Title

Towards Precision Quantum Simulation of Lattice Gauge Theories

ISBN

9798189275821

Author

Hariprakash, Siddharth

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Towards Precision Quantum Simulation of Lattice Gauge Theories

by

Siddharth Hariprakash

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Dr. Christian Bauer, Co-chair
Dr. Raphael Bousso, Co-chair
Dr. Luca Iliesiu
Dr. Lin Lin

Spring 2026

Towards Precision Quantum Simulation of Lattice Gauge Theories

Abstract

Towards Precision Quantum Simulation of Lattice Gauge Theories

by

Siddharth Hariprakash

Doctor of Philosophy in Physics

University of California, Berkeley

Dr. Christian Bauer, Co-chair

Dr. Raphael Bousso, Co-chair

Quantum simulation of lattice gauge theories offers a path to first-principles calculations of real-time dynamics and finite-density phenomena that are inaccessible to classical Euclidean Monte Carlo methods. Realizing this potential requires a pipeline of controlled approximations—lattice discretization, truncation, digitization, state preparation, time evolution, measurement, and continuum extrapolation—each contributing errors and computational costs. General algorithmic asymptotic complexity arguments capture only part of the cost picture: the practical resources required for end-to-end simulation depend on numerical pre-factors, the cost of the subroutines that various algorithms call, and structural features of the lattice theory, all of which the asymptotic scaling leaves unspecified. This thesis develops contributions that aim to make such dependencies explicit and quantitative.

We begin with a rigorous analysis of the asymptotic gate complexity of product formulas and Quantum Signal Processing (QSP) for the time evolution of general classes of lattice Hamiltonians, identifying the parameter regimes in which each method is favorable and estimating the associated pre-factors numerically. Since the cost of QSP-based simulation is dominated (in regimes of practical interest) by the block-encoding subroutine it calls, we also develop several novel block-encoding constructions tailored to bosonic lattice field theories and show that signal-processing techniques can be repurposed as highly efficient block-encoding subroutines.

We then turn to the problem of truncating the Kogut–Susskind electric-field Hilbert space. Existing error bounds rely on energy-conservation arguments and require the truncation to grow with simulation time. By exploiting Hilbert space fragmentation we derive the tightest error bounds known to date: they decay factorially in the truncation parameter and are independent of simulation time. We validate them numerically for U(1) pure gauge theory and the Schwinger model.

Finally, we address the continuum limit in the presence of approximate time evolution. For product formulas, we show that Trotter errors correspond to irrelevant operators that vanish in the continuum limit, which yields a renormalization trajectory independent of the Trotter step size. More generally, we introduce the Statistically-Bounded Time Evolution (SBTE) protocol, an algorithm-independent framework that treats approximation errors as a systematic uncertainty to be driven below the working statistical precision. This decouples the continuum-limit prescription from the choice of time-evolution algorithm and provides a basis for end-to-end cost comparisons across simulation methods.

To my parents

Contents

List of Figures

List of Tables

Acknowledgments

First, I would like to thank my advisors. I would like to thank Christian Bauer for leading my growth as a researcher, helping shape my approach to tackling hard problems, and all the invaluable advice and mentorship over the years. I also thank Raphael Bousso for the many valuable discussions during my time as a graduate student.

A huge thanks to Luca Iliesiu and Lin Lin for their help as part of my dissertation committee and also Hartmut Häffner for being a part of my qualifying exam committee.

It has been a privilege to learn from my collaborators throughout the years, including Michael Kreshchuk, Christopher Kane, Anthony Ciaveralla, Neel Modi, Jad Halimeh, Daan Camps, Katherine Klymko, Roel Van Beeumen, and others. I would like to thank the members of the HEP-QIS and LITP groups at LBNL and UC Berkeley for many stimulating conversations. I would also like to acknowledge the many funding agencies that have supported my work throughout the years.

Thanks to Jay Hauser and members of the UCLA CMS group for giving me my first opportunity as an undergraduate and introducing me to Physics research.

I would also like to thank all my close friends and family. An exhaustive list would take up way too much space. This thesis is a result of all their love and support. I would like to thank my mom for taking me to the park as a kid and for making me do 5th grade math worksheets as a 2nd grader. I thank my dad for printing out said worksheets at work and for taking me to Costco on the weekends. I would also like to thank my sister and congratulate her on all her own success as she continues on her own journey. A special thanks to MK for valuable discussions at the Bergen public library.

I acknowledge that the abstract and introduction in this thesis was written after significant interaction with CLAUDE

Chapter 1

Introduction

1.1 Gauge Theories and the Need for Ab-Initio Simulation

The fundamental interactions of nature (with the exception of gravity) are described by gauge field theories: quantum field theories whose Lagrangians are invariant under local transformations belonging to a symmetry group G. The Standard Model of particle physics is built on the gauge group $\mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$, where $\mathrm{SU}(3)_c$ governs the strong interaction through quantum chromodynamics (QCD), and $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ governs the electroweak interaction. The principle of local gauge invariance dictates the form of the interactions: it requires the introduction of gauge bosons—the gluons of QCD, and (after electroweak symmetry breaking) the $W^{\pm}$, Z, and the photon—which mediate forces between matter fields and constrain the allowed couplings among them. The remarkable success of this framework is evidenced by its agreement with experiment across an enormous range of energy scales.

1.1.1 Perturbation Theory and its Limitations

The standard analytical tool for extracting predictions from gauge theories is perturbation theory: one expands observables as a power series in the coupling constant and evaluates the resulting Feynman diagrams order by order. For quantum electrodynamics (QED), the U(1) gauge theory of the electromagnetic interaction, the fine-structure constant $\alpha \approx 1/137$ is small at experimentally accessible energies, and low-order perturbative calculations have achieved extraordinary agreement with experiment. The anomalous magnetic moment of the electron, for instance, has been computed and measured to better than one part in 10^{10}.

However, perturbation theory faces intrinsic limitations even in weakly coupled theories. The number of Feynman diagrams grows factorially with the loop order: the tenth-order (α^5) QED contribution to the electron anomalous magnetic moment requires the evaluation of 12,672 multi-loop diagrams and is one of the most computationally demanding multi-

loop calculations performed to date [1]. More fundamentally, the perturbative series in quantum field theory is generically asymptotic rather than convergent [2]: if the coupling were continued to negative values the theory would become unstable, suggesting that the radius of convergence of the perturbative expansion is zero. The series therefore fails to converge for any nonzero coupling, and its reliability rests on the coupling being small enough that an optimal truncation provides adequate precision.

Most importantly, perturbation theory is entirely blind to phenomena that are non-perturbative in nature: effects whose dependence on the coupling constant is of the form e^{-c/g^2} vanish to all orders in the perturbative expansion and cannot be captured at any finite order. Such effects include tunneling between topological sectors (instantons) and *dimensional transmutation*: the dynamical generation of a renormalization-group-invariant mass scale such as $\Lambda_{\mathrm{QCD}} \sim \mu \exp(-1/(2b_0 g^2(\mu)))$. Dimensional transmutation in turn drives confinement and *dynamical* chiral symmetry breaking in QCD, both governed by strong-coupling dynamics rather than by a fundamental scalar field.

1.1.2 Non-perturbative Phenomena in QCD

Due to asymptotic freedom, the strong coupling constant $\alpha_s(\mu)$ decreases logarithmically at high momentum scales μ, and perturbative QCD successfully describes hard scattering processes such as jet production at colliders and deep inelastic scattering. At low energies, however, α_s grows to $O(1)$ and the theory becomes strongly coupled, rendering perturbation theory inapplicable. The rich phenomenology of the strong interaction at low energies is governed by non-perturbative effects:

- *Confinement*: quarks and gluons are bound into color-neutral hadrons, and no isolated color charge has ever been observed. Confinement is supported by overwhelming numerical evidence but lacks a fully analytic proof. The static quark–antiquark potential grows linearly with separation, $V(r) \sim \sigma r$, where σ is the string tension.

- *Chiral symmetry breaking*: the approximate $\mathrm{SU}(N_f)_L \times \mathrm{SU}(N_f)_R$ chiral symmetry of the light quark sector is spontaneously broken to the diagonal $\mathrm{SU}(N_f)_V$, generating the pions as pseudo-Goldstone bosons and producing the bulk of the visible mass in the universe. The Higgs mechanism contributes only a few percent of the proton mass; the remainder arises from the non-perturbative dynamics of QCD.

- *The hadron spectrum*: the masses and quantum numbers of the hundreds of observed mesons, baryons, and exotic states (glueballs [3, 4, 5], tetraquarks, pentaquarks) are determined entirely by the non-perturbative dynamics of QCD.

- *Phase transitions*: at high temperatures, QCD undergoes a crossover to a deconfined quark-gluon plasma; at finite baryon density a rich and largely unexplored phase structure is expected, including a conjectured critical point and color-superconducting phases.

None of these phenomena can be computed from perturbation theory. They demand ab initio computational methods that work directly with the fundamental QCD Lagrangian without uncontrolled approximations—methods that can, in principle, be systematically improved to arbitrary precision.

1.1.3 Beyond QCD: the Broad Reach of Gauge Theories

The need for non-perturbative simulation extends well beyond QCD. Gauge theories with various gauge groups appear across physics, and in many cases their most interesting dynamics lies precisely in the non-perturbative regime. A few examples include

- *Grand unification*: theories based on groups such as SU(5) or SO(10) [6, 7] typically remain asymptotically free at high scales and are broken by the Higgs mechanism near the unification scale. The associated symmetry-breaking transitions can have cosmological consequences, including the formation of topological defects (monopoles, cosmic strings, domain walls) and contributions to baryogenesis, whose dynamics requires non-perturbative treatment.

- *Dark matter*: models involving confining hidden-sector gauge theories [8, 9, 10, 11] predict dark glueballs or dark baryons as dark matter candidates, whose masses, interaction cross-sections, and cosmological abundances are determined by the non-perturbative dynamics of the hidden gauge group.

- *Condensed matter*: emergent U(1) and $\mathbb{Z}_2$ gauge fields arise in the low-energy description of quantum spin liquids [12, 13, 14] and topological phases of matter [15]. The real-time dynamics of confinement and deconfinement transitions in these materials is directly probed by neutron scattering and other experimental techniques, and understanding these dynamics from first principles requires non-perturbative simulation.

- *Quantum gravity via holography*: the AdS/CFT correspondence posits a duality between certain gauge theories and theories of quantum gravity in anti-de Sitter space [16]. In the strong-coupling regime of the gauge theory—precisely the regime inaccessible to perturbation theory—the dual description is semiclassical gravity, and vice versa. Real-time simulation of strongly coupled gauge theories at finite temperature, for example, would provide direct access to the dynamics of black hole formation and thermalization in the dual gravitational theory [17]. The ability to simulate large-N gauge theories and matrix models non-perturbatively could therefore open a computational window into quantum gravity.

We highlight compact U(1) lattice gauge theory in particular, the simplest continuous gauge group, because it serves as a concrete test setting for the algorithmic developments of this thesis (Chapters 2 to 5). Despite its apparent simplicity, even this theory exhibits rich non-perturbative physics. In $2+1$ dimensions, Polyakov showed that magnetic monopole instantons generate a confining potential between static charges in the weak-coupling regime

via a dilute Coulomb gas of monopoles. Confinement was subsequently proved rigorously for all values of the coupling in the Villain-action lattice formulation [18], and a logarithmically growing quark–antiquark potential was recently established for Wilson-action lattice gauge theories whose gauge group contains a central U(1) [19]. The real-time dynamics of this theory, its finite-density phase structure when coupled to dynamical fermions, and its out-of-equilibrium behavior are all classically intractable for the reasons we discuss in Sec. 1.2.2.

The ability to simulate gauge theories from first principles—across gauge groups, energy scales, and physical contexts—is therefore of broad and pressing scientific importance.

1.2 Euclidean Lattice Gauge Theory

Among non-perturbative approaches, the lattice regularization of gauge theories, introduced in 1974 [20], provides a definition of these theories that is amenable to numerical computation. We describe in this section the Euclidean path-integral formulation, which underlies the dominant ab initio computational approach to gauge theories—Markov chain Monte Carlo—together with its fundamental limitations. The complementary Hamiltonian formulation, which forms the basis for quantum simulation, is presented in Sec. 1.3.

Quantum field theories are physically formulated in Minkowski spacetime, where the basic dynamical objects are vacuum expectation values of time-ordered products,

$$\langle 0|T\{\hat{O}(x_1)\cdots\hat{O}(x_n)\}|0\rangle \;=\; \frac{\int \mathcal{D}\phi\, O(x_1)\cdots O(x_n)\, e^{iS_M[\phi]}}{\int \mathcal{D}\phi\, e^{iS_M[\phi]}}\,, \tag{1.1}$$

where $S_M[\phi]$ is the classical action. The oscillatory weight e^{iS_M} is not directly amenable to stochastic evaluation: every field configuration contributes with unit modulus and an arbitrary phase, and naive Monte Carlo sampling of such an integrand fails by catastrophic cancellation. The standard remedy, originating with Wick, is to analytically continue to imaginary time, $t \to -i\tau$. Under this rotation, $iS_M \to -S_E$, where the Euclidean action $S_E[\phi]$ is real and bounded below for theories with a stable vacuum, so the oscillatory weight becomes the real, positive Boltzmann factor e^{-S_E}. Euclidean correlation functions

$$\langle \hat{O}(x_1)\cdots\hat{O}(x_n)\rangle_E \;=\; \frac{\int \mathcal{D}\phi\, O(x_1)\cdots O(x_n)\, e^{-S_E[\phi]}}{\int \mathcal{D}\phi\, e^{-S_E[\phi]}} \tag{1.2}$$

can then be computed by importance sampling on a classical computer. Real-time observables are formally encoded in the Euclidean correlators and can in principle be recovered by analytic continuation, though as we discuss in Sec. 1.2.2.2 this recovery is fundamentally ill-conditioned in practice. The Euclidean formulation is therefore naturally suited to *static, equilibrium* quantities—spectra, matrix elements, thermodynamic phase structure at zero chemical potential—which are determined by the long-Euclidean-time decay of correlators rather than by their analytic structure in real time. Within this domain, lattice Monte Carlo has been extraordinarily successful.

We now provide a short review of Wilson's discretization of the Euclidean path integral on a hypercubic lattice, which provides the gauge-invariant regularization underlying all subsequent lattice-gauge-theory computations.

1.2.1 Wilson's Lattice Construction

Continuous Euclidean spacetime is replaced by a hypercubic lattice $\Lambda \subset a\mathbb{Z}^4$ of spacing a, with matter fields defined on the sites $x \in \Lambda$ and gauge fields defined on the oriented links connecting nearest neighbors. The fundamental degree of freedom on the link from x to $x + a\hat{\mu}$ is the *link variable*

$$U_\mu(x) = \mathcal{P} \exp\left(ig \int_x^{x+a\hat{\mu}} A_\mu(y)\, dy^\mu \right) \in G\,, \tag{1.3}$$

the path-ordered exponential of the continuum gauge potential along the link, where $A_\mu = A_\mu^a T^a$ takes values in the Lie algebra of G (with Hermitian generators T^a normalized by $\mathrm{Tr}(T^a T^b) = \frac{1}{2}\delta^{ab}$), and g is the bare coupling. By construction, $U_\mu(x)$ is a group element of G rather than an algebra element. Under a gauge transformation parametrized by $\Omega(x) \in G$ at each site, the link variable transforms as

$$U_\mu(x) \;\longmapsto\; \Omega(x)\, U_\mu(x)\, \Omega^\dagger(x + a\hat{\mu})\,, \tag{1.4}$$

in direct analogy with the parallel transporter of differential geometry. The reverse-oriented link satisfies $U_{-\mu}(x + a\hat{\mu}) \equiv U_\mu^\dagger(x)$.

The simplest non-trivial gauge-invariant object built from link variables is the *plaquette*: the ordered product of link variables around an elementary unit square in the $\mu\nu$-plane,

$$U_{\mu\nu}(x) \;=\; U_\mu(x)\, U_\nu(x + a\hat{\mu})\, U_\mu^\dagger(x + a\hat{\nu})\, U_\nu^\dagger(x)\,. \tag{1.5}$$

Gauge invariance of $\mathrm{Tr}\, U_{\mu\nu}(x)$ is immediate from (1.4) together with the cyclic property of the trace. From the plaquette one constructs the *Wilson action*

$$S_W[U] \;=\; \frac{2N}{g^2} \sum_{x \in \Lambda} \sum_{\mu<\nu} \left[1 - \frac{1}{N} \,\mathrm{Re}\,\mathrm{Tr}\, U_{\mu\nu}(x) \right], \tag{1.6}$$

where N is the dimension of the fundamental representation of $G = \mathrm{SU}(N)$ and the sum runs over each elementary plaquette once. To verify that (1.6) is a discretization of the continuum Yang–Mills action, expand the link variable for small a as $U_\mu(x) = 1 + iga\, A_\mu(x) + O(a^2)$ and apply the Baker–Campbell–Hausdorff formula to the plaquette product, yielding

$$U_{\mu\nu}(x) \;=\; \exp\Big[iga^2 F_{\mu\nu}(x) + O(a^3) \Big], \qquad F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + ig[A_\mu, A_\nu]\,. \tag{1.7}$$

Using the tracelessness of $F_{\mu\nu}$ and $\mathrm{Tr}(F_{\mu\nu}F_{\mu\nu}) = \frac{1}{2} F^a_{\mu\nu} F^a_{\mu\nu}$, one finds

$$1 - \frac{1}{N} \,\mathrm{Re}\,\mathrm{Tr}\, U_{\mu\nu}(x) \;=\; \frac{g^2 a^4}{4N}\, F^a_{\mu\nu}(x)\, F^a_{\mu\nu}(x) + O(a^6)\,, \tag{1.8}$$

so that

$$S_W[U] \;=\; \frac{1}{4}\int d^4x\, F^a_{\mu\nu}(x)\, F^{a\,\mu\nu}(x) \;+\; O(a^2)\,, \tag{1.9}$$

recovering the Yang–Mills action up to $O(a^2)$ discretization artifacts that can be systematically reduced through the Symanzik improvement program [20, 21, 22, 23, 24].

The path-integral measure for the gauge field is the product of left- and right-invariant Haar measures on each link,

$$\mathcal{D}U \;=\; \prod_{x,\mu} dU_\mu(x)\,, \tag{1.10}$$

normalized so that $\int_G dU = 1$. Haar invariance, $\int_G dU\, f(VU) = \int_G dU\, f(UV) = \int_G dU\, f(U)$ for any $V \in G$, together with (1.4), implies that the measure (1.10) is gauge-invariant. The lattice ultraviolet cutoff at scale $1/a$ regulates the path integral non-perturbatively, with no need for gauge fixing: gauge orbits have finite volume in $G^{|\Lambda|}$ and contribute only an overall constant. The pure gauge-theory partition function is

$$Z_{\mathrm{pg}} \;=\; \int \mathcal{D}U\, e^{-S_W[U]}\,. \tag{1.11}$$

The continuum limit is recovered by taking $a \to 0$ while tuning the bare coupling $g(a)$ along a renormalization-group trajectory so that dimensionful physical observables remain fixed. For asymptotically free theories such as Yang–Mills, the trajectory has the celebrated form

$$a\Lambda_L \;=\; \exp\!\left(-\frac{1}{2b_0\, g^2(a)}\right)\left[1 + O(g^2(a))\right], \tag{1.12}$$

where $b_0 = (11N - 2N_f)/(48\pi^2)$ for $\mathrm{SU}(N)$ with N_f flavors, and Λ_L is a renormalization-group invariant mass scale—the lattice analogue of Λ_{QCD}. Thus $g(a) \to 0^+$ logarithmically as $a \to 0^+$: the continuum limit is a critical point at vanishing coupling, and the dimensionless lattice correlation length ξ/a diverges accordingly. This is the lattice realization of dimensional transmutation.

Coupling to fermions. Matter fields are placed on lattice sites: $\psi(x), \bar{\psi}(x)$ are Grassmann-valued spinors at $x \in \Lambda$, with gauge transformations $\psi(x) \mapsto \Omega(x)\,\psi(x)$. A naive nearest-neighbor discretization of the Dirac operator suffers from *fermion doubling*: in d Euclidean spacetime dimensions, the lattice propagator has 2^d low-energy poles instead of one, an obstruction made precise by the Nielsen–Ninomiya no-go theorem [25, 26]. Wilson resolved this by adding a dimension-five operator (the *Wilson term*) that gives the doublers a mass of order $1/a$, decoupling them in the continuum limit at the cost of breaking chiral symmetry explicitly. The resulting Wilson-Dirac operator is

$$M_{xy}[U] \;=\; (m + d\,r/a)\,\delta_{xy} \;-\; \frac{1}{2a}\sum_{\mu=\pm1}^{\pm d} (r - \gamma_\mu)\, U_\mu(x)\, \delta_{x+a\hat{\mu},\,y}\,, \tag{1.13}$$

where r is the Wilson parameter and $\gamma_{-\mu} \equiv -\gamma_\mu$, $U_{-\mu}(x) \equiv U_\mu^\dagger(x - a\hat{\mu})$. Several alternative formulations—staggered, domain-wall, overlap, and twisted-mass fermions—reduce the doubling problem in different ways and preserve different remnants of chiral symmetry. Integrating out the Grassmann fields,

$$\int \mathcal{D}\bar{\psi}\,\mathcal{D}\psi\, e^{-\bar{\psi}M[U]\psi} \;=\; \det M[U]\,, \tag{1.14}$$

yields the partition function for a gauge theory with N_f degenerate flavors of fermions,

$$Z \;=\; \int \mathcal{D}U \left(\det M[U]\right)^{N_f} e^{-S_W[U]}\,. \tag{1.15}$$

Markov chain Monte Carlo. The Wilson lattice action and the Haar measure are real, the gauge-field weight $e^{-S_W[U]}$ is positive, and the fermion determinant $\det M[U]$ is real at zero chemical potential by virtue of γ_5-Hermiticity of the Wilson–Dirac operator, $\gamma_5 M \gamma_5 = M^\dagger$. This is a structural property of the discretization rather than an external assumption, and it ensures that for an even number of degenerate flavors, $(\det M)^{N_f} \geq 0$, so the integrand of (1.15) can be interpreted as a probability density on the space of gauge configurations. Markov chain Monte Carlo (MCMC) algorithms generate sequences of gauge configurations distributed according to this density, and observables are estimated as sample averages [24, 27]. Over five decades, this approach has produced quantitative results of extraordinary precision and breadth, including the ab initio determination of the hadron mass spectrum, precise measurements of the QCD transition temperature, and increasingly refined determinations of quark masses and the strong coupling constant (see [28] for a recent review).

A canonical non-perturbative observable computable in this framework is the static quark–antiquark potential. The expectation value of a rectangular Wilson loop $W(R,T) = \mathrm{Tr}\prod_{(x,\mu)\in C_{R\times T}} U_\mu(x)$ along a closed contour of spatial extent R and temporal extent T in Euclidean time decays as $\langle W(R,T)\rangle \sim e^{-V(R)\,T}$ at large T, where $V(R)$ is the static potential. A confining theory exhibits the *area law*,

$$\langle W(R,T)\rangle \;\sim\; e^{-\sigma\,R\,T} \qquad (R,T\to\infty)\,, \tag{1.16}$$

with string tension $\sigma > 0$, equivalent to a linear potential $V(R) \sim \sigma R$. The numerical confirmation of the area law in SU(3) Yang–Mills theory is among the most enduring achievements of the lattice approach.

1.2.2 Limitations of Euclidean Monte Carlo

Despite these successes, Euclidean lattice Monte Carlo faces two fundamental limitations that together obstruct first-principles access to a broad class of physically important real-time and finite-density observables. The first is the *sign problem*, which prevents importance sampling whenever the path-integral weight becomes complex—arising at finite

chemical potential, in the presence of a θ-term, and most severely for direct Minkowski-time evaluation. The second is the *inverse problem*, which obstructs the extraction of real-time dynamical information from Euclidean correlators even when the weight is positive. Both limitations are intrinsic to the Euclidean path-integral formulation and motivate the search for a fundamentally different computational paradigm.

1.2.2.1 The Sign Problem

Importance sampling of (1.15) requires the integrand to be real and non-negative, a property guaranteed at zero chemical potential for even N_f by the γ_5-Hermiticity argument above. The sign problem arises whenever positivity is lost—that is, whenever the path integral weight acquires a complex phase that prevents its interpretation as a probability distribution. This problem has been shown to be NP-hard in full generality [29], providing strong evidence that no polynomial-time classical algorithm can solve it generically.

When the weight is complex, one can formally rewrite the expectation value of an observable by reweighting with respect to a positive reference ensemble:

$$\langle \hat{O} \rangle \;=\; \frac{\langle \hat{O}\, e^{i\varphi} \rangle_{\mathrm{ref}}}{\langle e^{i\varphi} \rangle_{\mathrm{ref}}}\,, \tag{1.17}$$

where $e^{i\varphi}$ is the complex phase of the weight relative to the reference measure. The denominator $\langle e^{i\varphi} \rangle_{\mathrm{ref}}$ is an average of a rapidly oscillating phase, which generically vanishes exponentially with the spacetime volume V: $\langle e^{i\varphi} \rangle_{\mathrm{ref}} \sim e^{-cV}$ for some constant $c > 0$ [30, 31]. Extracting the signal therefore requires exponentially many samples, rendering the computation intractable for large systems.

This obstruction manifests in several physically important settings.

Finite baryon density. At non-zero baryon chemical potential μ_B, the Wilson–Dirac operator loses γ_5-Hermiticity, and the fermion determinant acquires a complex phase,

$$\det M(\mu_B, U) \;=\; |\det M(\mu_B, U)|\, e^{i\varphi[U]}\,, \tag{1.18}$$

so the partition function (1.15) is no longer amenable to importance sampling. As a result, much of the QCD phase diagram, including the conjectured critical point at finite density, the transition to color-superconducting phases, and the equation of state governing the interior of neutron stars remains inaccessible to first-principles classical computation.

Topological effects. The QCD vacuum has a non-trivial topological structure parametrized by the θ-angle. Including the θ-term modifies the Euclidean action by $\Delta S_E = -i\theta\, Q_{\mathrm{top}}$, where $Q_{\mathrm{top}} = \frac{g^2}{32\pi^2} \int d^4x\, \epsilon^{\mu\nu\rho\sigma}\, \mathrm{Tr}(F_{\mu\nu} F_{\rho\sigma})$ is the (integer-valued) topological charge in the convention of Eq. (1.7). The factor of i renders the path integral weight complex. This is another instance of the sign problem, and it limits our ability to study CP violation in the strong sector from first principles [32].

Minkowski real-time evolution. The Wick rotation $t \to -i\tau$ that makes Euclidean Monte Carlo possible is also what severs direct access to real-time dynamics. Recovering the Minkowski path integral (1.1) requires either evaluating it directly with the oscillatory weight e^{iS_M}, or analytically continuing Euclidean results back to real time. Direct evaluation is the sign problem in its most severe form: the weight is a pure phase, every configuration contributes with equal magnitude but rapidly varying phase, and the cancellations are catastrophic [33, 34]. Analytic continuation, the only remaining classical route, is itself ill-conditioned, as we discuss next. Real-time observables are therefore beyond the reach of Euclidean Monte Carlo, and recovering them is one of the central motivations for quantum simulation.

1.2.2.2 Real-Time Dynamics and the Inverse Problem

Even when the path-integral weight is positive, recovering real-time dynamical information from Euclidean correlators is an ill-posed problem. The real-time Wightman two-point function $G(t) = \langle 0|\hat{O}(t)\,\hat{O}(0)|0\rangle$ and its Euclidean counterpart $C(\tau)$ are both encoded in the spectral function $\rho(\omega) = \sum_n |\langle n|\hat{O}|0\rangle|^2 \delta(\omega - E_n)$ via the Källén–Lehmann representations

$$C(\tau) = \int_0^\infty d\omega\, \rho(\omega)\, e^{-\omega\tau}\,, \qquad G(t) = \int_0^\infty d\omega\, \rho(\omega)\, e^{-i\omega t}\,. \tag{1.19}$$

Recovering $\rho(\omega)$ from noisy Monte Carlo data for $C(\tau)$ measured at a discrete set of Euclidean times $\tau = na$ requires inverting a Laplace transform, which is exponentially ill-conditioned: small statistical perturbations in $C(\tau)$ can produce $O(1)$ changes in the reconstructed $\rho(\omega)$ [35]. The information needed to reconstruct fine spectral features is exponentially suppressed by the Euclidean kernel $e^{-\omega\tau}$.

A closely related obstruction is that in infinite volume, on-shell scattering amplitudes away from threshold cannot be extracted directly from the asymptotic large-time behavior of Euclidean correlators [36], since the Euclidean kernel projects onto threshold. The finite-volume method [37, 38] circumvents this obstruction for elastic two-particle scattering by relating finite-volume energy levels to scattering phase shifts, and substantial progress has been made on three-particle and inelastic extensions [39, 40]. These methods remain computationally demanding and limited in the kinematic range they can reliably access.

The inverse problem and its scattering analogue affect a wide range of physically important real-time quantities that lie beyond reliable classical reach: the thermalization of the quark–gluon plasma in heavy-ion collisions, parton fragmentation and hadronisation [41, 42, 43], transport coefficients such as shear viscosity and conductivity [35, 44, 45], and nuclear processes such as neutrinoless double-beta decay [46].

1.3 Quantum Simulation of Lattice Gauge Theories

The obstructions described in Sec. 1.2.2 are rooted in the use of classical statistical sampling to evaluate the Euclidean path integral. A fundamentally different approach is to work

directly in the Hamiltonian formulation, where space is discretized but time remains continuous and the dynamical content of the theory resides in a quantum state evolving under a Hamiltonian operator. That state can then be represented on a quantum computer. We first develop the Hamiltonian formalism, introduce a scalar field testbed that runs alongside the gauge-theory setting throughout this thesis, then explain how the quantum representation circumvents the limitations of the Euclidean approach, and review the body of work to which this thesis contributes.

1.3.1 The Hamiltonian Formulation

The Hamiltonian formulation of lattice gauge theories was developed in 1975 [47], shortly after Wilson's Euclidean construction (see also [48] for a review). It can be derived from Wilson's Euclidean action by passing to the temporal gauge $A_0 = 0$, taking the temporal lattice spacing to zero while retaining a finite spatial spacing a, and constructing the transfer matrix; the result is a quantum-mechanical system defined on a d-dimensional spatial lattice Λ_s with continuous time. The formulation has long served as a theoretical tool for analytic strong-coupling expansions and, more recently, as the basis for classical tensor-network simulations of low-dimensional gauge theories [49, 50, 51]; here it serves as the foundation for quantum simulation.

The fundamental degrees of freedom on each oriented spatial link ℓ are a unitary matrix-valued operator $\hat{U}(\ell) \in \mathrm{SU}(N)$ and the canonically conjugate *electric field operators* $\hat{L}^a(\ell)$ and $\hat{R}^a(\ell)$, $a = 1, \ldots, N^2 - 1$, which generate left and right G-translations of $\hat{U}$ respectively. They satisfy the algebra

$$
\begin{aligned}
&[\,\hat{L}^a(\ell),\, \hat{L}^b(\ell')\,] = i\,\delta_{\ell\ell'}\, f^{abc}\,\hat{L}^c(\ell)\,,\\
&[\,\hat{R}^a(\ell),\, \hat{R}^b(\ell')\,] = -i\,\delta_{\ell\ell'}\, f^{abc}\,\hat{R}^c(\ell)\,,\\
&[\,\hat{L}^a(\ell),\, \hat{R}^b(\ell')\,] = 0\,,\\
&[\,\hat{L}^a(\ell),\, \hat{U}(\ell)\,] = -T^a\,\hat{U}(\ell)\,,\\
&[\,\hat{R}^a(\ell),\, \hat{U}(\ell)\,] = \hat{U}(\ell)\,T^a\,,
\end{aligned}
\tag{1.20}
$$

the canonical algebra of a quantum rigid rotor on the group manifold G [47, 48]. The link Hilbert space $\mathcal{H}_\ell \cong L^2(G,\, dU)$ decomposes, by the Peter–Weyl theorem, into a direct sum over the unitary irreducible representations of G,

$$
\mathcal{H}_\ell \;\cong\; \bigoplus_{R \in \widehat{G}} \mathcal{H}_R \otimes \mathcal{H}_R^{*}\,,
\tag{1.21}
$$

with basis states $|R; m, n\rangle$ labeled by the irrep R and indices $m, n = 1, \ldots, \dim R$ that transform under left and right G-action respectively. The two electric Casimirs coincide on each link, $\hat{L}^a\hat{L}^a = \hat{R}^a\hat{R}^a \equiv \hat{E}^2$, and act as a multiple of the identity on each Peter–Weyl block:

$$
\hat{E}^2\,|R; m, n\rangle \;=\; C_2(R)\,|R; m, n\rangle\,,
\tag{1.22}
$$

where $C_2(R)$ is the quadratic Casimir of the irrep R. For SU(2), the irreps are labeled by spin $j = 0, \frac{1}{2}, 1, \ldots$ with $C_2(j) = j(j+1)$; for SU(3), they are labeled by Dynkin indices (p, q) with $C_2(p, q) = \frac{1}{3}(p^2 + q^2 + pq + 3p + 3q)$. The link Hilbert space is therefore countably infinite-dimensional—a key structural feature that drives the truncation challenge for quantum simulation.

We now use d to denote the number of *spatial* dimensions, in a notational shift from Sec. 1.2.1 where d denoted Euclidean spacetime dimension; the corresponding spacetime is $(d+1)$-dimensional. The Kogut–Susskind Hamiltonian for a pure SU(N) gauge theory in d spatial dimensions is

$$\hat{H}_{\mathrm{KS}} \;=\; \frac{g^2}{2a^{d-2}} \sum_{\ell \in \Lambda_s^{(1)}} \hat{E}^2(\ell) \;-\; \frac{1}{a^{4-d} g^2} \sum_{p \in \Lambda_s^{(2)}} \mathrm{Re}\,\mathrm{Tr}\,\hat{U}_p \,, \tag{1.23}$$

where $\Lambda_s^{(1)}$ and $\Lambda_s^{(2)}$ denote the sets of spatial links and elementary plaquettes, and $\hat{U}_p$ is the oriented product of link operators around plaquette p. The first term, the *electric energy*, is diagonal in the electric basis (1.21) and dominates at large coupling; the second, the *magnetic energy*, is diagonal in the group-valued (magnetic) basis $\{|U\rangle\}$ and dominates at small coupling. The two non-commuting terms together generate the full dynamics, and their balance reproduces the continuum gauge theory in the limit $a \to 0$, $g(a) \to 0^+$ along the asymptotic-freedom trajectory (1.12). Coupling to fermions adds a mass term and a gauge-covariant nearest-neighbor hopping term constructed from $\hat{U}(\ell)$ and the matter operators $\hat{\psi}(x), \hat{\psi}^\dagger(x)$.

The role of (1.4) in the Euclidean theory is played, in the Hamiltonian formulation, by a constraint on physical states. The generators of local gauge transformations at a site $\vec{x}$ are

$$\hat{G}^a(\vec{x}) \;=\; \sum_{\hat{\jmath}} \Big[\, \hat{L}^a(\vec{x}, \hat{\jmath}) \;+\; \hat{R}^a(\vec{x} - a\hat{\jmath}, \hat{\jmath}) \,\Big] \;-\; \hat{\rho}^a(\vec{x}) \,, \tag{1.24}$$

where $\hat{\rho}^a(\vec{x})$ is the matter color-charge density and the sum runs over the d positive spatial directions. The lattice analogue of $\vec{\nabla} \cdot \vec{E} = \rho$, the operator $\hat{G}^a(\vec{x})$ generates $\Omega(\vec{x})$-rotations of all link variables touching $\vec{x}$. Physical states $|\psi_{\mathrm{phys}}\rangle$ are defined by Gauss's law,

$$\hat{G}^a(\vec{x})\,|\psi_{\mathrm{phys}}\rangle \;=\; 0 \qquad \forall\, \vec{x} \in \Lambda_s,\, \forall\, a \,. \tag{1.25}$$

Because the $\hat{G}^a(\vec{x})$ commute with $\hat{H}_{\mathrm{KS}}$, the physical subspace is preserved under time evolution, and gauge invariance is maintained dynamically by the Hamiltonian itself.

1.3.2 Scalar Field Theory as a Bosonic Testbed

Alongside the gauge-theory Hamiltonian, it will be useful throughout this thesis to work with a parallel scalar field theory testbed. Two of the four results chapters (Chapters 2 and 3) develop algorithmic methods first in the scalar setting before applying them to gauge

theories: scalar field theories retain key structural features of bosonic lattice theories such as infinite-dimensional local Hilbert spaces, non-commuting field and momentum operators, and the need for truncation and digitization, while being free of the additional group-theoretic structure of gauge theories. They therefore serve as a clean setting in which to develop and benchmark simulation algorithms. The lattice Hamiltonian for a scalar ϕ^4 theory in d spatial dimensions is [52]

$$\hat{H}_{\text{quartic}} \;=\; a^d \sum_{i=1}^{N_\Lambda} \left[\frac{1}{2}\hat{\pi}_i^2 + \frac{m^2}{2}\,\hat{\phi}_i^2 + \frac{\lambda}{4!}\,\hat{\phi}_i^4 \right] + a^d \sum_{\langle ij \rangle} \frac{(\hat{\phi}_i - \hat{\phi}_j)^2}{2a^2}\,, \tag{1.26}$$

where $\hat{\phi}_i$ and $\hat{\pi}_i$ are the field and conjugate momentum operators at site i, satisfying $[\hat{\phi}_i, \hat{\pi}_j] = ia^{-d}\delta_{ij}$, and $\langle ij \rangle$ denotes nearest neighbors.

1.3.3 Why Quantum Simulation?

The idea that quantum systems are most naturally and efficiently simulated by other quantum systems was first proposed in 1982 [53] and later formalized [54]. Working in the Hamiltonian formulation and representing the state of the lattice theory on a quantum computer offers two key advantages over the Euclidean Monte Carlo approach, each addressing one of the limitations identified in Sec. 1.2.2.

First, it accommodates the entanglement structure of generic quantum states without exponential overhead. A lattice gauge theory on N_Λ sites with a local Hilbert space of dimension D has a total Hilbert space of dimension D^{N_Λ}, and storing an arbitrary state in this space requires D^{N_Λ} complex amplitudes on a classical computer. Tensor network methods such as matrix product states (MPS) and projected entangled pair states (PEPS) [55, 56, 57] circumvent this exponential cost by representing states whose entanglement entropy across a spatial cut scales with the area of the cut rather than the volume it bounds. Ground states of gapped local Hamiltonians obey such an area law, and tensor networks represent them with a number of parameters polynomial in N_Λ; this is the regime in which tensor network methods have produced striking results for lattice gauge theories in low dimensions [50]. Real-time evolution from a low-entanglement initial state, however, generically drives the system out of the area-law regime: the entanglement entropy across any cut grows linearly in time [58] until it saturates at the volume-law value characteristic of generic states in Hilbert space. The bond dimension needed to track the state therefore grows exponentially in time, rendering tensor network simulation intractable on the timescales of physical interest. A quantum computer, by contrast, represents the state using only $N_\Lambda \log_2 D$ qubits, with the full exponentially large Hilbert space encoded natively in the quantum hardware, and is indifferent to whether the state obeys an area law or a volume law.

Second, the Hamiltonian formulation does not suffer from the sign problem. Time evolution is implemented directly: the operator $e^{-i\hat{H}t}$ is approximated by a sequence of quantum gates acting on the qubit register, evolving the state unitarily in real time. The real-time correlators that required ill-conditioned analytic continuation in Sec. 1.2.2 are now obtained

by applying $e^{-i\hat{H}t}$ to the prepared state, and finite-density calculations sidestep the complex Boltzmann weight entirely.

The seminal work of Jordan, Lee, and Preskill [52, 59] made this vision concrete by demonstrating that a scalar ϕ^4 theory could be simulated efficiently on a quantum computer, providing a polynomial-time algorithm spanning state preparation, time evolution, and particle detection. Since then, an extensive program has extended these ideas to Abelian and non-Abelian gauge theories with a range of gauge groups and spatial dimensions [60, 61, 62, 63, 64, 65, 66, 67, 68, 69]. Proof-of-principle demonstrations on quantum hardware have progressed from the first trapped-ion simulation of the Schwinger model on four ions [70] to recent vacuum preparation on 100 qubits [67] and hadron-dynamics simulations on 112 qubits [71]. For comprehensive reviews, see Refs. [28, 50, 72, 73].

1.4 The Quantum Simulation Pipeline

We now zoom in from the state of the field to the algorithmic stack underlying these demonstrations. Simulating a lattice gauge theory on a quantum computer requires a chain of controlled approximations, each introducing systematic errors and computational costs that must be quantified. The pipeline comprises lattice discretization, truncation, digitization, state preparation, time evolution, measurement, and continuum extrapolation, which we address in turn, establishing the notation and concepts used throughout this thesis.

Lattice discretization. Continuous space is replaced by a hypercubic lattice Λ of N_Λ sites with spacing a in d spatial dimensions. The continuum limit is recovered by taking $a \to 0$ while tuning the bare parameters of the Hamiltonian along a renormalization trajectory so that dimensionful physical observables remain fixed. Discretization errors are controlled by the Symanzik effective field theory [21, 22, 23]: for a lattice action with leading errors at $O(a^k)$, observables receive corrections that vanish as powers of the lattice spacing. Finite-volume errors for scattering observables computed on a quantum computer have recently been quantified in Ref. [74].

Truncation. A feature common to bosonic sectors of both gauge and scalar lattice field theories is that the local Hilbert space is infinite-dimensional and must be truncated to a finite dimension for simulation on a qubit register. For gauge theories in the electric basis, this means imposing a cutoff Λ on the irrep labels of each link, e.g., $j \leq \Lambda$ for SU(2) or $\max(p, q) \leq \Lambda$ for SU(3). For scalar field theories, the analogous step restricts the field range to $|\phi| \leq \phi_{\max}$. In both cases, the truncation error depends on the dynamics: a state initially well-represented by the truncated Hilbert space may develop support on truncated states during time evolution. Prior work [75] has shown that truncation errors decay super-exponentially in the truncation parameter, with Λ scaling polylogarithmically in ε^{-1}, provided Λ is allowed to grow with simulation time (linearly for lattice gauge theories, as $\sqrt{t}$ for generic bosonic systems); tighter, time-independent bounds are the subject of

Chapter 4. For gauge theories, the truncation problem is richer due to the group-theoretic structure of the link Hilbert space, and a variety of approaches have been developed: electric basis truncations [49, 61, 62, 69, 76, 77, 78, 79, 80, 81], discrete subgroups [82, 83, 84, 85, 86, 87], quantum link models [88, 89], hybrid electric/magnetic bases [90, 91, 92, 93, 94], and the orbifold lattice formulation [17, 95, 96]. Each of these schemes trades off qubit cost, gauge-invariance enforcement, and accuracy at strong versus weak coupling differently; the work in this thesis adopts the electric basis truncation, which is the most direct setting for the bounds developed in Chapter 4.

Digitisation. Once the local Hilbert space is truncated to dimension $N = 2^{n_q}$, the truncated field values are encoded into n_q qubits per lattice site. For a scalar field, a standard approach [97] discretizes the field eigenvalues into 2^{n_q} equally spaced values $\phi \in \{-\phi_{\max}, -\phi_{\max} + \delta\phi, \ldots, \phi_{\max}\}$ with spacing $\delta\phi = 2\phi_{\max}/(2^{n_q} - 1)$. The conjugate momentum operator in the field basis is obtained via the quantum Fourier transform on the n_q-qubit register, with $\pi_{\max} = \pi/\delta\phi$ [97, 98, 99, 100]. The digitization error of the low-lying spectrum converges doubly exponentially in n_q for an optimally chosen $\phi_{\max}$ [97]. Altogether, a lattice of N_Λ sites with n_q qubits per site requires $n = N_\Lambda n_q$ total qubits.

State preparation. An appropriate initial quantum state must be prepared on the qubit register before dynamics can begin. For scattering simulations as in [52], this involves preparing wavepacket states on top of the interacting vacuum. The free-theory vacuum can be prepared efficiently from a Gaussian state [52, 97]; reaching the interacting vacuum typically requires adiabatic state preparation or more advanced methods [67, 68, 101]. Recent work has demonstrated scalable vacuum preparation circuits for the Schwinger model on up to 100 qubits [67] and wavepacket preparation with circuit depth independent of the wavepacket size [71].

Time evolution. The core computational primitive is the approximation of the time evolution operator $U(t) = e^{-i\hat{H}t}$ for a time-independent Hamiltonian $\hat{H}$. The exponentiation error is defined as

$$\varepsilon \;\equiv\; \|U_{\text{approx}}(t) - e^{-i\hat{H}t}\|\,, \tag{1.27}$$

where $\|\cdot\|$ denotes the spectral norm. Product formulas (PFs) [102, 103, 104, 105, 106, 107] approximate the evolution as a product of local exponentials over many small time steps; a p-th order PF has per-step error $O(\delta t^{p+1})$, so reaching total time t with target error ε requires a number of steps that scales as $O(t^{1+1/p}\,\varepsilon^{-1/p})$, suppressing factors set by the Hamiltonian norm and the number of terms. Post-Trotter methods based on Quantum Signal Processing (QSP) [108, 109] and its generalizations [110] achieve the provably optimal scaling $O(\alpha_H t + \log(1/\varepsilon))$, where α_H is a measure of the Hamiltonian norm, at the cost of requiring a block encoding as a subroutine: a unitary U_H such that $(\langle 0|^{\otimes a} \otimes I)\, U_H\, (|0\rangle^{\otimes a} \otimes I) = \hat{H}/\alpha$ for some normalization $\alpha \geq \|\hat{H}\|$. The block encoding is typically constructed using the

Linear Combination of Unitaries (LCU) algorithm [111] or the sparse oracle approach [112]. Chapters 2 and 3 takes up the task of constructing efficient block encodings.

Measurement. Physical observables are extracted by performing measurements on the time-evolved state and averaging over N_{shots} repetitions of the full circuit. The statistical error scales as $\Delta\langle\hat{O}\rangle \sim 1/\sqrt{N_{\text{shots}}}$. The total cost of a simulation is the product of the per-shot circuit cost and the number of shots needed for the desired statistical precision.

Continuum limit. Lattice results must be extrapolated to the continuum limit $a \to 0$ by performing simulations at multiple lattice spacings, tuning bare parameters along a renormalization trajectory, performing scale setting, and extrapolating discretization artifacts. The interplay between approximate time evolution and this procedure is subtle: for product formulas, it has been shown [113] that the Trotter error can be absorbed into a modified renormalization trajectory related to the Euclidean transfer matrix, but this approach is specific to PFs and does not extend to post-Trotter algorithms. Developing an algorithm-independent framework is one of the contributions of Chapter 5.

The chapters of this thesis address several of these stages quantitatively: Chapters 2 and 3 focus on time evolution algorithms and the block-encoding subroutines they require, using scalar ϕ^4 theory and $U(1)$ gauge theory as concrete settings; Chapter 4 addresses truncation error control for lattice gauge theories; and Chapter 5 addresses the continuum limit. The methods and frameworks developed are designed to apply broadly to lattice gauge theories with arbitrary gauge groups.

1.5 Summary of Contributions

Chapter 2: Strategies for Simulating Time Evolution of Hamiltonian Lattice Field Theories [114]. We compute the asymptotic gate complexity of product formulas [105, 106] and Quantum Signal Processing [108, 109] for implementing the time evolution operator of two general classes of lattice Hamiltonians: k-site-local and k-geometrically-site-local theories. For each method, we derive the scaling with lattice size N_Λ, qubits per site n_q, simulation time t, and target error ε. We present analytical and numerical evidence for why PFs may perform better for realistic problem parameters even though QSP scales exponentially better with ε due to the overhead of constructing block encodings. As a case study, we apply these results to the scalar ϕ^4 theory in Eq. (1.26). We also introduce a Fourier-transform-based technique for constructing block encodings of bosonic Hamiltonians that exploits the quantum Fourier transform to switch between the field and momentum eigenbases, yielding an exponential improvement in both gate count and ancillary qubit overhead compared to the naive LCU approach [111].

Chapter 3: Block Encoding Bosons by Signal Processing [115]. We develop novel methods for constructing block encodings tailored to bosonic lattice field theories, demonstrating that signal processing techniques [116, 117, 118] can be repurposed as block encoding subroutines such that QSP based time evolution becomes more competitive with PFs. We introduce the following approaches:, QSVT-based [116], QETU-based [118], and an LCU variant called Linear-Operators-Via-Exponentiation (LOVE-LCU). We provide detailed asymptotic analysis and concrete gate counts for operators arising in scalar field and $U(1)$ theories. We find that LOVE-LCU is most efficient at moderate n_q, while QSVT-based methods offer superior asymptotic scaling.

Chapter 4: Truncation Uncertainties for Accurate Quantum Simulations of Lattice Gauge Theories [119]. We derive the tightest upper bounds known to date on the error introduced by truncating the electric field Hilbert space in the Kogut–Susskind Hamiltonian [47]. The key structural insight is that Hilbert space fragmentation [120, 121, 122] suppresses the amplitude for populating truncated states factorially in the truncation Λ. Crucially, these bounds do not require the truncation to grow with simulation time, improving upon previous results [75]. We validate the bounds numerically for U(1) pure gauge theory on infinite plaquette chains and for the Schwinger model on infinite lattices using matrix product state methods.

Chapter 5: Obtaining Continuum Physics from Dynamical Simulations of Hamiltonian Lattice Gauge Theories [123]. We address the interplay between approximate time evolution and the continuum limit. For product formula–based simulations, we show that the higher-order Trotter error terms appearing in the effective Hamiltonian (via the Baker–Campbell–Hausdorff expansion) are irrelevant operators that vanish in the continuum limit. This yields a simplified renormalization trajectory in which the bare parameters are tuned assuming exact time evolution—depending only on the lattice spacing a and not on the Trotter step size δ_t—while introducing only $O(a^p)$ systematic errors, where p is the order of the PF. This trajectory decouples the parameter tuning from the choice of simulation algorithm, allowing parameters determined by one method to be used with any other. More broadly, we develop the Statistically-Bounded Time Evolution (SBTE) protocol, an algorithm-independent framework for taking the continuum limit. The key observation is that no additional ultraviolet divergences arise in the limit of exact time evolution—in particular, no analogue of the fermion doubling problem is possible—so the error from approximate time evolution can be treated as a systematic uncertainty rather than requiring its own renormalization. We also show that the limitations of using classical Euclidean simulations to set renormalization parameters for real-time Hamiltonian evolution are more severe than previously understood [113].

Chapter 2

Strategies for Simulating Time Evolution of Hamiltonian Lattice Field Theories

2.1 Introduction

A major aspect of performing precise numerical simulations of Quantum Field Theories (QFTs) is the ability to efficiently simulate on a quantum computer the time evolution of a given Hamiltonian. In [52] Jordan, Lee, and Preskill made a detailed proposal to simulate a scalar field theory on quantum computers. Similarly to earlier research [54, 124, 125], their work relies on the Suzuki-Trotter approximation [102, 103, 104, 105, 106, 107] to compute the exponential of the Hamiltonian, and hence time evolve the scalar field theory.

This work compares the asymptotic resource requirements of various techniques to implement the time evolution operator of quantum field theories, including Suzuki-Trotter expansions and nearly-optimal simulation strategies, such as Quantum Signal Processing (QSP) [108, 109, 110, 126, 127] and the HHKL method named after the authors of [128]. For each method, we compute the asymptotic dependence with respect to various parameters characterizing the problem at hand, such as the error allowed in the approximation of the evolution operator, measures of Hamiltonian size / complexity, as well as the time for which the system is evolved. This comparison is carried out for two very general classes of lattice Hamiltonians, but we also provide more detailed results for a scalar quantum field theory as an explicit example.

Using nearly-optimal methods requires querying matrix elements of the Hamiltonian, typically in the form of the *Block Encoding* (BE) subroutine. In this work, we consider the BE for a bosonic degree of freedom, based on the Linear Combination of Unitaries (LCU) algorithm [111], and achieve a significant reduction in the number of gates and ancillary qubits by taking advantage of the quantum Fourier transform for switching between the eigenbases of the field and conjugate momentum operators. We do not consider BEs based

on the sparse oracle access input model, for two reasons. First, their construction is highly problem-specific [129, 130, 131, 132, 133]. Second, for local lattice field theories, which we consider as one of the major applications of the present work, our studies [133] indicate that the LCU-based approach is more efficient than that based on the sparse oracle access.

The work presented in this chapter is based on [114].

2.1.1 Definitions and Notations used in this work

We begin by presenting the standard mathematical notation used to describe the asymptotic behavior of functions

Definition 1. *Let $f, g : \mathbb{R} \to \mathbb{R}$ be functions of real variables. Then we say that $f = \mathcal{O}(g)$ if there exists a positive constant c such that $|f(x)| \leq c|g(x)|$ for all sufficiently large choices of x, i.e. $\forall x \geq x_0 \in \mathbb{R}$. We say that $f = \Omega(g)$ if $g = \mathcal{O}(f)$ and that $f = \Theta(g)$ if $f = \mathcal{O}(g)$ and $g = \mathcal{O}(f)$.*

We see that $\mathcal{O}(\cdot)$ and $\Omega(\cdot)$ denote asymptotic upper and lower bounds respectively, while $\Theta(\cdot)$ denotes an asymptotically *tight* bound. We will make use of these notations throughout this work to denote the asymptotic scaling / behavior of the variables we present in this section.

In HLFTs the degrees of freedom are typically the discretized quantum fields (in first quantization) or the eigenstates of the number operator (in second quantization) which live on the sites, links, or plaquettes of a lattice Λ.[1] For many Hamiltonians of interest, each term in the Hamiltonian couples fields belonging to a small subset of sites [47]. Depending on which fields get coupled together, one finds different types of Hamiltonians.

We first consider k-*site local* Hamiltonians of the form

$$\begin{aligned} H &= \sum_{J \in \mathcal{S}} H_J \,, \\ J = \{j_1, \ldots, j_{\mathsf{k}}\} \in \mathcal{S} \,, &\quad j_i \in \Lambda \,, \end{aligned} \tag{2.1}$$

where each H_J acts nontrivially only on at most k sites. The multi-index variable J, used to label each summand appearing in Eq. (2.1), consists of a subset (containing k elements) of the N_Λ total lattice sites. Each lattice site is comprised of n_q qubits, each term H_J acts on $\mathsf{k}\mathsf{n}_\mathsf{q}$ qubits, and the total system contains $\mathsf{n} = \mathsf{N}_\Lambda \mathsf{n}_\mathsf{q}$ qubits.[2] The cardinality of the set $\mathcal{S}$, which contains sets of indices used to label individual terms H_J, is also the total number

[1] In some cases, the lattice is naturally embedded into some metric space, e.g., when a lattice in real or momentum space is used. More generally, it can be considered as the space which the index enumerating the degrees of freedom belongs to. Furthermore, we use the term "lattice site" for any local ingredient of the lattice that can support a Hilbert space. While for fermions and scalar fields these are the actual lattice sites, for gauge fields this terminology also covers the links and plaquettes of the lattice.

[2] We note that our definition of k-site-local Hamiltonian reduces to the usual definition of a k-local Hamiltonian on n-qubits [106] if we set $\mathsf{n}_\mathsf{q} = 1$ and $\Lambda = \{1, 2, \ldots, \mathsf{n}\}$, in which case $\mathsf{N}_\mathsf{H} \sim \mathsf{n}^\mathsf{k}$ and $\mathsf{N}_\mathsf{P} \sim 4^\mathsf{k}$.

of individual terms (denoted by $\mathsf{N_H}$) that form the full Hamiltonian H. We assume that in general, for a k-site-local Hamiltonian, the following scaling relation holds:

$$\mathsf{N_H} \equiv |\mathcal{S}| = \mathcal{O}(\mathsf{N_\Lambda}^{\mathsf{m}}) \,, \tag{2.2}$$

where $\mathsf{m} \left(\in \mathbb{R}_{\geq 0}\right) \leq \mathsf{k}$ is a parameter fixed by the choice of a specific theory.

We note that each summand H_J can be expressed in the Pauli basis as

$$H_J = \sum_{i \in [\mathsf{N_P}^{(J)}]} c_i^{(J)} P_i^{(J)} \,, \tag{2.3}$$

where each $P_i^{(J)} \in \mathbb{C}^{2^n \times 2^n}$ represents an n-qubit Pauli operator, each $c_i^{(J)} \in \mathbb{C} \setminus \{0\}$ is a complex coefficient, and $\mathsf{N_P}^{(J)}$ enumerates the number of terms with non-zero coefficient in the Pauli decomposition of H_J. We introduce the quantity $\mathsf{N_P}$ as the maximum number of Pauli strings appearing in the decomposition of any single term H_J. We define it as follows:

$$\mathsf{N_P} \equiv \max_J \left(\mathsf{N_P}^{(J)}\right) . \tag{2.4}$$

In this way, the total number of Pauli strings appearing in a Hamiltonian with $\mathsf{N_H}$ summands can be upper bounded by $\mathsf{N_H N_P}$. Note that for a k-site-local Hamiltonian $\mathsf{N_P} = \mathcal{O}(4^{\mathsf{k n_q}})$ always holds as an upper bound.

An important quantity in asymptotic gate complexity estimates is the 1-norm of a Hamiltonian H, denoted by $\|H\|_1$, which we define as

$$\|H\|_1 \equiv \sum_{J \in \mathcal{S}} \|H_J\| \,, \tag{2.5}$$

where $\|\cdot\|$ represents the spectral / operator norm. The *induced* 1-norm of the Hamiltonian, denoted by $|||H|||_1$, can then be defined by maximizing the following sum over the qubit-index (i.e. the location within the multi-index variable $J \in \mathcal{S}$) $\ell \in \{1, \ldots, k\}$ and the site-index $j_\ell \in \Lambda$:

$$\begin{aligned} |||H|||_1 &\equiv \max_\ell \max_{j_\ell} \sum_{J_\ell^{(j_\ell)} \in \mathcal{S}_\ell^{(j_\ell)}} \|H_J\| \,, \\ J_\ell^{(j_\ell)} &= \{j_1, \ldots, j_{\ell-1}, j_{\ell+1}, \ldots, j_{\mathsf{k}}\} \in \mathcal{S}_\ell^{(j_\ell)} \,, \\ J &= \{j_1, \ldots, j_{\ell-1}, j_\ell, j_{\ell+1}, \ldots, j_{\mathsf{k}}\} \in \mathcal{S} \,. \quad j_i \in \Lambda \,. \end{aligned} \tag{2.6}$$

In other words, the induced 1-norm performs the sum over all norms $\|H_J\|$ which have the lattice site j_ℓ at the ℓ'th position, both of which are chosen to maximize the sum appearing in Eq. (2.6). This induced norm provides an upper bound on the *strength* of the Hamiltonian at any single site due to interactions with the remaining (allowed by locality restrictions and the specific theory under consideration) sites in the lattice, such that the qubit-index of that

site within the multi-index variable $J \in \mathcal{S}$ is fixed. We denote by $\mathsf{N}_{\mathsf{ind}}$ the number of terms in the sum appearing in Eq. (2.6):

$$\begin{aligned} \mathsf{N}_{\mathsf{ind}} &\equiv |S_{\ell^*}^{(j_{\ell^*})}|\,, \\ \ell^*, j_{\ell^*} &= \operatorname*{argmax}_{\ell} \operatorname*{argmax}_{j_\ell} \sum_{J_\ell^{(j_\ell)} \in \mathcal{S}_\ell^{(j_\ell)}} \|H_J\|\,. \end{aligned} \tag{2.7}$$

For a general k-site-local Hamiltonians, we write

$$\mathsf{N}_{\mathsf{ind}} = \mathcal{O}(\mathsf{N}_\Lambda{}^{\mathsf{m}_{\mathsf{ind}}})\,, \tag{2.8}$$

where $\mathsf{m}_{\mathsf{ind}}$ $(\in \mathbb{R}_{\geq 0}) \leq \mathsf{k} - 1$ is again a parameter fixed by the choice of a specific theory. We also define a quantity $\mathsf{c}_\mathsf{P} \in \mathbb{R}$ to be the maximum absolute value of the individual coefficients in the expansion of each H_J in terms of Pauli operators. Writing each summand H_J as in Eq. (2.3), one finds

$$\mathsf{c}_\mathsf{P} \equiv \max_{J,i}\{|c_i^{(J)}|\}\,, \tag{2.9}$$

so that $\|H\|_1 \leq \mathsf{c}_\mathsf{P}\mathsf{N}_\mathsf{P}\mathsf{N}_\mathsf{H}$ and $|||H|||_1 \leq \mathsf{c}_\mathsf{P}\mathsf{N}_\mathsf{P}\mathsf{N}_{\mathsf{ind}}$.

We define the *exponentiation error* ε as the difference between the exact time evolution operator e^{-iHt} and an approximate implementation $U(t)$

$$\varepsilon \equiv \|U(t) - \mathrm{e}^{-iHt}\|\,. \tag{2.10}$$

We summarize these results in a few definitions, that will be used throughout this paper

Definition 2. *A* k*-site-local Hamiltonian on a lattice* Λ *of size* $\mathsf{N}_\Lambda = |\Lambda|$ *is one of the form:*

$$\begin{aligned} H &= \sum_{J \in \mathcal{S}} H_J\,, \\ J &= \{j_1, \ldots, j_\mathsf{k}\} \in \mathcal{S}\,, \quad j_i \in \Lambda\,, \end{aligned}$$

where the number of terms in the set S *are given by* N_H*, each each* H_J *acts nontrivially only on at most* k *sites* j_i*, labeled by* $J = \{j_1, \ldots, j_\mathsf{k}\}$*, and each lattice site is comprised of* n_q *qubits. Let each term* H_J *take the form Eq.* (2.3) *and, as in Eq.* (2.4)*, let* N_P *be an upper bound on the number of distinct Pauli strings appearing in the decomposition of any single term* H_J*. Let* c_P *be the largest magnitude of any Pauli coefficient over all terms as in Eq.* (2.9)*, and* $\mathsf{N}_{\mathsf{ind}}$ *be the number of terms in the induced 1-norm of the Hamiltonian as in Eq.* (2.7)*.*

Many HLFT Hamiltonians have the property that they only couple fields at neighboring lattice locations together, giving rise to a *geometric* locality. Defining a distance metric to the different lattice sites, allows to define a k-geometrically-site-local Hamiltonian, which is a sum of operators, each involving only sites belonging to a connected set.

Definition 3. *A* k*-geometrically-site-local lattice Hamiltonian is a* k*-site-local lattice of Hamiltonian in which each individual term* H_J *acts on a connected set of sites* J*.*[3]

The restriction of geometric locality implies that each term H_J acts upon belong to a ball of diameter k, which imposes a tighter bound on $\mathsf{N}_{\mathsf{ind}}$:[4]

$$\mathsf{N}_{\mathsf{ind}} \leq \binom{\mathsf{k}^{\mathsf{d}}}{\mathsf{k}} = \mathcal{O}(1)\,, \tag{2.11}$$

i.e., $\mathsf{N}_{\mathsf{ind}}$ no longer scales with N_Λ (note that here we have assumed a constant value for the spatial dimension d). Furthermore, since each set J can be labeled by the center of the corresponding ball, the number of sets and hence terms in the Hamiltonian is given by

$$\mathsf{N}_{\mathsf{H}} = \mathcal{O}(\mathsf{N}_\Lambda)\,. \tag{2.12}$$

Hamiltonians of relativistic quantum field theories typically depend on the values of fields and their derivatives at a given point and thus one can expect that lattice formulations of these theories exist such that they are geometrically local. This is the case for the scalar field theory considered later in this paper for example. For gauge theories, gauge invariance needs to be included into the considerations, and the locality of the Hamiltonian now depends on what basis is chosen for the fields. In so-called electric bases, Gauss's law constraints are still local, such that the final Hamiltonian is geometrically local in this basis. On the contrary, magnetic bases, behaving better at small bare coupling (which is required in the continuum limit), require non-geometrically local interactions [90, 91, 93, 134, 135, 136].

As we will also consider Hamiltonians that act only on $\mathcal{O}(1)$ number of lattice sites, we will also provide the definition for them:

Definition 4. *Let* H *be a* $(\mathsf{k} = \mathsf{N}_\Lambda = \mathcal{O}(1))$*-site-local Hamiltonian, with* N_{P} *and* c_{P} *defined as in Eqs.* (2.4) *and* (2.9)*. Note that* $\mathsf{N}_\Lambda = \mathcal{O}(1)$ *also implies* $\mathsf{N}_{\mathsf{H}} = \mathcal{O}(1)$*, and that there is no difference between a site-local and geometrically site-local Hamiltonian in this case.*

2.1.2 Main results

The main results of this paper are summarized in Tables 2.2 to 2.4. In this section we summarize only the results, while the detailed derivations of these expressions will be given in Sec. 2.3.

We first consider the results for a lattice containing $\mathcal{O}(1)$ sites, and compare the implementation of the time evolution operator after a Product Formula (PF) based splitting to

[3] One typically assumes that the lattice is embedded into some metric space. In our work, we assume a hypercubic lattice with spacing 1. In principle, one could define locality without introducing a metric structure, e.g., by using a more general topology on $\mathcal{S}$.

[4] The precise upper bound in Eq. (2.11) is given for a hypercubic lattice, but statement on independence of $\mathsf{N}_{\mathsf{ind}}$ on N_Λ is general.

Λ	Lattice
$N_\Lambda = \lvert\Lambda\rvert$	Lattice size
J	Multi-index enumerating Hamiltonian terms
$H = \sum_{J\in S} H_J$	Hamiltonian
k	Hamiltonian site-locality
H_J	Term acting on k sites
S	Set containing sets of indices J
$N_H = \lvert S\rvert$	Number of summands forming H
n_q	Number of qubits per lattice site
$n = N_\Lambda n_q$	Total number of qubits
$N_p^{(J)}$	Number of Pauli terms representing a single term H_J
N_p	$\max_J(N_p^{(J)})$
c_P	Maximum Pauli coefficient in H
D	Hamiltonian geometric locality
d	Spatial dimension
ϵ	Exponentiation error
m	Parameter describing the scaling of N_H with respect to N_Λ, see Eq. (2.2)
m_{ind}	Parameter describing the scaling of N_{ind} with respect to N_Λ, see Eq. (2.8)
χ	Total number of elementary gates

Table 2.1: Notations used throughout the paper.

the implementation using Quantum Signal processing (QSP). Both of those techniques will be introduced and defined in Sec. 2.2. The important parameters that the time evolution operator depends on are the time t the system is evolved for, the exponentiation error ε, as well as the maximum number of Pauli terms $\mathsf{N_P}$ and the maximum coefficient $\mathsf{c_P}$ appearing in the Pauli decomposition of any single term in the Hamiltonian. The scaling with respect to these parameters is shown in Table 2.2, where p denotes the order of the PF used. While the scaling of the PF implementation on all parameters improves as p is increased, the number of gates needed is typically exponential in p, so the order of the PF is typically chosen as $p = \mathcal{O}(1)$. One can therefore see that QSP has a better scaling than PFs in all parameters, with the most dramatic improvement being in the scaling in ε, which is exponentially better.

Next, we include the asymptotic scaling with the size of the lattice $\mathsf{N_\Lambda}$ for a k-site local Hamiltonian. The Hamiltonian is a sum of terms H_J, each of which acts on $\mathsf{k} = \mathcal{O}(1)$ lattice sites and in general do not commute with one another. One approach is to assemble the

terms H_J using a PF, with the implementation of each H_J performed either using a PF or QSP, as just discussed. These two choices are compared in the first two rows of Table 2.3. One can see that the scaling is more favorable in all parameters when the individual terms H_J are themselves implemented using a PF.

A second approach is to use QSP for the entire Hamiltonian, with the corresponding scaling shown in the third row of Table 2.3. One can see that this leads to worse scaling (in general) with the number of lattice sites than the PF, while giving a much better scaling with ε. Therefore, the usage of PF may be preferential in situations when $\mathsf{N_H} \gg 1/\varepsilon$, which is often the case in field theory.

Local implementation	Complexity for $\mathcal{O}(1)$ sites
Product formula	$\mathcal{O}\left[(\mathsf{N_P})^{2+\frac{1}{p}} (\mathsf{c_P} t)^{1+\frac{1}{p}} \varepsilon^{-\frac{1}{p}}\right]$
QSP	$\mathcal{O}\left[\mathsf{N_P} \log \mathsf{N_P} \left(\mathsf{N_P c_P} t + \log(1/\varepsilon)\right)\right]$

Table 2.2: Scaling of the time-evolution operator of a Hamiltonian defined on $\mathcal{O}(1)$ lattice sites. The scaling is given in terms of the number of Pauli strings $\mathsf{N_P}$ of a single term in the Hamiltonian Eq. (2.4), the largest coefficient of each of these Pauli strings $\mathsf{c_P}$ Eq. (2.9), and the exponentiation error ε, defined in Eq. (2.10).

Finally, we consider the case of geometrically-local Hamiltonians which have an improved scaling in the number of lattice sites $\mathsf{N_\Lambda}$ due to the geometric locality. As before, the QSP approach wins in terms of dependence on ε while the PF-based wins in terms of scaling with $\mathsf{N_\Lambda}$. The geometric locality of the Hamiltonian, however, allows for the use of the HHKL algorithm, introduced in Sec. 2.2.3, to assemble the H_J terms, with each local term again implemented via either PF or QSP. The scaling of this approach is shown in the second and 5th rows of Table 2.4. We can see that this leads to a better scaling in $\mathsf{N_\Lambda}$, with a combination of HHKL with QSP leading to the best overall scaling. The scaling of course does not give any information on the prefactor, that could be considerably larger for HHKL than for product based approaches [105]. Furthermore, the scaling given here is not tight, so it could be that for realistic scenarios product based formulas still outperform HHKL based methods. We also note that the scaling of PFs is identical to that of HHKL combined with PFs.

As a case study, we consider a scalar lattice field theory and express the obtained asymptotic gate complexities in terms of model parameters. This requires specifying the BE used, and we consider two methods for constructing the BE subroutine U_H in the scalar lattice field theory. The first method uses the standard Linear Combination of Unitaries (LCU) algorithm [111] to prepare a BE of the Hamiltonian H of the entire system, which is efficient for arbitrary local Hamiltonians. The second method improves upon the first by using the Fourier Transform to switch between the eigenbases of position and momentum operators,

which exponentially improves the asymptotic cost of constructing each U_{H_J} in terms of the number of qubits $\mathsf{n_q}$ per site. The BE of the full Hamiltonian is then constructed using LCU to combine each U_{H_J}. The comparisons of costs for these approaches is shown in Fig. 2.9. Finally, we use the best available construction of BE in order to compare the gate counts between the PF and QSP approaches in application to a single site of the φ^4 theory, see Fig. 2.10.

	Simulation method	Asymptotic scaling
(a)	PF for both local terms and their assembly	$\mathsf{N_{ind}}\mathsf{N_P}\left(\mathsf{N_H}\mathsf{N_P}\mathsf{c_P}t\right)^{1+\frac{1}{p}}\varepsilon^{-\frac{1}{p}}$
(b)	QSP for local terms, assembly using a PF	$\mathsf{N_{ind}}\left(\mathsf{N_H}\mathsf{N_P}\mathsf{c_P}t\right)^{1+\frac{1}{p}}\varepsilon^{-\frac{1}{p}}\mathsf{N_P}\log\mathsf{N_P}\log(\mathsf{N_{ind}}\mathsf{N_H}\mathsf{N_P}\mathsf{c_P}t/\varepsilon)$
(c)	QSP for the entire system	$\mathsf{N_H}\mathsf{N_P}\log(\mathsf{N_H}\mathsf{N_P})(\mathsf{N_H}\mathsf{N_P}\mathsf{c_P}t+\log(1/\varepsilon))$

Table 2.3: Asymptotic resource scaling for site-local Hamiltonians defined in Definition 2, where for each entry the asymptotic notation $\mathcal{O}(\cdot)$ is understood and has been suppressed. By local term, we refer to exponentials of the form $\mathrm{e}^{-iH_J t}$.

	Simulation method	Asymptotic scaling
(d)	PF for both local terms and their assembly	$\mathsf{N_P}\left(\mathsf{N_\Lambda}\mathsf{N_P}\mathsf{c_P}t\right)^{1+\frac{1}{p}}\varepsilon^{-\frac{1}{p}}$
(e)	PF for local terms, assembly using HHKL	$\mathsf{N_P}\left(\mathsf{N_\Lambda}\mathsf{N_P}\mathsf{c_P}t\right)^{1+\frac{1}{p}}\varepsilon^{-\frac{1}{p}}$
(f)	QSP for local terms, assembly using a PF	$\left(\mathsf{N_\Lambda}\mathsf{N_P}\mathsf{c_P}t\right)^{1+\frac{1}{p}}\varepsilon^{-\frac{1}{p}}\mathsf{N_P}\log\mathsf{N_P}\log(\mathsf{N_\Lambda}\mathsf{N_P}\mathsf{c_P}t/\varepsilon)$
(g)	QSP for the entire system	$\mathsf{N_\Lambda}\mathsf{N_P}\log(\mathsf{N_\Lambda}\mathsf{N_P})\left(\mathsf{N_\Lambda}\mathsf{N_P}\mathsf{c_P}t+\log(1/\varepsilon)\right)$
(h)	QSP for local terms, assembly using HHKL	$\mathsf{N_\Lambda}{\mathsf{N_P}}^2\mathsf{c_P}t\left[\log(\mathsf{N_\Lambda}t/\varepsilon)\right]^{\mathsf{d}}\log(\mathsf{N_P}\log(\mathsf{N_\Lambda}t/\varepsilon))$

Table 2.4: Asymptotic resource scaling for geometrically site-local Hamiltonians defined in Definition 3, where for each entry the asymptotic notation $\mathcal{O}(\cdot)$ is understood and has been suppressed. As in Table 2.3, by local term we refer to exponentials of the form $\mathrm{e}^{-iH_J t}$. We note that cases d), f), and g) are obtained from cases a), b), and c), respectively, in Table 2.3 upon setting $\mathsf{N_H}=\mathcal{O}(\mathsf{N_\Lambda})$ and $\mathsf{N_{ind}}=\mathcal{O}(1)$.

2.2 Review of Simulation Algorithms

In this section, we review some algorithms commonly used in the literature for simulating the time evolution of generic quantum systems. We will focus on algorithms based on Product Formulas (PFs) [105, 106] and Quantum Signal Processing (QSP) [108, 109], both of which can be applied to small and spatially extended systems. We also review the HHKL algorithm [128], which may serve as a "nearly-optimal wrapper" of local evolution operators in geometrically-local systems. These algorithms will serve as the building blocks from which we express algorithms for Hamiltonian lattice field theories in later chapters.

The discussion in this section aims to be a standalone discussion that gives the reader a basic understanding of the different tools used to implement a time evolution operator. While we aim to make this section accessible to readers without much background in the topic, we will still provide a precise set of statements and theorems that will be used later. Note that this section is not detailed enough to re-derive all of these results, and for this we refer the reader to the papers cited in this section.

Note also that this discussion is kept very general, without restricting ourselves to the Hamiltonians considered in this work (except when stated explicitly otherwise).

2.2.1 Time evolution by Product Formulas

We begin by considering a time-independent Hamiltonian

$$H = \sum_{\gamma=1}^{\Gamma} H_\gamma \tag{2.13}$$

comprised of Γ summands, such that one can efficiently construct quantum circuits to implement the exponential $\mathrm{e}^{-iH_\gamma t}$ of each summand appearing on the RHS of Eq. (2.13). The unitary operator (called the time evolution operator for H) given by the exponential e^{-iHt} describes evolution under the full Hamiltonian H for some arbitrary time t. PFs allow us to construct an approximation to the full time evolution operator by decomposing it into a product of exponentials, each involving a single summand. The general form of a PF (see [106]) is given by

$$\mathrm{e}^{-iHt} \approx S(t) = \prod_{\upsilon=1}^{\Upsilon} \mathcal{P}_\upsilon \left(\prod_{\gamma=1}^{\Gamma} \mathrm{e}^{-ita_\upsilon^{(\gamma)} H_\gamma} \right), \tag{2.14}$$

such that each $a_\upsilon^{(\gamma)} \in \mathbb{R}$ and $\sum_{\upsilon=1}^{\Upsilon} a_\upsilon^{(\gamma)} = 1 \ \forall \gamma \in \mathcal{S} = [\Gamma] \equiv \{1, 2, \ldots, \Gamma\}$, where the set $\mathcal{S}$ was introduced in Eq. (2.1). The latter condition ensures that each individual term H_γ is evolved for the total time t. The parameter Υ determines the total number of stages in the PF, and for each stage $\upsilon \in [\Upsilon]$ there exists a corresponding permutation operator $\mathcal{P}_\upsilon$ that determines the ordering of the individual exponentials within that stage (the specific actions of the operators $\mathcal{P}_\upsilon$ depend on the chosen PF construction). PFs of the form Eq. (2.14)

provide good approximations to the full time evolution operator when the evolution time t is small, and in particular one defines a p^{th} order PF $S_p(t)$ as one that is correct up to p^{th} order in t:

Definition 5. *A p^{th} order PF $S_p(t)$ is a PF of the form Eq.* (2.14)*, such that it approximates the true exponential e^{-iHt} to p^{th} order in t, i.e.,*

$$\mathrm{e}^{-iHt} = S_p(t) + \mathcal{O}(t^{p+1})\,. \tag{2.15}$$

The simplest example of a PF arises from the use of the Baker-Campbell-Hausdorff formula

$$\mathrm{e}^{A+B} = \mathrm{e}^{A}\mathrm{e}^{B} + \mathcal{O}([A,B])\,, \tag{2.16}$$

and thus we can write

$$\mathrm{e}^{-iHt} \approx S_1(t) = \prod_{\gamma=1}^{\Gamma} \mathrm{e}^{-iH_\gamma t}\,. \tag{2.17}$$

This approximation is known as the first order Lie-Trotter formula. The higher (even) order Suzuki-Trotter formulas are defined recursively as

$$\begin{aligned} S_2(t) &\coloneqq \prod_{i=1}^{\Gamma} \mathrm{e}^{\frac{-iH_i t}{2}} \prod_{i=\Gamma}^{1} \mathrm{e}^{\frac{-iH_i t}{2}}\,, \\ S_{2k}(t) &\coloneqq S_{2k-2}(u_k t)^2 S_{2k-2}((1-4u_k)t) \times S_{2k-2}(u_k t)^2\,, \end{aligned} \tag{2.18}$$

where $u_k \equiv 1/\left(4 - 4^{1/(2k-1)}\right)$, and the product given by $\prod_{i=\Gamma}^{1}(\cdot)$ corresponds to reversing the order in which we choose the variable i with respect to the first product given by $\prod_{i=1}^{\Gamma}(\cdot)$. In theory, one could use such a construction to build arbitrarily large order PFs. However the number of stages Υ, and hence the number of elementary gate operations required to construct a circuit implementing such a PF, grows exponentially with the order p (see [106]). For this reason, we use relatively low order formulas for practical applications and our results assume that $p, \Upsilon(p) = \mathcal{O}(1)$. Note that Eq. (2.18) is only one possible form of Eq. (2.14), and investigating alternatives to it is an active area of research [137, 138, 139, 140, 141, 142, 143].

Since PFs have corrections that depend polynomially on the evolution time t, they give good approximations for small t. Given a PF valid for small t, however, we can obtain one for longer times. The general procedure is to divide the full evolution time t into r (known as the *Trotter number*) smaller time steps, and apply the chosen p^{th} order PF to each step, describing the evolution for a smaller time t/r, before multiplying together the results from the individual steps

$$S_p(t) = (S_p(t/r))^r, \tag{2.19}$$

and the associated exponentation error ε is given by

$$\varepsilon = \|(S_p(t/r))^r - \mathrm{e}^{-iHt}\| \,. \tag{2.20}$$

To work out the resource requirements to satisfy such an error bound, we ask for the required Trotter number r for a p^{th} order formula to satisfy Eq. (2.20) given a fixed value of ε. In general the answer depends on the asymptotic scaling on the evolution time t, as well as the norms of the individual summands H_γ and the commutators amongst them.

While Eq. (2.15) provides the correct asymptotic dependence of the error on t, [106] and [112] present bounds that show explicit dependence on the 1-norm of the full Hamiltonian (given by $\sum_{\gamma=1}^{\Gamma} \|H_\gamma\|$). Note however, that these bounds are not tight. For example, when all the summands H_γ commute, the exponentiation error ε becomes zero while these error bounds can become arbitrarily large. This was further addressed in [106], where the authors present another (tighter) bound that results from exploiting the commutativity properties of the summands H_γ:

Theorem 6. *Let $H = \sum_{\gamma=1}^{\Gamma} H_\gamma$ be an operator comprised of Γ Hermitian summands. Let $S_p(t)$ be a p^{th} order PF as defined by Eq. (2.15), with $t \geq 0$. Then, if we require an exponentiation error $\mathcal{O}(\varepsilon)$, as defined by Eq. (2.20), it suffices to choose*

$$r = \mathcal{O}\left(\frac{\tilde{\alpha}^{\frac{1}{p}} t^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}} \right), \tag{2.21}$$

where $\tilde{\alpha} \equiv \sum_{\gamma_1,\gamma_2,\ldots,\gamma_{p+1}=1}^{\Gamma} |||[H_{\gamma_{p+1}}, \ldots [H_{\gamma_2}, H_{\gamma_1}]]||$

Proof. This theorem follows directly from Theorem 11 and Corollary 12 of [106]. □

Note that not every Trotter number r satisfying the asymptotic upper bound Eq. (2.21) will yield the desired error bound. Rather, the above result merely provides an asymptotic upper bound on the *minimum* Trotter number r needed.

2.2.2 Time evolution by Qubitization

While most applications to date simulating relativistic quantum field theories use PFs, there is an alternate approach for simulating the time evolution of quantum systems that is based on a very general technique called *Quantum Signal Processing* (QSP) [108]. As described in more detail in Sec. 2.2.2.4, QSP can be used to apply broad classes of matrix functions of the Hamiltonian to a quantum state.

We begin by clarifying the terminology as confusion often arises in the literature. QSP is an algorithm which, given access to a circuit implementing a unitary operator W with eigenvalues $\{\mathrm{e}^{i\theta_\lambda}\}$, constructs a circuit implementing a unitary V with eigenvalues $\{\mathrm{e}^{ih(\theta_\lambda)}\}$, for a wide class of polynomials h. In the context of quantum simulation the unitary W is taken to be the so-called *Szegedy quantum walk operator* W_H associated with the Hamiltonian

H. The quantum walk operator W_H is a particular type of Hamiltonian *block encoding* operator, which can be constructed from more general block encodings U_H by means of the *Qubitization* [109] procedure. Confusingly, the term *Qubitization* is also often used for the entire process of simulating time evolution based on the usage of QSP and calls to W_H. Moreover, sometimes *Qubitization* serves as an umbrella term for all the simulation techniques which are not based on product formulæ and/or rely on some form of block encoding. On top of that, *Qubitization* has also been used in a completely unrelated context, for denoting the process of mapping physical degrees of freedom onto qubits [144, 145]. In this work, *Qubitization* will be used in the first, restricted sense, referring to the process of constructing W_H.

There are two main motivations for studying QSP-based quantum simulation algorithms. First, they have a better asymptotic complexity than PFs [108, 109]. Second, they are compatible with a wider class of Hamiltonians than those which can be efficiently mapped onto Pauli Hamiltonians, and for which the number of qubits is sublinear in the number of degrees of freedom in the system [131, 132, 146]. As with PFs, the aim of this section is a pedagogical overview of techniques based on QSP, however we do provide precise statements of the theorems needed in the rest of the paper. We will rely on results from the literature to provide proofs of these theorems.

Simulating the time evolution operator of a given Hamiltonian using QSP relies on a a few key insights, which we will now review in turn.

2.2.2.1 Block encodings

The key ingredient of most near-optimal simulation algorithms is the construction of the so-called *Block Encoding* (BE) of a given operator. BE typically refers to the embedding of an n-qubit operator $A \in \mathcal{C}^{2^n \times 2^n}$, acting on an n-qubit Hilbert space $\mathcal{H}_s$, in the (without loss of generality) principle block of a larger $(m+n)$-qubit unitary operator $U_A \in \mathcal{C}^{2^{n+m} \times 2^{n+m}}$, acting on a larger Hilbert space $\mathcal{H}_a \otimes \mathcal{H}_s$, where $\mathcal{H}_a$ is the m-qubit ancilla space. The BE U_A has the following form:

$$U_A = \begin{pmatrix} A/\alpha & * \\ * & * \end{pmatrix}, \tag{2.22}$$

where each $*$ represents some block to be chosen such that U_A is unitary and the quantity α, known as the scale factor, is chosen such that $||A/\alpha||_2 \leq 1$ since any sub-block of a unitary matrix must have its singular values upper bounded by 1. In this case, we call U_A the (α, m)-block-encoding of A. Eq. (2.22) therefore implies that U_A is, in fact, a BE for the entire equivalence class of matrices A differing by a constant factor. Thus, we see that for $|0\rangle_a \equiv |0^{\otimes m}\rangle_a \in \mathcal{H}_a$ and $|\psi\rangle_s \in \mathcal{H}_s$:

$$A/\alpha = (\langle 0|_a \otimes \mathbb{1}_s) U_A (|0\rangle_a \otimes \mathbb{1}_s)\,, \tag{2.23}$$

$$U_A|0\rangle_a|\psi\rangle_s = |0\rangle_a (A/\alpha|\psi\rangle_s) + |\perp\rangle_{as}\,, \tag{2.24}$$

where $|\perp\rangle_{as}$ is a vector perpendicular to $|0\rangle_a$ in the sense that $(\langle 0|_a \otimes \mathbb{1}_s)|\perp\rangle_{as} = 0$. This implies that we recover the action of the operator A on $|\psi\rangle_s$ only if the ancillary register is $|0\rangle_a$, see Fig. 2.1. The probability that we correctly recover the required action of A, known as the *success probability*, is given by $(||A\,|\psi\rangle_s||/\alpha)^2$.

We note that in general, one could use a different state $|g\rangle_a = G|0\rangle_a$ as the state that projects out the block of U_A containing the operator A/α to project out the sub-block containing the rescaled matrix of interest [109]. Thus, we may generalize Eqs. (2.23) and (2.24) as follows:

$$A/\alpha = (\langle g|_a \otimes \mathbb{1}_s) U_A (|g\rangle_a \otimes \mathbb{1}_s)\,, \tag{2.25}$$

$$U_A|g\rangle_a|\psi\rangle_s = |g\rangle_a(A/\alpha|\psi\rangle_s) + |\perp\rangle_{as}\,, \tag{2.26}$$

where $|\perp\rangle_{as}$ is a vector perpendicular to $|g\rangle_a$.

We summarize our discussion on BE thus far and present the following definition for block encodings:

Definition 7. *Let A be an operator acting on the n-qubit Hilbert space $\mathcal{H}_s \cong \mathbb{C}^{2^n}$. Let $k \geq 0$ be an integer and let U_A be a unitary operator on the joint Hilbert space $\mathcal{H}_a \otimes \mathcal{H}_s$, where $\mathcal{H}_a \cong \mathbb{C}^{2^k}$ denotes the Hilbert space of k ancillas. Then, for $\alpha \in \mathbb{R}^+$, the unitary operator U_A is called an (α, k)-block-encoding of A if for every state $|\psi\rangle \in \mathcal{H}_s$:*

$$\left(\mathbb{1}_s \otimes \langle g|_g\right) U_A \left(\mathbb{1}_s \otimes |g\rangle_a\right) = A/\alpha\,, \tag{2.27}$$

where $|g\rangle_a$ is some state in $\mathcal{H}_a$, $\mathbb{1}_s$ is the identity operator on $\mathcal{H}_s$. The value α chosen such that $\|A\|/\alpha \leq 1$ is called the scale factor of the BE.

In much of this work our goal will be to block encode a Hamiltonian H, and we call the corresponding BE U_H. The scale factor can then be absorbed by evolving a rescaled Hamiltonian

$$\tilde{H} = H/\alpha, \tag{2.28}$$

for the time

$$\tilde{t} = \alpha\, t\,, \tag{2.29}$$

so that

$$Ht = \tilde{H}\tilde{t}\,. \tag{2.30}$$

In the following discussion, we will choose a diagonal basis for the (rescaled) system Hamiltonian, such that basis vectors of the system Hilbert space are eigenstates of the Hamiltonian

$$\tilde{H}|\lambda\rangle_s = \lambda|\lambda\rangle_s\,, \tag{2.31}$$

with $|\lambda| < 1$. This basis is used only to simplify the discussion of the technique. Being able to diagonalize the Hamiltonian is not required for its use.

As already discussed, the vector $|g\rangle_a$ of the ancillary Hilbert space $\mathcal{H}_a$ defines the block in which the Hamiltonian is encoded in the unitary matrix U_H. In other words, the BE acts on the tensor product of a state $|\lambda\rangle_s$ and $|g\rangle_a$ as

$$U_H|g\rangle_a|\lambda\rangle_s = \lambda\,|g\rangle_a|\lambda\rangle_s + \sqrt{1-\lambda^2}|\perp^\lambda\rangle_{as}\,. \tag{2.32}$$

Thus, the action of $\tilde{H}$ can be extracted by acting with U_H on an enlarged Hilbert space and post-selecting on the ancilla qubit being in the $|g\rangle_a$ state, *i.e.*

$$\langle g|_a U_H|g\rangle_a|\lambda\rangle_s = \lambda\,|\lambda\rangle_s\,. \tag{2.33}$$

Note that if the ancillary Hilbert space $\mathcal{H}_a$ is 2-dimensional, one could choose $|g\rangle_a = |0\rangle_a$, which would imply $|\perp\rangle_{as} = |1\rangle_a|\chi\rangle_s$ for some $|\chi\rangle_s \in \mathcal{H}_s$. Also note that by direct computation one can use Eq. (2.23) and Eq. (2.32) to show that ${}_{as}\langle\perp^\lambda|\perp^{\lambda'}\rangle_{as} = {}_s\langle\lambda|\lambda'\rangle_s$.

Many different techniques for BE have been discussed in existing literature. Among many others, one can use general algorithms such as the FABLE algorithm [147], approaches that rely on a technique called linear combination of unitaries (LCU) [111, 117], a technique called QETU [118], and techniques using access to the matrix elements of the sparse matrix and compact mappings H [109, 117, 129, 131, 132, 148]. The choice of BE is typically dictated by the number of ancillary qubits required, the value of the scale factor α, the gate complexity required for its implementation, whether an approximate implementation of the Hamiltonian is sufficient, and whether it is applicable to general sparse Hamiltonians.

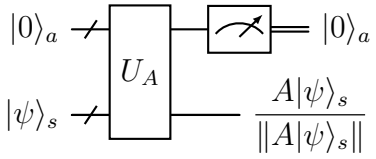

Figure 2.1: Circuit for implementing the block encoding U_A of an operator A, as defined in Eq. (2.23). Upon measuring the ancillary register in the state $|0\rangle_a$, the state of the system $|\psi\rangle_s$ is mapped to $A|\psi\rangle_s/\|A|\psi\rangle_s\|$.

2.2.2.2 Linear Combination of Unitaries

The Linear Combination of Unitaries (LCU) method enables the construction of BEs of operators by breaking an operator T acting on some system Hilbert space $\mathcal{H}_s$ into a sum of into unitary operators U_i, each with known implementation:

$$T = \sum_{i=0}^{K-1} c_i U_i\,, \tag{2.34}$$

where the coefficients c_i are chosen to be real and positive. To implement LCU, one adds an ancillary qubit register with $\lceil \log K \rceil$ qubits and defines a *select* oracle (denoted as SEL) as follows:

$$\text{SEL} \equiv \sum_{i=0}^{K-1} |i\rangle \langle i| \otimes U_i \,. \tag{2.35}$$

SEL implements each unitary U_i conditioned on the state $|i\rangle$ encoded as a binary string on the a ancillary qubits.

One also defines a *prepare* oracle (denoted as PREP) as follows:

$$\text{PREP} \left|0^{\otimes a}\right\rangle = \frac{1}{\sqrt{c}} \sum_{i=0}^{K-1} \sqrt{c_i} \, |i\rangle \,, \tag{2.36}$$

where $c \equiv \sum_{i=1}^{K-1} |c_i|$. Given the definitions of these oracles, one can verify that the quantum circuit given in Fig. 2.2 implements the action of the operator T on some input state $|\psi\rangle \in \mathcal{H}_s$ and thus prepares a BE of T. This result can be summarized with the following lemma

Lemma 8 (LCU block encoding)**.** *Let $U_0 \ldots, U_{K-1}$ be K unitary operators acting on a system Hilbert space $\mathcal{H}_s$ and let $T = \sum_{i=0}^{K-1} c_i U_i$, with $(c_0, \ldots, c_{K-1}) \in \mathbb{C}^K$, be an operator also acting on $\mathcal{H}_s$. Then, the unitary given by $\left(\text{PREP}^\dagger \otimes \mathbb{1}_s\right) \text{SEL} \left(\text{PREP} \otimes \mathbb{1}_s\right)$, where $\mathbb{1}_s \in \mathcal{H}_s$ is the identity operator, is a $(c, \log K)$ BE of the operator T, where $c \equiv \sum_{i=0}^{K-1} |c_i|$. The elementary gate complexity of constructing this BE asymptotically scales as $\mathcal{O}(K \log K)$.*

Proof. The proof of this Lemma can be found for example in [117]. The asymptotic cost stems from the complexity of implementing the SEL subroutine, since the cost of preparing an arbitrary state on $\log_2 M$ qubits in PREP is $\mathcal{O}(M)$ [149, 150], which is cheaper than the cost of SEL, as we explain next. The gate complexity of SEL is $\mathcal{O}(M \log M)$, which comes from sequentially implementing M unitaries, each being controlled on $k = \log M$ ancillary qubits. Decomposing k-controlled gates into elementary gates can be achieved at a cost linear in k, with the constant prefactor being smallest if k additional ancillary qubits are allowed to be used [130, 151]. For applications considered in this work, this gives a minor overhead in the qubit cost since k is logarithmic is the parameters of the system. □

While this approach to LCU relies on the binary encoding of terms on the RHS of Eq. (2.34) in the ancillary register, other encodings can be used as well. In particular, a variation of the LCU algorithm exists, which takes advantage of the Hamiltonian locality [152]. For an k-local $\mathsf{n_q}$-qubit Hamiltonian, it requires $2\mathsf{n_q}$ ancillary qubits yet only uses $2k$-controlled gates.

The LCU algorithm can be applied to a Hamiltonian provided in a form of a $2^{\mathsf{n_q}} \times 2^{\mathsf{n_q}}$ matrix, upon decomposing it into Pauli strings. This approach to constructing BEs is often combined with other methods that result in more efficient constructions. In Sec. 2.4.3.2 we shall consider one such modification which takes into account the structure of the particular lattice Hamiltonians we consider.

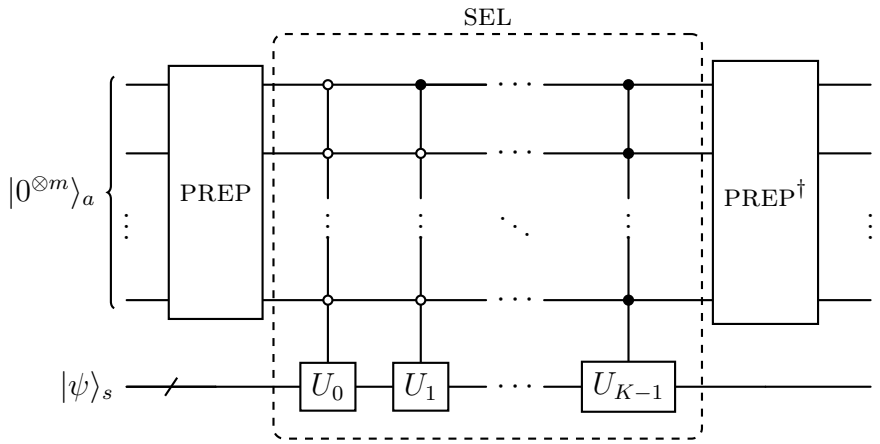

Figure 2.2: The Linear Combination of Unitaries (LCU) circuit [111] to implement the BE of an operator $T = \sum_{i=0}^{K-1} c_i U_i$, a linear combination of unitary operators with known implementations. Operators, PREP and SEL are defined in Eq. (2.36) and Eq. (2.35), correspondingly.

Lastly, we briefly mention that simulation techniques based on BE, considered in the next section, are compatible with a wide class of *sparse* Hamiltonians, including those which cannot be efficiently mapped onto Pauli strings. A general approach to constructing BEs for such systems is based on the usage of *sparse oracle* subroutines. As was mentioned in Sec. 2.1, their construction is highly problem-specific [129, 130, 131, 132, 133] and does not seem to provide advantage over LCU for local HLFTs [133], which are considered as the main application of this work.

2.2.2.3 Qubitization

In approaches to quantum simulation based on Quantum Signal Processing (QSP), one approximates the time evolution operator e^{-iHt} by constructing polynomials of the Hamiltonian operator. In what follows we shall base our algorithms upon the Qubitization technique [109], which utilizes a special *qubitized* form of block encoding W_H, also known as *Szegedy quantum walk operator* or the *iterate*).[5]

One issue with BEs discussed so far is that the action of the BE does not stay in the Hilbert space spanned by $\mathrm{span}\{|g\rangle_a|\lambda\rangle_s, |\perp^\lambda\rangle_{as}\}$. In other words, the action of U_H on $|\perp^\lambda\rangle_{as}$

[5]A different QSP-based approach to approximating e^{-iHt} amounts to using Quantum Singular Value Transformation [116, 117]. This method is advantageous in the sense that it avoids controlled calls to the block encoding, and uses directly U_H instead of W_H. However, as such an approach does not allow one to simultaneously implement both even and odd parts of e^{-iHt}, it requires adding them using LCU at the final stage. This, in turn, leads to success probabilities which (even upon potential amplitude amplification) cannot be made exponentially close to 1, as required by the HHKL algorithm.

can produce a state with different value of $\lambda' \neq \lambda$, such that in general

$$\begin{aligned} \langle g|_a \langle \lambda'|_s U_H |\perp^\lambda\rangle_{as} \neq 0\,, \\ \langle \perp^{\lambda'}|_{as} U_H |\perp^\lambda\rangle_{as} \neq 0\,. \end{aligned} \tag{2.37}$$

Therefore, repeated actions of the unitary U_H are not easily known. The qubitized block encoding W_H is designed so that its successive applications to states of the form $|g\rangle_a|\lambda\rangle_s$ produce a state belonging to a two-dimensional subspace $\text{span}\{|g\rangle_a|\lambda\rangle_s, |\perp^\lambda\rangle_{as}\}$, for all eigenstates $|\lambda\rangle$:

$$\begin{aligned} W_H|g_\lambda\rangle_{as} &= \lambda\,|g_\lambda\rangle_{as} + \sqrt{1-\lambda^2}|g_\lambda^\perp\rangle_{as}\,, \\ W_H|g_\lambda^\perp\rangle_{as} &= -\sqrt{1-\lambda^2}|g_\lambda\rangle_{as} + \lambda\,|g_\lambda^\perp\rangle_{as}\,, \end{aligned} \tag{2.38}$$

where $|g_\lambda\rangle_{as} \equiv |g\rangle_a|\lambda\rangle_s$ and $|g_\lambda^\perp\rangle_{as} \equiv |\perp^\lambda\rangle_{as}$. In other words, for each eigenstate $|\lambda\rangle$, the subspace $\text{span}\{|g_\lambda\rangle_{as}, |g_\lambda^\perp\rangle_{as}\}$ stays invariant under the repeated action of powers of W_H. In what follows, we will show that repeated applications of W_H, in fact, produce block encodings of Chebyshev polynomials of H, which would not be true for U_H.

For any dimension of $\mathcal{H}_a$, the orthogonal state $|g_\lambda^\perp\rangle_{as}$ can now be defined as

$$|g_\lambda^\perp\rangle_{as} = \frac{W_H - \lambda \mathbb{1}_{as}}{\sqrt{1-\lambda^2}}|g_\lambda\rangle_{as}\,, \tag{2.39}$$

which implies that for each eigenvalue λ of the Hamiltonian there is one specific transverse vector $|g_\lambda^\perp\rangle_{as}$ in $\mathcal{H}_a \otimes \mathcal{H}_s$.

From Eq. (2.38) it follows that the information relevant to block encoding of H is contained in two-dimensional blocks of $W_H^{(\lambda)}$ corresponding to subspaces

$$\hbar_a^{(\lambda)} = \text{span}\{|g_\lambda\rangle_{as}, |g_\lambda^\perp\rangle_{as}\} \in \mathcal{H}_a \otimes \mathcal{H}_s\,, \tag{2.40}$$

which we denote by $W_H^{(\lambda)}$:

$$\begin{aligned} W_H^{(\lambda)} &= \begin{pmatrix} \langle g_\lambda\ |_{as} W_H^{(\lambda)} |g_\lambda\rangle_{as} & \langle g_\lambda\ |_{as} W_H^{(\lambda)} |g_\lambda^\perp\rangle_{as} \\ \langle g_\lambda^\perp|_{as} W_H^{(\lambda)} |g_\lambda\rangle_{as} & \langle g_\lambda^\perp|_{as} W_H^{(\lambda)} |g_\lambda^\perp\rangle_{as} \end{pmatrix}_{\hbar_a^{(\lambda)}} \\ &= \begin{pmatrix} \lambda & -\sqrt{1-\lambda^2} \\ \sqrt{1-\lambda^2} & \lambda \end{pmatrix}_{\hbar_a^{(\lambda)}} \\ &= \exp(-iY\theta_\lambda)\,, \end{aligned} \tag{2.41}$$

where

$$\theta_\lambda = \arccos(\lambda)\,. \tag{2.42}$$

While the qubitized BE W_H gives the same result as the original BE U_H when acted on $|g_\lambda^\perp\rangle_{as}$, it has a much simpler structure when applied multiple times:

$$\langle g_\lambda|_{as} \left(W_H\right)^k |g_{\lambda'}\rangle_{as} = \left(W_H^{(\lambda)}\right)^k \delta_{\lambda\lambda'}\,. \tag{2.43}$$

Using Eq. (2.41), one can show that

$$\begin{aligned}\left(W_H^{(\lambda)}\right)^k &= \begin{pmatrix} \langle g_\lambda |_{as} \left(W_H^{(\lambda)}\right)^k |g_\lambda\rangle_{as} & \langle g_\lambda |_{as} \left(W_H^{(\lambda)}\right)^k |g_\lambda^\perp\rangle_{as} \\ \langle g_\lambda^\perp |_{as} \left(W_H^{(\lambda)}\right)^k |g_\lambda\rangle_{as} & \langle g_\lambda^\perp |_{as} \left(W_H^{(\lambda)}\right)^k |g_\lambda^\perp\rangle_{as} \end{pmatrix}_{\hbar_a^{(\lambda)}} \\ &= \begin{pmatrix} T_k(\lambda) & -\sqrt{1-\lambda^2}\, U_{k-1}(\lambda) \\ \sqrt{1-\lambda^2}\, U_{k-1}(\lambda) & T_k(\lambda) \end{pmatrix}_{\hbar_a^{(\lambda)}} \\ &= \exp\left(-iYk\theta_\lambda\right), \end{aligned} \tag{2.44}$$

where $T_k(\lambda)$ and $U_k(\lambda)$ are Chebyshev polynomials of the first and second kind, respectively. Given that Chebyshev polynomials are an excellent polynomial basis for approximating L^∞ functions on a finite interval, one can hope to construct a general function of the eigenvalues (and therefore of the Hamiltonian matrix itself) by combining powers of the W_H operator. Note that while the matrix of $W_H^{(\lambda)}$ may contain blocks of size larger than 2×2 when the ancillary register contains more than a single qubit, only its 2×2 blocks $W_H^{(\lambda)}$ are relevant for the block encodings of H and its powers. In the next section we discuss how this technique can be used for approximating the time evolution operator.

While so far we have only defined the unitary W_H through its action on a given state, a very important result is that it can be obtained from a BE of the Hamiltonian U_H using a process known as *qubitization*, which we now describe. In many cases of interest, as shown in [109], the iterate can be written as

$$W_H = \left(R_g \otimes \mathbb{1}_s\right)\left(S \otimes \mathbb{1}_s\right) U_H\,, \tag{2.45}$$

where $R_g = \left(2\,|g\rangle_a \langle g|_a - \mathbb{1}_a\right)$ implements reflection with respect to the state $|g\rangle_a$, if there exists an operator S satisfying

$$\left(\langle g|_a \otimes \mathbb{1}_s\right)\left(S \otimes \mathbb{1}_s\right) U_H \left(|g\rangle_a \otimes \mathbb{1}_s\right) = H/\alpha\,, \tag{2.46a}$$

$$\left(\langle g|_a \otimes \mathbb{1}_s\right)\left[\left(S \otimes \mathbb{1}_s\right) U_H\right]^2 \left(|g\rangle_a \otimes \mathbb{1}_s\right) = \mathbb{1}_s. \tag{2.46b}$$

The circuits for W_H and R_g are shown in Figs. 2.3 and 2.4.

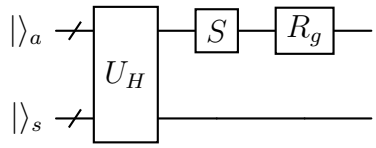

Figure 2.3: Circuit for the qubitized block encoding W_H defined in Eq. (2.45).

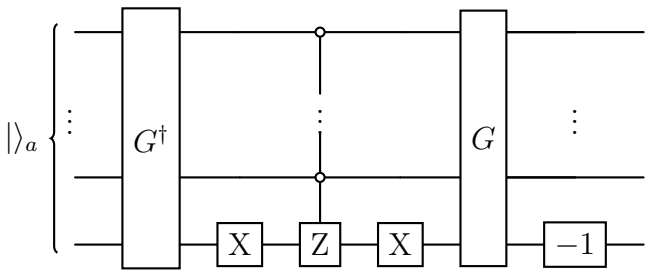

Figure 2.4: Circuit for the reflection operator R_g shown in Eq. (2.45), see [117] for alternative implementations. G is the oracle for preparing the state $|g\rangle_a$, i.e. $G|0\rangle_a = |g\rangle_a$. The $\boxed{-1}$ gate adds an overall phase to the circuit (negative sign), and can be implemented, e.g., as XZXZ.

One particularly simple case is when U_H is its own inverse, *i.e.* $U_H^2 = \mathbb{1}_{as}$. In this case, one can simply set $S = \mathbb{1}_a$. This construction can be applied to the case when the block encoding U_H is constructed via the LCU procedure [130]. To see this, we must modify our perspective slightly. In the original construction, we had $U_H = (\textsc{prep}^\dagger)\textsc{sel}(\textsc{prep})$ and $|g\rangle_a = |0\rangle_a$. Alternatively, we can regroup terms and instead take $U_H \to U_H^{\text{LCU}} \equiv \textsc{sel}$ and $|g\rangle_a \to |g^{\text{LCU}}\rangle_a \equiv \textsc{prep}|0\rangle_a = |\beta\rangle_a$. Because $(U_H^{\text{LCU}})^2 = \textsc{sel}^2 = \mathbb{1}_{as}$, we can simply set $S = \mathbb{1}_a$ to construct W_H:

$$\begin{aligned}
&\left(\langle 0|_a \otimes \mathbb{1}_s\right) U_H \left(|0\rangle_a \otimes \mathbb{1}_s\right) \\
&= \left(\langle 0|_a \otimes \mathbb{1}_s\right) (\textsc{prep}^\dagger)\textsc{sel}(\textsc{prep}) \left(|0\rangle_a \otimes \mathbb{1}_s\right) \\
&= \left(\langle \beta|_a \otimes \mathbb{1}_s\right) U_H^{\text{LCU}} \left(|\beta\rangle_a \otimes \mathbb{1}_s\right).
\end{aligned} \tag{2.47}$$

It is important to note that the BE U_H^{LCU} is defined with $|g\rangle_a = \textsc{prep}|0\rangle_a$.

The construction described above also applies to cases where the block encoding is based on Hamiltonian sparsity, since one can use sparse oracles to construct block encodings that are Hermitian and thus self-inverse [117]. In more general scenarios where this is not necessarily the case (e.g., as in [153]), the operator S' can always be constructed at the expense of introducing a single ancillary qubit and performing two controlled calls to U_H [109]. Note that while this procedure does not increase the asymptotic cost of preparing W_H, it does result in a significant increase of the overall prefactor. Importantly, for all methods used in this work, the BE's are also Hermitian, and the general procedure requiring an extra qubit was avoided. For completeness, the procedure for constructing W_H from an arbitrary U_H is presented in App. A.2.

2.2.2.4 QSP using Qubitization

Having just demonstrated that repeated applications of W_H construct block encodings of Chebyshev polynomials of the Hamiltonian, one natural strategy for constructing a polynomial approximation to e^{-itH} is as follows. One would first construct circuits with different

powers of W_H and then add these circuits with the appropriate coefficients using an LCU procedure. However, this would result in a block encoding approximating e^{-iHt} with a scale factor that grows with the degree of the polynomial approximation and cannot be made arbitrarily close to one. This is highly undesirable for the reasons previously discussed in Footnote 5, and below we discuss how this limitation can be overcome with the aid of QSP.

In order to use the operator W_H of the previous section to calculate the exponential of the Hamiltonian, we first note that within the subspaces $\mathrm{span}\{|g_\lambda\rangle_{as}, |g_\lambda^\perp\rangle_{as}\}$ the qubitized operator W_H can be diagonalized by defining the states

$$|\theta_\lambda^\pm\rangle_{as} \equiv \frac{1}{\sqrt{2}}\Big(|g_\lambda\rangle_{as} \pm i|g_\lambda^\perp\rangle_{as}\Big)\,. \tag{2.48}$$

Using Eq. (2.38) one can easily show that

$$W_H|\theta_\lambda^\pm\rangle_{as} = \mathrm{e}^{i\theta_\lambda^\pm}|\theta_\lambda^\pm\rangle_{as}\,, \tag{2.49}$$

with

$$\theta_\lambda^\pm = \mp\theta_\lambda\,, \tag{2.50}$$

where θ_λ is defined in Eq. (2.42).[6]

Given this form, one can now define a new operator $V_H(\phi)$, parameterised by ϕ, which can be constructed from the operator W_H. It requires one more ancillary qubit $|b\rangle \in \mathcal{H}_b$, such that the operator $V_H(\phi)$ acts on the Hilbert space $\mathcal{H}_{bas} \equiv \mathcal{H}_b \otimes \mathcal{H}_a \otimes \mathcal{H}_s$. It is defined by (see Fig. 2.6)

$$V_H(\phi) = (\mathrm{e}^{-i\phi Z_b/2} \otimes \mathbb{1}_{as}) W_H' (\mathrm{e}^{i\phi Z_b/2} \otimes \mathbb{1}_{as})\,, \tag{2.51}$$

where Z_b is the usual Pauli Z operator acting on the qubit $|b\rangle$, and W_H' is the W_H operator controlled on the qubit $|b\rangle$, which now acts on the Hilbert space $\mathcal{H}_{bas}$:

$$W_H' = |+\rangle_b\langle+|_b \otimes \mathbb{1}_{as} + |-\rangle_b\langle-|_b \otimes W_H\,. \tag{2.52}$$

One can show through explicit computation that

$$\big(\mathbb{1}_b \otimes \langle\theta_\pm^{\lambda'}|_{as}\big) V_H(\phi) \big(\mathbb{1}_b \otimes |\theta_\pm^{\lambda}\rangle_{as}\big) = V_H^{(\theta_\pm^\lambda)}(\phi)\,\delta_{\lambda\lambda'}\,, \tag{2.53}$$

with

$$\begin{aligned} V_H^{(\theta)}(\phi) &= \mathrm{e}^{i\frac{\theta}{2}} \begin{pmatrix} \cos\frac{\theta}{2} & -i\sin\frac{\theta}{2}\mathrm{e}^{-i\phi} \\ -i\sin\frac{\theta}{2}\mathrm{e}^{i\phi} & \cos\frac{\theta}{2} \end{pmatrix}_b \\ &= \mathrm{e}^{i\frac{\theta}{2}} \mathrm{e}^{-i\frac{\theta}{2}(X_b\cos\phi + Y_b\sin\phi)}\,, \end{aligned} \tag{2.54}$$

[6] The fact that W_H encodes information about the spectrum of the system via $\mathrm{spec}\, W_H^{(\lambda)} = \{\mathrm{e}^{\pm i\arccos(\lambda)}\}$ can be used for implementing a highly efficient version of Quantum Phase Estimation algorithm based using controlled calls to W_H [130].

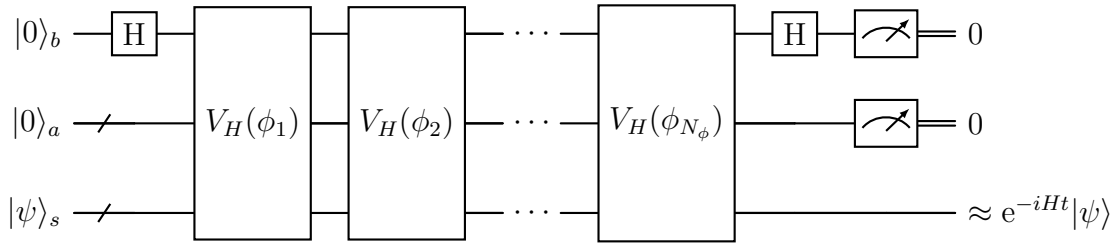

Figure 2.5: Quantum Signal Processing (QSP) circuit used for simulating time evolution, reproduced from [153]. The circuit prepares a block encoding for the approximation of the time evolution operator $e^{-iHt}|\psi\rangle$. Circuit for $V_H(\phi)$ is shown in Fig. 2.6. When all the ancillary qubits are measured in the zero states, the desired operator is applied to the system register $|\psi\rangle_s$. Without loss of generality, $|g\rangle_a$ is set to $|g\rangle_a = |0\rangle_a$.

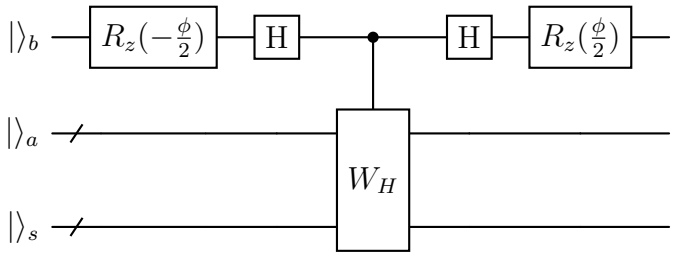

Figure 2.6: Circuit for $V_H(\phi)$, as defined is Eq. (2.51).

where X_b and Y_b denote the Pauli-X and Y matrices in $\mathcal{H}_b$.

What this construction has accomplished is to create a unitary operator $V_H(\phi)$ that is block diagonal in the original Hilbert space $\mathcal{H}_s$, with the entries of the 2×2 matrix in $\mathcal{H}_b$ given by functions of the eigenvalues of the Hamiltonian (through $\cos\theta_\lambda^\pm$), parameterized by some ϕ. By multiplying several of these operators in sequence (each with a different value of ϕ), one obtains an operator

$$V_H^{(\theta)}(\vec{\phi}) \equiv V_H^{(\theta)}(\phi_{N_\phi}) \dots V_H^{(\theta)}(\phi_1)\,, \tag{2.55}$$

that is still block diagonal in $\mathcal{H}_s$. The resulting operator is a 2×2 matrix in $\mathcal{H}_b$ which acquires the form of

$$V_H^{(\theta)}(\vec{\phi}) = \mathcal{A}(\theta)\mathbb{1}_b + i\mathcal{B}(\theta)X_b + i\mathcal{C}(\theta)Y_b + i\mathcal{D}(\theta)Z_b\,, \tag{2.56}$$

where the functions $\mathcal{A}(\theta) \dots \mathcal{D}(\theta)$ are all parameterized by the angles $\vec{\phi} = (\phi_1, \dots, \phi_{N_\phi})$ and, similarly to Eq. (2.44), can be expressed in terms of Chebyshev polynomials of λ upon substituting $\theta = \arccos(\lambda)$.

The classes of functions $\mathcal{A}(\theta)\ldots\mathcal{D}(\theta)$ that can be realized by the above procedure are described by the theory of Quantum Signal Processing. In particular, it was shown in [154] that one can construct the vector of angles $\vec{\phi}$ to obtain functions $\mathcal{A}(\theta)$ and $\mathcal{C}(\theta)$ of the form

$$\begin{aligned} |\mathcal{A}^2(\theta)| + |\mathcal{C}^2(\theta)| \leq 1\,, \qquad \mathcal{A}(0) = 1\,, \\ \mathcal{A}(\theta) = \sum_{k=0}^{N_\phi/2} a_k \cos(k\theta)\,, \\ \mathcal{C}(\theta) = \sum_{k=1}^{N_\phi/2} c_k \sin(k\theta)\,. \end{aligned} \tag{2.57}$$

The two other functions, $\mathcal{B}(\theta)$ and $\mathcal{D}(\theta)$ are not relevant to the calculation since projecting onto $|+\rangle_b$ leaves only $\mathcal{A}(\theta)$, $\mathcal{C}(\theta)$. Reversing this logic, one can start from some desired functions $\mathcal{A}(\theta)$, $\mathcal{C}(\theta)$, and determine phases $\vec{\phi}$ such that $V_H^{(\theta)}(\vec{\phi})$ reproduces these functions (approximately). This can be accomplished using an efficient classical procedure given in [154].[7]

We omit now a few final details, which can be found in [109], and simply state the relevant results. As can be confirmed from Eq. (2.56), letting the operator $V_H^{(\theta)}(\vec{\phi})$ act on the state $|+\rangle_b$ and post-selecting on $\langle+|_b$ picks out the function

$$\begin{aligned} \langle+|_b\langle g|_a\langle\lambda'|_s V_H(\vec{\phi})|+\rangle_b|g\rangle_a|\lambda\rangle_s \\ = \delta_{\lambda\lambda'}\left[A(\lambda) + iC(\lambda)\right]. \end{aligned} \tag{2.58}$$

The functions $A(\lambda)$ and $C(\lambda)$ are directly related to the functions $\mathcal{A}(\theta)$, $\mathcal{C}(\theta)$, with the detailed relation given in [109]. The success probability for measuring the system in the state $|+\rangle_b|g\rangle_a$ for a given value λ depends on the functions and is given by

$$p \geq |A^2(\lambda)| + |C^2(\lambda)|\,. \tag{2.59}$$

The final task is to determine which functions $A(\lambda)$ and $C(\lambda)$ (and therefore which phases $\vec{\phi}$) approximate the time evolution operator with uncertainty bounded by ε,

$$\left|\mathrm{e}^{i\lambda\tilde{t}} - [A(\lambda) + iC(\lambda)]\right| < \varepsilon\,. \tag{2.60}$$

Since $A(\lambda)$ and $C(\lambda)$ can be expressed in terms of Chebyshev polynomials, they can be used for approximating the real and imaginary part of $\mathrm{e}^{i\lambda\tilde{t}}$ through the rapidly converging Jacobi-Anger expansion [108]:

$$A(\lambda) = J_0(t) + 2\sum_{k\ \text{even}>0}^{N_\phi/2} (-1)^{k/2} J_k(\tilde{t}) T_k(\lambda)\,, \tag{2.61a}$$

[7] A recent version of QSP, *generalized QSP*, significantly improves the cost of calculating phases $\vec{\phi}$ [110]. and is considered in Chapter 3

$$C(\lambda) = 2 \sum_{k \text{ odd}>0}^{N_\phi/2} (-1)^{(k-1)/2} J_k(\tilde{t}) T_k(\lambda) \,, \tag{2.61b}$$

(note that N_ϕ is always chosen to be even). In this case, the circuit in Fig. 2.5 would implement the approximate BE of the evolution operator e^{-iHt}. The number of phases needed for a given value of ε is bounded by [108]

$$N_\phi = \mathcal{O}\big(\tilde{t} + \log(1/\varepsilon)\big) = \mathcal{O}\big(\alpha t + \log(1/\varepsilon)\big) \,. \tag{2.62}$$

Note also that Eqs. (2.59) and (2.60) imply that

$$p \geq 1 - \mathcal{O}(\varepsilon) \,, \tag{2.63}$$

since the time evolution operator has unit magnitude. Note that due to the $\log(1/\epsilon)$ scaling, one can require the success probability to be exponentially close to one for only a polynomial cost (see also Footnote 5).

It is worth mentioning that the success rate of QSP is always $\mathcal{O}(1)$, independent of the success rate $\mathcal{O}(1/\alpha)$ of the BE U_H used. This is due to the fact that we have used QSP to implement a unitary function (e^{-iHt}) which has unit norm. This implies that $A(\lambda)^2 + C(\lambda)^2 \approx 1$, which together with Eq. (2.59) gives a success probability that is close to one. While this is a desirable feature, the cost of a poor success probability of the BE is now reflected in the number of phases N_ϕ required to approximate the time evolution operator to an error ϵ, as seen by Eq. (2.62). For this reason, significant savings can be achieved by performing dedicated studies to minimize α. The following theorem summarizes the content of this section

Theorem 9. *Let U_H be a an (α, k)-block-encoding of a Hamiltonian H acting on a Hilbert space $\mathcal{H}_s$. Furthermore, let t and ε be positive real values. Then there is a quantum circuit that produces a $\big(1 \pm \mathcal{O}(\varepsilon), k+2\big)$-block-encoding of the time evolution operator e^{-iHt} up to exponentiation error ε, and possesses a complexity of*

$$\mathcal{O}\big(\alpha t + \log(1/\varepsilon)\big) \,, \tag{2.64}$$

relative to the oracle U_H. The probability for this BE to succeed is

$$p > 1 - \mathcal{O}(\varepsilon) \,. \tag{2.65}$$

Proof. See for instance [109]. The fact that the BE is $(\tilde{\alpha}, \tilde{k}) = (1 \pm \mathcal{O}(\varepsilon), k+2)$ can be easily understood. First, the BE of the time evolution operator requires two additional qubits, one ($|q\rangle$) for the qubitization and one ($|b\rangle$) to add the functional dependence on ϕ, giving $\tilde{k} = k + 2$. The fact that $\tilde{\alpha} = 1 \pm \mathcal{O}(\varepsilon)$ follows from the fact that the operator we are block-encoding is a unitary operator, and our approximation is good up to error ε. Since the norm of a unitary operator is 1, this gives that both $\tilde{\alpha}$ and the success probability are unity (up to $\mathcal{O}(\varepsilon)$). □

2.2.3 HHKL

The results presented so far (PFs and QSP) can be applied to general site-local Hamiltonians (see Definition 2). For geometrically-site-local Hamiltonians (see Definition 3) one can use the fact that information from one area of the lattice can only spread to areas not directly connected to it through repeated actions of the time evolution operators, which implies that the speed with which this information spreads is bounded. This is formalized by the Lieb-Robinson bound [155]. This idea was used in a paper by Haah, Hastings, Kothari and Low (HHKL) [128] to devise a scheme for simulating the time evolution of geometrically-site-local Hamiltonians.

The main ingredient to the HHKL algorithm is that the time evolution of a large quantum system with geometrically local Hamiltonian can be obtained from the time evolution of two smaller systems that share some overlap, which has to be of at least the size given by the ball of size $\mathcal{O}(1)$ used in Definition 3. Thus, the system Hilbert space $\mathcal{H}_s$ is divided into three parts, $\mathcal{H}_{a,b,c}$ such that $\mathcal{H}_s = \mathcal{H}_a \otimes \mathcal{H}_b \otimes \mathcal{H}_c$. In the discussion of [128] is assumed that the local Hamiltonian H_J such that $||H_J|| \leq 1$. Thus, in the remainder we work with

$$\tilde{H}_J = \frac{H_J}{||H_J||}\,, \qquad \tilde{t} = ||H_J||t\,. \tag{2.66}$$

First, one uses that the evolution of a system for some long time $\tilde{t}$ can be split into the product of unitaries over shorter times $\Delta t = \mathcal{O}(1)$:

$$U(\tilde{t},0) = U(\tilde{t},\tilde{t}-\Delta t)\cdots U(\Delta t,0)\,. \tag{2.67}$$

The total number of such unitaries is $\mathcal{O}(\tilde{t})$.

Lemma 6 of [128] states that the time evolution $U_s(\Delta t)$ of a Hamiltonian on the whole system over time Δt can be obtained by multiplying together the forward time evolution in the system $\mathcal{H}_{a\cup b} = \mathcal{H}_a \otimes \mathcal{H}_b$ and $\mathcal{H}_{b\cup c} = \mathcal{H}_b \otimes \mathcal{H}_c$ and the backwards evolution in the system $\mathcal{H}_b$ as

$$U_{a\cup b\cup c}(\Delta t) = U_{a\cup b}(\Delta t)U_b^\dagger(\Delta t)U_{b\cup c}(\Delta t) + \mathcal{O}(\delta)\,, \tag{2.68}$$

where the error due to this approximation is given by

$$\delta \equiv \delta(a,b,c) = \mathcal{O}\left(\mathrm{e}^{-\mu\,\mathrm{d}(a,c)} \sum_{X\in \mathrm{BD}(ab,c)} ||\tilde{H}_X||\right)\,. \tag{2.69}$$

For now we assume that each of the $U(\Delta t)$ on the RHS of Eq. (2.68) is implemented exactly. In Eq. (2.69) $\mu > 0$ is a constant, $\mathrm{d}(a,c)$ is the smallest distance between any point in a and any point in c, and the norm $||\tilde{H}_J||$ is summed over all terms that are on the boundary of $a\cup b$ and c, indicated by $\mathrm{BD}(ab,c)$. The uncertainty is therefore exponentially suppressed in the distance between the regions a and c. This implies that this approximation works well for large enough distance $\ell \equiv \mathrm{d}(a,c)$.

This result can be applied recursively, splitting the evolution into smaller and smaller pieces, each of which just have to be larger than the ball of size $\sim \mathcal{O}(1)$. If the size of the original system is given by L^{d} (where d represents the dimension of the Euclidean space we consider with size L in each dimension) and the final pieces are each of size ℓ^{d}, one ends up with

$$m = \mathcal{O}\left(\frac{L^{\mathsf{d}}}{\ell^{\mathsf{d}}}\right) \tag{2.70}$$

individual pieces. Taking into account the $\mathcal{O}(\tilde{t})$ evolution operators we started from, the final error $\mathcal{O}(\varepsilon)$ can be reached by requiring

$$\delta = \mathcal{O}\left(\frac{\varepsilon}{\tilde{t}\, m}\right). \tag{2.71}$$

We see that the sum of the norms of the boundary terms appearing in Eq. (2.69) is $\mathcal{O}(L^{\mathsf{d}-1})$ and thus from Eqs. (2.69) and (2.71) one finds it suffices to choose

$$\ell = \Theta\left(\log\left(\frac{L^{\mathsf{d}}\tilde{t}}{\varepsilon}\right)\right). \tag{2.72}$$

One can then estimate the required resources using the fact that one needs to compute

$$\mathcal{O}(\tilde{t}m) = \mathcal{O}\left(L^{\mathsf{d}}\tilde{t}\log^{-\mathsf{d}}\left(\frac{L^{\mathsf{d}}\tilde{t}}{\varepsilon}\right)\right) \tag{2.73}$$

time evolution operators on a system of size $\mathcal{O}(\ell^{\mathsf{d}})$.[8] The cost and the associated additional error for implementing each of these "local" pieces depends on the algorithm used for their implementation. In particular, if an algorithm with cost polynomial in ℓ and polylogarithmic in $1/\varepsilon$ is used (such as Qubitization), then the cost of the HHKL algorithm remains $\mathcal{O}(L^{\mathsf{d}}\tilde{t}\,\mathrm{polylog}(L^{\mathsf{d}}\tilde{t}/\varepsilon))$. Implementing the time evolution of local blocks using PF would lead to an algorithm whose asymptotic cost dependence on ε is given by that of the employed PF method, i.e., no longer polylogarithmic in ε. This can make sense if, for a given value of ε, the PF-based method leads to lower gate counts as compared to a nearly-optimal implementation – for a system of size ℓ but not of size L [105].

The following theorem summarizes the results of this section

Theorem 10. *Let Λ be a lattice of qubits embedded in d-dimensional Euclidean space of size L in each dimension, and let $\hat{H}$ be a geometrically-local Hamiltonian, normalized such*

[8] Note that decreasing ε is logarithmically decreasing the number of evolution operators. This number clearly has to be larger than unity, meaning ε cannot be exponentially small with system size.

that each local term has a norm of order unity. Furthermore, let $\tilde{t}$ and ε be positive real values. Then there exists an algorithm that approximates $\mathrm{e}^{-i\hat{H}\tilde{t}}$ to exponentiation error ε, which requires[9]

$$\mathcal{O}\left(\frac{L^{\mathsf{d}}\tilde{t}}{\log^{\mathsf{d}}\left(L^{\mathsf{d}}\tilde{t}/\varepsilon\right)}\right) \tag{2.74}$$

calls to a helper function that implements the time evolution for a system of size

$$\ell = \Theta\left(\log\left(L^{\mathsf{d}}\tilde{t}/\varepsilon\right)\right), \tag{2.75}$$

with error

$$\delta = \mathcal{O}\left(\frac{\varepsilon \log^{\mathsf{d}}(L^{\mathsf{d}}\tilde{t}/\varepsilon)}{L^{\mathsf{d}}\tilde{t}}\right) \tag{2.76}$$

Proof. This result has been proven in [128]. □

2.3 Simulating Time Evolution on the Lattice

In this section, we apply the algorithms and methods of Sec. 2.2 to the case of Hamiltonians defined in Sec. 2.1.1. In the following sections (organized by primary simulation technique), we present a series of theorems that quantify the asymptotic gate complexities required to perform approximate time evolution for these Hamiltonians. For each of these sections, we consider two case: (1) $\mathcal{O}(1)$-site-local HLFTs defined on $\mathcal{O}(1)$-many sites, which behave very similarly to generic quantum systems, and will serve as a review of the relationship between our notation and other common notations (2) k-site-local and k-geometrically-site-local Hamiltonians, whose simulation can roughly be categorized as being composed of simulation methods for $\mathcal{O}(1)$-site-local Hamiltonians that are "glued" via PFs, LCU, or HHKL, depending on the flavor of locality whose preservation is desired.

2.3.1 Product Formulas

We begin by presenting the two following lemmas that we will leverage to investigate the asymptotic scaling of PFs for the lattice Hamiltonians we consider in this work. The first Lemma relates the nested commutator $\tilde{\alpha}$ presented in Sec. 2.2.1 to the norm (and induced norm) of the Hamiltonian:

Lemma 11. *Let H be a k-site-local Hamiltonian, as defined in Definition 2. Then the nested commutator $\tilde{\alpha} \equiv \sum_{J_1,J_2,\ldots,J_{p+1}\in\mathcal{S}} |||[H_{J_{p+1}},\ldots[H_{J_2},J_{J_1}]]||$ can be asymptotically bounded as*

$$\tilde{\alpha} = \mathcal{O}\big(|||H|||_1^p ||H||_1\big). \tag{2.77}$$

[9] We assume that the error ε is such that the number of calls to the helper function exceeds one.

Proof. This result has been proven in Appendix G of [106]. □

The next Lemma presents a constraint on the unitaries used to approximate each of the exponentials that constitute the PF in order to (asymptotically) maintain some desired overall approximation error to the true time evolution operator. The PF introduced in Eq. (2.14) is used as an approximation to e^{-iHt}, and another approximation $U_\upsilon^{(J)}(t/r)$ is used to approximate each term of the form $\mathrm{e}^{-ita_\upsilon^{(J)}H_J/r}$:

Lemma 12. *Let H be a site-local Hamiltonian as defined by Definition 2. Let $S_p(t)$ be a p^{th} order PF for e^{-iHt} as given in Eq. (2.14), and let the unitary $U_\upsilon^{(J)}(t/r)$ be an approximation to $\mathrm{e}^{-ita_\upsilon^{(J)}H_J/r}$ for all choices of allowed pairs (υ, J), with $\upsilon \in \{1,\ldots,\Upsilon(p)\}$ and $J \in \mathcal{S}$. Define the PF $\tilde{S}_p(t)$ as the PF for e^{-iHt} obtained by replacing $\mathrm{e}^{-ita_\upsilon^{(J)}H_J/r}$ in Eq. (2.14) with $U_\upsilon^{(J)}(t/r)$*

$$\tilde{S}_p(t/r) \equiv \prod_{\upsilon=1}^{\Upsilon(p)} \mathcal{P}_\upsilon \left(\prod_{J\in\mathcal{S}} U_\upsilon^{(J)}(t/r) \right) . \tag{2.78}$$

If

$$\left\| U_\upsilon^{(J)}(t/r) - \mathrm{e}^{-ita_\upsilon^{(J)}H_J/r} \right\| = \mathcal{O}\left(\frac{\varepsilon}{\mathsf{N_H} r \Upsilon(p)} \right), \tag{2.79}$$

it is true that

$$\left\| \tilde{S}_p(t) - \mathrm{e}^{-iHt} \right\| = \mathcal{O}(\varepsilon) . \tag{2.80}$$

Proof. See App. A.1.1 □

We use these lemmas to present the following results for general site-local, $\mathcal{O}(1)$, and geometrically site-local Hamiltonians (respectively).

Theorem 13. *Let H be a site-local Hamiltonian as defined in Definition 2 or Definition 3. To simulate the time evolution operator e^{-iHt} up to exponentiation error $\mathcal{O}(\varepsilon)$ it suffices to choose a Trotter number r such that:*

$$r = \mathcal{O}\left(\mathsf{N_{ind}} \mathsf{N_H}^{1/p} \frac{(\mathsf{N_P} \mathsf{c_P} t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}} \right) . \tag{2.81}$$

Furthermore, if the Trotter number r is chosen with the tight asymptotic scaling

$$r = \Theta\left(\mathsf{N_{ind}} \mathsf{N_H}^{1/p} \frac{(\mathsf{N_P} \mathsf{c_P} t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}} \right), \tag{2.82}$$

then the associated elementary gate count χ can be asymptotically bounded as

$$\chi = \mathcal{O}\left(\mathsf{N}_{\mathsf{ind}}\mathsf{N}_{\mathsf{P}}\frac{(\mathsf{N}_{\mathsf{H}}\mathsf{N}_{\mathsf{P}}\mathsf{c}_{\mathsf{P}}t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right). \tag{2.83}$$

Proof. See App. A.1.2 □

Theorem 14. *Let H be an $\mathcal{O}(1)$ Hamiltonian as defined in Definition 4 Let $S_p(t)$ be a p^{th} order PF as given by Eq. (2.14) and Eq. (2.15). Then to simulate the time evolution operator e^{-iHt} up to exponentiation error $\mathcal{O}(\varepsilon)$ it suffices to choose a Trotter number r such that:*

$$r = \mathcal{O}\left(\frac{(\mathsf{N}_{\mathsf{P}}\mathsf{c}_{\mathsf{P}}t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right). \tag{2.84}$$

Furthermore, if the Trotter number r is chosen with the tight asymptotic scaling

$$r = \Theta\left(\frac{(\mathsf{N}_{\mathsf{P}}\mathsf{c}_{\mathsf{P}}t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right), \tag{2.85}$$

then the associated elementary gate count χ can be asymptotically bounded as

$$\chi = \mathcal{O}\left(\frac{\mathsf{N}_{\mathsf{P}}\,(\mathsf{N}_{\mathsf{P}}\mathsf{c}_{\mathsf{P}}t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right). \tag{2.86}$$

Proof. This result is obtained immediately by setting $\mathsf{N}_{\mathsf{H}} \sim \mathsf{N}_{\mathsf{ind}} = \mathcal{O}(1)$ in Theorem 13. □

Theorem 15. *Let H be a geometrically site-local Hamiltonian as defined in Definition 3. To simulate the time evolution operator e^{-iHt} up to exponentiation error $\mathcal{O}(\varepsilon)$ it suffices to choose a Trotter number r such that:*

$$r = \mathcal{O}\left(\mathsf{N}_{\Lambda}{}^{1/p}\frac{(\mathsf{N}_{\mathsf{P}}\mathsf{c}_{\mathsf{P}}t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right). \tag{2.87}$$

Furthermore, if the Trotter number r is chosen with the tight asymptotic scaling

$$r = \Theta\left(\mathsf{N}_{\Lambda}{}^{1/p}\frac{(\mathsf{N}_{\mathsf{P}}\mathsf{c}_{\mathsf{P}}t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right), \tag{2.88}$$

then the associated elementary gate count χ can be asymptotically bounded as

$$\chi = \mathcal{O}\left(\mathsf{N}_{\mathsf{P}}\frac{(\mathsf{N}_{\Lambda}\mathsf{N}_{\mathsf{P}}\mathsf{c}_{\mathsf{P}}t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right). \tag{2.89}$$

Proof. This result follows directly from Theorem 13 making the replacements $\mathsf{N}_{\mathsf{ind}} = \mathcal{O}(1)$ and $\mathsf{N}_{\mathsf{H}} = \mathcal{O}(\mathsf{N}_{\Lambda})$ as given in Eqs. (2.11) and (2.12). □

2.3.2 Quantum Signal Processing

As discussed in detail in Sec. 2.2, QSP provides a method to take the BE of a given Hamiltonian and obtain the approximation to the BE of its time evolution operator. In this section, we derive the general complexity of using QSP to implement the time evolution operator of the full system Hamiltonian.

We first establish the asymptotic cost of using LCU to prepare a BE for the lattice Hamiltonians of interest

Lemma 16. *Let H be a site-local Hamiltonian as as in Definition 2. There exists a quantum circuit acting on $\mathsf{n} \equiv \mathsf{N}_\Lambda \mathsf{n}_\mathsf{q}$-many qubits together with $k = \mathcal{O}\left(\log(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P})\right)$-many ancillas that produces an $\left(\mathcal{O}\left(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P}\mathsf{c}_\mathsf{P}\right), k\right)$-block-encoding of H on the n qubits, and requires a total of $\mathcal{O}\left(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P}\log(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P})\right)$-many elementary quantum gates.*

Proof. This follows directly from Lemma 8, if we observe that the total number of Pauli strings is given by $\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P}$ and the sum of the coefficients of all Pauli strings is bounded by $\beta < \mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P}$. □

Given this result now allows us to state the central QSP result we use in this work:

Theorem 17. *Let H be a site-local Hamiltonian as defined by Definition 2, and let t and ε be positive values. Then there exists a quantum circuit that serves as a $\left(1 \pm \mathcal{O}(\varepsilon), \mathcal{O}(\log(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P}))\right)$-block-encoding for e^{-iHt} up to exponentiation error ε, and requires an elementary gate count of*

$$\chi = \mathcal{O}\Big[\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P}\log(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P})\left(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P}\mathsf{c}_\mathsf{P}t + \log(1/\varepsilon)\right)\Big]. \tag{2.90}$$

Proof. This result follows from Lemma 16 and Theorem 9. From Lemma 16 one obtains that LCU give a $\left(\mathcal{O}\left(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P}\mathsf{c}_\mathsf{P}\right), \mathcal{O}\left(\log(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P})\right)\right)$ BE of the Hamiltonian. Using this result in Theorem 9 one finds that QSP gives a $\left(1 \pm \mathcal{O}(\varepsilon), \mathcal{O}(\log(\mathsf{N}_\mathsf{H}\mathsf{N}_\mathsf{P}))\right)$-block-encoding for e^{-iHt}. Using the resource scaling from Theorem 9, with the scale factor and resource scaling of LCU from Lemma 16, gives the resource requirements stated in the theorem. □

Remark 18. Note that we have assumed that the block-encoding is formed from a sum over all Pauli strings appearing in the Hamiltonian. When more structure is known about a given field theory, the factor in the complexity due to the block-encoding may be reduced.

Theorem 17 allows us to immediately present the following results for $\mathcal{O}(1)$ and geometrically site-local lattice Hamiltonians:

Theorem 19. *Let H be an $\mathcal{O}(1)$ Hamiltonian as defined in Definition 4, and let t and ε be positive values. Then there exists a quantum circuit that serves as a $\left(1 \pm \mathcal{O}(\varepsilon), \mathcal{O}(\log(\mathsf{N}_\mathsf{P}))\right)$-block-encoding for e^{-iHt} up to exponentiation error ε, and requires an elementary gate count of*

$$\chi = \mathcal{O}\Big[\mathsf{N}_\mathsf{P}\log\mathsf{N}_\mathsf{P}\left(\mathsf{N}_\mathsf{P}\mathsf{c}_\mathsf{P}t + \log(1/\varepsilon)\right)\Big]. \tag{2.91}$$

Proof. This follows directly from Theorem 17 by taking the limit $\mathsf{N_H} = \mathcal{O}(1)$. □

Theorem 20. *Let H be a geometrically site-local Hamiltonian as defined by Definition 2, and let t and ε be positive values. Then there exists a quantum circuit that serves as a $\big(1 \pm \mathcal{O}(\varepsilon), \mathcal{O}(\log(\mathsf{N_H N_P}))\big)$-block-encoding for e^{-iHt} up to exponentiation error ε, and requires an elementary gate count of*

$$\chi = \mathcal{O}\Big[\mathsf{N_\Lambda N_P} \log(\mathsf{N_\Lambda N_P}) \left(\mathsf{N_\Lambda N_P c_P} t + \log(1/\varepsilon)\right)\Big]. \tag{2.92}$$

Proof. This result follows directly from Theorem 17 making the replacements $\mathsf{N_H} = \mathcal{O}(\mathsf{N_\Lambda})$ as given in Eq. (2.12). □

We now consider the case of combining a PF with QSP. We will use Lemma 12 to write a PF for e^{-iHt}, but then use QSP for each $U_v^{(J)}(t/r)$, which serves as an approximation to each term of the form $\exp\left(-ita_v^{(J)} H_J/r\right)$ appearing in the PF.

Theorem 21. *Let H be a site-local Hamiltonian as defined by Definition 2, and let t and ε be positive values. By using QSP (see Theorem Theorem 19) to simulate each exponential term $\mathrm{e}^{-ia_v^{(J)} H_J t/r}$ appearing in the PF $S_p(t/r)$ to error $\varepsilon/(r\mathsf{N_H})$, and taking*

$$r = \Theta\left(\mathsf{N_{ind}} \mathsf{N_H}^{1/p} \frac{(\mathsf{N_P c_P} t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right), \tag{2.93}$$

we obtain a quantum circuit that serves as a $(1 \pm \mathcal{O}(\varepsilon), \log(\mathsf{N_H N_P}))$-block-encoding for the full time evolution operator e^{-iHt} to exponentiation error $\mathcal{O}(\varepsilon)$. The gate count for this circuit is given by

$$\chi = \mathcal{O}\Big[\mathsf{N_{ind}} \left(\mathsf{N_H N_P c_P} t\right)^{1+\frac{1}{p}} \varepsilon^{-\frac{1}{p}} \mathsf{N_P} \log \mathsf{N_P} \log(\mathsf{N_{ind} N_H N_P c_P} t/\varepsilon)\Big]. \tag{2.94}$$

Proof. See App. A.1.3 □

Remark 22. The results above were for a general site-local Hamiltonian. To obtain the results for a geometrically site-local Hamiltonian one makes the replacement $\mathsf{N_{ind}} = \mathcal{O}(1)$ and $\mathsf{N_H} = \mathcal{O}(\mathsf{N_\Lambda})$ from Eqs. (2.11) and (2.12).

2.3.3 HHKL

As discussed in Sec. 2.2.3 the HHKL technique is a method of "gluing" together implementations of a time evolution operator on a "local" lattice with size ℓ^{d}. In this section we consider HHKL with two approaches for the implementation of the local evolution operator, first using a PF and then QSP.

Theorem 23. *Let H be a geometrically site-local Hamiltonian as defined by Definition 3, and let t and ε be positive values. Using the HHKL algorithm with a PF for the helper function gives rise to an algorithm that approximates the time evolution operator e^{-iHt} with exponentiation error requiring a gate complexity of*

$$\chi = \mathcal{O}\left(\mathsf{N}_\mathsf{P}\left(\mathsf{N}_\Lambda \mathsf{N}_\mathsf{P} \mathsf{c}_\mathsf{P} t\right)^{1+\frac{1}{p}} \varepsilon^{-\frac{1}{p}}\right) . \tag{2.95}$$

Proof. See App. A.1.4 □

Theorem 24. *Let H be a geometrically site-local Hamiltonian as given by Definition 3, and let t and ε be positive values. Using the HHKL algorithm with QSP for the helper function gives rise to an algorithm that approximates the time evolution operator e^{-iHt} with exponentiation error requiring a gate complexity of*

$$\chi = \mathcal{O}\left(\mathsf{N}_\mathsf{P}{}^2 \mathsf{N}_\Lambda \mathsf{c}_\mathsf{P} t \log^{\mathsf{d}}\left(\mathsf{N}_\Lambda t/\varepsilon\right) \log\left(\mathsf{N}_\mathsf{P} \log\left(\mathsf{N}_\Lambda t/\varepsilon\right)\right)\right) . \tag{2.96}$$

Proof. See App. A.1.5 □

2.4 Lattice Quartic Theory

In this section, we specialize to a particular model of interacting quantum field theories: the lattice scalar field theory. From the computational perspective, this theory is similar to lattice gauge theories and so intuition can be drawn from its consideration. Furthermore, it is often used as a proverbial example in the studies of quantum simulation algorithms [52, 156, 157], and so detailed comparisons with past work can be performed.

In this section, we review the Hamiltonian of the theory and our procedure for digitizating the bosonic field. We derive the expressions for the asymptotic dependence of parameters of the qubit Hamiltonian (N_P, c_P, N_H, etc.) on the parameters of the physical model. These expressions can then be substituted into the complexities quoted in Tables 2.3 and 2.4 to thus determine the asymptotic gate complexities for implementing time evolution. To determine the coefficients in these asymptotic gate complexities, we explicitly construct circuits that implement time evolution for a single site. We first compare the gate cost of constructing controlled calls to the walk operator W_H obtained using two LCU based BE procedures. Then we compare the gate costs for time evolution simulated via 4th order PF and QSP.

2.4.1 Introduction

The model we consider is defined on a hypercubic lattice Λ of spatial dimension d with lattice spacing a. Assuming the simplest finite-difference approximation to the spatial deriva-

tive $(\nabla_\mu \hat{\varphi})^2$, the Hamiltonian is given by[10]

$$\hat{H}_{\text{quartic}} = \hat{H}_\varphi + \hat{H}_\pi \,, \tag{2.97a}$$

$$\hat{H}_\varphi = a^{\mathsf{d}} \left[\sum_{i=1}^{\mathsf{N}_\Lambda} \frac{m^2}{2} \hat{\varphi}_i^2 + \frac{\lambda}{4!} \hat{\varphi}_i^4 + \sum_{\langle ij \rangle} \frac{(\hat{\varphi}_j - \hat{\varphi}_i)^2}{2a^2} \right] , \tag{2.97b}$$

$$\hat{H}_\pi = a^{\mathsf{d}} \sum_{i=1}^{\mathsf{N}_\Lambda} \frac{1}{2} \hat{\pi}_i^2 \,, \tag{2.97c}$$

where the indices i and j label lattice sites, and the notation $\langle ij \rangle$ runs over all distinct links joining nearest neighbors in the lattice. In the remainder of this work, we will denote operators with a hat to differentiate between their digitized eigenvalues.

The Hamiltonian in Eq. (2.97) is 2-site-geometrically-local, which implies that $\mathsf{N}_{\mathsf{ind}} = \mathcal{O}(1)$. We see that the Hamiltonian is a sum of $\mathsf{N}_{\mathsf{H}} = \mathcal{O}(\mathsf{N}_\Lambda)$ terms.

The operators $\hat{\pi}_i$ and $\hat{\varphi}_i$ are conjugate operators satisfying $[\hat{\varphi}_j, \hat{\pi}_i] = i a^{-\mathsf{d}} \delta_{ij}$. This relation implies that the so-called field basis, in which the φ_i operators are diagonal, is related to the so-called momentum basis, where the π_i operators are diagonal, by the usual Fourier transform. The parameters m and λ are bare quantities, and are in general complicated non-perturbative functions of the lattice spacing, *i.e.* $m = m(a)$ and $\lambda = \lambda(a)$. In principle, this functional dependence could be included to determine the asymptotic complexity of a realistic physical calculation in terms of the lattice spacing. In order to keep our expressions as general as possible, however, we choose to leave the bare parameters m and λ as free variables in our quoted asymptotic complexities; this allows our results to be adapted in a straightforward way to any specific physical use case. For the remainder of this work, we use units where the lattice spacing $a = 1$.

The above Hamiltonian, although discretized on the lattice, is still an unbounded operator acting on an infinite dimensional Hilbert space. In order to obtain concrete time-evolution operators whose action can actually be simulated on a system of qubits, we must first truncate this Hamiltonian to a finite-dimensional operator. This is achieved by a procedure known as *digitization*, in which the Hilbert space at each lattice site is represented using n_{q} qubits; the details of our digitization procedure are given in the next section.

2.4.1.1 Digitization

The basic idea is that we wish to construct a sequence of matrix operators $H_{\mathsf{n}_{\mathsf{q}}}$ that can act on n_{q} qubits at a time, and whose spectrum converges to any prescribed single-site Hamiltonian H as $\mathsf{n}_{\mathsf{q}} \to \infty$. There are numerous ways to proceed with such a construction [98, 99, 100], but the one we follow is outlined in [97, 158].

[10] The following discussion can be generalized to the case of higher-order finite-difference approximations to the derivative operator. One common approach is to first perform integration by parts to get the term $\hat{\varphi} \nabla^2 \hat{\varphi}$ and then use a symmetric approximation for the second derivative.

The strategy is to *digitize* the field on a given lattice site into $N = 2^{\mathsf{n_q}}$ allowed field values

$$\varphi \in \{-\varphi_{\max}, -\varphi_{\max} + \delta\varphi, \ldots, \varphi_{\max} - \delta\varphi, \varphi_{\max}\}\,, \tag{2.98}$$

where the spacing $\delta\varphi$ is given by

$$\delta\varphi = \frac{2\varphi_{\max}}{N-1}\,. \tag{2.99}$$

This restriction of the field values is to be understood as a restriction of the Hilbert space at the given lattice site, identifying the eigenvalues of the restricted field operator precisely with this set of allowed field values. In other words, the digitized field operator in the field basis at site i becomes the diagonal matrix

$$\hat{\varphi}_i = \mathrm{diag}\left(\varphi_{\min}, \varphi_{\min} + \delta\varphi, \ldots, \varphi_{\max}\right)_i\,. \tag{2.100}$$

This single-site field operator acts nontrivially only on the $\mathsf{n_q}$-qubit Hilbert space attributed to the site labeled by i. Note that all lattice sites use the same digitized φ values.

Because $\hat{\pi}_i$ and $\hat{\varphi}_i$ are conjugate operators, the momentum operator in the field basis can be written as

$$\begin{aligned} \hat{\pi}_i &= \mathrm{FT}\cdot\hat{\pi}_i^{(p)}\cdot\mathrm{FT}^{-1}\,, \\ \hat{\pi}_i^{(p)} &= \mathrm{diag}\left(-\pi_{\max}, -\pi_{\max} + \delta\pi, \ldots, \pi_{\max}\right)_i\,, \end{aligned} \tag{2.101}$$

where the superscript (p) indicates the operator is represented in the momentum basis, FT denotes the *symmetric Fourier transform* [97] and

$$\pi_{\max} = \frac{\pi}{\delta\varphi}, \quad \delta\pi = \frac{2\pi_{\max}}{N-1}\,. \tag{2.102}$$

Note that this choice implies anti-periodic boundary conditions in field space. Using this definition, the momentum part of the Hamiltonian can be written as

$$\hat{H}_\pi = \sum_{i=1}^{\mathsf{N_\Lambda}} \frac{1}{2}\mathrm{FT}_i\cdot\left(\hat{\pi}_i^{(p)}\right)^2\cdot\mathrm{FT}_i^{-1}\,. \tag{2.103}$$

The final component of the digitization strategy is the choice of $\varphi_{\max}$. It was shown in [97] that, for a given value of $\mathsf{n_q}$ and λ, there is an optimal choice of $\varphi_{\max}$ that minimizes the digitization error of the low lying spectrum of the digitized Hamiltonian. The procedure for determining $\varphi_{\max}$ involved scanning over many values to minimize the digitization error. While this procedure is in general possible, an explicit formula for the optimal value of $\varphi_{\max}$ for the $\lambda = 0$ case was given in [98, 100, 159]. As argued in [97], because $\lambda \neq 0$ results in a smooth deformation of the wavefunction, this choice of $\varphi_{\max}$ is expected to work well even in the $\lambda \neq 0$ case. Therefore, we use the value

$$\varphi_{\max} = 2^{\mathsf{n_q}}\sqrt{\frac{2\pi}{2^{\mathsf{n_q}+1}+1}}\,, \tag{2.104}$$

as in [91].

Altogether, these definitions allow for the construction of $2^{\mathsf{n_q}} \times 2^{\mathsf{n_q}}$ Hermitian matrices that digitize each of the single-site terms in Hamiltonian Eq. (2.97) into operators that can act on $\mathsf{n_q}$ qubits at a time. An appropriate full digitization of Hamiltonian Eq. (2.97) is then obtained by directly substituting the definitions of the digitized single-site operators $\hat{\varphi}_i$ and $\hat{\pi}_i$ into Eq. (2.97), together with the link terms $\hat{\varphi}_i\hat{\varphi}_j = \hat{\varphi}_i \otimes \hat{\varphi}_j$.

We are now in a position to discuss the parameters of this theory necessary for determining the cost of the various methods described in Sec. 2.2. Recall that we found $\mathsf{N_H} = \mathcal{O}(\mathsf{N_\Lambda})$ and $\mathsf{N_{ind}} = \mathcal{O}(1)$. The final parameters needed are $\mathsf{N_P}$ and $\mathsf{c_P}$, which we now discuss.

To determine the maximum number of Pauli strings $\mathsf{N_P}$ appearing in any term in $\hat{H}$, we start by writing the various operators in terms of Pauli operators. Note that the $\mathsf{n_q}$ qubit operators $\hat{\varphi}_i$ and $\hat{\pi}_i^{(p)}$ are diagonal operators with evenly spaced eigenvalues. It was shown in [97] that operators of this form can be decomposed into $\mathsf{n_q}$ single Pauli-Z gates. For example, one has

$$\hat{\varphi} = \frac{\varphi_{\max}}{2^{\mathsf{n_q}} - 1} \sum_{j=0}^{\mathsf{n_q}-1} 2^j Z_j \,, \tag{2.105}$$

where Z_j is a single Pauli-Z gate acting on qubit j. Using this relation, we see that $\mathsf{N_P}$ for the various terms in H_φ are

$$\mathsf{N_P}^{(\varphi_i^2)} = \mathcal{O}(\mathsf{n_q^2}) \,, \tag{2.106a}$$

$$\mathsf{N_P}^{(\varphi_i^4)} = \mathcal{O}(\mathsf{n_q^4}) \,, \tag{2.106b}$$

$$\mathsf{N_P}^{(\varphi_i\varphi_j)} = \mathcal{O}(\mathsf{n_q^2}) \,. \tag{2.106c}$$

Turning now to $\hat{H}_\pi$, while the operator $(\hat{\pi}_i^{(p)})^2$ is a sum of $\mathcal{O}(\mathsf{n_q^2})$ gates, the same is not true for $\hat{\pi}_i^2 = \mathrm{FT}_i \cdot (\hat{\pi}_i^{(p)})^2 \cdot \mathrm{FT}_i^\dagger$; after using the Fourier transform to change to the field basis, the operator $\hat{\pi}_i^2$ is dense and requires $\mathcal{O}(3^{\mathsf{n_q}})$ Pauli strings (see App. A.3). Since the quantum Fourier transform requires $\mathcal{O}(\mathsf{n_q^2})$ gates, it is more cost effective to first decompose $(\hat{\pi}_i^{(p)})^2$ into $\mathcal{O}(\mathsf{n_q^2})$ Pauli strings, and then change to the field basis using the Fourier transform. Using this method results in

$$\mathsf{N_P}^{(\pi^2)} = \mathcal{O}(\mathsf{n_q^2}) \,. \tag{2.107}$$

From this we see that $\mathsf{N_P}$ for the entire Hamiltonian is dominated by the $\hat{\varphi}^4$ term, and is given by

$$\mathsf{N_P} = \mathcal{O}(\mathsf{n_q^4}) \,. \tag{2.108}$$

This result for $\mathsf{N_P}$ is significantly more favorable than the worst case scaling for an arbitrary 2-site-local Hamiltonian of $\mathcal{O}(4^{2\mathsf{n_q}})$.

The final piece is to determine the maximum coefficient $\mathsf{c_P}$ of any Pauli string in $\hat{H}$. Using the fact that our choice of digitization leads to $\pi_{\max} = \mathcal{O}(2^{\mathsf{n_q}/2})$ and $\varphi_{\max} = \mathcal{O}(2^{\mathsf{n_q}/2})$, we see

$$\mathsf{c_P}^{(\pi^2)} = \mathcal{O}(\pi_{\max}^2) = \mathcal{O}(2^{\mathsf{n_q}}), \tag{2.109a}$$

$$\mathsf{c}_{\mathsf{P}}{}^{(\varphi^2)} = \mathcal{O}(|m^2|\varphi^2_{\max}) = \mathcal{O}(|m^2|2^{\mathsf{n}_\mathsf{q}}), \tag{2.109b}$$

$$\mathsf{c}_{\mathsf{P}}{}^{(\varphi^4)} = \mathcal{O}(\lambda\varphi^4_{\max}) = \mathcal{O}(\lambda 4^{\mathsf{n}_\mathsf{q}})\,, \tag{2.109c}$$

$$\mathsf{c}_{\mathsf{P}}{}^{(\varphi_i\varphi_j)} = \mathcal{O}(\varphi^2_{\max}) = \mathcal{O}(2^{\mathsf{n}_\mathsf{q}})\,. \tag{2.109d}$$

We assume that the overall value of c_P is dominated by the $\hat{\varphi}^4$ term, and is then given by

$$\mathsf{c}_{\mathsf{P}} = \mathcal{O}(\lambda 4^{\mathsf{n}_\mathsf{q}})\,. \tag{2.110}$$

If this was not the case due to the relative size of m^2 and λ constants, Eq. (2.110) would be modified trivially by using the appropriate term from Eq. (2.109).

These values for N_P, c_P and N_H will be used in the following sections to determine the asymptotic gate complexity of the various time-evolution algorithms discussed in Sec. 2.2 for geometrically site-local Hamiltonians.

2.4.2 Product Formulas

In this section, we consider using a p^{th} order PF to time evolve the scalar field theory Hamiltonian (given by Eq. (2.97)) that has been digitized using the formalism presented in Sec. 2.4.1.1.

Specifically, we construct concrete error bounds and asymptotic gate complexity estimates as a special case of the results obtained in Sec. 2.3.1 for general geometrically site-local Hamiltonians. We use the result of Theorem 15 for the asymptotic bound on the elementary gate count χ (see Eq. (2.89)) and make the following substitutions based on the discussion provided in Sec. 2.4.1.1:

$$\mathsf{N}_\mathsf{H} = \mathcal{O}(\mathsf{N}_\Lambda), \tag{2.111}$$

$$\mathsf{N}_\mathsf{P} = \mathcal{O}(\mathsf{n}_\mathsf{q}^4), \tag{2.112}$$

$$\mathsf{c}_\mathsf{P} = \mathcal{O}(\lambda 4^{\mathsf{n}_\mathsf{q}}). \tag{2.113}$$

These follow from Eq. (2.12), Eq. (2.108), and Eq. (2.110) respectively. Putting this all together, the asymptotic gate complexity is thus given by

$$\chi = \mathcal{O}\left(\mathsf{n}_\mathsf{q}^{4\left(2+\frac{1}{p}\right)} \frac{\left(4^{\mathsf{n}_\mathsf{q}}\mathsf{N}_\Lambda\lambda t\right)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right). \tag{2.114}$$

2.4.3 Quantum Signal Processing

Obtaining gate counts for the case of pure QSP requires two logical steps. The first is to determine the number of calls to the local block encoding, and the second is to estimate the

complexity of implementing this block encoding. By Theorem 9, the query complexity for simulating with evolution time t to exponentiation error ε is given by

$$\mathcal{O}(\alpha t + \log(1/\varepsilon)) \,, \tag{2.115}$$

where α is the scale factor of the BE used.

We now consider two approaches to implement the local block encoding. The first approach is based on the fully general LCU procedure, and does not take into account any features of the model. The second approach is a modification of the former which takes advantage from splitting the Hamiltonian into the field and momentum parts, in analogy with how it was done in the product-formula-based approach. In this section, we consider only asymptotic gate complexities; detailed gate count comparisons of these techniques will be given in Sec. 2.4.5.

2.4.3.1 Block encoding via LCU

In this section, we determine the cost of simulating time-evolution via QSP when the BE is implemented using a fully general LCU procedure. This involves first digitizing the Hamiltonian $\hat{H}$ and then decomposing it into Pauli operators. In such a scenario, the dominant cost arises from the $\hat{\pi}_i^2$ operator, which is represented by $\mathcal{O}(3^{\mathsf{n_q}})$ Pauli strings (excluding Z gates in all combinations, see App. A.3). The value of $\mathsf{c_P}$ in this case is the same as in Eq. (2.110). Lastly, we have $\mathsf{N_H} = \mathcal{O}(\mathsf{N_\Lambda})$. Plugging these results into Theorem 17 leads to gate complexity

$$\chi = \mathcal{O}\Big[\mathsf{N_\Lambda} 3^{\mathsf{n_q}} \log\left(\mathsf{N_\Lambda} 3^{\mathsf{n_q}}\right) \left(\mathsf{N_\Lambda} 3^{\mathsf{n_q}} (\lambda 4^{\mathsf{n_q}}) t + \log(1/\varepsilon)\right)\Big] \,. \tag{2.116}$$

We see that the naive decomposition of the $\hat{\pi}_i^2$ operators into Pauli gates results in a poor scaling with $\mathsf{n_q}$.

By exploiting the fact that $\textsc{sel}^2 = \mathbb{1}$, the walk operator W_H can then be constructed without ancillary qubits using Eq. (2.47), as was discussed in Sec. 2.2.2.3.

The final step is to construct a controlled call to W_H. Rather than naively controlling every gate in W_H, one can achieve this by only controlling the SEL oracle and the multi-controlled Z gate in the reflector R; the PREP oracle does not need to be controlled [130]. This can be understood by noting that if one wanted to implement a controlled call to some operator $U = A^\dagger B A$ where $A^\dagger A = \mathbb{1}$, one only has to control B. The circuit for the controlled implementation of W_H is shown in Fig. 2.7.

2.4.3.2 Block encoding via LCU and FT

In order to take advantage of the decomposition Eq. (2.103) within the LCU construction, one needs to slightly modify the PREP and SEL oracles. To write $\hat{H}$ in the form required for LCU, we first write the momentum component of $\hat{H}$ in the momentum basis as a sum of M_π

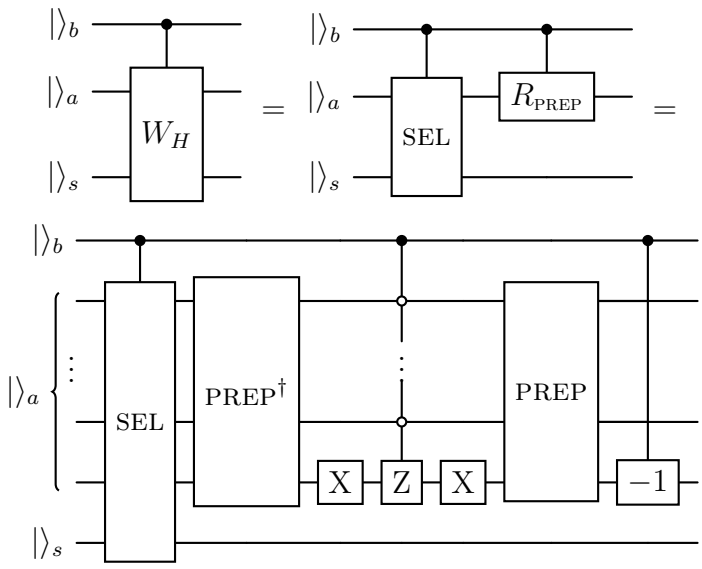

Figure 2.7: Circuit for controlled call to W_H when the BE is prepared using LCU. Note that $|g\rangle_a = \text{PREP}|0\rangle_a$ and so the reflector R_{PREP} (see Fig. 2.4) uses $G = \text{PREP}$.

unitary operators

$$\hat{H}_\pi^{(p)} = \sum_{j=0}^{M_\pi} \beta_j^{(\pi^{(p)})} \hat{U}_j^{(\pi^{(p)})} \,. \tag{2.117}$$

Next, we write the full Hamiltonian as

$$\hat{H} = \sum_{i=0}^{M_\varphi - 1} \beta_i^{(\varphi)} \hat{U}_j^{(\varphi)} + \text{FT} \left(\sum_{j=0}^{M_\pi - 1} \beta_j^{(\pi^{(p)})} \hat{U}_j^{(\pi^{(p)})} \right) \text{FT}^\dagger \,, \tag{2.118}$$

where it is understood that the Fourier transform is performed locally at each site. The Hamiltonian is now in a form suitable for the modified LCU procedure.

First, the PREP oracle is used to encode the real positive coefficients $\beta_i^{(\varphi)}$ and $\beta_j^{(\pi^{(p)})}$. The new SEL oracle is shown in Fig. 2.8 and is given by

$$\text{SEL} = (\mathbb{1}_a \otimes \text{FT}) \otimes \left(\sum_{j=0}^{M_\pi - 1} |j\rangle_a\langle j|_a \otimes \hat{U}_j^{(\pi^{(p)})} \right) (\mathbb{1}_a \otimes \text{FT}^\dagger) + \sum_{i=0}^{M_\varphi - 1} |i + M_\pi\rangle_a\langle i + M_\pi|_a \otimes \hat{U}_i^{(\varphi)} \,, \tag{2.119}$$

where, again, it is understood that the Fourier transform is performed locally at each site. It is important to note that the modified SEL operator still satisfies $\text{SEL}^2 = \mathbb{1}$, which enables the use of the efficient construction of W_H.

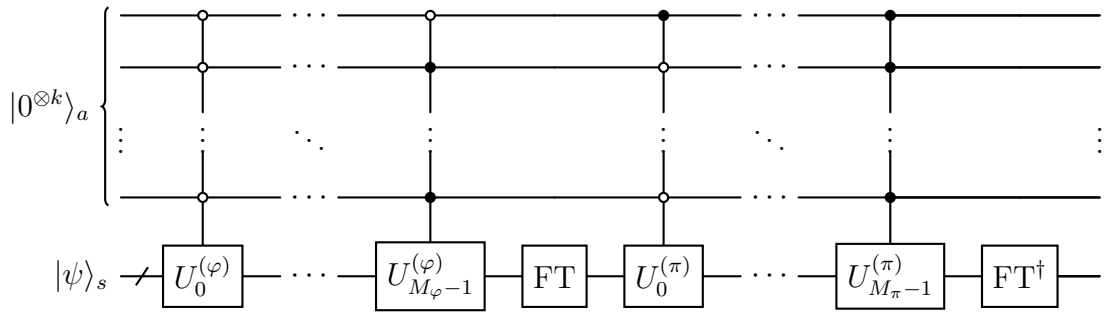

Figure 2.8: Circuit for SEL oracle in Eq. (2.119). As compared to $\mathcal{O}(3^{\mathsf{n_q}})$ and $\mathcal{O}(\mathsf{n_q})$ in the usual LCU implementation from Sec. 2.4.3.1, the usage of Fourier Transform described in Sec. 2.4.3.2 allows one to reduce the circuit depth and the number of ancillary qubits to $\mathcal{O}(\mathsf{n_q^4})$ and $\mathcal{O}(\log \mathsf{n_q})$, correspondingly. Note that while in the usual LCU procedure the gate complexity $\mathcal{O}(3^{\mathsf{n_q}})$ stems from the number of terms in $\hat{H}_\pi$, upon adding the Fourier Transform to the circuit, the $\mathcal{O}(\mathsf{n_q^4})$ complexity stems from the $\hat{\varphi}^4$ term in $\hat{H}_\varphi$.

Note that the SEL oracle prepared using this method encodes the same operators as the SEL oracle prepared using the naive LCU procedure in Sec. 2.4.3.1; the only difference is how it is implemented. This different implementation, however, results in a more favorable scaling for $\mathsf{N_P}$. This can be understood by noting that the $\mathsf{N_P}$ associated with the momentum Hamiltonian is reduced to $\mathsf{N_P}^{(\pi^{(p)})} = \mathcal{O}(\mathsf{n_q^2})$, implying that the value of $\mathsf{N_P}$ is dominated by the $\hat{\varphi}^4$ term. Substituting $\mathsf{N_P} = \mathcal{O}(\mathsf{n_q^4})$, $\mathsf{c_P} = \mathcal{O}(\lambda 4^{\mathsf{n_q}})$ and $\mathsf{N_H} = \mathcal{O}(\mathsf{N_\Lambda})$ into the result in Theorem 17 gives the gate complexity

$$\chi = \mathcal{O}\left[\mathsf{N_\Lambda n_q^4} \log(\mathsf{N_\Lambda n_q^4}) \left(\mathsf{N_\Lambda n_q^4}(\lambda 4^{\mathsf{n_q}})t + \log(1/\varepsilon)\right)\right], \tag{2.120}$$

which has a more favorable scaling with $\mathsf{n_q}$ than the naive LCU implementation in Sec. 2.4.3.1. Since $\mathrm{SEL}^2 = \mathbb{1}$, the controlled call to W_H can be constructed in the same way as the naive LCU case in Sec. 2.4.3.1.

2.4.4 HHKL

In order to apply the HHKL results formulated in our conventions in Theorem 23 and Theorem 24, note that we must scale down the Hamiltonian Eq. (2.97) by the appropriate factor that brings the local terms to at most unit spectral norm, and simultaneously rescale evolution time up by the same factor. This factor is precisely the maximum spectral norm

of any local term $\tilde{H}_J$,

$$\begin{aligned}\max_{J\in\mathcal{S}} \|\tilde{H}_J\| &= \left\| \frac{1}{2}\hat{\pi}_i^2 + \frac{1}{2}m^2\hat{\varphi}_i^2 + \frac{\lambda}{4!}\hat{\varphi}_i^4 + \sum_{j\in\langle ij\rangle} \frac{(\hat{\varphi}_j - \hat{\varphi}_i)^2}{2} \right\| \\ &= \mathcal{O}(\mathsf{N_P c_P}) = \mathcal{O}(\mathsf{n_q^4}(\lambda 4^{\mathsf{n_q}}))\,,\end{aligned} \tag{2.121}$$

where we have used $\mathsf{N_P} = \mathcal{O}(\mathsf{n_q^4})$ and $\mathsf{c_P} = \mathcal{O}(\lambda 4^{\mathsf{n_q}})$, and where any lattice site with index i in the interior of the lattice (or on the boundary if there are no sites in the interior) can be chosen to compute the spectral norm, and the sum runs over all lattice sites j neighboring i. We therefore perform the following rescalings

$$\tilde{H} = \frac{H}{\max_{J\in\mathcal{S}} \|H_J\|} \tag{2.122}$$

$$\tilde{t} = t \max_{J\in\mathcal{S}} \|H_J\| = \mathcal{O}(4^{\mathsf{n_q}}\mathsf{n_q^4}\lambda t) \tag{2.123}$$

For HHKL built out of local Trotter operations, Theorem 23 implies a gate complexity

$$\chi = \mathcal{O}\Big[\mathsf{n_q^4}(\mathsf{N_\Lambda n_q^4}(\lambda 4^{\mathsf{n_q}})t)^{1+\frac{1}{p}}\varepsilon^{-\frac{1}{p}}\Big]\,. \tag{2.124}$$

where we substituted $\mathsf{c_P} = \mathcal{O}(\lambda 4^{\mathsf{n_q}})$ and $\mathsf{N_P} = \mathcal{O}(\mathsf{n_q^4})$.

Similarly, for HHKL built out of local QSP operations, Theorem 24 then implies a gate complexity

$$\chi = \mathcal{O}\Big[(\mathsf{n_q^4})^2\mathsf{N_\Lambda}(\lambda 4^{\mathsf{n_q}})t \log^{\mathsf{d}}\left(\mathsf{N_\Lambda}t/\varepsilon\right) \log\left(\mathsf{n_q^4}\log\left(\mathsf{N_\Lambda}t/\varepsilon\right)\right)\Big]. \tag{2.125}$$

To conclude our discussion on HHKL scalings, we again stress an important point. As we have previously remarked, while the purely asymptotic scaling of the Trotter-based HHKL gate count is worse than that for QSP-based HHKL, the difference is not striking when the Trotter order p is taken to be large. This allows for the possibility of lower constant prefactors due to the smaller overhead involved in implementing Trotter operations compared to QSP operations, and may even be a prudent choice for a more practical implementation on near-term devices, if the exact gate counts follow suit before asymptotics take over.

2.4.5 Numerical results

In the previous section we worked out the asymptotic gate count scaling for several time-evolution algorithms. While these expressions are helpful for identifying which methods have the best scaling in a particular parameter, the practical question of the size of the prefactors requires a dedicated numerical study, which we perform in this section. In particular, we first compare the cost of the two BE methods described in the previous section. For the

comparison, we choose to compare the cost of constructing a single controlled call to the W_H operator. This choice allows a more direct comparison of the cost of each method as in some cases significant circuit simplifications occur compared to naively controlling every gate in the W_H operator. From there, we turn to time evolution and compare the cost of using QSP to that of a 4th order PF, focusing on the scaling with time t and error ε.

We choose to focus on a single site Hamiltonian given by

$$\hat{H} = \frac{1}{2}\hat{\pi}^2 + \frac{1}{2}\hat{\varphi}^2 + \frac{\lambda}{4!}\hat{\varphi}^4 , \tag{2.126}$$

where we set $\lambda = 32$ (similarly to [97]) as an instance of the strongly coupled regime where perturbation theory breaks down. While this will not provide an explicit scaling with the number of lattice sites, analyzing the single site case in detail will provide intuition for what values of t and ε QSP outperforms PF methods for $\mathsf{N}_\Lambda > 1$.

Before comparing gate counts of the various methods, we describe the details of constructing the circuits. All gate counts were found using QISKIT's transpile function with an optimization level of 1; higher optimization levels result in the dropping of gates with small arguments and were therefore not used. We use the universal gate-set containing R_x, R_z, and CNOT gates. When implementing multi-controlled gates, we compared the cost using two different methods. The first uses QISKIT's built-in functionality to decompose the multi-controlled gates into $\mathcal{O}(M^2)$ CNOT gates, where M is the number of qubits being controlled on. The second method uses $\mathcal{O}(M)$ additional work qubits to decompose the multi-controlled gates into $\mathcal{O}(M)$ gates [151]. For each value of $\mathsf{n_q}$, we compare the two and choose the method that results in the smallest number of gates. Lastly, it is important to note that, when implementing multi-controlled gates, global phases must also be controlled. Because using LCU requires the coefficients in the PREP oracle to be real and positive, any overall phase must necessarily be absorbed into the individual operators in the SEL oracle. Controlling these phases requires an additional controlled single-qubit rotation. Alternatively, one could use the generalized LCU method to allow the coefficients in the PREP oracle to be complex and set the phase of the operators in the SEL oracle to have zero overall phase (see e.g. [117]). While this avoids having to control the phases, it leads to a more complicated PREP oracle. Additionally, because the generalized LCU method uses circuits of the form $\widetilde{\text{PREP}} \times \text{SEL} \times \text{PREP}$ with $\widetilde{\text{PREP}} \times \text{PREP} \neq \mathbb{1}$, one will have to also control both PREP and $\widetilde{\text{PREP}}$ in the controlled W_H circuits in Fig. 2.7. For these reasons, we use the former method to control the overall phases.

We now compare the cost of implementing the controlled W_H operator using the two different BE methods. The top and middle plots in Fig. 2.9 show the CNOT gate count and rotation gate count, respectively, as a function of $\mathsf{n_q}$. The blue points show the gate counts using the LCU method where one naively decomposes the full digitized Hamiltonian into Pauli strings as described in Sec. 2.4.3.1. As expected, this method scales least favorably with $\mathsf{n_q}$. As previously discussed in Sec. 2.4.3.2, this poor gate scaling with $\mathsf{n_q}$ can be improved by utilizing the Fourier Transform to implement the momentum component of the

Hamiltonian. This improved scaling can be seen by looking at the orange squares in the top and middle plots in Fig. 2.9.

The bottom plot in Fig. 2.9 shows the number of ancillary qubits required to implement a controlled call to W_H for each method. As expected from the $\mathcal{O}(\mathsf{n_q})$ scaling, using the naive LCU method requires the most ancillary qubits. Using LCU combined with the FT reduces the scaling to $\mathcal{O}(\log \mathsf{n_q})$. Note that, for both methods, the overall coefficient in the number of ancillary qubits is large due to needing additional work qubits to efficiently implement multi-controlled gates.

Taken together, the gate count and ancillary qubit count comparisons show that using LCU combined with the FT is superior to the naive LCU method for all values of $\mathsf{n_q}$.

Now that we have identified the optimal method for preparing BEs, we turn to the cost of performing time evolution. The first method we implement is a 4th order PF. The second uses QSP, where the BE is prepared using the BE method that minimizes the number of rotation gates, *i.e.*, using LCU combined with the FT. The top (bottom) plot in Fig. 2.10 show the rotation (CNOT) gate count for both methods as a function of evolution time t and error ε for $\mathsf{n_q} = 3$. The value of the scale factor α for the QSP data is $\alpha = 134.5$. Due to QSP's large overall prefactor from implementing the BE, using the 4th order PF requires less rotation and CNOT gates for errors $\varepsilon \gtrsim 10^{-8}$. For errors $\varepsilon \lesssim 10^{-8}$, QSP wins due to its exponentially better scaling with respect to ε.

It is interesting to note that the QSP gate count scaling is almost independent of the desired error ε. This is due to the fact that the αt term in the asymptotic cost is significantly larger than the $\log(1/\varepsilon)$ term (see e.g. Eq. (2.115)). These findings therefore show that the additional cost of implementing time evolution to floating point precision using QSP is negligible compared to implementing time evolution approximately (not accounting for the increased cost from requiring higher precision rotation gates). In Sec. 2.5 we discuss the implications of these results for multi-site systems.

2.5 Discussion

In this section we discuss how results of Sec. 2.3 provide guidance for choosing a simulation algorithm for the particular physical model and then examine the results of Sec. 2.4.

The major parameters of interest defining the asymptotic complexities of time evolution algorithms are the evolution time t, the error of the time-evolved state ε, as well as some measure of Hamiltonian size/complexity. For lattice systems, this can be the number of terms in the Hamiltonian $\mathsf{N_H}$. For a particular theory, $\mathsf{N_H}$ can also be expressed in terms of lattice size $\mathsf{N_\Lambda}$. The dependence of $\mathsf{N_H}$ on $\mathsf{N_\Lambda}$ is linear for geometrically-site-local Hamiltonians and is, more generally, polynomial. For algorithms based on PFs, one additionally needs to consider a parameter $\mathsf{N_{ind}}$ related to the induced norm of the Hamiltonian. The asymptotic dependence on individual parameters for algorithms discussed in Secs. 2.2 and 2.3 is summarized in Table 2.5.

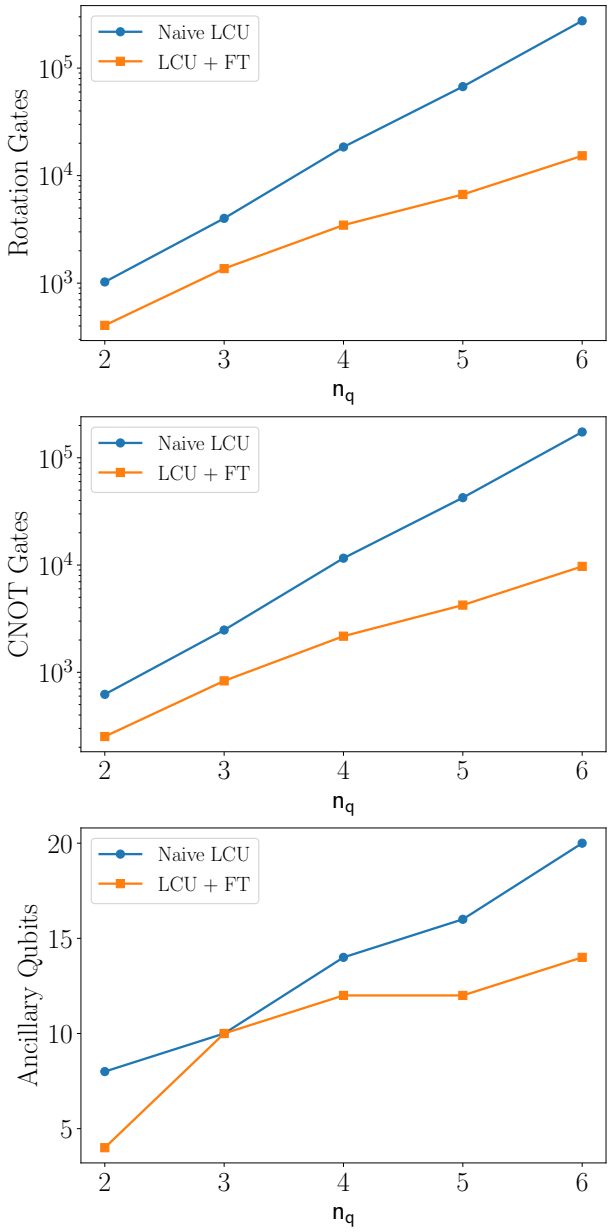

Figure 2.9: The top, middle, and bottom plots show the CNOT gate count, rotation gate count, and number of ancillary qubits needed to prepare a controlled call to the W_H operator corresponding to the single site Hamiltonian in Eq. (2.126) as a function of $\mathsf{n_q}$. Blue circles correspond to using the naive LCU method described in Sec. 2.4.3.1. Orange squares correspond to using LCU combined with the Fourier Transform as described in Sec. 2.4.3.2. Using LCU combined with the FT results in an exponential improvement with $\mathsf{n_q}$ in both the gate count and number of ancillary qubits required.

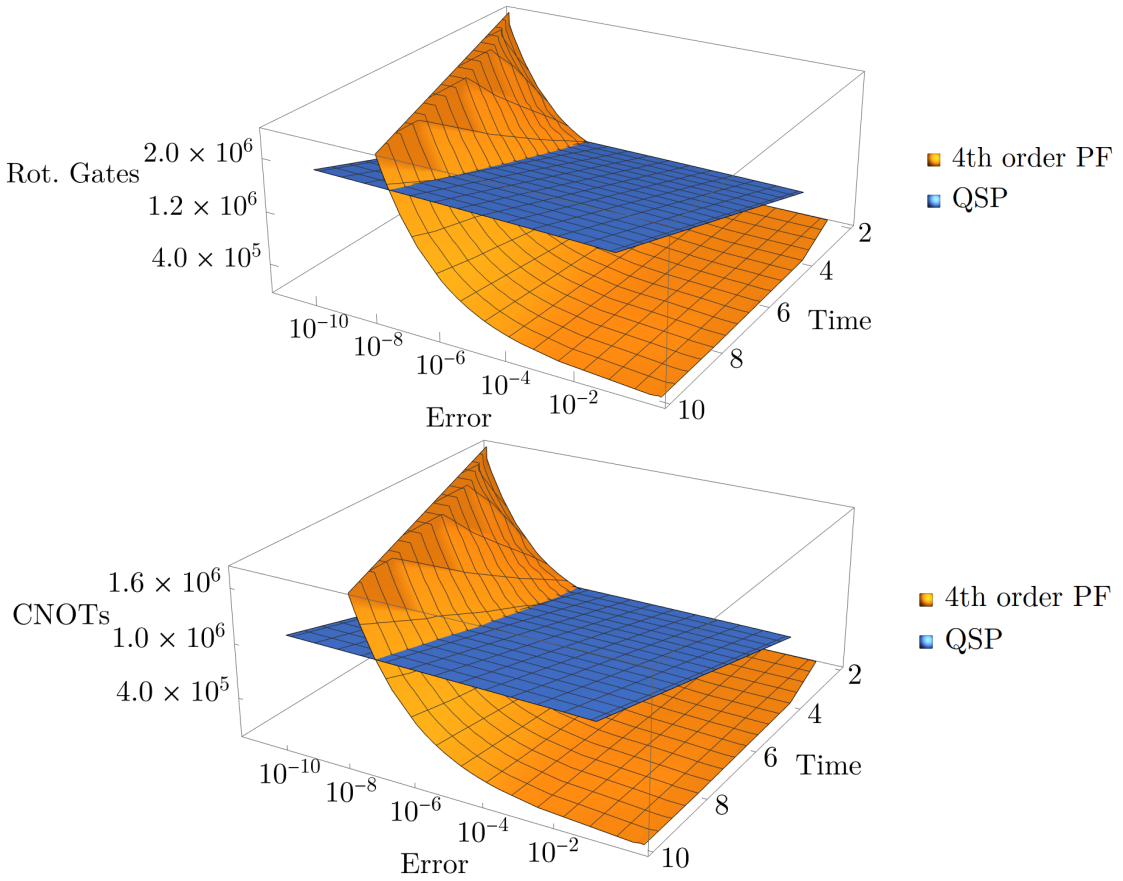

Figure 2.10: The top (bottom) plot shows the rotation (CNOT) gate count for implementing the time evolution operator for the single site Hamiltonian in Eq. (2.97) with $\mathsf{n_q} = 3$. Results are shown as a function of time t and error ε. The orange and blue colored surfaces show results using a 4th order PF and QSP, respectively. The BE used in the QSP calculation was implemented using the method in Sec. 2.4.3.2. The PF requires less rotation and CNOT gates for errors $\varepsilon \gtrsim 10^{-8}$. For error tolerances $\varepsilon \lesssim 10^{-8}$, QSP wins due to its exponentially better scaling with ε. Furthermore, the QSP gate count is almost independent of the error due to the optimal scaling with ε.

Non-geometrically-local

	ε	t	$\mathsf{N_H}$
a) PF	$\varepsilon^{-\frac{1}{p}}$	$t^{1+\frac{1}{p}}$	$\mathbf{\mathsf{N_H}^{1+\frac{1}{p}}\mathsf{N_{ind}}}$
b) QSP+PF	$\varepsilon^{-\frac{1}{p}}\log(1/\varepsilon)$	$t^{1+\frac{1}{p}}$	$\mathsf{N_H}^{1+\frac{1}{p}}\mathsf{N_{ind}}\log(\mathsf{N_H}\mathsf{N_{ind}})$
c) QSP	$\mathbf{\log(1/\varepsilon)}$	$\boldsymbol{t}$	$\mathbf{\mathsf{N_H}^2\log(\mathsf{N_H})}$

Geometrically-local

	ε	t	$\mathsf{N_\Lambda}$
d) PF	$\varepsilon^{-\frac{1}{p}}$	$t^{1+\frac{1}{p}}$	$\mathsf{N_\Lambda}^{1+\frac{1}{p}}$
e) PF+HHKL	$\varepsilon^{-\frac{1}{p}}$	$t^{1+\frac{1}{p}}$	$\mathsf{N_\Lambda}^{1+\frac{1}{p}}$
f) QSP+PF	$\varepsilon^{-\frac{1}{p}}\log(1/\varepsilon)$	$t^{1+\frac{1}{p}}$	$\mathsf{N_\Lambda}^{1+\frac{1}{p}}\log(\mathsf{N_\Lambda})$
g) QSP	$\mathbf{\log(1/\varepsilon)}$	$\boldsymbol{t}$	$\mathsf{N_\Lambda}^2$
h) QSP+HHKL	$\mathbf{\log(1/\varepsilon)}$	$\boldsymbol{t}$	$\mathbf{\mathsf{N_\Lambda}[\log(\mathsf{N_\Lambda})]^d\log\log(\mathsf{N_\Lambda})}$

Table 2.5: Asymptotic dependence ($\mathcal{O}$ omitted) on major problem parameters for algorithms in Tables 2.3 and 2.4. Boldface indicates best asymptotic scaling among applicable methods. While $\mathsf{N_{ind}}$ is constant for geometrically-local theories, it typically scales polynomially with $\mathsf{N_\Lambda}$ otherwise. For most non-geometrically-local theories, the order p can be chosen large enough so that PFs demonstrate better scaling with $\mathsf{N_\Lambda}$ than QSP.

The parameter for which the difference in asymptotic dependence is most striking for PF- and QSP-based algorithms is the error ε. Nevertheless, using algorithms with $\log(1/\varepsilon)$ dependence may not necessarily be an optimal choice for several reasons. First, their asymptotic cost comes with large constant prefactors stemming from the implementation of the BE subroutine. While progress is being made in this direction, the value of ε at which using such techniques clearly becomes advantageous is $\varepsilon \sim 10^{-8}$, which is much smaller than uncertainties coming from other sources (e.g., from digitizing a theory with bosonic degrees of freedom, finite lattice spacing and volume effects, to name just a few). Second, as discussed below, in some scenarios PF-based algorithms have better dependence on problem size than those based on QSP. Third, and least importantly (in the fault-tolerant regime), PF based algorithms do not require the usage of ancillary qubits, while QSP-based algorithms typically require a logarithmic (in problem parameters) number of those.

Recent research [133] indicates that the error tolerance at which near-optimal algorithms

start outperforming the PF-based ones is highly problem-dependent. In particular, when simulating lattice gauge theories, the complexity of PF-based simulation is likely to scale exponentially with the number of colors [63, 64], while the near-optimal algorithms can have polynomial scaling with the number of colors, which is achieved by leveraging the sparsity of gauge field operators. As a result, the usage of near-optimal techniques may lead to tremendous savings starting at error rates as low as $\varepsilon = 10^{-1}$ [133].

In terms of the asymptotic dependence on problem size, a clear winner is the HHKL algorithm, assuming that for implementing the time evolution of local blocks, an algorithm with $\log(1/\varepsilon)$ scaling is chosen, such as QSP. Note, however, that for higher-order PFs [105] the improvement reached by the HHKL algorithm in terms of problem size may not be particularly pronounced, especially, unless a very high precision is required, as discussed above.

For non-geometrically-local Hamiltonians, the situation is largely determined by the Hamiltonian structure defining the dependence of $\mathsf{N}_{\mathsf{ind}}$ on N_{Λ}. In particular, in the second-quantized formulation, this dependence will be defined by the choice of single-particle basis states (plane waves [131, 157], wavelets [160], harmonic oscillator eigenstates [161], etc.). Importantly, whether the Hamiltonian is geometrically-local or not, one can imagine scenarios in which, at a fixed value of ε, switching from PF to QSP would make sense as one increases the problem size. However, given the sharp difference in the dependence on ε for these approaches, such scenarios are not highly probable.

Lastly, we note that the dependence on t is likely to not be a crucial parameter in choosing the simulation algorithm. Nevertheless, a situation in which considering longer times can motivate one to switch from PF to QSP is possible as well.

In addition to factors considered in this work, other considerations should also be taken into account. Most notably, the compatibility of methods based on BE with general sparse Hamiltonians and modern simulation techniques [75, 162, 163] and the potential for applying the PF-based methods on early-fault-tolerant quantum computers [118, 164].

Construction of efficient block encodings (BEs) is critical for improving the efficiency of QSP-based algorithms. In [52, 97] it was shown how digitizing a bosonic degree of freedom in the eigenbases of position/momentum operator is advantageous from the perspective of using PF algorithms. There, the key factors are a) an economical qubit representation of arbitrary powers of position/momentum operators in their respective eigenbases and b) the efficiency of switching between the position and momentum bases with the aid of Fourier Transform implemented on the quantum computer. In Sec. 2.4.3.2 we showed how this approach can be equally well applied when constructing a BE with LCU.

2.6 Conclusion

In this work we studied various approaches to simulating time evolution in application to Hamiltonian Lattice Field Theories (HLFTs). We started off by providing a review of existing techniques based on Product Formula (PF) splitting [106], Quantum Signal Processing

(QSP) [108] and Qubitization [109], Linear Combination of Unitaries (LCU) [111], as well as the HHKL algorithm [128]. Following that, we derived asymptotic costs of these methods in terms of parameters describing site-local and geometrically-site-local Hamiltonians arising in HLFT, see Tables 2.3 and 2.4. While in this work we chose to contrast the most accessible and widely used algorithms based on PF with the methods having the best asymptotic scaling (QSP, HHKL), we note that there exist intermediate approaches. Most notably, this includes the direct usage of Taylor series and LCU without QSP [165].

Next, we focused on a particular type of HLFT, the quartic scalar field theory defined on a hypercubic lattice. Our specific choice was guided by the fact that all the fundamental interactions are mediated by bosons, and from the computational perspective the scalar field theory is largely reminiscent of lattice gauge theories [90, 91, 93, 134, 135, 136, 166]. Besides this, an extensive volume of literature exists focusing on simulations of fermionic systems [129, 130, 146, 167, 168, 169, 170, 171, 172, 173, 174]. Developing simulation algorithms for systems combining particles of both statistics is an exciting direction which was pursued in [98, 99, 175] where authors assumed the usage of PF-based methods, and is currently investigated in the context of nearly-optimal techniques [133].

For the scalar field theory, we worked out the query complexities of various simulation methods as well provided explicit gate counts (not the upper bounds) for the case of simulating the system on a single lattice site using PF and QSP. Next, we focused on block encoding a single site in scalar field theory and left to further work additional improvements, such as reducing the cost of SEL operator from $\mathcal{O}(\mathsf{N}_\Lambda \log \mathsf{N}_\Lambda)$ to $\mathcal{O}(\mathsf{N}_\Lambda)$ [130] or taking into account the structure of coefficients in the Hamiltonian operator while/in constructing PREP operator. We developed an improved version of the LCU block encoding for digitized bosons [52, 97] which utilizes the Quantum Fourier Transform circuit for switching between the field and momentum eigenbases, see Fig. 2.8. We then provided explicit gate counts for simulating time evolution which showed that PF-based simulation works exceptionally well for simple systems, such as anharmonic oscillator, and outperforms the QSP-based algorithm for $\varepsilon \gtrsim 10^{-9}$. We note, however, that considering more complex models (such as non-Abelian lattice gauge theories) and larger lattice system sizes may result in near-term algorithms outperforming PF for error tolerances as little as $\varepsilon \lesssim 10^{-1}$ [133].

Another important distinction between PF and QSP methods is the ease of quoting accurate error bounds for the final result. Starting with QSP methods, all that is needed to place an upper bound on the error is the scale factor α, the time t, and the degree of the polynomial approximation. While the error determined in this way is only an upper bound, because QSP methods have optimal scaling $\log(1/\varepsilon)$ in the error, the cost of overestimating the error is low compared to the contribution stemming from the αt term in Eq. (2.62).

For PF methods on the other hand, while rigorous error bounds also exist, they are known to generally significantly overestimate the actual error [106]. Choosing the value of the Trotter number r based on these loose error bounds would therefore result in significant extra computational costs due to the sub-optimal $(1/\varepsilon)^{1/p}$ scaling of PF methods. Tightening error bounds in this case implies either performing laborious problem-specific commutator calculations [106, 175] or performing numerical simulations at several values of the Trotter

number in order to extrapolate the results. As shown in [113], since using PFs to approximate the time evolution operator is equivalent to introducing a temporal lattice with spacing a_t, such an extrapolation would necessitate a renormalization procedure with a costly scale setting procedure at each value of a_t.[11]

[11] As was also shown in [113], to avoid a double extrapolation $a \to 0$ and $a_t \to 0$, one should work at fixed anisotropy $\xi \equiv a/a_t$ and take the limit $a \to 0$.

Chapter 3

Block Encoding Bosons by Signal Processing

3.1 Introduction

An important ingredient of post-Trotter methods such as Quantum Signal Processing (QSP) [108, 109] is to encode the Hamiltonian directly into a quantum circuit. Given that quantum circuits can only implement unitary, not Hermitian, operators, the Hamiltonian is typically encoded as a sub-block of a larger unitary matrix. This process is called Block Encoding (BE), and is the source for most of the prefactors in the algorithmic cost of post-Trotter algorithms. The commonly used approaches to constructing block encodings include the LCU algorithm and the sparse oracle approach [117, 131, 132, 176, 177]. Highly optimized constructions of BEs based on both approaches have been developed for various types of physical systems [130, 133, 148].

In this work, we present several novel approaches to constructing BEs for a certain class of Hamiltonians. We demonstrate that, while BEs are commonly seen as building blocks for QSP-based algorithms, *QSP-based algorithms themselves can be utilized for the highly efficient construction of BEs*. Our approaches apply to Hamiltonians which can be written as

$$\hat{H} = \sum_{\ell=0}^{N-1} \hat{H}_\ell \,, \tag{3.1}$$

where each $\hat{H}_\ell$ acts on at most n qubits. Many Hamiltonians have this structure, and we shall term them as *n-site-local* Hamiltonians. We will focus on the situation where each term in the Hamiltonian can be diagonalized via a known and efficiently implementable transformation

$$\hat{H}_\ell = U_\ell^\dagger \hat{H}_\ell^{(D)} U_\ell \tag{3.2}$$

For some of our methods, we will make a further assumption about the structure of the Hamiltonian, which is satisfied for certain bosonic lattice field theories, which have recently garnered attention for their crucial role in high-energy and low-energy nuclear physics [41, 42, 61, 62, 65, 66, 67, 71, 72, 76, 77, 91, 93, 114, 133, 135, 156, 166, 175, 177, 178, 179, 180, 181, 182, 183, 184]. Specifically, we consider theories for which the Hamiltonians are formulated in terms of conjugate operators $[\hat{\varphi}_i, \hat{\pi}_j] = i\delta_{ij}$ at lattice sites i, j. Using $\mathsf{n_q}$ qubits, the operators $\hat{\varphi}_i$ are digitized by sampling at equally spaced values. Since the operator $\hat{\pi}_i$ is conjugate to $\hat{\varphi}_i$, it can be diagonalized using a Fourier transform $\hat{\pi}_i = \widehat{\mathrm{FT}}^\dagger \hat{\pi}_i^{(D)} \widehat{\mathrm{FT}}$. The Hamiltonian can be written as functions of simple operators $\hat{\varphi}_i$ and $\hat{\pi}_i$, and possibly combinations like $\hat{\varphi}_i - \hat{\varphi}_j$ needed for gradient terms. We will use the general notation $\hat{\xi}_k$ ($k \in \{0, 1, 2\}$) to denote any of the operators

$$\hat{\xi}_0 \in \{\hat{\varphi}_i\}\,, \quad \hat{\xi}_1 \in \{\hat{\pi}_i^{(D)}\}\,, \quad \hat{\xi}_2 \in \{\hat{\varphi}_i - \hat{\varphi}_j\}\,, \tag{3.3}$$

such that each $\hat{H}_\ell^{(D)}$ is then given by a function of one such operator

$$\hat{H}_\ell^{(D)} = f_{k_\ell}(\hat{\xi}_{k_\ell})\,. \tag{3.4}$$

A key property of the operators $\hat{\xi}_{k_\ell}$ we exploit in this work is that each of them can be written as a linear combination of $\mathcal{O}(\mathsf{n_q})$ many Pauli Z gates.

Based on this set up and these assumptions, we present 3 different methods for the BE of the individual Hamiltonian terms $\hat{H}_\ell$, which are then combined via LCU, see Fig. 3.1. The first method first obtains a block encoding for the operator $\hat{\xi}_{k_\ell}$ and then uses QSVT to construct the BE for $\hat{H}_\ell^{(D)}$. The second method constructs the unitary operator $\mathrm{e}^{-i\tau\hat{\xi}_{k_\ell}}$ (or $e^{-2i\arccos(\hat{\xi}_k/\alpha)}$) and uses it as a building block to construct the BE for $\hat{H}_\ell^{(D)}$ using QETU. The final method, which we call Linear Operators Via Exponentiation and Linear Combination of Unitaries (LOVE-LCU) constructs the unitary operators $\mathrm{e}^{\pm i\arccos(f_{k_\ell}(\hat{\xi}_{k_\ell})/\alpha)}$, from which $\hat{H}_\ell^{(D)}$ can easily be constructed by adding the two terms via LCU. Following the construction of the BE for $\hat{H}_\ell^{(D)}$ we conjugate it by U_ℓ resulting in a BE for $\hat{H}_\ell$.

In Sec. 3.2, we briefly review the GQSP, QSVT, and QETU algorithms. In Sec. 3.3, we first review the general formulation of bosonic field theories considered in this work and then discuss the 3 novel block encoding methods in more detail. We also demonstrate that, due to the highly symmetric structure of QSVT, QETU, and LOVE-LCU circuits, the qubitized walk operator can be constructed without ancillary qubits in a straightforward way. In Sec. 3.4 we use the proposed methods to construct BEs for several models of interest, construct explicit circuits, and provide metrics such as gate count and the number of ancillary qubits. For a single site system, we find that using our methods allows to implement time evolution with less gates than PFs for errors as small as $\epsilon \sim 10^{-2}$. Our conclusions are presented in Sec. 3.5.

The work presented in this chapter is based on [115].

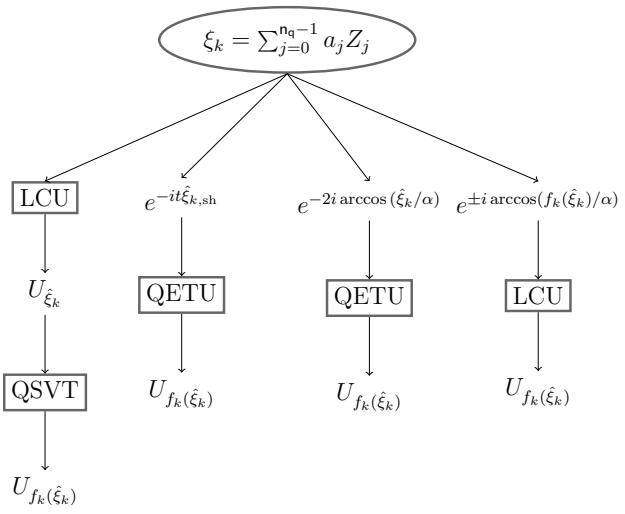

Figure 3.1: Various methods for block encoding operators in Hamiltonians of scalar field theories (see Sec. 3.3). In the first two approaches, the cost of constructing the building block is polynomial in $\mathsf{n_q}$, while in the latter two, the cost is exponential in $\mathsf{n_q}$. In the first three approaches, the query depth is determined by the convergence of the polynomial approximation to the desired function. In the last approach, the query depth is constant (two).

3.2 Algorithms

In this section, we provide brief reviews of the simulation algorithms and techniques we consider in this work, along with associated circuit diagrams. Readers familiar with these concepts are encouraged to skip directly to Sec. 3.3.

3.2.1 Simulating time evolution using Generalized QSP

In this section, we review an improved version of the original QSP method discussed in Sec. 2.2, known as Generalized QSP (GQSP). Provided access to some unitary U, QSP and GQSP allow one to construct circuits implementing the action of polynomials of U. In the context of simulating time evolution, one can choose U to be the Hamiltonian walk operator and use GQSP to implement polynomial approximations of the time evolution operator. The main result of [110], relevant to the problem of simulating time evolution, is given by:

Theorem 25 (Corollary 5 from [110])**.** *$\forall P \in \mathbb{C}[x]$ with deg(P) = d, if*

$$\forall x \in \mathbb{T}, \quad |P(x)|^2 \leq 1\,, \tag{3.5}$$

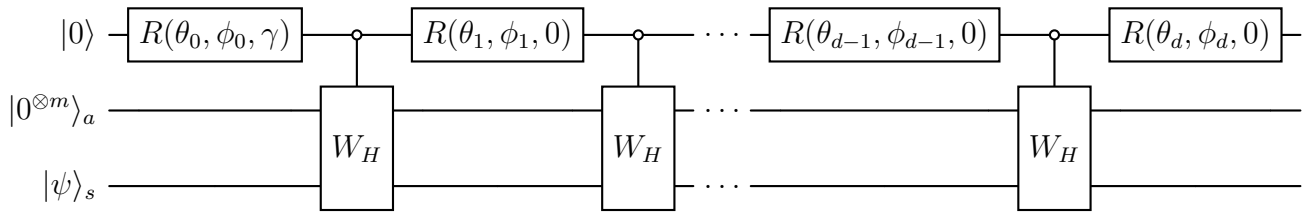

Figure 3.2: Circuit diagram for simulating time evolution using GQSP, adapted from [110]. Each $R(\theta_0, \phi_0, \gamma)$ is an SU(2) rotation, with the individual angle parameters shown in the diagram chosen to construct a polynomial of W_H in the context of this work.

where $\mathbb{T} = \{x \in \mathbb{C} : |x| = 1\}$ *is the unit circle in the complex plane, then* $\exists\, \vec{\theta}, \vec{\phi} \in \mathbb{R}^{d+1}$ *and* $\gamma \in \mathbb{R}$ *such that:*

$$\left(\prod_{j=1}^{d} R(\theta_j, \phi_j, 0)\, C_U\right) R(\theta_0, \phi_0, \gamma) = \begin{pmatrix} P(U) & * \\ * & * \end{pmatrix}, \tag{3.6}$$

where $R(\theta, \phi, \gamma)$ *represents an SU(2) rotation and* $C_U = (|0\rangle\langle 0| \otimes U) + (|1\rangle\langle 1| \otimes \mathbb{1})$ *is a 0-controlled application of* U *(the* signal operator*),*

Furthermore, [110] also provides an efficient classical algorithm to compute the required rotation angles $\vec{\theta}, \vec{\phi}, \gamma$ for a desired polynomial $P(x)$.

For the purpose of simulating time evolution, we note that the eigenvalues of the operator $W_H^{(\lambda)}$, defined by restricting the action of W_H to the subspace spanned by $\text{span}\{|0\rangle_a|\lambda\rangle_s, |\perp^{\lambda}\rangle_{as}\}$, are given by $\mathrm{e}^{\pm i \arccos(\lambda)1}$. Thus we set the unitary to which controlled calls are made in the GQSP circuit to be $U = W_H$ (see Fig. 3.2 adapted from [110]) and seek an implementation of the function given by $\mathrm{e}^{-it\cos(x)}$ to achieve the desired polynomial transformation given by $P(W_H^{(\lambda)}) = \mathrm{e}^{-it\lambda}$. This is achieved through the Jacobi-Anger expansion (see for example Refs. [109, 110]), given by:

$$\mathrm{e}^{-it\cos(x)} = \sum_{k=-\infty}^{\infty} (-i)^k J_k(t) \mathrm{e}^{ikx}, \tag{3.7}$$

where J_k is the k^{th} Bessel function of the first kind. It has been previously established the coefficients in this Fourier series decay exponentially and in particular, the number of SU(2) rotations required to achieve an ϵ-close approximation to this polynomial transformation is asymptotically bounded by $\mathcal{O}(t + \log(1/\epsilon)/\log\log(1/\epsilon))$ (see Refs. [109, 110]). We conclude by noting that, while using QSP for Hamiltonian simulation achieves the same asymptotically optimal scaling, using GQSP requires half as many controlled calls to W_H [110, 185].

[1]Throughout this work we will use $|g\rangle_a = |0\rangle_a$, see Sec. 2.2.2.4

3.2.2 QSVT

While QSP and GQSP are powerful algorithmic tools, in certain cases, the required polynomial transformations can be achieved using the simpler Quantum Singular Value Transformation (QSVT) [116]. The QSVT method is very similar to the QSP and GQSP methods in that it allows one to implement general matrix functions of some Hermitian operator A by performing repeated calls to its BE U_A. The major difference of QSVT has to do with the classes of matrix functions that can be implemented. In particular, QSVT implements real functions with definite parity, whereas QSP implements complex functions with the real and imaginary parts having opposite parity. Only requiring the real part of the matrix polynomial leads to several advantages. First, the fundamental building block is now the BE U_A, rather than a controlled call to the walk operator W_A, as in QSP or GQSP. Second, the conditions on implementable real polynomial are relaxed relative to the complex polynomials implemented using QSP. Third, the phases that parameterize the circuit can be solved using a straightforward optimization procedure for large polynomials using e.g. QSPPACK [186]. The QSVT theorem for Hermitian matrices is as follows [117]:

Theorem 26 (QSVT for Hermitian matrices with real polynomials)**.** *Let $A \in \mathbb{C}^{2^n \times 2^n}$ be a normalized Hermitian operator with a $(1, m)$ block encoding U_A. Given a degree d polynomial $P_{\mathrm{Re}}(x) \in \mathbb{R}[x]$ such that*

1. *P_{Re} has parity $d \pmod 2$,*
2. *$|P_{\mathrm{Re}}(x)| \leq 1 \ \forall x \in [-1, 1]$,*

then there exist symmetric phase factors $(\phi_0, \ldots, \phi_d) \in \mathbb{R}^{d+1}$ such that the circuit presented in Fig. 3.3 implements a $(1, m+1)$ block-encoding of $P_{\mathrm{Re}}(A)$.

One important property of QSVT circuits is that their controlled versions can be implemented with a small overhead [116]. For even polynomials, only the single-qubit R_z gates in the $\mathrm{CR}_{\tilde{\phi}_j}$ circuits (see Fig. 3.4) acting on the signal qubit need to be controlled. For odd polynomials, in addition to controlling the single-qubit R_z gates, one must also control a single call to the BE U_A. This property can be used to combine different QSVT circuits using LCU without introducing a large overall constant prefactor, which is important when using LCU to combine BEs of different local terms in the Hamiltonian constructed using QSVT, as explained in Sec. 3.3.3.

3.2.3 QETU

QETU is similar to the QSVT method described in Sec. 3.2.2 in that it allows one to implement a broad class of matrix functions of some Hermitian operator A by performing repeated calls to some fundamental building block. The major difference of QETU is that, instead of performing repeated calls to a BE U_A as with QSVT, the building block for QETU is the time-evolution operator $\mathrm{e}^{-i\tau A}$ for some value of τ. In a broad sense, QETU replaces

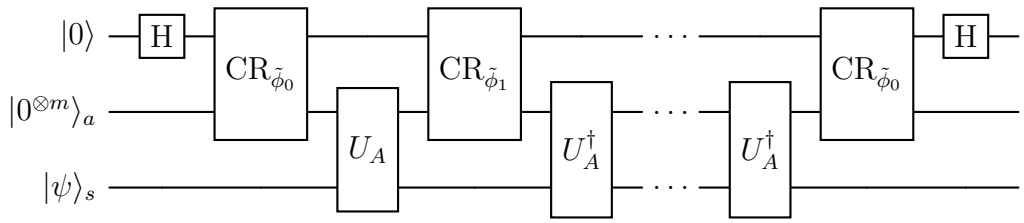

Figure 3.3: Circuit diagram for QSVT, reproduced from [117], using symmetric phases $(\tilde{\phi}_0, \tilde{\phi}_1, \ldots, \tilde{\phi}_1, \tilde{\phi}_0)$. The Hadamard gates serve to isolate the real part of the matrix function being implemented. The circuit for CR_ϕ is shown in Fig. 3.4.

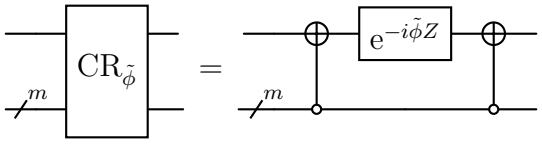

Figure 3.4: Diagram for the CR_ϕ circuit appearing in the QSVT circuit in Fig. 3.3.

the BE U_A in the QSVT circuit in Fig. 3.3 with the controlled time-evolution operator $\mathrm{e}^{-i\tau A}$. The QETU circuit for implementing even functions is shown in Fig. 3.5. Taken from [118], the QETU theorem is:

Theorem 27 (QETU)**.** *Let $U = \mathrm{e}^{-i\tau A}$ be a unitary operator, where A is an n-qubit Hermitian operator and τ is real parameter. For any real even polynomial $F(x)$ of degree d satisfying $|F(x)| \leq 1, \forall x \in [-1, 1]$, one can find a sequence of symmetric phase factors $(\varphi_0, \ldots, \varphi_d) \in \mathbb{R}^{d+1}$ such that the circuit in Fig. 3.5 denoted by $\mathcal{U}$ satisfies $(\langle 0| \otimes \mathbb{1})\mathcal{U}(|0\rangle \otimes \mathbb{1}) = F(\cos(\tau A/2))$.*

In practice, the QETU circuit is used for approximately implementing $\mathsf{f}(A)$ by realizing a transformation $F(\cos(\tau A/2))$, where $F(x)$ is a polynomial approximation to $\mathsf{F}(x) \equiv \mathsf{f}(\frac{2}{\tau}\arccos(x))$.

Similar to QSVT, the phases $\{\tilde{\varphi}_j\}$ appearing in the circuit in Fig. 3.5 can be calculated using QSPPACK [186]. As described in Appendix B of [118], the phases in the QETU circuit $\{\tilde{\varphi}_j\}$ are related to the phases calculated using QSPPACK $\{\tilde{\phi}_j\}$ by $\tilde{\varphi}_j = \tilde{\phi}_j + (2 - \delta_{j\,0})\pi/4$.

Completely analogous to the case with QSVT circuits, controlled calls to QETU circuits can also be implemented for a small additional overhead. This property can be used to combine BEs of different local terms in the Hamiltonian that were constructed using QETU; the procedure for constructing BEs using QETU is given in Sec. 3.3.4.

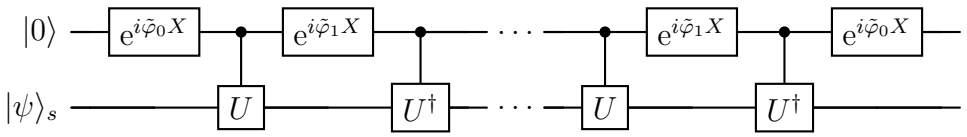

Figure 3.5: Circuit diagram for QETU. The top qubit is the control qubit and the bottom register is the state that the matrix function is applied to. Here $U = \mathrm{e}^{-i\tau A}$ is the time evolution circuit. If the control qubit is measured to in the zero state, one prepares the normalized quantum state $F(\cos(\tau A/2))|\psi\rangle/||F(\cos(\tau A/2))|\psi\rangle||$ for some even polynomial $F(x)$. For symmetric phase factors $(\tilde{\varphi}_0, \tilde{\varphi}_1, \ldots, \tilde{\varphi}_1, \tilde{\varphi}_0) \in \mathbb{R}^{d+1}$, the function $F(x)$ is a real even polynomial of degree d. The probability of measuring the control qubit in the zero state is $p = ||F(\cos(\tau A/2))|\psi\rangle||^2$.

3.3 Block encodings for a bosonic lattice field theory

In this section we describe how to construct block encodings for certain types of bosonic lattice field theories, and how to use those to simulate time evolution using GQSP. Our choice of GQSP is guided by the fact that, while this algorithm has the same asymptotic complexity as QSP or QSVT, but it comes with a smaller constant prefactor [110, 185].

The general strategy for using GQSP to simulate time evolution of the Hamiltonian in Eq. (3.1) is achieved with the following steps:

1. Construct the BEs $U_{f_k(\hat{\xi}_k)}$ for the individual terms of the form $f_k(\hat{\xi}_k)$.[2]
2. Construct $U_{f_1(\hat{\pi}_i)} = \widehat{\mathrm{FT}}_i^\dagger U_{f_1\left(\hat{\pi}_i^{(D)}\right)} \widehat{\mathrm{FT}}_i$ using the Quantum Fourier transform circuit.
3. Use LCU to combine these site-local BEs to obtain the BE for the full Hamiltonian U_H.
4. Construct the qubitized walk operator W_H.
5. Simulate time evolution using the GQSP circuit in Fig. 3.2.

After introducing the Hamiltonian of the scalar bosonic lattice field theory and discuss the digitization of the local bosonic Hilbert space, we present three different techniques for preparing BEs for a broad class of operators appearing in site-local Hamiltonians encountered in various formulations of lattice field theories. In all three techniques the BE of $f_k(\hat{\xi}_k)$ will start with an encoding of a basic building block, which is then combined to form $f_k(\hat{\xi}_k)$.

The first techniques uses QSVT, and starts from the simplest building block possible, namely the operator $\hat{\xi}_k$. Since this operator is not unitary, it needs to be block-encoded. As

[2]Since the functional form of f_{k_ℓ} is the same $\forall \ell \in \{0, 1, \ldots, N-1\}$ and fixed $k \in \{0, 1, 2\}$, in what follows we will drop the ℓ subscript on f_k when the statement is true / applies $\forall \ell$.

we will see, this encoding can be done with resources that are linear in $\mathsf{n_q}$, but comes with a large prefactor. Constructing the final function $f_k(\hat{\xi}_k)$ adds a constant multiplicative factor on the resource requirement.

The second method uses QETU, and directly uses the unitary operator $\mathrm{e}^{-i\tau\hat{\xi}_k}$ as the basic building block. The building block of QETU can still be implemented with resources that are linear in $\mathsf{n_q}$, but removes the large pre-factors arising in the previous method. The downside is that constructing $f_k(\hat{\xi}_k)$ introduces an exponential dependence on $\mathsf{n_q}$. As we will see, despite the exponential dependence on $\mathsf{n_q}$, the absence of a large pre-factor makes this method more efficient for $\mathsf{n_q} \lesssim 6$.

The final method takes a different approach, and chooses a more complicated building block, which requires resources with exponential scaling in $\mathsf{n_q}$ for its construction. This building block is chosen such to make the construction of $f_k(\hat{\xi}_k)$ as easy as possible, and is in fact trivial. This gives better scaling than the second method, and in fact outpeforms all methods $\mathsf{n_q} \lesssim 11$. We call this final method Linear Operators Via Exponentiation and Linear Combination of Unitaries (LOVE-LCU).

The comparison of different BEs is summarized in Table 3.1.

Algorithm	**Gate Complexity**	**Ancillary qubits**	**Indefinite parity**	**Multivariable functions**
Standard LCU	$\mathcal{O}(\mathsf{n_q}^d)$	$\mathcal{O}(d \log \mathsf{n_q})$	Yes	Yes
QSVT	$\mathcal{O}(d \cdot \mathsf{n_q} \log \mathsf{n_q})$	$\mathcal{O}(\log \mathsf{n_q})$	No	No
QETU with $\mathrm{e}^{i\tau\hat{\xi}}$	$\mathcal{O}(2^{\mathsf{n_q}} \cdot \mathsf{n_q})$	1	No	No
QETU with $\mathrm{e}^{i2\arccos\hat{\xi}}$	$\mathcal{O}(d \cdot 2^{\mathsf{n_q}})$	1	No	No
LOVE-LCU	$\mathcal{O}(2^{\mathsf{n_q}})$	1	Yes	Yes

Table 3.1: Properties of BE methods considered in this work. Gate count and ancillary qubit count are for constructing a BE of a degree d function $f(\hat{\xi})$ of an n-qubit diagonal Hermitian operator $\hat{\xi}$ that is a sum of $\mathcal{O}(n)$ single Z-gates (for operators considered in this work, $\hat{\xi} \in \{\hat{\varphi}_i, \hat{\pi}^{(D)}, \hat{\varphi}_i - \hat{\varphi}_j\}$). The QSVT- and QETU-based methods require f to have definite parity and be a function of a single variable. Standard LCU and LOVE-LCU can be applied to functions with indefinite parity.

3.3.1 General Form of Hamiltonians of Interest

A common approach to formulating quantum field theories in a way suitable for solving those on quantum computers amounts to 1) discretizing the real space so that the fields were defined on a finite subset $\mathcal{V}$ of sites of some lattice, e.g., $\mathcal{V} \subset a\mathbb{Z}^d \subset \mathbb{R}^d$, subject to an

appropriate boundary condition; 2) digitizing quantum degrees of freedom residing on lattice sites, *i.e.*, truncating the infinite-dimensional local Hilbert space to a finite-dimensional one. While alternative approaches are also subject to active investigation [17, 88, 131], the aforementioned paradigm has two major advantages: the locality of the discretized Hamiltonian (for local quantum field theories) and linear growth of the Hamiltonian norm with problem size.[3] Both of these are of great utility from the quantum computing perspective. While fermionic lattice systems have been subject to active investigation in the literature (for recent studies, see [71, 172] and references therein), in this work we focus on lattice bosons, which are an essential ingredient in nuclear [187] and high-energy particle physics [72].

Many physically relevant Hamiltonians, including scalar field theories [52, 97], dual basis formulations of U(1) [91, 188] and SU(2) [93] gauge theories, are (at least partially) formulated in terms of conjugate operators defined on each lattice site. For this work, similar to Chapter 2, we use the notation of a scalar field theory where we denote by $\hat{\varphi}_i$ and $\hat{\pi}_i$ the "field" and "momentum" operators at site i, respectively, satisfying the canonical commutation relations $[\hat{\varphi}_i, \hat{\pi}_j] = i\delta_{ij}$. The general idea is to exploit the fact that, for a particular highly efficient digitization scheme of such operators (similar to that from Chapter 2), the local operators $\hat{\varphi}_i$ and $\hat{\pi}_i$ take on simple forms and can therefore be used as low cost building blocks in the QSVT or QETU circuits. These can then be used to prepare low cost BEs of the individual terms appearing in the Hamiltonian; and the BE for the full Hamiltonian can be prepared by using LCU to combine the BEs of the site-local terms.

To be concrete, we consider Hamiltonians of the form

$$\hat{H} = \sum_{i=1}^{\mathsf{N}_\Lambda} \left(f_0\left(\hat{\varphi}_i\right) + f_1\left(\hat{\pi}_i\right) \right) + \sum_{\langle ij \rangle} f_2\left(\hat{\varphi}_i - \hat{\varphi}_j\right), \tag{3.8}$$

where N_Λ is the number of lattice sites and $\sum_{\langle ij \rangle}$ is a sum over all adjacent sites i and j. A discussion of applying our methods to more general forms of the Hamiltonian is given in Sec. 3.5.

To *digitize* the theory, the local bosonic Hilbert space at each lattice site j is represented using n_q qubits. Following Refs. [52, 97, 159], we work in the eigenbasis of the digitized field operators $\hat{\varphi}_j$, *i.e.*, we identify the computational basis states $|k\rangle_j$ at site j with the eigenstates $|\varphi_j^{(k)}\rangle$. The eigenvalues of the field operator $\hat{\varphi}_j$ are chosen as

$$\begin{gathered} \hat{\varphi}_j |\varphi_j^{(k)}\rangle = (-\varphi_{\max} + k\delta\varphi) |\varphi_j^{(k)}\rangle \,, \\ k = 0, \ldots, 2^{\mathsf{n}_\mathsf{q}} - 1 \,, \end{gathered} \tag{3.9}$$

where $\delta\varphi = 2\varphi_{\max}/(2^{\mathsf{n}_\mathsf{q}} - 1)$. The free parameter $\varphi_{\max}$ can be chosen to minimize the digitization errors and is theory dependent. Generally, by choosing $\varphi_{\max}$ to scale exponentially with n_q [52, 91, 93, 97, 136, 159], the digitization errors decrease exponentially with n_q.

[3] Here we do not touch upon the issue of renormalization which may affect the Hamiltonian norm via adjustment of coupling constants.

This fact implies that the number of qubits per site $\mathsf{n_q}$ can generally be kept small, which is important to keep in mind for the discussion moving forward.

Exploiting the fact that $\hat{\pi}_j$ is conjugate to $\hat{\varphi}_j$, the momentum operators are digitized as

$$\begin{aligned} \hat{\pi}_j^{(D)}|\pi_j^{(k)}\rangle = (-\pi_{\max} + k\delta\pi)|\pi_j^{(k)}\rangle\,, \\ k = 0, \ldots, 2^{\mathsf{n_q}} - 1\,, \end{aligned} \tag{3.10}$$

where

$$\hat{\pi}_j^{(D)} = \widehat{\mathrm{FT}}_j \hat{\pi}_j \widehat{\mathrm{FT}}_j^\dagger \tag{3.11}$$

is the conjugate momentum operator in the momentum basis, $\widehat{\mathrm{FT}}_j$ denotes the usual discrete Fourier transform (FT) at site j, $\pi_{\max} = \pi/\delta\varphi$ and $\delta\pi = 2\pi_{\max}/(2^{\mathsf{n_q}} - 1)$. The Hamiltonian in this digitization is then written through functions of the diagonal operators $\hat{\varphi}_j$ and $\hat{\pi}_j^{(D)}$ as

$$\hat{H} = \sum_{i=1}^{\mathsf{N}_\Lambda} \Big(f_0\left(\hat{\varphi}_i\right) + \widehat{\mathrm{FT}}_i^\dagger f_1(\hat{\pi}_i^{(D)}) \widehat{\mathrm{FT}}_i \Big) + \sum_{\langle ij \rangle} f_2\left(\hat{\varphi}_i - \hat{\varphi}_j\right). \tag{3.12}$$

With regards to implementing the field operators and their functions, we note that the operators $\hat{\varphi}_j$ and $\hat{\pi}_j^{(D)}$ are $2^{\mathsf{n_q}}$-dimensional diagonal operators with evenly spaced entries on the diagonal and hence can be expressed as a linear combination of $\mathsf{n_q}$ single-qubit Pauli-Z gates as shown in [97, 189] and Eq. (2.105). Aside from the overall prefactor, the operator $\hat{\pi}^{(D)}$ has an identical gate representation. Additionally, given that the operators $\hat{\xi}_k$ are diagonal, one can readily determine complicated functions $f_k(\hat{\xi}_k)$ using classical resources by simply applying the function f_k to the diagonal elements of $\hat{\xi}_k$.

3.3.2 LCU to Block Encode Individual Terms

In this section, we review the cost of using standard LCU techniques to BE the individual $f_k(\hat{\xi}_k)$ terms. For a more detailed analysis, see Chapter 2 and references therein. Consider first using LCU to BE $f_0(\hat{\varphi})$ and $f_1(\hat{\pi}^{(D)})$ From Eq. (2.105), we see that a general degree d function of $\hat{\varphi}$ (or $\hat{\pi}^{(D)}$) is a sum of $\mathcal{O}(\mathsf{n_q}^d)$ Pauli strings. It is important to note that the number of Pauli strings is upper bounded by $2^{\mathsf{n_q}}$. This, combined with Lemma 8, implies that the asymptotic gate cost to BE $f_0(\hat{\varphi})$ and $f_1(\hat{\pi}^{(D)})$ to be either $\mathcal{O}(\mathsf{n_q}^d \log \mathsf{n_q})$ or $\mathcal{O}(\mathsf{n_q} 2^{\mathsf{n_q}})$, depending on the degree d and value of $\mathsf{n_q}$ used[4]. Turning to the final term, for a degree d function, $f_2(\hat{\varphi}_i - \hat{\varphi}_j)$ is a sum of $\mathcal{O}(\mathsf{n_q}^d)$ terms. In this case, the number of Pauli strings is upper bounded by $2^{2\mathsf{n_q}}$. Using similar arguments, the asymptotic gate cost to BE $f_2(\hat{\varphi}_i - \hat{\varphi}_j)$ is either $\mathcal{O}(\mathsf{n_q}^d \log \mathsf{n_q})$ or $\mathcal{O}(\mathsf{n_q} 2^{2\mathsf{n_q}})$. Lastly, we discuss the scale factor when using LCU.The scale factor is the sum of the absolute values of the coefficients of the unitary terms being added together.Given a particular $f_k(\hat{\varphi}_k)$, the scale factor can be readily determined using using Eq. (2.105). The scale factors for the operators considered in Sec. 3.4 are derived in App. B.3.

[4]Given that the operators $\hat{\xi}_k$ are diagonal, one can readily determine complicated functions $f_k(\hat{\xi}_k)$ using classical resources by simply applying the function f_k to the diagonal elements of $\hat{\xi}_k$.

3.3.3 QSVT to Block Encode Individual Terms

In the context of quantum simulation, the QSVT algorithm is typically utilized as a way of constructing polynomial functions of Hamiltonian operator [116, 117]. In this section we use it to construct the functions $f_k(\hat{\xi}_k)$ and demonstrate that QSVT can also be used as a highly-efficient tool for constructing block encodings. The high level flow of this technique is depicted in Fig. 3.1. For this discussion we assume the functions f_k to have even parity.

The building block in the QSVT circuit is the BE of the simple operator $\hat{\xi}_k$. From Eq. (2.105) we see that each $\hat{\xi}_k$ is written as a sum of $\mathcal{O}(\mathsf{n_q})$ Pauli-Z gates, each of which are unitary. The BE $U_{\hat{\xi}_k}$ can therefore be constructed with LCU using $\mathcal{O}(\mathsf{n_q}\log\mathsf{n_q})$ gates and $\mathcal{O}(\log\mathsf{n_q})$ ancillary qubits.

If f_k is a polynomial function, QSVT can prepare the BE U_{f_k} using a constant number of calls to $U_{\hat{\xi}_k}$. For non-polynomial functions, one can express the function $f_k(\hat{\xi}_k)$ as a (potentially infinite) Chebyshev series of the operator $\hat{\xi}_k$, and then use QSVT to construct the BEs U_{f_k}.

To determine the Chebyshev expansion, we write

$$\sum_{j=0} c_{2j} T_{2j}\left(\frac{\hat{\xi}_k}{\alpha}\right) = \frac{1}{\beta} f_k(\hat{\xi}_k)\,, \tag{3.13}$$

where α is the scale factor used in the BE of $\hat{\xi}_k$ and β normalizes the function $f_k(\hat{\xi}_k)$ which can be obtained classically. The upper limit of the sum has been left blank to stress the series can be a sum of a finite or infinite number of terms. The coefficients c_{2j} can be then be found using the orthogonality relations of Chebyshev polynomials. The gate complexity depends on the convergence of the Chebyshev polynomial approximation, and for a degree-d polynomial the gate complexity is therefore given by using $\mathcal{O}(d\,\mathsf{n_q}\log\mathsf{n_q})$. The scale factor of the BE is $\beta = \|f_k(\hat{\xi}_k)\|$.

We conclude by pointing out that using $\hat{\varphi}_i$ and $\hat{\pi}_i$ as building blocks to construct a BE of the Hamiltonian is not a unique choice. One could equivalently use, e.g., $\hat{\varphi}_i^2$ and $\hat{\pi}_i^2$, and modify the Chebyshev polynomial used accordingly. However, because the number of terms, as well as the length of the largest Pauli string, in $\hat{\varphi}_i^n$ (and $\hat{\pi}_i^n$) scales as $\mathcal{O}(\mathsf{n_q^n})$, using the simpler building blocks $\hat{\varphi}_i$ and $\hat{\pi}_i$ will result in better asymptotic scaling with $\mathsf{n_q}$. We leave investigations of alternative constructions for future work.

3.3.4 QETU to Block Encode Individual Terms

While QETU was initially proposed to eliminate the need for block encoding in early fault-tolerant quantum algorithms [118], [101] has investigated its utility for preparing functions of Hermitian operators, using Gaussian states as an example. Below, we show that QETU can be utilized for constructing block encodings for Hamiltonians of the form Eq. (3.8).

Preparing BEs using QETU is conceptually similar to using QSVT. Following the procedure outlined in Fig. 3.1, the building block used is $\mathrm{e}^{-i\tau\hat{\xi}_k}$, which is the simplest unitary

operator containing information about $\hat{\xi}_k$. By performing repeated calls to $\mathrm{e}^{-i\tau\hat{\xi}_k}$, QETU can prepare a BE of $f_k(\hat{\xi}_k)$.

The remainder of this section is split into two parts. We first describe the technical details of using QETU to prepare general matrix functions. From there, we discuss the computational cost of preparing BEs of $f_k(\hat{\xi}_k)$.

3.3.4.1 Technical details

We begin by discussing QETU for a general function $f(\hat{A})$ of a Hermitian matrix $\hat{A}$ (not necessarily diagonal) with known eigenvalues $\hat{A}|a_j\rangle = a_j|a_j\rangle$, and will specialize to $f_k(\hat{\xi}_k)$ later. We will call the minimum and maximum eigenvalues $a_{\min}$ and $a_{\max}$, respectively. QETU allows to start from an implementation of the operator $\mathrm{e}^{-i\tau\hat{A}}$ and implement a BE of the function $F(\cos(\tau\hat{A}/2))$. Since one needs to shift the spectrum of the operator $\tau\hat{A}$ to lie in the range $[0, 2\pi]$, we therefore define a shifted operator

$$\hat{A}_{\mathrm{sh}} = c_1\hat{A} + c_2\mathbb{1}\,, \tag{3.14}$$

where

$$c_1 = \frac{\pi}{a_{\max} - a_{\min}}\,, \quad c_2 = -c_1 a_{\min}\,, \tag{3.15}$$

and require $\tau \leq 2$. In order to implement the desired function $f_k(\hat{A})/\beta$ (where β is required to normalize the function to allow a BE), one needs to choose

$$\mathsf{F}(x) = \frac{1}{\beta}\mathsf{f}\left(\frac{\frac{2}{\tau}\arccos(x) - c_2}{c_1}\right). \tag{3.16}$$

Note that the direct implementation of QETU applies only in cases when $\mathsf{F}(x)$ has definite parity. Since the function in Eq. (3.16) does not generally have definite parity, one would need LCU to combine even and odd functions together to obtain the general function $F(\hat{A})$. However, as shown in [101], the need for LCU to add the even and odd pieces can, in certain cases, be avoided altogether by choosing by choosing τ in such a way as to make $\mathsf{F}(x)$ have definite parity. To see this, note that if one chooses $\tau = \pi/c_2$, the function $\mathsf{F}(x)$ becomes

$$\mathsf{F}(x)\Big|_{\tau=\pi/c_2} = \frac{1}{\beta}\mathsf{f}\left(-\frac{2}{\pi}\frac{c_2}{c_1}\arcsin(x)\right) = \frac{1}{\beta}\mathsf{f}\left(\frac{2a_{\min}}{\pi}\arcsin(x)\right) \tag{3.17}$$

where we have used $\arcsin(x) = \pi/2 - \arccos(x)$. Because $\arcsin(x)$ has odd parity, the function $\mathsf{F}(x)$ has the same parity as the function $\mathsf{f}(x)$. This means that LCU is not required if the original function $f_k(\hat{A})$ one aims to implement has definite parity. Note if $\tau > 2$, this technique cannot be applied and one must construct the even and odd pieces separately. Importantly, all operators $\hat{\xi}_k$ considered in this work are digitized such that their minimum and maximum eigenvalues are equal, which implies $\tau = 2$.

Next, we consider the scale factor β, which is given by $\beta = \max_{x\in[-1,1]} |F(x)|$. Using the fact that $\arcsin(x) \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, we can rewrite the scale factor as

$$\beta = \max_{y\in[-a_{\min},a_{\min}]} |f(y)| . \tag{3.18}$$

We will use this to determine the scale factor for the BEs of interest in the next section.

The final observation we make regarding $\mathsf{F}(x)$ is that, for general functions $\mathsf{f}(x)$, its first derivative diverges at $x = \pm 1$ due to the presence of the $\arcsin(x)$ (or $\arccos(x)$) term. Chebyshev approximations of functions with m continuous derivatives on $x \in [-1, 1]$ are known to converge polynomially to the true function with the typical rate $(1/n_{\text{Ch}})^m$, where n_{Ch} is the number of Chebyshev polynomials used in the approximation [190]. This is to be contrasted with the exponential convergence of Chebyshev series to infinitely differentiable functions. If one needed to reproduce $\mathsf{F}(x)$ for all $x \in [-1, 1]$, this poor convergence could lead to a large gate cost.

As was shown in [101] this issue can be overcome by exploiting the fact that the spectrum of the operator $\hat{A}_{\text{sh}}$ is known exactly and therefore one only has to reproduce $\mathsf{F}(x)$ at those finitely many points. If $\hat{A}_{\text{sh}}$ is represented using $\mathsf{n_q}$ qubits, there are $\mathcal{O}(2^{\mathsf{n_q}})$ such values, and one can therefore reproduce the function at those points exactly by using $\mathcal{O}(2^{\mathsf{n_q}})$ Chebyshev polynomials. In this way, one avoids the poor polynomial convergence of the Chebyshev expansion, at the cost of introducing an exponential scaling with $\mathsf{n_q}$. This exponential scaling, however, does not pose a problem as long we work with small values of $\mathsf{n_q}$ and as previously discussed, this is precisely the case in lattice field theories owing to the fact that the digitization errors generally decrease exponentially with $\mathsf{n_q}$.

3.3.4.2 Cost to prepare block-encodings of bosonic operators

We now describe the cost of using QETU to prepare BEs of $f_k(\hat{\xi}_k)$, where, again, each f_k is assumed to have even parity. After constructing the shifted operator $\hat{\xi}_{k,\text{sh}}$ according to Eq. (3.14), the building block in the QETU circuit is a simple operator $e^{-i\tau\hat{\xi}_{k,\text{sh}}}$. Because each $\hat{\xi}_{k,\text{sh}}$ is still a diagonal operator with evenly spaced eigenvalues (or a Kronecker sum of two such operators in the case of $\hat{\xi}_2 \in \{\hat{\varphi}_i - \hat{\varphi}_j\}$), they can be written as a sum of $\mathcal{O}(\mathsf{n_q})$ single-qubit Pauli-Z gates, and $e^{-i\tau\hat{\xi}_{k,\text{sh}}}$ can be implemented using only $\mathcal{O}(\mathsf{n_q})$ single-qubit R_z gates. A single controlled call to $e^{-i\tau\hat{\xi}_{k,\text{sh}}}$ therefore requires $\mathcal{O}(\mathsf{n_q})$ of R_z and CNOT gates. As discussed, QETU can prepare the BE U_{f_k} exactly using a polynomial of degree equal to the number of times we sample the function f_k.

The operators $\hat{\varphi}_i$ and $\hat{\pi}_i^{(D)}$ are sampled $2^{\mathsf{n_q}}$ times, implying one must use a degree-$\mathcal{O}(2^{\mathsf{n_q}})$ polynomial to reproduce $\mathsf{F}_0(\hat{\varphi}_i) \sim \mathsf{f}_0(\arcsin(\hat{\varphi}_i))$ and $\mathsf{F}_1(\hat{\pi}_i^{(D)}) \sim \mathsf{f}_1(\arcsin(\hat{\pi}_i^{(D)}))$ exactly. Note that one can reduce the degree of the polynomial by a factor of two by exploiting the fact that the operators are sampled symmetrically about zero and that f_0 and f_1 are even functions. Therefore, the asymptotic gate cost of preparing the BEs U_{f_0} and U_{f_1} using QETU is $\mathcal{O}(\mathsf{n_q}2^{\mathsf{n_q}})$, independent of the degree d of the polynomial of f_0 and f_1. To determine the scale factor, note that the minimum eigenvalue of $\hat{\varphi}$ ($\hat{\pi}$) is $-\varphi_{\max}$ ($-\pi_{\max}$). Using Eq. (3.18),

the scale factor for U_{f_0} (U_{f_1}) is therefore $\max_{x\in[-\varphi_{\max},\varphi_{\max}]}|f_0(x)|$ ($\max_{x\in[-\pi_{\max},\pi_{\max}]}|f_1(x)|$. For many systems of physical interest, including scalar field theories (considered in Sec. 3.4), the maximum value of f_0, f_1 occur at the largest argument they are evaluated. In this situation, the scale factors for U_{f_0} and U_{f_1} are given by the smallest possible values, namely $\|f_0(\hat{\varphi})\|$ and $\|f_1(\hat{\varphi})\|$, respectively.

The final operator we need to block-encode is $\mathsf{f}_2(\hat{\varphi}_i - \hat{\varphi}_j)$. Because $(\hat{\varphi}_j - \hat{\varphi}_i)$ is a $2^{2\mathsf{n}_\mathsf{q}}$-dimensional matrix, one would naively expect that $\mathcal{O}(2^{2\mathsf{n}_\mathsf{q}})$ Chebyshev polynomials are needed to exactly reproduce $\mathsf{F}_2(\hat{\varphi}_j - \hat{\varphi}_i)$ at all the $2^{2\mathsf{n}_\mathsf{q}}$ values it is sampled. However, due to the symmetric digitization of $\hat{\varphi}$, the matrix $(\hat{\varphi}_j - \hat{\varphi}_i)$ only has $\mathcal{O}(2^{\mathsf{n}_\mathsf{q}})$ unique values, and can be implemented exactly using $\mathcal{O}(2^{\mathsf{n}_\mathsf{q}})$ Chebyshev polynomials. The asymptotic gate complexity of preparing the BE U_{f_2} using QETU is therefore $\mathcal{O}(\mathsf{n}_\mathsf{q}2^{\mathsf{n}_\mathsf{q}})$, independent of the degree d of the polynomial of f_2. Because the minimum eigenvalue of $(\hat{\varphi}_j - \hat{\varphi}_i)$ is $-2\varphi_{\max}$, the scale factor is $\max_{x\in[-2\varphi_{\max},2\varphi_{\max}]}$. As before, for many physical systems, the scale factor is given by the minimum value of $\|f_2(\hat{\varphi}_i - \hat{\varphi}_j)\|$.

Additional gate cost savings can be achieved if one uses the control-free version of QETU [118]. In practice, doing so reduces the number of R_z gates in the controlled call to $e^{-i\tau\hat{\xi}_{k,\mathrm{sh}}}$ by a factor of 2. A detailed procedure for implementing the control-free version of QETU for diagonal operators can be found in [101].

3.3.5 LOVE-LCU

In the previous section, we showed that, by employing $\mathrm{e}^{i\tau\hat{\xi}_k}$ as a building block, QETU can be used to construct BEs of $f_k(\hat{\xi}_k)$ with an asymptotic gate cost $\mathcal{O}(\mathsf{n}_\mathsf{q}2^{\mathsf{n}_\mathsf{q}})$. In this section, we study how the cost changes if different building blocks are utilized. First, we will argue that, in general, it is not possible to eliminate the exponential scaling with n_q. With this fundamental limitation in mind, we show that using more complicated building blocks can reduce the asymptotic cost to prepare a BE. Specifically, the cost of preparing $f(\hat{\xi}_k)$ is reduced from $\mathcal{O}(\mathsf{n}_\mathsf{q}2^{\mathsf{n}_\mathsf{q}})$ to $\mathcal{O}(d2^{\mathsf{n}_\mathsf{q}})$, where d is the degree of f. Finally, extending this logic further, we develop a conceptually simple framework that does not require QETU, and can prepare BEs of $f(\hat{\xi}_k)$ using $\mathcal{O}(2^{\mathsf{n}_\mathsf{q}})$ gates, independent of the degree d of f.

3.3.5.1 Main idea

To begin the discussion, recall that the exponential scaling in n_q arises from the need to work with functions whose derivatives diverge, such as $\arccos(x)$ or $\arcsin(x)$. This result generalizes, and any building block that requires working with these functions will exhibit the same exponential scaling.

To understand what functions fall into this class, consider the general building block $\mathrm{e}^{-i\tau g(\hat{\xi}_k)}$. Using this building block, QETU returns $F(\cos(\frac{\tau g(\hat{\xi}_k)}{2}))$, which implies that the function we need to implement using Chebyshev polynomials is $F(x) \sim f(g^{-1}(\frac{2}{\tau}\arccos(x)))$. Any function $g(x)$ that results in $f(g^{-1}(\frac{2}{\tau}\arccos(x)))$ having divergent derivatives for $x \in$

$[-1,1]$ will consequently require $\mathcal{O}(2^{\mathsf{n_q}})$-degree Chebyshev polynomials to produce $F(x)$ exactly.

Consider first choosing $g(x) = x^m$ for $m = \{1, 2, \dots\}$. As was the case with $m = 1$, the function $F(x)$ diverges as $x = \pm 1$ for, and therefore requires $\mathcal{O}(2^{\mathsf{n_q}})$ Chebyshev polynomials to reproduce $F(x)$ exactly. Because each call to $\mathrm{e}^{-i\tau g(\hat{\xi}_k)}$ requires $\mathcal{O}(\mathsf{n_q}^m)$ gates, the cost of preparing a BE of $f(\hat{\xi}_k)$ using $\mathrm{e}^{-i\tau g(\hat{\xi}_k)}$ as a building block is $\mathcal{O}(\mathsf{n_q}^m 2^{n_q})$.

Consider now using non-polynomial $g(x)$. Further suppose that some $g(x)$ exists (which we discuss next) that results in $F(x) \sim f(g^{-1}(\frac{2}{\tau}\arccos(x)))$ being infinitely differentiable. While this choice may reduce the number of Chebyshev polynomials required, the cost of implementing $\mathrm{e}^{-i\tau g(\hat{\xi}_k)}$ for non-polynomial $g(x)$ is generally $\mathcal{O}(2^{\mathsf{n_q}})$.

Taken together, these arguments imply that the cost of preparing BEs of $f_k(\hat{\xi}_k)$ using QETU generally scales exponentially with $\mathsf{n_q}$. With this in mind, the rest of this section focuses on finding a $g(x)$ that reduces the cost from $\mathcal{O}(\mathsf{n_q}2^{\mathsf{n_q}})$ to $\mathcal{O}(2^{\mathsf{n_q}})$.

3.3.5.2 Technical details

Our objective is to identify $g(x)$ such that $F(x) \sim f(g^{-1}(\frac{2}{\tau}\arccos(x)))$ is infinitely differentiable. The form of $F(x)$ suggests that $g(x) \sim \arccos(x)$ is a good starting point. The easiest way to accomplish this is to use the building block $\mathrm{e}^{-i2\arccos(\hat{\xi}_k/\alpha)}$, where $\alpha = \|\hat{\xi}_k\|$ to ensure the arccos is well defined. Setting $\tau = 1$ and replacing $\hat{A}$ in Theorem 27 with $2\arccos(\hat{\xi}_k/\alpha)$, we observe that the QETU circuit now returns $F(\hat{\xi}_k/\alpha)$, which is only a function of $\hat{\xi}_k$. The procedure for determining the Chebyshev expansion of $f_k(\hat{\xi}_k)$ is now identical to that explained in Sec. 3.2.2; degree-d polynomial functions of $\hat{\xi}_k$ can be implemented using $\mathcal{O}(d)$ controlled calls to $\mathrm{e}^{-i2\arccos(\hat{\xi}_k/\alpha)}$.

Implementing $\mathrm{e}^{-i2\arccos(\hat{\xi}_k/\alpha)}$, however, is more complicated than $\mathrm{e}^{-i\tau\hat{\xi}_k}$. While the operator $\arccos(\hat{\xi}_k/\alpha)$ is still diagonal, and can easily be determined classically, it is an infinite degree Chebyshev polynomial in its argument $\hat{\xi}_k/\alpha$. This implies that its Pauli decomposition consists of $\mathcal{O}(2^{\mathsf{n_q}})$ identity and R_z gates [189] and a controlled implementation of the operator $\mathrm{e}^{-i2\arccos(\hat{\xi}_k/\alpha)}$ requires $\mathcal{O}(2^{\mathsf{n_q}})$ CNOT and R_z gates. On the other hand, constructing degree-d polynomial functions $f_k(\hat{\xi}_k)$ out of this building block only requires an additional multiplicative factor of d, giving an overall asymptotic scaling of $\mathcal{O}(d2^{\mathsf{n_q}})$.

We thus see that choosing more complicated building blocks, such as $\mathrm{e}^{-i2\arccos(\hat{\xi}_k/\alpha)}$ which requires exponential (in $\mathsf{n_q}$) resources to implement, can potentially make constructing the functions $f_k(\hat{\xi}_k)$ simpler. One such method that we propose in this section, which we call Linear Operators Via Exponentiation and Linear Combination of Unitaries (LOVE-LCU), uses $\mathrm{e}^{\pm i\arccos(f_k(\hat{\xi}_k)/\beta)}$ as a building block. Since

$$\frac{f_k(\hat{\xi}_k)}{\beta} = \frac{\mathrm{e}^{i\arccos\left(\frac{f_k(\hat{\xi}_k)}{\beta}\right)} + \mathrm{e}^{-i\arccos\left(\frac{f_k(\hat{\xi}_k)}{\beta}\right)}}{2}, \tag{3.19}$$

the construction of $f_k(\hat{\xi}_k)$ only needs LCU to combine these two terms. This can be implemented using a simple LCU circuit, shown in Fig. 3.6, with $\text{PREP} = \text{H} \otimes \mathbb{1}_s$ and

$$\text{SEL} = |0\rangle\langle 0| \otimes \mathrm{e}^{i \arccos\left(\frac{f_k(\hat{\xi}_k)}{\beta}\right)} + |1\rangle\langle 1| \otimes \mathrm{e}^{-i \arccos\left(\frac{f_k(\hat{\xi}_k)}{\beta}\right)}. \tag{3.20}$$

The scale factor β can be chosen as the smallest possible value $\beta = \|f_k(\hat{\xi}_k)\|$.

The construction of the building block is very similar to before. The operators $\mathrm{e}^{\pm i \arccos\left(f_k(\hat{\xi}_k)/\beta\right)}$ are diagonal, which can be determined in a straightforward way using classical resources. It can then be decomposed into Pauli strings containing only the identity or Z-gates; in general, the decomposition will contain $\mathcal{O}(2^{\mathsf{n_q}})$ Pauli strings. The circuits for the diagonal unitary matrices $\mathrm{e}^{\pm i \arccos(\hat{A})}$ can then be implemented using at most[5] $2^{\mathsf{n_q}} - 1$ R_z gates and $2^{\mathsf{n_q}} - 2$ CNOT gates [189], and the entire LOVE-LCU circuit requires $\mathcal{O}(2^{\mathsf{n_q}})$ gates and a single ancillary qubit. This gate count can be further reduced by approximating the circuits for $\mathrm{e}^{i \arccos \hat{A}}$ in a systematic way by dropping rotation gates with small angles [68, 189]. LOVE-LCU can therefore prepare BEs of $f_k(\hat{\xi}_k)$ with resources that are independent of the degree d of the function f.

Another advantage of LOVE-LCU is that it is not restricted to BEs of single variables, but can also be applied to functions of multiple variables, *i.e.* the gate cost of using LOVE-LCU to construct a BE of the operator $f(\hat{\varphi}_1 - \hat{\varphi}_2)$ is the same as a general function $f(\hat{\varphi}_1, \hat{\varphi}_2)$. This implies that, instead of preparing BEs of $f_0(\hat{\varphi}_i)$, $f_0(\hat{\varphi}_j)$ and $f_2(\hat{\varphi}_i - \hat{\varphi}_j)$ separately and then adding them using another layer of LCU, one can directly use LOVE-LCU to prepare a BE of the sum $f_0(\hat{\varphi}_i) + f_0(\hat{\varphi}_j) + f_2(\hat{\varphi}_i - \hat{\varphi}_j)$.

The cost of LOVE-LCU can be further reduced by exploiting the structured form of the SEL oracle. In particular, SEL oracles of the form in Eq. (3.20) can be prepared with the techniques used in the control-free implementation of QETU. Rather than performing two controlled calls to $\mathrm{e}^{-i \arccos(f(\hat{\xi}_k/\alpha))}$, one can construct SEL for the cost of a single non-controlled circuit for $\mathrm{e}^{-i \arccos(f(\hat{\xi}_k/\alpha))}$ and an additional $2\mathsf{n_q}$ CNOT gates; the procedure for this control-free implementation can be found in Appendix B of [101].

While here we focused on using LOVE-LCU for block encoding diagonal operators, a similar construction can be applied to general Hermitian operators. The procedure is similar to the diagonal operator case; for an $\mathsf{n_q}$-qubit operator $\hat{A}$, one would evaluate $\mathrm{e}^{-i \arccos(\hat{A})}$ classically, and then exactly decompose the unitary operator into $\mathcal{O}(4^{\mathsf{n_q}})$ gates [149].

3.3.6 Constructing full Hamiltonian Block Encoding and Walk Operator

The previous sections discussed methods for constructing BEs of local terms in the Hamiltonian in Eq. (3.8). In this section, we discuss how to construct a BE of the full Hamiltonian

[5] If the operator $\hat{A}$ has definite parity, for example $\hat{A} = \hat{\varphi}^n$ for some polynomial n, then only half of the Pauli strings will have non-zero coefficients. This can be understood by noting that $\arccos(x) = \pi/2 - \arcsin(x)$, which implies that $\arccos(x)$ has the same parity of the argument up to an overall constant shift.

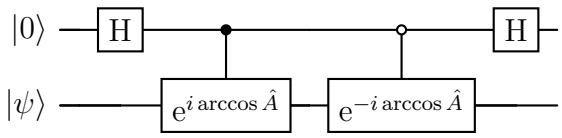

Figure 3.6: LCU circuit for block encoding of operator $\hat{A}$.

$U_{\hat{H}}$. Furthermore, we show that there exists a simple operator S satisfying the conditions in Eqs. (2.46a) and (2.46b), which can be used to construct the walk operator $W_{\hat{H}}$ while avoiding the large overall prefactor when using the general qubitization procedure in Ref [109]. Here we only provide a high-level procedure; the detailed proofs are given in App. B.1.

For this discussion, we write the Hamiltonian in the general form $\hat{H} = \sum_\ell \hat{H}_\ell$, where $\hat{H}_\ell$ is a local term in the full Hamiltonian $\hat{H}$ that it is diagonalized by an efficiently implementable unitary (see Eq. (3.2)). Each $\hat{H}_\ell^{(D)}$ is then given by

$$\hat{H}_\ell^{(D)} = f_{k_\ell}(\hat{\xi}_{k_\ell}) \tag{3.21}$$

The BE of the full Hamiltonian $U_{\hat{H}}$ is constructed using a standard LCU procedure to add BEs of local terms $U_{\hat{H}_\ell}$. We denote by χ_ℓ the number of gates required to construct a given $U_{\hat{H}_\ell}$, and define $\chi_{\text{site}} = \max_\ell \chi_\ell$. Using LCU to add each of the $\mathcal{O}(\mathsf{N}_\Lambda)$ local BEs $U_{\hat{H}_\ell}$ requires $\mathcal{O}(\log \mathsf{N}_\Lambda)$ ancillary qubits, the register of which we denote by a_Λ. Since the local terms $f_{k_\ell}(\hat{\xi}_{k_\ell})$ for a given k_ℓ all have identical coefficients, the PREP oracle will have a relatively simple circuit implementation and will be sub-leading in the asymptotic gate cost. The SEL oracle is given by $\text{SEL} = \sum_\ell |\ell\rangle\langle\ell| \otimes U_{\hat{H}_\ell}$, where each $U_{\hat{H}_\ell}$ is controlled on $\mathcal{O}(\log \mathsf{N}_\Lambda)$ qubits, each requiring $\mathcal{O}(\chi_{\text{site}} \log \mathsf{N}_\Lambda)$ gates. The gate cost of the full SEL oracle, and therefore the gate cost of constructing $U_{\hat{H}}$, is $\mathcal{O}(\chi_{\text{site}} \mathsf{N}_\Lambda \log \mathsf{N}_\Lambda)$.

We now describe how to construct the walk operator $W_{\hat{H}}$ in a way that avoids the expensive general qubitization procedure in [109]. The first step is to determine operators S for individual BEs $U_{\hat{H}_\ell}$ constructed using the methods in this work. Using this information, we then identify an S that can be used to construct $W_{\hat{H}}$ from the BE $U_{\hat{H}}$.

As shown in detail in App. B.1, for BEs $U_{\hat{H}_\ell}$ constructed using QSVT or QETU for even polynomials, $W_{\hat{H}_\ell}$ can be constructed by choosing S to be a single Z-gate acting on the signal and control qubit, respectively; this is due to the highly symmetric structure of QSVT and QETU circuits. Turning to LOVE-LCU, choosing S to be a single Z-gate acting on the ancillary qubit in the LOVE-LCU circuit also satisfies the relations in Eq. (2.46) necessary to construct $W_{\hat{H}_j}$. To see this, first observe that $\mathrm{Z}|0\rangle = |0\rangle$, which implies that Eq. (2.46a) is automatically satisfied. If we denote by U_A the circuit in Fig. 3.6, setting $U = \mathrm{e}^{i \arccos \hat{A}}$,

we see

$$\begin{aligned}
&(\mathrm{Z}\otimes\mathbb{1}_s)U_A \\
&\quad= (\mathrm{HX}\otimes\mathbb{1}_s)\left(|0\rangle\langle 0|\otimes U + |1\rangle\langle 1|\otimes U^\dagger\right)(\mathrm{H}\otimes\mathbb{1}_s) \\
&\quad= (\mathrm{H}\otimes\mathbb{1}_s)\left(|1\rangle\langle 1|\otimes U + |0\rangle\langle 0|\otimes U^\dagger\right)(X\mathrm{H}\otimes\mathbb{1}_s) \\
&\quad= U_A^\dagger(Z\otimes\mathbb{1}_s),
\end{aligned} \tag{3.22}$$

where we have used the relations $\mathrm{ZH} = \mathrm{HX}$, $\mathrm{X}|0\rangle\langle 0| = |1\rangle\langle 1|\mathrm{X}$ and $\mathrm{X}|1\rangle\langle 1| = |0\rangle\langle 0|\mathrm{X}$, which satisfies the relation in Eq. (2.46b).

The BE $U_{\hat{H}}$ for the full Hamiltonian is constructed using LCU to combine local BEs $U_{\hat{H}_\ell}$, which, for BE methods used in this work, share a common S operator. By using a specific qubit organization of the individual $U_{\hat{H}_\ell}$'s, these properties imply that $W_{\hat{H}}$ for the full Hamiltonian can be constructed from $U_{\hat{H}}$ using the same common operator S, namely a single Z-gate. In particular, each $U_{\hat{H}_\ell}$ shares the same ancillary qubit register, denoted by a; if each $U_{\hat{H}_j}$ uses $n_{\mathrm{anc}}^{(j)}$ ancillary qubits, then the register a contains $n_a = \max_j(n_{\mathrm{anc}}^{(j)})$ qubits. Additionally, we make the choice that, if $n_{\mathrm{anc}}^{(j)} < n_a$ for some j, the ancillary qubits used to construct $U_{\hat{H}_\ell}$ are $a_0, a_1, \ldots, a_{n_{\mathrm{anc}}^{(j)}-1}$. Under these assumptions, it is shown in App. B.1 that $W_{\hat{H}}$ can be constructed by choosing S to be a single Z, specifically $S = \mathbb{1}_{a_\Lambda}\otimes(\mathrm{Z}\otimes\mathbb{1})_a\otimes\mathbb{1}_s$.

We conclude by stressing that this simple construction of $W_{\hat{H}}$ is applicable to situations where one uses different methods to prepare different local BEs $U_{\hat{H}_\ell}$. This implies that one is free to choose the BE method that results in the lowest gate cost for each individual term $U_{\hat{H}_\ell}$.

3.4 Numerical demonstration

In Sec. 3.3 we described three main methods to block encode each local Hamiltonian $\hat{H}_\ell$, and provided the dependence of the required resources on the number of qubits $\mathsf{n_q}$ and degree of the polynomial of f_k. We did mention at the beginning of that section that some methods come with large pre-factors, which are not captured by the asymptotics provided. In this section we provide a more detailed numerical analysis of the resources required, including all required prefactors. Such an analysis requires making a detailed choice of the Hamiltonian, and we apply our techniques to the example of a scalar field theory Hamiltonian. Next, we compare the explicit gate cost of preparing BEs of local terms in the Hamiltonian using the methods described in Sec. 3.3. We then use the BE with lowest resource requirements and obtain explicit gate counts for using GQSP to time evolve a both a single and two-site system, and compare the cost to time evolution using 2^{nd} and 4^{th} order product formulas.

We consider a hypercubic lattice Λ of spatial dimension d and lattice spacing a. The Hamiltonians we consider can be written (using the notation from Sec. 3.3.1) as follows:

$$\hat{H} = \hat{H}_\varphi + \hat{H}_\pi\,, \tag{3.23a}$$

$$\hat{H}_{\varphi} = a^{\mathsf{d}} \left[\sum_{i=1}^{\mathsf{N}_\Lambda} V\left(\hat{\varphi}_i\right) + \sum_{\langle ij \rangle} \frac{(\hat{\varphi}_j - \hat{\varphi}_i)^2}{2a^2} \right], \tag{3.23b}$$

$$\hat{H}_{\pi} = a^{\mathsf{d}} \sum_{i=1}^{\mathsf{N}_\Lambda} \frac{1}{2} \hat{\pi}_i^2 = a^{\mathsf{d}} \sum_{i=1}^{\mathsf{N}_\Lambda} \frac{1}{2} \widehat{\mathrm{FT}}_i^{\dagger} \big(\hat{\pi}_i^{(D)}\big)^2 \widehat{\mathrm{FT}}_i \,, \tag{3.23c}$$

where $V\left(\hat{\varphi}_i\right)$ denotes the potential function. In this work, we consider the following potentials

$$V_1\left(\hat{\varphi}_i\right) = \frac{m^2}{2} \hat{\varphi}_i^2 + \frac{\lambda}{4!} \hat{\varphi}_i^4 \,, \tag{3.24}$$

$$V_2\left(\hat{\varphi}_i\right) = g \cos(\hat{\varphi}_i) \,, \tag{3.25}$$

where $V_1(\hat{\varphi}_i)$ is the standard φ^4 scalar field theory potential, and the periodic potential $V_2(\hat{\varphi}_i)$ is relevant for dual basis formulations of compact U(1) gauge theories [90, 91], and a mixed-basis formulation of SU(2) gauge theory [93]. Setting the lattice spacing $a = 1$, the functions $f_k(\hat{\xi}_k)$ for this Hamiltonian are

$$f_0^{(1)}(\hat{\varphi}_i) = V_1(\hat{\varphi}_i) = \frac{m^2}{2} \hat{\varphi}_i^2 + \frac{\lambda}{4!} \hat{\varphi}_i^4 \,, \tag{3.26a}$$

$$f_0^{(2)}(\hat{\varphi}_i) = V_2(\hat{\varphi}_i) = g \cos(\hat{\varphi}_i) \,, \tag{3.26b}$$

$$f_1(\hat{\pi}_i) = (\hat{\pi}_i^{(D)})^2 \,, \tag{3.26c}$$

$$f_2(\hat{\varphi}_i - \hat{\varphi}_j) = \frac{1}{2} (\hat{\varphi}_j - \hat{\varphi}_i)^2 \,. \tag{3.26d}$$

For all numerical results shown in this section, we set $m = 1, \lambda = 32$ and $g = 1$. To obtain the gate counts we use the basis gate-set containing CNOT, R_x, and R_z gates. There are several options for compiling multicontrolled gates. One choice is to use ancillary qubits as in [151] and then compile the resulting circuit down into CNOT, R_x, and R_z gates using, e.g., QISKIT transpilers. Another choice is to exploit the fact that all multicontrolled gates required for this work are diagonal matrices, which implies they can be compiled down into R_z and CNOT gates using the Walsh function formalism [189]. For the circuits studied in this work, we found that the latter choice resulted in smaller total gate counts, with significant reductions in some cases. All PREP oracles are implemented using QISKIT's exact state preparation algorithm. After compiling multicontrolled gates in this way, the final gate count in terms of CNOT, R_x and R_z gates is obtained using QISKIT transpiler with optimization level 1.

Following the discussion in Sec. 3.3, the scale factor for $f_0^{(1)}(\hat{\varphi}_i)$, $f_1(\hat{\pi}_i)$, and $f_2(\hat{\varphi}_i - \hat{\varphi}_j)$ is the same for all BE methods considered (see App. B.3 for a derivation of the scale factors when using standard LCU). Furthermore, the associated scale factors are given by their minimum values. As explained in more detail in App. B.3, the same is not true for $f_0^{(2)} = g\cos(\hat{\varphi})$; using standard LCU techniques results in a value of the scale factor $\sim 30\%$ larger than optimal. However, as we will show, LOVE-LCU has exponentially better complexity

in $\mathsf{n_q}$ compared to standard LCU, and this slight difference in scale factor does not change our qualitative findings.

Fig. 3.7 shows the number of rotation gates and ancillary qubits required to construct BEs of the functions in Eq. (3.26). Plots of CNOT gate counts are shown in App. B.2.

Considering first constructing BEs of $f_0^{(1)}(\hat{\varphi}_i)$ and $f_1(\hat{\pi}_i^{(D)})$, we find that, despite the asymptotically inefficient scaling with $\mathsf{n_q}$, using LOVE-LCU requires the fewest rotation gates for $\mathsf{n_q} \lesssim 11$; for $\mathsf{n_q} = 4$ qubits, LOVE-LCU requires only 13 rotation gates. For $\mathsf{n_q} \gtrsim 11$, due to the linear scaling with $\mathsf{n_q}$, using QSVT requires the fewest rotation gates. However, because digitization errors typically decrease exponentially with $\mathsf{n_q}$ [91, 97], realistic quantum simulations will not require so large a value of $\mathsf{n_q}$, implying that in some cases of practical interest LOVE-LCU will require the fewest rotation gates. Furthermore, QSVT requires $\mathcal{O}(\log \mathsf{n_q})$ ancillary qubits, while LOVE-LCU requires only a single ancillary qubit.

We now turn to preparing a BE of $f_2(\hat{\varphi}_i - \hat{\varphi}_j)$. Because this operator acts on two lattice sites and is of dimension $2^{2\mathsf{n_q}}$, the relative cost of the different methods are fundamentally different than the single qubit operators. In particular, the asymptotic gate cost of LOVE-LCU is now $\mathcal{O}(2^{2\mathsf{n_q}})$, which is the worst of all methods. Despite this fact, LOVE-LCU requires the fewest rotation gates for $\mathsf{n_q} \leq 5$. For $\mathsf{n_q} > 5$, using either QSVT or QETU with $\mathrm{e}^{-i\tau\hat{\varphi}}$ as a building block requires fewer gates than LOVE-LCU, with QSVT being the cheapest due to the $\mathcal{O}(\mathsf{n_q} \log \mathsf{n_q})$ asymptotic gate count scaling.

Because using $\mathsf{n_q} \sim 5$ is a much more realistic value for quantum simulations, it appears at first glance that using QSVT is the best method to prepare a BE of $f_2(\hat{\varphi}_i - \hat{\varphi}_j)$. However, LOVE-LCU has the advantage that it can prepare BEs of arbitrary $2\mathsf{n_q}$-qubit Hermitian operators for the same computational cost. This implies that one could use LOVE-LCU to prepare a BE of, e.g., $f_0^{(0)}(\hat{\varphi}_1) + f_0^{(0)}(\hat{\varphi}_2) + f_2(\hat{\varphi}_1 - \hat{\varphi}_2)$, for the same cost as just $f_2(\hat{\varphi}_1 - \hat{\varphi}_2)$. The same is not true for QSVT or the QETU based methods, which would require one to first prepare a BE of $f_0^{(0)}(\hat{\varphi}_1)$, $f_0^{(0)}(\hat{\varphi}_2)$, and $f_2(\hat{\varphi}_1 - \hat{\varphi}_2)$, and then add these operators using LCU. A comparison using different methods to prepare BE of the two-site Hamiltonian

$$\begin{aligned} \hat{H}_\varphi^{(2)} &= f_0^{(0)}(\hat{\varphi}_1) + f_0^{(0)}(\hat{\varphi}_2) + f_2(\hat{\varphi}_1 - \hat{\varphi}_2) \\ &= \sum_{j=1}^{2} \left(\frac{1}{2} m^2 \hat{\varphi}_j^2 + \frac{\lambda}{4!} \hat{\varphi}_j^2 \right) + \frac{1}{2} \left(\hat{\varphi}_1 - \hat{\varphi}_2 \right)^2 \end{aligned} \tag{3.27}$$

is shown in Fig. 3.8. Due to the need for an additional layer of LCU, we see that the cost of QSVT and QETU based methods has increased relative to standard LCU and LOVE-LCU methods; LOVE-LCU is now the best method for $\mathsf{n_q} \lesssim 6$, with QSVT requiring almost identical rotation gate counts for $\mathsf{n_q} = 6$.

The final operator we study is $\cos(\hat{\varphi})$. Using $\hat{A} = \cos(\hat{\varphi})$ is a special case of LOVE-LCU; the SEL oracle reduces to

$$\textsc{sel} = |0\rangle\langle 0| \otimes e^{-i\hat{\varphi}} + |1\rangle\langle 1| \otimes e^{i\hat{\varphi}}, \tag{3.28}$$

which can be implemented using the control-free procedure using $2\mathsf{n_q}$ of CNOT gates and $\mathsf{n_q}$ of R_z gates. While using QETU with $\mathrm{e}^{-i\tau\hat{\varphi}}$ can also achieve an asymptotic gate count of

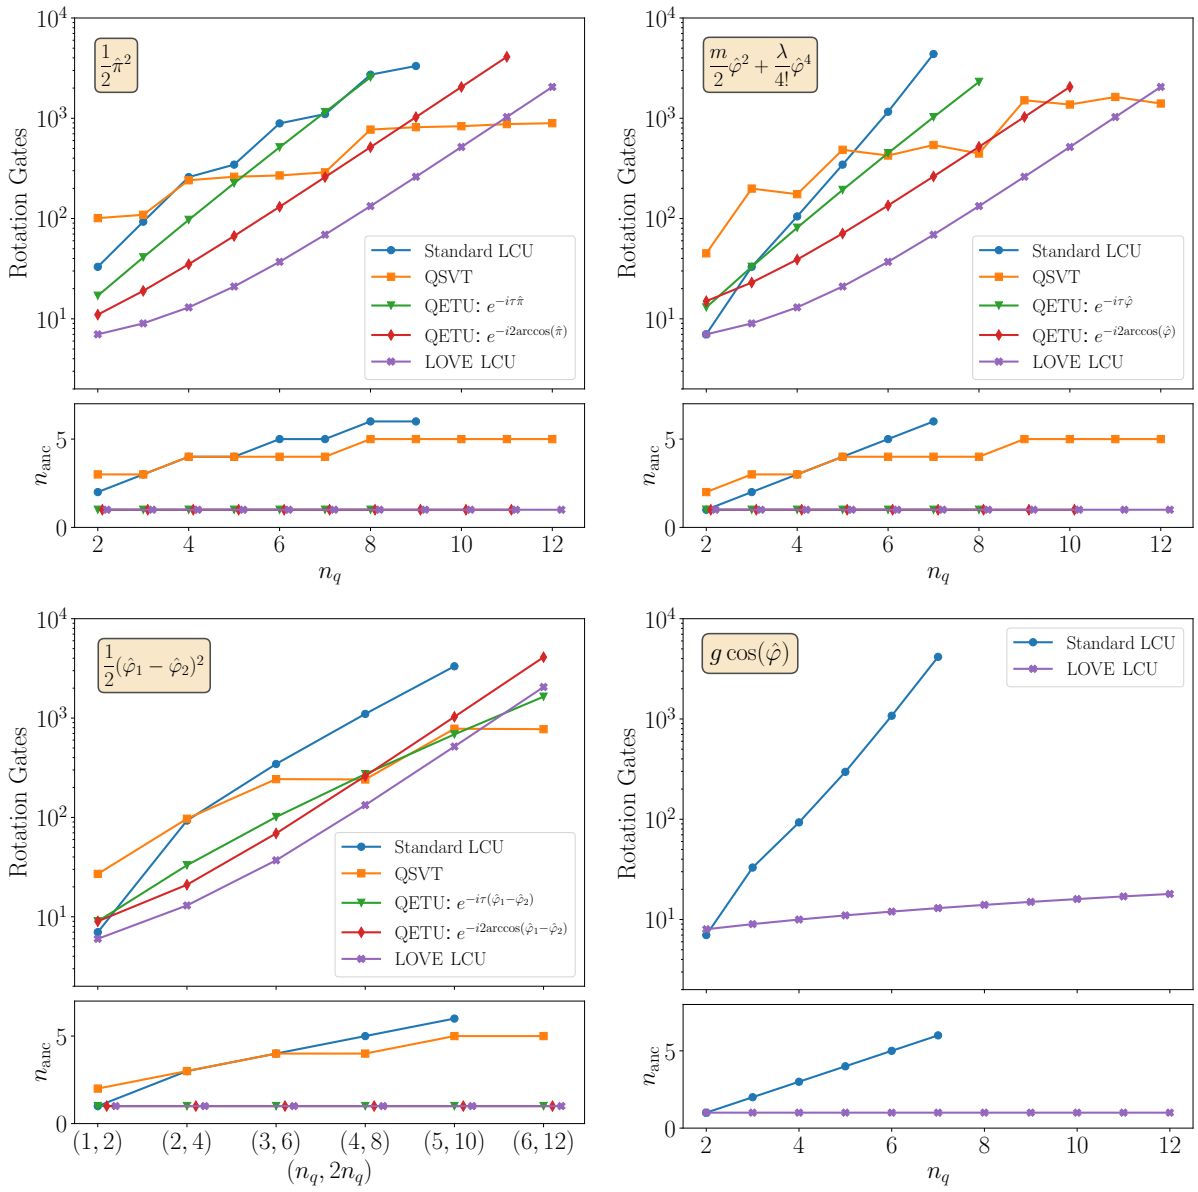

Figure 3.7: Rotation gate count and number of ancillary qubits required to block encode local bosonic operators. The top left, top right, bottom left, and bottom right plots show resource requirements to block encode $\frac{1}{2}\hat{\pi}^2$, $\frac{m}{2}\hat{\varphi}^2 + \frac{\lambda}{4!}\hat{\varphi}^4$, $\frac{1}{2}(\hat{\varphi}_1 - \hat{\varphi}_2)^2, g\cos(\hat{\varphi})$, respectively, for $m = 1, \lambda = 32, g = 1$. Different colored and shaped data points correspond to different methods to prepare the BE. Points with the same number of ancillary qubits have been shifted slightly for clarity. For the single-site operators $\hat{\pi}^2$ and $\frac{1}{m}\hat{\varphi}^2 + \frac{\lambda}{4!}\hat{\varphi}^4$, LOVE-LCU requires the fewest rotation gates for $\mathsf{n_q} \leq 11$. For the two-site operator $\frac{1}{2}(\hat{\varphi}_1 - \hat{\varphi}_2)^2$, LOVE-LCU requires the fewest rotation gates for $\mathsf{n_q} < 6$. LOVE-LCU constructs a BE of $\cos(\hat{\varphi})$ using the fewest gates for all $\mathsf{n_q}$.

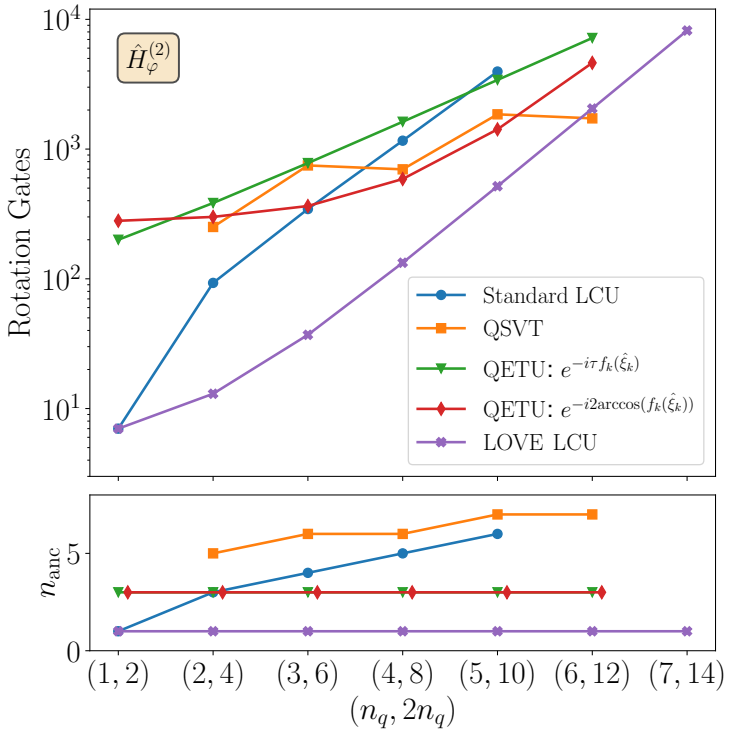

Figure 3.8: Rotation gate count and number of ancillary qubits required to block encode the two site Hamiltonian $\hat{H}_\varphi^{(2)}$ in Eq. (3.27). Different colored and shaped data points correspond to different methods to prepare the BE. Points with the same number of ancillary qubits have been shifted slightly for clarity. While QSVT and QETU based methods require using a final layer of LCU to add the local BEs $f_0^{(0)}(\hat{\varphi}_1), f_0^{(0)}(\hat{\varphi}_2), f_2(\hat{\varphi}_1 - \hat{\varphi}_2)$, LOVE-LCU does not. For this reason, LOVE-LCU requires only a single ancillary qubit. This advantage also leads to LOVE-LCU outperforming all other methods for $\mathsf{n_q} \lesssim 6$; for $\mathsf{n_q} > 6$, QSVT outperforms other methods due to the relative exponentially improved asymptotic scaling.

$\mathcal{O}(\mathsf{n_q})$, it will have a larger overall prefactor in the cost. Because $\cos(x)$ is an infinite degree polynomial, QSVT (or QETU using $\mathrm{e}^{-i2\arccos(\hat{\varphi}/\alpha)}$ as the building block), can only prepare approximate BEs of $\cos(\hat{\varphi})$; while one can use the efficient Jacobi-Anger expansion to do this, one must use a polynomial of degree $d = \mathcal{O}(\|\hat{\varphi}\| + \log(1/\epsilon)) = \mathcal{O}(2^{\mathsf{n_q}/2} + \log(1/\epsilon))$ to approximate to error $\mathcal{O}(\epsilon)$. Since $\cos(\hat{\varphi})$ is a sum of $\mathcal{O}(2^{\mathsf{n_q}})$ Pauli strings, the asymptotic gate cost of using LCU to BE $\cos(\hat{\varphi})$ requires $\mathcal{O}(\mathsf{n_q}2^{\mathsf{n_q}})$ gates. From this discussion, we see that using LOVE-LCU has both the best asymptotic gate scaling as well as the smallest overall prefactor due to the simplicity of the SEL oracle. A comparison of the gate counts as a result of using either LOVE-LCU and standard LCU to prepare a BE of $\cos(\hat{\varphi})$ is shown in the bottom right plot in Fig. 3.7; LOVE-LCU can prepare a BE of $\cos(\hat{\varphi})$ for $\mathsf{n_q} = 12$ using 18 rotation gates and a single ancillary qubit.

We now turn to the problem of simulating time evolution, considering both 2nd and 4th order PFs, and GQSP where the BE was constructed using LOVE-LCU. We study the gate cost to construct the time evolution operator as a function of the evolution time t and error ϵ; the error is defined as $\epsilon = \|U(t) - U_\epsilon(t)\|$, where $U(t)$ is the exact time evolution operator calculated using numerical exponentiation, and $U_\epsilon(t)$ is the approximate time evolution operator calculated using GQSP or PFs. Here we only show the rotation gate counts; CNOT gate counts are given in App. B.2.

The first system we study is the singe-site Hamiltonian

$$\begin{aligned}\hat{H}^{(1)} &= \mathcal{F}^\dagger f_1(\hat{\pi}^{(D)})\mathcal{F} + f_0^{(0)}(\hat{\varphi}) \\ &= \frac{1}{2}\hat{\pi}^2 + \frac{1}{2}m^2\hat{\varphi}^2 + \frac{\lambda}{4!}\hat{\varphi}^4.\end{aligned} \tag{3.29}$$

Fig. 3.9 shows the number of rotation gates as a function of t and ϵ for $\mathsf{n_q} = 3$. We find that GQSP requires less rotation gates compared to using PFs for errors as large as $\epsilon \lesssim 10^{-2}$, which improves upon the work in [114] and Chapter 2 by ~ 5 orders of magnitude; the majority of the savings leading to this significant improvement are a factor of ~ 13 gate reduction in the construction of controlled calls to W_H, and a factor of two gate reduction from using GQSP as opposed to standard QSP. Another notable observation is that, in the region of error ϵ where using GQSP outperforms PFs, the 2^{nd} order PF outperforms the 4^{th} order PF. This is due to the fact that the 4^{th} order PF has not yet reached a small enough value of ϵ for the expected $(1/\epsilon)^{1/4}$ scaling to take effect. While GQSP algorithms are typically thought of as appropriate in the era of full fault-tolerance, our result for the single-site system suggests that for certain systems post-Trotter methods can outperform the PF-based ones sooner than expected.

To better understand how the relative costs change with the number of lattice sites, we

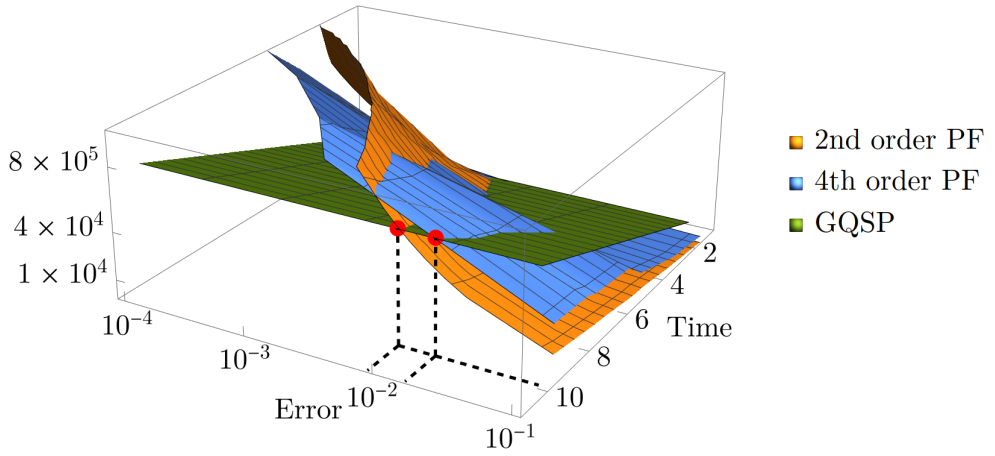

Figure 3.9: Rotation gate count as a function of error ϵ and simulation time t for simulating time-evolution of the single-site Hamiltonian in Eq. (3.29). The blue, orange, and green surfaces show results calculated using a 2^{nd} order PF, a 4^{th} order PF, and GQSP where the BE was constructing using LOVE-LCU. For $t = 10$, GQSP outperforms 2^{nd} and 4^{th} order PFs for $\epsilon \lesssim 10^{-2}$ and $\epsilon \lesssim 2 \times 10^{-2}$, respectively.

now consider the two-site Hamiltonian

$$\begin{aligned} \hat{H}^{(2)} &= \sum_{j=1}^{2} \left[\mathcal{F}_j^\dagger f_1(\hat{\pi}_j^{(D)}) \mathcal{F}_j + f_0^{(0)}(\hat{\varphi}_j) \right] + f_2(\hat{\varphi}_1 - \hat{\varphi}_2) \\ &= \sum_{j=1}^{2} \left(\frac{1}{2}\hat{\pi}_j^2 + \frac{1}{2}m^2\hat{\varphi}_j^2 + \frac{\lambda}{4!}\hat{\varphi}_j^2 \right) + \frac{1}{2}\left(\hat{\varphi}_1 - \hat{\varphi}_2\right)^2 . \end{aligned} \tag{3.30}$$

Fig. 3.10 shows the number of rotation gates as a function of t and ϵ for $\mathsf{n_q} = 3$. We see that GQSP now outperforms 2nd and 4th order PFs for $\epsilon \lesssim 5 \times 10^{-5}$ and $\epsilon \lesssim 10^{-7}$, respectively. While this result suggests one should use PFs for larger lattice sizes, it is likely that significant reductions in the GQSP gate count can be achieved by applying system specific optimizations for compiling the final LCU SEL and PREP oracles. Such gate reductions could make GQSP competitive with PFs for larger errors, and is an interesting direction for future work.

3.5 Discussion and Conclusion

In this work, we presented several new methods for preparing BEs for particular formulations of bosonic lattice field theories. The methods presented in this work relied on two properties of such digitizations, namely the simplicity of the local operators $\hat{\varphi}, \hat{\pi}^{(D)}$, and the

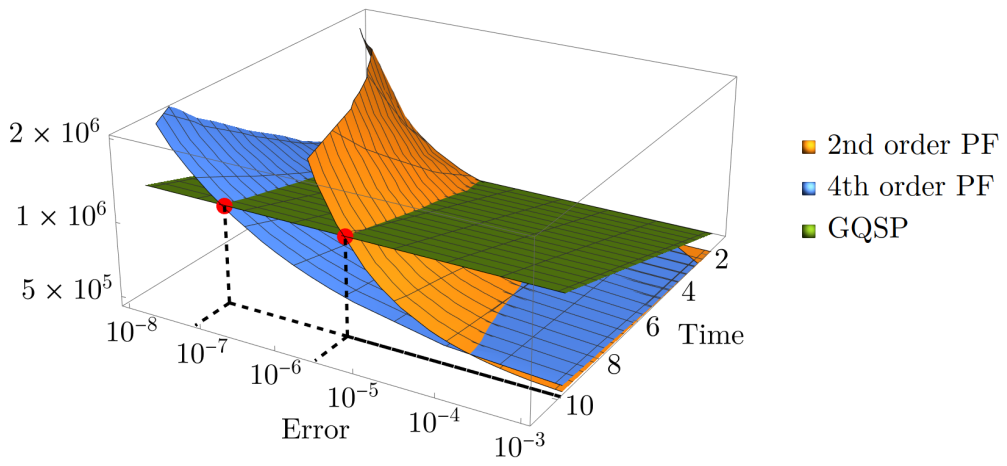

Figure 3.10: Rotation gate count as a function of error ϵ and simulation time t for simulating time-evolution of the two-site in Eq. (3.30). The blue, orange, and green surfaces show results calculated using a 2^{nd} order PF, a 4^{th} order PF, and GQSP where the BE was constructing using LOVE-LCU. For $t = 10$, GQSP outperforms 2^{nd} and 4^{th} order PFs for $\epsilon \lesssim 5 \times 10^{-6}$ and $\epsilon \lesssim 1 \times 10^{-7}$.

fact that only a small number of qubits per site $\mathsf{n_q}$ are required to achieve the digitization errors necessary for realistic simulations.

Exploiting the simplicity of local operators $\hat{\varphi}$ and $\hat{\pi}^{(D)}$, we showed that QSVT can prepare BEs of degree d functions of local terms in the Hamiltonian using $\mathcal{O}(d\mathsf{n_q} \log \mathsf{n_q})$ gates. For polynomial functions with fixed degree d (relevant for simulating $\hat{\varphi}^4$ scalar field theories), this QSVT based method scales asymptotically as $\mathcal{O}(\mathsf{n_q} \log \mathsf{n_q})$ and improves upon the methods in [191] (see Table 3.2 for a comparison), and is identical to the scaling of preparing BEs of operators of the form $\sim (\hat{\varphi}_i - \hat{\varphi}_j)^2$ in [192].

Next, we considered construction of BEs with the aid of the QETU algorithm. While the simple operator $\mathrm{e}^{-i\tau\hat{\varphi}}$ is a natural building block for QETU circuits, because QETU generally requires approximating functions with discontinuous derivatives, we found that high degree polynomials are required to prepare BEs to a high precision. This poor convergence was overcome by exploiting the fact that the spectrum of the local operators are known exactly, and one can use QETU to prepare exact BEs by reproducing the function only at those values, which requires $\mathcal{O}(\mathsf{n_q} 2^{\mathsf{n_q}})$ gates. Despite this technically inefficient scaling with $\mathsf{n_q}$, through explicit circuit constructions, we found that, due to QSVT having a relatively large prefactor in the gate cost, QETU outperforms QSVT for small values of $\mathsf{n_q}$. The comparison of different BEs considered in this work is summarized in Table 3.1.

Motivated by the observation that algorithms with inefficient scaling in $\mathsf{n_q}$ can require the fewest gates, we developed the conceptually simple LOVE-LCU approach. LOVE-LCU can prepare BEs of arbitrary Hermitian operators $\hat{A}$ acting on $\mathsf{n_q}$ qubits using only two

Operator	Complexity from [191]	Complexity from QSVT
$\hat{\varphi}_i\hat{\varphi}_j$	$\mathcal{O}(\mathsf{n_q})$	$\mathcal{O}(\mathsf{n_q}\log\mathsf{n_q})$
$\hat{\varphi}^2$	$\mathcal{O}(\mathsf{n_q}^2)$	$\mathcal{O}\left(\mathsf{n_q}\log\mathsf{n_q}\right)$
$\hat{\pi}^2$	$\mathcal{O}(\mathsf{n_q}^2)$	$\mathcal{O}\left(\mathsf{n_q}\log\mathsf{n_q}\right)$
$\hat{\varphi}^4$	$\mathcal{O}(\mathsf{n_q}^2)$	$\mathcal{O}\left(\mathsf{n_q}\log\mathsf{n_q}\right)$

Table 3.2: Comparison of complexities to block-encode terms for $\mathsf{N}_\Lambda = \mathcal{O}(1)$ sites in scalar field theory Hamiltonian. The second column are complexities based on [191] which uses a "comparator" operator to reduce the complexity of LCU operations; The third column is the relevant QSVT method described in Sec. 3.3 (taking polynomials of degree $d \leq 4$).

controlled calls to $\mathrm{e}^{-i\arccos(\hat{A})}$, requiring a single ancillary qubit and $\mathcal{O}(2^{\mathsf{n_q}})$ ($\mathcal{O}(4^{\mathsf{n_q}})$) gates for diagonal (non-diagonal) operators. Our numerical investigations demonstrated that LOVE-LCU outperforms all other methods for preparing BEs of operators acting on $\lesssim 11$ qubits, at which point QSVT wins due to its relative exponentially improved gate scaling, albeit requiring more ancillary qubits than LOVE-LCU.

Our findings indicate that, while understanding asymptotic gate complexities can serve as a useful guide, explicit numerical investigations are essential to compare different BE techniques for a specific application as their performance depends heavily on the considered Hamiltonian, problem size, and desired precision. We note that one can view the different approaches to BEs considered in this work as belonging to the same family of QSP-based methods but with different building blocks, which thus change the functions of interest and their implementations. We observe that methods utilizing asymptotically inefficient building blocks can potentially outperform the asymptotically optimal ones in practically relevant regimes.

In addition to developing new methods for constructing BEs using these techniques, we also developed methods for efficiently constructing the walk operator $W_{\hat{H}}$, required for using QSP-based methods for Hamiltonian simulation. In particular, when BEs of local terms are constructed using any of the new methods presented in this work, and the full Hamiltonian is then constructed using a final LCU procedure to add these local BEs, the walk operator can be constructed as in Eq. (2.45) where S is a single Z-gate. Using GQSP combined with LOVE-LCU, we found that, for a single site anharmonic oscillator, GQSP outperforms PF methods for errors as small as $\epsilon \sim 10^{-2}$, which is ~ 5 orders of magnitude improved relative to previous work [114]. Note that because PF methods have better asymptotic scaling with the number of lattice sites than GQSP [114], the error threshold at which GQSP outperforms PFs is expected to be lower for larger lattice sizes. To better understand how the relative cost changes with the number of lattice sites, we studied the two-site system and found that GQSP now outperforms PF methods for errors $\epsilon \sim 10^{-7}$. This dramatic change, however, likely stems (at least in part) from using a compiler that is agnostic to the structure of the

system studied, emphasizing the need for system specific circuit optimizations.

While we performed numerical gate count studies for a scalar field theory Hamiltonian, our methods can be directly applied to other bosonic theories formulated in terms of conjugate operators, including compact formulations of U(1) LGTs [91, 188] and the mixed-basis formulation of SU(2) LGT in [93]. Our methods can also be applied to LGTs that are not formulated in terms of conjugate variables, including the standard Kogut-Susskind Hamiltonian and Loop-String-Hadron formulations [76, 77, 78, 79]. Notably, this includes situations in which the Hamiltonian operator is non-local, *i.e.*, contains an exponential number of Pauli operators. For example, the Hamiltonian operator of the gauge-fixed U(1) LGT [90, 91, 135, 136, 188] includes a term of the form $\cos\left(\sum_j^N B_j\right)$, where the magnetic field operator at site j, B_j, is analogous to the field operator $\hat{\varphi}$. While this operator is a sum of a number of Pauli strings exponential in the number of lattice sites, it can be readily encoded via the LOVE-LCU construction in Fig. 3.6 with a cost linear in the number of sites; this exponential reduction in cost allows one to avoid the usage of costly sparse oracle routines. More generally, the methods discussed in this work could be useful for block encoding Hamiltonians involving operator functions of multiple variables. One potential application is the first-quantized chemistry Hamiltonians, which can expressed in terms of position and derivative operators, such as $1/r$ and ∂_x^2, acting on fermionic wavefunctions.

Our work leads naturally to several interesting research directions. Regarding using QETU with $\mathrm{e}^{-i\tau\hat{\varphi}_{\mathrm{sh}}}$ as a building block, it was found in [101] that limiting the spectrum of $\hat{\varphi}_{\mathrm{sh}}$ to be in the range $[\eta, \pi-\eta]$ and varying η allowed one to approximate Gaussian functions $\sim e^{-\hat{\varphi}^2}$ to a precision ϵ using only $\log(\mathsf{n_q}\log(1/\epsilon))$ gates. The same is not true, however, for polynomial operators of the form $\hat{\varphi}^n$ considered in this work; as $\eta \to \pi/2$ the scale factor of the BE goes to zero, leading to a significant increase in the cost of Hamiltonian simulation. It would be interesting to see if the method of varying η can achieve a $\mathcal{O}(\mathsf{n_q}\log(1/\epsilon))$ scaling for other functions that do not diverge for large argument, some relevant physical examples being central potentials of the form $1/r$, or $1/r^6$ and $1/r^{12}$ which appear in the Lennard-Jones potential, see for instance [192].

With an eye towards fault tolerance, it is essential to compile the circuits down to metrics suitable for the error correction protocols (e.g., T-gate counts) to assess the relative cost of the methods in this work. The QETU-based methods and LOVE-LCU inherently require the use of rotation gates, which implies they likely come with a large prefactor in the T-gate count. These methods, however, may serve useful for partial-fault tolerant implementations where only Clifford gates are implemented in a fault-tolerant manner, and rotation gates are implemented in a non-fault-tolerant way [193]. The QSVT-based methods are generally expected to require less T-gates, as one can use unary-LCU methods to reduce the cost of the SEL oracle, and the methods in [191] to construct a BE of $\hat{\varphi}$ avoiding the need for T-gates in the PREP oracle. It would be interesting to do a dedicated T-gate study using a combination the QSVT method in this work and the efficient T-gate constructions in Refs. [130, 191].

While the total T-gate count is usually considered for the total cost of a fault-tolerant simulation (see for instance [64, 133]), there have been suggestions that another possible metric is instead T-gate *depth* [194]. This is similar in spirit to Amdahl's Law in classical

computing, which states that the maximum possible speedup from parallelizing one's code is limited by components of the code that must be executed sequentially. In other words, Amdahl's Law describes the time a code takes to run in the limit of access to infinitely many classical computing nodes (barring communication time). By asking the same question in the context of using several fault-tolerant quantum computing nodes, T-gate depth is the metric for how parallelizable a circuit is. In scenarios where T-gate depth is the relevant cost, reductions in the depth of preparing BEs can be achieved using the modified LCU procedure in [195]. Using this or similar procedures, it is possible to imagine that the depth of constructing a BE of geometrically local Hamiltonians can be reduced to constant in the number of lattice sites. Such a method would result in the depth of QSP-based time evolution methods scaling linearly in the volume, which outperforms the volume scaling of the depth of PF based methods for geometrically local systems [106].

We conclude by pointing out that, while further improvements can likely be made within the paradigm of preparing BEs via QSP-based methods, dramatic reductions in the gate count will likely require considering new paradigms. Although our focus has been on constructing BEs using LCU for QSP-based time evolution, which typically requires long coherent circuits and full fault tolerance, our methods can be adapted for simulations relying solely on LCU with approaches like the Truncated Taylor Series [165]. To make these simulations more suitable for noisy intermediate-scale quantum (NISQ) devices, one promising direction is to explore recently developed randomized near-term implementations of LCU [196, 197], which could be combined with our constructions to enable more practical quantum simulations in the near term. Another promising avenue is to combine Trotter-based methods with near-optimal time evolution techniques [196, 198, 199], potentially leveraging the strengths of both approaches for improved efficiency. Additionally, the developing field of multi-variable QSP methods [184, 200, 201, 202, 203, 204] (this general class of methods includes multi-variable QETU), which leverages the structure of lattice field theories written as sums of many commuting variables, could result in significant cost reductions.

Chapter 4

Truncation Uncertainties for Accurate Quantum Simulations of Lattice Gauge Theories

4.1 Introduction

The large potential for quantum simulation and the recent development of noisy quantum computers has led to a large amount of work exploring how to use these devices to perform simulations. Lattice gauge theories with a number of different gauge groups in $1+1D$ have been simulated on quantum computers [46, 65, 66, 67, 70, 71, 180, 205, 206, 207, 208, 209, 210, 211, 212, 213, 214, 215, 216, 217, 218, 219, 220, 221, 222, 223, 224, 225, 226, 227, 228, 229, 230]. Limited simulations have been performed in higher spatial dimensions [45, 61, 62, 68, 231, 232, 233, 234, 235, 236, 237, 238, 239]. This is due to gauge fields having continuous degrees of freedom and quantum computers having finite, discrete degrees of freedom. This challenge is not present in $1+1D$ since Gauss's law can be used to integrate out gauge fields, leaving only fermions, which have a finite Hilbert space. This leads to long-range interactions; however, the range of this interaction can be truncated due to the exponential decay of correlations in low-energy states [71]. Mapping the gauge fields in higher dimensions onto quantum computers requires a truncation of the gauge fields. For the standard Kogut–Susskind Hamiltonian, a number of approaches have been developed, including electric basis truncations [32, 49, 60, 61, 62, 69, 76, 77, 78, 79, 80, 81, 166, 239, 240, 241, 242, 243, 244, 245, 246, 247, 248, 249], discrete subgroups [82, 83, 84, 85, 86, 87, 250, 251, 252, 253], and hybrid electric/magnetic bases [91, 92, 93, 94, 254, 255, 256]. Alternatives to the usual Kogut–Susskind Hamiltonian have also been developed, such as quantum link models [88, 89, 257, 258, 259, 260] and orbifold lattice gauge theory [17, 95, 96], which are conjectured to have QCD as their continuum limit.

Ultimately, simulations capable of making predictions with precision will need to have controllable uncertainties. This will have contributions from both the hardware and the

theoretical formulation of the simulation. The effects of a finite lattice spacing can be estimated using the same Symanzik action approach used in traditional lattice QCD [261]. This formalism can also be used to understand the error coming from discrete (possibly Trotterized) time steps [113, 123]. Recent work has developed the formalism necessary for estimating finite volume errors in a generic correlation function computed on a quantum computer [74]. The truncation of the gauge fields will also contribute to the theoretical uncertainty. Several works have shown a fast convergence of energy eigenstates with field truncation for a number of different theories [49, 51, 62, 80, 97, 145, 262, 263, 264, 265, 266]. For scalar field theories, the exponential convergence is a reflection of the Nyquist-Shannon sampling theorem [97], however a similar mechanism for non-Abelian gauge theories has not yet been identified. Previous work has shown that truncation errors in dynamics converge exponentially fast, provided one uses a truncation that grows linearly with time [75]. However, it is essential to have tight estimates of the truncation error, as overestimating errors will unnecessarily delay quantum simulations of scientific importance. In this work, the presence of Hilbert space fragmentation in the Kogut–Susskind Hamiltonian is used to obtain estimates of the truncation error for truncations in the electric basis. These error estimates do not require the truncation to grow with time to maintain accuracy and go to zero as a factorial in the size of the truncated link Hilbert space. Numerical simulations are performed for the Schwinger model and a pure U(1) lattice gauge theory on a plaquette ladder. These simulations show that the error estimates in this work correctly capture the dynamics of lattice gauge theories.

The work presented in this chapter is based on [119]

4.2 Truncation Errors in Pure Gauge Theories

We begin this section by introducing the $U(1)$ lattice gauge theory on a single plaquette which will serve as a motivating example for the subsequent sections. The Hamiltonian is given by

$$\hat{H} = 2g^2\hat{E}^2 + \frac{1}{2g^2}\left(2 - \hat{\square} - \hat{\square}^\dagger\right) \tag{4.1}$$

$$\hat{E} = \sum_{n=-\infty}^{\infty} n\,|n\rangle\langle n| \tag{4.2}$$

$$\hat{\square} = \sum_{n=-\infty}^{\infty} |n\rangle\langle n+1|\ , \tag{4.3}$$

where $\hat{E}$ is the electric operator and the states $|n\rangle$ are the eigenstates of $\hat{E}$ that form the standard electric basis of the Kogut-Susskind Hamiltonian. $\hat{\square} + \hat{\square}^\dagger$ is the plaquette term which for a single plaquette connects electric basis states that differ by one unit of electric flux. We note that the methods in the following sections apply more generally beyond just

the $U(1)$ theory as we will discuss. We will use $U(1)$ as a concrete example to demonstrate our methods.

4.2.1 Eigenstate Truncation Errors

To motivate why improvements in truncation error are expected to be possible, traditional perturbation theory will be applied to eigenstates of truncated Hamiltonians. We will consider a Hamiltonian, $\hat{H}$, defined on a single bosonic mode with basis states given by $|0\rangle, |1\rangle, |2\rangle, \ldots$ that are eigenstates of a number operator (similar to the electric basis states for the single plaquette $U(1)$ theory) and assume that the terms in the Hamiltonian do not change the boson number by more than 1, i.e. $\langle n+k|\hat{H}|n\rangle = 0$ for $k > 1$. For a lattice gauge theory, this would be analogous to a theory with a single plaquette and no matter. A truncated Hamiltonian, $\hat{H}_\Lambda$, can be defined by removing the off-diagonal elements of $\hat{H}$ above some truncation Λ. Explicitly, the relation between the truncated and untruncated Hamiltonian is given by

$$\hat{H} = \hat{H}_\Lambda + \hat{V}_\Lambda \,, \tag{4.4}$$

where $\langle n|\hat{H}|n\rangle = \langle n|\hat{H}_\Lambda|n\rangle \ \forall n$, $\langle n|\hat{H}|k\rangle = \langle n|\hat{H}_\Lambda|k\rangle$ if both $n, k \leq \Lambda$, and $\langle n|\hat{H}_\Lambda|k\rangle = 0$ otherwise. We note that by this construction, time evolution under $\hat{H}_\Lambda$ preserves the truncated Hilbert space. We denote the eigenstates and energies of the full Hamiltonian by $|\psi_n\rangle$ and E_n, and the eigenstates and energies of the truncated Hamiltonian by $\left|\psi_n^\Lambda\right\rangle$ and E_n^Λ. Thus

$$\begin{aligned} |\psi_n\rangle &= \left|\psi_n^\Lambda\right\rangle + \left|\delta\psi_n^\Lambda\right\rangle \\ \left|\delta\psi_n^\Lambda\right\rangle &= \left(1 - \left|\psi_n^\Lambda\right\rangle\left\langle\psi_n^\Lambda\right|\right)\frac{1}{E_n - \hat{H}_\Lambda}\hat{V}_\Lambda|\psi_n\rangle \\ E_n - E_n^\Lambda &= \left\langle\psi_n^\Lambda\right|\hat{V}_\Lambda\left|\delta\psi_n^\Lambda\right\rangle \end{aligned} \tag{4.5}$$

where $\left\langle\psi_n^\Lambda\right|\left|\delta\psi_n^\Lambda\right\rangle = 0$. Using standard perturbation theory, the leading order correction to $|\psi_n\rangle$ is given by[1]

$$\left|\delta\psi_n^\Lambda\right\rangle = \sum_{m>\Lambda}|m\rangle\frac{\langle m|\hat{V}_\Lambda\left|\psi_n^\Lambda\right\rangle}{E_n^\Lambda - \langle m|\hat{H}_\Lambda|m\rangle} \,. \tag{4.6}$$

Since $\hat{V}_\Lambda$ can only change the boson number by at most 1 and the state $\left|\psi_n^\Lambda\right\rangle$ only has support within the truncated Hilbert space, only $m = \Lambda + 1$ survives in the sum. Furthermore, since $|\Lambda\rangle$ is the only state within the truncated Hilbert space that is connected to $|\Lambda+1\rangle$ by $\hat{V}_\Lambda$, we can insert $|\Lambda\rangle\langle\Lambda|$ to write

$$\left|\delta\psi_n^\Lambda\right\rangle = |\Lambda+1\rangle\frac{\langle\Lambda+1|\hat{V}_\Lambda|\Lambda\rangle\langle\Lambda|\left|\psi_n^\Lambda\right\rangle}{E_n^\Lambda - \langle\Lambda+1|\hat{H}_\Lambda|\Lambda+1\rangle} \,. \tag{4.7}$$

[1]Note that $|\psi_n\rangle$ is not normalized to 1, but when expanding perturbatively, correcting this will only contribute at higher orders in perturbation theory.

We note that the magnitude of this correction term is controlled by the quantity ϵ_Λ defined as follows:

$$\epsilon_\Lambda := \frac{\left|\langle \Lambda+1| \hat{V}_\Lambda |\Lambda\rangle\right|}{\left|E_n^\Lambda - \langle \Lambda+1| \hat{H}_\Lambda |\Lambda+1\rangle\right|}. \tag{4.8}$$

The RHS of Eq. (4.7) depends on the overlap of the truncated eigenstate with the bosonic mode corresponding to the largest bosonic number kept in the truncated Hilbert space. To determine the scaling as the truncation is raised, we can determine this overlap by expanding perturbatively around a lower truncation Λ_0. First, we note that

$$\langle \Lambda_0+1| |\psi_n\rangle = \langle \Lambda_0+1| \left|\psi_n^{\Lambda_0}\right\rangle + \langle \Lambda_0+1| \left|\delta\psi_n^{\Lambda_0}\right\rangle \tag{4.9}$$

$$= \langle \Lambda_0+1| \left|\delta\psi_n^{\Lambda_0}\right\rangle \tag{4.10}$$

$$= \frac{\langle \Lambda_0+1| \hat{V}_{\Lambda_0} |\Lambda_0\rangle \langle \Lambda_0| \left|\psi_n^{\Lambda_0}\right\rangle}{E_n^{\Lambda_0} - \langle \Lambda_0+1| \hat{H}_{\Lambda_0} |\Lambda_0+1\rangle}, \tag{4.11}$$

where the second equality follows from the fact that $\psi_n^{\Lambda_0}$ only has support on $\{|0\rangle, |1\rangle, \ldots, |\Lambda_0\rangle\}$ and the third equality follows from Eq. (4.7). We also have that that

$$\langle \Lambda_0+1| |\psi_n\rangle = \langle \Lambda_0+1| \left|\psi_n^{\Lambda_0+1}\right\rangle + \langle \Lambda_0+1| \left|\delta\psi_n^{\Lambda_0+1}\right\rangle = \langle \Lambda_0+1| \left|\psi_n^{\Lambda_0+1}\right\rangle \tag{4.12}$$

since Eq. (4.7) implies that $\left|\delta\psi_n^{\Lambda_0+1}\right\rangle \propto |\Lambda_0+2\rangle$. Putting this together, we thus have that

$$\langle \Lambda_0+1| |\psi_n\rangle = \langle \Lambda_0+1| \left|\psi_n^{\Lambda_0+1}\right\rangle = \frac{\langle \Lambda_0+1| \hat{V}_{\Lambda_0} |\Lambda_0\rangle \langle \Lambda_0| \left|\psi_n^{\Lambda_0}\right\rangle}{E_n^{\Lambda_0} - \langle \Lambda_0+1| \hat{H}_{\Lambda_0} |\Lambda_0+1\rangle}. \tag{4.13}$$

Similarly, it is true that

$$\langle \Lambda_0+2| \left|\psi_n^{\Lambda_0+2}\right\rangle \tag{4.14}$$

$$= \frac{\langle \Lambda_0+2| \hat{V}_{\Lambda_0+1} |\Lambda_0+1\rangle \langle \Lambda_0+1| \left|\psi_n^{\Lambda_0+1}\right\rangle}{E_n^{\Lambda_0+1} - \langle \Lambda_0+2| \hat{H}_{\Lambda_0+1} |\Lambda_0+2\rangle} \tag{4.15}$$

$$= \langle \Lambda_0| \left|\psi_n^{\Lambda_0}\right\rangle \cdot \frac{\langle \Lambda_0+2| \hat{V}_{\Lambda_0+1} |\Lambda_0+1\rangle}{E_n^{\Lambda_0+1} - \langle \Lambda_0+2| \hat{H}_{\Lambda_0+1} |\Lambda_0+2\rangle} \cdot \frac{\langle \Lambda_0+1| \hat{V}_{\Lambda_0} |\Lambda_0\rangle}{E_n^{\Lambda_0} - \langle \Lambda_0+1| \hat{H}_{\Lambda_0} |\Lambda_0+1\rangle} \tag{4.16}$$

$$= \langle \Lambda_0| \left|\psi_n^{\Lambda_0}\right\rangle \cdot \frac{\langle \Lambda_0+2| \hat{V}_{\Lambda_0} |\Lambda_0+1\rangle}{E_n^{\Lambda_0+1} - \langle \Lambda_0+2| \hat{H}_{\Lambda_0} |\Lambda_0+2\rangle} \cdot \frac{\langle \Lambda_0+1| \hat{V}_{\Lambda_0} |\Lambda_0\rangle}{E_n^{\Lambda_0} - \langle \Lambda_0+1| \hat{H}_{\Lambda_0} |\Lambda_0+1\rangle} \tag{4.17}$$

$$\approx \langle \Lambda_0| \left|\psi_n^{\Lambda_0}\right\rangle \cdot \frac{\langle \Lambda_0+2| \hat{V}_{\Lambda_0} |\Lambda_0+1\rangle}{E_n^{\Lambda_0} - \langle \Lambda_0+2| \hat{H}_{\Lambda_0} |\Lambda_0+2\rangle} \cdot \frac{\langle \Lambda_0+1| \hat{V}_{\Lambda_0} |\Lambda_0\rangle}{E_n^{\Lambda_0} - \langle \Lambda_0+1| \hat{H}_{\Lambda_0} |\Lambda_0+1\rangle} \tag{4.18}$$

where we have treated ϵ_{Λ_0} as a small parameter, and the approximate equality above is to leading order in ϵ_{Λ_0} (since Eqs. (4.5) and (4.7) taken together imply that $E_n^{\Lambda_0+1} - E_n^{\Lambda_0} = \mathcal{O}(\epsilon_{\Lambda_0}^2)$). Continuing this process, we can compute $\langle \Lambda_0 + 3| \left|\psi_n^{\Lambda_0+3}\right\rangle$, $\langle \Lambda_0 + 4| \left|\psi_n^{\Lambda_0+4}\right\rangle$, and so on to leading order until we reach the target truncation of Λ. We thus arrive at the following result:

$$\langle \Lambda| \left|\psi_n^{\Lambda}\right\rangle = \langle \Lambda_0| \left|\psi_n^{\Lambda_0}\right\rangle \prod_{k=\Lambda_0}^{\Lambda-1} \frac{\langle k+1| \hat{V}_{\Lambda_0} |k\rangle}{E_n^{\Lambda_0} - \langle k+1| \hat{H}_{\Lambda_0} |k+1\rangle} \tag{4.19}$$

to leading order. Therefore, the leading corrections to the states and energies at a truncation of Λ are

$$\left|\delta\psi_n^{\Lambda}\right\rangle = |\Lambda+1\rangle \langle \Lambda_0| \left|\psi_n^{\Lambda_0}\right\rangle \frac{\langle \Lambda+1| \hat{V}_{\Lambda} |\Lambda\rangle}{E_n^{\Lambda} - \langle \Lambda+1| \hat{H}_{\Lambda} |\Lambda+1\rangle} \prod_{k=\Lambda_0}^{\Lambda-1} \frac{\langle k+1| \hat{V}_{\Lambda_0} |k\rangle}{E_n^{\Lambda_0} - \langle k+1| \hat{H}_{\Lambda_0} |k+1\rangle} \tag{4.20}$$

$$\begin{aligned} E_n - E_n^{\Lambda} = & \frac{1}{E_n^{\Lambda} - \langle \Lambda+1| \hat{H}_{\Lambda_0} |\Lambda+1\rangle} \\ & \times \left| \langle \Lambda_0| \left|\psi_n^{\Lambda_0}\right\rangle \langle \Lambda| \hat{V}_{\Lambda_0} |\Lambda+1\rangle \prod_{k=\Lambda_0}^{\Lambda-1} \frac{\langle k+1| \hat{V}_{\Lambda_0} |k\rangle}{E_n^{\Lambda_0} - \langle k+1| \hat{H}_{\Lambda_0} |k+1\rangle} \right|^2 \end{aligned} \tag{4.21}$$

Treating ϵ_{Λ_0} as a small parameter is valid when $\left|\langle \Lambda_0+1| \hat{V}_{\Lambda_0} |\Lambda_0\rangle\right|$ is small compared to $\left|E_n^{\Lambda_0} - \langle \Lambda_0+1| \hat{H}_{\Lambda_0} |\Lambda_0+1\rangle\right|$. For Kogut-Susskind formulations of lattice gauge theory Hamiltonians we note that the electric energy gap grows quadratically with the size of the representation whereas the plaquette matrix elements are bounded, which makes treating ϵ_{Λ_0} as a small parameter a better and better approximation as we increase the target truncation past Λ_0 assuming Λ_0 is chosen to be large enough. For the corrections Eq. (4.20) to systematically decrease as the truncation is raised, it is necessary for the energy of the eigenstate to be below $\langle \Lambda_0| \hat{H}_{\Lambda_0} |\Lambda_0\rangle$. For such low energy states, where $E_n^{\Lambda_0} \ll \langle k+1| \hat{H}_{\Lambda_0} |k+1\rangle$, the size of the corrections will scale as $\mathcal{O}\Big(\prod_{n<\Lambda} \frac{\langle n+1|\hat{V}_{\Lambda_0}|n\rangle}{\langle n+1|\hat{H}_{\Lambda_0}|n+1\rangle}\Big)$.

We can make this argument more concrete for the case of $U(1)$ on a single plaquette. We split up the Hamiltonian shown in Eq. (4.2) as follows:

$$\hat{H} = \hat{H}_{\Lambda} + \hat{V}_{\Lambda} \quad , \tag{4.22}$$

with

$$\hat{H}_{\Lambda} = 2g^2 \hat{E}^2 + \frac{1}{2g^2} \left(2 - \hat{\square}_{\Lambda} - \hat{\square}_{\Lambda}^{\dagger}\right)$$

$$\hat{\square}_{\Lambda} = \sum_{n=-\Lambda+1}^{\Lambda-1} |n\rangle \langle n+1|$$

$$\hat{V}_\Lambda = -\frac{1}{2g^2}\sum_{|n|\geq\Lambda} |n\rangle\langle n+1| + \text{h.c.}\,. \tag{4.23}$$

This implies that for $k \geq \Lambda_0$ we have

$$\begin{aligned}\langle k+1|\hat{V}_{\Lambda_0}|k\rangle &= -\frac{1}{2g^2}\\ \langle k+1|\hat{H}_{\Lambda_0}|k+1\rangle &= 2g^2(k+1)^2\,.\end{aligned} \tag{4.24}$$

Thus, the product shown in Eq. (4.20) becomes

$$\prod_{k=\Lambda_0}^{\Lambda-1} \frac{1}{4g^4(k+1)^2} = \mathcal{O}\left(\frac{1}{(\Lambda!)^2}\right). \tag{4.25}$$

This expansion is valid for eigenstates where

$$\epsilon_{\Lambda_0} = \frac{1}{2g^2 E_n^{\Lambda_0} - 4g^4(\Lambda_0+1)^2} \quad , \tag{4.26}$$

is a small parameter. For a theory with a boson creation operator and number operator in the Hamiltonian, we have that

$$\begin{aligned}\langle k+1|\hat{V}_{\Lambda_0}|k\rangle &\propto \sqrt{k+1}\\ \langle k+1|\hat{H}_{\Lambda_0}|k+1\rangle &\propto k+1\,,\end{aligned} \tag{4.27}$$

and hence the the product shown in Eq. (4.20) becomes

$$\prod_{k=\Lambda_0}^{\Lambda-1} \frac{1}{\sqrt{k+1}} = \mathcal{O}\left(\frac{1}{\sqrt{\Lambda!}}\right). \tag{4.28}$$

Note that these expressions were derived for a theory with a single bosonic mode, but it is expected that the general scaling of the truncation corrections should hold for larger systems since the scalings are determined by local properties of the Hamiltonian.

4.2.2 Time Dependent Truncation Errors with Fragmented States

4.2.2.1 Single Plaquettes

The results of the previous section show that truncation errors in eigenstates go to zero as a factorial of the field truncation, but do not provide quantitative estimates of the truncation error. In this section, it will be shown how time-dependent perturbation theory can be used to estimate truncation errors quantitatively. In the authors' previous work, it was demonstrated

that the Kogut–Susskind Hamiltonian generically demonstrates Hilbert space fragmentation (HSF) [122]. This is due to the quadratic nature of the electric energy terms, leading to large gaps in the spectrum for states with large electric fields. This will allow us to apply a strong coupling expansion to the dynamics of states with these large electric fields. In the following discussion, a single plaquette will be considered, allowing the entire state of the system to be specified by the electric field on a single link. Note that while errors in the time evolution will be estimated here, the results in this section can be used to estimate contributions to eigenstates by using an adiabatic switching procedure, such as the one used in the proof of the Gellmann-Low theorem [267]. The results of estimating truncation errors in such a manner are consistent with the previous section. Explicitly, to study the evolution of a state $\left|\phi(t)\right\rangle$, define the interaction picture state $\left|\phi_I(t)\right\rangle$ by

$$\left|\phi(t)\right\rangle = e^{-i\hat{H}_\Lambda t}\left|\phi_I(t)\right\rangle , \tag{4.29}$$

where we assume $\left|\phi(0)\right\rangle$ only has support on basis states with electric field below Λ. The state $\left|\phi_I(t)\right\rangle$ satisfies

$$\left|\phi_I(t)\right\rangle = \left|\phi(0)\right\rangle - i\int_0^t dt_0\ e^{i\hat{H}_\Lambda t_0}\hat{V}_\Lambda e^{-i\hat{H}_\Lambda t_0}\left|\phi_I(t_0)\right\rangle . \tag{4.30}$$

By repeatedly substituting $\left|\phi_I(t)\right\rangle$ into this equation, one can generate a series for $\left|\phi_I(t)\right\rangle$ in powers of $\hat{V}_\Lambda$. In particular, we have that

$$\left|\phi_I(t)\right\rangle = \left|\phi(0)\right\rangle - i\int_0^t dt_0\ e^{i\hat{H}_\Lambda t_0}\hat{V}_\Lambda e^{-i\hat{H}_\Lambda t_0}\left(\left|\phi(0)\right\rangle - i\int_0^{t_0} dt_1\ e^{i\hat{H}_\Lambda t_1}\hat{V}_\Lambda\left|\phi_I(t_1)\right\rangle\right) \tag{4.31}$$

$$= \left|\phi(0)\right\rangle - i\int_0^t dt_0\ e^{i\hat{H}_\Lambda t_0}\hat{V}_\Lambda e^{-i\hat{H}_\Lambda t_0}\left|\phi(0)\right\rangle + \mathcal{O}(\hat{V}_\Lambda^2) \tag{4.32}$$

Thus,

$$(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\left|\phi(0)\right\rangle = e^{-i\hat{H}_\Lambda t}(\left|\phi_I(t)\right\rangle - \left|\phi(0)\right\rangle) \tag{4.33}$$

$$= -ie^{-i\hat{H}_\Lambda t}\int_0^t dt_0\ e^{i\hat{H}_\Lambda t_0}\hat{V}_\Lambda e^{-i\hat{H}_\Lambda t_0}\left|\phi(0)\right\rangle \tag{4.34}$$

$$= -ie^{-i\hat{H}_\Lambda t}\left|\Lambda+1\right\rangle\int_0^t dt_0\ e^{i\langle\Lambda+1|\hat{H}_\Lambda|\Lambda+1\rangle t_0}\left\langle\Lambda+1\right|\hat{V}_\Lambda\left|\Lambda\right\rangle\left\langle\Lambda\right|e^{-i\hat{H}_\Lambda t_0}\left|\phi(0)\right\rangle \tag{4.35}$$

to leading order. To determine the size of this expression, the behavior of $\left\langle\Lambda\right|e^{-i\hat{H}_\Lambda t_0}\left|\phi(0)\right\rangle$ needs to be understood. For Λ large enough for fragmentation to occur, the phase of this term will be dominated by the diagonal piece of the Hamiltonian so we can approximate $\left\langle\Lambda\right|e^{-i\hat{H}_\Lambda t_0}\left|\phi(0)\right\rangle \approx c\ e^{-i\langle\Lambda|\hat{H}_\Lambda|\Lambda\rangle t_0}$ where c is a slowly varying function with $|c|^2 \leq 1$. With this approximation, we can evaluate this integral and find

$$\left|(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\left|\phi(0)\right\rangle\right| \approx |\left\langle\Lambda+1\right|\hat{V}_\Lambda\left|\Lambda\right\rangle|\left|\int_0^t dt_0 e^{i(\langle\Lambda+1|\hat{H}_\Lambda|\Lambda+1\rangle - \langle\Lambda|\hat{H}_\Lambda|\Lambda\rangle)t_0}\right| \tag{4.36}$$

$$= |\langle \Lambda+1| \hat{V}_\Lambda |\Lambda\rangle| \left| \frac{e^{i(\langle \Lambda+1|\hat{H}_\Lambda|\Lambda+1\rangle - \langle \Lambda|\hat{H}_\Lambda|\Lambda\rangle)t} - 1}{\langle \Lambda+1| \hat{H}_\Lambda |\Lambda+1\rangle - \langle \Lambda| \hat{H}_\Lambda |\Lambda\rangle} \right| \tag{4.37}$$

$$\leq \frac{2 \cdot |\langle \Lambda+1| \hat{V}_\Lambda |\Lambda\rangle|}{\langle \Lambda+1| \hat{H}_\Lambda |\Lambda+1\rangle - \langle \Lambda| \hat{H}_\Lambda |\Lambda\rangle}. \tag{4.38}$$

The above calculation is valid provided that we truncate at least at the first electric field strength, Λ_0, where HSF occurs. As in the previous section, the perturbative expansion performed here is valid when ϵ_{Λ_0} is small where

$$\epsilon_{\Lambda_0} = \frac{|\langle \Lambda_0+1| \hat{V}_{\Lambda_0} |\Lambda_0\rangle|}{\langle \Lambda_0+1| \hat{H}_{\Lambda_0} |\Lambda_0+1\rangle - \langle \Lambda_0| \hat{H}_{\Lambda_0} |\Lambda_0\rangle} \quad . \tag{4.39}$$

However, in practice, we will likely truncate at some $\Lambda > \Lambda_0$ and in this case, we can instead use time-dependent perturbation theory to compute the leading contribution to $\langle \Lambda| e^{-i\hat{H}_\Lambda t_0} \left|\phi(0)\right\rangle$. Restricting $\left|\phi(0)\right\rangle$ to only have support on basis states with electric field below Λ_0, we start by performing a Dyson series expansion as follows

$$e^{-i\hat{H}_\Lambda t_0} \left|\phi(0)\right\rangle = e^{-i\hat{H}_{\Lambda_0} t_0} \sum_{N=0}^{\infty} (-i)^N \int_0^{t_0} dt_1 \int_0^{t_1} dt_2 \cdots \int_0^{t_{N-1}} dt_N \tag{4.40}$$

$$\times \left(\hat{V}_{\Lambda_0,I}(t_1) \hat{V}_{\Lambda_0,I}(t_2) \ldots \hat{V}_{\Lambda_0,I}(t_N) \right) \left|\phi(0)\right\rangle, \tag{4.41}$$

where we have defined

$$\hat{V}_{\Lambda_0,I}(t_i) := e^{i\hat{H}_{\Lambda_0} t_i} (\hat{H}_\Lambda - \hat{H}_{\Lambda_0}) e^{-i\hat{H}_{\Lambda_0} t_i}. \tag{4.42}$$

We note that for the truncation levels accessed in these nested integrals, it is true that

$$\hat{V}_{\Lambda_0,I}(t_i) = e^{i\hat{H}_{\Lambda_0} t_i} \hat{V}_{\Lambda_0} e^{-i\hat{H}_{\Lambda_0} t_i} = e^{i\hat{H}_\Lambda t_i} \hat{V}_\Lambda e^{-i\hat{H}_\Lambda t_i}. \tag{4.43}$$

Since each application of $(\hat{H}_\Lambda - \hat{H}_{\Lambda_0})$ changes the boson number by at most 1, setting $N = \Lambda - \Lambda_0$ results in the leading order contribution to Eq. (4.40). Plugging this into the RHS of Eq. (4.34), we obtain

$$(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t}) \left|\phi(0)\right\rangle = -ie^{-i\hat{H}_\Lambda t} \int_0^t dt_0 \; e^{i\hat{H}_\Lambda t_0} \hat{V}_\Lambda e^{-i\hat{H}_\Lambda t_0} (-i)^{\Lambda-\Lambda_0}$$

$$\times \int_0^{t_0} dt_1 \int_0^{t_1} dt_2 \cdots \int_0^{t_{\Lambda-\Lambda_0-1}} dt_{\Lambda-\Lambda_0} \hat{V}_{\Lambda_0,I}(t_1) \hat{V}_{\Lambda_0,I}(t_2) \ldots \hat{V}_{\Lambda_0,I}(t_{\Lambda-\Lambda_0}) \left|\phi(0)\right\rangle \tag{4.44}$$

to leading order. This expression consists of integrals of exponentials with phases set by the change in energy from applying $\hat{V}_\Lambda$ or $\hat{V}_{\Lambda_0}$. These integrals can be explicitly evaluated as a function of t, using the fact that the operator $\hat{V}_{\Lambda_0}$ only acts on electric basis states $|n\rangle$ with

$n \geq \Lambda_0$ and for those states the Hamiltonian $\hat{H}_{\Lambda_0}$ is diagonal. In particular, the leading order action of $\hat{V}_{\Lambda_0,I}(t_j)$ on a basis state $|\Lambda - j\rangle$ is given by

$$\hat{V}_{\Lambda_0,I}(t_j)\,|\Lambda - j\rangle = \langle \Lambda - j + 1|\,\hat{V}_\Lambda\,|\Lambda - j\rangle\, e^{i\left(\langle \Lambda-j+1|\hat{H}_\Lambda|\Lambda-j+1\rangle - \langle \Lambda-j|H_\Lambda|\Lambda-j\rangle\right)t_j}\,|\Lambda - j + 1\rangle \tag{4.45}$$

$$:= \langle \Lambda - j + 1|\,\hat{V}_\Lambda\,|\Lambda - j\rangle\, e^{i\left(E_{\Lambda-j+1} - E_{\Lambda-j}\right)t_j}\,|\Lambda - j + 1\rangle\,, \tag{4.46}$$

and thus the norm of the LHS from Eq. (4.44) becomes

$$\left|(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\,|\phi(0)\rangle\right| \tag{4.47}$$

$$= \left|\int_0^t dt_0 e^{i\left(E_{\Lambda+1}-E_\Lambda\right)t_0} \cdots \int_0^{t_{\Lambda-\Lambda_0-1}} dt_{\Lambda-\Lambda_0} e^{i\left(E_{\Lambda_0+1}-E_{\Lambda_0}\right)t_{\Lambda-\Lambda_0}} \prod_{k=\Lambda_0}^{\Lambda} \langle k+1|\hat{V}_\Lambda|k\rangle\right| \tag{4.48}$$

$$= \left|\sum_{k=\Lambda_0}^{\Lambda} \left(e^{-i\langle \Lambda+1|\hat{H}_\Lambda|\Lambda+1\rangle t} - e^{-i\langle k|H_\Lambda|k\rangle t}\right) \prod_{\substack{l=\Lambda_0 \\ l\neq k}}^{\Lambda+1} \frac{1}{\langle k|H_\Lambda|k\rangle - \langle l|H_\Lambda|l\rangle}\right|. \tag{4.49}$$

These bounds will be independent of what the initial state was, provided that the initial state only has support on basis states with electric fields below Λ_0.

In practice, one needs to be able to estimate Λ_0. A rough way of doing this is to pick Λ_0 so that the leading order contribution to the leakage error from Eq. (4.34) is smaller than 1.

To make this discussion concrete, we will apply these results to estimate errors in the truncated simulation of a single plaquette in a $U(1)$ lattice gauge theory. In particular, using Eq. (4.24), we find that

$$\left|(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\,|\phi(0)\rangle\right| = \left|\left(\frac{1}{2g^2}\right)^{2(\Lambda+1-\Lambda_0)} \sum_{k=\Lambda_0}^{\Lambda} \left(e^{-i2g^2(\Lambda+1)^2 t} - e^{-i2g^2k^2t}\right) \prod_{\substack{l=\Lambda_0 \\ l\neq k}}^{\Lambda+1} \frac{1}{k^2-l^2}\right| \times \left|c_{\Lambda_0}\,|\Lambda+1\rangle + c_{-\Lambda_0}\,|-\Lambda-1\rangle\right|\,, \tag{4.50}$$

where $c_{\pm\Lambda_0}$ are the slowly varying norms of $\langle \pm\Lambda|\, e^{-i\hat{H}_\Lambda t_{\Lambda_0}}\,|\phi(0)\rangle$. Defining the leakage amplitude by

$$L(g,\Lambda,\Lambda_0,T) = \max_{t<T}\left|\int_0^t dt_0 e^{i\left(E_{\Lambda+1}-E_\Lambda\right)t_0} \cdots \int_0^{t_{\Lambda-\Lambda_0-1}} dt_{\Lambda-\Lambda_0} e^{i\left(E_{\Lambda_0+1}-E_{\Lambda_0}\right)t_{\Lambda-\Lambda_0}}\right| \tag{4.51}$$

$$= \max_{t<T}\left|\sum_{k=\Lambda_0}^{\Lambda}\left[\left(e^{-i2g^2(\Lambda+1)^2 t} - e^{-i2g^2k^2t}\right) \prod_{\substack{l=\Lambda_0 \\ l\neq k}}^{\Lambda+1} \frac{1}{k^2-l^2}\right]\right|, \tag{4.52}$$

where T is the maximum simulation time, the leading error in the truncated state is thus at most

$$\left|(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\left|\phi(0)\right\rangle\right| \leq 2L(g,\Lambda,\Lambda_0,T)\left(\frac{1}{2g^2}\right)^{\Lambda-\Lambda_0+1}. \tag{4.53}$$

To get an idea of how this scales, we can place an upper bound on $L(g,\Lambda,\Lambda_0,T)$. In particular, we note that each of the nested integrals from the integral representation of $L(g,\Lambda,\Lambda_0,T)$ can be bounded as

$$\left|\int_0^{t_j-1} dt_j e^{i\left(E_{\Lambda-j+1}-E_{\Lambda-j}\right)}\right| \leq \frac{2}{E_{\Lambda-j+1}-E_{\Lambda-j}} = \frac{1}{g^2(2(\Lambda-j)+1)}, \tag{4.54}$$

and thus we have that

$$L \leq \prod_{k=\Lambda_0}^{\Lambda} \frac{1}{g^2(2k+1)} = \left(\frac{1}{g^2}\right)^{\Lambda-\Lambda_0+1} \frac{(2\Lambda_0-1)!!}{(2\Lambda+1)!!}. \tag{4.55}$$

Note that this is a loose bound as it neglects destructive interference between many fast oscillating phases in the actual value of the integral. Regardless, this predicts that the leading error due to truncation converges as a factorial and is time-independent.

To understand the performance of these error estimates, we apply these techniques to estimate the error in the expectation of the electric energy $\hat{E}^2$. Using the leading correction to the states, it can be seen that the leading error in the expectation of the electric energy is

$$\left|\langle\phi| e^{i\hat{H}t}\hat{E}^2 e^{-i\hat{H}t}|\phi\rangle - \langle\phi| e^{i\hat{H}_\Lambda t}\hat{E}^2 e^{-i\hat{H}_\Lambda t}|\phi\rangle\right| \leq 2(\Lambda+1)^2\left(\frac{1}{4g^4}\right)^{\Lambda-\Lambda_0+1} L(g,\Lambda,\Lambda_0,T)^2 \tag{4.56}$$

We can now compare this bound against the numerical calculation of this expectation values. $\hat{E}^2$ was measured as a function of time for a single plaquette with $g = 0.5$ and maximum evolution time of $T = 30$. The evolution with $\Lambda = 20$ was treated as the exact evolution and was compared to the evolution with lower truncations. The maximum error in the expectation of $\hat{E}^2$ achieved when beginning in different electric basis states, $|n\rangle$, is shown in Fig. 4.1. As this figure shows, the bound on the error from Eq. (4.56) correctly upper bounds the actual error for all of these initial states. Note that the bound becomes tighter for initial states with larger electric energies.

4.2.2.2 Comparison to Previous Work

Previous work has derived resource estimates for bosonic theories through a combination of rigorous upper bounds on the Dyson series expansion in an interaction picture and constraints from energy conservation [75]. To understand how the current work compares to these results, we will study the analytic bounds of this previous work for the evolution of the electric vacuum state on a single plaquette with a $U(1)$ gauge field. Note that these

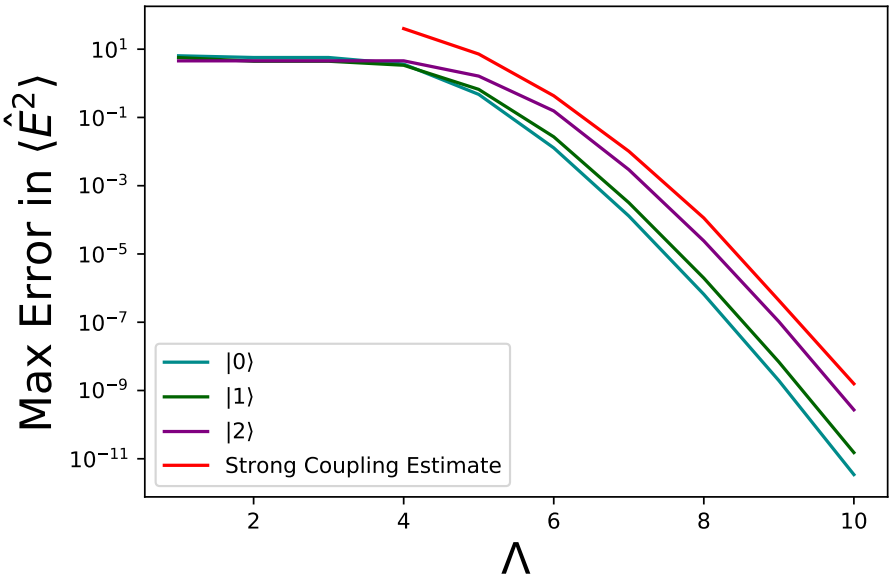

Figure 4.1: Maximum error in the expectation of $\hat{E}^2$ as a function of time on a single plaquette with $g = 0.5$ and a max evolution time of $T = 30$. The blue, green, and purple lines correspond to using different electric basis states as the initial state. The red curve is the error bound in Eq. (4.56), computed using $\Lambda_0 = 4$.

techniques apply to larger system sizes, but for comparison, we will restrict to a single plaquette. The previous bounds on truncation errors were derived by placing upper bounds on the expectation of $\hat{\Pi}_n = |n\rangle \langle n|$ as a function of time. These leakage bounds were then turned into a bound on error in the time evolution operator due to truncation.

One starts from the expectation value of the full Hamiltonian in the electric vacuum $|0\rangle$, which by energy conservation has to be time invariant. Working again with a U(1) pure gauge theory, the total energy is always greater than the electric energy; one therefore obtains

$$\langle 0| \hat{H} |0\rangle = \langle 0| e^{i\hat{H}t} \hat{H} e^{-i\hat{H}t} |0\rangle = \frac{1}{g^2} . \tag{4.57}$$

Since both the electric and magnetic terms in the Hamiltonian are positive definite, we can bound the total energy from above by the electric energy

$$\langle 0| e^{i\hat{H}t} \hat{H} e^{-i\hat{H}t} |0\rangle \geq \langle 0| e^{i\hat{H}t} \hat{H}_E e^{-i\hat{H}t} |0\rangle = 2g^2 \langle 0| e^{i\hat{H}t} \hat{E}^2 e^{-i\hat{H}t} |0\rangle , \tag{4.58}$$

Using the expression of the electric operator in Eq. (4.2), we can bound the electric energy from above by the contributions of a single bosonic state $|\Lambda\rangle$ such giving

$$\langle 0| e^{i\hat{H}t} \hat{E}^2 e^{-i\hat{H}t} |0\rangle \geq \Lambda^2 \langle 0| e^{i\hat{H}t} \hat{\Pi}_\Lambda e^{-i\hat{H}t} |0\rangle . \tag{4.59}$$

Combining everything together, we obtain the bound

$$\frac{1}{g^2} \geq 2g^2\Lambda^2 \langle 0| e^{i\hat{H}t}\hat{\Pi}_\Lambda e^{-i\hat{H}t} |0\rangle \ , \tag{4.60}$$

or

$$\langle 0| e^{i\hat{H}t}\hat{\Pi}_\Lambda e^{-i\hat{H}t} |0\rangle \leq \frac{1}{2g^4\Lambda^2} \, , \tag{4.61}$$

for all times t. Note that for a larger lattice, similar volume-independent bounds can be derived for translationally invariant states.

For short times, a time-dependent bound can be derived by upper-bounding the error in a Dyson series expansion. Explicitly, one works in the interaction picture where the free part of the Hamiltonian is given by $2g^2\hat{E}^2$. The Hamiltonian in the interaction picture is then given by

$$\hat{H}_I(t) = -\frac{1}{2g^2} \sum_n e^{i2g^2t(2n+1)} |n+1\rangle \langle n| + \text{h.c.} \, . \tag{4.62}$$

Denoting the electric vacuum evolved in the interaction picture by $\big|\phi_I(t)\big\rangle$, the expectation of the projector $\hat{\Pi}_\Lambda$ is given by

$$\langle 0| e^{i\hat{H}t}\hat{\Pi}_\Lambda e^{-i\hat{H}t} |0\rangle = \Big|\langle \Lambda| e^{-i\hat{H}t} |0\rangle\Big|^2 = \Big|\langle \Lambda| \big|\phi_I(t)\big\rangle\Big|^2 \, . \tag{4.63}$$

The Dyson series expansion is generated by inserting $\big|\phi_I(t)\big\rangle$ recursively into the equation

$$\big|\phi_I(T)\big\rangle = \big|\phi_I(0)\big\rangle - i\int_0^T dt \ \hat{H}_I(t) \big|\phi_I(t)\big\rangle \ . \tag{4.64}$$

The quantity $\langle \Lambda| \big|\phi_I(t)\big\rangle$ only receives a non-zero contribution from the Λ-th order which is given by

$$\langle \Lambda| \big|\phi_I(t)\big\rangle = (-i)^\Lambda \int_0^t dt_1 \int_0^{t_1} dt_2 \cdots \int_0^{t_{\Lambda-1}} dt_\Lambda \hat{H}_I(t_1)\hat{H}_I(t_2)\cdots\hat{H}_I(t_\Lambda) \big|\phi_I(t_\Lambda)\big\rangle \ . \tag{4.65}$$

Since the spectral norm of $H_I(t)$ is bounded by

$$\Big|\big|H_I(t)\big|\Big| \leq \frac{1}{g^2} \, , \tag{4.66}$$

the magnitude of this overlap is upper bounded by

$$\Big|\langle \Lambda| \big|\phi_I(t)\big\rangle\Big| \leq \Big|\big|H_I(t)\big|\Big|^\Lambda \int_0^t dt_1 \int_0^{t_1} dt_2 \cdots \int_0^{t_{\Lambda-1}} dt_\Lambda = \frac{t^\Lambda}{g^{2\Lambda}\Lambda!} \, . \tag{4.67}$$

We therefore find

$$\langle 0| e^{i\hat{H}t}\hat{\Pi}_\Lambda e^{-i\hat{H}t} |0\rangle \leq \frac{t^{2\Lambda}}{g^{4\Lambda}\Lambda!^2} \,. \tag{4.68}$$

Similar techniques can be used to show

$$\left|\hat{\Pi}_\Lambda e^{-i\hat{H}t}\hat{\Pi}_{\Lambda-\Delta}\right| \leq \frac{t^\Delta}{g^{2\Delta}\Delta!} \,. \tag{4.69}$$

These results are referred to as short-time bounds, as at long times they will exceed the bound derived from energy conservation.

Time dependent bounds that are valid for longer times were also derived by inserting intermediate projectors and bounding individual terms in the sum. For example, one can consider adding one intermediate projector, defining

$$\begin{aligned} \hat{\Pi}_{>\Lambda} &= \sum_{|n|>\Lambda} |n\rangle \langle n| \\ \hat{\Pi}_{\leq\Lambda} &= \sum_{|n|\leq\Lambda} |n\rangle \langle n| \,. \end{aligned} \tag{4.70}$$

This allows to derive

$$\begin{aligned} \left|\hat{\Pi}_\Lambda e^{-i\hat{H}t}\hat{\Pi}_{\Lambda-\Delta_1-\Delta_2}\right| &= \left|\hat{\Pi}_\Lambda e^{-i\hat{H}(t-t_1)}\left(\hat{\Pi}_{>\Lambda-\Delta_1}+\hat{\Pi}_{\leq\Lambda-\Delta_1}\right)e^{-i\hat{H}t_1}\hat{\Pi}_{\Lambda-\Delta_1-\Delta_2}\right| \\ &\leq \left|\hat{\Pi}_{>\Lambda-\Delta_1}e^{-i\hat{H}t_1}\hat{\Pi}_{\Lambda-\Delta_1-\Delta_2}\right| + \left|\hat{\Pi}_\Lambda e^{-i\hat{H}(t-t_1)}\hat{\Pi}_{\leq\Lambda-\Delta_1}\right| \\ &\leq \frac{t_1^{\Delta_2}}{g^{2\Delta_2}\Delta_2!} + \frac{(t-t_1)^{\Delta_1}}{g^{2\Delta_1}\Delta_1!} \,, \end{aligned} \tag{4.71}$$

and t_1 can be chosen to minimize this expression. To obtain a long-term bound for the expectation value $\langle 0| e^{i\hat{H}t}\hat{\Pi}_\Lambda e^{-i\hat{H}t} |0\rangle$ one can insert up to Λ projectors and optimize over the intermediate times. These are referred to as long-time bounds as they take longer to saturate the bound from energy conservation.

In summary, the work of [75] derived three types of bounds, namely time-independent bounds based on energy conservation, as well as time-dependent bounds valid at short and longer times. These can be combined into a more optimal time-dependent bound, by taking their minimum. To determine the tightness of the resulting bound, the time evolution of the electric vacuum state will be simulated numerically. Note that

$$\langle 0| e^{i\hat{H}t}\hat{\Pi}_\Lambda e^{-i\hat{H}t} |0\rangle \leq \left|\hat{\Pi}_\Lambda e^{-i\hat{H}t}\hat{\Pi}_0\right|^2 \,, \tag{4.72}$$

so these bounds directly translate to bounds on the expectation of operators. The electric vacuum was evolved under a Hamiltonian with a truncation of $\Lambda = 20$ and $g = 0.5$. The solid

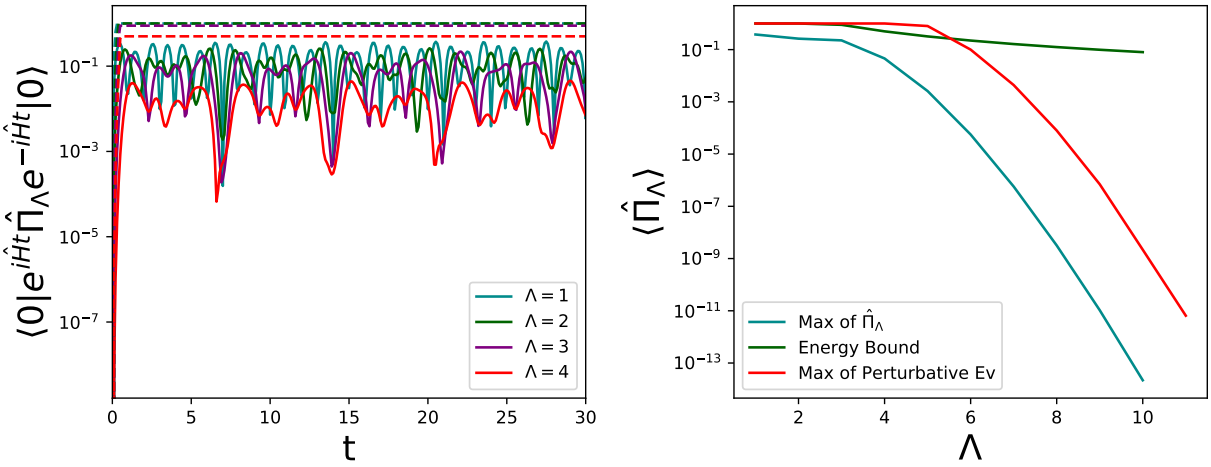

Figure 4.2: Evolution of the electric vacuum state on a single plaquette. Numerical simulations were performed with a maximum electric field of 20 and $g = 0.5$. The left panel shows the expectation of $\hat{\Pi}_\Lambda$ for various values of Λ as a function of time t. The solid curves are the exact time evolution, and the dashed lines are the rigorous long-time bounds. The right panel shows the maximum of the expectation of $\hat{\Pi}_\Lambda$ for the simulated time evolution. The blue curve is the exact result, the green curve is the bound from energy conservation, and the red curve is the leading contribution to the expectation obtained by calculating $L(g, \Lambda, \Lambda_0, T)$ with $\Lambda_0 = 4$.

lines in the left panel of Fig. 4.2 show the exact time evolution of the expectation of $\hat{\Pi}_\Lambda$ and the dashed lines show the bounds as derived above (the minimum of the three approaches discussed). As this figure shows, the long-time bounds quickly saturate and overestimate the expectation of $\hat{\Pi}_\Lambda$. The figure on the right shows the behavior of the maximum of the expectation value of the projection operator taken over large times, as a function of the truncation Λ. In blue we show the result from the numerical simulation, which shows that the maximum expectation for $\Lambda = 10$ is below 10^{-13}. The green curve is the result of [75] derived in this section, which clearly overestimates the true value by many orders of magnitude. We also show in red the new results of this work, which were derived in the previous section. One can see that these results provide a much tighter bound.

As discussed in [75], the bounds derived in that work imply that to guarantee that $\left|\langle k+\Lambda| e^{-i\hat{H}t} |k\rangle\right|$ is less than 0.01 for $g = \sqrt{3}$ for a maximum evolution time of $t = 8$ would require a cutoff satisfying $\Lambda > 100$. On the other hand, using the upper bound derived in this work for $L(g, \Lambda, \Lambda_0, T)$, given in Eq. (4.55), one estimates that $\left|\langle k+\Lambda| e^{-i\hat{H}t} |k\rangle\right| < 6\times10^{-308}$ for these parameters. This shows in a pretty dramatic fashion how much tighter the bounds derived in this work are. To verify this claim numerically, the evolution of the electric

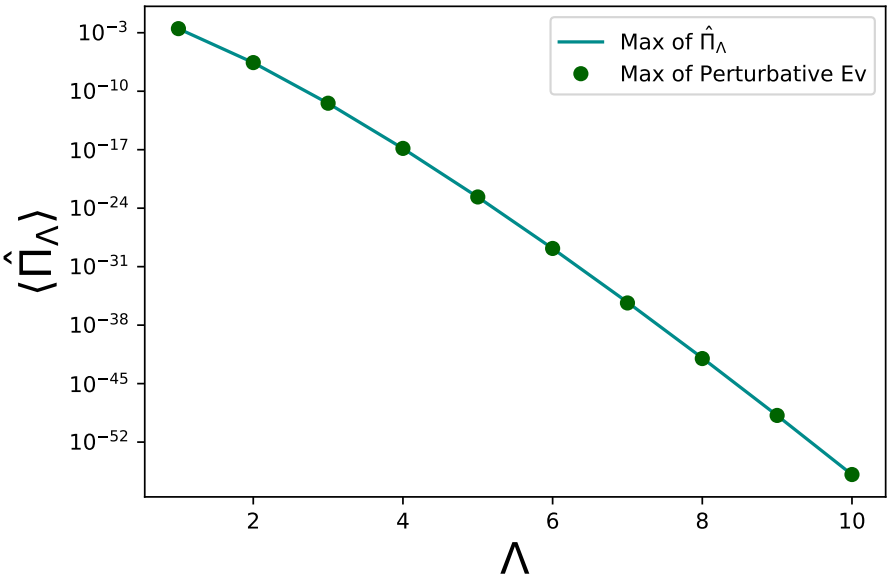

Figure 4.3: Evolution of the electric vacuum state on a single plaquette. Numerical simulations were performed with a maximum electric field of 20, $g = \sqrt{3}$, and a maximum evolution time of $t = 8$. The blue curve is the maximum of the expectation of $\hat{\Pi}_\Lambda$ during this evolution. The green points are the leading contribution to the expectation obtained by calculating $L(g, \Lambda, \Lambda_0, T)$ with $\Lambda_0 = 0$.

vacuum on a single plaquette was simulated for $g = \sqrt{3}$ and a maximum evolution time of $t = 8$. Fig. 4.3 shows the maximum value reached by the expectation of $\hat{\Pi}_\Lambda$ during this time evolution and the predicted maximum for the perturbative calculation done in this work. As this figure shows, the predicted maximum is in good agreement with the numerical simulation, indicating the tightness of the bounds derived in this work.

The bounds obtained in [75] are extremely loose due to the bounds being computed by applying the triangle inequality to upper-bound the Dyson series expansion. The Dyson series consists of a sum over many oscillating phases, and the actual physics of the system comes from how these phases constructively or destructively interfere. As the triangle inequality removes these phases, it should not be surprising that these bounds fail to come remotely close to the qualitative behavior of the system. The estimate for the maximum of $\hat{\Pi}_\Lambda$ obtained using the strong coupling expansion appears to fall off at the same rate as the actual system, but is still a loose overestimate. This is because the strong coupling expansion cannot be applied to states with small electric fields, and the amplitude for leakage into the fragmented sector of the theory was upper-bounded by 1.

4.2.3 Extension to Larger Systems

To extract new predictions from simulations of lattice gauge theories on quantum computers, it is necessary to simulate lattices with more than one plaquette and have estimates of the errors due to truncations of the gauge fields. This can be done using similar techniques to the single plaquette case. We will write the Hamiltonian for the untruncated theory, $\hat{H}$, as a sum of the truncated Hamiltonian, $\hat{H}_\Lambda$, plus a sum over $\hat{V}_\Lambda^{\vec{x}}$ which contains all off-diagonal couplings for electric fields above the truncation Λ for the plaquette at position $\vec{x}$, explicitly

$$\hat{H} = \hat{H}_\Lambda + \sum_x \hat{V}_\Lambda^{\vec{x}} \,. \tag{4.73}$$

Following the same approach as Section 4.2.2.1, the leading correction to the truncated evolution is given by

$$(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t}) \left|\phi(0)\right\rangle = \sum_{\vec{x}} -ie^{-i\hat{H}_\Lambda t} \int_0^t ds\ e^{i\hat{H}_\Lambda s} \hat{V}_\Lambda^{\vec{x}} e^{-i\hat{H}_\Lambda s} \left|\phi(0)\right\rangle \,. \tag{4.74}$$

The quantity $\hat{V}_\Lambda^{\vec{x}} e^{-i\hat{H}_\Lambda s} \left|\phi(0)\right\rangle$ can be estimated to leading order by applying time dependent perturbation theory in the fragmented subspace as before to give at leading order

$$\begin{aligned}
&(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t}) \left|\phi(0)\right\rangle \qquad (4.75)\\
&= -ie^{-i\hat{H}_\Lambda t} \sum_{\vec{x}} \int_0^t dt_0\ e^{i\hat{H}_\Lambda t_0} \hat{V}_\Lambda^{\vec{x}} e^{-i\hat{H}_\Lambda t_0} (-i)^{\Lambda-\Lambda_0} \int_0^{t_0} dt_1 e^{i\hat{H}_{\Lambda_0} t_1} \hat{V}_{\Lambda_0}^{\vec{x}} e^{-i\hat{H}_{\Lambda_0} t_1}\\
&\times \int_0^{t_1} dt_2 e^{i\hat{H}_{\Lambda_0} t_2} \hat{V}_{\Lambda_0}^{\vec{x}} e^{-i\hat{H}_{\Lambda_0} t_2} \cdots \int_0^{t_{\Lambda-\Lambda_0+1}} dt_{\Lambda-\Lambda_0} e^{i\hat{H}_{\Lambda_0} t_{\Lambda-\Lambda_0}} \hat{V}_{\Lambda_0}^{\vec{x}} e^{-i\hat{H}_{\Lambda_0} t_{\Lambda-\Lambda_0}} \left|\phi(0)\right\rangle \,. \qquad (4.76)
\end{aligned}$$

As before, the integrand in this expression consists of phases given by changes in energy from applying $\hat{V}_{\Lambda_0}^{\vec{x}}$ or $\hat{V}_\Lambda^{\vec{x}}$. These changes in energy can be approximated by the change in electric energy. This gives the leading corrections to the state vector and can be used to evaluate the leading error in the expectation of observables. Note that at leading order, local observables will only have errors coming from the $\hat{V}_\Lambda^{\vec{x}}$ terms that share support with the observable.

To make this discussion concrete, the truncation error on the expectation value of an operator will be estimated for a pure $U(1)$ lattice gauge theory on a plaquette chain. The Hamiltonian is given by

$$\hat{H} = \sum_l \frac{g^2}{2} \hat{E}_l^2 + \sum_p \frac{1}{2g^2} \left(2 - \hat{\square}_p - \hat{\square}_p^\dagger\right) \,, \tag{4.77}$$

where the sum over l corresponds to summing over all links on the lattice and the sum over p corresponds to summing over all plaquettes on the lattice. This Hamiltonian can be gauge-fixed so gauge-invariant states are given by specifying the electric field on a single plaquette

per link [122]. The gauge-fixed Hamiltonian is given by

$$\begin{aligned}\hat{H} &= \sum_p \left\{ g^2 \left[\hat{E}_p^2 + \frac{1}{2}\Big(\hat{E}_p - \hat{E}_{p+1}\Big)^2 \right] + \frac{1}{2g^2}\left(2 - \hat{\square}_p - \hat{\square}_p^\dagger\right) \right\} \\ \hat{E}_p &= \sum_{n_p=-\infty}^{\infty} n_p \left|n_p\right\rangle \left\langle n_p\right| \\ \hat{\square}_p &= \sum_{n_p=-\infty}^{\infty} \left|n_p\right\rangle \left\langle n_p+1\right| \,, \end{aligned} \tag{4.78}$$

and for a chain of length L the electric basis states are given by $|n_1, n_2, \cdots n_L\rangle$. The operator we will choose for this example will be the electric energy on the four links of a given plaquette, and for concreteness we choose the time-dependent expectation value of $\hat{E}_p^2$. This operator will receive contributions from plaquette p, as well as the neighboring plaquettes $p \pm 1$. This implies that the truncation errors can be estimated using a 3 plaquette chain, with a generic state given by $|n_1, n_2, n_3\rangle$. To simplify the notation in the calculation, only the error from the truncation on positive electric fields will be estimated (the contribution from negative electric fields will only contribute a factor of 2).

To start, we write as before

$$\left|\phi(0)\right\rangle = \sum_{n_1,n_2,n_3 \leq \Lambda_0} c_{n_1,n_2,n_3} \left|n_1, n_2, n_3\right\rangle \,. \tag{4.79}$$

Following similar steps as before, the time-evolved perturbation can be written as

$$e^{i\hat{H}_{\Lambda_0}t}\hat{V}_{\Lambda_0}^{(2)}e^{-i\hat{H}_{\Lambda_0}t}\left|\phi(0)\right\rangle = \frac{-1}{2g^2}\sum_{n_1,n_3} c_{n_1,\Lambda_0,n_3} e^{itg^2(4\Lambda_0 - n_1 - n_3 + 2)} \left|n_1, \Lambda_0 + 1, n_3\right\rangle \,. \tag{4.80}$$

Defining the function

$$A_{n,m}(g, \Lambda, \Lambda_0, T) = \left(\frac{-i}{2g^2}\right)^{\Lambda - \Lambda_0 + 1} \prod_{k=\Lambda_0}^{\Lambda} \int_0^{t_{k+1}} \mathrm{d}t_k \, e^{ig^2(4k - n - m + 2)t_k} \,, \tag{4.81}$$

with $T \equiv t_{\Lambda+1}$ the correction to the state vector is therefore given by

$$(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\left|\phi(0)\right\rangle = \sum_{n,m} A_{n,m}(g, \Lambda, \Lambda_0, t) c_{n,\Lambda_0,m} \left|n, \Lambda + 1, m\right\rangle \,. \tag{4.82}$$

Combining this, we obtain the correction to the expectation of the electric energy on the center plaquette ($\hat{E}_2^2$),

$$\left\langle\phi\right| e^{i\hat{H}t}\hat{E}_2^2 e^{-i\hat{H}t}\left|\phi\right\rangle - \left\langle\phi\right| e^{i\hat{H}_\Lambda t}\hat{E}_2^2 e^{-i\hat{H}_\Lambda t}\left|\phi\right\rangle$$

$$= \sum_{n,m\leq\Lambda_0} \left|A_{n,m}(g,\Lambda,\Lambda_0,t)\right|^2 \left[(\Lambda+1)^2\left|c_{n,\Lambda_0,m}\right|^2 + m^2\left(\left|c_{\Lambda_0,m,n}\right|^2 + \left|c_{n,m,\Lambda_0}\right|^2\right)\right] . \tag{4.83}$$

Using the fact that $\sum_{n_1,n_2,n_3\leq\Lambda_0}\left|c_{n_1,n_2,n_3}\right|^2 = 1$, the error can be upper bounded by

$$\begin{aligned}&\left|\langle\phi|\, e^{i\hat{H}t}\hat{E}_2^2 e^{-i\hat{H}t}\,|\phi\rangle - \langle\phi|\, e^{i\hat{H}_\Lambda t}\hat{E}_2^2 e^{-i\hat{H}_\Lambda t}\,|\phi\rangle\right| \leq \\ &\max_{\tau<t}\left\{(\Lambda+1)^2\left|A_{n,m}(g,\Lambda,\Lambda_0,\tau)\right|^2 : n,m\leq\Lambda_0\right\} \\ &+ 2\max_{\tau<t}\left\{n^2\left|A_{n,m}(g,\Lambda,\Lambda_0,\tau)\right|^2 : n,m\leq\Lambda_0\right\} .\end{aligned} \tag{4.84}$$

The above expression is only for truncating the positive electric fields with strength above Λ. When truncating the positive and negative electric fields with strength above Λ, the resulting expression for the leading error in the electric field evolution is given by

$$\begin{aligned}&\left|\langle\phi|\, e^{i\hat{H}t}\hat{E}_2^2 e^{-i\hat{H}t}\,|\phi\rangle - \langle\phi|\, e^{i\hat{H}_\Lambda t}\hat{E}_2^2 e^{-i\hat{H}_\Lambda t}\,|\phi\rangle\right| \\ &\leq 2\max_{\tau<t}\left\{(\Lambda+1)^2\left|A_{n,m}(g,\Lambda,\Lambda_0,\tau)\right|^2 : |n|,|m|\leq\Lambda_0\right\} \\ &+ 4\max_{\tau<t}\left\{n^2\left|A_{n,m}(g,\Lambda,\Lambda_0,\tau)\right|^2 : |n|,|m|\leq\Lambda_0\right\} .\end{aligned} \tag{4.85}$$

Note that while the expression was derived using a three-plaquette lattice, the result will generalize to a plaquette chain of arbitrary length. The only information used about the initial state was that it did not have support on electric fields above Λ_0. This bound could potentially be made tighter by using energy conservation arguments to upper bound the probability of having a state with an electric field of Λ_0 present. For translationally invariant states, these energy bounds will have no volume dependence. This bound is based on a perturbative expansion of the dynamics connecting different fragmented sectors of the Hilbert space. This is an expansion in powers of

$$\epsilon_{\Lambda_0} = \frac{1}{4g^4(\Lambda_0+1)} \quad . \tag{4.86}$$

Provided that Λ_0 is taken to be sufficiently large, this bound can be used to estimate truncation errors at any value of g, including when g is small enough to probe continuum physics.

To demonstrate the performance of these error estimates, an infinite MPS was used to simulate the time evolution of an infinite plaquette chain with $g = 0.8$ [268]. The system was initialized in the electric vacuum and evolved until $t = 8$. The details of the implementation are in Appendix C.1. Fig. 4.4 shows the expectation of $\hat{E}^2$ for a single plaquette on the lattice as a function of time. The truncation error in the time evolution is consistent with the estimate of the error computed in Eq. (4.85). This agreement suggests that truncation errors in local observables do not have a system size dependence as predicted from the perturbative calculation.

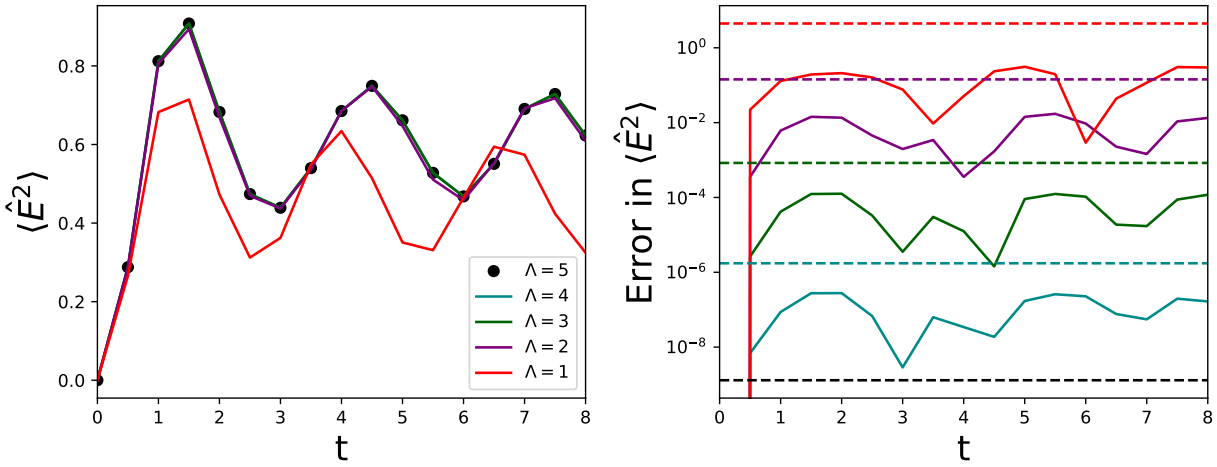

Figure 4.4: Simulation of an infinite plaquette chain with $g = 0.8$. The left panel shows the electric energy as a function of time for different truncations of the electric field. The solid curves in the right panel show the difference in electric energy between the $\Lambda = 5$ truncation and the lower truncations. The dashed lines show the predicted error in the electric energy from truncating the Hamiltonian, computed using Eq. (4.85) with $\Lambda_0 = 1$. See App. C.1 for details regarding the iMPS implementation.

The probability of exciting a given link on the lattice to a large electric field on a plaquette chain can be estimated using this framework. Denoting $\hat{\Pi}_{\Lambda,\vec{x}} = |\Lambda_{\vec{x}}\rangle \langle \Lambda_{\vec{x}}|$, the probability of exciting a link to electric field Λ when the initial state, $|\phi\rangle$ has all electric fields below Λ_0 is upper bounded by

$$\langle \phi | \, e^{i\hat{H}t} \hat{\Pi}_{\Lambda,\vec{x}} e^{-i\hat{H}t} \, | \phi \rangle \leq P(\Lambda, \Lambda_0, g, t) = \max_{\tau<t} \left\{ \left| A_{n,m}(g, \Lambda, \Lambda_0, \tau) \right|^2 : |n|, |m| \leq \Lambda_0 \right\} \quad (4.87)$$

provided that HSF is present for states with electric fields above Λ_0. From Eq. (4.81), it follows that $P(\Lambda, \Lambda_0, g, t)$ is upper bounded by

$$P(\Lambda, \Lambda_0, g, t) \leq \frac{1}{g^{8(\Lambda - \Lambda_0)}} \times f(\Lambda, \Lambda_0) \, , \quad (4.88)$$

where $f(\Lambda, \Lambda_0)$ is a g independent function. Fig. 4.5 shows a numerical demonstration of the validity of this upper bound. Using an infinite MPS we again simulate the time evolution of an infinite plaquette chain with four different values of $g = 1.0, 0.9, 0.8$, and 0.7. Starting from the electric vacuum, we perform time evolution until $t = 8$ and compute the expectation value $\langle \hat{\Pi}_{\Lambda,\vec{x}} \rangle$ as a function of time. We observe that our numerical results are consistent with Eq. (4.88). For further implementation details pertaining to the time evolution see App. C.1. Fig. 4.5 also shows a comparison between the bound shown in Eq. (4.88) and

the time independent energy conservation based bound shown in Eq. (4.61). In particular, we show that for the largest value of Λ considered for each value of g, the bound based on energy conservation is between 5 and 8 orders of magnitude larger than the bound based on Eq. (4.88).

4.3 Truncated Schwinger Model

Simulations of lattice QCD will require including quarks in addition to the gauge theory degrees of freedom. To understand how the inclusion of matter affects the truncation error, we will apply our framework to the Schwinger model. The Hamiltonian is given by

$$\begin{aligned}\hat{H} &= \hat{H}_K + \hat{H}_m + \hat{H}_E \\ \hat{H}_K &= \frac{1}{2}\sum_x \hat{\psi}^\dagger_{x+1}\hat{U}_x\hat{\psi}_x + \text{h.c.} \\ \hat{H}_m &= m\sum_x (-1)^x \hat{\psi}^\dagger_x\hat{\psi}_x \\ \hat{H}_E &= \frac{g^2}{2}\sum_x \hat{E}^2_x\,,\end{aligned} \tag{4.89}$$

where $\hat{\psi}_x$ is the fermion field at site x, $\hat{U}_x$ is the parallel transporter between sites x and $x+1$, and $\hat{E}_x$ is the corresponding electric field. Note that in this theory, all gauge fields are sourced by fermions somewhere on the lattice. To have an electric field of strength Λ, there needs to be at least Λ positively charged particles and Λ negatively charged particles. This means that to raise the maximum electric field in a state by Δ, Δ pairs of particles and anti-particles must be created. Furthermore, one can have at most one fermion on each lattice site, such that all Λ fermions need to sit on one side of the link with electric field Λ, while the anti-fermions sit on the other side. Starting from the electric vacuum, creating a single electric field requires a single hopping term. Creating an additional electric flux on this link requires first moving the fermion and anti-fermion created in the first step by two lattice units (because of staggering) to the left and right, requiring $2 \times 2 = 4$ more hopping terms, and then adding an additional fermion anti-fermion pair. This requires a total of 6 hopping terms. Creating 3 units of electric flux requires moving the two fermion anti-fermion pairs once more, requiring $4 \times 2 = 8$ before adding the extra pair, giving 15 hopping terms. In general, to create n units of flux requires at least $n(2n-1)$ hopping terms. This complicates the perturbative calculation of the leakage probabilities compared to the pure gauge case, but should also lead to a stronger suppression of states with large electric fields. To perform the perturbative estimate, the truncated Hamiltonian $\hat{H}_\Lambda$ is related to the untruncated Hamiltonian by

$$\hat{H} = \hat{H}_\Lambda + \sum_x \hat{V}^x_\Lambda$$

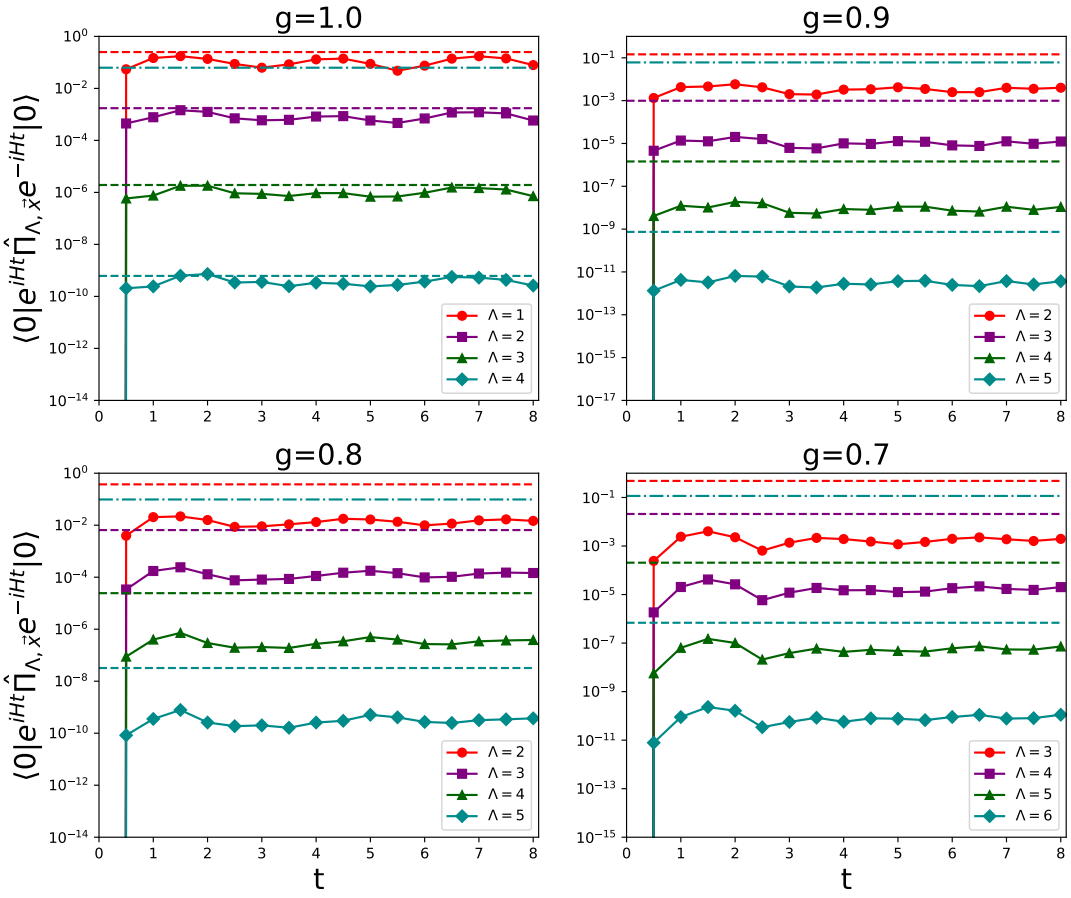

Figure 4.5: Simulation of an infinite plaquette chain with four different values of $g = 1.0, 0.9, 0.8,$ and 0.7 and maximum electric fields of $5, 6, 6,$ and 7 respectively. The solid curves in each panel show the expectation value of the projection operator $\hat{\Pi}_{\Lambda,\vec{x}}$ with respect to the states obtained by time evolving the electric vacuum for various choices of Λ. See App. C.1 for details regarding the iMPS time evolution. The dashed lines in each panel correspond to the upper bound on the probability given by Eq. (4.88) with $\Lambda_0 = 0, 1, 1,$ and 2 respectively for each choice of g. The dash-dotted lines in each panel show the upper bound based on energy conservation, computed using Eq. (4.61) for the largest value of Λ considered for each value of g.

$$\hat{V}_\Lambda^x = \frac{1}{2}\hat{\psi}_{x+1}^\dagger \hat{U}_x^\Lambda \hat{\psi}_x + \text{h.c.}$$
$$\hat{U}_x^\Lambda = \sum_{n>\Lambda} |n-1\rangle\langle n| + \sum_{n\leq -\Lambda} |n-1\rangle\langle n| \ . \tag{4.90}$$

Following the same approach as the previous sections, the leading correction to the truncated time evolution is given by

$$(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\left|\phi(0)\right\rangle = \sum_x -ie^{-i\hat{H}_\Lambda t}\int_0^t ds\ e^{i\hat{H}_\Lambda s}\hat{V}_\Lambda^x e^{-i\hat{H}_\Lambda s}\left|\phi(0)\right\rangle \ . \tag{4.91}$$

Assuming that Λ is chosen to be the electric field strength where HSF begins, the leakage amplitude to go from an electric field of Λ to $\Lambda+1$ can be estimated using a lattice with two staggered sites (one physical site), open boundary conditions, and an external electric field. For a system of this size, the gauge invariant states can be specified by basis states $|E, f_0, f_1\rangle$ where E is an external electric field and f_i are the fermion occupation numbers on each site. $f_0 = 1$ corresponds to a quark being present on site 0 and $f_1 = 0$ corresponds to an anti-quark being present on site 1. The leading order contribution to the truncation error is

$$(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\left|\phi(0)\right\rangle \tag{4.92}$$
$$= -ie^{-it(\frac{g^2}{2}(\Lambda+1)^2+2m)}|\Lambda+1,1,0\rangle\langle\Lambda+1,1,0|\int_0^t dt' e^{i\hat{H}_\Lambda t'}\hat{V}_\Lambda e^{-i\hat{H}_\Lambda t'}\left|\phi(0)\right\rangle$$
$$= \frac{-ie^{-it(\frac{g^2}{2}(\Lambda+1)^2+2m)}}{2}|\Lambda+1,1,0\rangle\int_0^t dt' c_{\Lambda,0,1}e^{it(2m+\frac{g^2}{2}(2\Lambda+1))} \tag{4.93}$$

where $c_{\Lambda,0,1}$ is defined analogously to the previous sections. This quantity is upper bounded by

$$\left|(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\left|\phi(0)\right\rangle\right| \leq \frac{1}{2m + g^2(\Lambda+1/2)} \ . \tag{4.94}$$

Note that the amplitude to increase the electric field is suppressed when either m is large or $g^2\Lambda$ is large. Increasing the electric field by two requires 6 hopping terms acting on 6 staggered sites. As before, the state of a lattice with 6 staggered sites is given by an external electric field and the fermion occupation number on 6 sites. Therefore, if HSF begins at $\Lambda-1$, the leading error in the state on a lattice with 6 staggered sites is given by the expression

$$(e^{-i\hat{H}t} - e^{-i\hat{H}_\Lambda t})\left|\phi(0)\right\rangle$$
$$= 14(-i)^6 e^{-i\hat{H}_\Lambda t}\left(\prod_{k=1}^6\int_0^{t_{k-1}} dt_k\right)\hat{V}_\Lambda^3(t_1)\hat{V}_{\Lambda-1}^4(t_2)\hat{V}_{\Lambda-1}^2(t_3)\hat{V}_{\Lambda-1}^5(t_4)\hat{V}_{\Lambda-1}^3(t_5)\hat{V}_{\Lambda-1}^1(t_6)\left|\phi(0)\right\rangle$$
$$= -14\,|\Lambda-1,1,1,1,1,0,0,0\rangle\left(\prod_{k=1}^6\int_0^{t_{k-1}} dt_k\right)$$

$$\times \langle \Lambda-1,1,1,1,0,0,0| \hat{V}^3_\Lambda(t_1)\hat{V}^4_{\Lambda-1}(t_2)\hat{V}^2_{\Lambda-1}(t_3)\hat{V}^5_{\Lambda-1}(t_4)\hat{V}^3_{\Lambda-1}(t_5)\hat{V}^1_{\Lambda-1}(t_6)\left|\phi(0)\right\rangle$$

$$= -14\,|\Lambda-1,1,1,1,0,0,0\rangle \left(\prod_{k=1}^{6}\int_0^{t_{k-1}} dt_k\right) \frac{1}{2^6} e^{it_1(2m+2g^2(\Lambda+1))} e^{it_2(-2m+g^2(\Lambda+\frac{1}{2}))}$$

$$\times e^{it_3(-2m+g^2(\Lambda+\frac{1}{2}))} e^{it_4(2m+g^2(\Lambda-\frac{1}{2}))} e^{it_5(2m+g^2(\Lambda-\frac{1}{2}))} e^{it_6(2m+g^2(\Lambda-\frac{1}{2}))} c_{\Lambda-1,0,1,0,1,0,1}$$

$$= A(\Lambda,g,m,t) c_{\Lambda-1,0,1,0,1,0,1}\,|\Lambda-1,1,1,1,0,0,0\rangle \tag{4.95}$$

where $\hat{V}^x_\Lambda(t) = e^{i\hat{H}_\Lambda t}\hat{V}^x_{\Lambda_0}e^{-i\hat{H}_\Lambda t}$, $t_0 = t$, and $A(\Lambda,g,m,t)$ has been defined to be the integral in the above expression. The factor of 14 comes from counting the number of possible orderings of operators that can be applied to raise the electric field on the center link by two. One possible ordering of operators to raise the field truncation is shown in Fig. 4.6. From this expression, it can be seen that the leading contribution to the error scales as $1/(g^2\Lambda)^6$. This is a much more rapid falloff than in the equivalent scenario in the pure gauge theory. In high spatial dimensions with both gauge fields and matter, the electric field on a link can be raised by applying either plaquette or link operators. The leading terms in the truncation error will contain at most one hopping operator, and the rest will be plaquette operators. Consequently, the truncation errors in the $1+1D$ Schwinger model will not be representative of the behavior of truncation errors in higher spatial dimensions. However, the more rapid convergence of truncation errors in $1+1D$ suggests that non-trivial simulations of $1+1D$ dynamics with controlled uncertainties may be within near-term reach.

As in the pure gauge case, these estimates of leakage errors can be used to determine the truncation error of local observables. The error in a local observable will be given by the expectation of the observable in the state created by applying $\sum_x \hat{V}^x_\Lambda$ to the truncated state weighted by the probability of reaching that state. If the expression in Eq. (4.94) is used to estimate the leakage probability, the error in the electric energy ($\hat{E}^2_l$) and chiral condensate ($\hat{\chi}_x = (-1)^x\hat{\psi}^\dagger_x\hat{\psi}_x$) is given by

$$\delta E^2 = (2\Lambda^2 + (\Lambda+1)^2) \times \left(\frac{1}{2m+g^2(\Lambda+1/2)}\right)^2$$
$$\delta\chi = 2 \times \left(\frac{1}{2m+g^2(\Lambda+1/2)}\right)^2. \tag{4.96}$$

Note that the electric energy receives a contribution from the link itself being excited to $\Lambda+1$ and when its neighbor is excited to $\Lambda+1$. For higher truncations, HSF will occur at electric field strengths below Λ, and the expression in Eq. (4.95) can be used to estimate the error. Defining

$$M(\Lambda,g,m,t) = \max_{\tau<t}\left|A(\Lambda,g,m,\tau)\right|, \tag{4.97}$$

the error in the electric energy and chiral condensate is given by

$$\delta E^2 = \left(2(\Lambda-1)^2 + 4\Lambda^2 + (\Lambda+1)^2\right) M^2(\Lambda,g,m,t)$$
$$\delta\chi = 4M^2(\Lambda,g,m,t), \tag{4.98}$$

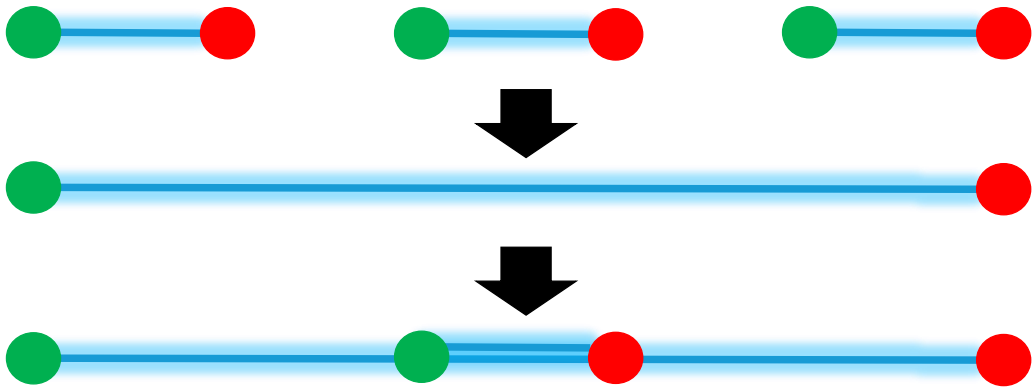

Figure 4.6: One possible way of increasing electric field strength by 2. The green circles correspond to positive charges, and the red circles correspond to negative charges. First, the hopping operators are applied to every other link to create three $q\bar{q}$ pairs. The other hopping operators are then applied to remove the $q\bar{q}$ pairs in the center of the lattice. In the last step, a $q\bar{q}$ pair is excited in the center of the lattice.

where $A(\Lambda, g, m, \tau)$ is defined in Eq. (4.95).

To demonstrate the performance of these error estimates, the dynamics of the Schwinger model were simulated using infinite MPS with $g = 0.8$ and $m = 0.1$. The details of the implementation are in Appendix C.1. The system was initialized in the electric vacuum state and evolved in time until $t = 8$. The left panel of Fig. 4.7 shows the expectation of the electric energy and chiral condensate as a function of time for different truncations of the electric field. The truncation error from these numerical simulations is consistent with the analytical estimates in Eq. (4.96) and Eq. (4.98).

4.4 Discussion

In this work, a formalism has been developed for estimating the errors due to electric basis truncations in the simulation of lattice gauge theories. This formalism is based on the presence of Hilbert space fragmentation in the Kogut–Susskind Hamiltonian. This enables the use of time-dependent perturbation theory to describe the dynamics of states with large electric fields. Consequently, the leading contributions to the truncation error fall off factorially with the field truncation. Numerical simulations were performed to demonstrate the performance of these error estimates. The details of the truncation error depend on the change in electric energy from applying the plaquette operator. This suggests lattices with

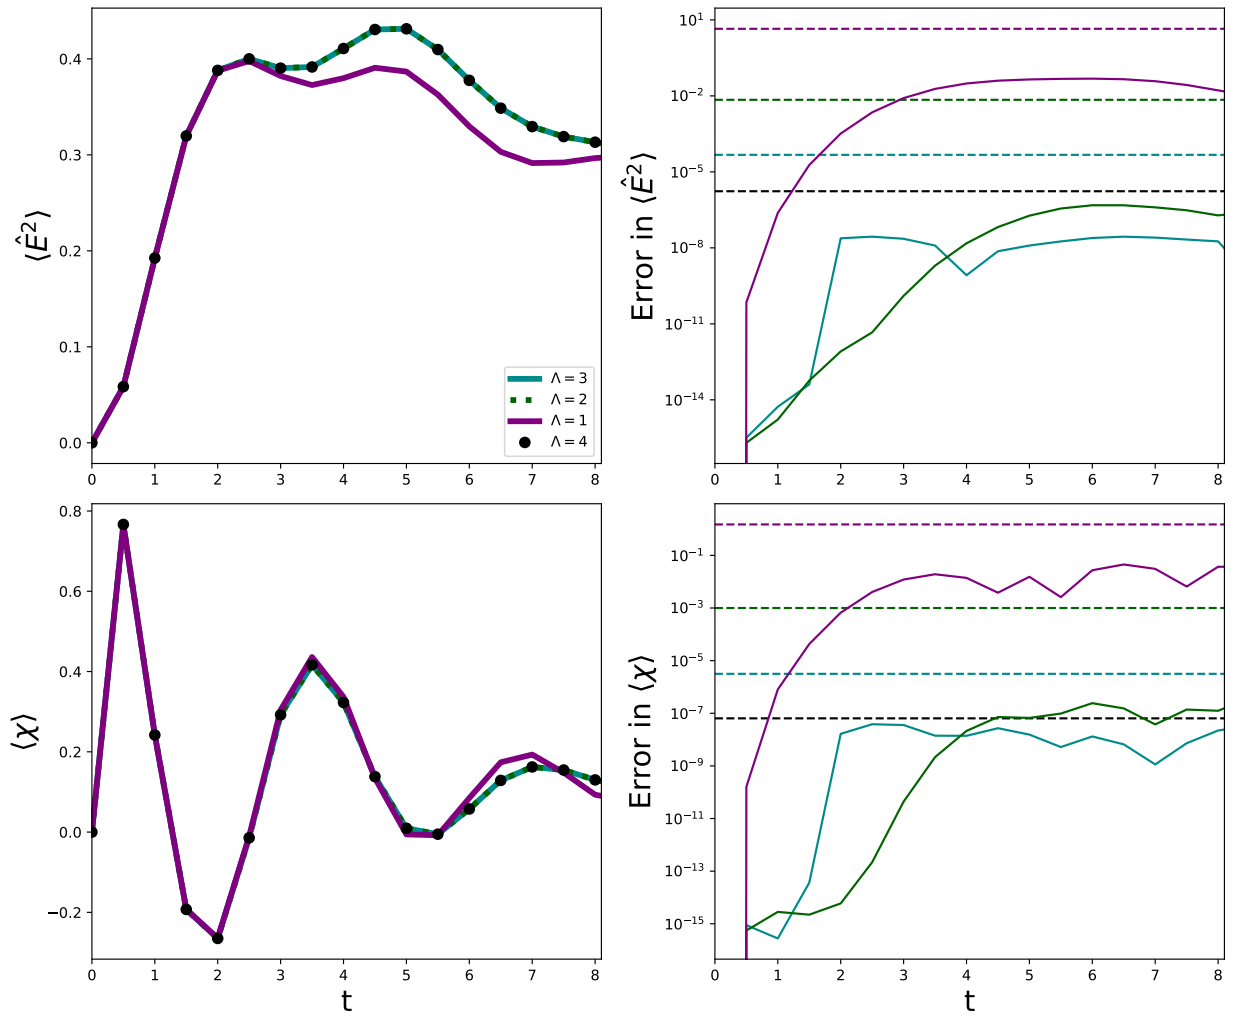

Figure 4.7: Evolution of the electric vacuum in the Schwinger model with $g = 0.8$ and $m = 0.1$. The left panel shows the electric energy per link and the chiral condensate per site as a function of time for different truncations. The right panel shows the difference in the expectation of the observable between the $\Lambda = 4$ and lower truncations. The dashed lines in the right panel show the predicted truncation error. For $\Lambda = 1$, the truncation error was computed using Eq. (4.96) and for larger Λ, Eq. (4.98) was used.

more links per plaquette, such as a honeycomb lattice [245, 249, 261, 266] or a triamond lattice [238, 239], will have a stronger suppression of truncation errors.

The truncation errors estimated in this work are significantly smaller than in previous work. This is due to previous work being based on a loose estimate of the remainder term in the Dyson series expansion, while this work was able to derive the form of the dominant error due to truncation. As a result, for some parameter choices, our results reduce the error estimate by an astronomical factor of 10^{306}. Our formalism does not apply to generic bosonic simulations; however, preliminary simulations in Appendix C.2 suggest that there may be other mechanisms that restrict the growth of boson numbers more generally. While the formalism was only applied to Abelian gauge theories in one dimension in this work, the generalization to non-Abelian gauge groups and higher spatial dimensions is clear. This formalism also applies to scalar field theories with a compact scalar field, such as the $O(3)$ non-linear σ model. The calculations performed in this work show that different observables have different dependencies on the truncation errors. One could potentially use this to construct observables with less sensitivity to truncation errors and potentially extract useful information from simulations with low truncations. This formalism can also be extended to improvement schemes using the similarity renormalization group to reduce the effects of truncation errors [224]. These extensions of this work will enable quantum simulations of lattice QCD with complete theoretical control over the uncertainties due to truncating the gauge fields.

Chapter 5

Obtaining Continuum Physics from Dynamical Simulations of Hamiltonian Lattice Gauge Theories

5.1 Introduction

One crucial ingredient in any LGT calculation is taking the continuum limit, *i.e.*, taking the lattice spacing $a \to 0$. Doing so requires several steps, including tuning the bare parameters in the Hamiltonian as a function of the lattice spacing a, renormalizing the operators being studied, and determining the lattice spacing through a scale setting procedure. Each of these steps in principle requires calculating a physical observable, which implies that the cost of each step could be comparable to the cost of state preparation and time evolution. While approaching the continuum limit has been studied in the context of U(1) quantum link models [262, 263] and strategies for reducing the quantum cost of renormalization through the use of classical Euclidean LGT simulations [113, 269, 270, 271, 272] and perturbation theory [273] have been proposed, the full cost of calculating continuum physics using quantum simulations of LGTs is not fully understood.

One important question to ask is how error from approximate time evolution impacts the continuum limit. This issue was studied in [113] in the context of simulations using product formulas (PFs). Using the Baker–Campbell–Hausdorff formula, it was shown that the effective Hamiltonian one actually time evolves according to with PF-based time evolution is closely related to the transfer matrix in Euclidean LGTs. Through this connection, approximate time evolution can be interpreted as exact evolution on a temporal lattice with spacing δ_t, which is proportional to the Euclidean temporal spacing a_0. The associated error from finite a_t can then be removed by extrapolating $\delta_t \to 0$ (or fixing the anisotropy $\xi \equiv a/\delta_t$ and taking $a \to 0$), following a procedure analogous to the continuum limit $a \to 0$.

While this is a valid approach, it is currently applicable only to PF-based simulations and does not extend to the broader class of post-Trotter simulation algorithms [108, 109,

110, 111, 128, 196, 199, 274, 275, 276, 277, 278]. This implies that it is currently not possible to perform direct comparisons of the full cost of taking the continuum limit when using different time evolution algorithms. This is particularly important given the recent efforts for determining which formulations and algorithms which result in the smallest gate counts [63, 64, 69, 81, 133]. In order to broaden the class of simulation algorithms at our disposal, including those not yet developed, an alternative approach is necessary.

In this work, we develop an alternative approach to controlling the impact of approximate time evolution on the continuum limit. The key insight is that, while approximate time evolution may be interpreted through the lens of renormalization, no additional UV divergences are introduced in the limit of exact time evolution. This suggests that, rather than invoking the full machinery of renormalization, as in [113], the error from approximate time evolution can simply be treated as a systematic uncertainty. As long as this error remains negligible compared to the statistical precision of the simulation, its effect on the continuum limit is negligible and can be ignored. In particular, we will argue that this applies to any unphysical behavior introduced from approximate time evolution, and that no problem analogous to the notorious fermion doubling problem is possible in the limit of exact time evolution. We now state the central question we address in this work:

For a general time evolution algorithm, what is the additional computational cost required to ensure that the time evolution error has a negligible affect on the continuum limit relative to the working statistical precision?

In this work, we provide a general framework that can be used to answer this question for any time evolution algorithm *a priori*, that is, before performing any simulations. We show that, as long as the systematic error incurred for each step in the renormalization procedure is below the working statistical precision, the need to treat the approximate time evolution error using renormalization is no longer necessary, which results in significant simplifications in taking the continuum limit. We coin this proposal Statistically-Bounded Time Evolution (SBTE) Protocol.

To demonstrate our procedure, we apply our formalism to simulations to two common simulation algorithms, namely product formulas and those based on Quantum Signal Processing (QSP). We find that, although the additional cost required to reduce the time evolution error below the working statistical precision can be significant when using Trotter-based methods, it becomes negligible with QSP due to its provably optimal scaling with the approximation error. While this result is noteworthy, a comprehensive cost analysis of the entire renormalization procedure is still necessary to determine which approach yields the lowest overall quantum cost. Looking ahead, our framework can be applied to any algorithm with known rigorous upper bounds on the error, and thus serves as a foundation for comparing the total cost of calculating continuum physics from quantum simulations of LGTs using different time evolution algorithms.

In addition to this general formalism, we propose an alternative approach to renormalization, differing slightly from that used in [113], when employing product-formulas. Our

approach relies on the observation that simulations using PFs are equivalent to simulations with extra irrelevant operators in the Hamiltonian that disappear in the continuum limit. Using this fact, we argue that, rather than tuning the bare parameters in the Hamiltonian as a function of both the lattice spacing a and Trotter step-size δt, one can simply tune the bare parameters assuming $\delta t = 0$ and only introduce systematic errors that scale as $\mathcal{O}(a^p)$, where p is the order of the PF used[1]. This new renormalization trajectory offers significant advantages. First, it simplifies the tuning process by making the bare parameters a function of lattice spacing a alone, removing their dependence on the Trotter step δt. Second, it decouples the tuning from the simulation method, allowing parameters to be determined with algorithms other than product formulas without introducing systematic mismatch errors. Finally, we show that if one does use Euclidean simulations to capture the $\delta t, a$ dependence of the Hamiltonian parameters, the limitations are more severe than previously understood [113].

The rest of this work is organized as follows. In Sec. 5.2, we provide a high-level review of the continuum limit and explain the main concept of our strategy. Then, in Sec. 5.3, we provide a more detailed review, aimed at non-experts, of the steps necessary to renormalize the theory and properly take the continuum limit assuming exact time evolution. In Sec. 5.4, we discuss the procedure in [113] for taking the simultaneous $a, \delta_t \to 0$ limit when using PFs for simulation, and then describe our alternative renormalization approach. We then show that, if one drives the time evolution error below the working statistical precision, it can be ignored and renormalization simplifies to the exact time evolution procedure. In Sec. 5.5.1, we show how to bound the additional cost of driving the time evolution error below the shot noise, and then apply our methods to simulation via product formulas and Quantum Signal Processing (QSP). We will show that, while this procedure requires non-negligible overhead when using PFs, the additional cost is entirely negligible when using QSP. We present our conclusions in Sec. 5.6.

The work presented in this chapter is based on [123].

5.2 Obtaining the continuum limit

Calculations in lattice field theories provide a numerical technique to calculate observables defined in a continuum quantum field theory. The lattice field theory can be viewed as an effective theory of the underlying continuum field theory one aims to simulate, where the lattice spacing and lattice size provide an ultraviolet (UV) and infrared cutoff (IR), respectively. Numerical results obtained from a lattice simulation therefore differ from the continuum results by systematic effects arising from the non-zero lattice spacing and finite lattice size. In addition, one needs to control systematic uncertainties from any other ap-

[1]While this renormalization trajectory was employed in [113], it was described as an approximation to one with fixed physical anisotropy. In this work, we determine the size of the errors, and develop a unique renormalization trajectory.

proximation that one might make in the lattice calculation, as well as statistical uncertainties made in the numerical lattice calculation.

The most straightforward way to deal with the systematic uncertainties is to make them subdominant to the statistical uncertainties. In particular, if a particular systematic effect can be bounded to be significantly below the statistical uncertainties, its effect can be neglected in the lattice calculation. For example, since the effects from finite lattice spacing are known to vanish polynomially with the size of the lattice spacing, discretization errors in a calculation performed at sufficiently small lattice spacing a are guaranteed to be below the statistical uncertainties.

In reality, many systematic effects cannot be treated this way. First, while one often knows the scaling of the systematic effects with certain parameters, the actual size of them can not be bounded reliably. Second, the computational cost for performing calculations for which certain systematic effects are negligible are often prohibitively expensive. For this reason, many systematic effects are dealt with by peforming extrapolations. For example, one typically performs calculations at various values of the lattice spacing a, all of which have non-neglible systematics. However, one can extrapolate these results to $a = 0$ using the known polynomial scaling in a.

An important property of quantum field theories is that they are UV divergent. This implies that physics from arbitrarily short distance effects contribute to a given observable, such that the integration over all possible energies contributing is divergent. This does not cause any problem with the predictability of such quantum field theories, since these divergences can be absorbed into the parameters of the theory, such as masses and coupling constants, rendering physical observables finite. As will be discussed in more detail in Sec. 5.3, the lattice spacing regulates the divergences at short distances, causing the parameters of the effective lattice theory to depend on the lattice spacing used. In other words, coupling constants and masses used in the lattice simulation depend on the lattice cutoff used

$$g \to g(a)\,, \quad m \to m(a)\,, \quad \dots\,, \tag{5.1}$$

and the value of these bare parameters need to be adjusted to yield physically meaningful results.

A systematic effect that arises in Hamiltonian lattice gauge theories is due to approximations that typically need to be made when obtaining the time-evolution operator from the Hamiltonian H of the lattice gauge theory. This is due to the fact that it is rarely possible to exponentiate the Hamiltonian exactly to obtain the unitary time evolution operator $\exp(-iHt)$. The most common approximation scheme used is based on product formulas (as discussed in Chapter 2), with the simplest approximation to the time evolution operator of a Hamiltonian

$$H = \sum_i H_i\,, \tag{5.2}$$

given by

$$e^{-iHt} = \left[e^{-iH\delta t}\right]^{N_{\rm PF}} = \left[\prod_i e^{-iH_i\delta t}\right]^{N_{\rm PF}} + \mathcal{O}\left(\delta t\right)^2 , \tag{5.3}$$

with

$$\delta t \equiv t/N_{\rm PF}, \text{ where } N_{\rm PF} \in \mathbb{N} . \tag{5.4}$$

The approximation above is a first-order ($p = 1$) product formula, and higher order product formulas with larger values of p are possible, with the deviation from the exact time-evolution operator scaling as $(\delta t)^p$ for fixed t.

Choosing a non-zero value of δt can affect the UV behavior of the theory meaning that in general the parameters of the lattice theory depend on this parameter as well

$$g \to g(a, \delta t) , \quad m \to m(a, \delta t) , \quad \ldots . \tag{5.5}$$

This implies that additional calculations might need to be performed to determine the correct parameter values, adding to the computational cost. In the limit $\delta t \to 0$ the product formula approximation of the time evolution operator becomes exact, such that one has

$$g(a, 0) = g(a) , \quad m(a, 0) = m(a) , \quad \ldots , \tag{5.6}$$

In other words, taking the limit of exact time evolution does not introduce any UV divergences. We will discuss the implications of this in Sec. 5.5.

One way to deal with this systematic uncertainty is therefore to choose a small enough value of δt (large enough value of $N_{\rm PF}$) to ensure this error is below the systematic uncertainties and therefore negligible. One issue with this approach is that, while one knows the scaling of the error with δt, the absolute size depends on the details of the Hamiltonian being simulated and is difficult to tightly bound *a priori*. Furthermore, the computational cost required to achieve an approximation to the exact time evolution operator with error ϵ grows as $\epsilon^{-1/p} \exp(p)$, where $\exp(p)$ comes from the fact that the implementation of a p'th order Suzuki-Trotter product formula requires exponential resources in p. This makes it difficult to reach a sufficiently small value of ϵ.

Given these complications, an alternative procedure to deal with the systematic effect from finite values of δt was proposed in [113]. In their proposal one keeps the ratio of δt and the lattice spacing fixed and takes the simultaneous limit $a \to 0$ and $\delta t \to 0$. This is similar to a Euclidean lattice calculation, where both space and Euclidean time are discretized in a space-time lattice, with a fixed ratio of the Euclidean time-like and space-like lattice spacing a_0 and a. While the additional dependence on δt can change the value of the bare parameters that have to be chosen, the authors of [113] went on to show that one can determine the values of these parameters from classical calculations on Euclidean lattices through analytic continuation. This however introduces additional systematic uncertainties that need to be

quantified. We discuss this approach in more detail in Sec. 5.4.3, and also introduce several variants to the approach proposed in [113].

It is important to realize, however, that approximations to the time evolution operator beyond product formulas exist. A notable example are methods that use the general idea of Quantum Signal Processing (QSP), which allows one to construct a wide class of polynomials of operators. These techniques give rise to an approximation to the time-evolution operator, with the computational complexity scaling only as $\log \frac{1}{\epsilon}$[108, 109, 110, 116].

A worry of the community has been[2] whether or not deviations of QSP techniques from the exact time evolution operator can be captured by an effective Hamiltonian picture, as is the case when using PFs. One possible issue is that a polynomial approximation to the time-evolution operator is generally not unitary (this issue is handled in QSP approaches through ancilla qubits and post-selection protocols), and it is therefore not obvious how to construct a Hermitian effective Hamiltonian describing this non-unitary evolution. In the absence of an effective theory description, it might be difficult to understand the dependence of the bare parameters of the lattice theory on the approximation used. One might therefore worry that a lattice calculation obtained using QSP techniques can not easily be related to the continuum physics one aims to extract.

However, as we argue in this work, if the error from the approximation to time evolution is significantly below all other systematic and statistical uncertainties in the problem, it can not affect the results of the simulation in a meaningful way. To the accuracy of the calculation, the bare parameters will therefore only have a dependence on the spatial lattice spacing a. This implies that there is no in-principle impediment to using QSP techniques, or any other time evolution algorithms, for approximating the time evolution of a quantum field theory. Whether such an approach is favorable compared to the more standard techniques based on product formulas depends on details of the approaches chosen, and need to be investigated carefully.

In the following sections we will explain the two basic topics raised in this section. In Sec. 5.3 we discuss the renormalization of the parameters of the theory and how their numerical values depend on the lattice spacing and the approximation to the time-evolution operator. In Sec. 5.5.1 we study the scaling of the computational cost with the deviation from the exact time-evolution operator for both product formula and QSP based approaches.

5.3 Renormalization in Hamiltonian Lattice Gauge Theory for Exact Time Evolution

In this section, we start with a high-level overview of renormalization in quantum field theories aimed at non-experts. We then more formally discuss how to renormalize LGT calculations and take the continuum limit where time evolution is performed exactly.

[2]We thank Henry Lamm for many useful discussions on this point.

The Kogut-Susskind (KS) Hamiltonian of a pure gauge theory is given by

$$\begin{aligned} H_{\mathrm{KS}} &= \frac{1}{a}\left[\frac{g_t^2}{2}\sum_{\ell}\left(\hat{E}_\ell\right)^2 - \frac{1}{g_s^2}\sum_{p}\mathrm{Re}\,\mathrm{Tr}\,\hat{P}_p\right] \\ &\equiv H_E + H_B\,, \end{aligned} \tag{5.7}$$

were hatted operators denote dimensionless operators, rescaled by appropriate powers of the only length-scale in the problem, namely the lattice spacing a. Due to the breaking of Lorentz invariance the dimensionless coupling constants of the electric (g_t) and magnetic (g_s) term are different, and the bare speed of light is defined as[3]

$$c \equiv \frac{g_t}{g_s}\,. \tag{5.8}$$

One can therefore write

$$\begin{aligned} H_{\mathrm{KS}} &= \frac{c}{a}\left[\frac{g^2}{2}\sum_{\ell}\left(\hat{E}_\ell\right)^2 - \frac{1}{g^2}\sum_{p}\mathrm{Re}\,\mathrm{Tr}\,\hat{P}_p\right] \\ &\equiv \frac{c}{a}\left[\hat{H}_E + \hat{H}_B\right] \\ &\equiv \frac{c}{a}\hat{H}_{\mathrm{KS}}\,, \end{aligned} \tag{5.9}$$

where we have defined

$$g = \sqrt{g_s g_t}\,. \tag{5.10}$$

Using this parameterization introduces a second dimensionful parameter into the theory, namely the time-scale

$$a_t \equiv \frac{a}{c}\,. \tag{5.11}$$

Interactions with fermions give rise to additional terms that involve the masses m_i of the fermions as well as a so-called hopping term describing the interactions between gauge bosons and fermions with coefficient κ.[4]

It is important to remember that the KS Hamiltonian is only equal to the full continuum Hamiltonian up to corrections that vanish as the lattice spacing approaches zero. The pure

[3]Note that, while the incorrect definition of the speed of light g_s/g_t was used in [113], the methods presented are still valid.

[4]To the best of our knowledge, the need for the additional bare parameter κ multiplying the fermionic kinetic term in the Hamiltonian formulation has not been discussed in the literature. The need for this parameter is discussed in Appendix D.1.

gauge Hamiltonian as written in Eq. (5.9) is correct up to corrections scaling as a^2, such that the continuum Yang-Mills (YM) Hamiltonian can be written as

$$\hat{H}_{\mathrm{YM}} = \hat{H}_{\mathrm{KS}} + \sum_i c_i \hat{O}_i^{(3)} + \dots \,, \tag{5.12}$$

where $\hat{O}_i^{(3)}$ denotes higher dimensional operators whose contributions to long-distance physics at distances L scale as a^3/L^3. Since the rescaled Hamiltonian $\hat{H}_{\mathrm{KS}}$ vanishes as a/L as just discussed, any calculation performed with the KS Hamiltonian will only reproduce the desired continuum physics up to relative corrections suppressed by a^2/L^2.

As discussed in more detail in the next section, any dimenionful parameter in the Hamiltonian will also be replaced by the corresponding dimensionless parameter by multiplying by the appropriate factors of the lattice spacing and the speed of light. The ruler with which temporal- and spatial-type quantities are measured with respect to are a/c and a, respectively[5]. Examples are fermion masses in the Hamiltonian, the time one evolves the system by, momenta, or the physical distance between correlators. Their dimenionless equivalents are given by

$$\hat{m} = a_t\, m\,, \qquad \hat{t} = \frac{1}{a_t} t\,, \qquad \hat{p} = ap, \qquad \hat{x} = \frac{1}{a} x\,. \tag{5.13}$$

In the limit of vanishing lattice spacing $a \to 0$, the KS Hamiltonian has to reproduce the spectrum of the continuum theory. This implies that the eigenvalues of the Hamiltonian H_{KS} have to be finite in that limit, which further implies that the eigenvalues of the rescaled Hamiltonian $\hat{H}_{\mathrm{KS}}$ have to vanish as $a \to 0$.

Calculating physical observables using quantum field theory requires integrating over all possible energy and momentum values of intermediate states that contribute to a process. These integrals, however, generally diverge as the integration range is taken to infinity. One intuitive explanation for this is that such calculations assume that the theory is valid at arbitrarily high energy scales; because current theories are incomplete (in the sense that they are not the final "theory of everything"), this is simply not true. This, however, does not give rise to an ill-defined theory or a loss of predictive power. The physics at short distances is captured using a process known as *renormalization*, in which the parameters of the effective theory (coupling strengths of interactions, masses of particles, etc.) are adjusted to reproduce the physics at long distances one aims to compute.

In general, renormalization requires two steps. The first is to *regulate* the theory by manually imposing some cutoff Λ on the energies or momenta included in the theory (in a lattice theory $\Lambda \sim 1/a$). While this renders calculations finite, it naively makes physical observables depend on the cutoff Λ. The key insight of renormalization is that this unphysical behavior can be removed by choosing the bare parameters in the theory to depend on the

[5]The fact that spatial-type quantities are measured with respect to the lattice spacing is reflected in the fact that the lattice momentum operator would have an overall prefactor of $1/a$ instead of c/a in the case of the Hamiltonian.

cutoff in precisely the correct way such that some physical observable is reproduced correctly and is independent of the cutoff. In particular, the bare coupling, speed of light and masses become Λ dependent $g \to g(\Lambda)$, $c \to c(\Lambda)$, $m_i \to m_i(\Lambda)$ and $\kappa \to \kappa(\Lambda)$. While this leads to the bare parameters diverging (or going to zero) as $\Lambda \to \infty$, this is not a problem as they are not physical observables in the theory.

While the dependence of the parameters on Λ can in some cases be computed perturbatively, in general this is not possible, and one has to determine this dependence by demanding that the effective theory reproduces a limited set of physical observables. After adjusting the bare parameters to absorb the divergences as $\Lambda \to \infty$ and to reproduce a small set of observables, the theory has been renormalized and all other predictions will be finite and equal to the physical values[6] in the $\Lambda \to \infty$ limit.

Since computers (both classical and quantum) can only deal with dimensionless numbers, one numerically simulates a rescaled Hamiltonian $\hat{H}_{\mathrm{KS}}$ introduced in Eq. (5.9), and from now on, all parameters of the lattice Hamiltonian are assumed to be dimensionless, and the parameter $a_t = a/c$ only appears as an overall constant. This has two consequences that are intimately connected and which we will discuss more later. First, lattice calculations can only compute dimensionless quantities, and in order to obtain dimensionful quantities observed in nature, such as masses, one needs to multiply by the appropriate powers of a in the end. Second, because one only chooses values of the rescaled parameters, the lattice spacing is not an explicit input in a simulation. Instead, its value is only obtained after the fact by comparing calculations to dimensionful quantities observed in nature. See Appendix D.1 for a more detailed discussion of this property of lattice field theories.

We distinguish two types of parameters in our theory. The first set are dimensionless, possibly rescaled parameters that determine the relative size of terms in the Hamiltonian. For the Kogut-Susskind Hamiltonian the only such parameter is the gauge coupling g, but when including fermions, rescaled masses are present as well. We denote these parameters collectively as

$$p_i \equiv \{g, \hat{m}_i, \hat{\kappa}\}\,. \tag{5.14}$$

where $\hat{\kappa} = \kappa/c$ as shown in Appendix D.1. The second set of parameters are those that set the scales in the problem, one scale for time a_t and one for length a, or equivalently a and the speed of light c. A common choice is the lattice spacing a and the speed of light c.

As already mentioned, the presence of a lattice spacing introduces a cutoff $\Lambda_a \sim 1/a$, since wavelength shorter than the lattice spacing a, and time scales shorter than a/c can not be resolved. From the previous discussion it therefore follows that renormalization demands that the bare parameters in the Hamiltonian depend on the lattice spacing. This gives a lattice gauge theory Hamiltonian $H(a)$ with all parameters being functions of the lattice spacing

$$p_i \to p_i(a)\,. \tag{5.15}$$

[6] If using perturbation theory, the predictions will be equal up to higher order corrections.

As we will discuss, the speed of light will also depend on the lattice spacing a.

At first sight, one might be worried that this makes it impossible to use a lattice field theory to predict physical observables. After all, it was just discussed that the value of the lattice spacing a can only be extracted at the end of a lattice simulation, and that the functional dependence $p_i(a)$ is non-perturbative and therefore not known analytically. Given this, how does one know what values to choose for the bare parameters $p_i(a)$ used in the calculation? The answer is that one needs to tune these parameters to reproduce experimentally accessible quantities, such as masses in the spectrum of the full theory. One typically views this as a two step process, as we will now discuss.

In a first step, one determines the dimensionless parameters of the Hamiltonian. This is accomplished by finding parameters for which the numerical calculation reproduce experimentally known dimensionless quantities. A common choice for such dimensionless ratios are ratios of masses

$$\frac{\frac{a}{c}M_1(p_i)}{\frac{a}{c}M_2(p_i)} = \frac{M_1^{(\text{phys})}}{M_2^{(\text{phys})}}\,. \tag{5.16}$$

Note that for a pure gauge Hamiltonian, this step is not necessary as all choices of the gauge coupling are valid, with different choices corresponding to different lattice spacings as determined in the next step. When fermions are included, however, the bare fermion masses have to be tuned by the above procedure. The parameter values are typically determined by iterating over a range of parameter values in the numerical calculation and using some fitting procedure for the final tuning.

Having reproduced dimensionless ratios, one determines the scales a_t and a by comparing the results obtained in the previous step to observed dimensionful quantities. For example, the scale a_t is can be obtained by taking the ratio of a dimensionless mass computed numerically to its observed value

$$a_t = \frac{\hat{M}(p_i)}{M^{(\text{phys})}}\,. \tag{5.17}$$

The scale a, or equivalently the speed of light, is obtained by demanding that observables depending on spatial distances are reproduced correctly. For example, one can demand that the energy of a massive particle with given lattice momentum reproduces the energy predicted from the relativistic dispersion relation.

Thus, for each parameter set p_i one finds corresponding values of a and a_t, or equivalently a and c. One can then interpret these parameter sets as labeled by the lattice spacing found and write

$$\{p_i, c, a\} \to \{p_i(a), c(a)\}\,. \tag{5.18}$$

One therefore obtains a line in parameter space, where each point corresponds to a set of parameter values and lattice spacing. This line is often called the renormalization trajectory. Given that this trajectory corresponds to exact time evolution, we denote it by

$$\text{traj}_0 \,:\, \{p_i(a), c(a)\}\,. \tag{5.19}$$

Due to the divergences in the underlying continuum theory, the parameter values corresponding to vanishing lattice spacing are typically either infinite or vanishing; because Lorentz invariance is restored in the continuum limit, the bare speed of light and the parameter $\kappa(a)$ obeys $\lim_{a\to 0} c(a) = \lim_{a\to 0} \kappa(a) = 1$. Note that different renormalization trajectories associated with different observables used in scale setting are possible, with each giving the same result in the $a \to 0$ limit. The dependence of the dimensionless coupling g on the lattice spacing is logarithmic.

Given this renormalization trajectory, one can now compute new observables of interest. After renormalization, a lattice calculation of low energy quantities differs from the continuum result by powers of the lattice spacing; in modern calculations the leading errors are typically $\mathcal{O}(a^2)$. One can not perform a calculation directly at $a = 0$, since as discussed above the parameter values at that point are either infinite or vanishing, and would also require an infinite number of lattice sites to maintain a constant physical volume. One therefore calculates the desired physical observables for various parameter choices along the renormalization trajectory, resulting in a value that depends on the corresponding lattice spacing. To obtain the continuum result, one performs a final extrapolation to $a = 0$.

Given that a major advantage of Hamiltonian lattice gauge theory simulated on quantum computers over traditional lattice gauge theory is the ability to calculate time-dependent observables, we now consider the calculation of expectation values of a time-dependent operator $\langle O(t)\rangle$[7] through the calculation of the expectation value of an operator defined on the lattice $\langle O(t,a)\rangle$. The dependence on a captures the fact that calculation is performed with the parameter values $\{p_i(a)\}$ of the renormalization trajectory corresponding to the value a. As before, the lattice only allows one to compute dimensionless quantities, such that for an operator of mass dimension γ one computes the expectation value $\langle \hat{O}(t,a)\rangle = \langle a^\gamma O(t,a)\rangle$.

Unless protected by a symmetry in the theory, the operator $\hat{O}(t,a)$ also needs to be renormalized. This is done by introducing renormalization factors $Z(a)$. For simplicity we assume multiplicative renormalization, but the arguments presented in this manuscript generalize to additive renormalization factors and operator mixing. We determine $Z(a)$ by demanding some other chosen matrix element of $Z(a)\hat{O}(t,a)$ is equal the associated physical value, *i.e.*,

$$Z(a) \equiv \frac{\langle \phi_1 | \, O(t,a) \, | \phi_2 \rangle}{[\langle \phi_1 | \, O(t) \, | \phi_2 \rangle]^{(\text{phys})}} \,, \tag{5.20}$$

for some states $|\phi_1\rangle$ and $|\phi_2\rangle$. Once $Z(a)$ has been determined, the physical observable is given in the limit

$$\langle O(t)\rangle = \lim_{a\to 0} Z(a) a^{-\gamma} \langle \hat{O}(t,a)\rangle \,. \tag{5.21}$$

Note that $Z(a)$ is time-independent in this case, as the renormalized operator is given by $Z(a)O(t,a) = e^{iH(a)t}\left(Z(a)O(0,a)\right)e^{-iH(a)t}$.

We now summarize the discussion given so far, by providing the steps required to take the continuum limit of a time-dependent observable. They are

[7]Note that our arguments also apply to the scenario where one is computing more general matrix elements of an operator between two different states $\langle \psi_f | \, O(t) \, | \psi_i \rangle$.

1. Determine the renormalization trajectory and therefore $g(a)$, $\hat{m}_i(a)$, $\hat{\kappa}(a)$, and $c(a)$
    a) Calculate a set of known dimensionless observables for a given parameter set p_i
    b) Adjust the parameters to until calculations match the known dimensionless values
    c) Determine a and a_t for this parameter set by comparing against a dimensionful observables (recall that converting between a and a_t is done by a trivial rescaling with $c(a)$)
    d) Repeat steps (1a)-(1c) until the renormalization trajectory is obtained
2. Compute the desired physical observable
    a) Pick a value of a by choosing a set of parameters $p_i(a)$ and $c(a)$
    b) If calculating a time-dependent observable, determine the appropriate value of $\hat{t} = c(a)t/a$
    c) Determine $Z(a)$ using the previously determined parameter values
    d) Calculate $\langle \hat{O}(\hat{t}, a) \rangle$ and rescale by powers of a to convert to physical units
    e) Repeat steps (2a)-(2d) for several values of a and extrapolate $Z(a)\langle O(\hat{t}, a) \rangle$ to zero lattice spacing

From the discussion in this section, it should be clear that while all parameters and observables in the above steps have explicit dependence on a, this dependence on a only labels the parameter values used in the given calculation.

We conclude this section by noting that, for a given lattice geometry and Hamiltonian, one only needs to tune the bare parameters and scale set once; once the renormalization trajectory has been determined, the relative cost of these steps can be reduced by reusing it for several observables.

5.4 Renormalization for approximate time evolution via product formulas

In this section, we discuss how the procedure of renormalization changes when time evolution is performed approximately using product formulas. We first review the previous approach developed in [113]. From there, we introduce new, simplified, renormalization trajectories. We conclude by discussing the limitations of using Euclidean LGT simulations to aid with renormalization of quantum simulations done via PFs.

5.4.1 Review of existing approach

In this section we review the methods developed in [113] for taking the continuum limit when using PFs.

Product formulas are based on the Baker-Campbell Haussdorff (BCH) formula

$$e^{-iH_1\delta_t}e^{-iH_2\delta_t} = e^{-i(H_1+H_2)\delta_t - \frac{\delta_t^2}{2}[H_1,H_2]+\dots}\,, \tag{5.22}$$

and approximate the time evolution operator by products of individual exponentials. For example, writing

$$H = H_1 + H_2 \tag{5.23}$$

the first and second order product formulas will approximate the time evolution operator $U = \exp[-iHt]$ as

$$\begin{aligned} S_1(t) &= \big[S_1(\delta_t)\big]^{N_{\rm PF}} = \Big[e^{-i\delta_t H_1}e^{-i\delta_t H_2}\Big]^{N_{\rm PF}} \\ S_2(t) &= \big[S_2(\delta_t)\big]^{N_{\rm PF}} = \Big[e^{-i\delta_t H_1/2}e^{-i\delta_t H_2}e^{-i\delta_t H_1/2}\Big]^{N_{\rm PF}}\,, \end{aligned} \tag{5.24}$$

where δ_t is given by Eq. (5.4). Note that time evolution using product formulas requires two independent parameters $N_{\rm PF}$ and δ_t, which need to combine to produce the desired physical time of interest

$$N_{\rm PF}\delta_t = t\,. \tag{5.25}$$

One way to use product formulas is to choose a value of $N_{\rm PF}$ that is large enough such that the systematic uncertainties to time evolution becomes negligible. This will be discussed in detail and generalized to other approximations to time evolution in Sec. 5.5, and requires to find a value of $N_{\rm PF}$ for which the the resulting systematic uncertainties are guaranteed to be below the statistical uncertainties made in the numerical calculation. In the rest of this section, we discuss how to handle the systematic uncertainties through renormalization [113].

Using the BCH formula, Eq. (5.24), time evolution via product formulas correspond to time evolution for a time step δ_t according to the effective Hamiltonians

$$S_k(\delta_t) = e^{-i\delta_t H_{\rm eff}^{(k)}}\,, \tag{5.26}$$

with

$$\begin{aligned} H_{\rm eff}^{(1)}(\delta_t) &= H_1 + H_2 - i\frac{\delta_t}{2}\,[H_1,H_2] + \dots \\ H_{\rm eff}^{(2)}(\delta_t) &= H_1 + H_2 + \frac{\delta_t^2}{24}\Big(\,\big[H_1,[H_1,H_2]\big] + 2\,\big[H_2,[H_1,H_2]\big]\,\Big) + \dots\,. \end{aligned} \tag{5.27}$$

We call the terms involving commutators of H_i BCH terms, since they arise from the higher order terms in the BCH formula. Rescaling all parameters by the appropriate factors of the lattice spacing a, including the appropriate factors of the speed of light

$$\hat{H}_i = \frac{a}{c}H_i\,, \qquad \hat{\delta}_t = \frac{\hat{t}}{N_{\rm PF}}\,, \tag{5.28}$$

where $\hat{t}$ is defined in Eq. (5.13) one therefore finds

$$\hat{H}^{(1)}_{\text{eff}}(\delta_t) = \hat{H}_1 + \hat{H}_2 - i\frac{\hat{\delta}_t}{2}\left[\hat{H}_1, \hat{H}_2\right] + \dots \tag{5.29}$$

$$\hat{H}^{(2)}_{\text{eff}}(\delta_t) = \hat{H}_1 + \hat{H}_2 + \frac{\hat{\delta}_t^2}{24}\Big(\left[\hat{H}_1, [\hat{H}_1, \hat{H}_2]\right] + 2\left[\hat{H}_2, [\hat{H}_1, \hat{H}_2]\right] \Big) + \dots . \tag{5.30}$$

For a k^{th} order product formula the leading BCH terms contain k powers of $\hat{\delta}_t$ and commutators involving $k+1$ factors of $\hat{H}_i$. We will discuss the implications of this in a later section.

Using the effective Hamiltonian picture, it is clear that using PFs changes the low energy physics of the theory being simulated, and so the renormalization procedure must be modified accordingly. One method to account for this [113] is to choose a several values for the parameter $\hat{\delta}_t$ and then perform the same steps as in Sec. 5.3 to determine the parameters p_i, as well as a and c (we discuss how to choose $\hat{\delta}_t$ in the next section). This results in a renormalization trajectory that depends on the choice of parameter $\hat{\delta}_t$, which we will denote as

$$\text{traj}_1 : \{p_i(a, \hat{\delta}_t), c(a, \hat{\delta}_t)\}\,. \tag{5.31}$$

By repeating this step for several pairs of values $(a, \hat{\delta}_t)$, one can safely take the $\lim_a \lim_{\delta t} \to 0$ limit.

While this is a valid approach, one downside is that it requires performing a double extrapolation in both a and δ_t, leading to a significant increase in the quantum cost. Fortunately, this limitation can be overcome by instead working at a fixed physical anisotropy $\xi = a/\delta_t$ and extrapolating $a \to 0$ [113]; we denote this renormalization trajectory by

$$\text{traj}_2 : \{p_i(a, \xi), c(a, \xi)\}\big|_{\xi \text{ fixed}}\,. \tag{5.32}$$

In this way, by adding ξ to the set of quantities that need to be tuned, taking $a \to 0$ implies $\delta_t \to 0$ as well. Note that fixing ξ requires choosing the number of Trotter steps $N_{\text{PF}} = \xi t/a$ to maintain a fixed physical time t.

As discussed in [113], it is possible to simplify this procedure further by noting that fixing ξ is equivalent to fixing $c/\hat{\delta}_t$. Up to corrections due to the renormalization of $c(a)$ (which are expected to be mild [113]), one can instead fix the bare step-size $\hat{\delta}_t$, removing the need for tuning the renormalized anisotropy ξ.

Despite this improvement, a key limitation remains: all parameter tuning must use observables obtained with the same effective Hamiltonian, $\hat{H}_{\text{eff}}(\delta_t)$, that governs the time evolution. Using observables associated with a different Hamiltonian, even with a different δ_t, will lead to systematic errors in the renormalization trajectory. In practice, the only approach to avoid such mismatch errors constrains is to use same PF for both procedures, preventing the use of other, potentially more computationally efficient, methods. In the next section, we propose an alternative renormalization trajectory that simultaneously allows using the exact lattice Hamiltonian $\hat{H}(a)$ to tune parameters for Trotter simulations, while maintaining the beneficial property of only extrapolating $a \to 0$ at fixed $\hat{\delta}_t$.

5.4.2 Improved renormalization trajectories

In this section, we show that the BCH terms in $H_{\text{eff}}^{(k)}(\delta_t)$ are simply irrelevant operators, and use this fact to develop a new set of simple and flexible renormalization trajectories.

As discussed briefly in the previous section, for a k^{th} order product-formula, the leading BCH terms contain k powers of $\hat{\delta}_t$ and commutators involving $k+1$ factors of $\hat{H}_i$. Writing the effective Hamiltonian in terms of the bare step-size $\hat{\delta}_t$ and the dimensionful Hamiltonian terms, the effective Hamiltonian takes the form

$$\begin{aligned} H_{\text{eff}}^{(k)}(\delta_t) &= H_1 + H_2 + \hat{\delta}_t^{k+1} \left(\frac{a}{c}\right)^k \mathcal{C}_k(H_1, H_2) + \dots \\ &= H_1 + H_2 + \mathcal{O}\left(\hat{\delta}_t^{k+1} \left(\frac{a}{c}\right)^k\right), \end{aligned} \tag{5.33}$$

where $\mathcal{C}_k(H_1, H_2)$ is a function that contains nested commutators of H_1 and H_2 [106]. Written in this form, one sees that BCH terms in the effective Hamiltonian are irrelevant operators with mass dimension k.

One immediate consequence of this observation is that (for fixed $\hat{\delta}_t$), if one uses a lattice Hamiltonian with discretization errors scaling as $\mathcal{O}(a^\gamma)$, one must use a PF with order $k \geq \gamma$ to avoid introducing larger lattice spacing errors. For example, using a first order PF could introduce $\mathcal{O}(a)$ errors. On the other hand, if one makes sure that $k \geq \gamma$, this immediately implies that the BCH operators contribute at the same (or higher) order as the other higher dimensional operators that are already neglected in the lattice Hamiltonian $H(a)$.

This implies that including or not including the BCH contributions only changes numerical simulations at a given lattice spacing only up to $\mathcal{O}(a^k)$ corrections, at which level other effects not included contribute as well.

These considerations lead directly to a simplified renormalization trajectory. In particular, after fixing the bare Trotter step-size $\hat{\delta}_t$, it is possible to use the renormalization trajectory associated with exact time evolution (see Eq. (5.19)) in simulations with $H_{\text{eff}}^{(k)}(\delta_t)$ while only incurring $\mathcal{O}((a/c)^k)$ errors. We formally state this renormalization trajectory as follows:

$$\text{traj}_3 : \{p_i(a), c(a)\}\big|_{\hat{\delta}_t \text{ fixed}}. \tag{5.34}$$

Similar to the trajectories in Sec. 5.4.1, fixing the physical evolution time t requires choosing the number of Trotter steps to scale as $N_{\text{PF}} = (ct/a)\hat{\delta}_t$.

It is important to note that because the observables calculated using $H_{\text{eff}}^{(k)}(\delta_t)$ differ from $H(a)$ by $\mathcal{O}(\hat{\delta}_t^{k+1}(a/c)^k)$ errors, the bare parameters for these trajectories also differ by $\mathcal{O}(\hat{\delta}_t^{k+1}(a/c)^k)$ errors. For large values of a or $\hat{\delta}t$, this mismatch could lead to enhanced discretization errors from any non-linear behavior in $p(a)$, e.g. the non-linear dependence of the lattice spacing on the gauge coupling $g(a)$. However, for small lattice spacing a and bare steps-size $\hat{\delta}_t$, the mismatch errors in the bare parameters will only lead to $\mathcal{O}(\hat{\delta}_t^{k+1}(a/c)^k)$ errors.

In the regime of small a and $\hat{\delta}_t$, this trajectory has several advantages relative to those discussed in the previous section. First, because $\hat{\delta}_t$ is a bare parameter, it can be trivially set to a fixed value. Second, it only requires extrapolating the lattice spacing to zero. Lastly, because we are using the bare parameters associated with the exact Hamiltonian $H(a)$, it is possible to calculate quantities required to tune the bare parameters using algorithms beyond PFs, offering flexibility to consider the most cost-effective approach. It is possible to simplify this procedure even further by neglecting the renormalization of speed of light and κ altogether, since their differences from their continuum value $c = \kappa = 1$ are themselves of $\mathcal{O}(a^2)$.

To make more concrete the simplifications this renormalization trajectory leads to, consider simulations of the Kogut-Susskind Hamiltonian. Using a $k = 4$ PF with fixed $\hat{\delta}_t$, the Trotter errors appear at $\mathcal{O}(a^4)$ and go to zero as $a \to 0$. Setting $c = \kappa = 1$ can be done at the cost of $\mathcal{O}(a^2)$ errors, which are the same size as errors in the KS Hamiltonian already. In this way, one only has to tune the bare quark masses and determine the lattice spacing through scale setting, which can be done using any algorithm that simulates the KS Hamiltonian.

Looking forward, this renormalization trajectory leads naturally to a possible future Hamiltonian LGT program in which parameter tuning can be reused among many different simulations for many different algorithms. We conclude this section by noting that, while this renormalization trajectory is conceptually and computationally simple, understanding the absolute size of the $\mathcal{O}(\hat{\delta}_t^{k+1}(a/c)^k)$ errors, and how they propagate relative to those in the existing trajectories discussed in Sec. 5.4.1 will require dedicated studies, which we leave for future work.

5.4.3 Limitations of Euclidean simulations to perform renormalization of real-time simulations

A different approach to renormalization is to use Euclidean simulations performed on classical computers to obtain the renormalization trajectory. We will discuss this approach, which was first proposed in [113], and expand upon its limitations, in the rest of this section. The starting point is the fact that the Euclidean action used in classical lattice Monte-Carlo calculations and the Hamiltonian in Hamiltonian lattice simulations can be related to one another using the transfer matrix formalism.

This formalism, first proposed in [24], is usually used to show that one can obtain the Kogut-Susskind Hamiltonian from the Wilson action by taking the limit of vanishing Euclidean time-like lattice spacing. The starting point is to define a transfer matrix operator T, such that the partition function Z can be written as

$$Z \equiv \int \mathrm{D}U \, e^{-S_E} = \mathrm{Tr} T^N \,, \tag{5.35}$$

where the Euclidean Wilson action is given by

$$S_E = -\beta_s \sum_s \mathrm{Re}\,\mathrm{Tr} P_s - \beta_s \sum_t \mathrm{Re}\,\mathrm{Tr} P_t \,, \tag{5.36}$$

where P_s (P_t) denotes the plaquette operators defined on space-like (time-like) plaquettes. The parameters β_i are defined as

$$\beta_s = \frac{a_0}{g_s^2 a}\,, \qquad \beta_t = \frac{a}{g_t^2 a_0}\,, \tag{5.37}$$

where a_0 is the lattice spacing in the Euclidean time direction.

The transfer matrix is defined as the matrix element of the transfer operator T

$$T_{UU'} \equiv \langle U|T|U'\rangle\,. \tag{5.38}$$

Defining the gauge configurations U_i as those that correspond to the fixed time-slice i on the periodic Euclidean lattice (such that $U_0 = U_N$), one can therefore write

$$T = \int \mathrm{d}U_0 \ldots \mathrm{d}U_{N-1}\, T_{U_0U_1}\, T_{U_1U_2} \ldots T_{U_{N-1}U_0}\,. \tag{5.39}$$

Demanding that the matrices $T_{UU'}$ are Hermitian one finds that the operator

$$T = \exp\left[-a_0\frac{H_B}{2}\right] \exp\left[-a_0 H_E\right] \exp\left[-a_0\frac{H_B}{2}\right]\,, \tag{5.40}$$

reproduces the partition function corresponding to the Wilson action in Eq. (5.36) in the limit $a_0 \to 0$.

Since the transfer operator provides the relation between different time slices on the Euclidean lattice, it is related to the Hamiltonian of the system which governs the differential time evolution of the system. In particular, by taking the $a_0 \to 0$ limit one obtains

$$T = \exp[-a_0 H_{\mathrm{KS}}]\,, \tag{5.41}$$

up to higher order corrections in a_0, producing the result that Eq. (5.40) reproduces the Kogut Susskind Hamiltonian in this limit[8].

It was pointed out in [33, 34, 113] that the operator T in Eq. (5.40) (associated with the heat-kernel action [279]) is equal exactly to the approximation of the time evolution operator of the pure gauge Kogut-Susskind Hamiltonian when using a second order product formula with the same breakup into H_E and H_E as in Eq. (5.40). In other words, one finds

$$T = S_2\left(-ia_0\frac{c}{a}\right)\,, \tag{5.42}$$

[8]Since this transfer operator only reproduces the Wilson action in the $a_0 \to 0$ limit, this transfer operator is not unique, implying one has freedom in defining which lattice Hamiltonian to simulate. For example, one could choose a transfer operator that uses more products of the various exponentials, such that it resembles a higher order representation of the Baker Campbell Hausdorff formula. One could also use an operator that do not contain the operators $\hat{H}_B$ and $\hat{H}_E$ in the exponentials, but rather operators that include higher order corrections in a_0. We are grateful to Wanqiang Liu for discussions on this topic.

without any higher order corrections in a_0 or $\hat{\delta}_t$ so long as $H_1 = H_B$ and $H_2 = H_E$ in Eq. (5.24). This implies that calculations performed in a Hamiltonian lattice gauge theory with Kogut-Susskind Hamiltonian and this second order Trotter evolution is exactly equal to the analytical continuation of a Euclidean path integral calculation performed with the heat kernel action. It is important to keep in mind that this works only for this particular second order Trotter formula. Using other Trotter expansions, such as first order, second order with a different breakup of the Hamiltonian or a different order of terms, of working higher-order product-formulas, however, is a difficult task, as one must start from the desired higher-order transfer matrix and work backwards to determine the associated Euclidean lattice action.

Another issue with this approach is that the exact relationship between the classical Euclidean action and a Trotterized time evolution operator in the Hamiltonian formulation is spoiled when including fermions, and to our knowledge no classical Euclidean action is known that reproduces a Trotterized version of a corresponding Hamiltonian without introducing $\mathcal{O}(a_0)$ corrections. This implies that the approach of using classical calculations to obtain the renormalization trajectory required for Hamiltonian simulations no longer works to the desired accuracy in the presence of fermions.

To summarize, the discussion above shows that one can use classical simulation methods to determine the renormalization trajectory [113, 273] for pure Yang-Mills simulations. In this approach, avoiding a systematic mismatch at $\mathcal{O}(a^2)$ requires performing classical calculations with the heat kernel action, which is considerably more difficult to use than the Wilson action, and no Euclidan action is currently known that allows to extend this approach to include fermions, as is required to simulate the strong interaction.

In the next section we discuss an alternative approach, which does not suffer from the shortcomings of the renormalization approach, and works generically for any lattice Hamiltonian. This approach takes the ordered limit $\lim_{a\to 0}\lim_{\delta_t\to 0}$ in a way that incurs negligible systematic uncertainties from approximate time evolution. The key insight is that, as all numerical computations performed, whether for tuning of parameters or to compute the final observables, have associated statistical uncertainties, and as long as the systematic uncertainties from the approximate time evolution can be guaranteed to be below these statistical uncertainties, they can be neglected.

5.5 Statistically-Bounded Time Evolution Protocol

In the previous section we have outlined a general procedure to numerically compute observables in Hamiltonian lattice gauge theories, including the requirements to renormalize the bare parameters of the Hamiltonian. In particular, we discussed how to perform the two necessary steps when an approximation to the time evolution operator is being used, as is typically always required. In this section we will discuss how to deal with the systematic uncertainties that arise from this approximation to the time evolution operator, in particular how to ensure that these systematic uncertainties have no effect on the continuum limit. We

refer to this general protocol as the Statistically-Bounded Time Evolution (SBTE) protocol. An important aspect of our discussion will be the computational cost of the SBTE protocol.

From the previous discussion on renormalization when approximating time evolution using product formulas, it should be clear that all required numerical calculations are performed with a fixed set of parameters in the Hamiltonian. In the first step, these calculations are used to determine the renormalization trajectory, while in the second step calculations are performed for parameters that lie on this trajectory. It is important to realize that any numerical calculation will come with an associated calculational uncertainty, unrelated to any systematic uncertainty due to underlying theoretical assumptions. Such calculational uncertainty can arise from noisy gates on quantum hardware, but even in the absence of such noise, calculations have statistical uncertainties, either from using Monte-Carlo integrations on classical computers or from measuring expectation values of operators from repeated circuits on quantum computers. This implies that calculating the expectation value of any observable is only possible to a given uncertainty

$$\langle \hat{O} \rangle_{\text{calc}} = \langle \hat{O} \rangle \pm \sigma_{\hat{O}} \,. \tag{5.43}$$

Now consider calculating this expectation value using some approximation to the time evolution operator, where the approximation is controlled by some parameter δ, with exact time evolution recovered in the limit $\delta \to 0$. For approximations using product formulas the parameter $\delta = \delta_t$ is the bare Trotter step size used, but other approximations are possible as well. We will later discuss the case of quantum signal processing (QSP), for which the analogous parameter characterizes the accuracy of the Jacobi-Anger expansion of the time evolution operator in terms of Chebyshev polynomials up to some finite order. For both of these approaches, bounds exist that guarantee that the accuracy of the approximation

$$\left| \langle \hat{O}(\delta) \rangle - \langle \hat{O} \rangle \right| = \epsilon_\delta \,. \tag{5.44}$$

is below a certain value. For product formulas this is achieved by keeping the number of Trotter steps N_{PF} to be above a certain value, while for QSP it requires the number of Chebyshev polynomials k used in the polynomial approximation to exceed a certain value.

As already mentioned in Sec. 5.2 no divergences appear in the limit $\delta \to 0$, even though bare parameters can depend on the value of δ. This implies that choosing a sufficiently large value of δ guarantees that $\epsilon_\delta \ll \sigma_{\hat{O}}$, ensuring that the resulting systematic uncertainties are negligible compared to $\sigma_{\hat{O}}$. Therefore all steps of the previous procedure go through as if one had used exact time evolution.

As will be discussed in Sec. 5.5.1, ensuring a small enough error with product formulas can lead to prohibitively expensive gate costs for two reasons. The first is that, while rigorous error bounds exist when using PFs, they are generally loose, see, *e.g.*, [105]. The second is the polynomial scaling $\mathcal{O}(\epsilon_\delta^{-1/p})$ of the gate cost. These facts explain the desire to deal with this uncertainty using the renormalization procedure discussed, which could result in smaller gate costs at the expense of a more complicated renormalization procedure.

For QSP, the scaling of computational resources with their accuracy is however much more favorable in comparison to product formulas, being proportional to $\log(1/\epsilon_\delta)$, and is in fact provably optimal in both total simulation time and error.

5.5.1 Algorithmic Complexity

In this section we compute the algorithmic cost required to control the approximation accuracy ϵ_δ with respect to the statistical uncertainty $\sigma_{\hat{O}}$ in order to ensure the renormalization steps outlined in Sec. 5.3 can be carried out as if one were performing exact time evolution. In particular, for a specific choice of time evolution algorithm, we ask how to choose the algorithmic parameter δ that controls the approximation error ϵ_δ incurred by the algorithm to ensure that $\epsilon_\delta \leq \sigma_{\hat{O}}/\beta$ where $\beta \in \mathbb{R}^+$ such that $\beta \geq 1$. In this work, we do so for two common (and conceptually simple) time evolution algorithm i.e. Product Formula (PF) as well as for Quantum Signal Processing (QSP) based algorithms.

Let H be the Hamiltonian governing the dynamics of the theory we are interested in simulating. Let $\exp(-iHt)$ and $U_{\text{sim}}(t)$ be exact and approximate time evolution operators, where the latter is the operator implemented by the algorithm of choice. We define the operator Δ_{sim} as follows:

$$\Delta_{\text{sim}} := \mathrm{e}^{-iHt} - U_{\text{sim}}(t)\,, \tag{5.45}$$

such that the spectral norm $\|\Delta_{\text{sim}}\|$ is the exponentiation error incurred by the time evolution algorithm. The following lemma provides a necessary condition on the quantity $\|\Delta_{\text{sim}}\|$ required while computing operator expectation values of the form $\langle\hat{O}(t,a)\rangle$:

Lemma 28. *For any choice of time evolution algorithm and any $\beta \in \mathbb{R}^+$ such that $\beta \geq 1$, choosing algorithmic parameters that ensure*

$$\|\Delta_{\text{sim}}\| \leq \frac{\sigma_{\hat{O}}}{2\beta\|\hat{O}(0,a)\|}\,, \tag{5.46}$$

will result in $\epsilon_\delta \leq \sigma_{\hat{O}}/\beta$.

Proof. From Eq. (5.44), we have that

$$\epsilon_\delta = \|\mathrm{e}^{iHt}\hat{O}(0,a)\mathrm{e}^{-iHt} - U_{\text{sim}}(t)\hat{O}(0,a)U_{\text{sim}}^{-1}(t)\| \tag{5.47}$$

$$= \|\mathrm{e}^{iHt}\hat{O}(0,a)\mathrm{e}^{-iHt} - \left(\mathrm{e}^{iHt} - \Delta_{\text{sim}}^\dagger\right)\hat{O}(0,a)\left(\mathrm{e}^{-iHt} - \Delta_{\text{sim}}\right)\| \tag{5.48}$$

$$= \|\Delta_{\text{sim}}^\dagger\hat{O}(0,a))\mathrm{e}^{-iHt} + \mathrm{e}^{iHt}\hat{O}(0,a)\Delta_{\text{sim}} - \Delta_{\text{sim}}^\dagger\hat{O}(0,a)\Delta_{\text{sim}}\|\,. \tag{5.49}$$

We can then apply the triangle inequality to write:

$$\epsilon_\delta \leq 2\|\Delta_{\text{sim}}\| \cdot \|\hat{O}(0,a)\| + \mathcal{O}(\|\Delta_{\text{sim}}\|)^2\,, \tag{5.50}$$

where we ignore the quadratic correction term and upper bound the RHS of Eq. (5.50) with $\sigma_{\hat{O}}/\beta$ to arrive at Eq. (5.46). □

Equipped with this Lemma, we can now consider specific instances of time evolution algorithms and analyze the algorithmic gate complexity required to appropriately bound the correction term Δ_{sim}. While we consider the cost to do so using rigorous error bounds in the next section, another approach (as we will discuss) is to perform the simulation at several values of decreasing δ and showing that the result does not change within the working statistical precision.

We first consider the complexity of the SBTE protocol when performing time evolution with PFs. As given for example in [106, 114] and Chapter 2, the total number of exponentials appearing in a PF formula grows exponentially with the PF order p. Due to this, for practical applications it is common to choose "low" values of the order and hence we will assume the most common choice of $p = 2$ for the remainder of this work.

In general, for a fixed value of the PF order p, increasing N_{PF} polynomially decreases the exponentiation error (see [106]) given by:

$$\|\Delta_{\text{sim}}\| = \| \left(S_p(t/N_{\text{PF}})\right)^{N_{\text{PF}}} - \mathrm{e}^{-iHt} \| \,. \tag{5.51}$$

We thus have that the total number of Pauli rotations, and hence elementary gate complexity which we denote by χ, required to implement a 2^{nd}-order PF scales linearly with the total number of terms present in the PF formula:

$$\chi = \mathcal{O}\left(N_{\text{PF}} L\right) \,. \tag{5.52}$$

For the rest of this work we will assume that $L = \mathcal{O}(1)$ and hence Eq. (5.52) reduces to

$$\chi = \mathcal{O}(N_{\text{PF}}) \,. \tag{5.53}$$

Thus, under our assumptions, the parameter N_{PF} fully determines the asymptotic scaling of both the exponentiation error $\|\Delta_{\text{sim}}\|$ and the algorithmic gate complexity χ. During an actual simulation procedure the question arises as to how to actually choose an *explicit* value of N_{PF} such that $\epsilon_\delta \leq \sigma_{\hat{O}}/\beta$ for 2^{nd} order PFs. To this end, we use Proposition 16 from [106] with $t \to t/N_{\text{PF}}$ and use the triangle inequality N_{PF}-many times to arrive at the following upper bound for $\|\Delta_{\text{sim}}\|$:

$$\|\Delta_{\text{sim}}\| \leq \frac{f t^3}{N_{\text{PF}}^2} \,. \tag{5.54}$$

where $f \in \mathbb{C}$ defined as follows:

$$\begin{aligned} f := \frac{1}{12} \sum_{\ell_1=1}^{L} \left\| \left[\sum_{\ell_3=\ell_1}^{L} H_{\ell_3}, \left[\sum_{\ell_2=\ell_1+1}^{L} H_{\ell_2}, H_{\ell_1} \right] \right] \right\| \\ + \frac{1}{24} \sum_{\ell_1=1}^{L} \left\| \left[H_{\ell_1}, \sum_{\ell_2=\ell_1+1}^{L} H_{\ell_2} \right] \right\| \,. \end{aligned} \tag{5.55}$$

We can then upper bound the RHS of the inequality Eq. (5.54) with the RHS of the inequality from Lemma 28 to arrive at the following sufficient choice for N_{PF}

$$N_{\mathrm{PF}} = \sqrt{\frac{2\beta\|\hat{O}(0,a)\|ft^3}{\sigma_{\mathrm{O}}}}\,, \tag{5.56}$$

Thus, Eq. (5.56) provides a sufficient choice for the Trotter number N_{PF} required to drive the error stemming from approximating the true time evolution operator with a PF formula below the statistical error $\sigma_{\hat{O}}$ by a factor of β. This allows one to use all the machinery described in Sec. 5.3 and proceed as if the time evolution is exact.

If we instead choose to perfrom time evolution using QSP we can use the machinery developed for instance in Chapters 2 and 3, which we discuss again using slightly different notation. The theory of QSP allows us to implement bounded finite degree polynomial transformations of unitary operators. In particular, let U be a unitary operation and let $P(U)$ be a degree-d polynomial such that

$$P(x)^2 \leq 1 \;\; \forall x \in \{y \in \mathbb{C} : |y| = 1\}. \tag{5.57}$$

From Corollary 5 of [110], we are then guaranteed the existence of a set of $(d+1)$-many pairs of phases $\mathcal{A} = \{(\theta_1, \phi_1), (\theta_2, \phi_2), \ldots, (\theta_d, \phi_d), (\theta_{d+1}, \phi_{d+1})\}$ (where $(\theta_i, \phi_i) \in \mathbb{R}^2$) and a scalar $\lambda \in \mathbb{R}$ such that the following is true:

$$U_{\mathcal{A}}(C_{U,0}) := \left(\prod_{i=2}^{d+1} R(\theta_i, \phi_i, 0) C_{U,0}\right) R(\theta_1, \phi_1, \lambda) \tag{5.58}$$

$$= \begin{pmatrix} P(U) & * \\ * & * \end{pmatrix}, \tag{5.59}$$

where each $*$ represents a sub-block with entries chosen to make $U_{\mathcal{A}}(C_{U,0})$ unitary, $C_{U,0}$ is the zero-controlled application of U, i.e.

$$C_{U,0} := \begin{pmatrix} U & 0 \\ 0 & I \end{pmatrix}, \tag{5.60}$$

and $R(\theta_i, \phi_i, \lambda)$ represents the following $SU(2)$ rotation:

$$R(\theta_i, \phi_i, \lambda) = \begin{pmatrix} e^{i(\lambda+\phi_i)}\cos(\theta_i) & e^{i\phi_i}\sin(\theta_i) \\ e^{i\lambda}\sin(\theta_i) & -\cos(\theta_i) \end{pmatrix} \otimes I. \tag{5.61}$$

In other words, the unitary $U_{\mathcal{A}}(C_{U,0})$ encodes the polynomial $P(U)$ in its top left block. Thus, the goal of QSP is to find the appropriate set of phases $\mathcal{A}$ and scalar λ such that the resulting QSP unitary encodes a desired polynomial. As shown in [110], the set $\mathcal{A}$ can be determined efficiently and accurately using an optimization algorithm for d up to 2^{24}.

Our approach to apply QSP to the task of time evolution involves the following steps:

1. Construct a unitary that provides query access to the eigenvalues of the Hamiltonian H.

2. Determine the set $\mathcal{A}$ necessary to implement a polynomial transformation of these eigenvalues such that the resulting polynomial sufficiently approximates the true time evolution operator

We accomplish the first of these tasks as follows. Given a Hamiltonian H, the *block-encoding* of H (denoted by U_H) is a unitary that encodes H in (typically) its top left block:

$$U_H = \begin{pmatrix} H/\alpha_H & * \\ * & * \end{pmatrix}, \tag{5.62}$$

where $\alpha_H \in \mathbb{R}^+$ is defined such that $\|H/\alpha_H\| \leq 1$ and $*$ refers to a block of entries that make U_H unitary. The size of the block $*$ corresponds to the number of ancilla qubits used to construct U_H. Note that Eq. (5.58) provides another example of a block-encoding, specifically of the polynomial $P(U)$. The most general construction of U_H can be performed by the linear combination of unitaries (LCU) method [111]. However, this can prove to be computationally challenging. In many instances, including Hamiltonian LGTs, the structure of H can lend itself to more efficient block-encoding constructions [114, 115, 133, 280].

Given U_H, we then use it as input to the *qubitization* procedure presented in [109] to construct a unitary W_H (known as the walk operator or iterate) designed such that the eigenphases of W_H encode the eigenvalues of the rescaled Hamiltonian H/α_H and thus provide the desired query access to the eigenvalues of H. Thus, by setting $U \to W_H$ in Eq. (5.58) we can construct polynomial transformations of W_H and hence also of H/α_H. In particular, we seek a polynomial transformation that satisfies:

$$P(W_H) = \exp(-iHt) = \exp\left(-i\left(\frac{H}{\alpha_H}\right)\alpha_H t\right). \tag{5.63}$$

If we let $x \in \text{spec}(H/\alpha_H)$, one can use the Jacobi-Anger expansion [108, 116] to obtain the following polynomial decomposition

$$\begin{aligned} \mathrm{e}^{-i\alpha_H x t} = J_0(\alpha_H\, t) + 2 \sum_{k \text{ even},>0}^{\infty} (-i)^{k/2}\, J_k(\alpha_H\, t) T_k(x) \\ + 2i \sum_{k \text{ odd},>0}^{\infty} (-i)^{(k-1)/2}\, J_k(\alpha_H\, t) T_k(x)\,, \end{aligned} \tag{5.64}$$

where J_k is the k^{th} order Bessel function of the first kind and T_k is the k^{th} Chebyshev polynomial of the first kind. If we construct a finite degree approximation to this polynomial decomposition, we can then encode an approximation to the true time evolution operator in

the top left block of a unitary of the form Eq. (5.58). To accomplish this, we consider the following truncated decomposition:

$$\begin{aligned} \mathrm{e}^{-i\alpha_H x t} = J_0(\alpha_H\, t) + 2 \sum_{k \text{ even},>0}^{N_{\mathrm{QSP}}} (-i)^{k/2}\, J_k(\alpha_H\, t) T_k(x) \\ + 2i \sum_{k \text{ odd},>0}^{N_{\mathrm{QSP}}} (-i)^{(k-1)/2}\, J_k(\alpha_H\, t) T_k(x)\,. \end{aligned} \tag{5.65}$$

Thus, N_{QSP} determines the degree of the polynomial used to approximate the true time evolution operator, and from Eq. (5.58) we see that the elementary gate complexity χ scales linearly with N_{QSP}:

$$\chi = \mathcal{O}(N_{\mathrm{QSP}}). \tag{5.66}$$

Thus, similar to the parameter N_{PF}, N_{QSP} is the key algorithmic parameter for QSP that determines both the exponentiation error $\|\Delta_{\mathrm{sim}}\|$ (as shown for instance in [108], $\|\Delta_{\mathrm{sim}}\|$ decreases *exponentially* as we increase N_{QSP} and χ. In the language of Sec. 5.5, we identify $\delta = 1/N_{\mathrm{QSP}}$.

We now determine an explicit choice for the value of N_{QSP} required to satisfy the condition $\epsilon_\delta \leq \sigma_{\hat{O}}/\beta$ and hence perform an actual simulation. From [109, 116, 281], we have that

$$N_{\mathrm{QSP}} \geq \frac{e\alpha_H t}{2} + \log\left(\frac{1}{\|\Delta_{\mathrm{sim}}\|}\right). \tag{5.67}$$

for which the proof is based on the super-exponentially decaying upper bound on the norms of Bessel functions with respect to the order. We can then use Lemma 28 to arrive at the following sufficient choice for N_{QSP}:

$$N_{\mathrm{QSP}} \geq \frac{e\alpha_H t}{2} + \log\left(2\beta\|\hat{O}(0,a)\|\right) + \log\left(\frac{1}{\sigma_{\mathrm{O}}}\right). \tag{5.68}$$

5.5.2 Comparisons

The most striking differences between the two algorithmic kernels we consider for Hamiltonian simulation are associated with how the rigorous error bounds scale with key parameters N_{PF} and N_{QSP} functionally depend on t, $\sigma_{\hat{O}}$, β, and $\|\hat{O}(0,a)\|$. From Eqs. (5.56) and (5.68) we see that N_{QSP} asymptotically scales more favorably than N_{PF} with respect to all these quantities. In particular from Eq. (5.56) we see that N_{PF} scales polynomially with $1/\sigma_{\hat{O}}$ whereas from Eq. (5.68) we see that N_{QSP} scales *logarithmically* with $1/\sigma_{\hat{O}}$. Moreover, while N_{PF} scales polynomially with the product $\beta\|\hat{O}(0,a)\|$, N_{QSP} again scales *logarithmically* with the product $\beta\|\hat{O}(0,a)\|$. Furthermore, the additional logarithmic overhead when using QSP is *additive* relative to the cost of simulating time evolution to error $\sigma_{\hat{O}}$. Because

the scale factor must satisfy $\alpha_H \geq \|H\|$, it generally grows with volume, which implies that the overhead is entirely negligible for any physically interesting operator.

To better understand the additional overhead for both methods, we consider two physically relevant classes of operators $\hat{O}(0,a)$. In particular, consider the overhead for local operators, *e.g.*, the local quark condensate, and operators that involve summing over the entire lattice, *e.g.*, the hadronic tensor. In the case of local operators, the operator norm $\|\hat{O}(0,a)\|$ is independent of the volume and may introduce only a small overhead when using PFs. Simulating non-local operators whose norm $\|\hat{O}(0,a)\|$ grows with the volume, however, will result in a significant cost increase for simulations via PFs. While it is possible that this cost scaling could be improved with more detailed error analysis, if one instead uses QSP, the overhead can be guaranteed to be entirely negligible *a priori*.

Additionally, as we show for the case of a ϕ^4 Hamiltonian in App. D.2, the QSP error bound used to prove Eq. (5.68) is generally known to be tighter than the analogous PF error bound. Thus, employing the SBTE protocol we propose in this work, the QSP error bound would then enable us to choose smaller values of N_{QSP} than the values of N_{PF} implied by the PF bound relative to the ideal values of these quantities (i.e. the smallest values of $N_{\text{QSP}}, N_{\text{PF}}$ for which the desired error bound is satisfied).

While employing rigorous error bounds to determine the value of δ that guarantees $\epsilon_\delta \leq \sigma_{\hat{O}}/\beta$ has the advantage of only needing to perform a single simulation, an alternative approach is to simply perform the calculation at smaller values of δ until the value $\langle\hat{O}(\delta)\rangle$ no longer changes within error bars. Such an approach could result in savings if the error bounds are known to be loose, as is the case with PFs.

From this discussion it should be clear that, for a given problem, it is not immediately obvious which technique would be better suited for performing the SBTE protocol proposed in this work and result in an overall lower gate complexity. The answer to this question requires knowledge of the prefactors appearing in Eqs. (5.53) and (5.66) as well as the quantities f (Eq. (5.55)) and α_H (Eq. (5.62)). The major contribution to the QSP prefactor is associated with the task of constructing the block-encoding for a given Hamiltonian (which also determines α_H). This task is highly problem dependent and hence we expect the more efficient algorithmic kernel to also change depending on the specific Hamiltonian as well as the specific values of $\|\Delta_{\text{sim}}\|, t$, and the size of the lattice. The work in [114, 115] and Chapters 2 and 3 present, for example, values of $\|\Delta_{\text{sim}}\|$ and t for which the *combined* task of block-encoding and time evolution via QSP proves to be more efficient than time evolution by a PF for the case of the ϕ^4 Hamiltonian at specific lattice sizes. Thus, for a specific problem/Hamiltonian one can take advantage of any structure inherent to the problem to construct efficient block-encodings and either analytically or numerically (by extrapolation) determine the parameter ranges corresponding to the either QSP or PF being the more efficient algorithmic choice.

5.6 Discussion and Conclusion

In this work, we have presented an approach, called the SBTE protocol, to taking the continuum limit in dynamical simulations of Hamiltonian lattice field theories. Our framework is conceptually simpler than previous methods based on a renormalization procedure, which were shown to allow classical calculations using the heat-kernel action to obtain the renormalization trajectory for time evolution in pure Yang-Mills gauge theory. One advantage of this new protocol is that it works for any approximation to the time evolution operator, but more important is that it works for any Hamiltonian, in particular for the strong interaction including fermions, which is one of the most important use cases for Hamiltonian simulations. The SBTE protocol is based on the idea that as long as the systematic error from approximate time evolution is negligible compared to the statistical uncertainty of the simulation, its effect on the continuum limit can be safely ignored. By exploiting rigorous upper bounds on the error incurred from approximate time evolution, one can *a priori* determine the quantum gate cost required to achieve this.

To demonstrate our approach, we considered simulation via product formulas and methods based on Quantum Signal Processing. Using existing error bounds, we showed that, while driving the time evolution error below the statistical threshold can result in significant overhead when using PFs due to the sub-optimal scaling $\mathcal{O}(t^{1+1/p}\epsilon^{-1/p})$, this additional cost is negligible when using QSP based methods due to its provably optimal scaling $\mathcal{O}(\alpha_H t + \log\frac{1}{\epsilon})$, in particular due to the *additive* $\log\frac{1}{\epsilon}$ scaling with error; for algorithms with multiplicative error scaling of the form $\sim t\log\frac{1}{\epsilon}$, the overhead will be smaller than for PFs, but still non-negligible.

Rather than using error bounds in the SBTE protocol, we discussed an alternative approach where one simply performs the calculation with smaller values of δ until the value $\langle\hat{O}(\delta)\rangle$ no longer changes within error bars. This approach is most likely to be beneficial when the error bounds are known to be loose, as is the case with product-formulas. In general, the computational cost in the SBTE protocol is simply given by the computational cost of ensuring the systematic uncertainties introduced by an approximation to the time-evolution operator are below the statistical ones.

In addition to developing this general alternative approach, we developed a simplified renormalization trajectory for PF simulations. Exploiting the fact that BCH error terms in the effective Hamiltonian are irrelevant operators who disappear in the continuum limit, one can simply use the bare parameters tuned assuming exact time evolution while only introducing $\mathcal{O}(a^k)$ errors, where k is the order of the PF used. This new trajectory has several advantages, including the simplicity of working at a fixed value of the bare step-size $\hat{\delta}_t$, and the ability to tune the bare parameters using algorithms beyond PFs. Whether the computational simplicity of this approach outweighs the potentially larger lattice spacing errors will be interesting to study in future work.

In the context of PF simulations, we also provided a detailed discussion of the proposal for using classical resources [113] to help reduce the quantum cost of determining the renormalization trajectory. In this discussion we highlighted that for pure gauge theories these

methods require classical simulations with complicated actions. Once fermions are included, no classical actions are known for this approach to work. It is also important to keep in mind that, because performing time evolution requires quantum computers with similar capabilities as would be needed for tuning bare parameters and scale setting, such methods do not allow one to get away with using smaller or noisier quantum devices. These approaches can, however, trade total run-time on quantum computers for run-time on classical computers.

Our findings suggest several directions for future research. One promising avenue is to develop an explicit effective Hamiltonian description for QSP-based time evolution, which could provide deeper insight into the renormalization process and potentially inspire new strategies similar to the PF-based approach. Additionally, a comprehensive end-to-end renormalization procedure—including parameter tuning, scale setting, and operator renormalization—remains to be established for general simulation algorithms. There are many different renormalization schemes one can imagine using, and detailed comparisons of the computational cost are necessary to understand which approach is best suited for a given problem.

Ultimately, our method lays the groundwork for systematic and fair comparisons between different quantum simulation algorithms for lattice field theory simulations, enabling one to assess the total quantum resources required to obtain continuum physics. This, in turn, will be crucial for guiding the development of quantum algorithms and hardware toward achieving quantum advantage in this challenging domain.

Furthermore, our results have significant implications for future Hamiltonian LGT quantum simulation programs. Drawing inspiration from Euclidean lattice QCD simulations, it is highly desirable to be able to reuse as many components of a quantum simulation as possible. Prior to our work, the effective Hamiltonians one actually evolves according to when using approximate time evolution were viewed as different lattice regularizations, each requiring their own tuning process. By identifying renormalization trajectories associated with exact Hamiltonians, our work enables the reuse of costly computations across different algorithms and observables.

Bibliography

[1] Tatsumi Aoyama, Toichiro Kinoshita, and Makiko Nio. "Theory of the Anomalous Magnetic Moment of the Electron". In: *Atoms* 7.1 (2019), p. 28. DOI: 10.3390/atoms7010028.

[2] F. J. Dyson. "Divergence of perturbation theory in quantum electrodynamics". In: *Phys. Rev.* 85 (1952), pp. 631–632. DOI: 10.1103/PhysRev.85.631.

[3] Colin J. Morningstar and Mike J. Peardon. "The Glueball spectrum from an anisotropic lattice study". In: *Phys. Rev. D* 60 (1999), p. 034509. DOI: 10.1103/PhysRevD.60.034509. arXiv: hep-lat/9901004.

[4] Markus Q. Huber, Christian S. Fischer, and Hèlios Sanchis-Alepuz. "Spectrum of scalar and pseudoscalar glueballs from functional methods". In: *Eur. Phys. J. C* 80.11 (2020), p. 1077. DOI: 10.1140/epjc/s10052-020-08649-6. arXiv: 2004.00415 [hep-ph].

[5] Wolfgang Ochs. "The Status of Glueballs". In: *J. Phys. G* 40 (2013), p. 043001. DOI: 10.1088/0954-3899/40/4/043001. arXiv: 1301.5183 [hep-ph].

[6] H. Georgi and S. L. Glashow. "Unity of All Elementary Particle Forces". In: *Phys. Rev. Lett.* 32 (1974), pp. 438–441. DOI: 10.1103/PhysRevLett.32.438.

[7] John C. Baez and John Huerta. "The Algebra of Grand Unified Theories". In: *Bull. Am. Math. Soc.* 47 (2010), pp. 483–552. DOI: 10.1090/S0273-0979-10-01294-2. arXiv: 0904.1556 [hep-th].

[8] Lindsay Forestell, David E. Morrissey, and Kris Sigurdson. "Non-Abelian Dark Forces and the Relic Densities of Dark Glueballs". In: *Phys. Rev. D* 95.1 (2017), p. 015032. DOI: 10.1103/PhysRevD.95.015032. arXiv: 1605.08048 [hep-ph].

[9] David McKeen et al. "Dark matter from dark glueball dominance". In: *Phys. Rev. D* 111.1 (2025), p. 015044. DOI: 10.1103/PhysRevD.111.015044. arXiv: 2406.18635 [hep-ph].

[10] Pierluca Carenza et al. "Glueball Dark Matter Revisited". In: *Phys. Rev. Lett.* 129.26 (2022), p. 261302. DOI: 10.1103/PhysRevLett.129.261302. arXiv: 2207.13716 [hep-ph].

[11] Edward Hardy et al. "Big Bang Synthesis of Nuclear Dark Matter". In: *JHEP* 06 (2015), p. 011. DOI: 10.1007/JHEP06(2015)011. arXiv: 1411.3739 [hep-ph].

[12] Lucile Savary and Leon Balents. "Quantum spin liquids: a review". In: *Rept. Prog. Phys.* 80.1 (2017), p. 016502. DOI: 10.1088/0034-4885/80/1/016502. arXiv: 1601.03742 [cond-mat.str-el].

[13] Simon Trebst and Ciarán Hickey. "Kitaev materials". In: *Phys. Rept.* 950 (2022), pp. 1–37. DOI: 10.1016/j.physrep.2021.11.003.

[14] Yukitoshi Motome et al. "Hunting Majorana Fermions in Kitaev Magnets". In: *J. Phys. Soc. Jap.* 89.1 (2020), p. 012002. DOI: 10.7566/JPSJ.89.012002. arXiv: 1909.02234 [cond-mat.str-el].

[15] Michael A. Levin and Xiao-Gang Wen. "String net condensation: A Physical mechanism for topological phases". In: *Phys. Rev. B* 71 (2005), p. 045110. DOI: 10.1103/PhysRevB.71.045110. arXiv: cond-mat/0404617.

[16] Juan Martin Maldacena. "The Large N limit of superconformal field theories and supergravity". In: *Adv. Theor. Math. Phys.* 2 (1998), pp. 231–252. DOI: 10.4310/ATMP.1998.v2.n2.a1. arXiv: hep-th/9711200.

[17] Alexander J. Buser et al. "Quantum simulation of gauge theory via orbifold lattice". In: *JHEP* 09 (2021), p. 034. DOI: 10.1007/JHEP09(2021)034. arXiv: 2011.06576 [hep-th].

[18] M. Göpfert and G. Mack. "Proof of confinement of static quarks in three-dimensional U(1) lattice gauge theory for all values of the coupling constant". In: *Commun. Math. Phys.* 82 (1982), pp. 545–606. DOI: 10.1007/BF01961240.

[19] Sourav Chatterjee. "A short proof of confinement in three-dimensional lattice gauge theories with a central U(1)". In: *arXiv preprint* (2026). arXiv: 2602.00436.

[20] Kenneth G. Wilson. "Confinement of quarks". In: *Phys. Rev. D* 10 (8 Oct. 1974), pp. 2445–2459. DOI: 10.1103/PhysRevD.10.2445. URL: https://link.aps.org/doi/10.1103/PhysRevD.10.2445.

[21] K. Symanzik. "Continuum limit and improved action in lattice theories: (I). Principles and ϕ^4 theory". In: *Nucl. Phys. B* 226 (1983), pp. 187–204. DOI: 10.1016/0550-3213(83)90468-6.

[22] K. Symanzik. "Continuum limit and improved action in lattice theories: (II). $O(N)$ non-linear sigma model in perturbation theory". In: *Nucl. Phys. B* 226 (1983), pp. 205–227. DOI: 10.1016/0550-3213(83)90469-8.

[23] M. Lüscher and P. Weisz. "On-shell improved lattice gauge theories". In: *Commun. Math. Phys.* 97 (1985). Erratum: Commun. Math. Phys. 98, 433 (1985), pp. 59–77. DOI: 10.1007/BF01206178.

[24] Michael Creutz. *Quarks, Gluons and Lattices*. Oxford University Press, 1983. ISBN: 978-1-009-29039-5, 978-1-009-29038-8, 978-1-009-29037-1, 978-0-521-31535-7. DOI: 10.1017/9781009290395.

[25] Holger Bech Nielsen and M. Ninomiya. "Absence of Neutrinos on a Lattice. 1. Proof by Homotopy Theory". In: *Nucl. Phys. B* 185 (1981). [Erratum: Nucl. Phys. B 195, 541 (1982)], p. 20. DOI: 10.1016/0550-3213(81)90361-8.

[26] Holger Bech Nielsen and M. Ninomiya. "Absence of Neutrinos on a Lattice. 2. Intuitive Topological Proof". In: *Nucl. Phys. B* 193 (1981), p. 173. DOI: 10.1016/0550-3213(81)90524-1.

[27] Colin Morningstar. "The Monte Carlo method in quantum field theory". In: *21st Annual Hampton University Graduate Studies Program (HUGS 2006)*. Feb. 2007. arXiv: hep-lat/0702020.

[28] Zohreh Davoudi et al. "Report of the Snowmass 2021 Topical Group on Lattice Gauge Theory". In: *Snowmass 2021*. Sept. 2022. arXiv: 2209.10758 [hep-lat].

[29] Matthias Troyer and Uwe-Jens Wiese. "Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations". In: *Phys. Rev. Lett.* 94 (2005), p. 170201. DOI: 10.1103/PhysRevLett.94.170201. arXiv: cond-mat/0408370.

[30] E. Y. Loh et al. "Sign problem in the numerical simulation of many-electron systems". In: *Phys. Rev. B* 41 (1990), pp. 9301–9307. DOI: 10.1103/PhysRevB.41.9301.

[31] Gaopei Pan and Zi Yang Meng. "Sign Problem in Quantum Monte Carlo Simulation". In: (Apr. 2022). DOI: 10.1016/B978-0-323-90800-9.00095-0. arXiv: 2204.08777 [cond-mat.str-el].

[32] Angus Kan et al. "Investigating a (3+1)D topological θ-term in the Hamiltonian formulation of lattice gauge theories for quantum and classical simulations". In: *Phys. Rev. D* 104.3 (2021), p. 034504. DOI: 10.1103/PhysRevD.104.034504. arXiv: 2105.06019 [hep-lat].

[33] Gurtej Kanwar and Michael L. Wagman. "Real-time lattice gauge theory actions: Unitarity, convergence, and path integral contour deformations". In: *Phys. Rev. D* 104.1 (2021), p. 014513. DOI: 10.1103/PhysRevD.104.014513. arXiv: 2103.02602 [hep-lat].

[34] Hiroki Hoshina, Hirotsugu Fujii, and Yoshio Kikukawa. "Schwinger-Keldysh formalism for Lattice Gauge Theories". In: *PoS* LATTICE2019 (2020), p. 190. DOI: 10.22323/1.363.0190.

[35] Gert Aarts et al. "Electrical conductivity and charge diffusion in thermal QCD from the lattice". In: *JHEP* 02 (2015), p. 186. DOI: 10.1007/JHEP02(2015)186. arXiv: 1412.6411 [hep-lat].

[36] L. Maiani and M. Testa. "Final state interactions from Euclidean correlation functions". In: *Phys. Lett. B* 245 (1990), pp. 585–590. DOI: 10.1016/0370-2693(90)90695-3.

[37] M. Luscher. "Volume Dependence of the Energy Spectrum in Massive Quantum Field Theories. 2. Scattering States". In: *Commun. Math. Phys.* 105 (1986), pp. 153–188. DOI: 10.1007/BF01211097.

[38] M. Lüscher. "Two particle states on a torus and their relation to the scattering matrix". In: *Nucl. Phys. B* 354 (1991), pp. 531–578. DOI: 10.1016/0550-3213(91)90366-6.

[39] John Bulava and Maxwell T. Hansen. "Scattering amplitudes from finite-volume spectral functions". In: *Phys. Rev. D* 100.3 (2019), p. 034521. DOI: 10.1103/PhysRevD.100.034521. arXiv: 1903.11735 [hep-lat].

[40] Mattia Bruno and Maxwell T. Hansen. "Variations on the Maiani-Testa approach and the inverse problem". In: *JHEP* 06 (2021), p. 043. DOI: 10.1007/JHEP06(2021)043. arXiv: 2012.11488 [hep-lat].

[41] Christian W. Bauer et al. "Quantum Algorithm for High Energy Physics Simulations". In: *Phys. Rev. Lett.* 126.6 (2021), p. 062001. DOI: 10.1103/PhysRevLett.126.062001. arXiv: 1904.03196 [hep-ph].

[42] Christian W. Bauer, So Chigusa, and Masahito Yamazaki. "Quantum parton shower with kinematics". In: *Phys. Rev. A* 109.3 (2024), p. 032432. DOI: 10.1103/PhysRevA.109.032432. arXiv: 2310.19881 [hep-ph].

[43] Christian W. Bauer. "Efficient use of quantum computers for collider physics". In: (Mar. 2025). arXiv: 2503.16602 [hep-ph].

[44] Thomas D. Cohen et al. "Quantum algorithms for transport coefficients in gauge theories". In: *Phys. Rev. D* 104.9 (2021), p. 094514. DOI: 10.1103/PhysRevD.104.094514. arXiv: 2104.02024 [hep-lat].

[45] Francesco Turro, Anthony Ciavarella, and Xiaojun Yao. "Classical and quantum computing of shear viscosity for (2+1)D SU(2) gauge theory". In: *Phys. Rev. D* 109.11 (2024), p. 114511. DOI: 10.1103/PhysRevD.109.114511. arXiv: 2402.04221 [hep-lat].

[46] Ivan A. Chernyshev et al. "Pathfinding Quantum Simulations of Neutrinoless Double-β Decay". In: (June 2025). arXiv: 2506.05757 [quant-ph].

[47] John B. Kogut and Leonard Susskind. "Hamiltonian Formulation of Wilson's Lattice Gauge Theories". In: *Phys. Rev. D* 11 (1975), pp. 395–408. DOI: 10.1103/PhysRevD.11.395.

[48] John B. Kogut. "An Introduction to Lattice Gauge Theory and Spin Systems". In: *Rev. Mod. Phys.* 51 (1979), p. 659. DOI: 10.1103/RevModPhys.51.659.

[49] Mari Carmen Bañuls et al. "Efficient basis formulation for 1+1 dimensional SU(2) lattice gauge theory: Spectral calculations with matrix product states". In: *Phys. Rev. X* 7.4 (2017), p. 041046. DOI: 10.1103/PhysRevX.7.041046. arXiv: 1707.06434 [hep-lat].

[50] M. C. Bañuls et al. "Simulating Lattice Gauge Theories within Quantum Technologies". In: *Eur. Phys. J. D* 74.8 (2020), p. 165. DOI: 10.1140/epjd/e2020-100571-8. arXiv: 1911.00003 [quant-ph].

[51] Falk Bruckmann, Karl Jansen, and Stefan Kühn. "O(3) nonlinear sigma model in 1+1 dimensions with matrix product states". In: *Phys. Rev. D* 99.7 (2019), p. 074501. DOI: 10.1103/PhysRevD.99.074501. arXiv: 1812.00944 [hep-lat].

[52] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. "Quantum Computation of Scattering in Scalar Quantum Field Theories". In: *Quant. Inf. Comput.* 14 (2014), pp. 1014–1080. arXiv: 1112.4833 [hep-th].

[53] Richard P. Feynman. "Simulating physics with computers". In: *Int. J. Theor. Phys.* 21 (1982). Ed. by L. M. Brown, pp. 467–488. DOI: 10.1007/BF02650179.

[54] Seth Lloyd. "Universal Quantum Simulators". In: *Science* 273.5278 (1996), p. 1073. DOI: 10.1126/science.273.5278.1073.

[55] Steven R. White. "Density matrix formulation for quantum renormalization groups". In: *Phys. Rev. Lett.* 69 (1992), pp. 2863–2866. DOI: 10.1103/PhysRevLett.69.2863.

[56] Ulrich Schollwöck. "The density-matrix renormalization group in the age of matrix product states". In: *Annals Phys.* 326 (2011), pp. 96–192. DOI: 10.1016/j.aop.2010.09.012. arXiv: 1008.3477 [cond-mat.str-el].

[57] F. Verstraete, V. Murg, and J. I. Cirac. "Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems". In: *Adv. Phys.* 57.2 (2008), pp. 143–224. DOI: 10.1080/14789940801912366.

[58] Pasquale Calabrese and John Cardy. "Evolution of entanglement entropy in one-dimensional systems". In: *J. Stat. Mech.* 0504 (2005), P04010. DOI: 10.1088/1742-5468/2005/04/P04010. arXiv: cond-mat/0503393.

[59] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. "Quantum Algorithms for Quantum Field Theories". In: *Science* 336 (2012), pp. 1130–1133. DOI: 10.1126/science.1217069. arXiv: 1111.3633 [quant-ph].

[60] Tim Byrnes and Yoshihisa Yamamoto. "Simulating lattice gauge theories on a quantum computer". In: *Phys. Rev. A* 73 (2006), p. 022328. DOI: 10.1103/PhysRevA.73.022328. arXiv: quant-ph/0510027.

[61] Natalie Klco, Jesse R. Stryker, and Martin J. Savage. "SU(2) non-Abelian gauge field theory in one dimension on digital quantum computers". In: *Phys. Rev. D* 101.7 (2020), p. 074512. DOI: 10.1103/PhysRevD.101.074512. arXiv: 1908.06935 [quant-ph].

[62] Anthony Ciavarella, Natalie Klco, and Martin J. Savage. "Trailhead for quantum simulation of SU(3) Yang-Mills lattice gauge theory in the local multiplet basis". In: *Phys. Rev. D* 103.9 (2021), p. 094501. DOI: 10.1103/PhysRevD.103.094501. arXiv: 2101.10227 [quant-ph].

[63] Angus Kan and Yunseong Nam. "Lattice Quantum Chromodynamics and Electrodynamics on a Universal Quantum Computer". In: (July 2021). arXiv: 2107.12769 [quant-ph].

[64] Zohreh Davoudi, Alexander F. Shaw, and Jesse R. Stryker. "General quantum algorithms for Hamiltonian simulation with applications to a non-Abelian lattice gauge theory". In: *Quantum* 7 (2023), p. 1213. DOI: 10.22331/q-2023-12-20-1213. arXiv: 2212.14030 [hep-lat].

[65] Roland C. Farrell et al. "Preparations for quantum simulations of quantum chromodynamics in 1+1 dimensions. I. Axial gauge". In: *Phys. Rev. D* 107.5 (2023), p. 054512. DOI: 10.1103/PhysRevD.107.054512. arXiv: 2207.01731 [quant-ph].

[66] Roland C. Farrell et al. "Preparations for quantum simulations of quantum chromodynamics in 1+1 dimensions. II. Single-baryon β-decay in real time". In: *Phys. Rev. D* 107.5 (2023), p. 054513. DOI: 10.1103/PhysRevD.107.054513. arXiv: 2209.10781 [quant-ph].

[67] Roland C. Farrell et al. "Scalable Circuits for Preparing Ground States on Digital Quantum Computers: The Schwinger Model Vacuum on 100 Qubits". In: *PRX Quantum* 5.2 (2024), p. 020315. DOI: 10.1103/PRXQuantum.5.020315. arXiv: 2308.04481 [quant-ph].

[68] Zhiyao Li, Dorota M. Grabowska, and Martin J. Savage. "Sequency Hierarchy Truncation (SeqHT) for Adiabatic State Preparation and Time Evolution in Quantum Simulations". In: (July 2024). arXiv: 2407.13835 [quant-ph].

[69] Praveen Balaji et al. *Quantum Circuits for SU(3) Lattice Gauge Theory*. Mar. 2025. arXiv: 2503.08866 [hep-lat].

[70] E. A. Martinez et al. "Real-time dynamics of lattice gauge theories with a few-qubit quantum computer". In: *Nature* 534 (2016), pp. 516–519. DOI: 10.1038/nature18318. arXiv: 1605.04570 [quant-ph].

[71] Roland C. Farrell et al. "Quantum simulations of hadron dynamics in the Schwinger model using 112 qubits". In: *Phys. Rev. D* 109.11 (2024), p. 114510. DOI: 10.1103/PhysRevD.109.114510. arXiv: 2401.08044 [quant-ph].

[72] Christian W. Bauer et al. "Quantum Simulation for High-Energy Physics". In: *PRX Quantum* 4.2 (2023), p. 027001. DOI: 10.1103/PRXQuantum.4.027001. arXiv: 2204.03381 [quant-ph].

[73] Christian W. Bauer et al. "Quantum simulation of fundamental particles and forces". In: *Nature Rev. Phys.* 5.7 (2023), pp. 420–432. DOI: 10.1038/s42254-023-00599-8. arXiv: 2404.06298 [hep-ph].

[74] Ivan M. Burbano et al. "Real-time Estimators for Scattering Observables: A full account of finite volume errors for quantum simulation". In: (June 2025). arXiv: 2506.06511 [hep-lat].

[75] Yu Tong et al. "Provably accurate simulation of gauge theories and bosonic systems". In: *Quantum* 6 (2022), p. 816. DOI: 10.22331/q-2022-09-22-816. arXiv: 2110.06942 [quant-ph].

[76] Indrakshi Raychowdhury and Jesse R. Stryker. "Solving Gauss's Law on Digital Quantum Computers with Loop-String-Hadron Digitization". In: *Phys. Rev. Res.* 2.3 (2020), p. 033039. DOI: 10.1103/PhysRevResearch.2.033039. arXiv: 1812.07554 [hep-lat].

[77] Indrakshi Raychowdhury and Jesse R. Stryker. "Loop, string, and hadron dynamics in SU(2) Hamiltonian lattice gauge theories". In: *Phys. Rev. D* 101.11 (2020), p. 114502. DOI: 10.1103/PhysRevD.101.114502. arXiv: 1912.06133 [hep-lat].

[78] Saurabh V. Kadam, Indrakshi Raychowdhury, and Jesse R. Stryker. "Loop-string-hadron formulation of an SU(3) gauge theory with dynamical quarks". In: *Phys. Rev. D* 107.9 (2023), p. 094513. DOI: 10.1103/PhysRevD.107.094513. arXiv: 2212.04490 [hep-lat].

[79] Saurabh V. Kadam et al. "Loop-string-hadron approach to SU(3) lattice Yang-Mills theory: Gauge invariant Hilbert space of a trivalent vertex". In: (July 2024). arXiv: 2407.19181 [hep-lat].

[80] Torsten V. Zache, Daniel González-Cuadra, and Peter Zoller. "Quantum and Classical Spin-Network Algorithms for q-Deformed Kogut-Susskind Gauge Theories". In: *Phys. Rev. Lett.* 131.17 (2023), p. 171902. DOI: 10.1103/PhysRevLett.131.171902. arXiv: 2304.02527 [quant-ph].

[81] Anthony N. Ciavarella, I. M. Burbano, and Christian W. Bauer. "Efficient Truncations of SU(N_c) Lattice Gauge Theory for Quantum Simulation". In: (Mar. 2025). arXiv: 2503.11888 [hep-lat].

[82] Henry Lamm, Scott Lawrence, and Yukari Yamauchi. "General Methods for Digital Quantum Simulation of Gauge Theories". In: *Phys. Rev. D* 100.3 (2019), p. 034518. DOI: 10.1103/PhysRevD.100.034518. arXiv: 1903.08807 [hep-lat].

[83] Andrei Alexandru et al. "Gluon Field Digitization for Quantum Computers". In: *Phys. Rev. D* 100.11 (2019), p. 114501. DOI: 10.1103/PhysRevD.100.114501. arXiv: 1906.11213 [hep-lat].

[84] Erik J. Gustafson et al. "Primitive quantum gates for an SU(2) discrete subgroup: Binary tetrahedral". In: *Phys. Rev. D* 106.11 (2022), p. 114501. DOI: 10.1103/PhysRevD.106.114501. arXiv: 2208.12309 [quant-ph].

[85] Erik J. Gustafson, Henry Lamm, and Felicity Lovelace. "Primitive quantum gates for an SU(2) discrete subgroup: Binary octahedral". In: *Phys. Rev. D* 109.5 (2024), p. 054503. DOI: 10.1103/PhysRevD.109.054503. arXiv: 2312.10285 [hep-lat].

[86] Erik J. Gustafson et al. "Primitive quantum gates for an SU(3) discrete subgroup: $\Sigma(36\times3)$". In: *Phys. Rev. D* 110.3 (2024), p. 034515. DOI: 10.1103/PhysRevD.110.034515. arXiv: 2405.05973 [hep-lat].

[87] Benoît Assi and Henry Lamm. "Digitization and subduction of SU(N) gauge theories". In: *Phys. Rev. D* 110.7 (2024), p. 074511. DOI: 10.1103/PhysRevD.110.074511. arXiv: 2405.12204 [hep-lat].

[88] R. Brower, S. Chandrasekharan, and U. J. Wiese. "QCD as a quantum link model". In: *Phys. Rev. D* 60 (1999), p. 094502. DOI: 10.1103/PhysRevD.60.094502. arXiv: hep-th/9704106.

[89] R. Brower et al. "D theory: Field quantization by dimensional reduction of discrete variables". In: *Nucl. Phys. B* 693 (2004), pp. 149–175. DOI: 10.1016/j.nuclphysb.2004.06.007. arXiv: hep-lat/0309182.

[90] Jan F. Haase et al. "A resource efficient approach for quantum and classical simulations of gauge theories in particle physics". In: *Quantum* 5 (2021), p. 393. DOI: 10.22331/q-2021-02-04-393. arXiv: 2006.14160 [quant-ph].

[91] Christian W. Bauer and Dorota M. Grabowska. "Efficient representation for simulating U(1) gauge theories on digital quantum computers at all values of the coupling". In: *Phys. Rev. D* 107.3 (2023), p. L031503. DOI: 10.1103/PhysRevD.107.L031503. arXiv: 2111.08015 [hep-ph].

[92] Timo Jakobs et al. "Canonical momenta in digitized Su(2) lattice gauge theory: definition and free theory". In: *Eur. Phys. J. C* 83.7 (2023), p. 669. DOI: 10.1140/epjc/s10052-023-11829-9. arXiv: 2304.02322 [hep-lat].

[93] Irian D'Andrea et al. "New basis for Hamiltonian SU(2) simulations". In: *Phys. Rev. D* 109.7 (2024), p. 074501. DOI: 10.1103/PhysRevD.109.074501. arXiv: 2307.11829 [hep-ph].

[94] Dorota M. Grabowska, Christopher F. Kane, and Christian W. Bauer. *A Fully Gauge-Fixed SU(2) Hamiltonian for Quantum Simulations*. Sept. 2024. arXiv: 2409.10610 [quant-ph].

[95] Georg Bergner et al. "Toward QCD on quantum computer: orbifold lattice approach". In: *JHEP* 05 (2024), p. 234. DOI: 10.1007/JHEP05(2024)234. arXiv: 2401.12045 [hep-th].

[96] Jad C. Halimeh et al. *A universal framework for the quantum simulation of Yang-Mills theory*. Nov. 2024. arXiv: 2411.13161 [quant-ph].

[97] Natalie Klco and Martin J. Savage. "Digitization of scalar fields for quantum computing". In: *Phys. Rev. A* 99.5 (2019), p. 052335. DOI: 10.1103/PhysRevA.99.052335. arXiv: 1808.10378 [quant-ph].

[98] Alexandru Macridin et al. "Electron-Phonon Systems on a Universal Quantum Computer". In: *Phys. Rev. Lett.* 121.11 (2018), p. 110504. DOI: 10.1103/PhysRevLett.121.110504. arXiv: 1802.07347 [quant-ph].

[99] Alexandru Macridin et al. "Digital quantum computation of fermion-boson interacting systems". In: *Phys. Rev. A* 98.4 (2018), p. 042312. DOI: 10.1103/PhysRevA.98.042312. arXiv: 1805.09928 [quant-ph].

[100] Alexandru Macridin et al. "Bosonic field digitization for quantum computers". In: *Phys. Rev. A* 105.5 (2022), p. 052405. DOI: 10.1103/PhysRevA.105.052405. arXiv: 2108.10793 [quant-ph].

[101] Christopher F. Kane, Niladri Gomes, and Michael Kreshchuk. "Nearly-optimal state preparation for quantum simulations of lattice gauge theories". In: (Oct. 2023). arXiv: 2310.13757 [quant-ph].

[102] Joel E Cohen et al. "Eigenvalue inequalities for products of matrix exponentials". In: *Linear Algebra and its Applications* 45 (1982), pp. 55–95.

[103] Masuo Suzuki. "Generalized Trotter's formula and systemaftic approximants of exponential operators and inner derivations with applications to many-body problems". In: *Communications in Mathematical Physics* 51.2 (1976), pp. 183–190.

[104] Masuo Suzuki. "General theory of fractal path integrals with applications to many-body theories and statistical physics". In: *J. Math. Phys.* 32.2 (1991), p. 400. DOI: 10.1063/1.529425.

[105] Andrew M Childs and Yuan Su. "Nearly optimal lattice simulation by product formulas". In: *Physical review letters* 123.5 (2019), p. 050503.

[106] Andrew M. Childs et al. "Theory of Trotter Error with Commutator Scaling". In: *Phys. Rev. X* 11.1 (2021), p. 011020. DOI: 10.1103/physrevx.11.011020. arXiv: 1912.08854 [quant-ph].

[107] Naomichi Hatano and Masuo Suzuki. "Finding Exponential Product Formulas of Higher Orders". In: *Lect. Notes Phys.* 679 (2005), p. 37. DOI: 10.1007/11526216_2. arXiv: math-ph/0506007.

[108] Guang Hao Low and Isaac L. Chuang. "Optimal Hamiltonian Simulation by Quantum Signal Processing". In: *Phys. Rev. Lett.* 118.1 (2017), p. 010501. DOI: 10.1103/PhysRevLett.118.010501. arXiv: 1606.02685 [quant-ph].

[109] Guang Hao Low and Isaac L. Chuang. "Hamiltonian Simulation by Qubitization". In: *Quantum* 3 (2019), p. 163. DOI: 10.22331/q-2019-07-12-163. arXiv: 1610.06546 [quant-ph].

[110] Danial Motlagh and Nathan Wiebe. "Generalized Quantum Signal Processing". In: (Aug. 2023). arXiv: 2308.01501 [quant-ph].

[111] Andrew M. Childs and Nathan Wiebe. "Hamiltonian Simulation Using Linear Combinations of Unitary Operations". In: *Quant. Inf. Comput.* 12 (2012), pp. 0901–0924. DOI: 10.26421/QIC12.11-12-1. arXiv: 1202.5822 [quant-ph].

[112] Dominic W. Berry et al. "Efficient Quantum Algorithms for Simulating Sparse Hamiltonians". In: *Commun. Math. Phys.* 270.2 (2007), pp. 359–371. DOI: 10.1007/s00220-006-0150-x. arXiv: quant-ph/0508139.

[113] Marcela Carena et al. "Lattice renormalization of quantum simulations". In: *Phys. Rev. D* 104.9 (2021), p. 094519. DOI: 10.1103/PhysRevD.104.094519. arXiv: 2107.01166 [hep-lat].

[114] Siddharth Hariprakash et al. "Strategies for simulating the time evolution of Hamiltonian lattice field theories". In: *Phys. Rev. A* 111.2 (2025), p. 022419. DOI: 10.1103/PhysRevA.111.022419. arXiv: 2312.11637 [quant-ph].

[115] Christopher F. Kane et al. "Block encoding bosons by signal processing". In: *Quantum* 9 (2025), p. 1747. DOI: 10.22331/q-2025-05-15-1747. arXiv: 2408.16824 [quant-ph].

[116] András Gilyén et al. "Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics". In: *51st Annual ACM SIGACT Symposium on Theory of Computing*. June 2018. DOI: 10.1145/3313276.3316366. arXiv: 1806.01838 [quant-ph].

[117] Lin Lin. "Lecture Notes on Quantum Algorithms for Scientific Computation". In: (Jan. 2022). arXiv: 2201.08309 [quant-ph].

[118] Yulong Dong, Lin Lin, and Yu Tong. "Ground-State Preparation and Energy Estimation on Early Fault-Tolerant Quantum Computers via Quantum Eigenvalue Transformation of Unitary Matrices". In: *PRX Quantum* 3.4 (2022), p. 040305. DOI: 10.1103/PRXQuantum.3.040305. arXiv: 2204.05955 [quant-ph].

[119] Anthony N. Ciavarella et al. "Truncation uncertainties for accurate quantum simulations of lattice gauge theories". In: (July 2025). arXiv: 2508.00061 [quant-ph].

[120] Pablo Sala et al. "Ergodicity-breaking arising from Hilbert space fragmentation in dipole-conserving Hamiltonians". In: *Phys. Rev. X* 10.1 (2020), p. 011047. DOI: 10.1103/PhysRevX.10.011047. arXiv: 1904.04266.

[121] Vedika Khemani, Michael Hermele, and Rahul Nandkishore. "Localization from Hilbert space shattering: From theory to physical realizations". In: *Phys. Rev. B* 101.17 (2020), p. 174204. DOI: 10.1103/PhysRevB.101.174204. arXiv: 1904.04815.

[122] Anthony N. Ciavarella, Christian W. Bauer, and Jad C. Halimeh. "Generic Hilbert Space Fragmentation in Kogut–Susskind Lattice Gauge Theories". In: (Feb. 2025). arXiv: 2502.03533 [quant-ph].

[123] Christopher F. Kane, Siddharth Hariprakash, and Christian W. Bauer. "Obtaining continuum physics from dynamical simulations of Hamiltonian lattice gauge theories". In: (June 2025). arXiv: 2506.16559 [hep-lat].

[124] Christof Zalka. "Simulating quantum systems on a quantum computer". In: *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 454.1969 (1998), pp. 313–322.

[125] Stephen Wiesner. "Simulations of many body quantum systems by a quantum computer". In: (Mar. 1996). arXiv: quant-ph/9603028.

[126] Sina Zeytinoğlu and Sho Sugiura. "Error-robust quantum signal processing using Rydberg atoms". In: *Phys. Rev. Res.* 6.1 (2024), p. 013003. DOI: 10.1103/PhysRevResearch.6.013003. arXiv: 2201.04665 [quant-ph].

[127] Yuta Kikuchi et al. "Realization of quantum signal processing on a noisy quantum computer". In: *npj Quantum Inf.* 9 (2023), p. 93. DOI: 10.1038/s41534-023-00762-0. arXiv: 2303.05533 [quant-ph].

[128] Jeongwan Haah et al. "Quantum algorithm for simulating real time evolution of lattice Hamiltonians". In: (Jan. 2018). DOI: 10.1137/18M1231511. arXiv: 1801.03922 [quant-ph].

[129] Borzu Toloui and Peter J. Love. "Quantum Algorithms for Quantum Chemistry based on the sparsity of the CI-matrix". In: (Dec. 2013). arXiv: 1312.2579 [quant-ph].

[130] Ryan Babbush et al. "Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity". In: *Phys. Rev. X* 8.4 (2018), p. 041015. DOI: 10.1103/PhysRevX.8.041015.

[131] Michael Kreshchuk et al. "Quantum simulation of quantum field theory in the light-front formulation". In: *Phys. Rev. A* 105.3 (2022), p. 032418. DOI: 10.1103/PhysRevA.105.032418. arXiv: 2002.04016 [quant-ph].

[132] William M. Kirby et al. "Quantum simulation of second-quantized Hamiltonians in compact encoding". In: *Phys. Rev. A* 104.4 (2021), p. 042607. DOI: 10.1103/PhysRevA.104.042607. arXiv: 2105.10941 [quant-ph].

[133] Mason L. Rhodes, Michael Kreshchuk, and Shivesh Pathak. "Exponential Improvements in the Simulation of Lattice Gauge Theories Using Near-Optimal Techniques". In: *PRX Quantum* 5.4 (2024), p. 040347. DOI: 10.1103/PRXQuantum.5.040347. arXiv: 2405.10416 [quant-ph].

[134] Yao Ji, Henry Lamm, and Shuchen Zhu. "Gluon Field Digitization via Group Space Decimation for Quantum Computers". In: *Phys. Rev. D* 102.11 (2020), p. 114513. DOI: 10.1103/PhysRevD.102.114513. arXiv: 2005.14221 [hep-lat].

[135] Dorota M. Grabowska et al. "Overcoming exponential scaling with system size in Trotter-Suzuki implementations of constrained Hamiltonians: 2+1 U(1) lattice gauge theories". In: (Aug. 2022). arXiv: 2208.03333 [quant-ph].

[136] Christopher Kane et al. "Efficient quantum implementation of 2+1 U(1) lattice gauge theories with Gauss law constraints". In: (Nov. 2022). arXiv: 2211.10497 [quant-ph].

[137] M. B. Hastings et al. "Improving Quantum Algorithms for Quantum Chemistry". In: (Mar. 2014). arXiv: 1403.1539 [quant-ph].

[138] David Poulin et al. "The Trotter step size required for accurate quantum simulation of quantum chemistry". In: *Quant. Inf. Comput.* 15.5-6 (2015), pp. 0361–0384. DOI: 10.26421/QIC15.5-6-1. arXiv: 1406.4920 [quant-ph].

[139] Andrew Tranter et al. "A comparison of the Bravyi–Kitaev and Jordan–Wigner transformations for the quantum simulation of quantum chemistry". In: *Journal of chemical theory and computation* 14.11 (2018), pp. 5617–5630.

[140] Yi Hu et al. "Greedy algorithm based circuit optimization for near-term quantum simulation". In: *Quantum Sci. Technol.* 7.4 (2022), p. 045001. DOI: 10.1088/2058-9565/ac796b.

[141] Andrew Tranter et al. "Ordering of trotterization: Impact on errors in quantum simulation of electronic structure". In: *Entropy* 21.12 (2019), p. 1218.

[142] Albert T. Schmitz et al. "Graph Optimization Perspective for Low-Depth Trotter-Suzuki Decomposition". In: (Mar. 2021). arXiv: 2103.08602 [quant-ph].

[143] Johann Ostmeyer. "Optimised Trotter decompositions for classical and quantum computing". In: *J. Phys. A* 56.28 (2023), p. 285303. DOI: 10.1088/1751-8121/acde7a. arXiv: 2211.02691 [quant-ph].

[144] Xin-Yu Huang et al. *Qubitization of Bosons*. May 2021. arXiv: 2105.12563 [quant-ph].

[145] Andrei Alexandru et al. "Qubitization strategies for bosonic field theories". In: *Phys. Rev. D* 107.3 (2023), p. 034503. DOI: 10.1103/PhysRevD.107.034503. arXiv: 2209.00098 [hep-lat].

[146] Ryan Babbush et al. "Quantum simulation of chemistry with sublinear scaling in basis size". In: *npj Quantum Inf.* 5 (2019), p. 92. DOI: 10.1038/s41534-019-0199-y.

[147] Daan Camps and Roel Van Beeumen. "Fable: Fast approximate quantum circuits for block-encodings". In: *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE. 2022, pp. 104–113.

[148] Ryan Babbush et al. "Exponentially more precise quantum simulation of fermions in the configuration interaction representation". In: *Quantum Sci. Technol.* 3.1 (2017), p. 015006. DOI: 10.1088/2058-9565/aa9463.

[149] V. V. Shende, S. S. Bullock, and I. L. Markov. "Synthesis of quantum-logic circuits". In: *IEEE Trans. Comput. Aided Design Integr. Circuits Syst.* 25.6 (2006), pp. 1000–1010. DOI: 10.1109/tcad.2005.855930.

[150] Martin Plesch and Časlav Brukner. "Quantum-state preparation with universal gate decompositions". In: *Phys. Rev. A* 83.3 (2011), p. 032302. DOI: 10.1103/PhysRevA.83.032302.

[151] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, June 2012. DOI: 10.1017/cbo9780511976667.

[152] William Kirby, Mario Motta, and Antonio Mezzacapo. "Exact and efficient Lanczos method on a quantum computer". In: *Quantum* 7 (2023), p. 1018. DOI: 10.22331/q-2023-05-23-1018. arXiv: 2208.00567 [quant-ph].

[153] I. Novikau, E. A. Startsev, and I. Y. Dodin. "Quantum signal processing for simulating cold plasma waves". In: *Phys. Rev. A* 105.6 (2022), p. 062444. DOI: 10.1103/PhysRevA.105.062444. arXiv: 2112.06086 [physics.plasm-ph].

[154] Guang Hao Low, Theodore J. Yoder, and Isaac L. Chuang. "Methodology of Resonant Equiangular Composite Quantum Gates". In: *Phys. Rev. X* 6.4 (2016), p. 041067. DOI: 10.1103/PhysRevX.6.041067.

[155] E. H. Lieb and D. W. Robinson. "The finite group velocity of quantum spin systems". In: *Commun. Math. Phys.* 28 (1972), pp. 251–257. DOI: 10.1007/BF01645779.

[156] Stephen P. Jordan et al. "BQP-completeness of Scattering in Scalar Quantum Field Theory". In: *Quantum* 2 (2018), p. 44. DOI: 10.22331/q-2018-01-08-44. arXiv: 1703.00454 [quant-ph].

[157] Junyu Liu et al. "Towards a variational Jordan–Lee–Preskill quantum algorithm". In: *Mach. Learn. Sci. Tech.* 3.4 (2022), p. 045030. DOI: 10.1088/2632-2153/aca06b. arXiv: 2109.05547 [quant-ph].

[158] Terry Farrelly and Julien Streich. "Discretizing quantum field theories for quantum simulation". In: (Feb. 2020). arXiv: 2002.02643 [quant-ph].

[159] Christian W. Bauer, Marat Freytsis, and Benjamin Nachman. "Simulating Collider Physics on Quantum Computers Using Effective Field Theories". In: *Phys. Rev. Lett.* 127.21 (2021), p. 212001. DOI: 10.1103/PhysRevLett.127.212001. arXiv: 2102.05044 [hep-ph].

[160] Gavin K. Brennen et al. "Multiscale quantum simulation of quantum field theory using wavelets". In: *Phys. Rev. A* 92.3 (2015), p. 032315. DOI: 10.1103/PhysRevA.92.032315. arXiv: 1412.0750 [quant-ph].

[161] J. P. Vary et al. "Hamiltonian light-front field theory in a basis function approach". In: *Phys. Rev. C* 81 (2010), p. 035205. DOI: 10.1103/PhysRevC.81.035205. arXiv: 0905.1411 [nucl-th].

[162] Guang Hao Low and Nathan Wiebe. "Hamiltonian Simulation in the Interaction Picture". In: (May 2018). arXiv: 1805.00675 [quant-ph].

[163] Mária Kieferová, Artur Scherer, and Dominic W. Berry. "Simulating the dynamics of time-dependent Hamiltonians with a truncated Dyson series". In: *Phys. Rev. A* 99 (4 Apr. 2019), p. 042314. DOI: 10.1103/PhysRevA.99.042314. URL: https://link.aps.org/doi/10.1103/PhysRevA.99.042314.

[164] Lin Lin and Yu Tong. "Heisenberg-Limited Ground-State Energy Estimation for Early Fault-Tolerant Quantum Computers". In: *PRX Quantum* 3.1 (2022), p. 010318. DOI: 10.1103/PRXQuantum.3.010318. arXiv: 2102.11340 [quant-ph].

[165] Dominic W. Berry et al. "Simulating Hamiltonian Dynamics with a Truncated Taylor Series". In: *Phys. Rev. Lett.* 114.9 (2015), p. 090502. DOI: 10.1103/PhysRevLett.114.090502. arXiv: 1412.4687 [quant-ph].

[166] Anthony N. Ciavarella and Christian W. Bauer. "Quantum Simulation of SU(3) Lattice Yang Mills Theory at Leading Order in Large N". In: (Feb. 2024). arXiv: 2402.10265 [hep-ph].

[167] Sergey B. Bravyi and Alexei Yu. Kitaev. "Fermionic Quantum Computation". In: *Annals Phys.* 298 (2002), pp. 210–226. DOI: 10.1006/aphy.2002.6254. arXiv: quant-ph/0003137.

[168] F. Verstraete and J. I. Cirac. "Mapping local Hamiltonians of fermions to local Hamiltonians of spins". In: *J. Stat. Mech.* 0509 (2005), P09012. DOI: 10.1088/1742-5468/2005/09/P09012. arXiv: cond-mat/0508353.

[169] Andrew Tranter et al. "The B ravyi–K itaev transformation: Properties and applications". In: *International Journal of Quantum Chemistry* 115.19 (2015), pp. 1431–1441.

[170] Ryan Babbush et al. "Chemical basis of Trotter-Suzuki errors in quantum chemistry simulation". In: *Phys. Rev. A* 91.2 (2015), p. 022311. DOI: 10.1103/PhysRevA.91.022311.

[171] Alán Aspuru-Guzik et al. "Exponentially more precise quantum simulation of fermions in second quantization". In: *New J. Phys.* 18.3 (2016), p. 033032. DOI: 10.1088/1367-2630/18/3/033032.

[172] Kanav Setia and James D. Whitfield. "Bravyi-Kitaev Superfast simulation of electronic structure on a quantum computer". In: *J. Chem. Phys.* 148.16 (2018), p. 164104. DOI: 10.1063/1.5019371.

[173] William Kirby et al. "Second-Quantized Fermionic Operators with Polylogarithmic Qubit and Gate Complexity". In: *PRX Quantum* 3.2 (2022), p. 020351. DOI: 10.1103/PRXQuantum.3.020351. arXiv: 2109.14465 [quant-ph].

[174] Kazuki Sakamoto et al. "End-to-end complexity for simulating the Schwinger model on quantum computers". In: *Quantum* 8 (2024), p. 1474. DOI: 10.22331/q-2024-09-17-1474. arXiv: 2311.17388 [quant-ph].

[175] James D. Watson et al. "Quantum Algorithms for Simulating Nuclear Effective Field Theories". In: (Dec. 2023). arXiv: 2312.05344 [quant-ph].

[176] Lewis W. Anderson et al. "Solving lattice gauge theories using the quantum Krylov algorithm and qubitization". In: *Quantum* 9 (2025), p. 1669. DOI: 10.22331/q-2025-03-25-1669. arXiv: 2403.08859 [quant-ph].

[177] Weijie Du and James P. Vary. "Systematic input scheme for many-boson Hamiltonians via quantum walk". In: (July 2024). arXiv: 2407.13672 [quant-ph].

[178] Hsuan-Hao Lu et al. "Simulations of Subatomic Many-Body Physics on a Quantum Frequency Processor". In: *Phys. Rev. A* 100.1 (2019), p. 012320. DOI: 10.1103/PhysRevA.100.012320. arXiv: 1810.03959 [quant-ph].

[179] So Chigusa and Masahito Yamazaki. "Quantum simulations of dark sector showers". In: *Phys. Lett. B* 834 (2022), p. 137466. DOI: 10.1016/j.physletb.2022.137466. arXiv: 2204.12500 [hep-ph].

[180] N. Klco et al. "Quantum-classical computation of Schwinger model dynamics using quantum computers". In: *Phys. Rev. A* 98.3 (2018), p. 032331. DOI: 10.1103/PhysRevA.98.032331. arXiv: 1803.03326 [quant-ph].

[181] João Barata et al. "Single-particle digitization strategy for quantum computation of a ϕ^4 scalar field theory". In: *Phys. Rev. A* 103.4 (2021), p. 042410. DOI: 10.1103/PhysRevA.103.042410. arXiv: 2012.00020 [hep-th].

[182] Bo Peng et al. "Quantum simulation of boson-related Hamiltonians: techniques, effective Hamiltonian construction, and error analysis". In: *Quantum Sci. Technol.* 10.2 (2025), p. 023002. DOI: 10.1088/2058-9565/adbf42. arXiv: 2307.06580 [quant-ph].

[183] Lento Nagano, Aniruddha Bapat, and Christian W. Bauer. "Quench dynamics of the Schwinger model via variational quantum algorithms". In: *Phys. Rev. D* 108.3 (2023), p. 034501. DOI: 10.1103/PhysRevD.108.034501. arXiv: 2302.10933 [hep-ph].

[184] Niladri Gomes, Hokiat Lim, and Nathan Wiebe. "Multivariable QSP and Bosonic Quantum Simulation using Iterated Quantum Signal Processing". In: (Aug. 2024). arXiv: 2408.03254 [quant-ph].

[185] Dominic W. Berry et al. "Doubling the efficiency of Hamiltonian simulation via generalized quantum signal processing". In: *Phys. Rev. A* 110.1 (2024), p. 012612. DOI: 10.1103/PhysRevA.110.012612. arXiv: 2401.10321 [quant-ph].

[186] *QSPPACK*. https://github.com/qsppack/QSPPACK.

[187] Martin J. Savage. "Quantum computing for nuclear physics". In: *EPJ Web Conf.* 296 (2024), p. 01025. DOI: 10.1051/epjconf/202429601025. arXiv: 2312.07617 [nucl-th].

[188] Julian Bender and Erez Zohar. "Gauge redundancy-free formulation of compact QED with dynamical matter for quantum and classical computations". In: *Phys. Rev. D* 102.11 (2020), p. 114517. DOI: 10.1103/PhysRevD.102.114517. arXiv: 2008.01349 [quant-ph].

[189] Jonathan Welch et al. "Efficient quantum circuits for diagonal unitaries without ancillas". In: *New Journal of Physics* 16.3 (Mar. 2014), p. 033040. DOI: 10.1088/1367-2630/16/3/033040. URL: https://dx.doi.org/10.1088/1367-2630/16/3/033040.

[190] Lloyd N. Trefethen. *Approximation Theory and Approximation Practice, Extended Edition*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2019. DOI: 10.1137/1.9781611975949. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611975949. URL: https://epubs.siam.org/doi/abs/10.1137/1.9781611975949.

[191] Andrew Hardy et al. "Optimized Quantum Simulation Algorithms for Scalar Quantum Field Theories". In: (July 2024). arXiv: 2407.13819 [quant-ph].

[192] Tyler Kharazi et al. "Explicit block encodings of boundary value problems for many-body elliptic operators". In: (July 2024). arXiv: 2407.18347 [quant-ph].

[193] Yutaro Akahoshi et al. "Partially Fault-Tolerant Quantum Computing Architecture with Error-Corrected Clifford Gates and Space-Time Efficient Analog Rotations". In: *PRX Quantum* 5 (1 Mar. 2024), p. 010337. DOI: 10.1103/PRXQuantum.5.010337. URL: https://link.aps.org/doi/10.1103/PRXQuantum.5.010337.

[194] Matthew Amy et al. "A Meet-in-the-Middle Algorithm for Fast Synthesis of Depth-Optimal Quantum Circuits". In: 32.6 (June 2013), pp. 818–830. ISSN: 0278-0070. DOI: 10.1109/TCAD.2013.2244643. URL: https://doi.org/10.1109/TCAD.2013.2244643.

[195] Gregory Boyd. "Low-Overhead Parallelisation of LCU via Commuting Operators". In: (Dec. 2023). arXiv: 2312.00696 [quant-ph].

[196] Pei Zeng et al. "Simple and High-Precision Hamiltonian Simulation by Compensating Trotter Error with Linear Combination of Unitary Operations". In: *PRX Quantum* 6.1 (2025), p. 010359. DOI: 10.1103/PRXQuantum.6.010359. arXiv: 2212.04566 [quant-ph].

[197] Shantanav Chakraborty. "Implementing any Linear Combination of Unitaries on Intermediate-term Quantum Computers". In: *Quantum* 8 (2024), p. 1496. DOI: 10.22331/q-2024-10-10-1496. arXiv: 2302.13555 [quant-ph].

[198] Paul K. Faehrmann et al. "Randomizing multi-product formulas for Hamiltonian simulation". In: *Quantum* 6 (2022), p. 806. DOI: 10.22331/q-2022-09-19-806. arXiv: 2101.07808 [quant-ph].

[199] Shantanav Chakraborty et al. "Quantum singular value transformation without block encodings: Near-optimal complexity with minimal ancilla". In: (Apr. 2025). arXiv: 2504.02385 [quant-ph].

[200] Zane M. Rossi and Isaac L. Chuang. "Multivariable quantum signal processing (M-QSP): prophecies of the two-headed oracle". In: *Quantum* 6 (2022), p. 811. DOI: 10.22331/q-2022-09-20-811. arXiv: 2205.06261 [quant-ph].

[201] Zane M. Rossi, Jack L. Ceroni, and Isaac L. Chuang. "Modular quantum signal processing in many variables". In: (Sept. 2023). arXiv: 2309.16665 [quant-ph].

[202] Hitomi Mori, Kaoru Mizuta, and Keisuke Fujii. "Comment on "Multivariable quantum signal processing (M-QSP): prophecies of the two-headed oracle"". In: *Quantum* 8 (2024), p. 1512. DOI: 10.22331/q-2024-10-29-1512. arXiv: 2310.00918 [quant-ph].

[203] Balázs Németh et al. "On variants of multivariate quantum signal processing and their characterizations". In: (Dec. 2023). arXiv: 2312.09072 [quant-ph].

[204] Lorenzo Laneve and Stefan Wolf. "On multivariate polynomials achievable with quantum signal processing". In: *Quantum* 9 (2025), p. 1641. DOI: 10.22331/q-2025-02-20-1641. arXiv: 2407.20823 [quant-ph].

[205] Frederik Görg et al. "Realization of density-dependent Peierls phases to engineer quantized gauge fields coupled to ultracold matter". In: *Nature Phys.* 15.11 (2019), pp. 1161–1167. DOI: 10.1038/s41567-019-0615-4. arXiv: 1812.05895 [cond-mat.quant-gas].

[206] Alexander Mil et al. "A scalable realization of local U(1) gauge invariance in cold atomic mixtures". In: *Science* 367.6482 (2020), pp. 1128–1130. DOI: 10.1126/science.aaz5312. arXiv: 1909.07641 [cond-mat.quant-gas].

[207] Bing Yang et al. "Observation of gauge invariance in a 71-site Bose–Hubbard quantum simulator". In: *Nature* 587.7834 (2020), pp. 392–396. DOI: 10.1038/s41586-020-2910-8. arXiv: 2003.08945 [cond-mat.quant-gas].

[208] Zhao-Yu Zhou et al. "Thermalization dynamics of a gauge theory on a quantum simulator". In: *Science* 377.6603 (2022), abl6277. DOI: 10.1126/science.abl6277. arXiv: 2107.13563 [cond-mat.quant-gas].

[209] Zhan Wang et al. "Observation of emergent Z2 gauge invariance in a superconducting circuit". In: *Phys. Rev. Res.* 4.2 (2022), p. L022060. DOI: 10.1103/PhysRevResearch.4.L022060. arXiv: 2111.05048 [quant-ph].

[210] Nhung H. Nguyen et al. "Digital Quantum Simulation of the Schwinger Model and Symmetry Protection with Trapped Ions". In: *PRX Quantum* 3.2 (2022), p. 020324. DOI: 10.1103/PRXQuantum.3.020324. arXiv: 2112.14262 [quant-ph].

[211] Guo-Xian Su et al. "Observation of many-body scarring in a Bose-Hubbard quantum simulator". In: *Phys. Rev. Res.* 5.2 (2023), p. 023010. DOI: 10.1103/PhysRevResearch.5.023010. arXiv: 2201.00821 [cond-mat.quant-gas].

[212] Han-Yi Wang et al. "Interrelated Thermalization and Quantum Criticality in a Lattice Gauge Simulator". In: *Phys. Rev. Lett.* 131.5 (2023), p. 050401. DOI: 10.1103/PhysRevLett.131.050401. arXiv: 2210.17032 [cond-mat.quant-gas].

[213] Wei-Yong Zhang et al. "Observation of microscopic confinement dynamics by a tunable topological θ-angle". In: *Nature Phys.* 21.1 (2025), pp. 155–160. DOI: 10.1038/s41567-024-02702-x. arXiv: 2306.11794 [cond-mat.quant-gas].

[214] Takis Angelides et al. "First-order phase transition of the Schwinger model with a quantum computer". In: *npj Quantum Inf.* 11.1 (2025), p. 6. DOI: 10.1038/s41534-024-00950-6. arXiv: 2312.12831 [hep-lat].

[215] Julius Mildenberger et al. "Confinement in a $\mathbb{Z}_2$ lattice gauge theory on a quantum computer". In: *Nature Phys.* 21.2 (2025), pp. 312–317. DOI: 10.1038/s41567-024-02723-6. arXiv: 2203.08905 [quant-ph].

[216] Clement Charles et al. "Simulating Z2 lattice gauge theory on a quantum computer". In: *Phys. Rev. E* 109.1 (2024), p. 015307. DOI: 10.1103/PhysRevE.109.015307. arXiv: 2305.02361 [hep-lat].

[217] Niklas Mueller et al. "Quantum Computation of Dynamical Quantum Phase Transitions and Entanglement Tomography in a Lattice Gauge Theory". In: *PRX Quantum* 4.3 (2023), p. 030323. DOI: 10.1103/PRXQuantum.4.030323. arXiv: 2210.03089 [quant-ph].

[218] Arinjoy De et al. *Observation of string-breaking dynamics in a quantum simulator*. Oct. 2024. arXiv: 2410.13815 [quant-ph].

[219] Zohreh Davoudi, Chung-Chun Hsieh, and Saurabh V. Kadam. "Scattering wave packets of hadrons in gauge theories: Preparation on a quantum computer". In: *Quantum* 8 (2024), p. 1520. DOI: 10.22331/q-2024-11-11-1520. arXiv: 2402.00840 [quant-ph].

[220] Yibin Guo et al. *Concurrent VQE for Simulating Excited States of the Schwinger Model*. July 2024. arXiv: 2407.15629 [quant-ph].

[221] Yasar Y. Atas et al. "SU(2) hadrons on a quantum computer via a variational approach". In: *Nature Commun.* 12.1 (2021), p. 6499. DOI: 10.1038/s41467-021-26825-4. arXiv: 2102.08920 [quant-ph].

[222] Yasar Y. Atas et al. "Simulating one-dimensional quantum chromodynamics on a quantum computer: Real-time evolutions of tetra- and pentaquarks". In: *Phys. Rev. Res.* 5.3 (2023), p. 033184. DOI: 10.1103/PhysRevResearch.5.033184. arXiv: 2207.03473 [quant-ph].

[223] Anton T. Than et al. *The phase diagram of quantum chromodynamics in one dimension on a quantum computer*. Dec. 2024. arXiv: 2501.00579 [quant-ph].

[224] Anthony N. Ciavarella. "Quantum simulation of lattice QCD with improved Hamiltonians". In: *Phys. Rev. D* 108.9 (2023), p. 094513. DOI: 10.1103/PhysRevD.108.094513. arXiv: 2307.05593 [hep-lat].

[225] Anthony N. Ciavarella. "String breaking in the heavy quark limit with scalable circuits". In: *Phys. Rev. D* 111.5 (2025), p. 054501. DOI: 10.1103/PhysRevD.111.054501. arXiv: 2411.05915 [quant-ph].

[226] Jiunn-Wei Chen, Yu-Ting Chen, and Ghanashyam Meher. "Parton Distributions on a Quantum Computer". In: (June 2025). arXiv: 2506.16829 [hep-lat].

[227] Zi-Hang Zhu et al. "Probing false vacuum decay on a cold-atom gauge-theory quantum simulator". In: (Nov. 2024). arXiv: 2411.12565 [cond-mat.quant-gas].

[228] Gaurav Gyawali et al. "Observation of disorder-free localization using a (2+1)D lattice gauge theory on a quantum processor". In: (Oct. 2024). arXiv: 2410.06557 [quant-ph].

[229] Julian Schuhmacher et al. "Observation of hadron scattering in a lattice gauge theory on a quantum computer". In: (May 2025). arXiv: 2505.20387 [quant-ph].

[230] Zohreh Davoudi, Chung-Chun Hsieh, and Saurabh V. Kadam. "Quantum computation of hadron scattering in a lattice gauge theory". In: (May 2025). arXiv: 2505.20408 [quant-ph].

[231] Danny Paulson et al. "Simulating 2D Effects in Lattice Gauge Theories on a Quantum Computer". In: *PRX Quantum* 2.3 (2021), p. 030334. DOI: 10.1103/PRXQuantum.2.030334. arXiv: 2008.09252 [quant-ph].

[232] Anthony N. Ciavarella and Ivan A. Chernyshev. "Preparation of the SU(3) lattice Yang-Mills vacuum with variational quantum methods". In: *Phys. Rev. D* 105.7 (2022), p. 074504. DOI: 10.1103/PhysRevD.105.074504. arXiv: 2112.09083 [quant-ph].

[233] Sarmed A Rahman et al. "Self-mitigating Trotter circuits for SU(2) lattice gauge theory on a quantum computer". In: *Phys. Rev. D* 106.7 (2022), p. 074502. DOI: 10.1103/PhysRevD.106.074502. arXiv: 2205.09247 [hep-lat].

[234] Emanuele Mendicelli et al. "Real time evolution and a traveling excitation in SU(2) pure gauge theory on a quantum computer." In: *PoS* LATTICE2022 (2023), p. 025. DOI: 10.22323/1.430.0025. arXiv: 2210.11606 [hep-lat].

[235] Jad C. Halimeh et al. *Disorder-Free Localization for Benchmarking Quantum Computers*. Oct. 2024. arXiv: 2410.08268 [quant-ph].

[236] Sabhyata Gupta et al. "Simulation of a Rohksar–Kivelson ladder on a NISQ device". In: *Sci. Rep.* 14.1 (2024), p. 29276. DOI: 10.1038/s41598-024-79480-2. arXiv: 2401.16326 [quant-ph].

[237] Arianna Crippa, Karl Jansen, and Enrico Rinaldi. *Analysis of the confinement string in (2 + 1)-dimensional Quantum Electrodynamics with a trapped-ion quantum computer*. Nov. 2024. arXiv: 2411.05628 [hep-lat].

[238] Ali H. Z. Kavaki and Randy Lewis. *False vacuum decay in triamond lattice gauge theory*. Mar. 2025. arXiv: 2503.01119 [hep-lat].

[239] Ali H. Z. Kavaki and Randy Lewis. "From square plaquettes to triamond lattices for SU(2) gauge theory". In: *Commun. Phys.* 7.1 (2024), p. 208. DOI: 10.1038/s42005-024-01697-4. arXiv: 2401.14570 [hep-lat].

[240] Erez Zohar and J. Ignacio Cirac. "Removing Staggered Fermionic Matter in $U(N)$ and $SU(N)$ Lattice Gauge Theories". In: *Phys. Rev. D* 99.11 (2019), p. 114511. DOI: 10.1103/PhysRevD.99.114511. arXiv: 1905.00652 [quant-ph].

[241] Randy Lewis and R. M. Woloshyn. *A qubit model for U(1) lattice gauge theory*. May 2019. arXiv: 1905.09789 [hep-lat].

[242] Jesse Stryker and Indrakshi Raychowdhury. "Tailoring Non-Abelian Gauge Theory for Digital Quantum Simulation". In: *PoS* LATTICE2019 (2020), p. 144. DOI: 10.22323/1.363.0144.

[243] Saurabh Vasant Kadam, Indrakshi Raychowdhury, and Jesse R. Stryker. "Loop-string-hadron formulation of an SU(3) gauge theory with dynamical quarks". In: *PoS* LATTICE2022 (2023), p. 373. DOI: 10.22323/1.430.0373.

[244] Anthony Ciavarella, Natalie Klco, and Martin J. Savage. *Some Conceptual Aspects of Operator Design for Quantum Simulations of Non-Abelian Lattice Gauge Theories*. Mar. 2022. arXiv: 2203.11988 [quant-ph].

[245] Berndt Müller and Xiaojun Yao. "Simple Hamiltonian for quantum simulation of strongly coupled (2+1)D SU(2) lattice gauge theory on a honeycomb lattice". In: *Phys. Rev. D* 108.9 (2023), p. 094505. DOI: 10.1103/PhysRevD.108.094505. arXiv: 2307.00045 [quant-ph].

[246] Tomoya Hayata and Yoshimasa Hidaka. "q deformed formulation of Hamiltonian SU(3) Yang-Mills theory". In: *JHEP* 09 (2023), p. 123. DOI: 10.1007/JHEP09(2023)123. arXiv: 2306.12324 [hep-lat].

[247] Marco Rigobello et al. *Hadrons in (1+1)D Hamiltonian hardcore lattice QCD*. Aug. 2023. arXiv: 2308.04488 [hep-lat].

[248] Pierpaolo Fontana, Marc Miranda Riaza, and Alessio Celi. *An efficient finite-resource formulation of non-Abelian lattice gauge theories beyond one dimension*. Sept. 2024. arXiv: 2409.04441 [quant-ph].

[249] Marc Illa, Martin J. Savage, and Xiaojun Yao. "Improved Honeycomb and Hyper-Honeycomb Lattice Hamiltonians for Quantum Simulations of Non-Abelian Gauge Theories". In: (Mar. 2025). arXiv: 2503.09688 [hep-lat].

[250] M. Sohaib Alam et al. "Primitive quantum gates for dihedral gauge theories". In: *Phys. Rev. D* 105.11 (2022), p. 114501. DOI: 10.1103/PhysRevD.105.114501. arXiv: 2108.13305 [quant-ph].

[251] Yao Ji, Henry Lamm, and Shuchen Zhu. "Gluon digitization via character expansion for quantum computers". In: *Phys. Rev. D* 107.11 (2023), p. 114503. DOI: 10.1103/PhysRevD.107.114503. arXiv: 2203.02330 [hep-lat].

[252] Erik J. Gustafson and Henry Lamm. *Robustness of Gauge Digitization to Quantum Noise*. Jan. 2023. arXiv: 2301.10207 [hep-lat].

[253] Tobias Hartung et al. "Digitising SU(2) gauge fields and the freezing transition". In: *Eur. Phys. J. C* 82.3 (2022), p. 237. DOI: 10.1140/epjc/s10052-022-10192-5. arXiv: 2201.09625 [hep-lat].

[254] Marco Garofalo et al. "Testing the SU(2) lattice Hamiltonian built from S_3 partitionings". In: *PoS* LATTICE2023 (2024), p. 231. DOI: 10.22323/1.453.0231. arXiv: 2311.15926 [hep-lat].

[255] I. M. Burbano and Christian W. Bauer. *Gauge Loop-String-Hadron Formulation on General Graphs and Applications to Fully Gauge Fixed Hamiltonian Lattice Gauge Theory*. Sept. 2024. arXiv: 2409.13812 [hep-lat].

[256] Timo Jakobs et al. *Dynamics in Hamiltonian Lattice Gauge Theory: Approaching the Continuum Limit with Partitionings of SU*(2). Mar. 2025. arXiv: 2503.03397 [hep-lat].

[257] Hanqing Liu et al. *Phases of 2d massless QCD with qubit regularization*. Dec. 2023. arXiv: 2312.17734 [hep-lat].

[258] Shailesh Chandrasekharan. "Qubit Regularization of Quantum Field Theories". In: *41st International Symposium on Lattice Field Theory*. Feb. 2025. arXiv: 2502.05716 [hep-lat].

[259] Shailesh Chandrasekharan, Rui Xian Siew, and Tanmoy Bhattacharya. *Monomer-dimer tensor-network basis for qubit-regularized lattice gauge theories*. Feb. 2025. arXiv: 2502.14175 [hep-lat].

[260] Andrei Alexandru et al. "Fuzzy gauge theory for quantum computers". In: *Phys. Rev. D* 109.9 (2024), p. 094502. DOI: 10.1103/PhysRevD.109.094502. arXiv: 2308.05253 [hep-lat].

[261] Marc Illa, Martin J. Savage, and Xiaojun Yao. "Dynamical Local Tadpole-Improvement in Quantum Simulations of Gauge Theories". In: (Apr. 2025). arXiv: 2504.21575 [quant-ph].

[262] Torsten V. Zache et al. "Toward the continuum limit of a 1 + 1D quantum link Schwinger model". In: *Phys. Rev. D* 106.9 (2022), p. L091502. DOI: 10.1103/PhysRevD.106.L091502. arXiv: 2104.00025 [hep-lat].

[263] Jad C. Halimeh et al. "Achieving the quantum field theory limit in far-from-equilibrium quantum link models". In: *Quantum* 6 (2022), p. 878. DOI: 10.22331/q-2022-12-19-878. arXiv: 2112.04501 [cond-mat.quant-gas].

[264] Jack Y. Araz, Sebastian Schenk, and Michael Spannowsky. "Toward a quantum simulation of nonlinear sigma models with a topological term". In: *Phys. Rev. A* 107.3 (2023), p. 032619. DOI: 10.1103/PhysRevA.107.032619. arXiv: 2210.03679 [quant-ph].

[265] Anthony N. Ciavarella et al. "Preparation for quantum simulation of the (1+1)-dimensional O(3) nonlinear σ model using cold atoms". In: *Phys. Rev. A* 107.4 (2023), p. 042404. DOI: 10.1103/PhysRevA.107.042404. arXiv: 2211.07684 [quant-ph].

[266] Xiaojun Yao et al. "Testing eigenstate thermalization hypothesis for non-Abelian gauge theories". In: *EPJ Web Conf.* 296 (2024), p. 13008. DOI: 10.1051/epjconf/202429613008. arXiv: 2312.13408 [hep-lat].

[267] Murray Gell-Mann and Francis Low. "Bound states in quantum field theory". In: *Phys. Rev.* 84 (1951), pp. 350–354. DOI: 10.1103/PhysRev.84.350.

[268] G. Vidal. "Efficient simulation of one-dimensional quantum many-body systems". In: *Phys. Rev. Lett.* 93 (2004), p. 040502. DOI: 10.1103/PhysRevLett.93.040502. arXiv: quant-ph/0310089.

[269] Giuseppe Clemente, Arianna Crippa, and Karl Jansen. "Strategies for the determination of the running coupling of (2+1)-dimensional QED with quantum computing". In: *Phys. Rev. D* 106.11 (2022), p. 114511. DOI: 10.1103/PhysRevD.106.114511. arXiv: 2206.12454 [hep-lat].

[270] Arianna Crippa et al. "Towards determining the (2+1)-dimensional Quantum Electrodynamics running coupling with Monte Carlo and quantum computing methods". In: (Apr. 2024). arXiv: 2404.17545 [hep-lat].

[271] Lena Funcke et al. "Hamiltonian limit of lattice QED in 2+1 dimensions". In: *PoS* LATTICE2022 (2023), p. 292. DOI: 10.22323/1.430.0292. arXiv: 2212.09627 [hep-lat].

[272] C. F. Groß et al. "Matching Lagrangian and Hamiltonian Simulations in (2+1)-dimensional U(1) Gauge Theory". In: (Mar. 2025). arXiv: 2503.11480 [hep-lat].

[273] Marcela Carena et al. "Gauge theory couplings on anisotropic lattices". In: *Phys. Rev. D* 106.11 (2022), p. 114504. DOI: 10.1103/PhysRevD.106.114504. arXiv: 2208.10417 [hep-lat].

[274] Earl Campbell. "Random Compiler for Fast Hamiltonian Simulation". In: *Phys. Rev. Lett.* 123.7 (2019), p. 070503. DOI: 10.1103/PhysRevLett.123.070503. arXiv: 1811.08017 [quant-ph].

[275] Ying Li and Simon C. Benjamin. "Efficient Variational Quantum Simulator Incorporating Active Error Minimization". In: *Phys. Rev. X* 7.2 (2017), p. 021050. DOI: 10.1103/physrevx.7.021050. arXiv: 1611.09301 [quant-ph].

[276] Xiao Yuan et al. "Theory of variational quantum simulation". In: *Quantum* 3 (2019), p. 191. DOI: 10.22331/q-2019-10-07-191. arXiv: 1812.08767 [quant-ph].

[277] James D. Watson and Jacob Watkins. "Exponentially Reduced Circuit Depths Using Trotter Error Mitigation". In: (Aug. 2024). arXiv: 2408.14385 [quant-ph].

[278] James D. Watson. "Randomly Compiled Quantum Simulation with Exponentially Reduced Circuit Depths". In: (Nov. 2024). arXiv: 2411.04240 [quant-ph].

[279] P. Menotti and E. Onofri. "The Action of SU(N) Lattice Gauge Theory in Terms of the Heat Kernel on the Group Manifold". In: *Nucl. Phys. B* 190 (1981), pp. 288–300. DOI: 10.1016/0550-3213(81)90560-5.

[280] William A. Simon et al. "Ladder Operator Block-Encoding". In: (Mar. 2025). arXiv: 2503.11641 [quant-ph].

[281] Dominic W. Berry, Andrew M. Childs, and Robin Kothari. "Hamiltonian Simulation with Nearly Optimal Dependence on all Parameters". In: Jan. 2015. DOI: 10.1109/FOCS.2015.54. arXiv: 1501.01715 [quant-ph].

[282] Alberto Di Meglio et al. "Quantum Computing for High-Energy Physics: State of the Art and Challenges". In: *PRX Quantum* 5.3 (2024), p. 037001. DOI: 10.1103/PRXQuantum.5.037001. arXiv: 2307.03236 [quant-ph].

[283] Takashi Umeda, Kouji Nomura, and Hideo Matsufuru. "Charmonium at finite temperature in quenched lattice QCD". In: *Eur. Phys. J. C* 39S1 (2005), pp. 9–26. DOI: 10.1140/epjcd/s2004-01-002-1. arXiv: hep-lat/0211003.

[284] Kouji Nomura, Takashi Umeda, and Hideo Matsufuru. "Numerical study of staggered fermion on anisotropic lattices". In: *Nucl. Phys. B Proc. Suppl.* 129 (2004). Ed. by S. Aoki et al., pp. 390–392. DOI: 10.1016/S0920-5632(03)02591-X. arXiv: hep-lat/0312010.

Appendix A

Strategies for Simulating Time Evolution of Hamiltonian Lattice Field Theories

A.1 Proofs of Lemmas and Theorems

A.1.1 Proof of Lemma 12

Proof. Demanding that the new PF still reproduces the time evolution operator with the same exponentiation error we can write

$$\begin{aligned}
&\left\| \mathrm{e}^{-iHt} - \tilde{S}_p(t) \right\| \\
&= \left\| \mathrm{e}^{-iHt} - S_p(t) + \left[S_p(t/r) \right]^r - \left[\tilde{S}_p(t/r) \right]^r \right\| \\
&\leq \mathcal{O}(\varepsilon) + \left\| \left[\tilde{S}_p(t/r) \right]^r - \left[S_p(t/r) \right]^r \right\| \\
&= \mathcal{O}(\varepsilon) + r\, \mathsf{N}_\mathsf{H} \Upsilon(p) \left\| U_v^{(J)}(t/r) - \mathrm{e}^{-it/r a_{vJ} H_J} \right\| \\
&= \mathcal{O}(\varepsilon)\,.
\end{aligned} \tag{A.1}$$

□

A.1.2 Proof of Theorem 13

Proof. Using Theorem 6 and Lemma 11, one finds that the Trotter number r is asymptotically given by

$$r = \mathcal{O}\left(\frac{|||H|||_1 ||H||_1^{\frac{1}{p}} t^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}} \right)\,. \tag{A.2}$$

Using previously established notation, we have

$$||H||_1 = \mathcal{O}\left(\mathsf{N_H}\sum_J ||H_J||\right) = \mathcal{O}(\mathsf{N_H}\mathsf{N_P}\mathsf{c_P})\,, \tag{A.3}$$

$$|||H|||_1 = \mathcal{O}(\mathsf{N_{ind}}\mathsf{N_P}\mathsf{c_P})\,, \tag{A.4}$$

which immediately gives

$$r = \mathcal{O}\left(\mathsf{N_{ind}}\mathsf{N_H}^{1/p}\frac{(\mathsf{N_P}\mathsf{c_P}t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right). \tag{A.5}$$

Assume now that we wish to compute the total gate count χ for a quantum circuit implementing an approximation of e^{-iHt} to total exponentiation error $\mathcal{O}(\varepsilon)$ based on the splitting $\big(S_p(t/r)\big)^r$, where the Trotter number r in fact satisfies Eq. (2.82) (i.e., the asymptotic scaling is tight) so that the spectral norm $\| \big(S_p(t/r)\big)^r - \mathrm{e}^{-iHt}\| = \mathcal{O}(\varepsilon)$. This splitting takes the following form (see Eq. (2.14))

$$\Big[S_p(t/r)\Big]^r \equiv \left(\prod_{\upsilon=1}^{\Upsilon(p)} \mathcal{P}_\upsilon\left(\prod_{J\in\mathcal{S}} \exp\left(\frac{-ita_\upsilon^{(J)}H_J}{r}\right)\right)\right)^r. \tag{A.6}$$

Each $S_p(t/r)$ thus consists of $\mathcal{O}\left(\mathsf{N_H}\Upsilon(p)\right)$-many exponentials of the form $\exp\Big(-ita_\upsilon^{(J)}H_J/r\Big)$. Each of these exponentials can be approximated using a different PF $\tilde{S}_{\tilde{p}}^{(\upsilon,J)}(t/r\tilde{r})$ based on the Pauli decomposition of each term H_J (see Eq. (2.3)). This splitting takes the form

$$\Big[\tilde{S}_{\tilde{p}}^{(\upsilon,J)}(t/(r\tilde{r}))\Big]^{\tilde{r}} \equiv \left(\prod_{\upsilon=1}^{\Upsilon(\tilde{p})} \tilde{\mathcal{P}}_\upsilon\left(\prod_{i=1}^{\mathsf{N_P}^{(J)}} \exp\left(\frac{-it\tilde{a}_\upsilon^{(i)}c_i^{(J)}P_i^{(J)}}{r\tilde{r}}\right)\right)\right)^{\tilde{r}}. \tag{A.7}$$

From Lemma 12 we know that to ensure that the total error remains $\mathcal{O}(\varepsilon)$, the accuracy of this PF has to be

$$\begin{aligned}&\left\|\Big(\tilde{S}_{\tilde{p}}^{(\upsilon,J)}(t/r\tilde{r})\Big)^{\tilde{r}} - \exp\left(\frac{-ita_\upsilon^{(J)}H_J}{r}\right)\right\| \\ &= \mathcal{O}\Big(\frac{\varepsilon}{\mathsf{N_H}r\Upsilon(p)}\Big)\,.\end{aligned} \tag{A.8}$$

Substituting this into the result of Theorem 6 we therefore find

$$\begin{aligned}\tilde{r} &= \Omega\left(\left(\frac{\mathsf{N_P c_P} t}{r}\right)^{1+\frac{1}{\tilde{p}}}\left(\frac{\varepsilon}{\mathsf{N_H}\Upsilon(p) r}\right)^{-\frac{1}{\tilde{p}}}\right) \\ &= \Omega\left(\frac{\Upsilon(p)}{\mathsf{N_{ind}}}\left(\frac{\mathsf{N_P c_P} t}{\mathsf{N_H}\varepsilon}\right)^{\frac{p}{\tilde{p}}}\right).\end{aligned} \tag{A.9}$$

We choose $\tilde{p} = p$ and use that $\Upsilon(p) = \mathcal{O}(1)$. We also have that $\tilde{r}, \mathsf{N_{ind}} \in \mathbb{N}$ have to both be at least 1, and hence one finds $\tilde{r} = \Omega(1)$.

Using that $\Upsilon(\tilde{p}) = \mathcal{O}(1)$, we have that the resource scaling is thus given by

$$\chi = \mathcal{O}(\mathsf{N_H N_P} r) = \mathcal{O}\left(\mathsf{N_{ind} N_P}\frac{(\mathsf{N_H N_P c_P} t)^{1+\frac{1}{p}}}{\varepsilon^{\frac{1}{p}}}\right). \tag{A.10}$$

□

A.1.3 Proof of Theorem 21

Proof. From Lemma 12 we know that the simulation of a single exponential of the form $\mathrm{e}^{-ia_v^{(J)} H_J t/r}$ acting on $\mathcal{O}(1)$ sites is needed to accuracy $\varepsilon/(\mathsf{N_H} r)$. Performing this approximation using QSP (see Theorem 19) requires a gate count

$$\begin{aligned}\chi_J &= \mathcal{O}\left(\mathsf{N_P}\log\mathsf{N_P}\left(\mathsf{N_P c_P}\frac{t}{r} + \log(\mathsf{N_H} r/\varepsilon)\right)\right) \\ &= \mathcal{O}\left(\mathsf{N_P}\log\mathsf{N_P}\log(\mathsf{N_H} r/\varepsilon)\right),\end{aligned} \tag{A.11}$$

where the first term can be discarded because the asymptotic assumption on r given by Eq. (2.93) indicates that $\mathsf{N_P c_P} t/r = \mathcal{O}(\varepsilon/(\mathsf{N_P c_P} t)^{1/p})$, which is asymptotically bounded from above by a constant.

The total gate count from the $(\mathsf{N_H} r)$-many applications of this QSP subroutine is then given by

$$\chi = \mathsf{N_H}\, r\, \chi_J = \mathcal{O}\Big[\mathsf{N_{ind}}\,(\mathsf{N_H N_P c_P} t)^{1+\frac{1}{p}}\,\varepsilon^{-\frac{1}{p}}\mathsf{N_P}\log\mathsf{N_P}\log(\mathsf{N_{ind} N_H N_P c_P} t/\varepsilon)\Big]. \tag{A.12}$$

Since the block-encoding produced by QSP has a scale factor $1 \pm \mathcal{O}(\varepsilon/(r\mathsf{N_H}))$, the final block-encoding produced by multiplying $(\Upsilon(p)\mathsf{N_H} r)$-many such block-encodings will again be a block-encoding, this time with scale-factor $1 \pm \mathcal{O}(\varepsilon)$. This completes the proof. □

A.1.4 Proof of Theorem 23

Proof. From Theorem 10 we need the helper function to implement a time evolution operator for time $\tilde{t} = \|H_J\|t = \mathcal{O}(\mathsf{N_P c_P} t)$ on a system of size $N_\ell = \mathcal{O}(\ell^d)$ with exponentiation error δ, where ℓ and δ are given in Eqs. (2.75) and (2.76).

From Theorem 13 we find that a PF to implement the time evolution on a system with N_ℓ lattice sites and exponentiation error δ requires resources (taking $t = \mathcal{O}(1)$ in Eq. (2.90))

$$\begin{aligned}\chi_{\mathrm{PF}} &= \mathcal{O}\Big(\mathsf{N_P}\left(\mathsf{N}_\ell \mathsf{N_P c_P}\right)^{1+\frac{1}{p}} \delta^{-\frac{1}{p}}\Big) \\ &= \mathcal{O}\left(\mathsf{N_P}\left(\mathsf{N_P c_P}\right)^{1+\frac{1}{p}} \left[\log\left(\mathsf{N}_\Lambda t/\varepsilon\right)\right]^{\mathsf{d}\left(1+\frac{1}{p}\right)} \left(\frac{\mathsf{N}_\Lambda t}{\varepsilon \log^d\left(\mathsf{N}_\Lambda t/\varepsilon\right)}\right)^{\frac{1}{p}}\right) \\ &= \mathcal{O}\left(\mathsf{N_P}\left(\mathsf{N_P c_P}\right)^{1+\frac{1}{p}} \log^{\mathsf{d}}\left(\frac{\mathsf{N}_\Lambda t}{\varepsilon}\right)\left(\frac{\mathsf{N}_\Lambda t}{\varepsilon}\right)^{\frac{1}{p}}\right).\end{aligned} \tag{A.13}$$

Furthermore, we know from Theorem 10 that the number of calls to this helper function is given by Eq. (2.74), such that the final gate complexity is

$$\begin{aligned}\chi &= \mathcal{O}\left(\frac{\mathsf{N}_\Lambda t}{\log^{\mathsf{d}}\left(\mathsf{N}_\Lambda t/\varepsilon\right)}\chi_{\mathrm{PF}}\right) \\ &= \mathcal{O}\Big(\mathsf{N_P}\left(\mathsf{N}_\Lambda \mathsf{N_P c_P} t\right)^{1+\frac{1}{p}} \varepsilon^{-\frac{1}{p}}\Big).\end{aligned} \tag{A.14}$$

as desired. □

A.1.5 Proof of Theorem 24

Proof. Theorem 10 states that we need the helper function to implement a time evolution operator for time $\tilde{t} = \|H_J\|t = \mathcal{O}(\mathsf{N_P c_P} t)$ on a system of size $N_\ell = \mathcal{O}(\ell^d)$ with exponentiation error δ, where ℓ and δ are given in Eqs. (2.75) and (2.76). From Theorem 17 we find that implementing the time evolution on a system with N_ℓ lattice sites and exponentiation error δ with QSP requires resources (taking $t = \mathcal{O}(1)$ in Eq. (2.90))

$$\begin{aligned}\chi_{\mathrm{QSP}} &= \mathcal{O}\Big(\mathsf{N}_\ell \mathsf{N_P} \log\left(\mathsf{N}_\ell \mathsf{N_P}\right)\left(\mathsf{N}_\ell \mathsf{N_P c_P} + \log(1/\delta)\right)\Big] \\ &= \mathcal{O}\left(\log^{\mathsf{d}}\left(\mathsf{N}_\Lambda t/\varepsilon\right) \mathsf{N_P} \log\left[\mathsf{N_P} \log\left(\mathsf{N}_\Lambda t/\varepsilon\right)\right]\left[\mathsf{N_P c_P} \log^{\mathsf{d}}\left(\mathsf{N}_\Lambda t/\varepsilon\right) + \log\left(\frac{\mathsf{N}_\Lambda t}{\varepsilon \log^{\mathsf{d}}\left(\mathsf{N}_\Lambda t/\varepsilon\right)}\right)\right]\right).\end{aligned} \tag{A.15}$$

In the limit of large system size, the second term in the last square bracket (the one originating from the $\log(1/\delta)$ term) is subdominant to the first, which allows us to write

$$\chi_{\mathrm{QSP}} = \mathcal{O}\left(\mathsf{N_P}^2 \mathsf{c_P} \log^{2\mathsf{d}}\left(\mathsf{N}_\Lambda t/\varepsilon\right) \log\left[\mathsf{N_P} \log\left(\mathsf{N}_\Lambda t/\varepsilon\right)\right]\right). \tag{A.16}$$

Using the number of calls to this helper function from Eq. (2.74), the final gate complexity is

$$\begin{aligned}\chi &= \mathcal{O}\left(\frac{\mathsf{N}_\Lambda t}{\log^{\mathsf{d}}\left(\mathsf{N}_\Lambda t/\varepsilon\right)}\chi_{\mathrm{QSP}}\right) \\ &= \mathcal{O}\left(\mathsf{N}_\mathsf{P}{}^2\mathsf{N}_\Lambda \mathsf{c}_\mathsf{P} t \log^{\mathsf{d}}\left(\mathsf{N}_\Lambda t/\varepsilon\right)\log\left[\mathsf{N}_\mathsf{P}\log\left(\mathsf{N}_\Lambda t/\varepsilon\right)\right]\right),\end{aligned} \tag{A.17}$$

as desired. □

A.2 General qubitization procedure

In all the cases considered in this work the quantum walk operator W_H can be constructed from the block encoding U_H using the recipe given in Eq. (2.45). For completeness, below we provide a general approach from Ref. [109] which allows one to assemble the circuit for the W_H operator using an ancillary qubit and two controlled calls to U_H. This is necessary in cases when one is unable to find an operator S satisfying Eq. (2.46).

The general construction requires adding one qubit $|q\rangle$ to the ancillary Hilbert space, such that $\mathcal{H}_a \to \mathcal{H}_{qa} = \mathcal{H}_q \otimes \mathcal{H}_a$. In this new ancillary Hilbert space we define

$$|g\rangle_{qa} = |+\rangle_q |g\rangle_a, \quad \text{where} \quad \left(\langle g|_a \otimes \mathbb{1}_s\right) U_H \left(|g\rangle_a \otimes \mathbb{1}_s\right) = H/\alpha\,, \tag{A.18}$$

and the operator W_H is defined on the Hilbert space $\mathcal{H}_{qas} = \mathcal{H}_q \otimes \mathcal{H}_a \otimes \mathcal{H}_s$. As was shown

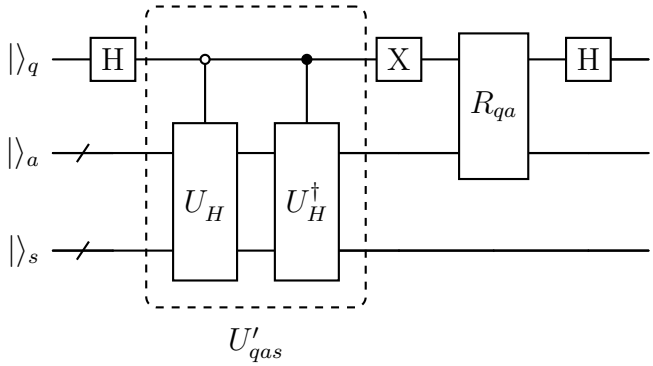

Figure A.1: Circuit for the qubitized block encoding W_H, given an arbitrary block encoding U_H. The circuit for R_{qa} is shown in Fig. A.2.

in Section 4 of Ref. [109], it is given by

$$W_H = (R_{qa} \otimes \mathbb{1}_s)(\mathrm{X}_q \otimes \mathbb{1}_{as}) U'_{qas}\,, \tag{A.19}$$

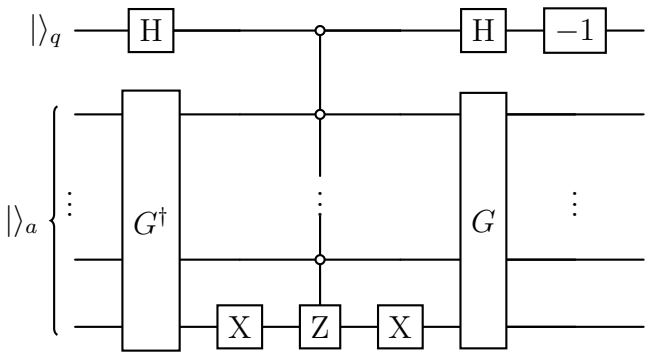

Figure A.2: Circuit for the reflection operator R_{qa} defined in Eq. (A.20) and used in Eq. (A.19), see [117] for alternative implementations. G is the oracle for preparing the state $|g\rangle_a$, i.e. $G|0\rangle_a = |g\rangle_a$. The $\boxed{-1}$ gate adds an overall phase to the circuit (negative sign), and can be implemented, e.g., as XZXZ.

where X_q is the X gate on the qubit $|q\rangle$, R_{qa} is a standard reflection operator about $|g\rangle_{qa} = |+\rangle_q|g\rangle_a$,

$$R_{qa} = 2|g\rangle_a|+\rangle_q\langle+|_q\langle g|_a - \mathbb{1}_{qa}\,, \tag{A.20}$$

and

$$U'_{qas} = |0\rangle_q\langle 0|_q U_H + |1\rangle_q\langle 1|_q U_H^\dagger\,, \tag{A.21}$$

is an operator that applies the U_H or $U_H^\dagger$ operator, depending on the state of $|q\rangle$;

One can show through explicit computation that

$$\left(\langle g|_{qa}\otimes\mathbb{1}_s\right)\left(X_q\otimes\mathbb{1}_a\otimes\mathbb{1}_s\right)U'_{qas}\left(|g\rangle_{qa}\otimes\mathbb{1}_s\right) = H/\alpha \tag{A.22}$$

$$\left(\langle g|_{qa}\otimes\mathbb{1}_s\right)\left[\left(X_q\otimes\mathbb{1}_a\otimes\mathbb{1}_s\right)U'_{qas}\right]^2\left(|g\rangle_{qa}\otimes\mathbb{1}_s\right) = \mathbb{1}_s \tag{A.23}$$

verifying that Eq. (A.19) gives indeed a qubitized BE of H. The circuits for operators W_H, U'_{qas}, and R_{qa} are shown in Figs. A.1 and A.2.

A.3 Pauli decomposition of Fourier transforms of diagonal operators

In this appendix, we show that, if $D = \mathrm{diag}(d_0, d_1, \ldots, d_{2^{n_q}-1})$ is an n_q-qubit diagonal Hermitian operator, then $\mathcal{F}^\dagger D\mathcal{F}$ can be decomposed into at most $\mathcal{O}(3^{n_q})$ Pauli-gates. To do so, we will first write D as

$$D = \sum_{j=0}^{2^{n_q}-1} d_j\,|j\rangle\,\langle j|\,. \tag{A.24}$$

The key relation we use is that that the action of $\mathcal{F}$ on a computational basis state factorizes,

$$\mathcal{F}\left|j\right\rangle = \frac{1}{2^{n_q/2}} \sum_{k=0}^{2^{n_q}-1} e^{2\pi i \frac{jk}{2^{n_q}}} \left|k\right\rangle = \left|\phi_n(j)\right\rangle \otimes \cdots \otimes \left|\phi_1(j)\right\rangle, \tag{A.25}$$

where

$$\left|\phi_k(j)\right\rangle = \frac{1}{\sqrt{2}} \left(\left|0\right\rangle + e^{2\pi i \frac{j}{2^k}} \left|1\right\rangle\right). \tag{A.26}$$

Using this, we can write

$$\mathcal{F}^\dagger D \mathcal{F} = \sum_{j=0}^{2^{n_q}-1} d_j \mathcal{F}^\dagger \left|j\right\rangle \left\langle j\right| \mathcal{F} = \frac{1}{2^{n_q}} \sum_{j=0}^{2^{n_q}-1} d_j \left|\phi_n(j)\right\rangle \left\langle \phi_n(j)\right| \otimes \cdots \otimes \left|\phi_1(j)\right\rangle \left\langle \phi_1(j)\right|. \tag{A.27}$$

Decomposing the 2×2 matrices $\left|\phi_k(j)\right\rangle \left\langle \phi_k(j)\right|$ into Pauli gates we find

$$\begin{aligned} \left|\phi_k(j)\right\rangle \left\langle \phi_k(j)\right| &= \frac{1}{2} \left(\left|0\right\rangle \left\langle 0\right| + e^{-2\pi i \frac{j}{2^k}} \left|0\right\rangle \left\langle 1\right| + e^{2\pi i \frac{j}{2^k}} \left|1\right\rangle \left\langle 0\right| + \left|1\right\rangle \left\langle 1\right|\right) \\ &= \mathbb{1} + \cos\left(\frac{2\pi j}{2^k}\right) \mathrm{X} + \sin\left(\frac{2\pi j}{2^k}\right) \mathrm{Y}, \end{aligned} \tag{A.28}$$

which shows that the Pauli decomposition of $\mathcal{F}^\dagger D \mathcal{F}$ will contain no Pauli-strings with Z-gates. Furthermore, for $\left|\phi_1(j)\right\rangle \left\langle \phi_1(j)\right|$, the coefficient of the Y-gate, $\sin(\pi j)$, is zero for all j, indicating that, for general matrices D, the matrix $\mathcal{F}^\dagger D \mathcal{F}$ is given by a sum of $\frac{2}{3} \times 3^{n_q}$ Pauli-strings; this will not be true in general if one uses shifted Fourier transforms. It is possible that matrices D with a particular structure could lead to cancellations in the sum over j and result in less Pauli-strings.

Appendix B

Block Encoding Bosons by Signal Processing

B.1 Constructing the walk operator

In this appendix, we provide detailed proofs for the selection of the operator S (see Sec. 3.3) used in constructing the walk operator when the associated BE is implemented using QSVT (Lemma 29) or QETU (Lemma 30) for even polynomials. Specifically, we demonstrate that for both QSVT and QETU, the operator S can be chosen as a single Z-gate acting on the signal and control qubit, respectively, to satisfy the conditions in Eq. (2.46). We then consider the scenario where the BE for the full Hamiltonian $U_{\hat{H}}$ is constructed using LCU to combine several local $U_{\hat{H}_j}$ BEs, with each $U_{\hat{H}_j}$ constructed using either QSVT, QETU, or LOVE-LCU. We demonstrate that, as long as the circuit organization discussed in detail in Sec. 3.3.6 is used, the associated walk operator $W_{\hat{H}}$ can be constructed by choosing S to be a single Z-gate.

The following Lemma shows how to construct the walk operator given that that the BE was constructed using QSVT for even polynomials.

Lemma 29. *Let $\hat{A}$ and $f(\hat{A})$ be Hermitian operators acting on the Hilbert space $\mathcal{H}_s$, and, without loss of generality, assume $\|\hat{A}\| \leq 1$. Let U_A be a $(1, m)$ block encoding of $\hat{A}$, acting on the Hilbert space $\mathcal{H}_a \otimes \mathcal{H}_s$, where $\mathcal{H}_a$ is an m-qubit ancillary register. Let $U_{f(\hat{A})}$ be an $(\alpha, m+1)$ block encoding of $f(\hat{A})$, constructed using the even QSVT circuit (see Fig. 3.3) with the building block U_A, acting on the Hilbert space $\mathcal{H}_c \otimes \mathcal{H}_a \otimes \mathcal{H}_s$, where $\mathcal{H}_c$ is the single-qubit signal qubit. Then, the operator $S' = Z_c \otimes \mathbb{1}_a \otimes \mathbb{1}_s$ satisfies*

$$\left(\langle 0|_c \otimes \langle 0|_a \otimes \mathbb{1}_s\right) S' U_{f(\hat{A})} \left(|0\rangle_c \otimes |0\rangle_a \otimes \mathbb{1}_s\right) = f(\hat{A})/\alpha\,, \tag{B.1}$$

$$\left(\langle 0|_c \otimes \langle 0|_a \otimes \mathbb{1}_s\right) S' U_{f(\hat{A})} S' U_{f(\hat{A})} \left(|0\rangle_c \otimes |0\rangle_a \otimes \mathbb{1}_s\right) = \mathbb{1}_s\,. \tag{B.2}$$

Proof. Using the fact that $S' \left|0\right\rangle_c \otimes \mathbb{1}_{as} = \left|0\right\rangle_c \otimes \mathbb{1}_{as}$, we see that

$$\left(\left\langle 0\right|_c \otimes \left\langle 0\right|_a \otimes \mathbb{1}_s\right) S' U_{f(\hat{A})} \left(\left|0\right\rangle_c \otimes \left|0\right\rangle_a \otimes \mathbb{1}_s\right) = \left(\left\langle 0\right|_c \otimes \left\langle 0\right|_a \otimes \mathbb{1}_s\right) U_{f(\hat{A})} \left(\left|0\right\rangle_c \otimes \left|0\right\rangle_a \otimes \mathbb{1}_s\right) \tag{B.3}$$

$$= f(\hat{A})/\alpha, \tag{B.4}$$

where the final equality in the above equation is the definition of the block encoding $U_{f(\hat{A})}$. The first relation is therefore satisfied.

To prove the second relation is satisfied, we show that $S' U_{f(\hat{A})} = U^\dagger_{f(\hat{A})} S'$. To see this, we use two circuit identities: $ZH = HX$ and $X_c \otimes \mathbb{1}_a \mathrm{CR}_{\tilde{\phi}} = \mathrm{CR}^\dagger_{\tilde{\phi}} X_c \otimes \mathbb{1}_a$ (see Fig. 3.4 for the circuit for $\mathrm{CR}_{\tilde{\phi}}$). The latter follows from the facts that X gates commute with the target of multi-controlled NOT gate and $XR_z(\theta) = R_z(\theta)^\dagger X$. Repeated use of these identities imply that commuting the Z gate in S' past the QSVT circuit flips the sign of the phases in each $\mathrm{CR}_{\tilde{\phi}}$. However, because the phases $\tilde{\phi}$ are symmetric, combined with the fact that one alternates calls to U_A and $U_A^\dagger$, flipping sign of the phases in each $\mathrm{CR}_{\tilde{\phi}}$ is equivalent to taking the Hermitian conjugate and thus $S' U_{f(\hat{A})} = U^\dagger_{f(\hat{A})} S'$. This implies that $S' U_{f(\hat{A})} S' U_{f(\hat{A})} = S' U_{f(\hat{A})} U^\dagger_{f(\hat{A})} S' = (S')^2 = \mathbb{1}_{cas}$ and so the second relation is satisfied.

□

The following Lemma shows how to construct the walk operator given that that the BE was constructed using QETU for even polynomials.

Lemma 30. *Let $\hat{A}$ and $f(\hat{A})$ be Hermitian operators acting on the Hilbert space $\mathcal{H}_s$. Let $U_{f(\hat{A})}$ be an $(\alpha, 1)$ block-encoding of $f(\hat{A})$, constructed using the QETU circuit for even polynomials (see Fig. 3.5) with the building block $\mathrm{e}^{-i\tau\hat{A}}$, acting on the Hilbert space $\mathcal{H}_c \otimes \mathcal{H}_s$ where $\mathcal{H}_c$ is the Hilbert space of the control qubit in the QETU circuit. Then, the operator $S' = Z_c \otimes \mathbb{1}_s$ satisfies*

$$\left(\left\langle 0\right|_c \otimes \mathbb{1}_s\right) S' U_{f(\hat{A})} \left(\left|0\right\rangle_c \otimes \mathbb{1}_s\right) = f(\hat{A})\,, \tag{B.5}$$

$$\left(\left\langle 0\right|_c \otimes \mathbb{1}_s\right) S' U_{f(\hat{A})} S' U_{f(\hat{A})} \left(\left|0\right\rangle_c \otimes \mathbb{1}_s\right) = \mathbb{1}_s\,. \tag{B.6}$$

Proof. Using the fact that $S' \left|0\right\rangle_c \otimes \mathbb{1}_s = \left|0\right\rangle_c \otimes \mathbb{1}_s$, we see that

$$\left(\left\langle 0\right|_c \otimes \mathbb{1}_s\right) S' U_{f(\hat{A})} \left(\left|0\right\rangle_c \otimes \mathbb{1}_s\right) = \left(\left\langle 0\right|_c \otimes \mathbb{1}_s\right) U_{f(\hat{A})} \left(\left|0\right\rangle_c \otimes \mathbb{1}_s\right) = f(\hat{A})\,, \tag{B.7}$$

where the final equality in the equation above is the definition of the block encoding $U_{\hat{H}}$. The first relation is therefore satisfied.

To show that the second relation is satisfied, we will demonstrate that $S' U_{f(\hat{A})} = U^\dagger_{f(\hat{A})} S'$. We denote by C_U the controlled call to the building block $\mathrm{e}^{-i\tau\hat{A}}$. Note that, due to the structure of alternating calls to C_U and $C_U^\dagger$, combined with the fact that the rotation angles

$\{\tilde{\varphi}_j\}$ are symmetric, the circuit for $U^\dagger_{f(\hat{A})}$ is equivalent to negating the phases in the single-qubit R_x gates in the original QETU circuit. We will now show that commuting $Z_c \otimes \mathbb{1}_s$ past $U_{\hat{H}}$ has the effect of negating the signs of the phases in the R_x gates. We note the following circuit identities: commuting a Z gate past an $R_x(\theta)$ gate flips the sign of the rotation angle, *i.e.*, $Ze^{i\theta X} = e^{-i\theta X}Z$, and Z gates acting on the control qubit commute with controlled calls to $e^{\pm i\tau\hat{A}}$. This can be seen by noting that diagonal matrices commute, which implies $[Z, |0\rangle\langle 0|] = [Z, |1\rangle\langle 1|] = 0$. These identities imply that commuting the Z gate past the QETU circuit flips the sign of all rotation angles, and therefore $S'U_{f(\hat{A})} = U^\dagger_{f(\hat{A})}S'$. Using this result, we see that $S'U_{f(\hat{A})}S'U_{f(\hat{A})} = S'U_{f(\hat{A})}U^\dagger_{f(\hat{A})}S' = (S')^2 = \mathbb{1}_{cs}$, and so the second relation is satisfied. □

The following lemma shows how to construct the walk operator $W_{\hat{H}}$ for the full Hamiltonian from $U_{\hat{H}}$ when $U_{\hat{H}}$ is constructed using LCU to combine local $U_{\hat{H}_j}$'s, where each $U_{\hat{H}_j}$ shares a common S operator for constructing $W_{\hat{H}_j}$.

Lemma 31. *Let $\hat{H} = \sum_{j=0}^{M-1}\beta_j\hat{H}_j$ be a Hamiltonian acting on a Hilbert space $\mathcal{H}_s$ with $\beta_j \in \mathbb{R}^+$. Without loss of generality, assume $\|\hat{H}_j\| \leq 1, \forall j$. Let $U_{\hat{H}_j}$ be a $(1, m_j)$ block-encoding of $\hat{H}_j$, acting on the joint Hilbert space $\mathcal{H}_a \otimes \mathcal{H}_s$, where $\mathcal{H}_a$ is the Hilbert space of ancillary qubit register containing $\max_j(m_j)$ qubits. Furthermore, let $U_{\hat{H}} = (\text{PREP}^\dagger \otimes \mathbb{1}_{as})\text{SEL}(\text{PREP} \otimes \mathbb{1}_{as})$ be a $\left(\|\vec{\beta}\|_1, \lceil\log_2 M\rceil + \max_j(m_j)\right)$ block-encoding of $\hat{H}$ acting on the joint Hilbert space $\mathcal{H}_{a_\Lambda} \otimes \mathcal{H}_a \otimes \mathcal{H}$, where $\mathcal{H}_{a_\Lambda}$ is the Hilbert space of the ancillary register containing $\lceil\log_2 M\rceil$ qubits, $\|\vec{\beta}\|_1 \equiv \sum_{j=0}^{M-1}|\beta_j|$, $\text{SEL} = \sum_{j=0}^{M-1}(|j\rangle\langle j|)_{a_\Lambda} \otimes U_{\hat{H}_j}$ and $\text{PREP}|0\rangle_{a_\Lambda} = \frac{1}{\sqrt{\|\vec{\beta}\|_1}}\sum_{j=0}^{M-1}\sqrt{\beta_j}|j\rangle_{a_\Lambda}$. Define $S' = \mathbb{1}_{a_\Lambda} \otimes S_a \otimes \mathbb{1}_s$. If for each $U_{\hat{H}_j}$ the operator S_a satisfies*

$$\left(\langle 0|_a \otimes \mathbb{1}_s\right)(S_a \otimes \mathbb{1}_s)U_{\hat{H}_j}\left(|0\rangle_a \otimes \mathbb{1}_s\right) = \hat{H}_j\,, \tag{B.8}$$

$$\left(\langle 0|_a \otimes \mathbb{1}_s\right)\left[S_a \otimes \mathbb{1}_s)U_{\hat{H}_j}\right]^2\left(|0\rangle_a \otimes \mathbb{1}_s\right) = \mathbb{1}_s\,, \tag{B.9}$$

then

$$\left(\langle 0|_{a_\Lambda} \otimes \langle 0|_a \otimes \mathbb{1}_s\right)S'U_{\hat{H}}\left(|0\rangle_{a_\Lambda} \otimes |0\rangle_a \otimes \mathbb{1}_s\right) = \frac{\hat{H}}{\|\vec{\beta}\|_1}\,, \tag{B.10}$$

$$\left(\langle 0|_{a_\Lambda} \otimes \langle 0|_a \otimes \mathbb{1}_s\right)S'U_{\hat{H}}S'U_{\hat{H}}\left(|0\rangle_{a_\Lambda} \otimes |0\rangle_a \otimes \mathbb{1}_s\right) = \mathbb{1}_s\,. \tag{B.11}$$

Proof. First, note that

$$\begin{aligned} S'U_{\hat{H}} &= \left(\mathbb{1}_{a_\Lambda} \otimes S_a \otimes \mathbb{1}_s\right)(\text{PREP}^\dagger \otimes \mathbb{1}_{as})\left(\sum_{j=0}^{M-1}(|j\rangle\langle j|)_{a_\Lambda} \otimes U_{\hat{H}_j}\right)(\text{PREP} \otimes \mathbb{1}_{as}) \\ &= \sum_{j=0}^{M-1}\left[\text{PREP}^\dagger(|j\rangle\langle j|)_{a_\Lambda}\text{PREP}\right] \otimes \left[(S_a \otimes \mathbb{1}_s)U_{\hat{H}_j}\right]. \end{aligned} \tag{B.12}$$

From Eq. (B.12) it immediately follows that the first relation is satisfied:

$$
\begin{aligned}
&(\langle 0|_{a_\Lambda} \otimes \langle 0|_a \otimes \mathbb{1}_s) S' U_{\hat{H}} \left(|0\rangle_{a_\Lambda} \otimes |0\rangle_a \otimes \mathbb{1}_s\right) \\
&\quad = \left(\langle 0|_{a_\Lambda} \otimes \langle 0|_a \otimes \mathbb{1}_s\right) \sum_{j=0}^{M-1} \left[\text{PREP}^\dagger (|j\rangle\langle j|)_{a_\Lambda} \text{PREP}\right] \otimes \left[(S_a \otimes \mathbb{1}_s) U_{\hat{H}_j}\right] \left(|0\rangle_{a_\Lambda} \otimes |0\rangle_a \otimes \mathbb{1}_s\right) \\
&\quad = \sum_{j=0}^{M-1} \langle 0|_{a_\Lambda} \left[\text{PREP}^\dagger (|j\rangle\langle j|)_{a_\Lambda} \text{PREP}\right] |0\rangle_{a_\Lambda} \otimes \left(\langle 0|_a \otimes \mathbb{1}_a\right) \left[(S_a \otimes \mathbb{1}_s) U_{\hat{H}_j}\right] \left(|0\rangle_a \otimes \mathbb{1}_a\right) \\
&\quad = \sum_{j=0}^{M-1} \langle 0|_{a_\Lambda} \left[\text{PREP}^\dagger (|j\rangle\langle j|)_{a_\Lambda} \text{PREP}\right] |0\rangle_{a_\Lambda} \hat{H}_j \\
&\quad = \frac{1}{\|\vec{\beta}\|_1} \sum_{j,k,l=0}^{M-1} \sqrt{\beta_k \beta_\ell} \langle k|j\rangle\langle j|l\rangle \hat{H}_j \\
&\quad = \frac{1}{\|\vec{\beta}\|_1} \sum_{j=0}^{M-1} \beta_j \hat{H}_j \\
&\quad = \frac{\hat{H}}{\|\vec{\beta}\|_1} .
\end{aligned}
\tag{B.13}
$$

The second relation also follows in a straightforward way:

$$
\begin{aligned}
&(\langle 0|_{a_\Lambda} \otimes \langle 0|_a \otimes \mathbb{1}_s) S' U_{\hat{H}} S' U_{\hat{H}} \left(|0\rangle_{a_\Lambda} \otimes |0\rangle_a \otimes \mathbb{1}_s\right) \\
&= (\langle 0|_{a_\Lambda} \otimes \langle 0|_a \otimes \mathbb{1}_s) \left[\sum_{j=0}^{M-1} \left[\text{PREP}^\dagger (|j\rangle\langle j|)_{a_\Lambda} \text{PREP}\right] \otimes \left[(S_a \otimes \mathbb{1}_s) U_{\hat{H}_j}\right]\right]^2 \left(|0\rangle_{a_\Lambda} \otimes |0\rangle_a \otimes \mathbb{1}_s\right) \\
&= (\langle 0|_{a_\Lambda} \otimes \langle 0|_a \otimes \mathbb{1}_s) \left[\sum_{j=0}^{M-1} \left[\text{PREP}^\dagger (|j\rangle\langle j|)_{a_\Lambda} \text{PREP}\right] \otimes \left[(S_a \otimes \mathbb{1}_s) U_{\hat{H}_j}\right]^2\right] \left(|0\rangle_{a_\Lambda} \otimes |0\rangle_a \otimes \mathbb{1}_s\right) \\
&= \sum_{j=0}^{M-1} \langle 0|_{a_\Lambda} \text{PREP}^\dagger (|j\rangle\langle j|)_{a_\Lambda} \text{PREP} |0\rangle_{a_\Lambda} \left(\langle 0|_a \otimes \mathbb{1}_s\right) \left[(S_a \otimes \mathbb{1}_s) U_{\hat{H}_j}\right]^2 \left(|0\rangle_a \otimes \mathbb{1}_s\right)\Big) \\
&= \frac{1}{\|\vec{\beta}\|_1} \sum_{j,k,l=0}^{M-1} \sqrt{\beta_k \beta_\ell} \langle k|j\rangle\langle j|l\rangle \\
&= \frac{1}{\|\vec{\beta}\|_1} \sum_{j=0}^{M-1} \beta_j \\
&= \mathbb{1}_s .
\end{aligned}
\tag{B.14}
$$

□

B.2 CNOT gate counts for scalar field theory simulation

In this appendix, we present the associated CNOT gate counts from the numerical studies conducted in Sec. 3.4. All results are qualitatively similar to the rotation gate counts discussed in the main text. Fig. B.1 shows the CNOT gate counts, associated with the rotation gate counts in Fig. 3.7, for BEs of the local terms $\frac{1}{2}\hat{\pi}^2$, $\frac{m}{2}\hat{\varphi}^2 + \frac{\lambda}{4!}\hat{\varphi}^4$, $\frac{1}{2}(\hat{\varphi}_1 - \hat{\varphi}_2)$, and $g\cos(\hat{\varphi})$. Fig. B.2 shows the CNOT gate counts, associated with the rotation gate counts in Fig. 3.8, for BEs of the two-site $\hat{\varphi}$ Hamiltonian in Eq. (3.27). Fig. B.3 shows the CNOT gate counts, associated with the rotation gate counts in Fig. 3.9, for simulating time evolution of the single site Hamiltonian in Eq. (3.29) using a 2nd order PF, a 4th order PF, and GQSP where W_H was constructed using LOVE-LCU. Fig. B.4 shows the CNOT gate counts, associated with the rotation gate counts in Fig. 3.10, for simulating time evolution of the two-site Hamiltonian in Eq. (3.29) using a 2nd order PF, a 4th order PF, and GQSP where W_H was constructed using LOVE-LCU.

B.3 Scale factor for Block Encodings constructed using LCU

In this appendix, we derive the scale factors for block encodings of $f_k(\hat{\xi}_k)$, given in Eq. (3.26) in the main text, when prepared using standard LCU techniques.

To begin, we consider the scale factor for $\hat{\varphi}^d$; the result is analogous for $\left(\hat{\pi}^{(D)}\right)^d$. Using Eq. (2.105) we have

$$\begin{aligned}\hat{\varphi}^d &= (-\varphi_{\max})^d \left(\frac{2^{\mathsf{n}_\mathsf{q}-1}}{2^{\mathsf{n}_\mathsf{q}}-1}\right)^d \left(\sum_{m=0}^{\mathsf{n}_\mathsf{q}-1} 2^{-m} Z_m\right)^d \\ &= (-\varphi_{\max})^d \left(\frac{2^{\mathsf{n}_\mathsf{q}-1}}{2^{\mathsf{n}_\mathsf{q}}-1}\right)^d \sum_{m_1,m_2,\ldots,m_d=0}^{\mathsf{n}_\mathsf{q}-1} 2^{-(m_1+m_2+\cdots+m_d)} Z_{m_1} Z_{m_2} \ldots Z_{m_d}\,.\end{aligned} \tag{B.15}$$

Since $Z_{m_j}^2 = \mathbb{1}$, the sum over m_j includes multiple instances of the same Pauli string. Constructing a BE using LCU requires grouping these repeated terms and determining the total coefficient for each unique Pauli string. However, because the coefficient of each Pauli string in the sum has the same sign, no cancellations in the final values of the coefficients occur. This fact implies that, even if there are repeated Pauli strings in the sum, the scale factor can be found by simply adding the coefficients of each Pauli string in the sum. Using

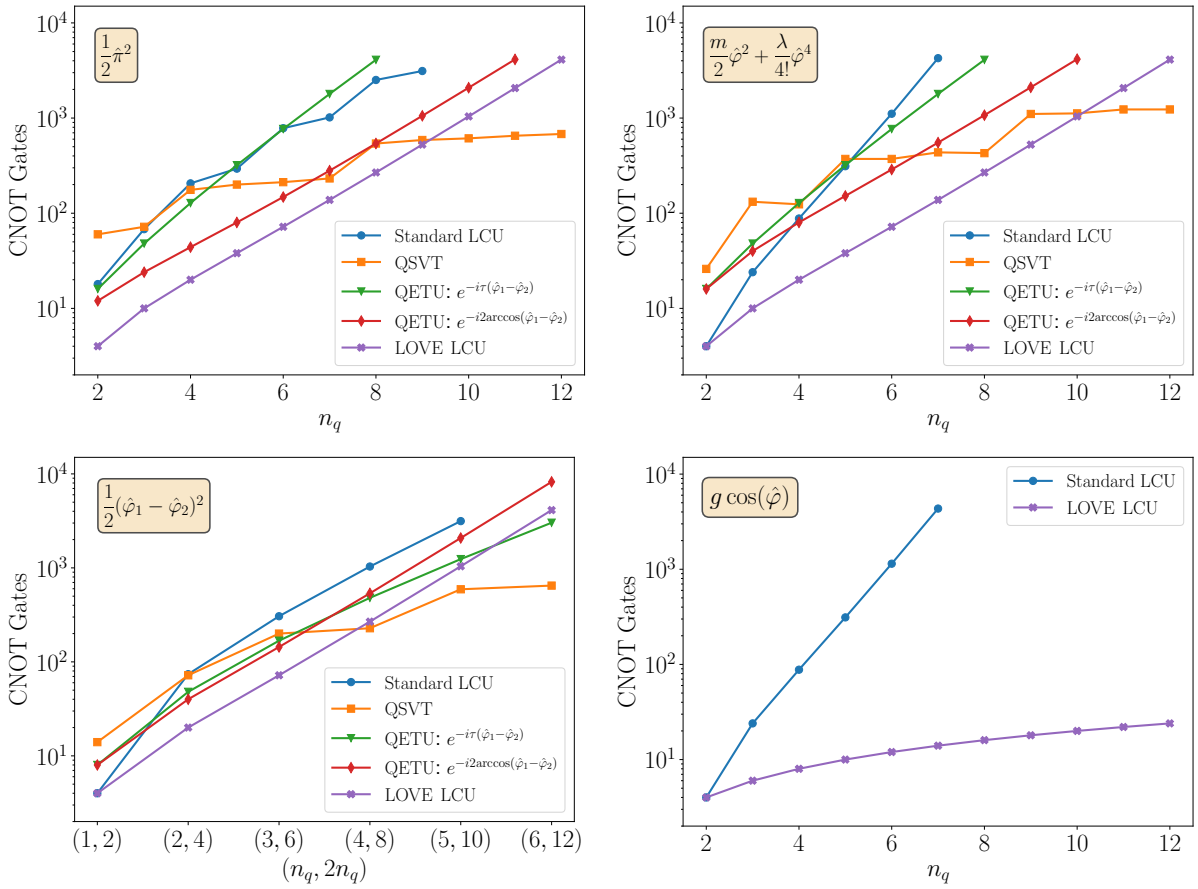

Figure B.1: CNOT gate count and number of ancillary qubits required to block encode local bosonic operators. The top left, top right, bottom left, and bottom right plots show resource requirements to block encode $\frac{1}{2}\hat{\pi}^2$, $\frac{m}{2}\hat{\varphi}^2 + \frac{\lambda}{4!}\hat{\varphi}^4$, $\frac{1}{2}(\hat{\varphi}_1 - \hat{\varphi}_2)^2$, $g\cos(\hat{\varphi})$, respectively, for $m = 1, \lambda = 32, g = 1$. Different colored and shaped data points correspond to different methods to prepare the BE. For the single-site operators $\hat{\pi}^2$ and $\frac{m}{2}\hat{\varphi}^2 + \frac{\lambda}{4!}\hat{\varphi}^4$, LOVE-LCU requires the fewest rotation gates for $\mathsf{n_q} \lesssim 9$. For the two-site operator $\frac{1}{2}(\hat{\varphi}_1 - \hat{\varphi}_2)^2$, LOVE-LCU requires the fewest rotation gates for $\mathsf{n_q} \lesssim 4$. LOVE-LCU constructs a BE of $g\cos(\hat{\varphi})$ using the fewest gates for all $\mathsf{n_q}$.

this fact, the scale factor is given by

$$\begin{aligned}
\beta_{\hat{\varphi}^d} &= \varphi_{\max}^d \left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^d \sum_{m_1,m_2,\ldots,m_d=0}^{\mathsf{n_q}-1} 2^{-(m_1+m_2+\cdots+m_d)} \\
&= \varphi_{\max}^d \left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^d \left(\sum_{m=0}^{\mathsf{n_q}-1} 2^{-m}\right)^d \\
&= \varphi_{\max}^d \left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^d \left(\frac{2^{\mathsf{n_q}}-1}{2^{\mathsf{n_q}-1}}\right)^d \\
&= \varphi_{\max}^d \,,
\end{aligned} \tag{B.16}$$

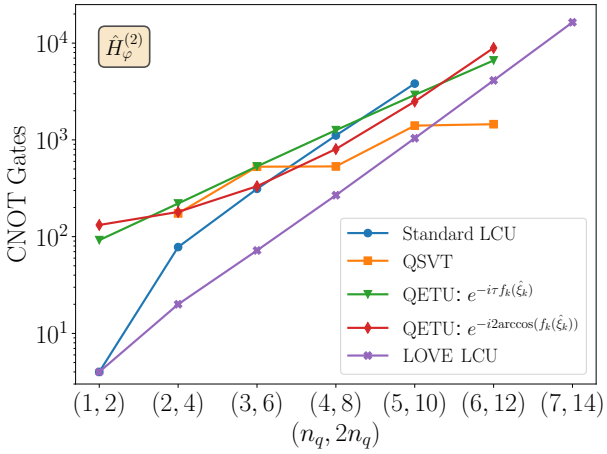

Figure B.2: CNOT gate count required to block encode the two site Hamiltonian $\hat{H}_\varphi^{(2)}$ in Eq. (3.27). Different colored and shaped data points correspond to different methods to prepare the BE. Points with the same number of ancillary qubits have been shifted slightly for clarity. While QSVT- and QETU-based methods require using a final layer of LCU to add the local BEs $f_0^{(0)}(\hat{\varphi}_1), f_0^{(0)}(\hat{\varphi}_2), f_2(\hat{\varphi}_1 - \hat{\varphi}_2)$, LOVE-LCU does not. This advantage also leads to LOVE-LCU outperforming all other methods for $\mathsf{n_q} \lesssim 5$; for $\mathsf{n_q} > 5$, QSVT outperforms other methods due to the relative exponentially improved asymptotic scaling.

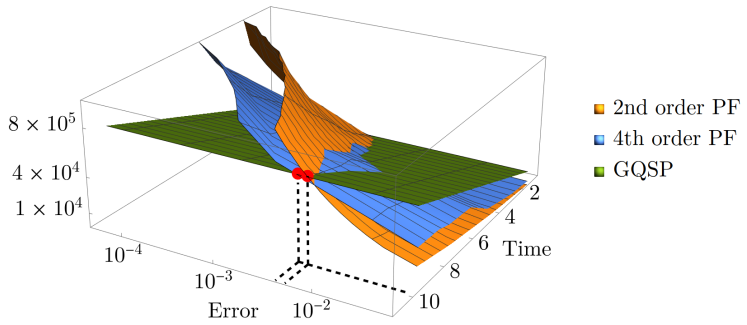

Figure B.3: CNOT gate count as a function of error ϵ and simulation time t for simulating time-evolution of the single-site and two-site Hamiltonian in Eq. (3.29). The blue, orange, and green surfaces show results calculated using a 2^{nd} order PF, a 4^{th} order PF, and GQSP where the BE was constructing using LOVE-LCU. For $t = 10$, GQSP outperforms both 2^{nd} and 4^{th} order PFs for $\epsilon \lesssim 5 \times 10^{-3}$.

where, going from the second to the third line, we used the identity $\sum_{m=0}^{\mathsf{n_q}-1} 2^{-m} = (2^{\mathsf{n_q}} -$

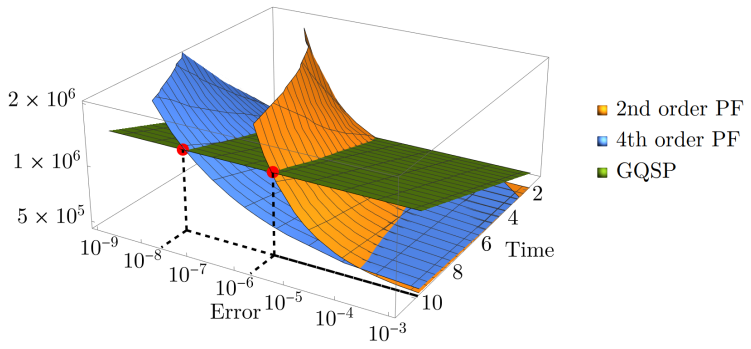

Figure B.4: CNOT gate count as a function of error ϵ and simulation time t for simulating time-evolution of the two-site Hamiltonian in Eq. (3.30). The blue, orange, and green surfaces show results calculated using a 2^{nd} order PF, a 4^{th} order PF, and GQSP where the BE was constructing using LOVE-LCU. For $t = 10$, GQSP outperforms 2^{nd} and 4^{th} order PFs for $\epsilon \lesssim 1 \times 10^{-6}$ and $\epsilon \lesssim 5 \times 10^{-8}$.

$1)/2^{\mathsf{n_q}-1}$. From this result we can directly see the scale factor for $f_1(\hat{\pi}^{(D)}) = \frac{1}{2}\big(\hat{\pi}^{(D)}\big)^2$ is

$$\beta_{f_1(\hat{\pi}^{(D)})} = \frac{1}{2}\pi_{\max}^2\,, \tag{B.17}$$

which is the smallest possible scale factor to BE this operator.

Next, we turn to $f_0^{(1)}(\hat{\varphi}) = \frac{m^2}{2}\hat{\varphi}^2 + \frac{\lambda}{4!}\hat{\varphi}^4$, whose Pauli decomposition is

$$\begin{aligned} f_0^{(1)}(\hat{\varphi}) &= \frac{m^2}{2}\hat{\varphi}^2 + \frac{\lambda}{4!}\hat{\varphi}^4 \\ &= (-\varphi_{\max})^2 \left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^2 \frac{m^2}{2} \sum_{m_1,m_2=0}^{\mathsf{n_q}-1} 2^{-(m_1+m_2)} Z_{m_1} Z_{m_2} + \\ &\quad + (-\varphi_{\max})^4 \left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^4 \frac{\lambda}{4!} \sum_{l_1,l_2,l_3,l_4=0}^{\mathsf{n_q}-1} 2^{-(l_1+l_2+l_3+l_4)} Z_{l_1} Z_{l_2} Z_{l_3} Z_{l_4}\,. \end{aligned} \tag{B.18}$$

Assuming $\lambda > 0$, we see that, similar to the previous case, the coefficient of each Pauli string in the decomposition of $f(\hat{\varphi})$ has the same sign, and the scale factor is simply the sum of the coefficients

$$\begin{aligned} \beta_{f_0^{(1)}(\hat{\varphi})} &= \frac{m^2}{2}\varphi_{\max}^2 \left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^2 \left(\sum_{m=0}^{\mathsf{n_q}-1} 2^{-m}\right)^2 + \frac{\lambda}{4!}\varphi_{\max}^4 \left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^4 \left(\sum_{m=0}^{\mathsf{n_q}-1} 2^{-m}\right)^4 \\ &= \frac{m^2}{2}\varphi_{\max}^2 + \frac{\lambda}{4!}\varphi_{\max}^4\,, \end{aligned} \tag{B.19}$$

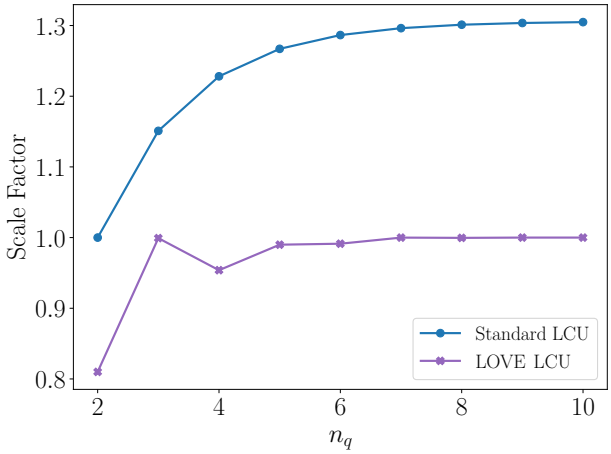

Figure B.5: Scale factor of the BE of $\cos(\hat{\varphi})$ as a function of $\mathsf{n_q}$. The blue circles and purple crosses correspond to using standard LCU and LOVE-LCU to construct the BE, respectively. The value of $\varphi_{\max}$ was set to $\varphi_{\max} = \pi$ for all values of $\mathsf{n_q}$.

which saturates the lower bound given by the operator norm.

The situation is more complicated for $f_0^{(2)}(\hat{\varphi}) = g\cos(\hat{\varphi})$ as its Pauli decomposition consists of terms with opposite signs, requiring careful accounting for cancellations. For this reason, we study the scale factor numerically. Fig. B.5 presents the scale factor of the BE of $\cos(\hat{\varphi})$ constructed using both standard LCU and LOVE-LCU. For this comparison, we set $\varphi_{\max} = \pi$. With LOVE-LCU, the scale factor approaches the optimal value of 1, whereas with standard LCU, it remains generally larger and asymptotes to ~ 1.3.

Lastly, we consider $f_2(\hat{\varphi}_1 - \hat{\varphi}_2) = \frac{1}{2}(\hat{\varphi}_1 - \hat{\varphi}_2)^2$. To study the scale factor, we first rewrite it as $f_2(\hat{\varphi}_1 - \hat{\varphi}_2) = \frac{1}{2}(\hat{\varphi}_1^2 + \hat{\varphi}_2^2 - 2\hat{\varphi}_1\hat{\varphi}_2)$. While this operator is a sum of terms with different signs, this does not affect the determination of the scale factor since the terms with differing signs have no Pauli strings in common. This can be verified by directly examining the decomposition of $f_2(\hat{\varphi}_1 - \hat{\varphi}_2)$. If we assign $\hat{\varphi}_1$ to act on qubits $0, 1, \ldots, \mathsf{n_q} - 1$ and $\hat{\varphi}_2$ to act on qubits $\mathsf{n_q}, \mathsf{n_q} + 1, \ldots, 2\mathsf{n_q} - 1$, the decomposition is

$$
\begin{aligned}
f_2(\hat{\varphi}_1 - \hat{\varphi}_2) &= \frac{1}{2}(\hat{\varphi}_1^2 + \hat{\varphi}_2^2 - 2\hat{\varphi}_1\hat{\varphi}_2) \qquad \text{(B.20)}\\
&= \frac{1}{2}\varphi_{\max}^2 \left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^2 \\
&\times \left(\sum_{m_1,m_2=0}^{\mathsf{n_q}-1} 2^{-(m_1+m_2)} Z_{m_1} Z_{m_2} + \sum_{l_1,l_2=0}^{\mathsf{n_q}-1} 2^{-(l_1+l_2)} Z_{\mathsf{n_q}+l_1} Z_{\mathsf{n_q}+l_2} - 2\sum_{m,l=0}^{\mathsf{n_q}-1} 2^{-(m+l)} Z_m Z_{\mathsf{n_q}+l} \right).
\end{aligned}
$$

From this expression, we see that all Pauli strings with opposite signs are distinct. Thus,

the scale factor can be determined by summing the coefficients. Doing so gives

$$\begin{aligned}\beta_{f_2(\hat{\varphi}_1-\hat{\varphi}_2)} &= \frac{1}{2}\varphi_{\max}^2\left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^2\Big(\sum_{m_1,m_2=0}^{\mathsf{n_q}-1}2^{-(m_1+m_2)}+\sum_{l_1,l_2=0}^{\mathsf{n_q}-1}2^{-(l_1+l_2)}+2\sum_{m,l=0}^{\mathsf{n_q}-1}2^{-(m+l)}\Big)\\ &= \frac{1}{2}\varphi_{\max}^2\left(\frac{2^{\mathsf{n_q}-1}}{2^{\mathsf{n_q}}-1}\right)^2 4\Big(\sum_{m=0}^{\mathsf{n_q}-1}2^{-m}\Big)^2\\ &= 2\varphi_{\max}^2.\end{aligned} \tag{B.21}$$

Since the operator $\hat{\varphi}$ is sampled symmetrically from $-\varphi_{\max}$ to $\varphi_{\max}$, this scale factor attains its smallest possible value.

Appendix C

Truncation Uncertainties for Accurate Quantum Simulations of Lattice Gauge Theories

C.1 Infinite MPS implementation

The simulations of the infinite plaquette chain were performed using the time-evolved block decimation algorithm (TEBD) [268]. A Trotter step size of $\Delta t = 0.01$ was used. For the results shown in Fig. 4.4, the time evolution was performed with a fixed bond dimension. The bond dimension was then increased by 10, and the full simulation was rerun. This process was repeated until the difference in the electric energy between two consecutive runs was an order of magnitude smaller than the predicted truncation error. The resulting bond dimensions are shown in Table C.1. The results shown in Fig. 4.5 were obtained by similarly performing the evolution with a fixed bond dimension and then again after increasing the bond dimension by 25. We repeated this process until the difference in the results from consecutive runs was 2 orders of magnitude below the upper bound computed using Eq. (4.88) (corresponding to the dashed horizontal lines shown in Fig. 4.5). We note that the simulation for all values of g and Λ converged within a bond dimension of 180.

The simulations of the Schwinger model were also performed using TEBD. The Hilbert space for the fermion on each site and the right neighboring link were combined into a single site, represented by a single tensor index in the matrix product state representation. A Trotter step size of $\Delta t = 0.01$ was used. It was found that performing simulations with a large fixed bond dimension led to numerical instabilities. For this reason, the simulation was performed using a bond dimension of 50 until $t = 1.5$. The bond dimension was then increased, and the same criteria was used as in the plaquette chain simulation to determine when the simulation had converged to sufficient precision. The resulting maximum bond dimensions are shown in Table C.2.

Λ	Bond Dimension
1	60
2	90
3	100
4	170
5	170

Table C.1: Bond dimension required for convergence of the electric energy for an infinite plaquette chain with $g = 0.8$ and truncation Λ shown in Fig. 4.4.

Λ	Bond Dimension
1	90
2	90
3	110
4	110

Table C.2: Bond dimension required for convergence of the electric energy for the Schwinger model on an infinite lattice with $g = 0.8$, $m = 0.1$, and truncation Λ.

C.2 Hubbard-Holstein Model

Previous work studying gauge field truncation errors for quantum simulation of lattice gauge theories also applied their results to bosonic systems [75]. The Hubbard-Holstein model was used as an example system. This theory is a model of electron-phonon interactions. The Hamiltonian is given by

$$\begin{aligned} \hat{H} &= \hat{H}_F + g\hat{H}_{BF} + \omega\hat{H}_B \\ \hat{H}_F &= \frac{1}{2}\sum_i \hat{\psi}_{i+1}^\dagger \hat{\psi}_i + \hat{\psi}_i^\dagger \hat{\psi}_{i+1} \\ \hat{H}_{BF} &= \sum_i \hat{\psi}_i^\dagger \hat{\psi}_i (\hat{a}_i + \hat{a}_i^\dagger) \\ \hat{H}_F &= \sum_i \hat{a}_i^\dagger \hat{a}_i \end{aligned} \tag{C.1}$$

where $\hat{\psi}_i$ is a fermion annihilation operator at site i, $\hat{a}_i$ is a boson annihilation operator at site i, ω is the phonon frequency and g is the electron-phonon coupling. Note that this theory does not exhibit the same Hilbert space fragmentation that was used to derive truncation error estimates for lattice gauge theories. It was estimated that to simulate the dynamics

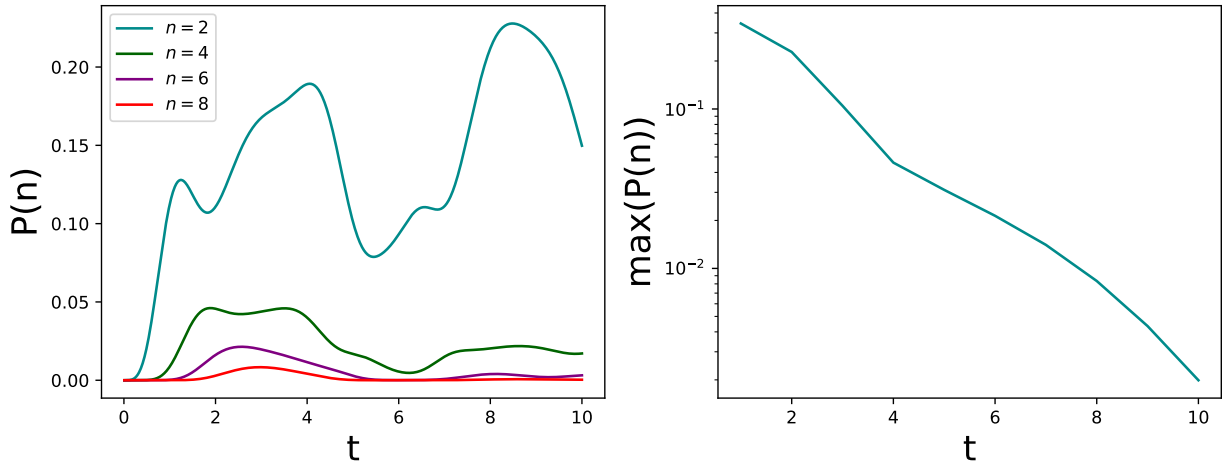

Figure C.1: Simulation of the Hubbard-Holstein model on two sites with $g = 1$ and $\omega = 1$ for a time $t = 10$. The left panel shows the probability of the boson mode on the left site having an occupation number of n as a function of time. The right panel shows the maximum probability of occupying each boson number during this time evolution.

of this model with $g = 1$ for a time $t = 1$, a boson truncation of 100 was necessary to keep leakage errors below $\epsilon = 0.2$. As a probe of this estimate, a simulation was performed of this model on two sites with $g = 1$, $\omega = 1$, and a single fermion present. In this simulation, the boson number at each site was truncated at $\Lambda = 100$. The system was initialized with the boson number at both sites equal to zero, and the fermion was placed on the left site. The system was evolved for a time $t = 10$. The left panel of Figure C.1 shows the probability of the boson mode on the left site having an occupation number of n as a function of time. The right panel shows the maximum probability of occupying each boson number during this time evolution. As this figure shows, the leakage probability is significantly smaller than previous work estimated.

Appendix D

Obtaining Continuum Physics from Dynamical Simulations of Hamiltonian Lattice Gauge Theories

D.1 Scale setting and the Kogut-Susskind Hamiltonian

In this Appendix, we provide a pedagogical discussion regarding the need for scale setting in lattice gauge theory simulations. After expressing the Kogut-Susskind Hamiltonian in terms of dimensionless variables, we comment on the bare parameters of the theory, and show that one needs to include an additional parameter in the fermionic kinetic term.

To show that one must determine the lattice spacing by a scale setting procedure, we will use the following two observations:

1. Computers (classical and quantum) can only work with dimensionless variables, *i.e.* all inputs into any numerical simulation are dimensionless.

2. The value of the lattice spacing a is not independent of the bare parameters in the Hamiltonian, but is related through renormalization group equations.

To isolate the logic of these two steps, we first consider the simple case of the harmonic oscillator and show that no such scale setting procedure is necessary. While this example may appear trivial, it highlights the key difference between standard quantum mechanical Hamiltonians and lattice field theory Hamiltonians where the lattice is used as a UV regulator. We then consider a U(1) LGT and show that the lattice spacing must be determined through a scale setting procedure.

The dimensionful harmonic oscillator Hamiltonian (with $\hbar = 1$) is given by

$$H = \omega(b^\dagger b + \frac{1}{2}) \,, \tag{D.1}$$

where b and $b^\dagger$ are lowering and raising operators that satisfy $[b, b^\dagger] = 1$. Aside from the Hamiltonian, the only dimensionful parameter is the frequency ω. If we define $\omega = \hat{\omega}/a$ where a is a dimensionful parameter with constant magnitude (with the symbol a being suggestively chosen for the LGT case), and $\hat{\omega}$ is a dimensionless parameter with magnitude $a\omega$, then the dimensionless Hamiltonian that we actually simulate is

$$\hat{H} \equiv aH = \hat{\omega}(a^\dagger a + \frac{1}{2}) \,. \tag{D.2}$$

Note that the parameter a acts as a physical scale of the system. The above equations shows that one actually calculates the energies $\hat{E}_n \equiv aE_n$ of the rescaled Hamiltonian $\hat{H}\,|n\rangle = \hat{E}_n\,|n\rangle$. Converting to the physical energies of H in this case can be done trivially by dividing by the known parameter a

$$E_n = \hat{E}_n/a \,. \tag{D.3}$$

Indeed, this step is so obvious and intuitive that it is usually done implicitly by simply setting the magnitude of $a = 1$, which implies one can directly interpret the dimensionless energies $\hat{E}_n$ as the physical energies E_n.

Consider now the more complicated scenario where the parameter $a = a(\hat{\omega})$ is an *unknown* function of $\hat{\omega}$ (or, equivalently, a function of ω). In order to convert $\widetilde{E}_n$ to the physical energies E_n, one must somehow determine the scale $a(\hat{\omega})$. One possibility is to perform a *scale setting* procedure, where one determines $a(\hat{\omega})$ by demanding the calculated dimensionless energy $\hat{E}_{n^*} = a(\hat{\omega})E_{n^*}$ is equal to the experimentally known value $E^{\text{phys}}_{n^*}$,

$$a(\hat{\omega}) \equiv \frac{[a(\hat{\omega})E_{n^*}]}{E^{\text{phys}}_{n^*}} \,. \tag{D.4}$$

As we will now discuss, this is analogous the case for LGTs, with the dimesionless bare parameters playing the role of $\hat{\omega}$ and the lattice spacing playing the role of the dimensionful scale $a(\hat{\omega})$.

To demonstrate that one does not actually choose the lattice spacing in a LGT simulation, it is sufficient to consider a U(1) LGT with a single quark flavor; this analysis generalizes to SU(N) LGTs with additional quark flavors. Here we only give the level of detail required to demonstrate the need to scale set, further details can be found in, *e.g.*, [282]. In $d \geq 2$ dimensions, the Hamiltonian is a sum of 4 terms

$$H = H_E + H_B + H_M + H_K \,, \tag{D.5}$$

where H_E and H_B are the electric and magnetic Hamiltonians governing the dynamics of the gauge fields, H_M is the fermionic mass Hamiltonian, and H_K is the fermionic kinetic (hopping) Hamiltonian. Because no symmetry demands it, the coefficients multiplying each of these 4 terms are in general different, which implies we will have 4 bare parameters in the theory; we discuss the need for a 4th bare parameter multiplying the fermionic kinetic term

in more detail later in this appendix. The individual Hamiltonian terms are

$$H_E = \frac{\widetilde{g}_t(a)^2}{2a^{d-2}} \sum_{\vec{x}} \sum_{j=1}^{d} E(\vec{x}, j)^2 \,, \tag{D.6}$$

$$H_B = -\frac{1}{2a^{4-d}\widetilde{g}_s(a)^2} \sum_{p} (P_p + P_p^\dagger) \,, \tag{D.7}$$

$$H_M = m(a) \sum_{\vec{x}} (-1)^{\vec{x}} \psi^\dagger(\vec{x}) \psi(\vec{x}) \,, \tag{D.8}$$

$$H_K = \frac{\kappa(a)}{a} \sum_{\vec{x}} \sum_{j=1}^{d} \eta(\vec{x}) \left(\psi^\dagger(\vec{x}) U(\vec{x}, j) \psi(\vec{x} + \hat{e}_j) + \text{h.c.} \right) \,, \tag{D.9}$$

where a is the lattice spacing, $\hat{E}(\vec{x}, j)$ and is an electric field operator at site $\vec{x}$ pointing in the direction of the $\hat{e}_j$ axis, $\hat{U}(\vec{x}, j)$ is a gauge link operator at site $\vec{x}$ pointing in the direction of the $\hat{e}_j$ axis, $\hat{\psi}(\vec{x})(\hat{\psi}^\dagger(\vec{x}))$ is the fermonic annihilation (creation) operator, and $\eta(\vec{x})$ is the staggered sign factor. The magnetic Hamiltonian is the sum over all independent plaquettes p on the lattice; a plaquette originating at site $\vec{x}$ oriented in the j, k plane is given by

$$P_p = U(\vec{x}, j) U(\vec{x} + \hat{e}_j, k) U^\dagger(\vec{x} + \hat{e}_k, j) U^\dagger(\vec{x}, k) \,. \tag{D.10}$$

Analogous to the simple harmonic oscillator case, we want to write a dimensionless Hamiltonian $\hat{H}$ in terms of dimensionless quantities. We write $[A] = M^\gamma$ to denote that the operator/parameter A has mass dimension γ. By definition, we have $[H] = M$, $[a] = M^{-1}$, and $[U(\vec{x}, j)] = M^0$, which implies the fermionic operators and electric field operators are dimensionless, *i.e.*, $[\psi] = M^0, [E(\vec{x}, j)] = M^0$, and the bare parameters have dimensions

$$[\widetilde{g}_t(a)] = M^{\frac{3-d}{2}}, \quad [\widetilde{g}_s(a)] = M^{\frac{3-d}{2}}, \quad [m(a)] = M^1, \quad [\kappa(a)] = M^0 \,. \tag{D.11}$$

The parameters $\widetilde{g}_t(a), \widetilde{g}_s(a)$ and $m(a)$ can be expressed in terms of the dimensionless bare parameters

$$g_t(a)^2 = a^{3-d} \widetilde{g}_t(a)^2 \,, \qquad g_s(a)^2 = a^{3-d} \widetilde{g}_s(a)^2 \,, \qquad \overline{m}(a) = a m(a) \,. \tag{D.12}$$

Rather than g_t and g_s, we reparameterize (as done in the main text) using $g^2 = g_s g_t$ and $c = g_t / g_s$. In this way, the Hamiltonian can be written as

$$H_{\text{KS}} = \frac{c}{a} \hat{H}_{\text{KS}} = \frac{c}{a} \left(\hat{H}_E + \hat{H}_B + \hat{H}_M + \hat{H}_K \right) \tag{D.13}$$

where the individual rescaled terms are given by

$$\hat{H}_E = \frac{g(a)^2}{2} \sum_{\vec{x}} \sum_{j=1}^{d} \hat{E}(\vec{x}, j)^2 \tag{D.14}$$

$$\hat{H}_B = -\frac{1}{2g(a)^2}\sum_p (\hat{P}_p + \hat{P}_p^\dagger) \tag{D.15}$$

$$\hat{H}_M = \hat{m}(a)\sum_{\vec{x}} (-1)^{\vec{x}} \hat{\psi}^\dagger(\vec{x})\hat{\psi}(\vec{x}) \tag{D.16}$$

$$\hat{H}_K = \hat{\kappa}(a)\sum_{\vec{x}}\sum_{j=1}^{d} \eta(\vec{x})\left(\hat{\psi}^\dagger(\vec{x})\hat{U}(\vec{x},j)\hat{\psi}(\vec{x}+\hat{e}_j) + \text{h.c.}\right). \tag{D.17}$$

and the rescaled bare parameters are $\hat{m}(a) = \frac{a}{c}m(a)$ and $\hat{\kappa}(a) = \kappa(a)/c$. Note that the bare speed of light appears in the denominator of the expressions in [113] because the incorrect definition of g_s/g_t was used. Written in this way, it is clear that $a_t = a/c$ is merely an overall scale factor of the Hamiltonian, and the dimensionless Hamiltonian we actually simulate is $\hat{H}_{\text{KS}}$. From this we see that one can only calculate dimensionless energies $\hat{E}_n = a_t E_n$. Because $a_t = a_t(g, c, \hat{m}, \hat{\kappa})$ is an unknown function of the coupling and quark mass, one must determine it through a scale setting procedure, analogous to the previous simple harmonic oscillator example. It is important to note that this logic applies not only to energies, but to any quantity calculated on the lattice. For example, if the operator O has units $[O] = M^b$, then the operator one actually implements in a simulation is $\hat{O} = a^b O$, and the physical result is calculated via $O = [\hat{O}]/a^b$, where the lattice spacing a is assumed to have been already determined via some scale setting procedure.

We now discuss the appearance of the bare parameter $\kappa(a)$ in the fermionic kinetic term, which is analogous to the need for different gauge couplings g_t and g_s multiplying H_E and H_B. The intuition for including $\kappa(a)$ can be drawn from studying lattice gauge theory actions on anisotropic Euclidean lattices, *i.e.*, the spatial lattice spacing a is different from the temporal lattice spacing a_0. For pure gauge theory, this anisotropy requires using different gauge couplings multiplying purely spatial plaquettes and plaqeuttes oriented in the temporal direction; one can equivalently work with a common gauge coupling g and a gauge anisotropy factor γ_g. For fixed g, one can tune γ_g to determine a renormalization trajectory with fixed physical anisotropy $\xi = a/a_0$. By carefully keeping track of these couplings, one can use the transfer matrix formalism to show the need for two independent gauge couplings in the Euclidean theory implies one also needs two different couplings, g_t and g_s, in the pure gauge Kogut-Susskind Hamiltonian.

In a similar way, as discussed in, *e.g.* [35, 283, 284], anisotropic simulations including fermions introduces a bare fermionic anisotropy factor γ_f, controlling the relative size of the fermionic hopping terms in the spatial and Euclidean-time directions. The fermionic anisotropy γ_f is generally different from the gauge anisotropy γ_g, and both parameters need to be tuned such that the renormalized anisotropy $\xi = a/a_0$ "seen" by fermionic and gauge degrees of freedom are equal; this can be done by, *e.g.*, studying the dispersion relation of gluonic states and fermionic states. Analogous to the pure gauge case, by keeping track of the anisotropy factor in the transfer matrix derivation, one finds that the prefactor in front of the fermionic kinetic term $\kappa(a)$ in the Kogut-Susskind Hamiltonian is different from unity and needs to be included in the renormalization procedure.

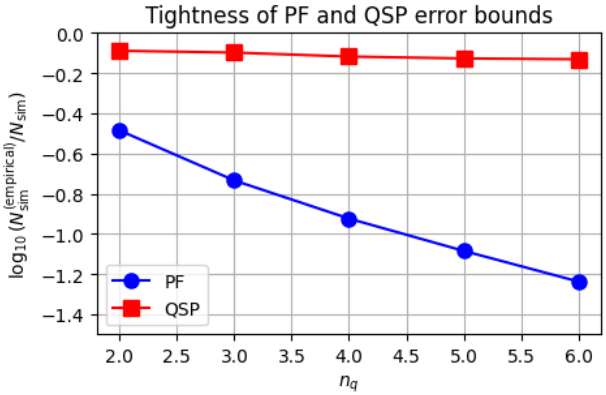

Figure D.1: Comparison of the ratios of the analytically determined algorithmic parameters N_{sim} with its numerically optimized counterpart $N_{\text{sim}}^{(\text{empirical})}$ for PF and QSP at a fixed target exponentiation error $\|\Delta_{\text{sim}}\| = 10^{-4}$ and simulation time $t = 1$.

D.2 Comparing numerical tightness of PF and QSP error bounds

In this section, we compare the numerical tightness of the PF and QSP error bounds used to prove Eqs. (5.56) and (5.68) respectively for the specific choice of a single site anharmonic oscillator with the following Hamiltonian:

$$H = \frac{\phi^2}{2} + \frac{32\phi^4}{34} + \frac{\pi^2}{2} := H_\phi + H_\pi \,, \tag{D.18}$$

where ϕ denotes the field operator and π denotes the conjugate momentum. Given a choice for number of qubits n_q and hence the size of the Hilbert space ($2^{n_q} \times 2^{n_q}$), we discretize H using the scheme presented for instance in [91, 97, 114]. Let N_{sim} be used to denote either of the algorithmic parameters N_{PF} and N_{QSP}. Thus, for a desired exponentiation error $\|\Delta_{\text{sim}}\|$ (see Eq. (5.45)), and simulation time t we can use Eqs. (5.54) and (5.67) to analytically compute sufficient choices N_{sim} for both PF and QSP such that the desired exponentiation error is achieved. We then empirically determine the lowest possible value of N_{sim}, denoted by $N_{\text{sim}}^{(\text{empirical})}$ that results in the target exponentiation error $\|\Delta_{\text{sim}}\|$. Fig. D.1 compares the ratio $N_{\text{sim}}^{(\text{empirical})}/N_{\text{sim}}$ for PF and QSP as a function of the number of qubits n_q used to represent H for the fixed values $\|\Delta_{\text{sim}}\| = 10^{-4}$ and $t = 1$. The key observation we make is that while both ratios deviate from the ideal value of 1 as we increase n_q, the QSP error bound is closer to the numerically determined result than the PF bound. This numerical experiment is thus a simple example illustrating that the QSP error bounds are generally tighter than the PF error bounds used in this work.