the ‘proton’ fills the lattice, and the lowest-lying excited states are the back-to-back low-momentum states of any two fermions. (Massless glue excitations, which do not exist at all in the Valentiner theory, can be removed with appropriate boundary conditions). The gap in this regime is proportional to L^{-1} , implying that L^2 time steps are needed. Each evolution step requires $O(V)$ operations, so the effort required to prepare the proton is $O(L^5)$.

5.5.3 Hadronic Tensor

The hadronic tensor characterizes the response of a hadron to a perturbative probe. In the context of an electromagnetic probe, it is defined by

$$W^{\mu\nu}(q) = \text{Re} \int d^3x e^{iqx} \langle P | T \{ J^\mu(x) J^\nu(0) \} | P \rangle, \quad (5.32)$$

where $J^\mu = \bar{\psi}\gamma^\mu\psi$ is the current associated to electromagnetic charge. As an example of the use of this object, the inclusive cross section of an electron scattering off of a hadron is given at leading order by [110]

$$\frac{d^2\sigma}{dx dy} = \frac{\alpha^2 y}{Q^4} L_{\mu\nu} W^{\mu\nu} \quad (5.33)$$

where the leptonic tensor $L_{\mu\nu}$, calculable in QED perturbation theory, is:

$$2(k_\mu k'_\nu + k_\nu k'_\mu - g_{\mu\nu} k \cdot k'). \quad (5.34)$$

In both equations above, $Q^2 = -q^2$, $x = Q^2/2P \cdot q$, $y = P \cdot q/P \cdot k$, and $k' = k - q$. Note that the hadronic tensor itself captures the nonperturbative physics of the proton; the perturbative expansion in the QED coupling appears only because the interaction between the proton and electron is treated perturbatively.

The hadronic tensor (and its various limits, particularly the parton distribution function discussed below) has been the target of extensive Euclidean lattice calculations [111–127]. Ultimately, all approaches suffer from difficulties stemming from the need to pass from a calculation performed in imaginary time to a quantity defined in real Minkowski time. However, the fact that the hadronic tensor is defined in Minkowski time makes it particularly amenable to computation with a quantum computer [109]. Indeed, the method of Section 5.3.1 may be directly

applied to compute the Hadronic tensor, using the perturbed Hamiltonian

$$H_x(t) = H_0 + \epsilon_x J^\mu(x) \tag{5.35}$$

and measuring the observable $J_\mu(x)$ after some time evolution.

5.5.3.1 Parton Distribution Function

The response of a hadron to a probe with large momentum transfer is characterized by an apparently simpler object, the parton distribution function (PDF):

$$f(x) = \int dy; e^{ixP^+y} \langle P | \bar{\psi}(y) \gamma^+ W \psi(0) | P \rangle \tag{5.36}$$

where $u^+ = \frac{1}{\sqrt{2}}(u^0 + u^1)$ denotes the lightcone component of a vector u , and W is a lightlike Wilson line between the origin and y . The PDF can be extracted from the Hadronic tensor, but one may consider calculating it directly. A procedure for directly computing the PDF is given in [109]. Here we will only discuss why this is a bad idea.

Momentarily ignoring gauge fields, note that one of the tools of Section 5.3.1 is inapplicable to the correlation function in the integrand. Because ψ is not Hermitian, the desired correlation function cannot be expressed in terms of linear response. We must instead decompose the operator $\bar{\psi}(y) \gamma^+ \psi(0)$ as a linear combination of unitaries, each to be measured individually. This is a minor inconvenience and a major inelegance, but not fatal.

For a gauge theory, the situation is dire. The operator $\psi(0)$ is not gauge-invariant in isolation. The lightlike Wilson line involves operators at many points in space and many points in time. The resulting algorithm would require one order of finite differencing for every lattice link included in the Wilson line. This is not a practical routine even in the absence of statistical noise from the quantum computer. The hadronic tensor is protected from these issues precisely because it is a correlator of physical operators — operators which can be coupled to external sources in the Hamiltonian.

5.6 Avoiding State Preparation

In terms of number of quantum gates required (that is, the algorithmic time complexity), the study of QCD with a quantum computer is dominated by the process of preparing a suitable ground state. It may be practical, particularly for near-term quantum computers, to avoid doing so by coupling a classical Euclidean lattice calculation to a quantum computer [106, 128].

We will consider a physical system that begins in the thermal state $\rho = e^{-\beta H_0}$ of an initial Hamiltonian H_0 , and evolves for some time t under a different Hamiltonian H , at which points we measure an observable \mathcal{O} . Note that this is not limited to linear response. The expectation value desired is given by

$$\langle \mathcal{O}(t) \rangle = \frac{\sum_{i,j} \rho_{ji} \mathcal{O}(t)_{ij}}{\sum_i \rho_{ii}} = \left(\frac{\sum_{i,j} \rho_{ji} \mathcal{O}(t)_{ij}}{\sum_{i,j} \rho_{ij}} \right) \left(\frac{\sum_i \rho_{ii}}{\sum_{i,j} \rho_{ij}} \right)^{-1} \equiv \frac{\langle \mathcal{O}(t) \rangle_\rho}{\langle \delta_{ij} \rangle_\rho} \quad (5.37)$$

where ρ_{ij} and $\mathcal{O}(t)_{ij}$ denote matrix elements of the density matrix and time-evolved

operator, respectively. It is important that the basis states be cheap to prepare on the quantum processor. Eigenstates of field operators are a good choice. The notation $\langle \cdot \rangle_\rho$ denotes expectation values sampled from the distribution ρ_{ij} .

The normalization $\langle \delta_{ij} \rangle_\rho$ can be disregarded if we restrict ourselves to looking at ratios of expectation values. (It also happens that it can be efficiently computed; see the appendix of [128].) The distribution ρ_{ij} may be efficiently sampled by the standard methods of Euclidean lattice field theory, with one important difference. Ordinarily, we are evaluating $\langle \mathcal{O} \rangle$ for some observable \mathcal{O} which is diagonal in the fiducial basis used for the lattice calculation. This means that we can disregard all off-diagonal elements of the density matrix and sample only along the diagonal; this is how the periodic boundary conditions of the Euclidean lattice come into being. The operator $\mathcal{O}(t)$, however, does not vanish off the diagonal, and so we must sample the full density matrix.

By treating the operator $\mathcal{O}(t)$ as an indivisible entity, we have avoided introducing the sign problem associated with the lattice Schwinger-Keldysh method. However, the classical computer has no way of accessing the matrix elements of $\mathcal{O}(t)$. Fortunately, this is precisely the task for which a quantum computer is best suited. We have assumed that the basis states $|\Psi_i\rangle$ are easily prepared on the quantum processor; this implies that we can also prepare the states

$$|+_{ij}\rangle = \frac{1}{\sqrt{2}}(|\Psi_i\rangle + |\Psi_j\rangle) \quad \text{and} \quad |-_{ij}\rangle = \frac{1}{\sqrt{2}}(|\Psi_i\rangle - |\Psi_j\rangle) \quad (5.38)$$

with the aid of an ancillary qubit. The expectation value of $\mathcal{O}(t)$ in each of these

states is readily measured, and the desired matrix element is given by

$$\frac{1}{2} [\mathcal{O}(t)_{ij} + \mathcal{O}(t)_{ji}] = \mathcal{O}(t)_{++} - \mathcal{O}(t)_{--}. \quad (5.39)$$

Because ρ is Hermitian, this is the only linear combination needed.

This method encounters a signal-to-noise problem, which may be alleviated as detailed in [128]. Whether these techniques are sufficient to make this method practical is a matter for further empirical study, which awaits the creation of intermediate-scale quantum computers on which they can be tested.

Appendix A: Grassmann Numbers

Grassmann numbers are anticommuting objects: two Grassmann numbers η_1 and η_2 obey $\eta_1\eta_2 = -\eta_2\eta_1$. The square of a Grassmann number vanishes: $\eta_1^2 = 0$. The Grassmann algebra on N Grassmann numbers is constructed by considering complex linear combinations of Grassmann numbers. An object in the Grassmann algebra of N Grassmann numbers has 2^N complex coefficients, one for each possible combination of Grassmann numbers. Addition is performed as it would be for a vector in \mathbb{C}^{2^N} , and the multiplication rule is fixed by the fact that it distributes over addition, the anticommutativity of Grassmann numbers, and the fact that Grassmann numbers commute with complex numbers.

A.1 Integration

The Berezin integral [129] is defined by the rule

$$\int d\eta(a + \eta b) = b, \tag{A.1}$$

where a and b are elements of the Grassmann algebra that do not contain η itself. Note that this integral is only a formal manipulation — there is no sense in

which it can be approximated by a limit of finite sums, as in Riemann integration. Integrations over Grassmann numbers, like the numbers themselves, anticommute.

The only form of the Berezin integral that will be relevant for physical applications discussed here is when all Grassmann variables are integrated over together. This is a linear map from the Grassmann algebra to the complex numbers, defined by taking the coefficient of the term containing all N Grassmann numbers

A.2 Coherent States

The fermionic path integral is derived through the use of coherent states, defined as

$$|\psi\rangle = e^{\psi^\dagger c^\dagger} |0\rangle \quad \text{and} \quad \langle\psi| = \langle 0| e^{c\psi}, \quad (\text{A.2})$$

where c (c^\dagger) is the annihilation (creation) operator acting on a fermionic mode, and the state $|0\rangle$ is the state in which that mode is unoccupied. This coherent state can be used to construct the identity operator on the Hilbert space of the theory:

$$I = \int d\psi^\dagger d\psi e^{\psi^\dagger \psi} |\psi\rangle \langle\psi|. \quad (\text{A.3})$$

The derivation of the fermionic path integral proceeds, at this point, in the usual way, with the insertion of many copies of this identity operator into the expression $\text{Tr} e^{-\beta H}$ for the partition function.

Note that this so-called “coherent state” $|\psi\rangle$ does not sit in the Hilbert space \mathbb{C}^2 of the theory at all: it is not equal to any complex linear combination of $|0\rangle$ and $|1\rangle$.

Formally, we have allowed coefficients to be Grassmann-valued, thus extending the Hilbert space to a larger module. Crucially, the expression (A.3) is only the identity operator when acting on the original Hilbert space; it annihilates all objects in the module not part of that original space.

Bibliography

- [1] Thomas DeGrand and Ethan Neil. Composite phenomenology as a target for lattice qcd. *arXiv preprint arXiv:1809.03473*, 2018.
- [2] Michael R Douglas. Report on the status of the yang-mills millenium prize problem. 2004.
- [3] Christof Gattringer and Christian Lang. *Quantum chromodynamics on the lattice: an introductory presentation*, volume 788. Springer Science & Business Media, 2009.
- [4] István Montvay and Gernot Münster. *Quantum fields on a lattice*. Cambridge University Press, 1997.
- [5] C Monroe, R Raussendorf, A Ruthven, KR Brown, P Maunz, L-M Duan, and J Kim. Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Physical Review A*, 89(2):022317, 2014.
- [6] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.
- [7] Seth Lloyd. Universal quantum simulators. *Science*, pages 1073–1078, 1996.
- [8] Richard P Feynman. Simulating physics with computers. *Int. J. Theor. Phys*, 21(6/7), 1999.
- [9] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. Quantum Computation of Scattering in Scalar Quantum Field Theories. 2011. [Quant. Inf. Comput.14,1014(2014)].
- [10] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to quantum field theory*. Addison-Wesley, Reading, USA, 1995.
- [11] Ulli Wolff. Precision check on triviality of ϕ^4 theory by a new simulation method. *Phys. Rev.*, D79:105002, 2009.

- [12] John Hubbard. Electron correlations in narrow energy bands. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 276(1365):238–257, 1963.
- [13] DJ Scalapino. Numerical studies of the 2d hubbard model. In *Handbook of High-Temperature Superconductivity*, pages 495–526. Springer, 2007.
- [14] Holger Bech Nielsen and Masao Ninomiya. No-go theorem for regularizing chiral fermions. Technical report, Science Research Council, 1981.
- [15] Masuo Suzuki. Generalized trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Communications in Mathematical Physics*, 51(2):183–190, 1976.
- [16] Hale F Trotter. On the product of semi-groups of operators. *Proceedings of the American Mathematical Society*, 10(4):545–551, 1959.
- [17] F. J. Wegner. Duality in Generalized Ising Models and Phase Transitions Without Local Order Parameters. *J. Math. Phys.*, 12:2259–2272, 1971.
- [18] John Kogut and Leonard Susskind. Hamiltonian formulation of wilson’s lattice gauge theories. *Physical Review D*, 11(2):395, 1975.
- [19] Michael Creutz. Gauge fixing, the transfer matrix, and confinement on a lattice. *Physical Review D*, 15(4):1128, 1977.
- [20] Henry Lamm, Scott Lawrence, and Yukari Yamauchi. General Methods for Digital Quantum Simulation of Gauge Theories. *Phys. Rev.*, D100(3):034518, 2019.
- [21] M. Luscher. Construction of a Selfadjoint, Strictly Positive Transfer Matrix for Euclidean Lattice Gauge Theories. *Commun. Math. Phys.*, 54:283, 1977.
- [22] V. K. Mitrjushkin. Transfer matrix and nonperturbative renormalization of fermionic currents in lattice QCD. 2002.
- [23] Michael Creutz. Transfer matrices and lattice fermions at finite density. *Found. Phys.*, 30:487–492, 2000.
- [24] Sergio Caracciolo and Fabrizio Palumbo. Transfer matrix for Kogut-Susskind fermions in the spin basis. *Phys. Rev.*, D87(1):014507, 2013.
- [25] Michael Creutz. Gauge Fixing, the Transfer Matrix, and Confinement on a Lattice. *Phys. Rev.*, D15:1128, 1977. [,132(1976)].
- [26] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21:1087–1092, 1953.

- [27] James M Lattimer. The nuclear equation of state and neutron star masses. *Annual Review of Nuclear and Particle Science*, 62:485–515, 2012.
- [28] MC Miller, Frederick K Lamb, AJ Dittmann, Slavko Bogdanov, Zaven Arzoumanian, Keith C Gendreau, S Guillot, AK Harding, WCG Ho, JM Lattimer, et al. Psr j0030+ 0451 mass and radius from nicer data and implications for the properties of neutron star matter. *The Astrophysical Journal Letters*, 887(1):L24, 2019.
- [29] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations. *Phys. Rev. Lett.*, 94:170201, 2005.
- [30] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- [31] Leonid V Keldysh et al. Diagram technique for nonequilibrium processes. *Sov. Phys. JETP*, 20(4):1018–1026, 1965.
- [32] Andrei Alexandru, Gokce Basar, Paulo F. Bedaque, Sohan Vartak, and Neill C. Warrington. Monte Carlo Study of Real Time Dynamics on the Lattice. *Phys. Rev. Lett.*, 117(8):081602, 2016.
- [33] Andrei Alexandru, Gokce Basar, Paulo F. Bedaque, and Gregory W. Ridgway. Schwinger-Keldysh formalism on the lattice: A faster algorithm and its application to field theory. *Phys. Rev. D*, 95(11):114501, 2017.
- [34] Stephen P Jordan, Hari Krovi, Keith SM Lee, and John Preskill. Bqp-completeness of scattering in scalar quantum field theory. *Quantum*, 2:44, 2018.
- [35] Michael J Bremner, Richard Jozsa, and Dan J Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2011.
- [36] Scott Aaronson, Adam Bouland, Greg Kuperberg, and Saeed Mehraban. The computational complexity of ball permutations. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 317–327, 2017.
- [37] N Yu Astrakhantsev, VV Braguta, and A Yu Kotov. Temperature dependence of shear viscosity of su (3)-gluodynamics within lattice simulation. *Journal of High Energy Physics*, 2017(4):101, 2017.
- [38] G. Parisi. The Strategy for Computing the Hadronic Mass Spectrum. *Phys. Rept.*, 103:203–211, 1984.

- [39] G. Peter Lepage. The Analysis of Algorithms for Lattice Field Theory. In *Boulder ASI 1989:97-120*, pages 97–120, 1989.
- [40] Abhijit Majumder and Berndt Müller. Hadron mass spectrum from lattice qcd. *Physical review letters*, 105(25):252002, 2010.
- [41] William Detmold, Gurtej Kanwar, Michael L. Wagman, and Neill C. Warrington. Path integral contour deformations for noisy observables. 3 2020.
- [42] Wilhelm Wirtinger. Zur formalen theorie der funktionen von mehr komplexen veränderlichen. *Mathematische Annalen*, 97(1):357–375, 1927.
- [43] Andrei Alexandru, Gökçe Başar, Paulo F. Bedaque, Henry Lamm, and Scott Lawrence. Finite Density QED_{1+1} Near Lefschetz Thimbles. *Phys. Rev.*, D98(3):034506, 2018.
- [44] Edward Witten. Analytic continuation of chern-simons theory. 2011.
- [45] Marco Cristoforetti, Francesco Di Renzo, and Luigi Scorzato. New approach to the sign problem in quantum field theories: High density QCD on a Lefschetz thimble. *Phys. Rev.*, D86:074506, 2012.
- [46] Marco Cristoforetti, Francesco Di Renzo, Abhishek Mukherjee, and Luigi Scorzato. Monte carlo simulations on the lefschetz thimble: Taming the sign problem. *Physical Review D*, 88(5):051501, 2013.
- [47] Yuya Tanizaki, Yoshimasa Hidaka, and Tomoya Hayata. Lefschetz-thimble analysis of the sign problem in one-site fermion model. *New Journal of Physics*, 18(3):033002, 2016.
- [48] Scott Lawrence. Beyond Thimbles: Sign-Optimized Manifolds for Finite Density. *PoS, LATTICE2018*:149, 2018.
- [49] Andrei Alexandru, Gökçe Basar, and Paulo Bedaque. Monte Carlo algorithm for simulating fermions on Lefschetz thimbles. *Phys. Rev. D*, 93(1):014504, 2016.
- [50] Andrei Alexandru, Gokce Basar, Paulo F. Bedaque, Gregory W. Ridgway, and Neill C. Warrington. Sign problem and Monte Carlo calculations beyond Lefschetz thimbles. *JHEP*, 05:053, 2016.
- [51] Andrei Alexandru, Gokce Basar, Paulo F. Bedaque, Gregory W. Ridgway, and Neill C. Warrington. Fast estimator of Jacobians in the Monte Carlo integration on Lefschetz thimbles. *Phys. Rev. D*, 93(9):094514, 2016.
- [52] Andrei Alexandru, Paulo F. Bedaque, Henry Lamm, and Scott Lawrence. Deep Learning Beyond Lefschetz Thimbles. *Phys. Rev.*, D96(9):094505, 2017.
- [53] S. Ruder. An overview of gradient descent optimization algorithms. *ArXiv e-prints*, September 2016.

- [54] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *ArXiv e-prints*, December 2014.
- [55] Andrei Alexandru, Paulo F. Bedaque, Henry Lamm, and Scott Lawrence. Finite-Density Monte Carlo Calculations on Sign-Optimized Manifolds. *Phys. Rev.*, D97(9):094510, 2018.
- [56] Yuto Mori, Kouji Kashiwa, and Akira Ohnishi. Application of a neural network to the sign problem via the path optimization method. *Progress of Theoretical and Experimental Physics*, 2018(2):023B04, 2018.
- [57] Walter E Thirring. A soluble relativistic field theory. *Annals of Physics*, 3(1):91–112, 1958.
- [58] Kenneth G. Wilson. Confinement of Quarks. pages 45–59, 2 1974.
- [59] Andrei Alexandru, Gokce Basar, Paulo F. Bedaque, and Neill C. Warrington. Tempered transitions between thimbles. *Phys. Rev. D*, 96(3):034513, 2017.
- [60] Andrei Alexandru, Gokce Basar, Paulo F. Bedaque, Gregory W. Ridgway, and Neill C. Warrington. Monte Carlo calculations of the finite density Thirring model. *Phys. Rev. D*, 95(1):014502, 2017.
- [61] Hirotugu Fujii, Syo Kamata, and Yoshio Kikukawa. Monte Carlo study of Lefschetz thimble structure in one-dimensional Thirring model at finite density. *JHEP*, 12:125, 2015. [Erratum: *JHEP* 09, 172 (2016)].
- [62] Hirotugu Fujii, Syo Kamata, and Yoshio Kikukawa. Lefschetz thimble structure in one-dimensional lattice Thirring model at finite density. *JHEP*, 11:078, 2015. [Erratum: *JHEP* 02, 036 (2016)].
- [63] Francesco Di Renzo, Simran Singh, and Kevin Zambello. One-thimble regularisation of lattice field theories: is it only a dream? In *37th International Symposium on Lattice Field Theory*, 2019.
- [64] Andrei Alexandru, Paulo F. Bedaque, Henry Lamm, Scott Lawrence, and Neill C. Warrington. Fermions at Finite Density in 2+1 Dimensions with Sign-Optimized Manifolds. *Phys. Rev. Lett.*, 121(19):191602, 2018.
- [65] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.
- [66] A. Yu Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191–1249, Dec 1997.
- [67] R Solovay. Lie groups and quantum circuits. *preprint*, 1999.
- [68] Christopher M. Dawson and Michael A. Nielsen. The solovay-kitaev algorithm, 2005.

- [69] David Elieser Deutsch, Adriano Barenco, and Artur Ekert. Universality in quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 449(1937):669–677, 1995.
- [70] Michael A Nielsen and Isaac L Chuang. Quantum computation and quantum information.
- [71] Tommaso Toffoli. Reversible computing. In *Proceedings of the 7th Colloquium on Automata, Languages and Programming*, pages 632–644, Berlin, Heidelberg, 1980. Springer-Verlag.
- [72] Mehdi Saeedi and Massoud Pedram. Linear-depth quantum circuits for n-qubit toffoli gates with no ancilla. *Physical Review A*, 87(6):062318, 2013.
- [73] Adriano Barenco, Charles H Bennett, Richard Cleve, David P DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Physical review A*, 52(5):3457, 1995.
- [74] Vivek V Shende and Igor L Markov. On the cnot-cost of toffoli gates. *arXiv preprint arXiv:0803.2316*, 2008.
- [75] Markus Püschel, Martin Rötteler, and Thomas Beth. Fast quantum fourier transforms for a class of non-abelian groups. In *International Symposium on Applied Algebra, Algebraic Algorithms, and Error-Correcting Codes*, pages 148–159. Springer, 1999.
- [76] Don Coppersmith. An approximate fourier transform useful in quantum factoring. *arXiv preprint quant-ph/0201067*, 2002.
- [77] Peter W Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th annual symposium on foundations of computer science*, pages 124–134. Ieee, 1994.
- [78] Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, 1996.
- [79] Andrei Alexandru, Paulo F. Bedaque, and Scott Lawrence. Quantum Algorithms For Disordered Physics. 2019.
- [80] P. W. Anderson. Absence of Diffusion in Certain Random Lattices. *Phys. Rev.*, 109:1492–1505, 1958.
- [81] Elihu Abrahams. *50 Years of Anderson Localization*. WORLD SCIENTIFIC, 2010.
- [82] Ferdinand Evers and Alexander D. Mirlin. Anderson transitions. *Rev. Mod. Phys.*, 80:1355–1417, Oct 2008.

- [83] Tobias Brandes and Stefan Kettmann. *Anderson localization and its ramifications: Disorder, phase coherence, and electron correlations*, volume 630. Springer Science & Business Media, 2003.
- [84] U.S. Department of Commerce and National Institute of Standards and Technology. *Secure Hash Standard - SHS: Federal Information Processing Standards Publication 180-4*. CreateSpace Independent Publishing Platform, USA, 2012.
- [85] John K Salmon, Mark A Moraes, Ron O Dror, and David E Shaw. Parallel random numbers: as easy as 1, 2, 3. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, page 16. ACM, 2011.
- [86] George Marsaglia. The marsaglia random number cdrom including the diehard battery of tests of randomness. <http://www.stat.fsu.edu/pub/diehard/>, 2008.
- [87] Robert G Brown. <https://webhome.phy.duke.edu/~rgb/General/dieharder.php>, 2017.
- [88] Gadi Aleksandrowicz, Thomas Alexander, Panagiotis Barkoutsos, Luciano Bello, Yael Ben-Haim, D Bucher, FJ Cabrera-Hernández, J Carballo-Franquis, A Chen, CF Chen, et al. Qiskit: An open-source framework for quantum computing. *Accessed on: Mar, 16, 2019*.
- [89] Robert S Smith, Michael J Curtis, and William J Zeng. A practical quantum instruction set architecture. *arXiv preprint arXiv:1608.03355*, 2016.
- [90] Andrei Alexandru, Paulo F. Bedaque, Henry Lamm, and Scott Lawrence. σ Models on Quantum Computers. *Phys. Rev. Lett.*, 123(9):090501, 2019.
- [91] J. S. Pedernales, R. Di Candia, I. L. Egusquiza, J. Casanova, and E. Solano. Efficient quantum algorithm for computing n -time correlation functions. *Phys. Rev. Lett.*, 113:020505, Jul 2014.
- [92] G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme. Quantum algorithms for fermionic simulations. *Phys. Rev.*, A64:022319, 2001.
- [93] Richard Cleve, Artur Ekert, Chiara Macchiavello, and Michele Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1969):339–354, 1998.
- [94] Pascual Jordan and Eugene P. Wigner. About the Pauli exclusion principle. *Z. Phys.*, 47:631–651, 1928.
- [95] Sergey B Bravyi and Alexei Yu Kitaev. Fermionic quantum computation. *Annals of Physics*, 298(1):210–226, 2002.

- [96] James D Whitfield, Vojtěch Havlíček, and Matthias Troyer. Local spin operators for fermion simulations. *Physical Review A*, 94(3):030301, 2016.
- [97] Frank Verstraete and J Ignacio Cirac. Mapping local hamiltonians of fermions to local hamiltonians of spins. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(09):P09012, 2005.
- [98] Hermann Valentiner. *De endelige transformations-gruppers teori: avec un résumé en français*, volume 2. Bianco Lunos, 1889.
- [99] Henrik Flyvbjerg. Character table for the 1080-element point-group-like subgroup of $su(3)$. *Journal of mathematical physics*, 26(11):2985–2989, 1985.
- [100] Andrei Alexandru, Paulo F. Bedaque, Siddhartha Harmalkar, Henry Lamm, Scott Lawrence, and Neill C. Warrington. Gluon Field Digitization for Quantum Computers. *Phys. Rev.*, D100(11):114501, 2019.
- [101] G. Bhanot and C. Rebbi. Monte Carlo Simulations of Lattice Models With Finite Subgroups of $SU(3)$ as Gauge Groups. *Phys. Rev.*, D24:3319, 1981.
- [102] Martin Luscher. Properties and uses of the Wilson flow in lattice QCD. *JHEP*, 08:071, 2010. [Erratum: JHEP03,092(2014)].
- [103] A. Francis, O. Kaczmarek, M. Laine, T. Neuhaus, and H. Ohno. Critical point and scale setting in $SU(3)$ plasma: An update. *Phys. Rev.*, D91(9):096002, 2015.
- [104] Masakiyo Kitazawa, Takumi Iritani, Masayuki Asakawa, Tetsuo Hatsuda, and Hiroshi Suzuki. Equation of State for $SU(3)$ Gauge Theory via the Energy-Momentum Tensor under Gradient Flow. *Phys. Rev.*, D94(11):114512, 2016.
- [105] David B. Kaplan, Natalie Klco, and Alessandro Roggero. Ground States via Spectral Combing on a Quantum Computer. 2017.
- [106] Henry Lamm and Scott Lawrence. Simulation of Nonequilibrium Dynamics on a Quantum Computer. *Phys. Rev. Lett.*, 121(17):170501, 2018.
- [107] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. Quantum Algorithms for Fermionic Quantum Field Theories. 2014.
- [108] Albert Messiah. *Quantum mechanics, volume II*. North-Holland Publishing Company, 1969.
- [109] Henry Lamm, Scott Lawrence, and Yukari Yamauchi. Parton Physics on a Quantum Computer. 2019.
- [110] M. Tanabashi et al. Review of particle physics. *Phys. Rev. D*, 98:030001, Aug 2018.

- [111] Keh-Fei Liu and Shao-Jing Dong. Origin of difference between anti-d and anti-u partons in the nucleon. *Phys. Rev. Lett.*, 72:1790–1793, 1994.
- [112] U. Aglietti, Marco Ciuchini, G. Corbo, E. Franco, G. Martinelli, and L. Silvestrini. Model independent determination of the shape function for inclusive B decays and of the structure functions in DIS. *Phys. Lett.*, B432:411–420, 1998.
- [113] Keh-Fei Liu. Parton degrees of freedom from the path integral formalism. *Phys. Rev.*, D62:074501, 2000.
- [114] Xiangdong Ji. Parton Physics on a Euclidean Lattice. *Phys. Rev. Lett.*, 110:262002, 2013.
- [115] A. J. Chambers, R. Horsley, Y. Nakamura, H. Perlt, P. E. L. Rakow, G. Schierholz, A. Schiller, K. Somfleth, R. D. Young, and J. M. Zanotti. Nucleon Structure Functions from Operator Product Expansion on the Lattice. *Phys. Rev. Lett.*, 118(24):242001, 2017.
- [116] Gunnar S. Bali, Vladimir M. Braun, Benjamin Gläbke, Meinulf Göckeler, Michael Gruber, Fabian Hutzler, Piotr Korcyl, Andreas Schäfer, Philipp Wein, and Jian-Hui Zhang. Pion distribution amplitude from Euclidean correlation functions: Exploring universality and higher-twist effects. *Phys. Rev.*, D98(9):094507, 2018.
- [117] A. V. Radyushkin. Quasi-parton distribution functions, momentum distributions, and pseudo-parton distribution functions. *Phys. Rev.*, D96(3):034025, 2017.
- [118] Kostas Orginos, Anatoly Radyushkin, Joseph Karpie, and Savvas Zafeiropoulos. Lattice QCD exploration of parton pseudo-distribution functions. *Phys. Rev.*, D96(9):094503, 2017.
- [119] Joseph Karpie, Kostas Orginos, and Savvas Zafeiropoulos. Moments of Ioffe time parton distribution functions from non-local matrix elements. *JHEP*, 11:178, 2018.
- [120] Joseph Karpie, Kostas Orginos, Alexander Rothkopf, and Savvas Zafeiropoulos. Reconstructing parton distribution functions from Ioffe time data: from Bayesian methods to Neural Networks. *JHEP*, 04:057, 2019.
- [121] Krzysztof Cichy, Luigi Del Debbio, and Tommaso Giani. Parton distributions from lattice data: the nonsinglet case. 2019.
- [122] Constantia Alexandrou, Krzysztof Cichy, Martha Constantinou, Kyriakos Hadjiyiannakou, Karl Jansen, Aurora Scapellato, and Fernanda Steffens. Systematic uncertainties in parton distribution functions from lattice QCD simulations at the physical point. *Phys. Rev.*, D99(11):114504, 2019.

- [123] Raza Sabbir Sufian, Joseph Karpie, Colin Egerer, Kostas Orginos, Jian-Wei Qiu, and David G. Richards. Pion Valence Quark Distribution from Matrix Element Calculated in Lattice QCD. *Phys. Rev.*, D99(7):074507, 2019.
- [124] Yan-Qing Ma and Jian-Wei Qiu. Extracting Parton Distribution Functions from Lattice QCD Calculations. *Phys. Rev.*, D98(7):074021, 2018.
- [125] Yan-Qing Ma and Jian-Wei Qiu. QCD Factorization and PDFs from Lattice QCD Calculation. *Int. J. Mod. Phys. Conf. Ser.*, 37:1560041, 2015.
- [126] Yan-Qing Ma and Jian-Wei Qiu. Exploring Partonic Structure of Hadrons Using ab initio Lattice QCD Calculations. *Phys. Rev. Lett.*, 120(2):022003, 2018.
- [127] Jian Liang, Terrence Draper, Keh-Fei Liu, Alexander Rothkopf, and Yi-Bo Yang. Towards the nucleon hadronic tensor from lattice QCD. 2019.
- [128] Siddhartha Harmalkar, Henry Lamm, and Scott Lawrence. Quantum simulation of field theories without state preparation.
- [129] F. A. Berezin. The method of second quantization. *Pure Appl. Phys.*, 24:1–228, 1966.