

Lawrence Berkeley National Laboratory

LBL Publications

Title

General quantum algorithms for Hamiltonian simulation with applications to a non-Abelian lattice gauge theory

Permalink

<https://escholarship.org/uc/item/2p7551fs>

Authors

Davoudi, Zohreh
Shaw, Alexander F
Stryker, Jesse R

Publication Date

2023-12-20

DOI

10.22331/q-2023-12-20-1213

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

General quantum algorithms for Hamiltonian simulation with applications to a non-Abelian lattice gauge theory

Zohreh Davoudi^{1,2,3,4}, Alexander F. Shaw^{1,3}, and Jesse R. Stryker^{1,2,5}

¹Department of Physics, University of Maryland, College Park, MD 20742, USA

²Maryland Center for Fundamental Physics, University of Maryland, College Park, MD 20742, USA

³Joint Center for Quantum Information and Computer Science, National Institute of Standards and Technology and University of Maryland, College Park, MD 20742, USA

⁴The NSF Institute for Robust Quantum Simulation, University of Maryland, College Park, Maryland 20742, USA

⁵Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

With a focus on universal quantum computing for quantum simulation, and through the example of lattice gauge theories, we introduce rather general quantum algorithms that can efficiently simulate certain classes of interactions consisting of correlated changes in multiple (bosonic and fermionic) quantum numbers with non-trivial functional coefficients. In particular, we analyze diagonalization of Hamiltonian terms using a singular-value decomposition technique, and discuss how the achieved diagonal unitaries in the digitized time-evolution operator can be implemented. The lattice gauge theory studied is the $SU(2)$ gauge theory in 1+1 dimensions coupled to one flavor of staggered fermions, for which a complete quantum-resource analysis within different computational models is presented. The algorithms are shown to be applicable to higher-dimensional theories as well as to other Abelian and non-Abelian gauge theories. The example chosen further demonstrates the importance of adopting efficient theoretical formulations: it is shown that an explicitly gauge-invariant formulation using loop, string, and hadron degrees of freedom simplifies the algorithms and lowers the cost compared with the standard formulations based on angular-momentum as well as the Schwinger-boson degrees of freedom. The loop-string-hadron formulation further retains the non-Abelian gauge symmetry despite the inexactness of the digitized simulation, without the need for costly controlled operations. Such theoretical and algorithmic considerations are likely to be essential in quantumly simulating other complex theories of relevance to nature.

Contents

1	Introduction	2
2	Methods: Strategies for simulating product formulas	4
2.1	Summand diagonalization	6
2.2	Implementing diagonal unitaries	9
2.2.1	Near term	9
2.2.2	Far term	10
3	Application: $SU(2)$ lattice gauge theory in 1+1 D	10
3.1	The Kogut-Susskind framework	10
3.1.1	Schwinger-boson formulation	12
3.1.2	Loop-string-hadron formulation	15
3.2	Circuit decomposition of the propagators	18
3.2.1	Schwinger-boson propagators	18
3.2.2	Loop-string-hadron propagators	31
3.3	Error-bound analysis and simulation cost	39
3.3.1	Schwinger-boson formulation	40
3.3.2	Loop-string-hadron formulation	43

Zohreh Davoudi: davoudi@umd.edu

Alexander F. Shaw: alexandershaw.f@gmail.com

Jesse R. Stryker: jstryker@lbl.gov

3.4 Discussion of the results	45
4 Conclusions and outlook	46
Acknowledgments	48
Appendix A Arithmetic algorithmic routines	50
Appendix B Deriving the loop-string-hadron hopping subterms	58
Appendix C Second-order commutator bounds	59
C.0.1 Schwinger-boson commutator bounds	61
C.0.2 Loop-string-hadron commutator bounds	63
Appendix D Gate-cost tables	64
References	64

1 Introduction

Motivation and brief overview of the work.—A strong case for the promised quantum advantage offered by quantum computing is the simulation of physical systems at an exponentially reduced cost [1–4]. Possibilities are countless for advancing various disciplines of theoretical and applied sciences if robust large-scale fault-tolerant universal quantum hardware becomes a reality. Such possibilities in the area of quantum chemistry and material science have led to a vigorous program in quantum-algorithm design and implementation [5–11]. Furthermore, the promise of substantially speeding up computations with quantum-computing resources has been driving a plethora of quantum-based research and development in nuclear and high-energy physics in recent years [12–15], from first-principles approaches rooted in quantum field theories of nature [13, 16–88] to effective (field) theory descriptions of strongly interacting systems such as nuclei [36, 45, 89–100].

Simulation algorithms on digital quantum computers, regardless of the theory under study, share a number of general features. Most importantly, these require a digitized approximation to the system’s evolution in real time. An immediate advantage of using quantum bits (qubits) is an exponentially more compact encoding of the degrees of freedom (DOFs) compared with classical encodings. Nonetheless, the efficiency of the simulation relies on how the number of costly operations varies as a function of error tolerance, the system’s size, and model’s parameters. While general statements can be made regarding the efficiency of algorithms for local or nearly local interactions [2, 17, 101], only an exact account of the type and the number of operations in connection to the accuracy goal of the computation can determine the viability of the algorithms. Such an analysis is particularly important in light of the limited capacity of the hardware and the imperfect fidelity of quantum operations (gates) in any practical implementation.

With the ultimate goal of quantifying quantum-resource requirements of complex theories of nature described by gauge field theories, in this paper we tackle the following questions: i) how to “best” digitize time evolution in theories with simultaneous changes in various types of quantum numbers due to interactions, ii) how to avoid quantumly evaluating certain operations to facilitate the simulation, and iii) how to take advantage of better formulations of the simulated theory to simplify operations and retain symmetries? The answers, as will be demonstrated, lie in a geometrical intuition, algebraic tricks, classical pre-processing, and rethinking a theory’s formulation of its DOFs and constraints.

Framework and concrete objectives.—Product formulas use the Trotter-Suzuki expansion to decompose the time-evolution operator e^{-iHt} (for a system with Hamiltonian H and evolution time t) into products of efficiently implementable exponentials in various ways [102–104]. For example, the (first-order) Lie-Trotter formula amounts to implementing e^{-iHt} with $H = \sum_{j=1}^{\gamma} H_j$ as

$$V_1(t) \equiv \left[\prod_{j=1}^{\gamma} e^{-itH_j/s} \right]^s, \quad (1)$$

up to an error that scales as $\mathcal{O}(t^2/s)$. Product formulas, given their simplicity, have been the primary digitization method in quantum simulation, and they require no extra qubits beyond what is needed to encode the physical DOFs. There exists a range of other quantum-simulation algorithms that take advantage of techniques such as Taylor series expansion and linear combinations of unitaries [105, 106], quantum signal processing [107],

qubitization and block encodings [108, 109], singular-value transformation [110], off-diagonal Hamiltonian expansion [111], and hybrid algorithms [112]. Compared with product formulas, the gate complexities of these algorithms generally scale better with the error tolerance and the system’s size and parameters asymptotically, but they involve more complicated circuits and often have a non-negligible ancilla-qubit overhead. It is, in fact, known empirically that in certain problems, product formulas perform better than what theoretical bounds on them indicate [101, 113, 114], pointing to the fact that such bounds are generally not tight, or that they involve small pre-factors, which can impact concrete resource estimates. For concreteness, we focus on product formulas in analyzing the time-evolution operator in this work, but many of the ideas to be introduced are applicable to other quantum-simulation methods as well.

Now given the time-evolution algorithm, a first estimate of resource requirements is obtained by counting the number of costly operations needed to guarantee an error tolerance $\epsilon > 0$, defined as $\|V(t) - e^{-iHt}\| \leq \epsilon$, given the system’s size and parameters. Here, $V(t)$ is the product-formula approximation to the exact time-evolution operator and $\|\cdot\|$ denotes the spectral norm. This estimate is not complete as the cost of state preparation and observable measurement need to be included subsequently. However, as many preparation and observable-evaluation routines require implementing e^{-iHt} as an ingredient, the cost estimate of time evolution is a useful indicator of the gate complexity of the full simulation. Two computational models are often considered. In the near term, when fault tolerance is out of reach but noise-mitigation strategies may ameliorate the accuracy loss [3], the number of two-qubit entangling gates (e.g., CNOT gates) needs to be minimized as such gates exhibit lower fidelities. In the far-term, when the gate errors meet certain thresholds, it becomes possible to correct anticipated qubit errors [115]. However, T gates are known to require costly error-correction encodings and hence it is the T-gate count that needs to be minimized in the far term.

Diagonalization, shearing transformation, and singular-value decomposition.—The question of how to best decompose exponentiated operators representing steps of time evolution, namely the “propagators”, in theories with interactions involving numerous DOFs is, for example, of paramount importance in the context of (lattice) gauge theories. There, multiple types of fermionic and bosonic fields may interact, and further the interactions are locally constrained by Gauss’s laws. Consider the propagator $e^{-iH_j t}$ where H_j is a term or appropriate collection of terms in the Hamiltonian, chosen such that $e^{-iH_j t}$ can be decomposed exactly to a universal set of gates. If H_j is a diagonal operator when expressed in the computational basis of qubit registers (that represent the DOFs in the original theory), then single- and two-qubit Pauli-Z rotations provide a complete basis for implementing $e^{-iH_j t}$, making its circuit synthesis rather straightforward. On the other hand, for non-diagonal H_j one needs to proceed with a simultaneous diagonalization in the basis states of all types of quantum registers involved in H_j , which is a non-trivial task in general, and may lead to approximations, hence a potential violation of original symmetries. An example of this is the implementation of the fermion-gauge-boson hopping propagator in the lattice Schwinger model, where the decomposition proposed in Ref. [54], while being efficient, introduces violations of Gauss’s law.

A subsequent work [116] guided by a geometrical interpretation of the transitions in the space of quantum numbers showed that a better diagonalization in the lattice Schwinger model is achievable via a “shearing” transformation, requiring negligible additional quantum resources over the earlier approach, but with the advantage of preserving the local constraints. In Sec. 2, we introduce an algebraic procedure that is a generalization of the shearing transformation. The approach reduces the diagonalizing of H_j and subsequently $e^{-iH_j t}$ to a modest distortion of the exact propagator (if any) through the use of singular-value decompositions (SVDs) of relevant matrices in the space of quantum numbers. While in the example of the Schwinger model, a fully gauge-invariant implementation of the local hopping propagator is achievable, in the example of the non-Abelian SU(2) lattice gauge theory (LGT) coupled to fermions that is studied in this work, the hopping propagator can be efficiently circuitized while maintaining some but not necessarily all of the local constraints, which is still an improvement over existing algorithms [69]. This diagonalization methodology, as will be shown, is applicable to more complex propagators such as those corresponding to magnetic interactions in gauge theories.¹ Furthermore, the procedure is relevant beyond implementing the time-evolution operator. For example, to measure a Hamiltonian’s expectation value, one may diagonalize its summands and estimate the expectation value of each summand individually, then add the results. Minimizing the length of this sum reduces the number of measurements required.

Phase evaluation and classical pre-processing.—After transforming the propagators to a computational basis, the remaining diagonal operations involve, at their core, a set of Z-rotations. In scenarios where the rotation angle (phase) is a constant, implementation amounts to a circuit synthesis with known counts of ancilla qubits and CNOT gates or T gates, given the algorithm used and the synthesis accuracy aimed. On the other hand, in many instances, the phases are non-trivial functions of dynamical quantum numbers, which despite being diagonal in the computational basis of their respective registers, need to be evaluated at each step of the

¹When the gauge DOFs are digitized in the group’s irreducible representation (irrep) basis.

evolution. As is known, while such function evaluations can proceed via a “phase-kickback” algorithm [117] and embedded quantum arithmetic routines [118, 119], they are prohibitively costly in the near term, requiring abundant ancillary registers and controlled operations to evaluate, store, and reprocess function values. This will likely be the biggest bottleneck to quantum simulations of non-Abelian gauge theories, as the non-Abelian algebra (in the irrep basis) involves Clebsch-Gordan coefficients (and generalizations of), which find their way to the exponent of diagonal propagators for both fermion-gauge hopping and the magnetic interactions.

Instead of quantumly evaluating these functions on the go, as a near-term strategy to be described in this work, one can locally decompose them to strings of Pauli- Z operations with now fixed coefficients, and within a fixed tolerance, systematically neglect rotations with small angles. The classical pre-processing involved in obtaining the Pauli decomposition of diagonal functions scales exponentially with the number of qubits per lattice site. It is, therefore, only a function of the cutoff on the gauge-field quantum numbers in the LGT example studied, but is independent of system’s size. Our strategy is similar in spirit to that proposed in Ref. [68] which uses the pre-processed local operations in simulating propagators in a pure SU(3) LGT. Nonetheless, despite the strategy in Ref. [68], our method does not require a hard encoding of the Clebsch-Gordan coefficients in the circuit *a priori* using controlled rotations with pre-set angles, but rather pre-evaluates the functional form for any input, as will be discussed in Sec. 2.2.1.

Basis/formulation considerations and symmetry imposition.—It is known that equivalent expressions of the Hilbert space may lead to different algorithmic complexity in quantum simulation. A known example is the use of position-space versus momentum-space wave functions (see. e.g., Ref. [113]), and the first- or second-quantized formulations (see e.g., Refs. [120–125]) in simulating electronic systems and quantum chemistry. In the context of this work, the question of what is the best basis to represent the DOFs not only concerns the potential reduction of the qubit cost of encoding, but also the simplicity of interactions (e.g., the type and the number of quantum registers to be operated on simultaneously), of phase-function evaluations, and of the expression of symmetries and constraints.

We make the need for wiser choices of basis evident through the SU(2) LGT example considered. In the irrep basis, where the electric Hamiltonian is diagonal at the cost of off-diagonal hopping and magnetic Hamiltonians, the states are characterized by their fermionic and gauge-flux content. The electric-field hence gauge-link DOFs can be represented, for example, in the standard angular-momentum formulation [126], the Schwinger-boson (prepotential) formulation [127–130], or the recently-developed loop-string-hadron (LSH) formulation [131]. As will be demonstrated, while the qubit count and the gate complexity of the diagonalization routine for the various formulations are found to be comparable, the absolute cost of performing diagonal functions is reduced substantially in the LSH formulation.

Furthermore, as already discussed, decomposing the various propagators to a set of gates in product formulas amounts to breaking up a collection of terms in the Hamiltonian that are only gauge invariant together, and this can potentially break the Gauss’s laws throughout the evolution. In the case of the SU(2) theory, this problem is circumvented all together in the LSH formulation that builds the physical Hilbert space *a priori* using a complete set of gauge-invariant local operators [131], eliminating the need for encoding a large unphysical Hilbert space [132]. A remaining link-local Abelian constraint can be easily retained by our diagonalization algorithm, making it possible to achieve an evolution that satisfies both non-Abelian and Abelian Gauss’s law constraints, improving upon quantum algorithms of Ref. [69]. Such a reformulation of the theory also eliminates the need for various symmetry-protection protocols [133–137], Gauss’s law verification circuits [38, 39], or expensive controlled operations [68] to enforce the gauge symmetries. For example, while Ref. [68] introduces a valuable strategy in simulating an SU(3) LGT by constraining local transitions to those satisfying the Gauss’s law using controlled operations, it is conceivable that generalizations of the LSH formulation to the case of SU(3) [138] will simplify the evolution there as well.

Generalizations and future applications.—Finally, in light of lessons learned from the investigation of this work for a non-Abelian LGT in 1+1 dimensions (D), and given the rather general algorithms and strategies proposed, one can explore future directions, with an eye on generalization to theories in 3+1 D, and other gauge groups such as quantum chromodynamics (QCD). Such discussions will follow in Sec. 4, including remarks on the applicability of our approach to other quantum-simulation algorithms, and ideas for improving upon the algorithmic error bounds obtained in this work.

2 Methods: Strategies for simulating product formulas

The goal of this section is to introduce generic methods to derive quantum circuits that approximate the unitary e^{-itH} , where H is a Hermitian operator acting simultaneously on multiple qubit registers and $t > 0$ is a real parameter. H could represent the full Hamiltonian of a physical system or a term (or collection of terms) that

is exponentiated separately in a product-formula approximation to the evolution operator.

Given a product formula, the quantum circuit that implements it can be derived using the identity

$$e^{-itH} = \mathcal{U}^\dagger e^{-it\mathcal{D}} \mathcal{U}, \quad (2)$$

where \mathcal{U} is a unitary that diagonalizes H and \mathcal{D} is the diagonalized form of H in the computational basis of the qubit registers. This reduces the task of simulating H to the task of providing quantum circuits for \mathcal{U} and $e^{-it\mathcal{D}}$, which turned out to be more tractable. An overarching task of quantum simulation is to implement Eq. (2) to within some fixed accuracy ϵ . The error can be considered as the spectral norm of the difference between the approximate and ideal unitaries, which is a useful quantity to bound the error in observables given an arbitrary initial state [54]. The error bound can be improved by assuming a specific input state or an input-state distribution [139–141] but such improvements will not be considered here.

When deriving quantum circuits for e^{-itH} via Eq. (2) for a Hamiltonian $H = \sum_j H_j = \sum_j \mathcal{U}_j^\dagger \mathcal{D}_j \mathcal{U}_j$, a balance must be found in minimizing i) the cost of implementing exponential of each summand H_j , ii) the number of costly summand exponentials, and iii) the Trotter error. This is because some \mathcal{U}_j and $e^{-it\mathcal{D}_j}$ could dominate the quantum-computational cost of the simulation, so it is important to split the Hamiltonian to the summands H_j properly. This split, on the other hand, should be informed by Trotter-error considerations since this error depends on the commutators among the summands. For example, consider $H_j = H_j^{(1)} + H_j^{(2)}$. It may be that the sum cost of implementing $\mathcal{U}_j^{(1)}$, $e^{-it\mathcal{D}_j^{(1)}}$, $\mathcal{U}_j^{(2)}$, and $e^{-it\mathcal{D}_j^{(2)}}$ is higher than that of \mathcal{U}_j and $e^{-it\mathcal{D}_j}$, in which case, it is beneficial to not split H_j , as long as diagonalization of H_j and an inexpensive implementation of \mathcal{U}_j and $e^{-it\mathcal{D}_j}$ can be found. If instead, $\mathcal{U}_j^{(1)}$, $e^{-it\mathcal{D}_j^{(1)}}$, $\mathcal{U}_j^{(2)}$, and $e^{-it\mathcal{D}_j^{(2)}}$ can be implemented more economically, it may make sense to consider the split to $H_j^{(1)}$ and $H_j^{(2)}$, notwithstanding this will increase the Trotter error if $[H_j^{(1)}, H_j^{(2)}] \neq 0$. Another consideration in choosing the summands is minimizing the symmetry violation in Trotter evolution, which may offer some benefits. For example, the more symmetries kept throughout the evolution, the more error diagnostics and noise-mitigation tools at one’s disposal to (partially) verify and/or correct noisy intermediate-scale quantum simulations.

Analytically optimizing this problem in search for the best choice of summands is not trivial in general, nor is numerically finding the optimized choice given the sheer dimensionality of the operators. While systematic strategies such as efficient numerical approaches for navigating this complex optimization problem are desired, finding the balance among the goals mentioned here may be possible on a case-by-case basis. As will be discussed in the example of the non-Abelian LGT studied in Sec. 3, the knowledge of the Hamiltonian structure and symmetries, and the commutator algebra involved, allow for decomposition choices that are more optimal than the others.

Given the summand H_j , one may wonder if classical-computing methods can be used to circuitize either e^{-itH_j} directly or \mathcal{U}_j and $e^{-it\mathcal{D}_j}$. Circuitizing e^{-itH_j} directly generally amounts to finding the Pauli decomposition of H_j . Since, in general, H_j is not diagonal in the computational basis of the qubits, one would need to determine $4^{\lceil \log_2(d) \rceil}$ coefficients, with d being the dimensionality of the relevant Hilbert space. Thus, this method is costly for large Hilbert-space sizes, and is not scalable. Furthermore, simulating e^{-itH_j} by Pauli decomposition introduces significant Trotter error, as e^{-itH_j} must be simulated through applications of a product formula, and Pauli strings do not necessarily commute. As a result, the Pauli-decomposition method should not be considered as the method of choice whenever alternative methods can be found that are exact and use circuits of comparable or lower cost.

The method based on the diagonalized form of H_j , first of all, requires diagonalizing H_j to find \mathcal{U}_j and \mathcal{D}_j , a problem that classically scales poorly with the dimensionality of the Hilbert space. For k -sparse Hamiltonians with $k \ll d$, which is the case for most physical Hamiltonians of interest, H_j only acts on a small part of the Hilbert space and efficient numerical methods can ameliorate the scaling problem. With the proper choice of H_j for local or semi-local Hamiltonians, such a diagonalization may be achievable far more efficiently numerically, or even analytically, as will be demonstrated in the method of this work. Implementing \mathcal{U}_j can take advantage of classical circuit-synthesis methods but the bottleneck is dealing with large dense matrices. $e^{-it\mathcal{D}_j}$ can be similarly circuitized using a Pauli decomposition but since it is a diagonal unitary, it requires obtaining $2^{\lceil \log_2(d) \rceil}$ coefficients, as only tensor products of Pauli- Z and identity operators are needed, see Sec. 2.2. More importantly, the implementation of these strings introduces no Trotter error. Numerical strategies for finding a circuit that simulates diagonal Hermitian operators are simpler than they are for arbitrary Hermitian matrices, and so in the far-term, simpler logic synthesis² can be used to find a circuit decomposition of $e^{-it\mathcal{D}_j}$, as will be detailed in the

²We are using the term ‘circuit synthesis’ to indicate a classical routine that decomposes a unitary matrix defined on a set of qubits to a set of quantum gates that effect the exact (or an approximation of the) matrix, while the term ‘logic synthesis’ is meant to imply a program that converts an abstract specification of the desired circuit behavior into a circuit implementation in terms of a set of logic gates [142]. An example of the former is the Pauli decomposition of diagonal unitaries for near-term applications, and an example of the latter is the implementation of diagonal unitaries using a phase-kickback algorithm and Newton’s method-based function evaluation for fault-tolerant applications, both discussed in Sec. 3.2.




Figure 1: Diagonalization of $A + A^\dagger$, where $A^2 = 0$. ‘Direct diagonalization’ of $A + A^\dagger$ is expressed by $A + A^\dagger \xrightarrow{\mathcal{U}} \mathcal{D}$, where \mathcal{D} is diagonal in the computational basis, \mathcal{U} is unitary, and the notation $O \xrightarrow{\mathcal{U}} O'$ means $O' = \mathcal{U} O \mathcal{U}^\dagger$. We suppose that an SVD of A is given by $A = \mathcal{V} \mathcal{S} \mathcal{W}^\dagger$. The property $A^2 = 0$ also implies $A = \mathcal{P}_0 A \mathcal{P}_1$, where \mathcal{P}_0 and \mathcal{P}_1 are projectors onto orthogonal subspaces \mathbb{H}_0 and \mathbb{H}_1 of the full Hilbert space $\mathbb{H} = \mathbb{H}_0 \oplus \mathbb{H}_1$. Introducing an ancillary qubit labelled as \mathbf{x} and tensoring it with the original space \mathbb{H} , unitaries \mathcal{P} and \mathcal{Q} are defined such that $\mathcal{P} |0\rangle_{\mathbf{x}} (\mathcal{P}_b |\psi\rangle) = |b\rangle_{\mathbf{x}} (\mathcal{P}_b |\psi\rangle)$ (for $b \in \{0, 1\}$) and $\mathcal{Q} = |0\rangle\langle 0|_{\mathbf{x}} \mathcal{V}^\dagger + |1\rangle\langle 1|_{\mathbf{x}} \mathcal{W}^\dagger$. The operators at the right end of the diagram are diagonal in the computational bases.

LGT example in Sec. 3. In summary, if the structure of the Hamiltonian and the proper choice of summands allow an analytical determination of \mathcal{U}_j , \mathcal{D}_j , and the circuit decomposition of \mathcal{U}_j , the only remaining task is circuitizing $e^{-it\mathcal{D}_j}$, which can benefit from less demanding classical pre-processing approaches and more straightforward circuit-synthesis methods in both near and far terms.

The following section details an approach to selecting H_j in order to derive analytical circuits that diagonalize them. It is a useful tool for navigating the goals stated earlier in this section for certain Hamiltonians.

2.1 Summand diagonalization

In this section, we introduce a strategy based on linear algebra for breaking given Hamiltonians into a sum of terms that can be diagonalized by simple quantum circuits. It will become clear through the examples that will follow which class of Hamiltonians can benefit from the proposed method.

First suppose that a summand H_j can be written as $A + A^\dagger$ such that $A^2 = A^{\dagger 2} = 0$. Equivalently, let the Hilbert space \mathbb{H} associated with the linear map A be expressed as the direct sum $\mathbb{H}_0 \oplus \mathbb{H}_1$, where \mathbb{H}_0 is the nullspace of A , i.e., $\mathbb{H}_0 \equiv \ker(A)$, and \mathbb{H}_1 is its complement space, i.e., $\mathbb{H}_1 \equiv (\ker(A))^\perp$. Then A can be written as $A = \mathcal{P}_0 A \mathcal{P}_1$, where \mathcal{P}_b is the projector to subspace \mathbb{H}_b for $b \in \{0, 1\}$. Similarly, $A^\dagger = \mathcal{P}_1 A^\dagger \mathcal{P}_0$. If deducing the unitary \mathcal{U} that diagonalizes $A + A^\dagger$ via $A + A^\dagger = \mathcal{U}^\dagger \mathcal{D} \mathcal{U}$ is not straightforward, or is known but costly to implement in a quantum circuit, one may alternatively proceed by using an SVD of A instead, which for square operators that are considered here is a diagonal square matrix of real non-negative elements, and should be easier to circuitize than the original non-diagonal form. This is provided that the SVD unitaries \mathcal{V} and \mathcal{W} , defined as $\mathcal{S} = \mathcal{V}^\dagger A \mathcal{W}$ ($= \mathcal{S}^\dagger = \mathcal{W}^\dagger A^\dagger \mathcal{V}$) for the singular-value matrix \mathcal{S} , can be found easily and their circuit implementation is efficient. As will be demonstrated shortly, such efficient SVD can be worked out via simple quantum circuits for certain (common) Hamiltonians.

Assuming that the SVD matrices \mathcal{V} and \mathcal{W} are found and can be implemented straightforwardly, the general quantum circuit that leads to \mathcal{S} can be formed as follows. One may first introduce an ancillary qubit \mathbf{x} which is prepared in the state $|0\rangle_{\mathbf{x}}$. The operator \mathcal{P} , defined via $\mathcal{P} |0\rangle_{\mathbf{x}} |\psi\rangle_{\mathbb{H}_b} = |b\rangle_{\mathbf{x}} |\psi\rangle_{\mathbb{H}_b}$ for $b \in \{0, 1\}$, can be introduced to apply the transformation $\mathcal{P} [|0\rangle\langle 0|_{\mathbf{x}} (A + A^\dagger)] \mathcal{P}^\dagger = |0\rangle\langle 1|_{\mathbf{x}} A + |1\rangle\langle 0|_{\mathbf{x}} A^\dagger$. Next the operator \mathcal{Q} , defined as $\mathcal{Q} = |0\rangle\langle 0|_{\mathbf{x}} \mathcal{V}^\dagger + |1\rangle\langle 1|_{\mathbf{x}} \mathcal{W}^\dagger$, applies the transformation $\mathcal{Q} (|0\rangle\langle 1|_{\mathbf{x}} A + |1\rangle\langle 0|_{\mathbf{x}} A^\dagger) \mathcal{Q}^\dagger = X_x \mathcal{S}$, where $X_x = |1\rangle\langle 0|_{\mathbf{x}} + |0\rangle\langle 1|_{\mathbf{x}}$ is the Pauli- X operator acting in the Hilbert space of the ancillary qubit. Finally, a Hadamard transformation H_x on the ancillary qubit leads to $H_x X_x \mathcal{S} H_x = Z_x \mathcal{S}$, where $Z_x = |0\rangle\langle 0|_{\mathbf{x}} - |1\rangle\langle 1|_{\mathbf{x}}$ is the Pauli- Z operator acting on the ancillary qubit. This final form is, therefore, diagonal in both the ancillary-qubit Hilbert space and in \mathbb{H} , as desired. These steps are schematically shown in Fig. 1.

If the SVD unitaries are unknown, or not implementable by cheap quantum circuits, it may be possible to further split $A + A^\dagger$ to $\sum_k (A_k + A_k^\dagger)$ such that simple diagonalizing circuits can be found for each $A_k + A_k^\dagger$. However, such further splitting increases the total number of unitary operations that are required to construct the product formula. It may also increase the Trotter error. As discussed before, an optimized splitting would balance circuit cost and error tolerance.

In the following, we demonstrate two examples from an Abelian LGT that can take advantage of the SVD algorithm above, as well as an example of a banded Hamiltonian, to give an idea of the general characteristics

of the Hamiltonians that may benefit from the method of this section. The case of SU(2) LGT that is examined thoroughly in Sec. 3 provides another example, suggesting that the SVD algorithm will have wide applicability to simulating LGT Hamiltonians.

- ◇ *Example 1: Coupled bosonic and fermionic incrementers.* This interaction type describes the hopping Hamiltonian in the U(1) LGT in the staggered formulation [126]. Its form can be generalized to other Abelian and non-Abelian LGTs upon appropriate modifications to the type and/or the number of fermionic and bosonic operators, see e.g., Sec. 3 for the example of a SU(2) LGT. After mapping the nearest-neighbor fermionic interaction in the U(1) LGT to qubits, the Hamiltonian on a single link can be written as

$$H_{\text{hop}} = |0\rangle\langle 1|_{\mathbf{x}} |1\rangle\langle 0|_{\mathbf{y}} U_{\mathbf{p}} + \text{H.c.}, \quad (3)$$

where \mathbf{x} and \mathbf{y} denote the two-dimensional Hilbert spaces of the qubit registers associated with fermions at two adjacent sites, and \mathbf{p} denotes a collection of qubits that encode the bosonic Hilbert space. $U_{\mathbf{p}}$ is a ladder operator that acts on the bosonic space \mathbf{p} , and is defined as $U_{\mathbf{p}} \equiv \sum_{j=-\infty}^{\infty} c_j |j-1\rangle\langle j|_{\mathbf{p}}$, with j being an integer and $c_j = 1$ for the U(1) LGT. These coefficients may generally depend on j as is the case in the SU(2) LGT. They do not affect the diagonalization procedure to be outlined here but are relevant when the resulting diagonal operator is to be implemented, as will be discussed in Sec. 2.2.

Since $(|1\rangle\langle 0|_{\mathbf{y}} U_{\mathbf{p}})^2 = 0$, one can use the procedure of this section to diagonalize $e^{-itH_{\text{hop}}}$. Here, the ancillary qubit and the \mathcal{P} operator are not needed, since the presence of $|0\rangle\langle 1|_{\mathbf{x}}$ and its Hermitian conjugate in the Hamiltonian provides the form needed for the application of the SVD routine. Noting that $X_{\mathbf{y}} |1\rangle\langle 0|_{\mathbf{y}} \mathcal{I}_{\mathbf{y}} = |0\rangle\langle 0|_{\mathbf{y}}$ and $\lambda_{\mathbf{p}}^+ |j-1\rangle\langle j|_{\mathbf{p}} \mathcal{I}_{\mathbf{p}} = |j\rangle\langle j|_{\mathbf{p}}$, with the incrementer/decrementor operator defined as $\lambda_{\mathbf{p}}^{\pm} |j\rangle_{\mathbf{p}} = |j \pm 1\rangle_{\mathbf{p}}$, and \mathcal{I} being the identity operator on the corresponding registers, one arrives at

$$|1\rangle\langle 0|_{\mathbf{y}} U_{\mathbf{p}} + \text{H.c.} = X_{\mathbf{y}} \lambda_{\mathbf{p}}^- (|0\rangle\langle 0|_{\mathbf{y}} \mathcal{D}_{\mathbf{p}}) \mathcal{I}_{\mathbf{y}} \mathcal{I}_{\mathbf{p}} + \text{H.c.}, \quad (4)$$

where $\mathcal{D}_{\mathbf{p}} \equiv \sum_{j=-\infty}^{\infty} c_j |j\rangle\langle j|_{\mathbf{p}}$. As a result, the singular-value unitaries \mathcal{V} and \mathcal{W} can be identified as $\mathcal{V} = X_{\mathbf{y}} \lambda_{\mathbf{p}}^-$ and $\mathcal{W} = \mathcal{I}_{\mathbf{y}} \mathcal{I}_{\mathbf{p}}$. Then according to the procedure depicted in Fig. 1, the Hamiltonian in Eq. (3) is diagonalized as:

$$H_{\text{hop}} = \mathcal{W}^\dagger (Z_{\mathbf{x}} |0\rangle\langle 0|_{\mathbf{y}} \mathcal{D}_{\mathbf{p}}) \mathcal{W}, \quad (5)$$

with the diagonalizing transformation

$$\mathcal{W} = H_{\mathbf{x}} (|0\rangle\langle 0|_{\mathbf{x}} X_{\mathbf{y}} \lambda_{\mathbf{p}}^+ + |1\rangle\langle 1|_{\mathbf{x}} \mathcal{I}_{\mathbf{y}} \mathcal{I}_{\mathbf{p}}). \quad (6)$$

On a quantum computer, the \mathcal{W} operation involves basic addition primitives and Pauli gates.³ In practice, the bosonic Hilbert space must be truncated and the incrementer (and decrementors) must be modified at the edge of the Hilbert space, requiring implementing modular additions (and subtractions), which are known operations in quantum circuitry. Such details will be dealt with more closely when we present algorithms for the SU(2) LGT in various formulations.

The decomposition for the U(1) hopping term presented in this example is exact, and hence the $e^{-itH_{\text{hop}}}$ operator can be implemented without violating the local Gauss's law. This feature was absent in the algorithm of Ref. [54], where the $U_{\mathbf{p}}$ operator was split to the sum of two terms, upon a (non-exact) periodic wrapping, and the exponential of each of the terms was shown to be implementable on a quantum computer efficiently. While the periodic wrapping can be mitigated at the cost of introducing multi-controlled operations, the splitting of the hopping term introduces Gauss's-law-violating Trotter errors even for a single-link term. On the other hand, the improved algorithm of Ref. [116] introduces shear transformations in the space of fermion and boson quantum numbers to effectively move the non-trivial operation onto one of the three registers (\mathbf{x} , \mathbf{y} , or \mathbf{p}), and implements the U(1) hopping term locally in an exact manner. This geometric picture inspired the present algorithm which finds the shear transformations systematically using an SVD.

- ◇ *Example 2: Multiple coupled bosonic incrementers along the edges of a square.* This type of interaction, called a plaquette interaction, corresponds to the magnetic Hamiltonian in a U(1) LGT in higher than one spatial dimensions. The strategy for diagonalizing this term can be generalized straightforwardly to other LGTs. The single-plaquette Hamiltonian is:

$$H_{\text{plaq.}} = U_{\mathbf{p}} U_{\mathbf{q}} U_{\mathbf{s}}^\dagger U_{\mathbf{t}}^\dagger + \text{H.c.}, \quad (7)$$

where \mathbf{p} , \mathbf{q} , \mathbf{s} , \mathbf{t} denote the bosonic qubit registers associated with the four links of the plaquette, and the link operators are defined as in the previous example. Here, the condition $(U_{\mathbf{p}} U_{\mathbf{q}} U_{\mathbf{s}}^\dagger U_{\mathbf{t}}^\dagger)^2 = 0$ does not hold

³For the projectors in Eq. (6), note that $Z = |0\rangle\langle 0| - |1\rangle\langle 1|$, so $|0\rangle\langle 0| = \frac{I+Z}{2}$ and $|1\rangle\langle 1| = \frac{I-Z}{2}$.

in general, and so an exact simulation of this terms via the SVD algorithm will not be possible. However, as mentioned before, one can keep splitting the operator such that the resulting subterms satisfy the required condition. In the case of the plaquette interaction, only one such splitting is needed to achieve the desired form for the application of the SVD algorithm.

Let $\mathcal{E}_p = \sum_{k=-\infty}^{\infty} |2k\rangle \langle 2k|_p$ (with integer k) and $\mathcal{O}_p = \mathcal{I}_p - \mathcal{E}_p$ be the projection operators onto the even and odd quantum numbers in the p register, respectively. With these projectors, the plaquette Hamiltonian can be written as

$$H_{\text{plaq.}} = H_{\text{plaq.}}^{(e)} + H_{\text{plaq.}}^{(o)}, \quad (8)$$

with

$$H_{\text{plaq.}}^{(e)} = (\mathcal{O}_p U_p \mathcal{E}_p) U_q U_s^\dagger U_t^\dagger + \text{H.c.}, \quad (9a)$$

$$H_{\text{plaq.}}^{(o)} = (\mathcal{E}_p U_p \mathcal{O}_p) U_q U_s^\dagger U_t^\dagger + \text{H.c.} \quad (9b)$$

$H_{\text{plaq.}}^{(e)}$ and $H_{\text{plaq.}}^{(o)}$ then will be exponentiated separately in the product formula. Let us apply the SVD algorithm to diagonalize $H_{\text{plaq.}}^{(e)}$. Diagonalization of $H_{\text{plaq.}}^{(o)}$ follows analogously upon $\mathcal{E} \leftrightarrow \mathcal{O}$.

First note that $\mathcal{O}_p U_p \mathcal{E}_p$ in the first term of Eq. (9a) and $(\mathcal{O}_p U_p \mathcal{E}_p)^\dagger$ in the Hermitian conjugate term act, respectively, on two disjoint Hilbert spaces, \mathbb{H}_1 for the even quantum numbers and \mathbb{H}_0 for the odd quantum numbers of the p register. Hence, we identify $\ker(\mathcal{O}_p U_p \mathcal{E}_p) = \mathbb{H}_0$. After introducing an ancillary qubit x , the operator $|0\rangle \langle 0|_x H_{\text{plaq.}}^{(e)}$ is in exactly the form on which the SVD algorithm can be applied, starting with the transformation \mathcal{P} as defined before. The \mathcal{Q} transformation depends on the singular-value unitaries \mathcal{V} and \mathcal{W} , which are easy to guess given an SVD of the link operator as obtained in the previous example. One then finds that

$$(\mathcal{O}_p U_p \mathcal{E}_p) U_q U_s^\dagger U_t^\dagger = \lambda_p^- \lambda_q^- [(\mathcal{D}_p \mathcal{E}_p) \mathcal{D}_q \mathcal{D}_s \mathcal{D}_t] \lambda_s^+ \lambda_t^+, \quad (10)$$

with the incrementer and decrementor operators defined in the previous example. With the identification of $\mathcal{V} = \lambda_p^- \lambda_q^-$ and $\mathcal{W} = \lambda_s^- \lambda_t^-$, the diagonalization proceeds as:

$$|0\rangle \langle 0|_x H_{\text{plaq.}}^{(e)} = \mathcal{W}^\dagger (Z_x (\mathcal{D}_p \mathcal{E}_p) \mathcal{D}_q \mathcal{D}_s \mathcal{D}_t) \mathcal{W}, \quad (11)$$

with the diagonalizing operator \mathcal{W} fully specified:

$$\mathcal{W} = H_x (|0\rangle \langle 0|_x \lambda_p^+ \lambda_q^+ \mathcal{I}_s \mathcal{I}_t + |1\rangle \langle 1|_x \mathcal{I}_p \mathcal{I}_q \lambda_s^+ \lambda_t^+) \mathcal{P}. \quad (12)$$

This unitary can be implemented on a quantum computer using standard operations. The modifications arising from the truncated Hilbert space of the links can be dealt with similar to the hopping-term example, which will be discussed in detail in Sec. 3.

There exists another method for diagonalizing $e^{-itH_{\text{plaq.}}}$ in the U(1) LGT [58, 69]. The U(1) LGT, when truncated in the irrep basis, maintains the group structure in the group-element basis and is isomorphic to some Z_n group. Therefore, one can use a quantum Fourier transform (QFT) over Z_n (for which many quantum circuits exist) to diagonalize the truncated U(1) plaquette, but only if one allows direct non-zero transitions between the positive and negative cutoff states. When the proper Fourier transform is not known, like with the continuous non-Abelian groups, or the unphysical transitions are to be avoided, the strategy presented here will be advantageous. Note that the algorithm of this work requires breaking only one of the link operators and, therefore, introduces fewer Trotter commutators compared with an algorithm presented in Ref. [69], in which all four link operators are split and implemented separately.

- ◇ *Example 3: Banded Hamiltonians.* Suppose Hamiltonian H is a $2^N \times 2^N$ matrix in a basis denoted $|n\rangle$, for $n = 0, 1, \dots, 2^N - 1$. Furthermore, suppose that H has non-zero entries only along the off-diagonal bands up to a modular distance k from the main diagonal, that is, k is the minimum non-negative integer such that $\langle m|H|n\rangle = 0$ whenever $|(m-n) \bmod 2^N| > k$. Then H can be decomposed into a sum of $2k+1$ terms which are systematically diagonalizable. To see this, consider that

$$H = \mathcal{D}_0 + \sum_{j=1}^k \mathcal{D}_j (\lambda^+)^j + \text{H.c.}, \quad (13)$$

for some diagonal matrices \mathcal{D}_j . Since the dimension is 2^N , one can argue that, for any $j > 0$, the operator $(\lambda^+)^j$ maps computational basis states across some bipartition of the Hilbert space. This bipartition can be

seen in the action of addition by j modulo 2^N in the qubit position corresponding to the least significant, non-zero bit of j . In graph-theoretical terms, let the computational basis vectors define a set of vertices, where vertices m and n are connected by an edge if $m = (n \pm j) \bmod 2^N$. For all $j > 0$, addition modulo 2^N generates an even-length cycle, since a theorem from number theory guarantees the cycle length must divide 2^N .

The even and odd steps along this cycle correspond to the bipartition. This bipartition is similar to the bipartition in Example 2, and results in a similar diagonalization procedure, where each $\mathcal{D}_j(\lambda^+)^j + \text{H.c.}$ for $j > 0$ is expressed as a sum of two diagonalizable summands. Therefore, H can be expressed as a sum of $2k + 1$ systematically diagonalizable summands.

2.2 Implementing diagonal unitaries

Once $e^{-iH_j t}$ is diagonalized via the strategy of the previous section, the remaining task is to implement a diagonal operator of the form

$$e^{-it\mathcal{D}(\hat{n}_1, \hat{n}_2, \dots, \hat{n}_\gamma)}, \quad (14)$$

where each \hat{n}_j is a number operator, which is diagonal in the computational basis. Two generic avenues can be explored to construct quantum circuits which implement such a unitary. They may be identified as near-term and far-term strategies, as the former emphasizes the use of multi-qubit diagonal rotations, while the latter involves reversibly computing $\mathcal{D}(\hat{n}_1, \hat{n}_2, \dots, \hat{n}_\gamma)$ using ancillary registers. Nonetheless, these methods are not mutually exclusive. One may blend them depending on the target unitary and available computational resources.

2.2.1 Near term

With the computational register in binary, the number operator on register j with η_j qubits can be written as $\hat{n}_j = \sum_{k=0}^{\eta_j-1} 2^k \hat{n}_{j,k}$, where $\hat{n}_{j,k}$ is the number operator of the k^{th} qubit of register j , returning a value 0 or 1. Arbitrary Hermitian diagonal matrices in $\mathbb{C}_{2^N} \times \mathbb{C}_{2^N}$ can be decomposed into tensor products of Pauli $Z_{j,k} (= \mathcal{I}_{j,k} - 2n_{j,k})$ matrices,

$$\begin{aligned} \mathcal{D}(\hat{n}_1, \hat{n}_2, \dots, \hat{n}_\gamma) \\ \equiv c_{\mathcal{I}} \mathcal{I} + c_{\{1,1\}} Z_{1,1} + \dots + c_{\{1,2\},\{1,3\}} Z_{1,2} Z_{1,3} + \dots + c_{\{1,2\},\{1,3\},\{1,4\}} Z_{1,2} Z_{1,3} Z_{1,4} + \dots, \end{aligned} \quad (15)$$

which is equivalent to the Walsh series of the function [143]. With at most 2^N non-zero constant coefficients $c_{\{j,k\}, \dots}$, the diagonal unitary $e^{-it\mathcal{D}(\hat{n}_1, \hat{n}_2, \dots, \hat{n}_\gamma)}$ can be implemented straightforwardly by exponentiating terms in this expansion individually. Such an implementation does not introduce any Trotter error as any pairs of terms in Eq. (15) commute. The coefficients $c_{\{j,k\}, \dots}$, sometimes known as Walsh coefficients, are rotation angles that can be determined via a classical pre-processing by solving the following equation

$$c_{\{j,k\},\{i,l\}, \dots} = \frac{1}{2^N} \text{Tr}(\mathcal{D} Z_{j,k} Z_{i,l} \dots). \quad (16)$$

A fast-Walsh-transform algorithm, for example, can compute the coefficients in $\mathcal{O}(N 2^N)$ floating-point operations [144].

Finally, the exponentiation of each term is performed by realizing the following identity:

$$e^{icZ_{j_1} \otimes \dots \otimes Z_{j_p}} = e^{ic\pi_{j_1, j_2, \dots, j_p}}. \quad (17)$$

Here, $\pi_{j_1, j_2, \dots, j_p}$ is 1 (-1) if the state of the qubit string composed of registers $\{j_1, j_2, \dots, j_p\}$ has an even (odd) number of $|1\rangle$ states. So this operation is done by a single Z rotation on a qubit storing the parity of p qubits.⁴ However, computing the parity of p qubits takes $p - 1$ entangling CNOT gates. Additionally, at worst, there are 2^N such rotations, where N is the total number of qubits, i.e., $N = \sum_{j=1}^{\gamma} \eta_j$. Therefore, while optimized algorithms to reduce the number of entangling gates exist in certain cases [54, 145, 146], quantum circuits for a diagonal operator remain costly if N is large and \mathcal{D} involves many non-zero coefficients in Eq. (15). This means that if one wants to avoid saturating the upper bound of 2^N sufficient rotations, either the function \mathcal{D} better have a small number of terms in its Z -string decomposition, or an approximation to \mathcal{D} is made such that small rotations are dropped, trading accuracy for lower cost of implementation on a quantum computer. The problem of finding the minimal-length Walsh-series approximation to a function with discrete argument, and the analytic bound on the error made, has been addressed in the context of quantum simulation in recent years, see e.g., Refs. [145, 146]. In Sec. 3, we provide empirical conclusions regarding such approximations in the context of simulating time dynamics of a non-Abelian LGT.

⁴This qubit can be one of the main qubit registers and so no ancillary qubit is necessary.

2.2.2 Far term

In the far term, one may avoid the potentially exponential scaling of Z rotations (which are assumed to be costly in this scenario) by unitarily computing $|\mathcal{D}(n_1, n_2, \dots, n_\gamma)\rangle$ from the registers $|n_1\rangle|n_2\rangle\cdots|n_\gamma\rangle$. One then extracts the computed phase with a number of single-qubit Z rotations equal to the bit precision desired, and then finishes by uncomputing $|\mathcal{D}(n_1, n_2, \dots, n_\gamma)\rangle$. This procedure is well known as “phase kickback”, and a detailed application of it is presented in Sec. 3. Depending on the form of the phase functions, a range of classical arithmetic algorithms can be generalized to quantum algorithms to enable the phase-kickback protocol. Newton’s method, addition, and multiplication-table algorithms are among the routines used to construct the diagonal phase functions in the non-Abelian LGT example of the next section. A brief description of these routines is presented in Appendix A.

3 Application: SU(2) lattice gauge theory in 1+1 D

Hamiltonian simulation of the SU(2) LGT has been the focus of theory and algorithmic developments in recent years, from theoretical studies to cast it in more suitable representations [129–132, 147–151], to the first tensor-network simulations of its static and dynamical properties [152–154], to the first quantum-simulation algorithms and experiments to study its spectrum and evolution [43, 62, 63, 69]. In the context of this work, the SU(2) LGT provides an ideal example for the application of the diagonalization and phase-evaluation methods introduced in the previous section: it exhibits interactions involving changes in several quantum numbers with non-trivial (functional) coefficients.

The technical details of any quantum simulation algorithm for a LGT are foremost decided by the choice of formulation, digitization of continuous bosonic DOFs, and the implementation of lattice fermions. Among the formulations of the SU(2) LGT [132] is the Kogut-Susskind formulation [126] and the Schwinger-bosons and LSH forms which are derived from it. Here, we limit our analysis to an electric eigenbasis given the naturally discrete nature of the eigenvalues in the electric basis, and their suitability in expressing Gauss’s laws, which are constraints on the local flux of electric fields.

For the SU(2) LGT in the Kogut-Susskind formulation, the electric basis is also known as the irrep or angular-momentum basis. The angular-momentum basis has been digitized and studied in Ref. [69] following algorithms developed for the U(1) LGT in Ref. [54]. Here, we apply the new algorithmic approach of this work, with added benefits, to two competing formulations: the Schwinger-boson formulation and its derivative, LSH. Only a 1+1-dimensional theory will be studied in detail although the algorithms of this work are equally applicable to higher-dimensional theories, as was demonstrated for the example of a U(1) magnetic Hamiltonian in the previous section. The ultimate goal of this section is to arrive at a rigorous comparison of resource requirements in each formulation, and to evaluate to what degree the symmetries are preserved in each simulation.

3.1 The Kogut-Susskind framework

In 1+1 D, the Kogut-Susskind Hamiltonian describing SU(2) gauge fields interacting with one flavor of staggered fermions is given by [126]

$$\hat{H} = \hat{H}_M + \hat{H}_E + \hat{H}_I. \quad (18)$$

The simplest contribution is the fermion self-energy,

$$\hat{H}_M = \mu \sum_{r=0}^{L-1} (-1)^r \hat{\psi}^\dagger(r) \hat{\psi}(r), \quad (19)$$

where $\hat{\psi} = \begin{pmatrix} \hat{\psi}_1 \\ \hat{\psi}_2 \end{pmatrix}$ is an SU(2) doublet in the fundamental representation and each component of $\hat{\psi}$ is a one-component field that satisfies fermionic statistics.⁵ The alternating sign $(-1)^r$ corresponds to the usage of staggered fermions. Next, the electric Hamiltonian \hat{H}_E is associated with the energy stored in the left, $\hat{E}_i^L(r)$, and right, $\hat{E}_i^R(r)$, electric fields defined on the link connecting site r to $r+1$. The index $i = 1, 2, 3$ corresponds to the three generators of SU(2). On each link, these satisfy an “Abelian Gauss’s Law” (AGL) condition:⁶ $\sum_{i=1}^3 (\hat{E}_i^L(r))^2 \equiv (\hat{E}^L(r))^2 = \sum_{i=1}^3 (\hat{E}_i^R(r))^2 \equiv (\hat{E}^R(r))^2$. In words, the Casimirs at each end of the link are

⁵Here and in the following, the site dependence of operators, states, and quantum numbers will be dropped for brevity, unless its specification is necessary for clarity.

⁶Strictly speaking, this is not a ‘Gauss’s law’ as it does not concern the flux of electric field at sites, and rather enforces the SU(2) group property at the links connecting the sites. For convenience, we choose to call this an ‘AGL’ throughout.

equal. Finally, the electric energy is directly expressed in terms of these Casimirs as

$$\hat{H}_E = \sum_{r=0}^{L-2} (\hat{E}(r))^2. \quad (20)$$

The left and right electric fields are the conjugate variables to the gauge-link variable,

$$[\hat{E}_i^L, \hat{U}] = \hat{T}_i \hat{U}, \quad (21a)$$

$$[\hat{E}_i^R, \hat{U}] = \hat{U} \hat{T}_i, \quad (21b)$$

and further satisfy the commutation relations of the SU(2) Lie algebra at each link,

$$[\hat{E}_i^L, \hat{E}_j^L] = -i\epsilon_{ijk} \hat{E}_k^L, \quad (22a)$$

$$[\hat{E}_i^R, \hat{E}_j^R] = i\epsilon_{ijk} \hat{E}_k^R, \quad (22b)$$

$$[\hat{E}_i^L, \hat{E}_j^R] = 0. \quad (22c)$$

Here, $T_i = \frac{1}{2}\sigma_i$, σ_i is the i^{th} Pauli matrix, and ϵ_{ijk} is the Levi-Civita tensor. The commutation relations for operators at different links vanish. The gauge-matter interaction Hamiltonian

$$\hat{H}_I = x \sum_{r=0}^{L-2} \hat{\psi}^\dagger(r) \hat{U}(r) \hat{\psi}(r+1) + \text{H.c.}, \quad (23)$$

consists of the hopping of a staggered fermion at site r , $\hat{\psi}(r)$, to an adjacent site via interactions with the gauge link $\hat{U}(r)$ originating from site r (and its Hermitian conjugate). \hat{U} is therefore realized as an element of the SU(2) group in the fundamental representation: it is a 2×2 matrix consisting of bosonic-field-operator elements. L denotes the number of lattice points and x is the hopping strength. Open boundary conditions (OBCs) are assumed here and throughout this work. The case of periodic boundary conditions (PBCs) requires minimal modifications to the algorithms presented.

For convenience, the Hamiltonian in Eq. (18) is written in dimensionless form, in which the original dimensionfull Hamiltonian is rescaled by $\frac{2}{a_s g^2}$, with a_s being the spatial lattice spacing and g being the gauge coupling. The dimensionless parameters x and μ are related to the original dimensionfull parameters via $x = \frac{1}{a_s^2 g^2}$ and $\mu = \frac{2m}{a_s g^2}$, where m is the fermion mass. The ‘‘strong-coupling vacuum’’ is associated with the ground state of the theory in the limit $x \rightarrow 0$ ($a_s g \rightarrow \infty$). The continuum limit is achieved by taking the double-ordered limit $\lim_{x \rightarrow \infty} \lim_{L \rightarrow \infty}$ for a fixed $\frac{m}{g}$ [155]. The electric Hamiltonian is diagonal in electric or irrep basis. It is, therefore, a more suitable basis in the strong-coupling limit, and may be less efficient towards the continuum limit, that is achieved in the weak-coupling limit. Alternative bases for SU(N) LGTs such as group-element basis [25] and dual bases [147, 148] are either not fully developed for continuous non-Abelian groups such as SU(2) and/or are not suitable for representing the electric Hamiltonian.

The Hamiltonian in Eq. (18) commutes with the Gauss’s law operators,

$$\hat{G}_i(r) = -\hat{E}_i^L(r) + \hat{E}_i^R(r-1) + \hat{\psi}^\dagger(r) T_i \hat{\psi}(r). \quad (24)$$

The physical sector of the Hilbert space corresponds to the zero eigenvalue of $\hat{G}_i(r)$ at every site r . Since $[\hat{G}_i(r), \hat{G}_j(r)] \neq 0$, specifying the physical sector of the Hilbert space is more complex than in the Abelian theories [132]. Moreover, checking the non-Abelian Gauss’s laws through a checker subroutine in the quantum circuits [38] will be non-trivial and costly.

Given the commutation relations in Eq. (22c), the left and right electric fields can be mapped to the body-frame ($\hat{\mathbf{J}}^b$) and space-frame ($\hat{\mathbf{J}}^s$) angular momenta of a rigid body. Explicitly, $\hat{\mathbf{E}}^L = -\hat{\mathbf{J}}^b (\equiv -\hat{\mathbf{J}}^L)$ and $\hat{\mathbf{E}}^R = \hat{\mathbf{J}}^s (\equiv \hat{\mathbf{J}}^R)$, satisfying $\mathbf{J}^2 \equiv (\hat{\mathbf{J}}^L)^2 = (\hat{\mathbf{J}}^R)^2$ on each link r . In such an angular-momentum basis, the Hilbert space of each link can be characterized by the basis state: ⁷

$$|J, m^L, m^R\rangle, \quad J = 0, \frac{1}{2}, 1, \frac{3}{2}, \dots, \quad -J \leq m^L, m^R \leq J. \quad (25)$$

The angular-momentum operators act on the basis states through the standard relations

$$(\hat{\mathbf{J}}^L)^2 |J, m^L, m^R\rangle = J(J+1) |J, m^L, m^R\rangle, \quad (26a)$$

$$(\hat{\mathbf{J}}^R)^2 |J, m^L, m^R\rangle = J(J+1) |J, m^L, m^R\rangle, \quad (26b)$$

$$\hat{J}_3^L |J, m^L, m^R\rangle = m^L |J, m^L, m^R\rangle, \quad (26c)$$

$$\hat{J}_3^R |J, m^L, m^R\rangle = m^R |J, m^L, m^R\rangle. \quad (26d)$$

⁷Matrix elements for the other angular-momentum components, $\hat{J}_1^{L/R}$ and $\hat{J}_2^{L/R}$, can be derived as in any standard treatment of angular momentum but are not needed below.

The link operator present in the hopping term acts on the basis states at link r as:

$$\hat{U}^{(\alpha,\beta)} |J, m^L, m^R\rangle = \sum_{j=\{0, \frac{1}{2}, 1, \dots\}} \sqrt{\frac{2J+1}{2j+1}} \langle J, m^L; \frac{1}{2}, \alpha | j, m^L + \alpha \rangle \langle J, m^R; \frac{1}{2}, \beta | j, m^R + \beta \rangle |j, m^L + \alpha, m^R + \beta\rangle, \quad (27)$$

where $\alpha, \beta = \pm \frac{1}{2}$ and $\hat{U}_{11} = \hat{U}^{(\frac{1}{2}, -\frac{1}{2})}$, $\hat{U}_{12} = \hat{U}^{(-\frac{1}{2}, -\frac{1}{2})}$, $\hat{U}_{21} = U^{(\frac{1}{2}, \frac{1}{2})}$, $\hat{U}_{22} = \hat{U}^{(-\frac{1}{2}, \frac{1}{2})}$.

The fermions Hilbert space is defined on site,

$$|f_1, f_2\rangle, \quad f_1 = 0, 1, \quad f_2 = 0, 1, \quad (28)$$

consisting of two fermionic quantum numbers f_1 and f_2 corresponding to the occupation number of the two components of the (anti)matter field, ψ_1 and ψ_2 , each taking values 0 and 1. These correspond to the absence and presence of (anti)matter, respectively:

$$\hat{\psi}_1 |f_1, f_2\rangle = (1 - \delta_{f_1,0}) |f_1 - 1, f_2\rangle, \quad (29a)$$

$$\hat{\psi}_1^\dagger |f_1, f_2\rangle = (1 - \delta_{f_1,1}) |f_1 + 1, f_2\rangle, \quad (29b)$$

$$\hat{\psi}_2 |f_1, f_2\rangle = (-1)^{f_1} (1 - \delta_{f_2,0}) |f_1, f_2 - 1\rangle, \quad (29c)$$

$$\hat{\psi}_2^\dagger |f_1, f_2\rangle = (-1)^{f_1} (1 - \delta_{f_2,1}) |f_1, f_2 + 1\rangle. \quad (29d)$$

Here, δ denotes the Kronecker-delta symbol.

The physical states are those that can be represented as a direct product of proper linear combinations of the local basis states $|J, m^R\rangle_{r-1} |f_1, f_2\rangle_r |J, m^L\rangle_r$ such that each linear combination satisfies Gauss's laws at site r , and that the AGL is satisfied so that the left and right total angular momenta on the link are the same. Gauss's laws in this basis amounts to ensuring the net angular momentum at each site is zero: $\mathbf{J}^L(r) + \mathbf{J}^R(r-1) + \mathbf{J}^f(r) = 0$, where $J^f(r) = \frac{1}{2}$ if $f_1(r) + f_2(r) = 1$ and $J^f(r) = 0$ if $f_1(r) + f_2(r) = 0 \pmod{2}$.

The Hilbert space of each link should be truncated to allow for simulations with finite capacity. The truncation can be implemented by imposing $J \leq \Lambda_J$, where $2\Lambda_J$ is an integer. The limit of $\Lambda_J \rightarrow \infty$ must be realized via an extrapolation procedure from finite but sufficiently large values of Λ_J . Different observables will have different sensitivity to Λ_J but previous work reveals that all observables eventually fall into a scaling region in which they asymptote to the $\Lambda_J \rightarrow \infty$ limit exponentially fast [68, 132, 156]. For any finite Λ_J , any raising operators in the Hamiltonian must be redefined in order to ensure they cannot raise a state beyond the Λ_J irrep. In later sections, this will often be done by introducing appropriate projection operators.

3.1.1 Schwinger-boson formulation

In an equivalent representation of the Kogut-Susskind theory, one may consider the single rotor in the body and fixed frames as two uncoupled rotors in one frame, with the requirement that $J(r) \equiv J^L(r) = J^R(r)$. Applying Schwinger's oscillator model of angular momentum, the left (right) rotor can be imagined as a collection of $n_1^L(r) + n_2^L(r)$ ($n_1^R(r) + n_2^R(r)$) spin- $\frac{1}{2}$ particles, with $n_1^L(r)$ ($n_1^R(r)$) spin-up and $n_2^L(r)$ ($n_2^R(r)$) spin-down particles. The benefit of this representation is that transitions between states are expressed using simple combinations of creation and annihilation operators acting on the spin-up and spin-down populations. The annihilation operators can be conveniently organized in SU(2) doublets as $\begin{pmatrix} \hat{a}_1(r) \\ \hat{a}_2(r) \end{pmatrix}$ for the left oscillators and $\begin{pmatrix} \hat{b}_1(r) \\ \hat{b}_2(r) \end{pmatrix}$ for the right oscillators (with the creation operators being Hermitian conjugates of these doublets).

Degrees of freedom.—The site-local fermionic DOFs are carried over unchanged from the Kogut-Susskind formulation, but the link Hilbert space becomes the tensor product space of two simple harmonic oscillators per side of the link (four bosonic modes per link). Thus, associated to the left (right) end of each link, there is a complete set of commuting observables \hat{n}_1^L and \hat{n}_2^L (\hat{n}_1^R and \hat{n}_2^R) that each have a bosonic spectrum (with occupation 0, 1, 2, \dots). The \hat{n}_1^L , \hat{n}_2^L , \hat{n}_1^R , and \hat{n}_2^R occupation-number operators are associated with harmonic-oscillator annihilation operators \hat{a}_1 , \hat{a}_2 , \hat{b}_1 , and \hat{b}_2 , respectively,

$$\hat{n}_1^L(r) = \hat{a}_1^\dagger(r) \hat{a}_1(r), \quad \hat{n}_2^L(r) = \hat{a}_2^\dagger(r) \hat{a}_2(r), \quad (30a)$$

$$\hat{n}_1^R(r) = \hat{b}_1^\dagger(r) \hat{b}_1(r), \quad \hat{n}_2^R(r) = \hat{b}_2^\dagger(r) \hat{b}_2(r), \quad (30b)$$




Figure 2: a) A physical site along the spatial direction is split to two staggered sites in the Kogut-Susskind Hamiltonian. These sites are connected by a gauge link. Corresponding to each staggered site, there is a two-component fermionic field. Furthermore, in the Schwinger-boson formulation, there is a SU(2)-doublet of oscillators associated with the left of the link to the right of the site, and a SU(2)-doublet of oscillators associated with the right of the link to the left of the next site. The operators involved in the non-Abelian and Abelian Gauss's laws are enclosed with corresponding boxes. b) The Schwinger-boson Hilbert space at each site consists of four distinct bosonic and two distinct fermionic Hilbert spaces that can then be mapped to the corresponding qubit registers, upon truncating the bosonic occupations.

which satisfy the usual commutation relations for bosons. The tensored fermionic-bosonic Hilbert space at each site can be constructed from an orthonormal basis defined by

$$|n_1^R, n_2^R\rangle_{r-1} |f_1, f_2\rangle_r |n_1^L, n_2^L\rangle_r = (\hat{\psi}_1^\dagger(r))^{f_1(r)} (\hat{\psi}_2^\dagger(r))^{f_2(r)} \frac{(\hat{a}_1^\dagger(r))^{n_1^L(r)}}{\sqrt{n_1^L(r)!}} \frac{(\hat{a}_2^\dagger(r))^{n_2^L(r)}}{\sqrt{n_2^L(r)!}} \frac{(\hat{b}_1^\dagger(r-1))^{n_1^R(r-1)}}{\sqrt{n_1^R(r-1)!}} \frac{(\hat{b}_2^\dagger(r-1))^{n_2^R(r-1)}}{\sqrt{n_2^R(r-1)!}} |0\rangle. \quad (31)$$

where $|0\rangle$ is the normalized simultaneous vacuum ket of all six modes. The Schwinger-boson DOFs are depicted in Fig. 2(a).

For future reference, the standard matrix elements of the number and ladder operators are collected below

(suppressing the site index):

$$\hat{n}_i^L |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = n_i^L |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle, \quad (32a)$$

$$\hat{n}_i^R |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = n_i^R |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle, \quad (32b)$$

$$\hat{a}_i |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = \sqrt{n_i^L} |n_1^L - \delta_{i,1}, n_2^L - \delta_{i,2}\rangle |n_1^R, n_2^R\rangle, \quad (32c)$$

$$\hat{b}_i |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = \sqrt{n_i^R} |n_1^L, n_2^L\rangle |n_1^R - \delta_{i,1}, n_2^R - \delta_{i,2}\rangle, \quad (32d)$$

$$\hat{a}_i^\dagger |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = \sqrt{n_i^L + 1} |n_1^L + \delta_{i,1}, n_2^L + \delta_{i,2}\rangle |n_1^R, n_2^R\rangle, \quad (32e)$$

$$\hat{b}_i^\dagger |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = \sqrt{n_i^R + 1} |n_1^L, n_2^L\rangle |n_1^R + \delta_{i,1}, n_2^R + \delta_{i,2}\rangle, \quad (32f)$$

for $i = 1, 2$, and

$$\hat{\psi}_1 |f_1, f_2\rangle |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = (1 - \delta_{0,f_1}) |f_1 - 1, f_2\rangle |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle, \quad (33a)$$

$$\hat{\psi}_1^\dagger |f_1, f_2\rangle |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = (1 - \delta_{1,f_1}) |f_1 + 1, f_2\rangle |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle, \quad (33b)$$

$$\hat{\psi}_2 |f_1, f_2\rangle |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = (1 - \delta_{0,f_2}) (-1)^{f_1} |f_1, f_2 - 1\rangle |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle, \quad (33c)$$

$$\hat{\psi}_2^\dagger |f_1, f_2\rangle |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle = (1 - \delta_{1,f_2}) (-1)^{f_1} |f_1, f_2 + 1\rangle |n_1^L, n_2^L\rangle |n_1^R, n_2^R\rangle. \quad (33d)$$

Composite fields.—The Hamiltonian of the Kogut-Susskind formulation is constructed in terms of $\hat{\psi}$, $\hat{E}^{L/R}$, and \hat{U} fields. The fermionic field $\hat{\psi}$ in the Schwinger-boson formulation is the same as in the the Kogut-Susskind formulation, but there is a translation of the bosonic fields. The left and right electric-field operators are given by

$$\hat{E}_i^L(r) = -\hat{a}^\dagger(r) T_i \hat{a}(r), \quad \hat{E}_i^R(r) = \hat{b}^\dagger(r) T_i \hat{b}(r), \quad (34)$$

for $i = 1, 2, 3$, while the gauge-link operator can be written as

$$\hat{U}(r) = \frac{1}{\sqrt{\hat{N}_L(r) + 1}} \begin{pmatrix} -\hat{a}_1(r) \hat{b}_2(r) + \hat{a}_2^\dagger(r) \hat{b}_1^\dagger(r) & \hat{a}_1(r) \hat{b}_1(r) + \hat{a}_2^\dagger(r) \hat{b}_2^\dagger(r) \\ -\hat{a}_2(r) \hat{b}_2(r) - \hat{a}_1^\dagger(r) \hat{b}_1^\dagger(r) & \hat{a}_2(r) \hat{b}_1(r) - \hat{a}_1^\dagger(r) \hat{b}_2^\dagger(r) \end{pmatrix} \frac{1}{\sqrt{\hat{N}_R(r) + 1}}, \quad (35)$$

with

$$\hat{N}^L(r) \equiv \hat{n}_1^L(r) + \hat{n}_2^L(r), \quad (36a)$$

$$\hat{N}^R(r) \equiv \hat{n}_1^R(r) + \hat{n}_2^R(r). \quad (36b)$$

Constraints.—The Kogut-Susskind identity $J^L(r) = J^R(r)$ translates into a constraint along links in the Schwinger-boson formulation:

$$\hat{N}^R(r) - \hat{N}^L(r) = 0. \quad (37)$$

Any basis state not satisfying this constraint is considered a part of the unphysical Hilbert space and has no counterpart in the Kogut-Susskind formulation. The constraint is known as the *Abelian Gauss's law* because the associated generators $\hat{N}^R(r) - \hat{N}^L(r)$ all commute. In addition to the Abelian generators, there are also the ordinary non-Abelian Gauss's law generators in Eq. (24), which carry over to the Schwinger-boson formulation with $\hat{E}^{L/R}$ translated as above. As before, this restricts the physical space to appropriate linear combinations of the basis states (including the fermionic DOFs) at each site such that the net angular momentum is zero. Furthermore, because of the AGL, the number of independent DOFs associated with the two \hat{a} -type and two \hat{b} -type oscillators at each link reduces from four to three. These can be identified as $n_1^L(r) + n_2^L(r) (= n_1^R(r) + n_2^R(r))$, $n_1^L(r) - n_2^L(r) = 2J_3^L(r)$, and $n_1^R(r) - n_2^R(r) = 2J_3^R(r)$.

Hamiltonian.—The Hamiltonian can be expressed in terms of the Schwinger-boson operators as well as fermionic operators. The fermion self-energy, being in all ways identical to the Kogut-Susskind formulation, is

$$\hat{H}_M^{\text{SB}} = \mu \sum_{r=0}^{L-1} (-1)^r (\hat{\psi}_1^\dagger(r) \hat{\psi}_1(r) + \hat{\psi}_2^\dagger(r) \hat{\psi}_2(r)). \quad (38)$$

The electric Hamiltonian can be derived by inserting the Schwinger-boson expressions into $\sum_{i=1}^3 \hat{E}_i^L(r) \hat{E}_i^L$ and using the canonical commutation relations, with the end result

$$\hat{H}_E^{\text{SB}} = \sum_{r=0}^{L-2} \frac{\hat{N}^L(r)}{2} \left(\frac{\hat{N}^L(r)}{2} + 1 \right) = \sum_{r=0}^{L-2} \frac{\hat{n}_1^L(r) + \hat{n}_2^L(r)}{2} \left(\frac{\hat{n}_1^L(r) + \hat{n}_2^L(r)}{2} + 1 \right), \quad (39)$$

where the left Casimirs for the electric energy are used by choice. The interacting (hopping) Hamiltonian can be written as before with the identification of the link operator as in Eq. (35), leading to eight distinct terms each with two types of fermionic operators and two types of Schwinger-boson operators:

$$\begin{aligned} \hat{H}_I^{\text{SB}} = x \sum_{r=0}^{L-2} \frac{1}{\sqrt{\hat{N}^L(r)+1}} & [\hat{\psi}_1^\dagger(r) \hat{\psi}_2(r+1) \hat{a}_1(r) \hat{b}_1(r) - \hat{\psi}_1^\dagger(r) \hat{\psi}_1(r+1) \hat{a}_1(r) \hat{b}_2(r) \\ & - \hat{\psi}_2^\dagger(r) \hat{\psi}_1(r+1) \hat{a}_2(r) \hat{b}_2(r) + \hat{\psi}_2^\dagger(r) \hat{\psi}_2(r+1) \hat{a}_2(r) \hat{b}_1(r) \\ & + \hat{\psi}_2(r) \hat{\psi}_1^\dagger(r+1) \hat{a}_1(r) \hat{b}_1(r) - \hat{\psi}_1(r) \hat{\psi}_1^\dagger(r+1) \hat{a}_2(r) \hat{b}_1(r) \\ & - \hat{\psi}_1(r) \hat{\psi}_2^\dagger(r+1) \hat{a}_2(r) \hat{b}_2(r) + \hat{\psi}_2(r) \hat{\psi}_2^\dagger(r+1) \hat{a}_1(r) \hat{b}_2(r)] \frac{1}{\sqrt{\hat{N}^L(r)+1}} + \text{H.c.} \quad (40) \end{aligned}$$

Here, the AGL is used to replace $\frac{1}{\sqrt{\hat{N}^R(r)+1}}$ on the RHS of the link operator \hat{U} in Eq. (35) with $\frac{1}{\sqrt{\hat{N}^L(r)+1}}$. The reason is that each term in the square bracket in Eq. (40) preserves the AGL.

Truncation.—The infinite-dimensional Hilbert space of the Schwinger bosons needs to be truncated to allow simulations with finite resources. One can introduce an upper cutoff integer Λ for each bosonic-oscillator quantum number, i.e., $0 \leq n_i^L, n_i^R \leq \Lambda$ for $i = 1, 2$. The exact SU(2) LGT is recovered as $\Lambda \rightarrow \infty$. Various relations need to be modified for any finite Λ by inclusion of an appropriate projection to ensure \hat{a}_1^\dagger (\hat{a}_1) acting on the kets (bras) with $n_1^L = \Lambda$ vanishes, and similarly for the other three modes.

The cutoff in the angular-momentum basis, Λ_J , is related to that in the Schwinger-boson basis, Λ , via the relation $2J = n_1^L + n_2^L (= n_1^R + n_2^R)$. However, there is a subtlety involved, in that our truncation scheme will leave incomplete irreps near the cutoff. The problem can be illustrated by considering $\Lambda = 1$ and focusing on one end (e.g., left) of a link. In this case, there is a Schwinger-boson configuration $(n_1^L, n_2^L) = (1, 1)$ corresponding to $(J^L, m^L) = (1, 0)$, but no configurations corresponding to $(J^L, m^L) = (1, \pm 1)$. Consequently, one can define “the” cutoff angular momentum to be the highest J for which the complete irrep is accounted for. This turns out to be $J = \frac{\Lambda}{2}$.

3.1.2 Loop-string-hadron formulation

The Schwinger-boson formulation can be used to derive a closely-related, but distinct, formulation known as the LSH formulation. Considering that both the Schwinger bosons and the matter fields associated with a given site transform in the fundamental representation of the local SU(2) group, various bilinear operators can be formed each transforming as a singlet under SU(2). These provide gauge-invariant operators that generate the physical Hilbert space out of the vacuum, that is the state with no matter and gauge excitations. These operators include segments of electric flux loops, quarks starting or ending at bosonic strings, and hadrons.

Degrees of freedom.—To each site of the 1D lattice, the LSH formulation associates a complete set of commuting observables $\{\hat{n}_\ell, \hat{n}_i, \hat{n}_o\}$, where \hat{n}_ℓ is a Hermitian operator with bosonic eigenvalues $(0, 1, 2, \dots)$, and \hat{n}_i and \hat{n}_o are Hermitian operators with fermionic eigenvalues (0 and 1 only). Normalized ladder operators $\hat{\Gamma}, \hat{\chi}_i, \hat{\chi}_o$, and their adjoints are introduced such that

$$[\hat{n}_\ell, \hat{\Gamma}] = -\hat{\Gamma}, \quad [\hat{n}_\ell, \hat{\Gamma}^\dagger] = \hat{\Gamma}^\dagger, \quad (41a)$$

$$[\hat{n}_{q'}, \hat{\chi}_q] = -\hat{\chi}_q \delta_{q',q}, \quad [\hat{n}_{q'}, \hat{\chi}_q^\dagger] = \hat{\chi}_q^\dagger \delta_{q',q}, \quad (41b)$$

for $q, q' \in \{i, o\}$, in addition to the usual anticommuting statistics of quark modes:

$$\{\hat{\chi}_{q'}, \hat{\chi}_q^\dagger\} = \delta_{q',q}, \quad (42a)$$

$$\{\hat{\chi}_{q'}, \hat{\chi}_q\} = \{\hat{\chi}_{q'}^\dagger, \hat{\chi}_q^\dagger\} = 0. \quad (42b)$$

The Hilbert space can be constructed from on-site, orthonormal basis states defined by

$$|n_\ell, n_i, n_o\rangle = (\hat{\Gamma}^\dagger)^{n_\ell} (\hat{\chi}_i^\dagger)^{n_i} (\hat{\chi}_o^\dagger)^{n_o} |0\rangle, \quad (43)$$




Figure 3: (a) Corresponding to each staggered site, there are a number of SU(2)-invariant fermionic and bosonic operators that act on the on-site Hilbert space of the LSH theory. The operators involved in the AGL are enclosed with a corresponding box. The non-Abelian Gauss's laws have already been satisfied by construction, as all operators enclosed in the corresponding box are individually gauge invariant. (b) The LSH Hilbert space at each site consists of one distinct bosonic and two distinct fermionic Hilbert spaces that can then be mapped to the corresponding qubit registers, upon truncating the bosonic occupation.

where $|0\rangle$ is the normalized vacuum ket of all three modes. The LSH DOFs are depicted in Fig. 3(a).

For future reference, the matrix elements of the number and normalized ladder operators are collected below:

$$\hat{n}_\ell |n_\ell, n_i, n_o\rangle = n_\ell |n_\ell, n_i, n_o\rangle, \quad (44a)$$

$$\hat{n}_i |n_\ell, n_i, n_o\rangle = n_i |n_\ell, n_i, n_o\rangle, \quad (44b)$$

$$\hat{n}_o |n_\ell, n_i, n_o\rangle = n_o |n_\ell, n_i, n_o\rangle, \quad (44c)$$

$$\hat{\Gamma}^\dagger |n_\ell, n_i, n_o\rangle = |n_\ell + 1, n_i, n_o\rangle, \quad (44d)$$

$$\hat{\Gamma} |n_\ell, n_i, n_o\rangle = (1 - \delta_{n_\ell, 0}) |n_\ell - 1, n_i, n_o\rangle, \quad (44e)$$

$$\hat{\chi}_i^\dagger |n_\ell, n_i, n_o\rangle = (1 - \delta_{n_i, 1}) |n_\ell, n_i + 1, n_o\rangle, \quad (44f)$$

$$\hat{\chi}_i |n_\ell, n_i, n_o\rangle = (1 - \delta_{n_i, 0}) |n_\ell, n_i - 1, n_o\rangle, \quad (44g)$$

$$\hat{\chi}_o^\dagger |n_\ell, n_i, n_o\rangle = (-1)^{n_i} (1 - \delta_{n_o, 1}) |n_\ell, n_i, n_o + 1\rangle, \quad (44h)$$

$$\hat{\chi}_o |n_\ell, n_i, n_o\rangle = (-1)^{n_i} (1 - \delta_{n_o, 0}) |n_\ell, n_i, n_o - 1\rangle. \quad (44i)$$

Composite fields.—Unlike the Kogut-Susskind and Schwinger-boson formulations, the LSH formulation does not have realizations of the $\hat{E}_i^{L/R}$ and \hat{U} fields. Instead, dynamics is generated by site-localized “loop,” “string,”

and “hadron” operators based on which the formulation is named. The string operators are given in factorized form by

$$\hat{S}_{o(i)}^{++} = \hat{\chi}_{o(i)}^\dagger (\hat{\Gamma}^\dagger)^{\hat{n}_{i(o)}} \sqrt{\hat{n}_\ell + 2 - \hat{n}_{i(o)}}, \quad (45a)$$

$$\hat{S}_{o(i)}^{--} = \hat{\chi}_{o(i)} \hat{\Gamma}^{\hat{n}_{i(o)}} \sqrt{\hat{n}_\ell + 2(1 - \hat{n}_{i(o)})}, \quad (45b)$$

$$\hat{S}_{o(i)}^{+(-+)} = \hat{\chi}_{o(i)}^\dagger \hat{\Gamma}^{1-\hat{n}_{o(i)}} \sqrt{\hat{n}_\ell + 2\hat{n}_{o(i)}}, \quad (45c)$$

$$\hat{S}_{o(i)}^{-+(+-)} = \hat{\chi}_{i(o)} (\hat{\Gamma}^\dagger)^{1-\hat{n}_{o(i)}} \sqrt{\hat{n}_\ell + 1 + \hat{n}_{o(i)}}. \quad (45d)$$

We omit the respective relations for the “loop” and “hadron” operators because the Hamiltonian, to be discussed shortly, does not make explicit use of them.

Constraints.—For this site-local representation to recover the original Kogut-Susskind theory, an AGL among the quantum numbers of two adjacent sites must be satisfied:

$$\hat{N}^L(r) - \hat{N}^R(r) = 0, \quad (46)$$

where

$$\hat{N}^L(r) \equiv \hat{n}_\ell(r) + \hat{n}_o(r)(1 - \hat{n}_i(r)), \quad (47a)$$

$$\hat{N}^R(r) \equiv \hat{n}_\ell(r+1) + \hat{n}_i(r+1)(1 - \hat{n}_o(r+1)). \quad (47b)$$

Incidentally, \hat{N}^L and \hat{N}^R correspond exactly to the total left and right occupation numbers defined in the Schwinger-boson formulation. The non-Abelian Gauss’s laws are built in to the LSH Hilbert space such that no constraints beyond the AGL need to be imposed on the Hilbert space *a posteriori*. This makes the expression of physical states simple in the basis of Eq. (43), leading to locally 1-sparse interactions, contrary to the Hamiltonians expressed in the angular-momentum and Schwinger-boson bases.

Hamiltonian.—The Hamiltonian of the SU(2) theory in 1+1 D with staggered quarks can then be expressed in terms of the LSH number and ladder operators. The mass and electric Hamiltonians are simply related to gauge-invariant operators $\hat{\psi}^\dagger(r) \cdot \hat{\psi}(r) = \hat{n}_i(r) + \hat{n}_o(r)$ and $\hat{N}^L(r) = \hat{N}^R(r)$ [131]:

$$\hat{H}_M^{\text{LSH}} = \mu \sum_{r=0}^{L-2} (-1)^r (\hat{n}_i(r) + \hat{n}_o(r)), \quad (48)$$

and

$$\hat{H}_E^{\text{LSH}} = \sum_{r=0}^{L-2} \frac{\hat{N}^L(r)}{2} \left(\frac{\hat{N}^L(r)}{2} + 1 \right) = \sum_{r=0}^{L-2} \frac{\hat{n}_\ell(r) + \hat{n}_o(r)(1 - \hat{n}_i(r))}{2} \left[\frac{\hat{n}_\ell(r) + \hat{n}_o(r)(1 - \hat{n}_i(r))}{2} + 1 \right], \quad (49)$$

where the left Casimirs for the electric energy are used by choice. The gauge-matter interaction Hamiltonian in the LSH formulation, which is constructed using the string operators, can be taken as

$$\begin{aligned} \hat{H}_I^{\text{LSH}} = & x \sum_{r=0}^{L-2} \left\{ \left[\hat{\chi}_i^\dagger(r+1) \hat{\Gamma}^\dagger(r+1)^{\hat{n}_o(r+1)} \right] \left[\hat{\chi}_i(r) \hat{\Gamma}^\dagger(r)^{1-\hat{n}_o(r)} \right] \sqrt{\frac{\hat{n}_\ell(r) + 1 + \hat{n}_o(r)}{\hat{n}_\ell(r) + 1 + \hat{n}_o(r+1)}} + \right. \\ & \left. \left[\hat{\chi}_o^\dagger(r) \hat{\Gamma}^\dagger(r)^{\hat{n}_i(r)} \right] \left[\hat{\chi}_o(r+1) \hat{\Gamma}^\dagger(r+1)^{1-\hat{n}_i(r+1)} \right] \sqrt{\frac{\hat{n}_\ell(r+1) + 1 + \hat{n}_i(r+1)}{\hat{n}_\ell(r+1) + 1 + \hat{n}_i(r)}} + \text{H.c.} \right\}. \quad (50) \end{aligned}$$

The above expressions are not directly expressed using the string operators in Eqs. (45) as prescribed in Ref. [131]. Instead, they have been derived from the original prescription by further invoking the AGL and simplifying the square-root functions, and the details of this manipulation are provided in Appendix B. The above forms are preferred in this work because they yield cost-saving benefits from the viewpoint of quantum algorithms.

Truncation.—As with the Schwinger-boson formulation, it is necessary to truncate the infinite-dimensional Hilbert spaces of bosons for simulation with finite resources. In the LSH formulation, this truncation can be realized by introducing an upper cutoff Λ_ℓ on the flux quantum numbers n_ℓ , i.e., $0 \leq n_\ell \leq \Lambda_\ell$. The limit of $\Lambda_\ell \rightarrow \infty$ must be taken to recover the Kogut-Susskind theory. With a finite cutoff on the n_ℓ quantum numbers

in effect, one must ensure $\hat{\Gamma}^\dagger$ ($\hat{\Gamma}$) acting on a ket (bra) with $n_\ell = \Lambda_\ell$ vanishes with the use of corresponding projectors.

The relation between the cutoff Λ_ℓ and those introduced in the angular-momentum and Schwinger-boson formulations is again realized through the equations for \hat{N}^L (\hat{N}^R). The subtlety in this formulation is a matter of some missing states near the cutoff, although it has nothing to do with incomplete irreps. Rather, there are certain angular momenta for which some but not all states are present. A simple example is illustrated by taking $\Lambda_\ell = 0$. This truncation admits states with $J = \frac{1}{2}$ arising from the string-end configurations such as $(n_\ell, n_i, n_o) = (0, 0, 1)$, since $N^L = 1 = 2J$ for this configuration. But not all $J = \frac{1}{2}$ states are accounted for, such as $(n_\ell, n_i, n_o) = (1, 0, 0)$, which produces the same value of N^L . As with the Schwinger-boson formulation, “the” cutoff angular momentum can be defined to be the highest J for which all states up to J are accounted for. This again turns out to be $\frac{\Lambda_\ell}{2}$. Our strategy in the following is to fix the cutoff in the angular-momentum basis, which given the definitions above, leads to the same cutoff on the Schwinger-boson and the loop excitations. For this reason, $\Lambda_\ell = \Lambda$ and no distinction will be needed in the following between the cutoffs in the two formulations.

To conclude this section, we emphasize that among all the formulations, only the LSH formulation constructs a Hilbert space with non-Abelian gauge invariance built in. Both the site-local LSH and Schwinger-boson bases require an AGL to be imposed on the quantum numbers of two adjacent sites, but this constraint is more easier to satisfy in the quantum-simulation algorithms than the (non-commuting) non-Abelian Gauss’s laws, as will be demonstrated shortly.

3.2 Circuit decomposition of the propagators

The general methods of Sec. 2 can be applied straightforwardly to decompose the SU(2) LGT within the formulations introduced in Sec. 3.1, once the mapping of the Hilbert space to qubits is specified and the simulable unitaries are identified. Since a quantum algorithm for simulating SU(2) dynamics within the angular-momentum formulation of the Kogut-Susskind theory has already been developed [69] (albeit using a different strategy than that of this work), in the following we exclusively discuss algorithms for simulating the time-evolution operator in the Schwinger-boson and the LSH bases. The application of our general method within the angular-momentum basis will be briefly discussed in Section 4 to provide a qualitative cost comparison. For notational brevity, the overhead hat symbols for operators will be dropped from this point on.

3.2.1 Schwinger-boson propagators

To decompose time evolution in the gate model of quantum computation, the field operators introduced in the Hamiltonian must be replaced by qubit operators. The fermionic modes on each lattice site require two qubits, one for each SU(2) component of the fermions, with no Hilbert space truncation needed. In 1+1 D, Jordan-Wigner transformation is the simplest way to replace the fermionic operators with operators that act on the two-dimensional qubit Hilbert space $\{|0\rangle, |1\rangle\}$. The transformation maps the L fermionic doublets onto $2L$ spin variables. Explicitly, we choose a Jordan-Wigner transformation that maps the fermionic modes $\psi(r)$ and $\psi^\dagger(r)$ according to⁸

$$\psi_1(r) \rightarrow \left(\prod_{k=0}^{2r-1} Z_k \right) \sigma_{2r}^+, \quad \psi_2(r) \rightarrow \left(\prod_{k=0}^{2r} Z_k \right) \sigma_{2r+1}^+, \quad (51a)$$

$$\psi_1^\dagger(r) \rightarrow \left(\prod_{k=0}^{2r-1} Z_k \right) \sigma_{2r}^-, \quad \psi_2^\dagger(r) \rightarrow \left(\prod_{k=0}^{2r} Z_k \right) \sigma_{2r+1}^-. \quad (51b)$$

The Hilbert space of the Schwinger-boson operators is built from basis states defined in Eq. (31), with the corresponding operators defined in Eqs. (32) and (33). With a qubit register \mathbf{p} of finite size η , only the first 2^η modes can be encoded, putting a truncation $\Lambda \equiv 2^\eta - 1$ on the occupation number of each bosonic mode: $\hat{N}_{\mathbf{p}} |p\rangle = p |p\rangle$ with $0 \leq p \leq \Lambda$. Moreover, the corresponding raising operator must be modified by a projection operator such that it annihilates the state with occupation Λ , that is $a_{\mathbf{p}}^\dagger |p\rangle = (1 - \delta_{p,\Lambda}) \sqrt{p+1} |p+1\rangle$. The state $|p\rangle$ is mapped to qubits via a binary encoding, that is

$$|p\rangle = \bigotimes_{j=0}^{\eta-1} |p_j\rangle \quad \text{with} \quad p = \sum_{j=0}^{\eta-1} p_j 2^j, \quad (52)$$

⁸Note that our convention set in Sec. 2 leads (counter-intuitively) to the relations $\sigma^+ = |0\rangle\langle 1|$ and $\sigma^- = |1\rangle\langle 0|$ for the spin- $\frac{1}{2}$ raising and lowering operators.

where $p_j = 0, 1$ are the coefficients of the binary representation of the integer p . Note that the $j = 0$ index is designated as the least significant bit in the binary representation of p . For the Schwinger bosons on a 1D chain with OBCs, there are $4(L - 1)$ such bosonic operators, associated with a bosonic register for each $SU(2)$ component of the left and right oscillators on the link, see Sec. 3.1.1. These registers are all assumed to have size η . Figure 2(b) depicts the Hilbert spaces associated with the DOFs of the Schwinger-boson formulation.

Next, the Hamiltonian needs to be represented in terms of the operators acting on these qubit registers. Since the Hamiltonian is the sum of (site- or link-)local terms, in this section we focus on deriving efficient circuit decompositions for exponentiating each of such local terms, and will consider product formulas for approximating the full time-evolution operator of the system within a fixed error tolerance later in Sec. 3.3.

—Near term—

Implementing mass propagators

The mass Hamiltonian is $\sum_{r=0}^{L-1} H_M^{\text{SB}}(r)$, with site-local mass terms decomposed into two commuting “sub-terms” as

$$H_M^{\text{SB}}(r) = \sum_{j=1}^2 H_M^{\text{SB}(j)}(r), \quad (53a)$$

$$H_M^{\text{SB}(1)}(r) = \frac{\mu}{2} (-1)^{r+1} Z_{2r}, \quad (53b)$$

$$H_M^{\text{SB}(2)}(r) = \frac{\mu}{2} (-1)^{r+1} Z_{2r+1}, \quad (53c)$$

after the Jordan-Wigner transformation in Eq. (51), and upon neglecting constant terms that only introduce time-dependent but otherwise constant phases in the dynamics. The circuit decomposition of this propagator amounts to single-qubit Z rotations on each fermionic register. With two qubits indexed by $2r$ and $2r + 1$, corresponding to ‘ $\psi_1(r)$ ’ and ‘ $\psi_2(r)$ ’ registers, respectively, the circuit implements

$$R_{2r}^Z ((-1)^{r+1} \mu t) R_{2r+1}^Z ((-1)^{r+1} \mu t), \quad (54)$$

with $R_j^Z(\theta)$ defined as $R_j^Z \equiv e^{-i\theta Z_j/2}$.

Lemma 3.1. *Using a 2-qubit register with no ancilla qubits, $e^{-itH_M^{\text{SB}}(r)}$ can be implemented without approximation, up to a phase, with no CNOT gates required.*

Proof. The circuit implementing the mass-term propagator in the near-term scenario is simply a straightforward implementation of Eq. (54) with exact rotation angles applied, and hence requires no additional ancilla qubits. As a result, the mass-term propagator is essentially ‘free’ in a near-term implementation (costing no entangling gates). This trivial cost is recorded in Table 1. \square

Implementing electric propagators

The electric Hamiltonian is $\sum_{r=0}^{L-2} H_E^{\text{SB}}(r)$, with link-local electric terms (Casimir operators) decomposed into two commuting subterms as

$$H_E^{\text{SB}}(r) = \sum_{j=1}^2 H_E^{\text{SB}(j)}(r), \quad (55a)$$

$$H_E^{\text{SB}(1)}(r) = \frac{1}{2} N^L(r), \quad (55b)$$

$$H_E^{\text{SB}(2)}(r) = \frac{1}{4} (N^L(r))^2. \quad (55c)$$

Note that the Casimir operators $H_E^{\text{SB}}(r)$ are already diagonalized in the chosen basis, so other divisions of $H_E(r)$ into subterms are possible with terms that all commute. In other words, our splitting into $H_E^{\text{SB}(1)}(r)$ and $H_E^{\text{SB}(2)}(r)$ is not unique, but a choice is made to provide the concrete gate-count analysis that follows. Nonetheless, a naive Pauli decomposition of $H_E(r)$ only involves $O(\eta^2)$ terms, so any choice of subterms will not have a CNOT-gate scaling worse than this. Regardless of the division into subterms, there are two conclusions one can draw: 1) $e^{-itH_E^{\text{SB}}(r)}$ is readily circuitized without any theoretical approximation, and using established

Schwinger-boson mass propagator subroutine	CNOT count
$e^{-itH_M^{\text{SB}(1)}(r)}$ and $e^{-itH_M^{\text{SB}(2)}(r)}$ gate	0
Full $e^{-itH_M^{\text{SB}}(r)}$	0
Schwinger-boson electric propagator subroutine	CNOT count
$e^{-itN_1^L/2}$ or $e^{-itN_2^L/2}$	0
$e^{-it(N_1^L)^2/4}$ or $e^{-it(N_2^L)^2/4}$	$(\eta + 2)(\eta - 1)/2$
$e^{-itN_1^L N_2^L/2}$	$2\eta^2$
Full $e^{-itH_E^{\text{SB}}(r)}$	$3\eta^2 + \eta - 2$

Table 1: Summary of the cost associated with the near-term implementation of diagonal operators in the Schwinger-boson Hamiltonian.

algorithms (see below). 2) Over-optimizing the circuitization of $e^{-itH_E^{\text{SB}}(r)}$ is unnecessary because, as will be seen later, its cost will ultimately be dwarfed by that of $e^{-itH_I^{\text{SB}}(r)}$.

The division into subterms as in Eqs. (55) involves two distinct types of subterms, one linear and one quadratic in $N^L = N_1^L + N_2^L$, where the number operator for either individual oscillator is related to the binary register's qubits via

$$N_{\mathbf{p}} |p\rangle = \left[\frac{1}{2}(2^\eta - 1)\mathcal{I}_{\mathbf{p}} - \frac{1}{2} \sum_{j=0}^{\eta-1} 2^j Z_{\mathbf{p}_j} \right] |p\rangle. \quad (56)$$

Here, $\mathcal{I}_{\mathbf{p}} = \prod_{j=0}^{\eta-1} \mathcal{I}_{\mathbf{p}_j}$ is the identity operator on the state of \mathbf{p} register, i.e., $\mathcal{I}_{\mathbf{p}} |p\rangle = |p\rangle$.

Lemma 3.2. *Using a 2η -qubit register with no ancilla qubits, $e^{-itH_E^{\text{SB}}(r)}$ can be implemented without approximation, up to a phase, using $3\eta^2 + \eta - 2$ CNOT gates (and a number of single-qubit rotations).*

Proof. This cost is obtained as follows. i) The operator N^L is a linear combination of the identity and every single-qubit Pauli- Z operators across the ‘ a_1 ’ and ‘ a_2 ’ registers. Therefore, simulating e^{-iN^L} amounts to a global phase and 2η single-qubit Z rotations, but no CNOT gates. ii) The operator $(N^L)^2$ may be further decomposed as

$$(N^L)^2 = (n_1^L)^2 + (n_2^L)^2 + 2n_1^L n_2^L.$$

The $e^{-it(n_i^L)^2/4}$ term can be simulated like an ‘ e^{-itE^2} ’ term in the U(1) theory, with the difference between the two operators being terms linear in single-qubit Z rotations (which cost no CNOT gates). According to Lemma 2 of Ref. [54], this can be done with $(\eta + 2)(\eta - 1)/2$ CNOT gates. The remaining $e^{-itn_1^L n_2^L/2}$ contribution can be simulated by naively circuitizing its Pauli decomposition consisting of η^2 distinct two-qubit Z rotations ($Z \otimes Z$ operators), each requiring two CNOT gates. The gate counts of i) and ii) are summarized in Table 1, adding up to the stated total CNOT count of $2 \times (\eta + 2)(\eta - 1)/2 + 2\eta^2 = 3\eta^2 + \eta - 2$. \square

Implementing hopping propagators

The gauge-matter interaction Hamiltonian is $\sum_{r=0}^{L-2} H_I^{\text{SB}}(r)$, with link-local hopping terms $H_I^{\text{SB}}(r)$ that are off-diagonal in the electric basis, see Eqs. (40). The hopping terms can be decomposed into a set of subterms but as with the mass and electric terms, the division of $H_I^{\text{SB}}(r)$ into subterms is not unique. However, unlike the mass and electric terms, the hopping terms are off-diagonal and do not readily split into commuting subterms. For this reason, we resort to an approximation $e^{-itH_I^{\text{SB}}(r)} \approx \prod_j e^{-itH_I^{\text{SB}(j)}(r)}$, in which each subterm is simulated without further splitting. Clearly, one important consequence of the non-commuting subterms is that the Trotter error bound will grow, as will be seen in Sec. 3.3. Furthermore, different choices of $H_I^{\text{SB}(j)}(r)$ could have different simulation costs, which could significantly change the total cost of the simulation as implementing $e^{-itH_I^{\text{SB}(j)}(r)}$ is seen to dominate the gate count. Even for a similar cost of each individual $H_I^{\text{SB}(j)}(r)$, the total number of hopping subterms has immediate implications for the Trotter error and gate count, so having less number of simulable terms is desirable. The number of hopping subterms will be denoted by ν in the following, with a ‘SB’ superscript when the Schwinger-boson formulation is concerned.

A choice of splitting is to the eight individual terms shown in Eq. (40), along with their Hermitian conjugates. To translate the hopping terms into the language of a quantum computer, one applies the Jordan-Wigner transformations in Eqs. (51), along with the truncation to η -bit registers for each bosonic oscillator mode. Then, the interaction Hamiltonian acting in the qubit space is:

$$H_I^{\text{SB}}(r) = \sum_{j=1}^8 H_I^{\text{SB}(j)}(r), \quad (57a)$$

$$H_I^{\text{SB}(1)}(r) = x \sigma_{2r}^- Z_{2r+1} Z_{2r+2} \sigma_{2r+3}^+ \lambda_{L1}^- \lambda_{R1}^- \mathcal{D}^{\text{SB}}(n_1^L, n_1^R, n_2^L) + \text{H.c.}, \quad (57b)$$

$$H_I^{\text{SB}(2)}(r) = -x \sigma_{2r}^- Z_{2r+1} \sigma_{2r+2}^+ \lambda_{L1}^- \lambda_{R2}^- \mathcal{D}^{\text{SB}}(n_1^L, n_2^R, n_2^L) + \text{H.c.}, \quad (57c)$$

$$H_I^{\text{SB}(3)}(r) = -x \sigma_{2r+1}^- \sigma_{2r+2}^+ \lambda_{L2}^- \lambda_{R2}^- \mathcal{D}^{\text{SB}}(n_2^L, n_2^R, n_1^L) + \text{H.c.}, \quad (57d)$$

$$H_I^{\text{SB}(4)}(r) = x \sigma_{2r+1}^- Z_{2r+2} \sigma_{2r+3}^+ \lambda_{L2}^- \lambda_{R1}^- \mathcal{D}^{\text{SB}}(n_2^L, n_1^R, n_1^L) + \text{H.c.}, \quad (57e)$$

$$H_I^{\text{SB}(5)}(r) = -x \sigma_{2r+1}^+ \sigma_{2r+2}^- \lambda_{L1}^- \lambda_{R1}^- \mathcal{D}^{\text{SB}}(n_1^L, n_1^R, n_2^L) + \text{H.c.}, \quad (57f)$$

$$H_I^{\text{SB}(6)}(r) = x \sigma_{2r}^+ Z_{2r+1} \sigma_{2r+2}^- \lambda_{L2}^- \lambda_{R1}^- \mathcal{D}^{\text{SB}}(n_2^L, n_1^R, n_1^L) + \text{H.c.}, \quad (57g)$$

$$H_I^{\text{SB}(7)}(r) = x \sigma_{2r}^+ Z_{2r+1} Z_{2r+2} \sigma_{2r+3}^- \lambda_{L2}^- \lambda_{R2}^- \mathcal{D}^{\text{SB}}(n_2^L, n_2^R, n_1^L) + \text{H.c.}, \quad (57h)$$

$$H_I^{\text{SB}(8)}(r) = -x \sigma_{2r+1}^+ Z_{2r+2} \sigma_{2r+3}^- \lambda_{L1}^- \lambda_{R2}^- \mathcal{D}^{\text{SB}}(n_1^L, n_2^R, n_2^L) + \text{H.c.} \quad (57i)$$

The diagonal function \mathcal{D}^{SB} is defined as⁹

$$\mathcal{D}^{\text{SB}}(p, q, p') \equiv \sqrt{\frac{pq}{(p+p')(p+p'+1)}}. \quad (58)$$

Here for brevity, we have adopted the shorthand notation $p \equiv N_p$, where N_p is the occupation-number operator on the bosonic register \mathbf{p} ¹⁰ (and similarly for q and p').

The action of the cyclic ladder operators λ^\pm on the bosonic modes should be realized as $\lambda_{\mathbf{p}}^+ |p\rangle = (1 - \delta_{p,\Lambda}) |p+1\rangle + \delta_{p,\Lambda} |0\rangle$ and $\lambda_{\mathbf{p}}^- |p\rangle = (1 - \delta_{p,0}) |p-1\rangle + \delta_{p,0} |\Lambda\rangle$. The reason such modified ladder operators still reproduce the original Schwinger-boson interaction Hamiltonian is the presence of $D(p, q, p')$ factors. Consider two situations that are affected by the new definition of the ladder operators: i) Terms with $\lambda_{\mathbf{p}}^- \lambda_{\mathbf{q}}^- D(p, q, p')$ operator structure: consider, for example, $\lambda_{\mathbf{p}}^- D(p, q, p') |0\rangle_{\mathbf{p}}$. While $\lambda_{\mathbf{p}}^- |0\rangle_{\mathbf{p}} = |\Lambda\rangle_{\mathbf{p}}$ for the cyclic operator, since $D(0, q, p') = 0$, the amplitude for this process remains zero, which reproduces what is expected from the non-cyclic ladder operator. Similarly, the same argument holds for $p \rightarrow q$. ii) Terms with $D(p, q, p') \lambda_{\mathbf{p}}^+ \lambda_{\mathbf{q}}^+$ operator structure: consider, for example, $D(p, q, p') \lambda_{\mathbf{p}}^+ |\Lambda\rangle_{\mathbf{p}}$. While $\lambda_{\mathbf{p}}^+ |\Lambda\rangle_{\mathbf{p}} = |0\rangle_{\mathbf{p}}$ for the cyclic operator, since $D(0, q, p') = 0$, the amplitude for this process remains zero, which reproduces what is expected from the non-cyclic but truncated ladder operator. Similarly, the same argument holds for $p \rightarrow q$.¹¹

The terms in Eqs. (57) have been put into the form of (shifting operators) \times (diagonal operators) to match the example Hamiltonians studied in Sec. 2.1, so that the SVD diagonalization methods are readily applied. In the language of Sec. 2.1, each term $H_I^{\text{SB}(j)}$ involves two spin raising and lowering operators (that act on fermionic modes at two adjacent sites), two bosonic raising or lowering operators (acting at opposite ends of a common link), and a non-trivial diagonal function (of three bosonic quantum numbers). The general structure of the Schwinger-boson hopping subterms is encapsulated by

$$H_I^{\text{SB}(j)} = \pm x \sigma_{\mathbf{x}'}^+ (Z_{\mathbf{x}'}) (Z_{\mathbf{y}'}) \sigma_{\mathbf{y}'}^- \mathcal{D}^{\text{SB}}(p, q, p') \lambda_{\mathbf{p}}^+ \lambda_{\mathbf{q}}^+ + \text{H.c.}, \quad (59)$$

with the labels $\{\mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}'\}$ introduced for fermionic modes, and $\{\mathbf{p}, \mathbf{q}, \mathbf{p}'\}$ for bosonic modes. The precise mapping to modes for each subterm is displayed in Table 2. Here and from now on, we denote the two Pauli-Z operators on registers \mathbf{x}' and \mathbf{y}' in parentheses in gray color, to remind that one or both of them may be absent in some of the hopping terms, as is seen in Eqs. (57).

Now recalling the discussion of Sec. 2.1, the combination $(Z_{\mathbf{x}'}) (Z_{\mathbf{y}'}) \sigma_{\mathbf{y}'}^- \mathcal{D}^{\text{SB}}(p, q, p') \lambda_{\mathbf{p}}^+ \lambda_{\mathbf{q}}^+$ squares to zero, allowing it to be identified with the ‘A’ operator of Sec. 2.1. There is also no need to introduce an ancilla qubit because the \mathbf{x} qubit can function as the control to the SVD-transformation gates (\mathcal{V} , \mathcal{W} , or their adjoints). A solution to the SVD transformations is $\mathcal{V} = \mathcal{I}_y \mathcal{I}_p \mathcal{I}_q$ and $\mathcal{W} = X_y \lambda_{\mathbf{p}}^- \lambda_{\mathbf{q}}^-$. The full diagonalizing unitary for the generalized subterm in Eq. (59) thus works out to be

$$\mathcal{U}_{\text{SVD}}^{\text{SB}} \equiv \mathbf{H}_{\mathbf{x}} (|0\rangle \langle 0|_{\mathbf{x}} + |1\rangle \langle 1|_{\mathbf{x}} X_y \lambda_{\mathbf{p}}^+ \lambda_{\mathbf{q}}^+). \quad (60)$$

⁹ $\mathcal{D}^{\text{SB}}(0, q, 0) \equiv 0$ because this would correspond to lowering one of the oscillators beyond its lower cutoff.

¹⁰ p as an eigenvalue versus an operator can be deduced from the context.

¹¹The situation is different for the U(1) theory, where the action of the ladder operators on any state is never equivalent to zero in the untruncated theory, and hence the cyclic wrapping will create a mixing between the lower and upper cutoff states, unless extra operations are introduced to prevent it, see e.g., Ref. [54].

Label	$H_I^{\text{SB}(j)}$ translation							
j	1	2	3	4	5	6	7	8
x	$2r$	$2r$	$2r+1$	$2r+1$	$2r+2$	$2r+2$	$2r+3$	$2r+3$
y	$2r+3$	$2r+2$	$2r+2$	$2r+3$	$2r+1$	$2r$	$2r$	$2r+1$
x'	$2r+1$	$2r+1$	$2r$	$2r$	$2r+3$	$2r+3$	$2r+2$	$2r+2$
y'	$2r+2$	$2r+3$	$2r+3$	$2r+2$	$2r$	$2r+1$	$2r+1$	$2r$
p	$L1$	$L1$	$L2$	$L2$	$L1$	$L2$	$L2$	$L1$
q	$R1$	$R2$	$R2$	$R1$	$R1$	$R1$	$R2$	$R2$
p'	$L2$	$L2$	$L1$	$L1$	$L2$	$L1$	$L1$	$L2$

Table 2: Label associations between registers of a generalized Schwinger-boson subterm, Eq. (59), and the eight subterms of Eqs. (57). For terms that do not involve the Pauli-Z operators on both registers \mathbf{x}' and \mathbf{y}' , the assignment of these registers to the qubit labels is arbitrary, but a choice is being made for concreteness.

When applied to the hopping subterms, one ultimately obtains

$$\mathcal{W}_{\text{SVD}}^{\text{SB}} H_I^{\text{SB}(j)} \mathcal{W}_{\text{SVD}}^{\text{SB} \dagger} = \pm x Z_{\mathbf{x}}(Z_{\mathbf{x}'}) (Z_{\mathbf{y}'}) |1\rangle \langle 1|_{\mathbf{y}} \mathcal{D}^{\text{SB}}(p, q, p'). \quad (61)$$

Lemma 3.3. *Using a $(4\eta + 4)$ -qubit register and no ancilla qubits, $e^{-itH_I^{\text{SB}(r)}}$ can be simulated approximately as $\prod_{j=1}^8 e^{-itH_I^{\text{SB}(j)}(r)}$ using at most $16 \times 8^\eta + 64\eta^2 + 64\eta + 32$ CNOT gates (and a number of single-qubit Z rotations).*

Proof. This cost consists of the (near-term) cost of $\mathcal{W}_{\text{SVD}}^{\text{SB}}$, plus the (near-term) cost of implementing a diagonalized subterm. The $\mathcal{W}_{\text{SVD}}^{\text{SB}}$ circuit, shown in Fig. 7, consists of one CNOT, two controlled- λ^+ ($C(\lambda^+)$) gates, and a Hadamard gate. A single $C(\lambda^+)$ gate is identical to the uncontrolled incrementer on $\eta + 1$ qubits (followed by a bit flip on the “control” qubit), which requires $2\eta(\eta + 1)$ CNOT gates, according to Lemma A.1. The total CNOT-gate cost of a single $\mathcal{W}_{\text{SVD}}^{\text{SB}}$ execution is then $1 + 2 \times 2\eta(\eta + 1) = 4\eta^2 + 4\eta + 1$, as reported in Table 3. The dominant CNOT-gate cost, however, is associated with implementing the diagonalized subterms, which will be discussed next.

The diagonalized subterms to be simulated, according to Eq. (61), are of the form:

$$\pm x Z_{\mathbf{x}}(Z_{\mathbf{x}'}) (Z_{\mathbf{y}'}) |1\rangle \langle 1|_{\mathbf{y}} \mathcal{D}^{\text{SB}}(p, q, p').$$

The function $\mathcal{D}^{\text{SB}}(p, q, p') = \sqrt{pq/((p+p')(p+p'+1))}$, in general, cannot be expressed by a low-degree polynomial, and the exact Pauli decomposition requires up to $(\Lambda + 1)^3 = 2^{3\eta}$ terms, which can be solved for numerically, see Eq. (16). Note that if a fixed accuracy is sought in evaluating \mathcal{D}^{SB} , one may consider replacing the complete expansion by an approximate one that uses less terms. Later we will consider two approaches to such approximation: “small-angle truncation” and “input truncation”, to be discussed below. For now, we will cost the simulation of the diagonal hopping propagator in the worst-case scenario, that is assuming the Pauli decomposition of the phase function \mathcal{D}^{SB} is not truncated, and that all the $2^{3\eta}$ terms in the expansion come with non-zero coefficients.

To find an efficient gate implementation of a full Pauli decomposition, first consider an individual Pauli- Z string $\prod_{i=1}^m Z_i$, mapped to an N -bit binary string ($m \leq N$), with a ‘0’ (‘1’) for bits that are acted on by ‘1’ (Z). In order to implement the rotation $e^{-i\alpha \prod_{i=1}^m Z_i}$, one has to compute the parity of the corresponding collection of bits labelled as 1. This can be done by computing that parity using m CNOT gates, each one controlled by one of the ‘1’ qubits and all with a shared target ancilla (initialized to 0), then applying $e^{-i\alpha Z}$ on that ancilla, and finally uncomputing the parity by reapplying the same m CNOT gates. To effect the full Pauli decomposition, it is sufficient to implement $e^{-i\alpha \prod_{i=1}^m Z_i}$ back to back for all possible collections of Z rotations on m qubits and for all values of m . In practice, the order in which the various Pauli- Z string operators are simulated is carefully chosen such that the overwhelming majority of CNOT gates involved with computing and uncomputing parities cancel each other out, leading to significant savings over applying each rotation, with all the CNOT gates involved, in isolation.

This cost reduction can be explained and quantified in terms of a Hamiltonian cycle on a hypercube. Consider $B = \{0, 1\}^N$ to be the set of bit strings, where the bit string $(b_{N-1}, b_{N-2}, \dots, b_1, b_0)$ is associated with the diagonal operator $Z_{N-1}^{b_{N-1}} Z_{N-2}^{b_{N-2}} \dots Z_1^{b_1} Z_0^{b_0}$. Each bit string $b \in B$ can be visualized as a corner on an N -dimensional hypercube. For any N -dimensional hypercube, there always exists a Hamiltonian cycle starting and ending at the all-zeroes vertex. Such a Hamiltonian cycle through the hypercube provides a Gray code for iterating

Schwinger-boson subterm SVD subroutine	CNOT count
‘Explicit’ CNOT	1
$C(\lambda^\pm)$	$2\eta(\eta + 1)$
Overall $\mathcal{U}_{\text{SVD}}^{\text{SB}}$ circuit for each subterm	$4\eta^2 + 4\eta + 1$
<hr/>	
Schwinger-boson subterm diagonal-rotations subroutine	CNOT count
Compute J-W parity of $\tau = 0, 1, 2$ number of Pauli-Z operators	τ
Hamiltonian cycle through ‘hypercube’ of Pauli parities	$2^{3\eta+1}$
Uncompute J-W parity	τ
Overall diagonal-rotations circuit	$2^{3\eta+1} + 2\tau$
<hr/>	
Schwinger-boson hopping-propagator routine	CNOT count
Subterm SVD	$4\eta^2 + 4\eta + 1$
Subterm diagonal rotations (Pauli-Z-string length of τ)	$2^{3\eta+1} + 2\tau$
Subterm SVD ⁻¹	$4\eta^2 + 4\eta + 1$
Full $e^{-itH_I^{\text{SB}}(r)}$ circuit	$16 \times 8^\eta + 64\eta^2 + 64\eta + 32$

Table 3: Summary of the CNOT-gate cost associated with near-term simulation of off-diagonal operators in the Schwinger boson Hamiltonian, as explained in the text.

through all the bit strings in B , and by extension a ‘‘Gray code’’ of the diagonal Pauli operators.¹² Now, each vertex is associated with a diagonal Pauli operator, and to simulate that operator one needs the associated parity. The starting vertex of all zeroes corresponds to the identity, for which one applies a global phase. Then, each step along the Hamiltonian cycle between adjacent corners b and b' corresponds to a CNOT gate in the following way: if the bit strings b and b' differ in the k^{th} bit, then the parity of b' is derived from that of b by applying a CNOT controlled by qubit k onto the target ancilla. Between each step (i.e., each CNOT implementation), one applies $e^{-i\alpha Z}$ on the ancilla to effect the rotation of a distinct operator. The cost of traversing all 2^N bit strings and returning to the origin will then be 2^N CNOT gates and $2^N - 1$ non-trivial ancilla rotations. Note that one can adapt the above procedure for when an ancilla is not desired, but the overall cost will be identical asymptotically.¹³

In the case at hand, the diagonal operator $|1\rangle\langle 1|_{\text{y}} \mathcal{D}^{\text{SB}}(p, q, p')$ depends on $3\eta + 1$ qubits, implying a cost of $2^{3\eta+1}$ CNOT gates. However, this has not yet accounted for the $Z_{\mathbf{x}}(Z_{x'}) (Z_{y'})$ factor on the diagonalized subterm. The \mathbf{x} qubit can simply act as the ancilla for all parity evaluations, costing no additional CNOT gates. If $\tau = 0, 1, 2$ number of operators in $\{Z_{x'}, Z_{y'}\}$ are called for, then an additional overhead of 2τ CNOT gates is needed. This leads to the total cost of $2^{3\eta+1} + 2\tau$ for a given diagonalized Schwinger-boson subterm, as reported in Table 3.

The cumulative CNOT cost associated with a complete hopping term is obtained as 16 times the cost of a single $\mathcal{U}_{\text{SVD}}^{\text{SB}}$ circuit, plus the individual costs of all eight diagonalized subterms. This works out to $16 \times 8^\eta + 64\eta^2 + 64\eta + 32$, also reported in Table 3. \square

Given the significant cost of implementing $e^{-itZ_{\mathbf{x}}(Z_{x'}Z_{y'})\mathcal{D}^{\text{SB}}(p,q,p')}$ exactly, one may seek approximations to the exact implementation in the near-term scenario, such as the cases considered in Refs. [145, 146] where the gate-count scaling is polynomial in the number of qubits and the reciprocal error tolerance. To that end, we explored the following two possible means of truncating the Pauli decompositions.

Small-angle truncation.—Let $\mathcal{D}_k^{\text{SB}}(p, q, p')$ denote the Pauli decomposition of $\mathcal{D}^{\text{SB}}(p, q, p')$ where only the first k terms with the largest coefficients are retained, a method that we here term as ‘‘small-angle truncation’’. If the goal is to match the ideal evolution up to a fixed spectral-norm error, the Pauli decomposition usually can only be truncated minimally, or not at all. Figure 4 plots the error trade-off per number of Pauli-Z strings retained (dropping the fermionic Z -rotations for simplicity) for $\eta = 2, 3$ (or $\eta = 3, 7$ for the LSH formulation) and at a fixed simulation time $t = x^{-1}$, where x is the hopping strength. Note that for the Schwinger-boson formulation (only), the diagonal function can be modified to incorporate knowledge of the AGL constraint:

¹²Gray codes for N -bit strings can be generated efficiently using, e.g., the recursive algorithm presented in Ref. [157]. The Gray code for $N = 1$ bit is $(0, 1)$. The $(N + 1)$ -bit code is obtained by concatenating one copy of the N -bit code with a reversed copy of the N -bit code, prepending a 0 to the first copy, and prepending a 1 to the reverse copy. Hence, for $N = 2$ one has $(0, 1) \rightarrow (0, 1, 1, 0) \rightarrow (00, 01, 11, 10)$, for $N = 3$ one obtains $(000, 001, 011, 010, 110, 111, 101, 100)$, and so on.

¹³For similar approaches on efficient and exact implementation of diagonal functions, see e.g., Ref. [158], in which an optimal cost is reported to be $2^N - 2$ CNOT gates and $2^N - 1$ single- R^Z rotations.




Figure 4: The spectral-norm error of approximating $e^{-ixt\mathcal{D}}$ at $xt = 1$ as function of k , the number of operators kept from the full Pauli decomposition of \mathcal{D} , defined in Eq. (58) (Eq. (80)) for the Schwinger-boson (LSH) hopping propagator, after truncating the smallest rotations. The Schwinger-boson diagonal function is modified for these plots to vanish for quantum numbers that cannot satisfy the constraint $q \leq p + p'$ imposed by the AGL. The truncated diagonal functions containing the k largest rotations are called \mathcal{D}_k , and the error here is defined as the spectral norm of $e^{-ixt\mathcal{D}} - e^{-ixt\mathcal{D}_k}$. η denotes the number of qubits per bosonic quantum number in either formulation. The diagonal functions considered are six-qubit (red) or nine-qubit (blue) operators. In other words, unequal η values (or different cutoffs) are chosen among the two formulations in order to enable direct comparisons of the accuracy achieved for truncating Pauli decompositions of equal sizes.

input states with $q > p + p'$ are inconsistent with the AGL, so \mathcal{D}^{SB} can be set to zero on such states. The results obtained in Fig. 4 correspond to such a modified function. It is observed that the spectral norm of the difference $e^{-ixt\mathcal{D}^{\text{SB}}(p,q,p')} - e^{-ixt\mathcal{D}_k^{\text{SB}}(p,q,p')}$ decreases rather slowly as k is increased until k is close to its maximal value $2^{3\eta}$, hence making it challenging to achieve high accuracy for any $k \ll 2^{3\eta}$. For example, consider a lattice with $L = 4$ and $tx = 1$. There are 48 s applications of the diagonal hopping propagator required in a second-order product formula with s Trotter steps. For $\eta = 2$ and for only a single Trotter step ($s = 1$), if the overall budget for the simulation is equal to 0.1, then the error budget per individual hopping subterm is only $0.1/48 \approx 2.08 \times 10^{-3}$. According to Fig. 4, the Pauli decomposition would then have to be truncated at $k \geq 62$ (out of 64 possible terms). Note that in this example, error tolerance does not even account for the Trotterization error.

Besides the fact that the truncation of small-angle rotations may not be an efficient way of reducing the simulation cost for the examples of this work, as was demonstrated in Fig. 4, there exist other problems with this method. Importantly, truncation of small angles alone does not specify *a priori* what Pauli operators will have to be simulated. The angles will have to be pre-computed classically,¹⁴ and the circuits that will have to be generated may, in principle, involve all qubits. For general functions, no systematic structure is guaranteed for the retained k terms, making it difficult to place a tight upper bound on the number of CNOT gates called for, as well as to give an explicit algorithm for cycling through all k rotations.

Input truncation.—To address these problems, an alternative truncation scheme may be used, which we term “input truncation”. This amounts to taking the integer arguments to the diagonal function, e.g., p , q , and p' in the Schwinger-boson formulation, and truncating some number of the least significant bits of those inputs. The resulting Pauli decomposition can only involve Pauli- Z operators for the retained bits of the input, with each truncated bit giving a factor of two reduction in the total set of operators. As a simple example, consider the matrix $\text{diag}(a, b, c, d)$, a diagonal function on two qubits whose Pauli decomposition involves $\mathcal{I}_1 \otimes \mathcal{I}_2$, $Z_1 \otimes \mathcal{I}_2$, $\mathcal{I}_1 \otimes Z_2$, and $Z_1 \otimes Z_2$. If dependence on the least significant bit is dropped, this matrix becomes $\text{diag}(a, a, c, c)$, which is a linear combination of $\mathcal{I}_1 \otimes \mathcal{I}_2$ and $Z_1 \otimes \mathcal{I}_2$ only.

Following input truncation, one can use the same Hamiltonian-cycle method previously described for the case of all qubits, but now being limited to a hypercube in a lower-dimensional space. The 2^N scaling of the CNOT gates for this method implies that the number of required CNOT gates will be halved for each dropped bit of the input. In Fig. 5, the spectral-norm errors resulting from the input-truncation method are shown for the

¹⁴Nonetheless, efficient (polynomial-time) algorithms have been developed to obtain a finite number of terms in the series expansion of functions [159]—algorithms that may be employed in this context to mitigate the classical-computational cost of finding the truncated Pauli expansion of functions.




Figure 5: The spectral-norm error of approximating e^{-ixtD} at $xt = 1$ as function of k , the number of operators in the Pauli decomposition of D after the least significant bits of its bosonic arguments are truncated. The Schwinger-boson diagonal function was modified for these plots to vanish for quantum numbers that cannot satisfy the constraint $q \leq p+p'$ imposed by the AGL. In the input-truncation scheme, k is a power of two. The truncated diagonal functions depending on only $\log_2(k)$ qubits are called \mathcal{D}_k , and the error here is defined as the spectral norm of $e^{-ixtD} - e^{-ixtD_k}$. η denotes the number qubits per bosonic quantum number in either formulation. The diagonal functions considered are all nine-qubit operators for both the Schwinger-boson and the LSH formulations. In other words, unequal η values (or different cutoffs) are chosen among the two formulations in order to enable direct comparisons of the accuracy achieved for truncating Pauli decompositions of equal sizes.

case of diagonal operators acting on nine total qubits for the Schwinger-boson (blue diamonds) and the LSH (red circles) diagonal functions. To obtain the results in Fig. 5, as was the case for Fig. 4, the diagonal function of the Schwinger-boson formulation (only) is modified to incorporate knowledge of the AGL constraint: input states with $q > p+p'$ are inconsistent with the AGL, so \mathcal{D}^{SB} was modified to be zero on such states. For ease of comparison, data from small-angle truncation of the nine-qubit operators in Fig. 4 are also displayed in Fig. 5. It is evident that, at least for the examples considered, the error bounds tend to be looser from this method, indicating a trade-off between tightening the theoretical error bound and lowering the number of the required CNOT gates systematically. Better characterization of the pros and cons of the two methods outlined here will be needed in future work on the near-term implementations of the SU(2) and similar theories, in which the computation of a non-trivial dynamical phase is necessary.

—Far term—

Implementing mass propagators

Lemma 3.4. *The mass propagator $e^{-itH_M^{\text{SB}}(r)}$ can be implemented, without approximation, using no ancilla qubits, two R^Z gates, and no additional T gates.*

Proof. The R^Z gate count follows directly from Eq. (54). It must be noted that in the far-term scenario, the synthesis of each R^Z calls for a number of T gates that depends on the desired accuracy in the rotation angle. We will assume there exists a cost function $\mathcal{C}_z(\epsilon)$ giving the number of T gates required to simulate an R^Z gate (of any input angle) within a spectral-norm error of ϵ . The repeat-until-success (RUS) method of Ref. [160] has an average cost $\mathcal{C}_z(\epsilon) \approx 1.15 \log_2(2/\epsilon) + 9.2$, giving a T -count scaling of $\log(1/\epsilon)$ in the limit of high precision. The trivial cost of a mass propagator is recorded in Table 4. \square

Implementing electric propagators

Lemma 3.5. *The electric propagator $e^{-itH_E^{\text{SB}}(r)}$ can be implemented, without approximation, using $5\eta+6$ ancilla qubits, $3\eta+3$ R^Z gates, and an additional $16\eta^2+32\eta+8$ T gates.*

Proof. The procedure leading to this cost consists of the following steps:

1. Compute N^L .

Schwinger-boson mass propagator subroutine	T gates	Workspace	Scratch space
Each $E^{-itH_M^{\text{SB}(j)}}$ rotation	$\mathcal{C}_z(\epsilon)$	0	0
Full $e^{-itH_M^{\text{SB}(r)}}$ circuit	$2\mathcal{C}_z(\epsilon)$	0	0

Schwinger-boson electric propagator subroutine	T gates	Workspace	Scratch space
(Un)compute N^L	4η	η	1
Copy	0	0	$\eta + 1$
$(\eta + 1) \times (\eta + 1)$ mult.	$8\eta^2 + 12\eta + 4$	$2\eta + 2$	$2\eta + 2$
$H_E^{\text{SB}(1)}$ rotations	$(\eta + 1)\mathcal{C}_z(\epsilon)$	0	0
$H_E^{\text{SB}(2)}$ rotations	$(2\eta + 2)\mathcal{C}_z(\epsilon)$	0	0
Full $e^{-itH_E^{\text{SB}(r)}}$ circuit	$16\eta^2 + 32\eta + 8 + (3\eta + 3)\mathcal{C}_z(\epsilon)$	$2\eta + 2$	$3\eta + 4$

Table 4: Summary of the costs associated with the far-term simulation of the diagonal operators in the Schwinger-boson Hamiltonian, as explained in the text.

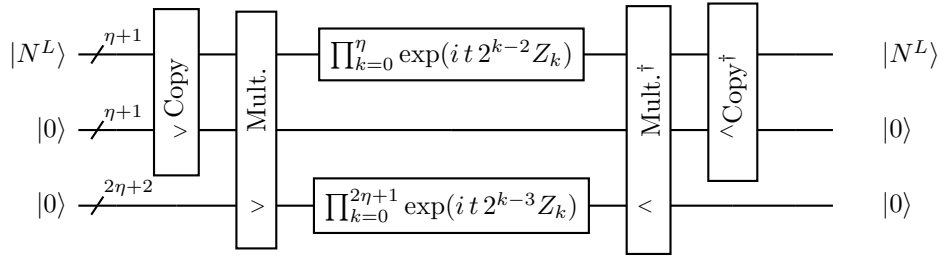


Figure 6: A quantum circuit to realize the phase kickback for $H_E(r)$. The same circuit is applicable to both the Schwinger-boson and the LSH formulations, with the only difference being the evaluation of N^L (the subcircuit that evaluates N^L is not shown). The $>$ symbol denotes that the obtained value in the subcircuit is stored in the corresponding qubit register, and the $<$ symbol indicates that the corresponding register is cleared from the stored values as a result of the action of the inverse subcircuit.

2. Compute $(N^L)^2$.
3. Effect phase kickback via the registers containing N^L and $(N^L)^2$.
4. Uncompute $(N^L)^2$ and N^L .

In step (1), N^L can be computed as

$$|n_1^L\rangle |n_2^L\rangle |0\rangle^{\otimes(1+\eta)} \mapsto |n_1^L\rangle |N^L\rangle |0\rangle^{\otimes\eta} \quad (62)$$

using an η -bit in-place adder. The adder, according to Lemma A.2, calls for η workspace qubits and costs 4η T gates. For (2), $(N^L)^2$ is computed as $|N^L\rangle |0\rangle^{\otimes(5\eta+5)} \mapsto |N^L\rangle |N^L\rangle |0\rangle^{\otimes(4\eta+4)} \mapsto |N^L\rangle |N^L\rangle |(N^L)^2\rangle |0\rangle^{\otimes(2\eta+2)}$, by first copying N^L to an $(\eta + 1)$ -bit register using CNOT gates, and then multiplying the two copies of N^L . According to Lemma A.3, the multiplier costs $8\eta^2 + 12\eta + 4$ T gates and $2\eta + 2$ bits of workspace. In step (3), the $N^L(r)$ and $(N^L(r))^2$ terms of $H_E(r)$ are effectively simulated by applying single-qubit R^Z gates across the $|N^L\rangle$ and $|(N^L)^2\rangle$ registers (up to global phases that are dropped). Finally, in step (4), uncomputation involves reversing the gates of steps (2) and (1) and the associated costs are the same. Steps (2)-(4) are shown in Fig. 6. In total, the procedure outlined above involves $3\eta + 3$ R^Z gates, which can each be done to the desired precision using the RUS method mentioned above [160]. The costs associated with all the subroutines are summarized in Table 4. The final, quoted costs are obtained by adding up the workspace and scratch-space sizes. \square

Implementing hopping propagators

Lemma 3.6. *Let $\eta = \log_2(\Lambda + 1)$ be the number of qubits per Schwinger boson mode, $n > 0$ be the number of Newton's method iterations, and $m > 0$ be a fixed binary arithmetic precision. Then $\prod_{j=1}^8 e^{-itH_I^{\text{SB}(j)}(r)}$ can be implemented within an additive spectral-norm error of*

$$8xt \left[2^n (\sqrt{2} - 1)^{2^n} + 2^{2-m} \left(\left(\frac{3}{2} \right)^n - 1 \right) \right] \quad (63)$$

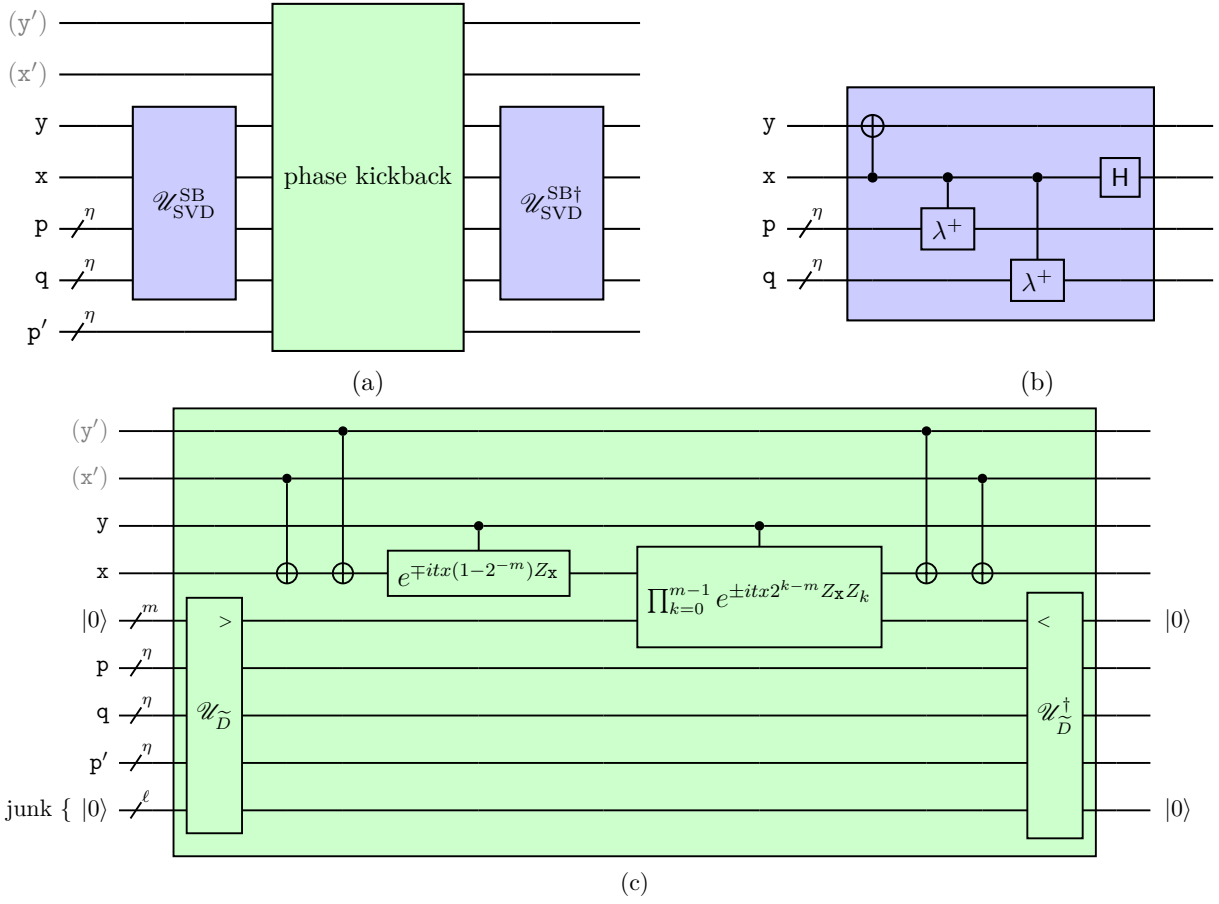


Figure 7: The circuit that implements the Schwinger-boson hopping propagator corresponding to each subterm in Eq. (57). (a) A high-level representation of the (far-term) circuit that realizes $e^{-itH_I^{\text{SB}(j)(r)}}$. (b) The diagonalization circuit $\mathcal{U}_{\text{SVD}}^{\text{SB}}$ for a Schwinger-boson hopping subterm. This circuit applies to both the near-term and far-term algorithms. (c) The phase-kickback circuit used in the far-term scenario. All circuits in (a-c) call for ancilla qubits that are not explicitly drawn but are discussed in the text and counted in the cost tables.

using

$$4n\eta + 27\eta + \max(2\eta + 2, m) + 3mn + 10m + 3n + 15 \quad (64)$$

ancilla qubits, $16(m+1)$ R^Z gates, and

$$2560\eta mn + 512m^2n + 768\eta^2 + 256\eta m + 256\eta n + 1408mn - 128n \max(4\eta + 2, 2m) + 896\eta - 64 \max(2\eta, m) - 64n \quad (65)$$

T gates.

Proof. To obtain this result, one must circuitize the propagators associated with each of the eight subterms of the hopping term. One subterm is simulated by applying an appropriate diagonalizing transformation $\mathcal{U}_{\text{SVD}}^{\text{SB}}$, implementing the diagonalized subterm, and finally changing basis back via $\mathcal{U}_{\text{SVD}}^{\text{SB}\dagger}$. In the far-term scenario, the diagonalized subterm will be implemented via phase kickback. The sequence for simulating a hopping subterm is shown schematically in Fig. 7(a).

Diagonalization.—The SVD transformation is as shown in Fig. 7(b): one CNOT plus two $C(\lambda^+)$ gates and a Hadamard gate. The $C(\lambda^+)$ gate, that is a controlled η -qubit incremter, can proceed as an uncontrolled $(\eta+1)$ -qubit incremter (followed by a bit flip on the “control” qubit). According to Lemma A.2, for an input of size $\eta+1$, the cyclic incremter costs $4\eta-4$ T gates and $\eta-2$ ancilla qubits, where the ancilla qubits can be reused for all subsequent implementations. Doubling this, the cost of $\mathcal{U}_{\text{SVD}}^{\text{SB}}$ (or $\mathcal{U}_{\text{SVD}}^{\text{SB}\dagger}$) is $8\eta-8$ T gates and $\eta-2$ workspace qubits, as reported in Table 5. Covering all eight subterms will further multiply the T-gate count by 16, for a total of $128\eta-128$ T gates per hopping term associated with diagonalizing transformations. The dominant cost of simulating a hopping term, however, comes from implementing the diagonalized subterms, which will be discussed next.

Phase kickback.—The implementation of the diagonalized subterm in Eq. (61) proceeds via a phase-kickback algorithm whose circuit implementation is depicted in Fig. 7(c). First, the value of the non-trivial diagonal bosonic operator \mathcal{D}^{SB} is computed, up to a known accuracy, via a unitary $\mathcal{U}_D^{\text{SB}}$, and the value is stored in an ancillary register of m bits, where $\tilde{\mathcal{D}}^{\text{SB}}$ denotes a fixed-precision approximation to \mathcal{D}^{SB} and m is the desired bit-precision. We will construct $\mathcal{U}_D^{\text{SB}}$ via Newton’s method shortly and will discuss its cost as a function of the approximation error. The fixed-precision arithmetic is another source of error, independent of the Trotter and R^Z synthesis errors, to be taken into account in the full error budget. Since the ancillary register holds the (approximate) value of \mathcal{D}^{SB} in binary form, the diagonal operator $e^{-itx|1\rangle\langle 1|_{y(Z_{x'})}(Z_{y'})Z_x\mathcal{D}^{\text{SB}}(p,q,p')}$ can be straightforwardly implemented by R^Z gates on the corresponding qubits, with the angles shown in the circuit. The ancillary register is set back to an all- $|0\rangle$ state by the inverse operations, and is used in subsequent implementations.

There are multiple ways of formulating the calculation of $\mathcal{D}^{\text{SB}}(p,q,p') = \sqrt{\frac{pq}{(p+p')(p+p'+1)}}$. In view of the identity

$$\sqrt{\frac{\text{numerator}}{\text{denominator}}} = \text{numerator} \times \frac{1}{\sqrt{\text{numerator} \times \text{denominator}}},$$

one way of formulating the calculation goes as follows:

1. Multiply p and q to get the numerator.
2. Compute $p + p'$ with an adder.
3. Multiply pq with $p + p'$ to obtain $pq(p + p')$.
4. Compute $p + p' + 1$ with an incrementer.
5. Multiply $pq(p + p')$ with $(p + p' + 1)$ to obtain $pq(p + p')(p + p' + 1)$. For later convenience, this product is referred to by

$$g^{\text{SB}}(p,q,p') \equiv pq(p + p')(p + p' + 1). \quad (66)$$

6. Compute the inverse square root of $g^{\text{SB}}(p,q,p')$.
7. Multiply pq with $g^{\text{SB}}(p,q,p')^{-1/2}$.

In this protocol, the most complicated and resource-intensive step is (6): the inverse-square-root evaluation.¹⁵ The unitary that implements the steps above is called $\mathcal{U}_D^{\text{SB}}$, and the evaluation only produces a fixed-precision approximation $\tilde{\mathcal{D}}^{\text{SB}}(p,q,p') \approx \mathcal{D}^{\text{SB}}(p,q,p')$. We will return shortly to address the details of using fixed-precision arithmetic and the associated errors. A circuit diagram for $\mathcal{U}_D^{\text{SB}}$ is provided in Fig. 8.

To evaluate $g^{-1/2}$ given some non-negative integer g , one may proceed with the following choice

$$F(y) \equiv \frac{1}{y^2} - g, \quad (67)$$

for the function to be used in the implementation of Newton’s method for inverse square root. Here, in the case at hand, $g = g^{\text{SB}}(p,q,p')$. The positive zero, y_z , of $F(y)$ is $g^{-1/2}$, which when is input to a final multiplier for the product pqy_z gives $\mathcal{D}^{\text{SB}}(p,q,p')$. Our method of calculation is motivated by the goal of computing $\mathcal{D}^{\text{SB}}(p,q,p')$ to high accuracy using the fewest number of quantum multiplication circuits. This choice also results in iterations that require only addition and multiplication (and division by two, which has no T-gate and ancilla-qubit cost in binary).

Given some approximation \tilde{y} of $g^{-1/2}$, the next approximation generated by Newton’s method is given by

$$\mathcal{N}_g(\tilde{y}) \equiv \frac{\tilde{y}(3 - \tilde{y}^2 g)}{2}. \quad (68)$$

With an approximation \tilde{y} to $g^{-1/2}$ already stored in some qubit register, a single Newton iteration to generate the next approximation $\mathcal{N}_g(\tilde{y})$ is performed via the circuit shown in Fig. 9. It uses known quantum-circuit routines such as copy, addition, and multiplication that are described in more detail in Appendix A. The total T-gate count of each iteration, as analyzed in Lemma A.4, is $32w^2 + 40kw + 4k + 8w - 8 \max(k, 2w) - 12$, and

¹⁵In Ref. [69], an alternative protocol was proposed that, when translated to the Schwinger-boson formulation, consists of the following steps: (1′) Multiply p and q to get the numerator. (2′) Compute $p + p'$ and $p + p' + 1$ using adders. (3′) Multiply $p + p'$ and $p + p' + 1$ to get the denominator. (4′) Compute the reciprocal of $(p + p')(p + p' + 1)$ from step (3′). (5′) Multiply the results from (1′) and (4′) to get the argument of the square root. (6′) Take the square root of the fractional quantity. In this protocol, the most complicated and resource-intensive steps are (4′) and (5′): the integer reciprocal and the square root. Our route has the benefit of evaluating the inverse square root in a single function-estimation step using Newton’s method.

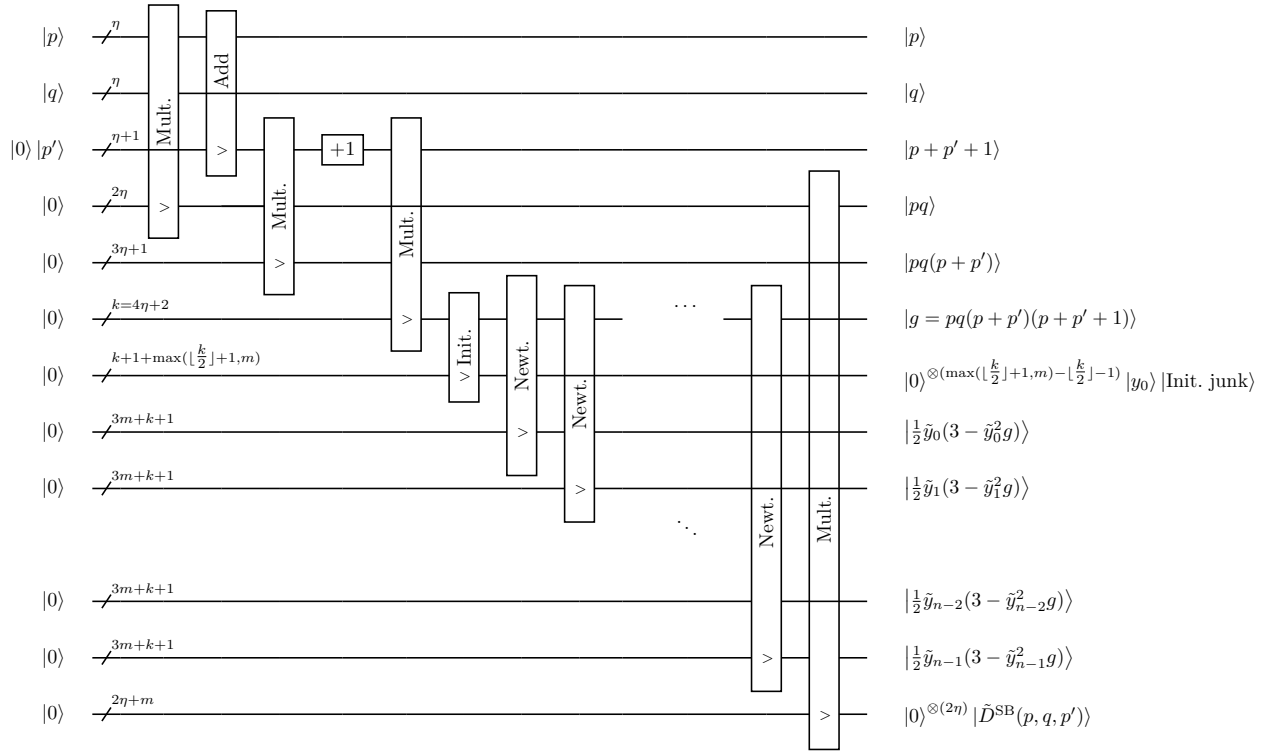


Figure 8: A quantum circuit for computing an m -bit approximation to diagonal function $\mathcal{D}^{\text{SB}}(p, q, p')$ defined in Eq. (58) with η -qubit bosonic registers p , q , and p' . Multiplication (Mult.), addition (Add), incrementers (+1), and initial-guess (Init.) subcircuits are described in Appendix A. The decomposition of a Newton-iteration step (Newt.) is depicted in Fig. 9. For each subcircuit, the line passing through the box indicates that the corresponding qubit register does not participate in the operations. The output of each unitary is marked by $>$ (or $<$ for the inverses). For example, for multiplication, $>$ indicates the register into which the product of the other two input registers is output. Workspace qubits are left implicit. $k \equiv 4\eta + 2$, n indexes the final Newton-iteration step, $y_0 = 2^{-s}$ with $s = \lfloor \log_4(pq(p+p')(p+p'+1)) + \frac{1}{2} \rfloor$, and the state [Init. junk] is stored in $k+1$ qubits, see Lemma A.5.




Figure 9: Newton-iteration circuit for computing the inverse square root of g . Workspace qubits are left implicit and the rest of the notation is as in Fig. 8.

the circuit calls for $9w + 3k + 3$ ancilla qubits, where k is the number of bits specifying g ($k = 4\eta + 2$ for the Schwinger-boson case) and w is the number of bits specifying \tilde{y} .

Conventionally, Newton's method involves generating some initial value y_0 , followed by successive iterates $y_n = \mathcal{N}_g(y_{n-1})$, up to some prescribed number of rounds. Defining the n th-order relative error $\delta_n = (y_n - y_z)/y_z$, it is straightforward to see that

$$\delta_{n+1} = -\frac{3}{2}\delta_n^2 - \frac{1}{2}\delta_n^3. \quad (69)$$

This illustrates that as long as $|\delta_0| < 1$, Newton's method converges fast and the relative error is bounded by $|\delta_n| \leq 2^n |\delta_0|^{2^n}$. Since the convergence rate depends on the initial guess y_0 , in order to guarantee quadratic convergence for any possible input g , instead of some arbitrary initial value, y_0 needs to be computed from g . A suitable choice is the inverse square root of the s^{th} power of four such that s is the closest integer to the logarithm of the input g in base four. A circuit implementation of the initial-value evaluation is presented in Appendix A. According to Lemma A.5, the initial value can be computed into a $(1 + \lfloor k/2 \rfloor)$ -bit register at a cost of $4k$ T gates and $k + 1$ 'junk' qubits, where again k is the number of bits specifying g .¹⁶

A practical issue with Newton's method as described above is that the cost of each Newton iteration grows extremely fast. It is more realistic to specify some moderate number of bits m for the precision that suffices for Newton's method to converge to the exact answer within a bounded error. Thus, we will instead seek approximations $\tilde{y}_0, \tilde{y}_1, \dots, \tilde{y}_n$ where \tilde{y}_i (for $i \geq 1$) is $\mathcal{N}_g(\tilde{y}_{i-1})$ truncated to $m - 1$ bits after the binary point and similarly for the initial value \tilde{y}_0 (note that as argued in Lemma A.6, $y_i \leq 1/\sqrt{g} \leq 1$ for $i \geq 0$). For concreteness, we will always apply the circuit of Fig. 9 with the input being implicitly truncated/padded to m bits as necessary, making the input size w effectively equal to m .

To summarize the phase evaluation, there exists a quantum circuit that implements the unitary $\mathcal{U}_D^{\text{SB}}$ defined as

$$|p\rangle |q\rangle |p'\rangle |0\rangle^{\otimes(1+m+\ell+\beta)} \xrightarrow{\mathcal{U}_D^{\text{SB}}} |p\rangle |q\rangle |p+p'+1\rangle |\tilde{D}^{\text{SB}}(p, q, p')\rangle |\text{junk}\rangle |0\rangle^{\otimes\beta}. \quad (70)$$

The state $|\tilde{D}^{\text{SB}}(p, q, p')\rangle$ is an m -bit fixed-precision encoding of $\mathcal{D}^{\text{SB}}(p, q, p')$ within absolute error $|\delta_{n,m}^{\text{Abs.}}| \leq 2^n (\sqrt{2}-1)^{2^n} + 2^{2-m} ((3/2)^n - 1)$ for n iterations of Newton's method. The state $|\text{junk}\rangle$, which holds intermediate values, occupies an ℓ -bit register. Furthermore, the circuit can be implemented using at most

$$160\eta mn + 32m^2 n + 48\eta^2 + 16\eta m + 16\eta n + 88mn - 8n \max(4\eta + 2, 2m) + 48\eta - 4 \max(2\eta, m) - 4n + 8 \quad (71)$$

T gates. The circuit consumes at most

$$\ell = 4\eta n + 3mn + 15\eta + \max(2\eta + 2, m) + 3n + 6 \quad (72)$$

in output junk qubits, and borrows

$$\beta = 12\eta + 9m + 9 \quad (73)$$

reusable workspace qubits. Lemma A.7 of Appendix A provides a derivation of these results.

Finally, with $\pm x \tilde{D}^{\text{SB}}$ evaluated, the diagonal operator $\pm x |1\rangle \langle 1|_y (Z_{x'}) (Z_{y'}) Z_x \mathcal{D}^{\text{SB}}(p, q, p')$ from Eq. (61) is effectively replaced by $|1\rangle \langle 1|_y (Z_{x'}) (Z_{y'}) Z_x N_{\text{aux}}$, where $N_{\text{aux}} = 2^{1-m} \sum_{k=0}^{m-1} 2^{k-1} (1 - Z_{\text{aux},k})$ is the fixed-point number operator on the m -bit auxiliary register holding \tilde{D}^{SB} (see Eq. (56)). With N_{aux} expanded out into m single-qubit Z operators plus a constant, the rotations can be split into $m + 1$ distinct commuting rotations. As remarked above, the far-term R^Z gates incur a T-gate cost associated with the tunable precision of the desired rotation angles, and one can use the RUS construction of Ref. [160]. These rotations are recorded in Table 5, where an additional factor of two is included to account for the control factor $|1\rangle \langle 1|_y = \frac{1}{2}(1 - Z_y)$.

Final tally.—Putting everything together, the circuit in Fig. 7(a) costs

$$4\eta n + 27\eta + \max(2\eta + 2, m) + 3mn + 10m + 3n + 15 \quad (74)$$

ancilla qubits for work and scratch spaces (where we have added m to account for the primary phase-kickback register), $2(m + 1) R^Z$ gates, and

$$320\eta mn + 64m^2 n + 96\eta^2 + 32\eta m + 32\eta n + 176mn - 16n \max(4\eta + 2, 2m) + 112\eta - 8 \max(2\eta, m) - 8n \quad (75)$$

¹⁶In Ref. [118], it is noted that the number of Newton iterations may be reduced by one for a small overhead in computing a better initial guess. We did not pursue this improvement due to its asymptotically identical cost.

Routine	T gates	Workspace	Scratch space
$\mathcal{U}_{\text{SVD}}^{\text{SB}}$ or $\mathcal{U}_{\text{SVD}}^{\text{SB}\dagger}$	$8\eta - 8$	$\eta - 2$	0
$\mathcal{U}_{\tilde{D}}^{\text{SB}}$ or $\mathcal{U}_{\tilde{D}}^{\text{SB}\dagger}$	$160\eta mn + 32m^2n + 48\eta^2 + 16\eta m + 16\eta n + 88mn - 8n \max(4\eta + 2, 2m) + 48\eta - 4 \max(2\eta, m) - 4n + 8$	$12\eta + 9m + 9$	$4\eta n + 3mn + 15\eta + \max(2\eta + 2, m) + m + 3n + 6$
Diagonal rotations	$2(m + 1)\mathcal{C}_z(\epsilon)$	0	0
Full $e^{-itH_I^{\text{SB}(j)}(r)}$ circuit	$320\eta mn + 64m^2n + 96\eta^2 + 32\eta m + 32\eta n + 176mn - 16n \max(4\eta + 2, 2m) + 112\eta - 8 \max(2\eta, m) - 8n + 2(m + 1)\mathcal{C}_z(\epsilon)$	$12\eta + 9m + 9$	$4\eta n + 3mn + 15\eta + \max(2\eta + 2, m) + m + 3n + 6$

Table 5: Summary of the costs associated with the far-term simulation of the off-diagonal operators in the Schwinger-boson Hamiltonian, as explained in the text. For itemized cost contributing to the cost of $\mathcal{U}_{\tilde{D}}^{\text{SB}}$, see Table 13 of Appendix A.

T gates. Adding the costs of doing eight structurally equivalent subterms, the full upper bound on implementing each Trotterized hopping propagator in the Schwinger-boson formulation comes to eight times the T-gate cost given in Eq. (75) and no more ancilla cost than what is stated in Eq. (74), as stated originally and summarized in Table 5. \square

3.2.2 Loop-string-hadron propagators

In the LSH formulation, the fermionic operators are $\chi_{i(o)}$ and $\chi_{i(o)}^\dagger$, and these can be directly mapped to Pauli spin operators via a Jordan-Wigner transformation. With an additional staggering analogous to what was introduced in the Schwinger-boson formulation, we choose the mapping:

$$\chi_i(r) \rightarrow \left(\prod_{k=0}^{2r-1} Z_k \right) \sigma_{2r}^+, \quad \chi_o(r) \rightarrow \left(\prod_{k=0}^{2r} Z_k \right) \sigma_{2r+1}^+, \quad (76a)$$

$$\chi_i^\dagger(r) \rightarrow \left(\prod_{k=0}^{2r-1} Z_k \right) \sigma_{2r}^-, \quad \chi_o^\dagger(r) \rightarrow \left(\prod_{k=0}^{2r} Z_k \right) \sigma_{2r+1}^-. \quad (76b)$$

The bosonic operators in the LSH formulation consist of n_ℓ (occupation-number) and Γ (ladder) operators, respectively. Assuming that there are η qubits available to encode the local occupation of bosons in binary, the n_ℓ quantum number at each site must be truncated to a finite value, $\Lambda \equiv 2^\eta - 1$. The only modification to the action of bosonic operators is that the raising operator Γ^\dagger must annihilate the state if the occupation of the state at the corresponding site is equal to Λ , that is: $\Gamma^\dagger |n_\ell, n_i, n_o\rangle = (1 - \delta_{n_\ell, \Lambda}) |n_\ell + 1, n_i, n_o\rangle$. The n_ℓ quantum number can be expressed in a binary representation and be mapped to qubit registers as in Eq. (52). Figure 2(b) depicts the Hilbert spaces associated with the DOFs of the LSH formulation.

An efficient circuit decomposition for exponentiating each site- or link-local term in the Hamiltonian will be presented in the following, and product formulas for approximating the full time-evolution operator of the system within a fixed error will be discussed later in Sec. 3.3.

—Near term—

Implementing mass propagators

The mass Hamiltonian is $\sum_{r=0}^{L-1} H_M^{\text{LSH}}(r)$, and each site-local mass term can be decomposed into two commuting subterms as

$$H_M(r) = \sum_{j=1}^2 H_M^{\text{LSH}(j)}(r), \quad (77a)$$

$$H_M^{\text{LSH}(1)}(r) = \frac{\mu}{2} (-1)^{r+1} Z_{2r}, \quad (77b)$$

$$H_M^{\text{LSH}(2)}(r) = \frac{\mu}{2} (-1)^{r+1} Z_{2r+1}, \quad (77c)$$

after the Jordan-Wigner transformation in Eq. (76), and upon neglecting constant terms that only introduce time-dependent but otherwise constant phases in the dynamics.

Lemma 3.7. *Using a 2-qubit register with no ancilla qubits, $e^{-itH_M^{\text{LSH}}(r)}$ can be implemented without approximation, up to a phase, with no CNOT gates required.*

Proof. The expressions in Eqs. (77) are identical to those of the Schwinger-boson formulation in Eqs. (53), and so the circuit decomposition of this propagator is trivial as with the Schwinger-boson case. Two qubits are required, each indexed by $2r$ and $2r+1$, corresponding to ‘ $n_i(r)$ ’ and ‘ $n_o(r)$ ’ registers, respectively. The circuit implements $R_{2r}^Z((-1)^{r+1}\mu t) R_{2r+1}^Z((-1)^{r+1}\mu t)$. \square

Implementing electric propagators

The electric Hamiltonian is $\sum_{r=0}^{L-2} H_E^{\text{LSH}}(r)$, with link-local Casimir operators decomposed into three commuting subterms as

$$H_E^{\text{LSH}}(r) = \sum_{j=1}^3 H_E^{\text{LSH}(j)}(r), \quad (78a)$$

$$H_E^{\text{LSH}(1)}(r) = \frac{1}{2}n_\ell(r), \quad (78b)$$

$$H_E^{\text{LSH}(2)}(r) = \frac{1}{4}n_\ell(r)^2, \quad (78c)$$

$$H_E^{\text{LSH}(3)}(r) = \frac{1}{2}(n_\ell(r) + \frac{3}{2})n_o(r)(1 - n_i(r)). \quad (78d)$$

As remarked for the Schwinger-boson electric propagators, there is a certain amount of arbitrariness involved with dividing $H_E(r)$ into commuting subterms, but any reasonable choice will do since slight difference in the cost associated with those choices will not matter given that simulating the interaction Hamiltonian will dominate the final cost.

The division into subterms in Eqs. (78) for the near-term simulation involves three kinds of subterms: one proportional to n_ℓ , one proportional to n_ℓ^2 , and the last being the coupling of n_ℓ to n_i and n_o . The CNOT-gate and ancilla-qubit requirements of implementing $e^{-itH_E^{\text{LSH}}(r)}$ in the near-term scenario can be derived following similar analysis to the Schwinger-boson case.

Lemma 3.8. *Using an $(\eta + 2)$ -qubit register with no ancilla qubits, $e^{-itH_E^{\text{LSH}}(r)}$ can be implemented without approximation, up to a phase, using $\frac{1}{2}\eta^2 + \frac{17}{2}\eta + 1$ CNOT gates (and a number of single-qubit rotations).*

Proof. This cost is obtained as follows. i) The operator n_ℓ is a linear combination of the identity and every single-qubit Pauli- Z operator across the ‘ n_ℓ ’ register. Simulating this term amounts to a global phase and η single-qubit Z rotations, but no CNOT gates. ii) The n_ℓ^2 term can be simulated like an ‘ E^2 ’ term in the U(1) theory [54], similar to what was described in the Schwinger-boson section. Following the same reasoning leads to a CNOT-gate cost of $(\eta + 2)(\eta - 1)/2$ for this term. iii) As for the remaining $(n_\ell + 3/2)n_o(r)(1 - n_i(r))$, a naive Pauli decomposition involves $2\eta + 1$ operators of the form $Z \otimes Z$, plus η of the form $Z \otimes Z \otimes Z$, leading to a cost of $2(2\eta + 1) + 4\eta = 8\eta + 2$ CNOT gates. The gate costs of i), ii), and iii) are summarized in Table 6, adding up to the stated total CNOT-gate count of $\frac{1}{2}\eta^2 + \frac{17}{2}\eta + 1$. \square

Implementing hopping propagators

The gauge-matter interaction Hamiltonian is $\sum_{r=0}^{L-2} H_I^{\text{LSH}}(r)$, with link-local hopping terms $H_I^{\text{LSH}}(r)$ that are off-diagonal in the electric basis. Like in the Schwinger-boson formulation, the hopping terms do not readily split into a set of mutually commuting subterms and we resort to an approximation $e^{-itH_I^{\text{LSH}}(r)} \approx \prod_j e^{-itH_I^{\text{LSH}(j)}(r)}$, in which each subterm is simulated without further splitting. Once again, the splitting into ν hopping subterms directly impacts the Trotter error bound and gate counts. In the LSH formulation, one is naturally led to a splitting into $\nu^{\text{LSH}} = 2$ subterms. After applying the Jordan-Wigner transformations in Eqs. (76), along with the truncation to η -bit registers for each bosonic oscillator mode, the two subterms in the qubit space are:

$$H_I^{\text{LSH}(1)}(r) = x \sigma_{2r}^+ Z_{2r+1} \sigma_{2r+2}^- \Gamma^\dagger(r)^{1-n_{2r+1}} \Gamma^\dagger(r+1)^{n_{2r+3}} \mathcal{D}^{\text{LSH}}(n_\ell(r), n_{2r+1}, n_{2r+3}) + \text{H.c.}, \quad (79a)$$

$$H_I^{\text{LSH}(2)}(r) = x \sigma_{2r+1}^- Z_{2r+2} \sigma_{2r+3}^+ \Gamma^\dagger(r)^{n_{2r}} \Gamma^\dagger(r+1)^{1-n_{2r+2}} \mathcal{D}^{\text{LSH}}(n_\ell(r+1), n_{2r+2}, n_{2r}) + \text{H.c.} \quad (79b)$$

LSH mass propagator subroutine	CNOT count
R^Z gate	0
Full $e^{-itH_M^{\text{LSH}}(r)}$ circuit	0

LSH electric propagator subroutine	CNOT count
Simulate $e^{-itn_\ell/2}$	0
Simulate $e^{-itn_\ell^2/4}$	$\frac{1}{2}(\eta+2)(\eta-1)$
Simulate $e^{-it(n_\ell+3/2)n_o(1-n_i)/2}$	$8\eta+2$
Full $e^{-itH_E^{\text{LSH}}(r)}$ circuit	$\frac{1}{2}\eta^2 + \frac{17}{2}\eta + 1$

Table 6: Summary of the costs associated with the near-term simulation of the diagonal operators in the LSH Hamiltonian, as explained in the text.

Label	$H_I^{\text{LSH}(j)}$ translation	
j	1	2
\mathbf{x}'	$(r, i) \sim 2r$	$(r+1, o) \sim 2r+3$
\mathbf{x}	$(r, o) \sim 2r+1$	$(r+1, i) \sim 2r+2$
\mathbf{y}'	$(r+1, i) \sim 2r+2$	$(r, o) \sim 2r+1$
\mathbf{y}	$(r+1, o) \sim 2r+3$	$(r, i) \sim 2r$
\mathbf{p}	(r, ℓ)	$(r+1, \ell)$
\mathbf{q}	$(r+1, \ell)$	(r, ℓ)

Table 7: Label associations between registers of a generalized LSH subterm, Eq. (81), and the two subterms of Eqs. (79).

The action of the LSH ladder operators, Γ and Γ^\dagger , on the corresponding LSH basis was introduced in Sec. 3.1.2, with the finite-cutoff modification introduced above, while the diagonal function \mathcal{D}^{LSH} is defined as

$$\mathcal{D}^{\text{LSH}}(p, n, n') \equiv \sqrt{\frac{p+1+n}{p+1+n'}}, \quad (80)$$

for $n, n' = 0, 1$ being fermionic quantum numbers.

The expressions in Eqs. (79) have been put into the form of (shifting operators) \times (diagonal operators) to match the example Hamiltonians studied in Sec. 2.1 so that the SVD diagonalization methods are readily applied. In the language of Sec. 2.1, each term $H_I^{\text{LSH}(j)}$ involves two spin raising and lowering operators (acting on the $n_{i(o)}$ quantum numbers at two adjacent sites), two conditional bosonic raising or lowering operators (acting on n_ℓ quantum numbers at the same adjacent sites), and a non-trivial diagonal function (of one bosonic and two fermionic occupation numbers). The general structure of the LSH hopping subterms is encapsulated by

$$\begin{aligned} H_I^{\text{LSH}(j)} &= x \sigma_{y'}^- Z_x \sigma_{x'}^+ (\Gamma_q^\dagger)^{n_y} (\Gamma_p^\dagger)^{1-n_x} \mathcal{D}^{\text{LSH}}(p, n_x, n_y) + \text{H.c.} \\ &= x Z_x (1 - |1\rangle \langle 1|_y |\Lambda\rangle \langle \Lambda|_p) (1 - |0\rangle \langle 0|_x |\Lambda\rangle \langle \Lambda|_p) \mathcal{D}^{\text{LSH}}(p, n_x, n_y) \sigma_{y'}^+ \sigma_{x'}^- (\lambda_q^-)^{n_y} (\lambda_p^-)^{1-n_x} + \text{H.c.} \end{aligned} \quad (81)$$

with the labels $\{\mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}'\}$ introduced for fermionic registers, and $\{\mathbf{p}, \mathbf{q}\}$ for bosonic registers. The precise mapping to modes for each subterm is displayed in Table 7. Note that $n_\ell(r)$ and $n_\ell(r+1)$ are effectively equal under the action of the Hamiltonian terms in Eqs. (79) because of the AGL, hence the replacement $|\Lambda\rangle \langle \Lambda|_q \rightarrow |\Lambda\rangle \langle \Lambda|_p$ in the second line of Eq. (81).

Now recalling the discussion of Sec. 2.1, the combination

$$x Z_x (1 - |1\rangle \langle 1|_y |\Lambda\rangle \langle \Lambda|_p) (1 - |0\rangle \langle 0|_x |\Lambda\rangle \langle \Lambda|_p) \mathcal{D}^{\text{LSH}}(p, n_x, n_y) \sigma_{x'}^- (\lambda_q^-)^{n_y} (\lambda_p^-)^{1-n_x}$$

squares to zero, allowing it to be identified with the ‘ A ’ operator of Sec. 2.1. There is also no need to introduce an ancilla because the \mathbf{y}' qubit can function as the control to the SVD-transformation gates (\mathcal{V} , \mathcal{W} , or their adjoints). A solution to the SVD transformations is $\mathcal{V} = \mathcal{I}_{x'} \mathcal{I}_p \mathcal{I}_q$ and $\mathcal{W} = X_{x'} (\lambda_q^+)^{n_y} (\lambda_p^+)^{1-n_x}$. The full diagonalizing unitary for the generalized subterm in Eq. (81) therefore works out to be

$$\mathcal{U}_{\text{SVD}}^{\text{LSH}} \equiv H_{y'} (|0\rangle \langle 0|_{y'} + |1\rangle \langle 1|_{y'}) X_{x'} (\lambda_q^-)^{n_y} (\lambda_p^-)^{1-n_x}. \quad (82)$$

LSH subterm SVD subroutine	CNOT count	Ancillas
‘Explicit’ CNOT	1	0
$C^2(\lambda^\pm)$	$2\eta^2 + 2\eta + 6$	1
Overall $\mathcal{U}_{\text{SVD}}^{\text{LSH}}$ circuit for each subterm	$4\eta^2 + 4\eta + 13$	1

LSH subterm diagonal-rotations subroutine	CNOT count	
Hamiltonian cycle through ‘hypercube’ of Pauli parities	$2^{\eta+3}$	0
Overall diagonal-rotations circuit	$2^{\eta+3}$	0

LSH hopping-propagator routine	CNOT count	
Subterm SVD	$4\eta^2 + 4\eta + 13$	1
Subterm diagonal rotations	$2^{\eta+3}$	0
Subterm SVD $^{-1}$	$4\eta^2 + 4\eta + 13$	1
Full $e^{-itH_I^{\text{LSH}}(r)}$ circuit	$16 \times 2^\eta + 16\eta^2 + 16\eta + 52$	1

Table 8: Summary of the costs associated with the near-term simulation of the off-diagonal operators in the LSH Hamiltonian, as explained in the text.

When applied to the hopping subterms, one ultimately obtains

$$\mathcal{U}_{\text{SVD}}^{\text{LSH}} H_I^{\text{LSH}(j)} \mathcal{U}_{\text{SVD}}^{\text{LSH}\dagger} = x |1\rangle \langle 1|_{\mathbf{x}'} Z_{\mathbf{y}'} Z_{\mathbf{x}} (1 - |1\rangle \langle 1|_{\mathbf{y}} |\Lambda\rangle \langle \Lambda|_{\mathbf{p}}) (1 - |0\rangle \langle 0|_{\mathbf{x}} |\Lambda\rangle \langle \Lambda|_{\mathbf{p}}) \mathcal{D}^{\text{LSH}}(p, n_{\mathbf{x}}, n_{\mathbf{y}}). \quad (83)$$

Lemma 3.9. *Using a $(2\eta + 4)$ -qubit register plus 1 ancilla qubit, $e^{-itH_I^{\text{LSH}}(r)}$ can be Trotterized using at most $16 \times 2^\eta + 16\eta^2 + 16\eta + 52$ CNOT gates (and a number of single-qubit Z rotations).*

Proof. This cost consists of the (near-term) cost of $\mathcal{U}_{\text{SVD}}^{\text{LSH}}$, plus the cost of implementing a diagonalized subterm. The $\mathcal{U}_{\text{SVD}}^{\text{LSH}}$ circuit, shown in Fig. 10, consists of one CNOT, two $C^2(\lambda^+)$ gates, and a Hadamard gate. A single $C^2(\lambda^+)$ gate can be achieved by performing the logical AND of the two control qubits into a blank ancilla qubit, then using that ancilla qubit as the single control to a $C(\lambda^\pm)$ gate, followed by uncomputing the AND. The AND gates can be done with 3 CNOT gates each using the construction of Ref. [161] (Sec. VI-B), while the $C(\lambda^\pm)$ gate can be done with $2\eta(\eta + 1)$ CNOT gates as explained previously (see Table 3). Thus, the cost of a single $C^2(\lambda^+)$ gate is $3 + 2\eta(\eta + 1) + 3 = 2\eta^2 + 2\eta + 6$ CNOT gates and one ancilla qubit. The total CNOT-gate cost of a single $\mathcal{U}_{\text{SVD}}^{\text{LSH}}$ execution is then $1 + 2 \times (2\eta^2 + 2\eta + 6) = 4\eta^2 + 4\eta + 13$, as reported in Table 8. As was the case with the Schwinger-boson formulation, the dominant CNOT cost is associated not with $\mathcal{U}_{\text{SVD}}^{\text{LSH}}$ but with implementing the diagonalized subterms, which will be discussed next.

The diagonalized subterms to be simulated are as given in Eq. (83):

$$x |1\rangle \langle 1|_{\mathbf{x}'} Z_{\mathbf{y}'} Z_{\mathbf{x}} (1 - |1\rangle \langle 1|_{\mathbf{y}} |\Lambda\rangle \langle \Lambda|_{\mathbf{p}}) (1 - |0\rangle \langle 0|_{\mathbf{x}} |\Lambda\rangle \langle \Lambda|_{\mathbf{p}}) \mathcal{D}^{\text{LSH}}(p, n_{\mathbf{x}}, n_{\mathbf{y}}).$$

The factors depending on $\{\mathbf{x}', \mathbf{y}, \mathbf{p}, \mathbf{x}\}$ can be treated as a single diagonal function of $\eta+3$ qubits to be decomposed to Pauli matrices. The only remaining qubit involved, \mathbf{y}' , can serve as the target for all parity evaluations of the terms in the $(\eta+3)$ -bit Pauli decomposition, and for the associated single-qubit R^Z gates. According to the Hamiltonian-cycle method described for the Schwinger-boson hopping propagators, implementing this diagonal Pauli decomposition will cost $2^{\eta+3}$ CNOT gates. This cost of a diagonalized LSH hopping subterm is recorded in Table 8.

The cumulative CNOT cost associated with a complete hopping term is obtained as four times the cost of a single $\mathcal{U}_{\text{SVD}}^{\text{LSH}}$ circuit, plus the individual costs of both diagonalized subterms. This works out to $16 \times 2^\eta + 16\eta^2 + 16\eta + 52$, also reported in Table 8. \square

The Pauli decomposition of \mathcal{D}^{LSH} can be truncated via both the small-angle and input-truncations methods described in the case of the Schwinger-boson propagators. The qualitative conclusion regarding the significant loss of accuracy by even minimal truncation remains the same with both methods, as demonstrated by the examples in Figs. 4 and 5. The implementation of the LSH diagonal function is seen to benefit more from both types of truncations than the Schwinger-model diagonal function. At severe truncations for the LSH functions, the error bounds become comparable for both input-truncation and small-angle truncation methods.

—Far term—

LSH mass-propagator subroutine	T gates	Workspace	Scratch space
$H_M^{\text{LSH}(j)}$ rotation	$\mathcal{C}_z(\epsilon)$	0	0
Full $e^{-itH_M^{\text{LSH}(r)}}$ circuit	$2\mathcal{C}_z(\epsilon)$	0	0

LSH electric-propagator subroutine	T gates	Workspace	Scratch space
(Un)compute N^L	$4\eta + 4$	$\eta - 1$	2
Copy	0	0	$\eta + 1$
$(\eta + 1) \times (\eta + 1)$ mult.	$8\eta^2 + 12\eta + 4$	$2\eta + 2$	$2\eta + 2$
$H_E^{\text{LSH}(1)}$ rotations	$(\eta + 1)\mathcal{C}_z(\epsilon)$	0	0
$H_E^{\text{LSH}(2)}$ rotations	$(2\eta + 2)\mathcal{C}_z(\epsilon)$	0	0
Full $e^{-itH_E^{\text{LSH}(r)}}$	$16\eta^2 + 32\eta + 16 + (3\eta + 3)\mathcal{C}_z(\epsilon)$	$2\eta + 2$	$3\eta + 5$

Table 9: Summary of the costs associated with the far-term simulation of the diagonal operators in the LSH Hamiltonian, as explained in the text.

Implementing mass propagators

Lemma 3.10. *The mass propagator $e^{-itH_M^{\text{LSH}(r)}}$ can be implemented, without approximation, using no ancilla qubits, two R^Z gates, and no additional T gates.*

Proof. The R^Z -gate count follows directly from Eqs. (77), which constitute a linear combination of two single-qubit Z operators. As discussed originally for the Schwinger-boson propagator, the far-term R^Z gates have a T-gate cost that scales with the desired synthesis error. The function $\mathcal{C}_z(\epsilon)$ has been introduced to refer to the number of T gates required to simulate an R^Z gate (of any input angle) within a spectral-norm error of ϵ , and the RUS method of Ref. [160] may be used for a T count scaling of $\log(1/\epsilon)$. The trivial cost of a mass propagator is recorded in Table 9. \square

Implementing electric propagators

In the far term, we use a splitting of H_E that is identical to the Schwinger-boson formulation:

$$H_E^{\text{LSH}(1)}(r) = \frac{1}{2}N^L(r), \quad (84a)$$

$$H_E^{\text{LSH}(2)}(r) = \frac{1}{4}(N^L(r))^2, \quad (84b)$$

where the only difference is the expression for N^L in terms of the LSH quantum numbers, i.e., $N^L(r) = n_\ell(r) + n_o(r)(1 - n_i(r))$.

Lemma 3.11. *The electric propagator $e^{-itH_E^{\text{LSH}(r)}}$ can be implemented, without approximation, using $5\eta + 6$ ancilla qubits, $3\eta + 3$ R^Z gates, and an additional $16\eta^2 + 32\eta + 16$ T gates.*

Proof. The procedure leading to this cost is essentially the same as in the Schwinger-boson formulation, with the exception of how N^L is computed. It consists of the following steps:

1. Compute N^L .
2. Compute $(N^L)^2$.
3. Realize a phase kickback via the registers containing N^L and $(N^L)^2$.
4. Uncompute $(N^L)^2$ and N^L .

In step (1), N^L can be computed as

$$|n_i\rangle |n_o\rangle |0\rangle |0\rangle |n_\ell\rangle |0\rangle^{\otimes(\eta-1)} \mapsto |n_i\rangle |n_o\rangle |n_o(1 - n_i)\rangle |N^L\rangle |0\rangle^{\otimes(\eta-1)},$$

where $|n_o(1 - n_i)\rangle$ is the output of a Toffoli gate and serves as the control bit to an incrementer on the $|0\rangle |n_\ell\rangle$ register. The Toffoli gate costs 4 T gates, while the controlled incrementer calls for 4η T gates and $\eta - 1$ workspace qubits (resulted from applying Lemma A.2 to an $(\eta + 2)$ -bit register). For (2), $(N^L)^2$ is computed

as $|N^L\rangle|0\rangle^{\otimes(5\eta+5)} \mapsto |N^L\rangle|N^L\rangle^{\otimes(4\eta+4)} \mapsto |N^L\rangle|N^L\rangle|(N^L)^2\rangle|0\rangle^{\otimes(2\eta+2)}$ by first copying N^L to an $(\eta+1)$ -bit register using CNOT gates and then multiplying the two copies of N^L . The multiplier costs $8\eta^2 + 12\eta + 4$ T gates and $2\eta + 2$ bits of workspace. In step (3), the $N^L(r)$ and $(N^L(r))^2$ terms of $H_E(r)$ are effectively simulated by applying single-qubit R^Z gates across the entire $|N^L\rangle$ and $|(N^L)^2\rangle$ registers (up to global phases that are dropped). Finally, in step (4), uncomputation involves reversing the gates of steps (2) and (1) and the associated costs are the same. Steps (2)-(4) are shown in Fig. 6. In total, the procedure outlined above involves $3\eta + 3$ R^Z gates, which can each be executed to the desired precision using the RUS method mentioned in previous sections [160]. The costs associated with all the subroutines are summarized in Table 9. The final, quoted costs are obtained by adding up the workspace and scratch-space sizes. \square

Implementing hopping propagators

Lemma 3.12. *Let $\eta = \log_2(\Lambda + 1)$ be the number of qubits per LSH bosonic register, $n > 0$ be the number of Newton’s method iterations, and $m > 0$ be a fixed binary arithmetic precision. Then $\Pi_{j=1}^2 e^{-itH_I^{\text{LSH}(j)}(r)}$ can be implemented within an additive spectral-norm error of*

$$2\sqrt{2}xt \left[2^n (\sqrt{2} - 1)^{2^n} + 2^{2-m} \left(\left(\frac{3}{2} \right)^n - 1 \right) \right] \quad (85)$$

using

$$2\eta n + 3mn + 12\eta + \max(\eta + 2, m) + 10m + 3n + 20 \quad (86)$$

ancilla qubits, $4(m+1)$ R^Z gates, and

$$320\eta mn + 128m^2n + 32\eta^2 + 32\eta m + 32\eta n + 352mn - 32n \max(2\eta + 2, 2m) + 224\eta + 32m - 16 \max(\eta + 1, m) - 16n + 64 \quad (87)$$

T gates.

Proof. To obtain this result, one must circuitize propagators associated with each of the two subterms of the hopping term. One subterm is simulated by applying an appropriate diagonalizing transformation $\mathcal{U}_{\text{SVD}}^{\text{LSH}}$, implementing the diagonalized subterm, and finally changing basis back via $\mathcal{U}_{\text{SVD}}^{\text{LSH}\dagger}$. In the LSH formulation, the diagonalized hopping subterm in Eq. (83) involves a logical control factor associated with the finite cutoff Λ , which will be computed as part of implementing the diagonalized subterm. The remainder of the diagonalized-subterm implementation follows the phase-kickback procedure used in the Schwinger-boson section. The full sequence for simulating a hopping subterm is shown schematically in Fig. 10(a).

Diagonalization.—The SVD transformation is as shown in Fig. 10(b): one CNOT plus two $C^2(\lambda^+)$ gates and a Hadamard gate. The $C^2(\lambda^+)$ operation, that is a doubly-controlled η -qubit incremter, can proceed as an uncontrolled $(\eta+2)$ -qubit incremter (followed by a CNOT and bit flips on the “control” qubits). According to Lemma A.2, for an input of size $\eta+2$, the cyclic incremter costs 4η T gates and $\eta-1$ ancilla qubits, where the ancilla qubits can be reused for all subsequent implementations. Doubling this, the cost of $\mathcal{U}_{\text{SVD}}^{\text{LSH}}$ is 8η T gates and $\eta-1$ workspace qubits, as reported in Table 10. Covering both subterms will further multiply the T count by four, for a total of 32η T gates per hopping term associated with diagonalizing transformations. The dominant cost of simulating a hopping term, however, comes from implementing the diagonalized subterms, which will be discussed next.

Cutoff control.—Before going into the phase-kickback procedure, the cutoff logical control (that was not explicitly needed in the Schwinger-boson formulation given the functional form of \mathcal{D}^{SB}) will be computed. The specific factor in Eq. (83) that is referred to here is $(1 - |1\rangle\langle 1|_y |\Lambda\rangle\langle \Lambda|_p)(1 - |0\rangle\langle 0|_x |\Lambda\rangle\langle \Lambda|_p)$, which amounts to an evaluation of $(1 - \delta_{n_y,1}\delta_{p,\Lambda})(1 - \delta_{n_x,0}\delta_{p,\Lambda})$. The cutoff logic evaluation is shown in circuit form in Fig. 10(c). The implementation requires four $C^{\eta+1}(X)$ gates and one Toffoli gate. According to Ref. [162] (Table 3), a $C^{\eta+1}(X)$ gate can be performed using 4η T gates, $\eta+1$ workspace qubits, and one output ancilla qubit. The outputs of the initial two $C^{\eta+1}(X)$ gates are combined into a single control using a Toffoli gate at an added cost of four T gates and one qubit for the final result. Overall, the cutoff logic circuit calls for $16\eta + 4$ T gates, $(\eta+1) + 2 = \eta + 3$ workspace qubits, and one qubit for the final output.

Phase kickback.—The remainder of the diagonalized hopping-subterm propagator consists of a phase-kickback algorithm that closely follows the one presented for the Schwinger-boson formulation. Its circuit implementation is depicted in Fig. 10(d). First, the value of the non-trivial diagonal bosonic operator \mathcal{D}^{LSH} is computed up to

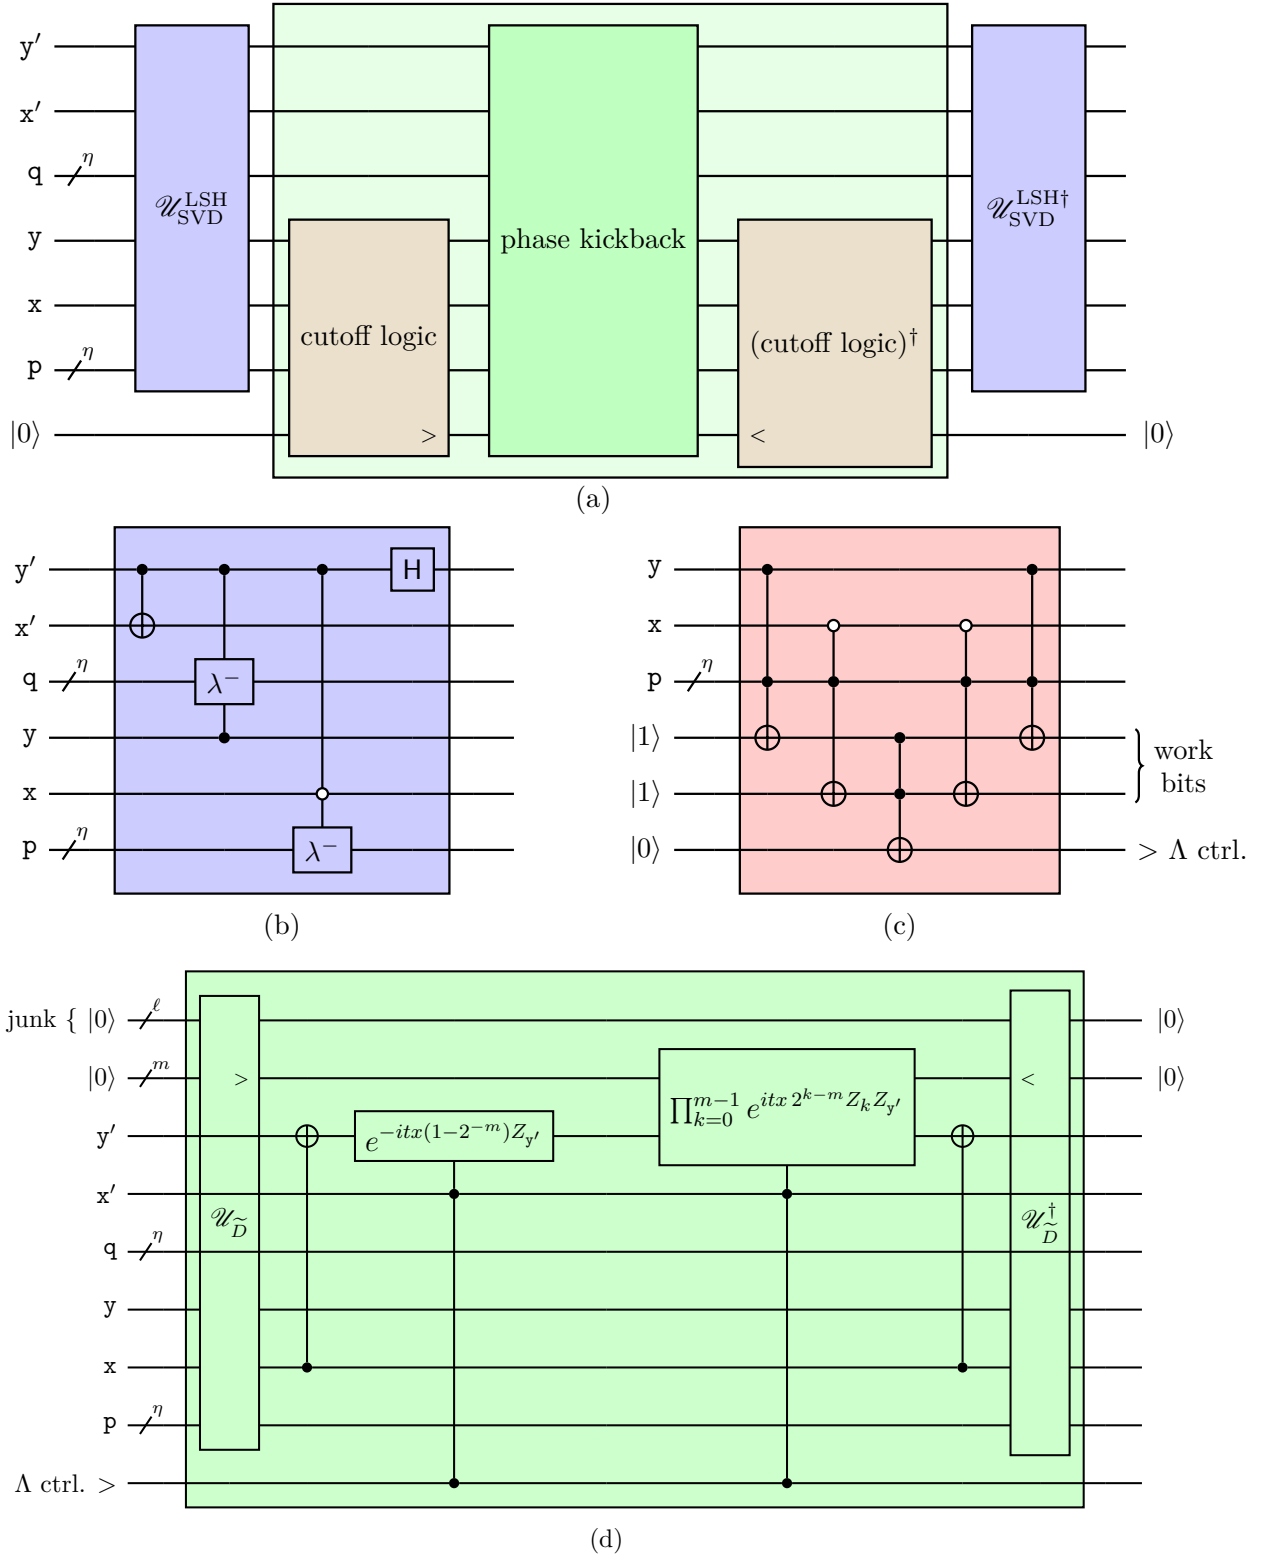


Figure 10: The circuit that implements the Schwinger-boson hopping propagator corresponding to each subterm in Eq. (79). (a) A high-level representation of the (far-term) circuit that realizes $H_I^{\text{LSH}(j)}(r)$. (b) The diagonalization circuit $\mathcal{U}_{\text{SVD}}^{\text{LSH}}$ for an LSH hopping subterm. This diagram applies to both the near-term and far-term algorithms. (c) The cutoff-control circuit. (d) The phase-kickback circuit. All circuits in (a-d) call for ancilla qubits that are not explicitly drawn but are discussed in the text and counted in the cost tables.

a known accuracy via a unitary $\mathcal{U}_D^{\text{LSH}}$ and the value is stored in an ancillary register of m bits, where $\tilde{\mathcal{D}}^{\text{LSH}}$ denotes a fixed-precision approximation to \mathcal{D}^{LSH} and m is the desired precision. We will construct $\mathcal{U}_D^{\text{LSH}}$ via Newton's method shortly and will discuss its cost as a function of the approximation error. Since the ancillary register holds the (approximate) value of \mathcal{D}^{LSH} in binary form, the diagonal operator

$$e^{-itx|1\rangle\langle 1|_{x'} Z_{y'} Z_x (1-|1\rangle\langle 1|_y |\Lambda\rangle\langle \Lambda|_p) (1-|0\rangle\langle 0|_x |\Lambda\rangle\langle \Lambda|_p)} \mathcal{D}^{\text{LSH}}(p, n_x, n_y)$$

can be straightforwardly implemented by R_Z gates on the corresponding qubits with the angles shown in the circuit. The ancillary register is set back to an all- $|0\rangle$ state by the inverse operations, and is used in subsequent implementations.

There are multiple ways of formulating the calculation of $\mathcal{D}^{\text{LSH}}(p, n, n') = \sqrt{\frac{p+1+n}{p+1+n'}}$. One way of formulating it goes as follows:

1. Increment p by one (after augmenting the register with a qubit).
2. Duplicate $p+1$ into an $(\eta+1)$ -bit register.
3. Use n and n' as controls to (separate) $C(\lambda^+)$ gates acting on the $|p+1\rangle$ registers, resulting in registers containing $p+1+n$ and $p+1+n'$.
4. Multiply the numerator and denominator into a $(2\eta+2)$ -bit register to obtain $(p+1+n)(p+1+n')$. For later convenience, this product is referred to by

$$g^{\text{LSH}}(p, n, n') \equiv (p+1+n)(p+1+n'). \quad (88)$$

5. Compute the inverse square root of $g^{\text{LSH}}(p, n, n')$.
6. Multiply $(p+1+n)$ with $g^{\text{LSH}}(p, n, n')^{-1/2}$.

In this protocol, the most complicated and resource-intensive step is (5): the inverse-square-root evaluation. The unitary that implements the steps above is called $\mathcal{U}_D^{\text{LSH}}$, and the evaluation only produces a fixed-precision approximation $\tilde{\mathcal{D}}^{\text{LSH}}(p, n, n') \approx \mathcal{D}^{\text{LSH}}(p, n, n')$. We will return shortly to address the details of using fixed-precision arithmetic and the associated errors. A circuit diagram for $\mathcal{U}_D^{\text{LSH}}$ is provided in Fig. 11.

The complete discussion of relative error as a function of Newton iterations n and fixed-bit precision m in calculating the inverse square root $g^{-1/2}$ carries through unchanged from the Schwinger-boson formulation to the LSH formulation. The only difference between the Schwinger-boson and LSH cases is the conversion from relative error bound to absolute error bound. To do this conversion in a way that is agnostic to the input state, one can multiply by the norm of the diagonal function. In Appendix C, it is shown that $\|\mathcal{D}^{\text{SB}}\| \leq 1$ and $\|\mathcal{D}^{\text{LSH}}\| = \sqrt{2}$, implying a factor of $\sqrt{2}$ needs to be included in the Newton's method approximation error for the LSH formulation.

To summarize the phase evaluation, there exists a quantum circuit that implements the unitary $\mathcal{U}_D^{\text{LSH}}$ defined as

$$|n\rangle |n'\rangle |p\rangle |0\rangle^{\otimes(1+m+\ell+\beta)} \xrightarrow{\mathcal{U}_D^{\text{LSH}}} |n\rangle |n'\rangle |p+1+n_q\rangle |\tilde{\mathcal{D}}^{\text{LSH}}(p, n, n')\rangle |\text{junk}\rangle |0\rangle^{\otimes\beta}. \quad (89)$$

The state $|\tilde{\mathcal{D}}^{\text{LSH}}(p, n, n')\rangle$ is an m -bit fixed-precision encoding of $\mathcal{D}^{\text{LSH}}(p, n, n')$ within absolute error $|\delta_{n,m}^{\text{Abs.}}| \leq \sqrt{2}[2^n(\sqrt{2}-1)^{2^n} + 2^{2-m}((3/2)^n - 1)]$ for n iterations of Newton's method. The state $|\text{junk}\rangle$, which holds intermediate values, occupies an ℓ -bit register. Furthermore, the circuit can be implemented using at most

$$80\eta mn + 32m^2n + 8\eta^2 + 8\eta m + 8\eta n + 88mn - 8n \max(2\eta + 2, 2m) + 32\eta - 4 \max(\eta + 1, m) + 8m - 4n + 8 \quad (90)$$

T-gates. Finally, the circuit consumes at most

$$\ell = 2\eta n + 3mn + 6\eta + \max(\eta + 2, m) + 3n + 7 \quad (91)$$

in output junk qubits and borrows

$$\beta = 6\eta + 9m + 9 \quad (92)$$

reusable workspace qubits. Lemma A.8 of Appendix A provides a derivation of these results.

Finally, with the cutoff-control factor and \mathcal{D}^{LSH} both evaluated, the diagonal operator

$$x|1\rangle\langle 1|_{x'} Z_{y'} Z_x (1-|1\rangle\langle 1|_y |\Lambda\rangle\langle \Lambda|_p) (1-|0\rangle\langle 0|_x |\Lambda\rangle\langle \Lambda|_p) \mathcal{D}^{\text{LSH}}(p, n_x, n_y)$$




Figure 11: A quantum circuit for computing the diagonal function \mathcal{D}^{LSH} defined in Eq. (80) with η -qubit bosonic register n_ℓ and one-qubit fermionic registers n_q and n'_q . Incrementers (+1), multiplication (Mult.), and initial-guess (Init.) subcircuits are described in Appendix A. The decomposition of a Newton-iteration step (Newt.) is depicted in Fig. 9. For each subcircuit, a line passing through the box indicates that the corresponding qubit register does not participate in the operations. The output of each unitary is marked by $>$ (or $<$ for the inverses). Workspace qubits are left implicit. $k \equiv 2\eta + 2$, n indexes the final Newton-iteration step, $y_0 = 2^{-s}$ with $s = \lceil \log_4((n_\ell + 1 + n_q)(n_\ell + 1 + n'_q)) + \frac{1}{2} \rceil$, and the state |Init. junk> is stored in $k + 1$ qubits, see Lemma A.5.

from Eq. (83) is effectively replaced by $x|1\rangle\langle 1|_x|1\rangle\langle 1|_{\Lambda\text{-ctrl}}Z_yZ_xN_{\text{aux}}$, where ‘ Λ -ctrl’ refers to the cutoff-control qubit and $N_{\text{aux}} = 2^{2-m} \sum_{k=0}^{m-1} 2^{k-1}(1 - Z_{\text{aux}kj})$ is the fixed-point number operator on the m -bit auxiliary register holding $\tilde{\mathcal{D}}^{\text{LSH}}$. With N_{aux} expanded out into m single-qubit Z operators plus a constant, the rotations can be split into $m + 1$ doubly-controlled R^Z rotations. To simplify the rotations, a scratch ancilla is introduced into which the result of logical AND of the two controls is input, adding a T-gate cost of eight (including the uncomputation). As remarked above, the far-term R^Z gates incur a T-gate cost associated with the tunable precision of the desired rotation angles and one can use the RUS construction of Ref. [160]. These rotations are recorded in Table 10, where a factor of two is included to account for now a single control on the R^Z gates.

Final tally.—Putting everything together, the simulation of the diagonalized hopping subterm costs

$$2\eta n + 3mn + 12\eta + \max(\eta + 2, m) + 10m + 3n + 18 \quad (93)$$

ancilla qubits (where m is added to account for the primary phase-kickback register), $2(m + 1)$ R^Z gates, and an additional

$$160\eta mn + 64m^2n + 16\eta^2 + 16\eta m + 16\eta n + 176mn - 16n \max(2\eta + 2, 2m) + 112\eta + 16m - 8 \max(\eta + 1, m) - 8n + 32 \quad (94)$$

T gates. Adding the costs of doing two structurally equivalent subterms, the full upper bound on implementing each Trotterized hopping propagator in the LSH formulation comes to twice the T-gate cost given in Eq. (94) and no more ancilla cost than what is stated in Eq. (93), as stated originally and summarized in Table 10. \square

3.3 Error-bound analysis and simulation cost

With a second-order Trotter-Suzuki formula $V_2(t)$, the Trotterization error scales cubically with the time duration of each Trotter step, improving upon the quadratic scaling of the first-order formula $V_1(t)$ at only twice the gate cost. Given the decompositions of the Hamiltonian in both the Schwinger-boson and the LSH bases

Routine	T gates	Workspace	Scratch space
$\mathcal{U}_{\text{SVD}}^{\text{LSH}}$ or $\mathcal{U}_{\text{SVD}}^{\text{LSH}\dagger}$	8η	$\eta - 1$	0
(Un)compute cutoff logic	$16\eta + 4$	$\eta + 3$	1
$\mathcal{U}_D^{\text{LSH}}$ or $\mathcal{U}_D^{\text{LSH}\dagger}$	$80\eta mn + 32m^2n + 8\eta^2 + 8\eta m + 8\eta n + 88mn - 8n \max(2\eta + 2, 2m) + 32\eta - 4 \max(\eta + 1, m) + 8m - 4n + 8$	$6\eta + 9m + 9$	$2\eta n + 3mn + 6\eta + \max(\eta + 2, m) + m + 3n + 7$
Diagonal rotations	$8 + 2(m + 1)\mathcal{C}_z(\epsilon)$	0	1
Overall $H_i^{\text{LSH}(j)}(r)$	$160\eta mn + 64m^2n + 16\eta^2 + 16\eta m + 16\eta n + 176mn - 16n \max(2\eta + 2, 2m) + 112\eta + 16m - 8 \max(\eta + 1, m) - 8n + 32 + 2(m + 1)\mathcal{C}_z(\epsilon)$	$6\eta + 9m + 9$	$2\eta n + 3mn + 6\eta + \max(\eta + 2, m) + m + 3n + 9$

Table 10: Summary of the costs associated with the far-term simulation of the off-diagonal operators in the LSH Hamiltonian, as explained in the text.

introduced in Sec. 3.2, and considering that in both formulations $H_E^{(j)}$ mutually commute, as do $H_M^{(j)}$, the second-order expansion of the time-evolution operator, $V_2(t)$, in the SU(2) theory can be formed in terms of the first-order expansion $V_1(t)$ as

$$V_1(t) = \prod_{r=0}^{L-2} \left[e^{-it(H_M(r)+H_E(r))} \prod_{j=1}^{\nu} e^{-itH_I^{(j)}(r)} \right] e^{-itH_M(L-1)}, \quad (95a)$$

$$V_2(t) = V_1(t/2)V_1(-t/2)^\dagger, \quad (95b)$$

where the product in Eq. (95a) is ordered left to right. Each exponential in the product can then be circuitized according to the near- or far-term strategies of Sec. 2.

Furthermore, the second-order product formula in Eq. (95b) can be bounded by the double-commutator relation [104]

$$\|V_2(t) - e^{-itH}\| \leq \frac{t^3}{24} \sum_{i=1}^{\Upsilon} \left\| \left[H_i, \left[H_i, \sum_{j=i+1}^{\Upsilon} H_j \right] \right] \right\| + \frac{t^3}{12} \sum_{i=1}^{\Upsilon} \left\| \left[\sum_{k=i+1}^{\Upsilon} H_k, \left[\sum_{j=i+1}^{\Upsilon} H_j, H_i \right] \right] \right\|, \quad (96)$$

where $H = \sum_{i=1}^{\Upsilon} H_i$ and $H_i \in \{H_M, H_E, H_I^{(j)}\}$ for $j = 1, \dots, \nu$. This bound is evaluated for the Schwinger-boson and the LSH formulations in Appendix C. The corresponding results will be used in the following to estimate the cost of simulation in each formulation, given the algorithms of Sec. 3.2 and the target error on the time-evolution operator.

In the fault-tolerant regime, in addition to the Trotterization error, there are additional known sources of error. The first of these is associated with the synthesis of arbitrary R^Z rotations, the accuracies of which are generally referred to as ϵ here. The second of these is the error due to finite precision in the diagonal phase-function evaluations. Specifically, this arises from truncating inverse-square-root evaluations to a finite number of steps n in Newton's method, as well as the truncation of the involved calculations to a fixed m -bit precision. The three sources of error will be summarized and combined in this section to derive the complete error bound in the far-term scenario.

3.3.1 Schwinger-boson formulation

Trotterization error.—In Appendix C.0.1, we apply Eq. (95b) to the complete Hamiltonian of the Schwinger-boson formulation to obtain the following result:

$$\|V(\theta) - e^{-i\theta H}\| \leq L\theta^3 \rho^{\text{SB}}(x, \Lambda, \mu), \quad (97a)$$

$$\rho^{\text{SB}}(x, \eta, \mu) \equiv \frac{1658x^3}{3} + 32\Lambda x^2 + \frac{218\mu x^2}{3} + 8x^2 + \frac{\Lambda^2 x}{3} + \frac{4\Lambda\mu x}{3} + \frac{\Lambda x}{6} + \frac{5\mu^2 x}{3} + \frac{\mu x}{3} + \frac{x}{48}, \quad (97b)$$

with $\Lambda = 2^\eta - 1$. For a total evolution time of T spread over s Trotter steps, this implies

$$\|V(T/s)^s - e^{-iT H}\| \leq \frac{LT^3}{s^2} \rho^{\text{SB}}(x, \eta, \mu). \quad (98)$$

When working with a fixed error budget Δ_{Trot} , one should take the number of Trotter steps s to be

$$s(x, \eta, \mu, L, T, \Delta_{\text{Trot}}) = \left\lceil \sqrt{\frac{LT^3}{\Delta_{\text{Trot}}} \rho^{\text{SB}}(x, \eta, \mu)} \right\rceil \quad (99)$$

in order to guarantee $\|V(T/s)^s - e^{-iTH}\| \leq \Delta_{\text{Trot}}$. Note that since dimensionless Hamiltonians have been used throughout, when considering the dimensionfull evolution time t , it should be converted to the scaled dimensionless time T via $t = 2xTa_s$, see discussions in Sec. 3.1.

Near-term simulation costs. In the near-term, the cost metric is identified as the number of CNOT gates due to their non-negligible error. Given some desired error Δ_{Trot} , one can use Eq. (99) to determine the minimal required number of Trotter steps necessary to bring the Trotterization error within the error budget. Then given the cost analysis of Sec. 3.2, the total number of CNOT gates required for the Schwinger-boson formulation is

$$2s(L-1)(16 \times 8^\eta + 67\eta^2 + 65\eta + 30) \quad (100)$$

Some representative values are tabulated for a range of modest simulation parameters in Table 16 of Appendix D. Furthermore, the CNOT-gate count as a function of η and evolution time are plotted in Fig. 12 for better visualization.

Finite-precision (Newton's method) error.—Turning to the simulation in the far-term scenario, in applying Newton's method to the evaluation of inverse-square-root functions, an error due to the finite number of steps, n , as well as the practical choice of using fixed m -bit precision occurred. As explained in the previous sections and in Appendix A, these truncations contribute an absolute additive error

$$x\theta \left[2^n (\sqrt{2} - 1)^{2^n} + 2^{2-m} \left(\left(\frac{3}{2} \right)^n - 1 \right) \right] \quad (101)$$

for an evolution by time θ of a single hopping subterm (cf. Eq. (63)). In the second-order Trotterization, one has $\theta \rightarrow T/(2s)$ and the total number of diagonal-function evaluations will be $2s\nu^{\text{SB}} = 16s$. The total error is, therefore,

$$\Delta_{\text{Newt.}} = x(L-1)T\nu^{\text{SB}} \left[2^n (\sqrt{2} - 1)^{2^n} + 2^{2-m} \left(\left(\frac{3}{2} \right)^n - 1 \right) \right]. \quad (102)$$

Given a fixed error budget in Newton's method truncation, Eq. (102) sets minimal sizes on n and m necessary to meet the error budget. However, the choices of n and m are not independent, since there are two correlated parameters but only one constraint. Therefore, a choice must be made for how to simultaneously choose n and m . To this end, we propose the following scheme: given that the vast majority of T gates are spent on implementing function evaluations, m and n are chosen such that the T count of $\mathcal{Q}_D^{\text{SB}}$ is minimized (subject to the constraint set by the error budget). The T count of $\mathcal{Q}_D^{\text{SB}}$ is as stated in Lemma A.7 and it is a function of η , n , and m . The advantage of this scheme is that it should optimize the overall T count of simulation. The drawback is that these implicit definitions do not lend themselves to closed-form expressions for $n(x, \eta, L, T, \Delta_{\text{Newt.}})$ and $m(x, \eta, L, T, \Delta_{\text{Newt.}})$. In practice, only a few (n, m) pairs have to be considered.

Synthesis error.—The total error due to synthesis of R^Z gates, $\Delta_{\text{synth.}}$, is the sum of individual synthesis errors ϵ of each rotation. Thus, the main quantity of interest is the number of R^Z gates associated with each type of subpropagator in the Trotter decomposition. These have been identified in the text of previous sections, and they are collected for convenience in Table 11. The total number of R^Z gates per second-order Trotter step is thus found to be

$$2[2L + (3\eta + 3)(L-1) + 8 \times 2(m+1)(L-1)] = -6\eta + 6\eta L + 32Lm + 42L - 32m - 38. \quad (103)$$

The total synthesis error is this value times s and ϵ :

$$\Delta_{\text{synth.}} = s\epsilon(-6\eta + 6\eta L + 32Lm + 42L - 32m - 38). \quad (104)$$

When there is a predefined error budget for rotation synthesis, Eq. (104) implies a maximum value for ϵ :

$$\epsilon(\eta, L, s, m, \Delta_{\text{synth.}}) = \frac{\Delta_{\text{synth.}}}{s(-6\eta + 6\eta L + 32Lm + 42L - 32m - 38)}. \quad (105)$$

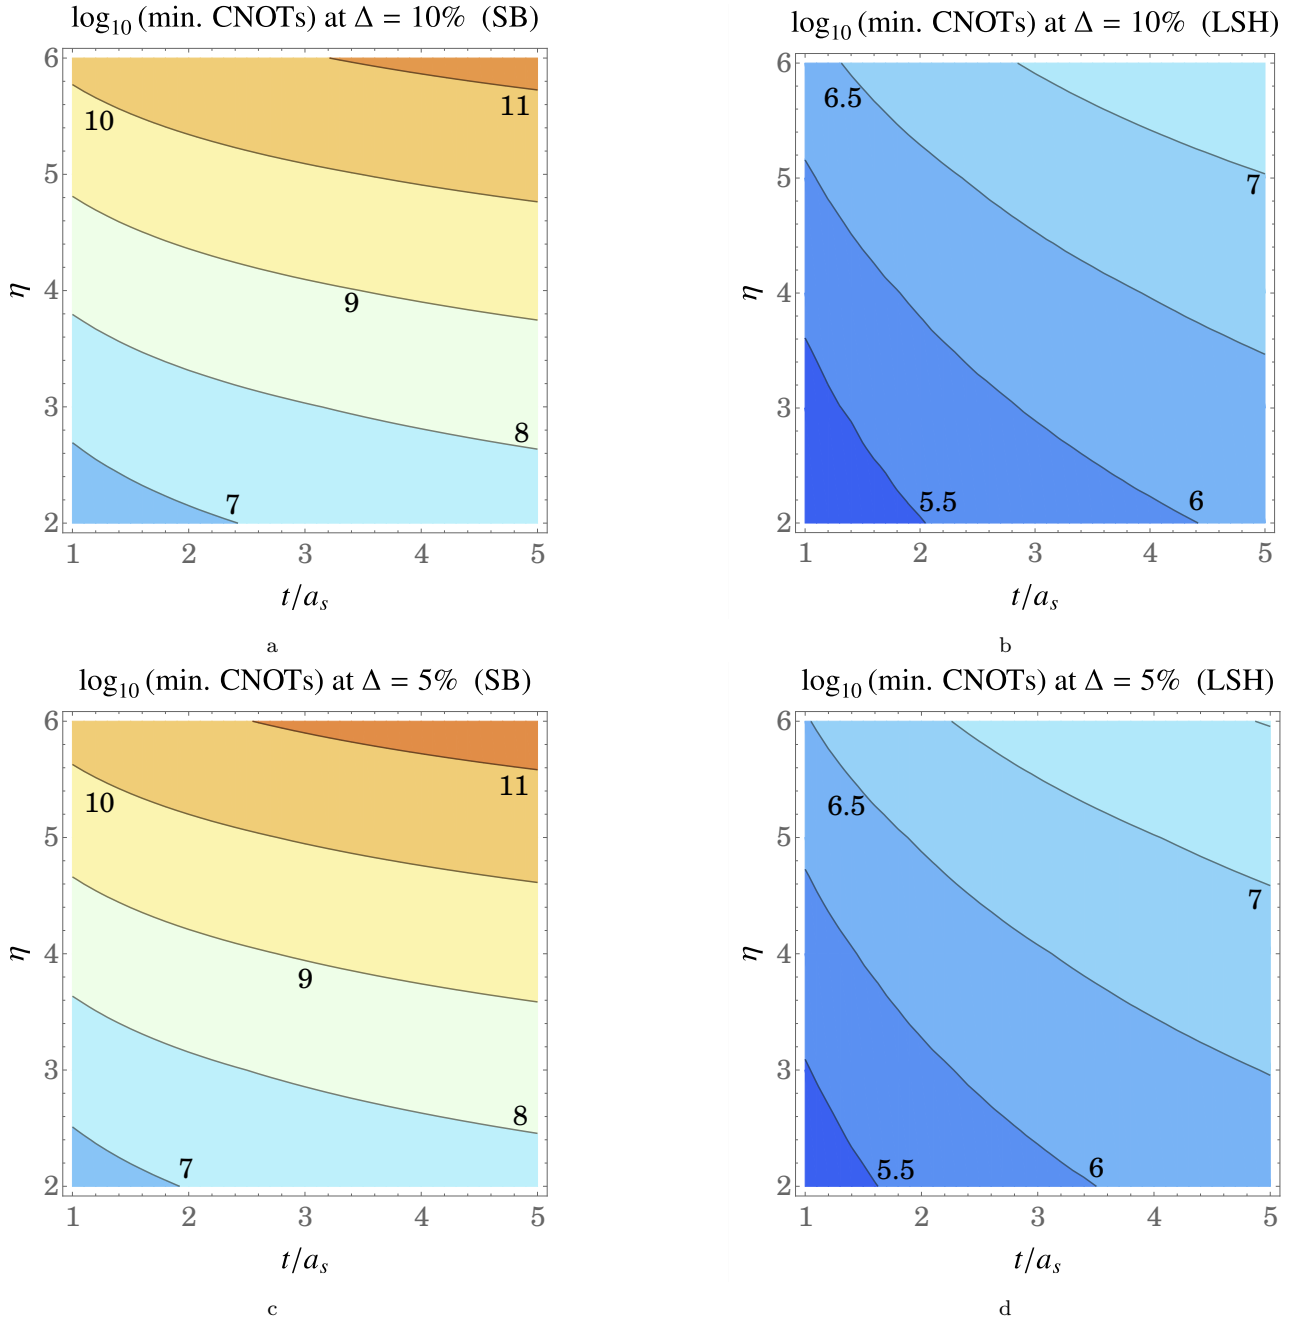


Figure 12: CNOT-gate costs at fixed $m/g = 1$. Other simulation parameters not explicitly shown are $x = 1$ and $L = 10$.

Far-term error-bounded simulation costs.—We are now ready to provide the full costs of the algorithms for simulating the $SU(2)$ LGT in the Schwinger-boson formulation as a function of model parameters and target accuracy. Here, x , $\Lambda = 2^\eta - 1$, μ , L , and T are taken as given parameters. Δ is then introduced to represent the complete error budget of the time-evolution operator, while $0 < \alpha_{\text{Trot.}} < 1$ represents the fraction allocated to Trotterization error and $0 < \alpha_{\text{Newt.}} < 1$ represents the fraction allocated to Newton-truncation errors. The fraction of Δ allocated for R^Z synthesis error is then $\alpha_{\text{synth.}} = 1 - \alpha_{\text{Trot.}} - \alpha_{\text{Newt.}}$ and must be positive. The truncation to $\Lambda < \infty$ is an additional known source of error that we do not quantify in this work.

With a target Trotterization error bound of $\alpha_{\text{Trot.}}\Delta$, the minimal number of Trotter steps s is derived as described above. Similarly, the target error bound on Newton’s method truncation error implies certain values of n and m according to the T-gate optimization scheme outlined above. Lastly, the R^Z gate accuracy ϵ is derivable given the error budget of $(1 - \alpha_{\text{Trot.}} - \alpha_{\text{Newt.}})\Delta$, once s and m are known.

Routine	R^Z gates
$e^{-i\theta H_M^{\text{SB}}(r)}$ or $e^{-i\theta H_M^{\text{LSH}}(r)}$	2
$e^{-i\theta H_E^{\text{SB}}(r)}$ or $e^{-i\theta H_E^{\text{LSH}}(r)}$	$3\eta + 3$
$e^{-i\theta H_I^{\text{SB}(j)}(r)}$ or $e^{-i\theta H_I^{\text{LSH}(j)}(r)}$	$2(m + 1)$
Full Trotter step (SB)	$-6\eta + 6\eta L + 32Lm + 42L - 32m - 38$
Full Trotter step (LSH)	$-6\eta + 6\eta L + 8Lm + 18L - 8m - 14$

Table 11: Summary of the number of the single-qubit R^Z gates involved in implementing the three types of (sub)propagators, and of the complete Trotter-step propagator.

The derived values of s , n , m , and ϵ are then inserted into the complete time-evolution cost formulas. The complete T-gate count is

$$2s \left[-64L \max(m, 2\eta) - 128Ln \max(2m, 4\eta + 2) + 64 \max(m, 2\eta) + 128n \max(2m, 4\eta + 2) - 784\eta^2 - 928\eta + 784\eta^2 L + 928\eta L + 512Lm^2 n + 256\eta Lm + 2560\eta Lmn + 1408Lmn + 256\eta Ln - 64Ln + 8L - 512m^2 n - 256\eta m - 2560\eta mn - 1408mn - 256\eta n + 64n - 8 + (-3\eta + 3\eta L + 16Lm + 21L - 16m - 19)C_z(\epsilon) \right]. \quad (106)$$

The ancilla-qubit count is just the ancilla-qubit count of a hopping propagator as it is the largest of all ancillary registers, and was found to be

$$\max(m, 2\eta + 2) + 27\eta + 3mn + 10m + 4\eta n + 3n + 15. \quad (107)$$

Some example values for the qubit and T-gate counts at given values of parameters are provided in Table 17 of Appendix D. Furthermore, the T-gate count as a function of x and L are plotted in Fig. 13 for better visualization of the resource requirement toward the continuum and bulk limits.

3.3.2 Loop-string-hadron formulation

Deriving the total error-bounded simulation costs in the near- and far-term scenarios for the LSH formulation follows closely that for the Schwinger-boson simulation. Without repeating the procedure, we only state the differences:

Trotterization error.—Instead of function ρ^{SB} , one uses ρ^{LSH} derived in Appendix 95b:

$$\rho^{\text{LSH}}(x, \eta, \mu) \equiv \frac{47\sqrt{2}x^3}{3} + 2\Lambda x^2 + \frac{25\mu x^2}{3} + 3x^2 + \frac{\Lambda^2 x}{24\sqrt{2}} + \frac{\Lambda\mu x}{3\sqrt{2}} + \frac{\Lambda x}{8\sqrt{2}} + \frac{5\mu^2 x}{6\sqrt{2}} + \frac{\mu x}{2\sqrt{2}} + \frac{3x}{32\sqrt{2}}, \quad (108)$$

with $\Lambda = 2^\eta - 1$.

Near-term simulation costs.—The total number of CNOT gates required is:

$$2s(L - 1) \left(16 \times 2^\eta + \frac{33}{2}\eta^2 + \frac{42}{2}\eta + 53 \right), \quad (109)$$

with s defined in Eq. (99) after substituting ρ^{SB} by ρ^{LSH} . Some representative values are tabulated for a range of modest simulation parameters in Table 16 of Appendix D. Furthermore, the CNOT-gate count as a function of η and evolution time are plotted in Fig. 12 for better visualization.

Finite-precision (Newton's method) error.—The total number of diagonal-function evaluations will be $2s\nu^{\text{LSH}} = 4s$, and one needs to replace x with $\sqrt{2}x$ in the expression for Δ_{Newt} . Eq. (102).

Synthesis error.—Given a different number of R^Z gates for the LSH formulation, the total synthesis error now is:

$$\Delta_{\text{synth.}} = s\epsilon(-6\eta + 6\eta L + 8Lm + 18L - 8m - 14). \quad (110)$$

This can be rearranged like Eq. (105) to give the maximum per-gate error $\epsilon(\eta, L, s, m, \Delta_{\text{synth.}})$ allowed within a given error budget.




Figure 13: T-gate costs at fixed $m/g = 1$. Other simulation parameters not explicitly shown are $\eta = 8$, $t/a_s = 1$, $\alpha_{\text{Trot}} = 90\%$, $\alpha_{\text{Newt.}} = 9\%$, and $\alpha_{\text{synth.}} = 1\%$.

Error-bounded simulation costs.—The complete time-evolution T-gate count in the LSH formulation is

$$\begin{aligned}
2s \left[-16L \max(m, \eta + 1) - 32Ln \max(2m, 2\eta + 2) + 16 \max(m, \eta + 1) + 32n \max(2m, 2\eta + 2) \right. \\
- 48\eta^2 - 256\eta + 48\eta^2 L + 256\eta L + 128Lm^2 n + 32\eta Lm + 320\eta Lmn + 352Lmn + 32Lm \\
+ 32\eta Ln - 16Ln + 80L - 128m^2 n - 32\eta m - 320\eta mn - 352mn - 32m - 32\eta n + 16n - 80 \\
\left. + (-3\eta + 3\eta L + 4Lm + 9L - 4m - 7)C_z(\epsilon) \right]. \quad (111)
\end{aligned}$$

The ancilla-qubit count is just the ancilla-qubit count of a hopping propagator, which was found to be

$$2\eta n + 3mn + 12\eta + \max(\eta + 2, m) + 10m + 3n + 18. \quad (112)$$

Some example values for the qubit and T-gate counts at given values of parameters are shown in Table 17

of Appendix D. Furthermore, the T-gate count as a function of x and L are plotted in Fig. 13 for better visualization of the resource requirement toward the continuum and bulk limits.

3.4 Discussion of the results

The cost analysis of this section resulted in the full simulation resource requirements of the SU(2) LGT in 1+1 D in two formulations. Here, we summarize the key features of the algorithms presented and the findings of this analysis:

◇ *Subdivision of hopping terms.*—One of the key insights of this work is the importance of carefully identifying subterms within a given hopping term that can be circuitized as a whole, the number of which we have referred to as ν . Firstly, the value of ν has a tangible impact on the gate costs because the total number of diagonal-function evaluations is proportional to ν . Therefore, any reductions in ν will carry over to the gate costs. Secondly, as shown in Appendix C, the second-order Trotter-error commutator bound is a cubic polynomial in the quantity νx , and thus a lower ν can dramatically tighten up the error-per-Trotter-step, especially in the weak-coupling (large- x) limit. In this regime, the minimal number of second-order Trotter steps (for a fixed error tolerance) scales as $\nu^{3/2}$, so reducing ν leads to a reduction in the number of required steps.

◇ *Symmetry considerations.*—Aside from the clear impact on gate costs, there is a connection between ν and local symmetry preservation. The limiting case of an algorithm with $\nu = 1$ is special in that all local symmetries would have to be preserved (because the hopping terms are gauge invariant cumulatively). As ν increases, the potential for individual subterms, hence the Trotterized hopping propagator, to violate symmetries also increases. Indeed, the LSH formulation naturally led to a $\nu = 2$ splitting that conserves all local symmetries, despite the Trotterized hopping propagator not being exact. In the Schwinger-boson formulation, it was natural to identify a splitting with $\nu = 8$, where each individual subterm can be seen to conserve $G_3(r)$ as well as the AGL, but not $G_1(r)$ and $G_2(r)$. This is an improvement compared with a naive $\nu = 64$ splitting (where all individual terms in the hopping Hamiltonian are separated).

It should be noted that the exact benefits afforded by conserving local symmetries are not well quantified. For example, while symmetry-breaking errors due to algorithms with bounded errors (such as Trotterization) can drive the evolution to outside the physical Hilbert space, these errors can be systematically eliminated (e.g., by increasing the number of Trotter steps). On the other hand, hardware noise can contribute to decoherence in the simulation, and so it appears that there may be a value in working with formulations such as LSH, and with simulation algorithms such as those presented here, that restrict the evolution to a substantially smaller Hilbert space.

◇ *Functional dependence on bosonic variables.*—Function-evaluation costs generally scale with the size and the number of arguments. In the near term, where we have considered naive Pauli decompositions, reducing the number of bosonic inputs called for by a function represents a potential for exponential savings. While the Schwinger-boson functions \mathcal{D}^{SB} depend on three bosonic quantum numbers (p , q , and p'), the LSH counterparts \mathcal{D}^{LSH} were shown to depend on only one bosonic variable (p), given the simplifications offered by gauge-invariant local DOFs and the AGL (see Appendix B). This led to relative CNOT-gate count per Trotter step of roughly $8^\eta + 4\eta^2$ (Schwinger boson) to $2^\eta + \eta^2$ (LSH).

In the far term, in addition to the number of bosonic inputs, generally the simplicity of the functional expressions is also important. In the LSH (Schwinger-boson) formulation, the diagonal function \mathcal{D}^{LSH} (\mathcal{D}^{SB}) could be expressed in terms of the ratio of two integers (products of integers), $p+1+n$ and $p+1+n'$ (pq and $(p+p')(p+p'+1)$). This ultimately requires the inverse square root of a product of two (four) integers that was called g^{LSH} (g^{SB}). Lemma A.4 confirms that the costly Newton iterations associated with the inverse square root scale linearly with the size of g , so the less bits needed for g , the better. Note that without the simplification of Appendix B, the LSH diagonal function would have ultimately involved four integers and likely had a comparable cost to the Schwinger-boson diagonal functions.

◇ *Function-evaluation choices.*—There are numerous possibilities to how one may obtain the irrational value $\sqrt{p/q}$ for integers p and q , but the approach that multiplies q with $1/\sqrt{pq}$ requires only a single approximate function evaluation: the inverse square root. It is expected that less function evaluations will generally lead to lower costs, although there are other factors to consider, such as the sizes of arguments and the rates of convergence for different functions. Additionally, the LSH diagonal functions may be better suited to other approximation methods since alternatively it can be written as

$$\mathcal{D}^{\text{LSH}}(p, n, n') = \sqrt{1 + \frac{n - n'}{p + 1 + n'}}, \quad (113)$$

where $|(n - n')/(p + 1 + n')| \leq 1$ provides a “small parameter”. Efficient schemes for evaluating the square root of one plus a small number in the far term may improve on the Newton’s method approach taken in this paper. This perspective could especially be helpful in the near-term, e.g., algorithms that allow for some arithmetic manipulations but exclude full-blown Newton’s method iterations that are costly.

- ◇ *Quantitative cost comparisons.*—According to the results of Sec. 3.3, the time evolution in the SU(2) LGT in 1+1 D for a range of parameters suitable for performing continuum- and bulk-limit extrapolations could be viable with $\sim 10^7$ (high-fidelity) CNOT gates. Over a wide range of realistic parameters, the LSH formulation reduces the minimum number of near-term Trotter steps by about 3-4 times. At low (high) cutoffs $\eta \leq 5$ ($\eta \geq 10$), the LSH formulation gives a $\sim 0.4 \times 4^n$ -fold ($\sim 4^n$ -fold) reduction in CNOT-gate count per Trotter step over a wide range of parameters. Far-term time evolution could be viable with $\sim 10^{12}$ T gates. Finally, over a wide range of parameters, the LSH formulation gives a ~ 20 -fold reduction in the T-gate count over the Schwinger-boson formulation.

4 Conclusions and outlook

Quantum simulation presents one of the most promising applications of digital quantum computers. As a result, research into algorithm and hardware technology that would enable scalable quantum simulation of physical systems continues to be a significant endeavor in the field of quantum information sciences and its intersection with other disciplines. With this motivation, we have provided a scalable algorithm that enables simulation of a class of physical Hamiltonians in a straightforward and systematic manner. The Hamiltonians studied are those involving operators that couple the quantum numbers of multiple (spin, bosonic, and/or fermionic) degrees of freedom. The algorithm works by diagonalizing the time-evolution operator in a chosen basis based on an algebraic approach rooted in singular-value decomposition. Strategies for simulating the diagonal functions were then presented. The algorithm is applied to a problem in nuclear and high-energy physics, namely simulating the non-Abelian SU(2) lattice gauge theory in 1+1 dimensions in two different formulations. The analysis of this work provides explicit circuit decompositions for the time-evolution operator in this theory using product formulas, and considers the near- and far-term scenarios that require different resource optimizations. The full error budget of the algorithms has been identified, and a comparative study of the two formulations are enabled. Such a study sets the standard for future algorithmic progress in gauge-theory simulations of relevance to the Standard Model of particle physics such as the SU(3) LGT in 3+1 D, and other physical models.

To conclude this work, we summarize some of our findings in relation to gauge-theory simulations, compare and contrast our approach with the existing developments, and remark on future directions and extensions:

- ◇ *Summary of pros and cons of the two formulations of the 1+1-D SU(2) theory.*—The LSH formulation offers significant advantages over the Schwinger-boson formulation, and likely other formulations, as concluded in Ref. [132] for classical Hamiltonian-simulation methods as well. LSH, being an alternative avenue to the same gauge-invariant dynamics contained by the other formulations, shares many features with Schwinger bosons, such as site-local DOFs, the decomposition of hopping terms into multiple structurally-identical and non-commuting subterms, and the need to deal with costly square-root functions. The important distinction is in simulating the hopping terms. In the near term, the minimal number of bosonic variables involved with each hopping operator makes the Pauli decomposition substantially cheaper. In the far term, the relative simplicity of the arguments to the square root means cheaper numerical evaluations. In either case, the ability to Trotterize a hopping term with $\nu^{\text{LSH}} = 2$ instead of $\nu^{\text{SB}} = 8$ subterms helps to reduce costly subroutines and make better use of the error budget. Lastly, for the 1+1-dimensional theory, LSH requires fewer qubits than either of the Schwinger-boson and the angular-momentum formulations, with a gain in simplicity of the interaction terms. Finally, the LSH Trotterization exactly conserves all local symmetries—a feature that has not been found for either the angular-momentum (e.g., the work of Ref. [69]) or Schwinger-boson (this work) formulation.
- ◇ *Comparison with another algorithm for the SU(2) theory.*—The formulations in this paper have been analyzed rather differently than how the pioneering work of Ref. [69] proposed simulating the Kogut-Susskind formulation in the standard angular-momentum basis. Our algorithms incorporate the steps necessary to eliminate mixing (“periodic wrapping”) between the upper and lower cutoff states—steps that reduce the errors in time evolution. Furthermore, the hopping-term simulation algorithm of Ref. [69] can be identified with a $\nu = 64$ splitting that increases the number of costly diagonal-function evaluations, enlarges the Trotter error, and breaks Gauss’s laws associated with $G_1(r)$ and $G_2(r)$ (and we suspect $G_3(r)$ as well, via wrap-around effects in the m^L and m^R quantum numbers).

While we have not applied our method to the angular-momentum formulation to allow a side-by-side comparison of resource requirements compared with Ref. [69], we remark that a simple reorganization of

the angular-momentum basis could be done to cast this formulation into a form that is essentially that of Schwinger bosons—but involving one less bosonic mode per link. With that reorganization, changing between the angular-momentum and the Schwinger-boson formulations is largely a matter of redefining quantum numbers and reshuffling where some subroutines are executed within circuits. Explicitly, instead of storing (J, m^L, m^R) (with m^L and m^R shifted to be non-negative), one could store $(J, J + m^L, J + m^R)$. This is slightly more space-efficient, and the Clebsch-Gordan coefficients are functions of $J \pm m^L$ and $J \pm m^R$ anyway ($J - m^L$ and $J - m^R$ would still be derived quantities). But the quantities $J + m^L$ and $J + m^R$ are identical to two of the four Schwinger-boson occupation numbers on the link. We suspect that the intrinsic operational differences between the two digitization schemes, when considered as a whole, are negligible because the dominating subroutines in each formulation will be so similar. The only concrete drawback to the Schwinger-boson formulation is the greater number of logical qubits than the angular-momentum basis (reorganized or not).

- ◇ *On the generalization of the algorithm to higher dimensions.*—In two or three space dimensions, one must simulate magnetic interactions, conventionally expressed as trace of four link operators around a square plaquette. Let us consider the Schwinger-boson formulation to demonstrate the generalization of the algorithms of this work to simulating plaquette operators, and further remark on how it compares with the algorithm of Ref. [69]. Generalization to higher dimensions for the LSH formulation can be worked out similarly.

Within the Schwinger-boson formulation, the plaquette operator expands out to 2^7 distinct terms, which can be understood as follows. A factor 2^4 comes from the 4 choices of indices a, b, c, d in the product of four link operators, $U_{ab}U_{bc}U_{cd}U_{da}$. Another factor of 2^4 comes from the choice of $J \rightarrow J \pm 1/2$ on each of the four links. Finally, a factor of $1/2$ is required to avoid double-counting terms that are related by Hermitian conjugation. Now, each of the 2^7 terms counted above involves a product of eight harmonic-oscillator ladder operators acting on distinct modes around the plaquette. The SVD algorithm of this work requires splitting only one of these ladder operators to operators acting on disjoint even and odd Hilbert spaces of the corresponding quantum number. The diagonalization procedure will follow straightforwardly after that, as demonstrated in the example of a plaquette interaction in the U(1) case in Sec. 2.1. Compared to the hopping interactions, therefore, the plaquette interactions not only involve a larger number of diagonal phases to be evaluated, but also the phases are more complex functions and depend on occupation numbers of twelve modes (three per link).

Finally, we remark that the algorithm of Ref. [69], in contrast, splits each of the eight ladder operators to even and odd parts, hence requiring 2^8 diagonal phases to be evaluated for each and every one of the 2^7 terms. This, and the increased Trotter error given such a dramatic splitting to non-commuting terms, contribute to a substantial increase in the computational cost compared to the method of this work. Further details are left to future work [163].

- ◇ *Classical pre-processing for near-term benefit.*—Research will need to continue to offer more efficient function-evaluation schemes, perhaps with the use of hybrid classical-quantum routines, to make the simulation of non-Abelian gauge theories more suitable for near-term quantum computing. Truncating Pauli decompositions by their coefficient size appears to not be an efficient near-term cost-reduction scheme for the functions encountered in this work. This is evident from the slow reduction in the spectral-norm error as more diagonal Pauli operators are kept until almost all the operators are being simulated. The same slow convergence was observed in the alternate scheme that truncates the size of the input registers into the diagonal functions before a complete Pauli decomposition is performed. There may exist more sophisticated truncation schemes that converge faster than ordering the Pauli operators by decreasing coefficient size, and those need to be explored in future work.

Another strategy that avoids calculating the non-trivial non-Abelian functions is to instead hard-code in the circuits the value of the Clebsch-Gordan coefficients associated with all allowed transitions. This approach requires as many controlled operations as there are for possibilities of such transitions, and the amplitude for those transitions is encoded in the rotation angles that depend on the corresponding Clebsch-Gordan coefficients. This approach is applied to the case of the pure SU(3) LGT in 2+1 D in Ref. [68]. Future work should determine a cost comparison between this and other near-term approaches, such as those outlined here for the similar theories.

- ◇ *Comparison to U(1) and remarks on simulation in the group-element basis.*—Compared to the Schwinger model, simulating the SU(2) model is substantially more complicated due to the properties of the non-Abelian link operator. In the electric basis, whose truncation and digitization are better understood, one must deal with Clebsch-Gordan coefficients, in one guise or another, and their irrational dependence on electric quantum numbers is a major departure from the trivial ones that characterize all non-vanishing

matrix elements of the $U(1)$ link operator. In either the near- or far-term scenario, the vast majority of gate costs are confirmed to result from the handling of the $SU(2)$ link operator in the hopping terms. In principle, one might try avoiding the Clebsch-Gordan coefficients by working in the group-element basis, but truncation schemes are less developed when the gauge group is infinite as with the $SU(2)$ theory, see e.g., Refs. [51, 164, 165] for recent studies. General algorithms exist for simulating gauge theories in the group-element basis [50] but more work is needed to augment those studies with a comparable level of algorithmic detail as the present work for the $SU(2)$ theory. Furthermore, parametrization of the $SU(2)$ group manifold could require function evaluations that may still ultimately dominate the gate cost. Finally, the electric Hamiltonian involves more complicated operator structure in the group-element basis, and will be more costly to implement than in the electric-field (irrep) basis. One could change the basis between group-element and irrep bases to only require implementing diagonal operators, nonetheless a discrete non-Abelian quantum Fourier transform for the digitized truncated gauge group must first be developed.

- ◇ *Remarks on post-Trotterization schemes.*—While this work considered product-formula algorithms for implementing the time evolution, there are two other major classes of simulation algorithms based on linear combinations of unitaries (LCU) and quantum signal processing (QSP). An in-depth comparison of quantum-resource requirements for simulating the Heisenberg model using each of these quantum algorithms in Ref. [166] revealed that product-formula algorithms, according to their “true performance”, will likely outperform LCU- and QSP-based algorithms in simulating spin Hamiltonians. The “true performance” was obtained by extrapolating the exact performance deduced from small system sizes, and so may not offer a rigorous performance guarantee. This performance is also considered a “worst case” in the sense that it uses no assumptions about the input state. Assumptions about the input state and symmetries are indeed shown to improve the error bounds [139, 141, 167, 168]. The toolbox one would need to access the “true performance” of product-formula simulations (whether rigorously or heuristically) via incorporating all the physically-motivated assumptions needs to be developed for all interesting problems including gauge theories.

The resource requirements of LCU- and QSP-based methods scale with the norm of the Hamiltonian, which is a disadvantage in simulating LGTs that exhibit unbounded Hamiltonians, requiring significant resources towards the continuum limit. One generic advantage of product formulas is that their error depends on the commutation among Hamiltonian terms as opposed to their norm, which is bounded by locality of the Hamiltonian, as was shown in this paper for the case of the $SU(2)$ LGT and in previous work [54, 69]. Hybrid approaches that combine two or more simulation algorithms to achieve better performance have been recently proposed [112]. Another desirable feature of product-formula-based algorithms is that they are simply an approximation to the exact time evolution, and will likely be of broader applicability in post-classical but pre-fault-tolerant era of quantum simulation where analog or hybrid digital-analog simulators will lead the way, see e.g., Ref. [64] for an example of a hybrid approach in the context of LGTs.

- ◇ *Remarks on systematic uncertainties and the choice of algorithms’ error tolerance.*—The error budget identified in this study is associated with the approximations made in the time digitization (Trotterization), function evaluation, and gate synthesis, but the value of total error tolerance in the select examples shown was chosen somewhat arbitrarily. Nonetheless, the value of the error tolerance cannot be chosen without regard for other systematic uncertainties in the simulation. For example, as one considers reproducing physical values in future quantum simulations of gauge theories of relevance to nature, one needs to account for uncertainties arising from finite discretization of space, finite extent of physical volume, finite truncation of bosonic fields, and other systematics associated with choices of input parameters that may deviate from those in nature. As a result, choosing an ultra-high algorithmic accuracy, which leads to substantial increase in the cost, may not be necessary given that other systematics will dominate the error. As a result, one ultimately needs a holistic approach to the resource-requirement analysis of LGTs, in which all sources of systematic uncertainties are taken into account as one takes the continuum infinite-volume limits, see e.g., Refs. [51, 68, 132, 164, 165, 169–171] for initial discussions on quantifying and controlling such systematics in the context of quantum simulation of lattice field theories.

Acknowledgments

We acknowledge fruitful discussions with Andrew Childs, Martin Savage, and Nathan Wiebe. We further acknowledge valuable discussions on a range of topics at two Quantum Simulation of Strong Interactions (QuaSI) workshops held virtually at the InQubator for Quantum Simulation (IQUS) at the University of Washington, Seattle in Spring and Summer of 2021, which solidified our vision for this work.

ZD, AFS, and JRS were supported by the U.S. Department of Energy’s Office of Science Early Career Award,

under award DE-SC0020271, for theoretical developments for simulating lattice gauge theories on quantum computers. ZD and JRS were further supported by the U.S. Department of Energy's Office of Science, Office of Advanced Scientific Computing Research, Accelerated Research in Quantum Computing program award DE-SC0020312, for algorithmic developments for quantum simulation. AFS was further supported by the Lanczos Fellowship of the University of Maryland and the National Institute for Standard and Technology's Joint Center for Quantum Information and Computer Science (QuICS), and National Science Foundation's Graduate Research Fellowship Program (GRFP). The IQUS workshops were supported by the U.S. Department of Energy's Office of Science, Office of Nuclear Physics, InQubator for Quantum Simulation, under award DE-SC0020970.

A Arithmetic algorithmic routines

This appendix contains explicit circuit constructions for various arithmetic routines used throughout this paper. While some of these routines are known, some are application or extension of the related existing results. Multiplication and addition are among the primary operations in the circuits presented in the main text. There are many algorithms for these operations, the choice of which depends on the goal of analysis, e.g., minimizing T-gate or ancilla-qubit count. We choose algorithms with low T-gate counts and employ them in subsequent arithmetic evaluations. These routines are described in the following lemmas, along with a near-term implementation of incrementer that is ancilla free. The rest of the appendix presents a step-by-step guide to evaluating the Schwinger-boson and the LSH diagonal phases via Newton's method.

Lemma A.1 (Ancilla-free incrementer). *There exists a quantum circuit that increments by one an η -qubit integer in the computational basis with a CNOT-gate cost $2\eta(\eta - 1)$ and without using any ancilla qubit.*

Proof. A straightforward inspection of the circuit shown in Fig. 14 shows that using a QFT circuit and its inverse, plus a number of Z rotations, an η -qubit integer can be incremented by one. Since each QFT (and its inverse) on an η -qubit register can be implemented with $\eta(\eta - 1)$ CNOT gates [172], the CNOT-gate count of the incrementer is $2\eta(\eta - 1)$. \square

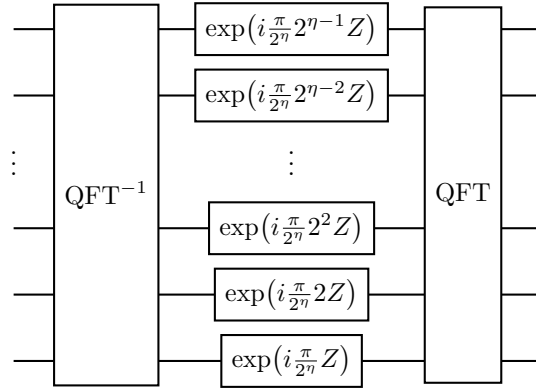


Figure 14: The circuit implementation of the ancilla-free incrementer on an η -qubit register based on the known quantum Fourier transform (QFT) circuit [54].

Lemma A.2 (Incrementer by a fixed integer in the fault-tolerant model). *Assume $|x\rangle$ is a computational basis vector encoded in a register of η qubits, i.e., $x \in \{0, \dots, 2^\eta - 1\}$, and y is a fixed integer such that $y \in \{0, \dots, 2^\eta - 1\}$. There exists a quantum circuit that maps $|x\rangle \rightarrow |x + y \bmod 2^\eta\rangle$ with a T-gate cost of $4\eta - 8$ and an ancilla-gate cost of $\eta - 3$. Additionally, an extra ancilla qubit may be needed for the execution of all the Toffoli gates.*

Proof. A circuit description for incrementation by a fixed integer can be obtained via slight modifications of the addition circuit presented in Ref. [173]. A proof is provided below for consistency of the presentation. For examples, the reader can refer to Fig. 15.

If x is the number being incremented, the output of the incrementer circuit should obtain in the k th-significant bit

$$x_k \oplus y_k \oplus \text{carry}_k, \quad (114)$$

where $\text{carry}_k = x_{k-1}y_{k-1} \oplus (x_{k-1} \oplus y_{k-1})\text{carry}_{k-1}$, and $\text{carry}_0 = 0$. To construct such a circuit, one first computes each carry_k for $k = 2, \dots, \eta - 2$ into an ancillary qubit. Note that $\text{carry}_0 = 0$, and $\text{carry}_1 = x_0y_0$ which is either x_0 or zero, so carry_0 and carry_1 do not need to be computed. Furthermore, the most significant bit of the output can be evaluated from $\text{carry}_{\eta-2}$ and $x_{\eta-2}$, thus $\text{carry}_{\eta-1}$ does not need to be computed and stored.

The value of carry_k for $k \geq 2$ can be computed as following. If $y_{k-1} = 0$, using a logical AND operation, the value of $x_{k-1}\text{carry}_{k-1}$, which is equal to carry_k in this case, is outputted to the ancilla qubit assigned to carry_k . If $y_{k-1} = 1$, first $\bar{x}_{k-1}\text{carry}_{k-1}$, with \bar{x}_{k-1} being the binary negation of x_{k-1} , is outputted to the ancilla qubit, then x_{k-1} is added to its value using a CNOT gate, obtaining carry_k . Computing all carry_k sequentially, starting from $k = 2$ up to $k = \eta - 2$, requires $\eta - 3$ logical AND computations, $\eta - 3$ logical AND uncomputations, and $\eta - 3$ ancilla qubits.




Figure 15: The circuit implementation of the incrementer by one (by three when including the gates in gray) on a five-qubit register in the fault-tolerant scenario is depicted in the left, with x_0, \dots, x_5 representing the binary representation of an integer from the least to the most significant digit. The wire corner denotes a logical AND, that is equivalent to a Toffoli gate with the ancilla qubit initialized in $|0\rangle$ as shown in the left. The uncomputation of logical AND proceeds by measuring the ancilla qubit as shown, hence requiring no T gates. The circuits are adopted from Ref. [173] upon necessary modifications.

Before the value of the carry qubits are reset by logical AND uncomputation, they are used to calculate the k th bit of the incremented output according to Eq. (114). This requires a CONT gate involving carry $_k$ qubit and the qubit holding either x_k or its binary negation depending on the value of y_k . This outputs in the qubits originally assigned to x_k for $k = 0, \dots, \eta - 2$ the corresponding incremented values in the end. For the most significant bit associated with $k = \eta - 1$, one does not need to store carry $_{\eta-1}$ in a designated ancilla qubit since $x_{\eta-1} \oplus y_{\eta-1} \oplus \text{carry}_{\eta-1}$ can be evaluated using a Toffoli gate controlled upon the values of $x_{\eta-2}$ (or its binary negation) and carry $_{\eta-2}$ (plus an additional CNOT gate if needed). Examples of incrementing a five-qubit integer by one and three are shown in the left-hand side of Fig. 15.¹⁷

The T-gate cost the circuit can be obtained by noting that as proposed by Gidney [173], uncomputing logical ANDs can proceed via a T-gate-free scheme that relies on the measurement of the ancilla qubits. Computing each local AND amounts to performing a Toffoli gate, and the Toffoli-gate compilation of Ref. [174] requires 4 T gates and a single ancilla qubit that can be reused. Adding the T-gate and ancilla-qubit counts gives the total cost stated in the Lemma. \square

Lemma A.3 (Multiplication of binary numbers in the fault-tolerant model). *Let a and b be integers such that $0 \leq a < 2^m$, and $0 \leq b < 2^n$. Then there exists a quantum circuit that performs*

$$|a\rangle |b\rangle |0\rangle^{\otimes(m+n+\gamma)} \rightarrow |a\rangle |b\rangle |ab\rangle |0\rangle^{\otimes\gamma} \quad (115)$$

(where $0 \leq a < 2^{m+n}$), using at most $\gamma = 2n$ ancillary qubits and $8mn - 4n$ T gates. Additionally, an extra ancilla qubit may be needed for the execution of all the Toffoli gates.

Proof. The basic idea is to compute ab using ‘schoolbook multiplication’ in the manner suggested by the following identity:

$$ab = \left(\sum_{i=0}^{m-1} 2^i a_i \right) b = a_0 \cdot b + 2a_1 \cdot b + \dots + 2^{m-1} a_{m-1} \cdot b.$$

The above corresponds to $m - 1$ applications of an n -bit addition subroutine.

For addition, one may use the T-gate-count-improved variation of the in-place, ripple-carry adder [175] as proposed by Gidney in Ref. [173]. The adder performs

$$|0\rangle |c\rangle |d\rangle |0\rangle^{\otimes\eta} \rightarrow |c+d\rangle |d\rangle |0\rangle^{\otimes\eta}, \quad (116)$$

where c and d are η -bit numbers, at a cost of 4η T gates and η workspace ancilla qubits. Note that since Gidney’s adder performs modular addition, we have padded $|c\rangle$ with a zero register from the left to effectively obtain a regular addition by performing an $(\eta + 1)$ -bit modular addition.

¹⁷This addition circuit and its η -qubit generalization performs modular addition. For the regular addition, one simply pads $|x\rangle$ with an extra qubit set in zero in the left, and performs the $\eta + 1$ modular fixed-number addition.

The $m + n$ bits which will ultimately contain $|ab\rangle$ will be referred to as the ‘output register.’ The algorithm is initiated by copying b to the n least significant bits of the output register, conditioned on the value of a_0 :

$$|a\rangle |b\rangle |0\rangle^{\otimes(m+n+\gamma)} \rightarrow |a\rangle |b\rangle (|0\rangle^{\otimes m} |a_0 \cdot b\rangle) |0\rangle^{\otimes \gamma}. \quad (117)$$

Conditionally copying an η -qubit register to another η -qubit register calls for η Toffoli gates.

Schoolbook multiplication then proceeds with $m - 1$ rounds of the following procedure, iterating j from 1 to $m - 1$:

1. Copy b to an auxiliary register of size n , conditioned on the value of a_j .
2. Use the n -bit, in-place adder to add the value of the ancillary register from Step 1 to the leading n bits of the currently-computed output string according to

$$\begin{aligned} & |a\rangle |b\rangle |0\rangle^{\otimes(m-j)} \left| \sum_{k=0}^{j-1} (a_k \cdot b) 2^k \right\rangle |a_j \cdot b\rangle |0\rangle^{\otimes \gamma - n} \\ & \rightarrow |a\rangle |b\rangle |0\rangle^{\otimes(m-j-1)} \left| \sum_{k=0}^j (a_k \cdot b) 2^k \right\rangle |a_j \cdot b\rangle |0\rangle^{\otimes \gamma - n}. \end{aligned} \quad (118)$$

3. Undo Step 1.

After round $j = m - 1$, the output register will indeed hold the value ab .

The ancillary qubit count of $\gamma = 2n$ is obtained as follows: Step (1) calls for n ancillary qubits, which temporarily hold logical ANDs until their uncomputation in Step (3). The adder of Step (2) calls for n ancilla qubits, which are also reset for use in the next iteration. Thus, $n + n = 2n$ ancillary qubits suffice to execute the complete multiplication algorithm.

The T-gate count is a combination of $4n$ for the Toffoli gates in the initialization operation of Eq. (117), plus $(m - 1) \times 4n$ for the Toffoli gates that implement all applications of Step 1 (noting that uncomputing temporary logical ANDs in Step 3 does not require any T gates according to Ref. [173]), plus $(m - 1) \times 4n$ for all applications of the addition circuit. This results in the total T count of $8mn - 4n$, as stated. \square

Lemma A.4 (Newton-iteration cost for an approximate evaluation of $1/\sqrt{g}$ in the fault-tolerant model). *Let g be any k -bit positive binary integer and y_n be a positive w -bit binary number, generated by the recurrence relation*

$$y_{n+1} = \frac{y_n(3 - y_n^2 g)}{2}, \quad (119)$$

with $\frac{|y_0 - 1/\sqrt{g}|}{1/\sqrt{g}} < 1$. The circuit depicted in Fig. 9, which implements

$$|g\rangle |y_n\rangle |0\rangle^{\otimes(3w+k+1+\beta)} \rightarrow |g\rangle |y_n\rangle |y_{n+1}\rangle |0\rangle^{\otimes \beta}, \quad (120)$$

may be constructed using

$$32w^2 + 40kw + 4k + 8w - 8 \max(k, 2w) - 12$$

T gates and

$$\beta = 9w + 3k + 3$$

auxiliary qubits. Additionally, an extra ancilla qubit may be needed for the execution of all the Toffoli gates.

Proof. With the multiplication cost obtained in Lemma A.3, we only need to cost the other three operations appearing in Fig. 9: Copy, $(3 - \#)$, and $\div 2$.

Copy only needs CNOT gates and no T gates. Similarly, the operation $\div 2$ can simply be done via SWAP gates and one ancilla qubit, as the numbers are in binary, which require no T gates. Alternatively, one just relabels the qubits at no computational cost, which is what is done here.

The operation $(3 - \#)$ is slightly more involved. Recognize that the input $\#$ is always such that $0 < \# < 3$ (and generally non-integer), so we will never have to work with signed arithmetic. This operation can be done using two’s-complement subtraction. As $\#$ is stored in bits labeled from the 0th binary place to the $-(2w + k - 1)$ th binary place, but 3 ranges over the 1st and 0th place, we must pad a single qubit on the register containing $\#$ in the 2nd place in order to use Lemma A.2. The Lemma A.2 is used to add $3 + 2^{-(2w+k-1)}$ to the one’s complement of $\#$ (where both numbers are interpreted as scaled by $2^{(2w+k-1)}$), which gives us $(3 - \#)$ in $2w + k + 1$ qubits that feed into the middle multiplication algorithm in Fig. 9. The cost of this operation is $8w + 4k - 4$ T gates and $2w + k - 2$ ancilla qubits, plus a single ancilla qubit that may be needed to execute the

Routine	T gates	Work bits	Scratch bits
Copy	0	0	w
y_n^2 : w -bit squaring	$8w^2 - 4w$	$2w$	$2w$
y_n^2g : $(2w \times k)$ -bit mult.	$16kw - 4 \max(k, 2w)$	$2 \max(k, 2w)$	$2w + k$
$3 - y_n^2g$ operation	$8w + 4k - 4$	$2w + k - 2$	1
$y_n(3 - y_n^2g)$: $w \times (2w + k + 1)$ -bit mult.	$16w^2 + 8kw - 4k - 4$	$4w + 2k + 2$	0
Division by 2	0	0	0
Overall circuit	$32w^2 + 40kw + 4k + 8w - 8 \max(k, 2w) - 12$	$4w + 2k + 2$	$5w + k + 1$

Table 12: Costs associated with an iteration of Newton’s method, i.e., the circuit depicted in Fig. 9 of the main text. Note that the overall T cost of the circuit accounts for two implementations of each of the w -bit squaring, $(2w \times k)$ -bit multiplication, and the $3 - \#$ operation.

Toffoli gates. Note that since the overflow bit is neglected in two’s-complement subtraction, the use of modular addition as in Lemma A.2 is acceptable.

Let us now determine the ancilla-qubit cost of each iteration of Newton’s method. The dominant cost is either associated with evaluating $(3 - \#)$ as discussed above or from the multiplication circuit which multiplies the numbers with the most bits. All other subroutines may use these workspace qubits. The ancilla cost of multiplications according to Lemma A.3 is $2w$ for the y_n^2 operation, $2 \max(k, 2w)$ for the y_n^2g multiplication, and $4w + 2k + 2$ for the $y_n(3 - y_n^2g)$ multiplication. This must be compared with the ancilla-qubit cost of the $(3 - \#)$ operation, that is $2w + k - 2$. So the number of needed ancilla qubits is $4w + 2k + 2$. Finally, one needs to add to this the number of ancilla qubits in the scratch space needed to store intermediate values, see Fig. 9 and Table 12, which adds to $5w + k + 1$. Combining these values gives the value of β as state above. Note that the scratch space is uncomputed in the end of each iteration and can be reused in other steps, hence the scratch space is counted toward the ancilla cost.

Finally, the T-gate cost is that of the five multiplications and two $(3 - \#)$ operations. Each y_n^2 multiplication costs $8w^2 - 4w$ T gates, each y_n^2g multiplication costs $16wk - 4 \max(k, 2w)$ T gates, and the $y_n(3 - y_n^2g)$ multiplication costs $16w^2 + 8kw - 4k - 4$ T gates. The T-gate count of each $(3 - \#)$ is stated before. Adding all these contributions, the T-gate count comes to the value stated above. The cost breakdown of this Lemma is summarized in Table 12. \square

Lemma A.5 (The initial guess for an approximate evaluation of $1/\sqrt{g}$ via Newton’s method). *A suitable initial value, y_0 , for Newton’s method iterations in Eq. (119) is the inverse square root of the s^{th} power of four such that s is the closest integer to the logarithm of the input g in base four. In other words,*

$$y_0(g) = \frac{1}{\sqrt{4^{s(g)}}} \quad \text{with} \quad s(g) = \left\lfloor \log_4 g + \frac{1}{2} \right\rfloor. \quad (121)$$

This is sufficient to guarantee an approximately quadratic convergence to the inverse square root of an η -qubit integer input g . There exists a quantum circuit which maps

$$|g\rangle |0\rangle^{\otimes(1+\lceil\eta/2\rceil+\ell+\beta)} \rightarrow |g\rangle |y_0\rangle |\text{junk}\rangle |0\rangle^{\otimes\beta}.$$

Here, $1 + \lceil\eta/2\rceil$ qubits are required to store the initial guess, the state $|\text{junk}\rangle$ occupies an ℓ -qubit register with $\ell = \eta + 1$. Furthermore, the circuit is implemented using 4η T gates, and $\beta = 1$ ancilla qubit is needed to implement all the Toffoli gates. When the input register is in the state $g = 0$, the circuit returns $y_0(0) = 0$.

Proof. As discussed in the main text, if the relative error $|\delta_0|$ in the initial value for Newton’s method is less than one, a roughly quadratic convergence is guaranteed, see Eq. (69).

Now let $y_0 = 2^{-s}$ be the initial value and let $1/\sqrt{g}$ be the exact value. So $g = 4^s + r$ where, based on the definition, one can obtain the relation $-\frac{1}{2}4^s \leq r < 4^s$. This implies that

$$|\delta_0| = \frac{|y_0 - 1/\sqrt{g}|}{1/\sqrt{g}} = \left| \sqrt{\frac{g}{4^s}} - 1 \right| = \left| \sqrt{1 + \frac{r}{4^s}} - 1 \right| \leq |\sqrt{2} - 1| < 1. \quad (122)$$

To compute this initial guess, one can use the generalization of the circuit shown in Fig. 16. The input g considered in the figure is a five-bit integer. To represent the zeroth iterate of an η -bit integer input requires that the output precision is at least $1 + \lceil\eta/2\rceil$ bits. The circuit works by inspecting the integer input g and finding the first occurrence of 1 from left to right, i.e., from the most significant bit to the least significant bit,




Figure 16: Circuit example for computing y_0 in Eq. (121), that is the zeroth approximation to inverse square root of an integer. Here, the input is encoded in $\eta = 5$ qubits and in each collection of qubit registers, the subscripts refer to the power of two in the binary representation of the integer values the registers hold.

in the binary form of g . This bit is flagged by a ‘flag’ qubit initialized in state $|1\rangle$. Once the first 1 bit is found, the shown operations in the figure allow to set the correct bit in the initially $y_0 = 0$ register to one, providing the final $y_0 = 2^{-s}$ value. Another η -sized register, noted as ‘work’ register in the figure, is used to ensure that the flag is set to zero once the first bit of the input with value 1 is found, which stops the rest of the operations and refrains from accumulating any value in the remaining y_0 bits. The circuit involves η Toffoli gates which can be compiled with 4η T gates and one ancilla qubit. A similar procedure for finding the initial-value guess for both integer and non-integer inputs is discussed in Ref. [118]. \square

Lemma A.6 (Bounding the absolute error in the evaluation of $1/\sqrt{g}$ via a fixed-bit-precision application of Newton’s method). *Let $g \geq 1$ and let n and m be two positive integers. Furthermore, let y_k for $1 \leq k \leq n$ denote the k th exact output of Newton’s iteration in Eq. (119), and \tilde{y}_k denote the truncated output of the Newton’s iteration applied to \tilde{y}_{k-1} such that \tilde{y}_k is an m -bit number. Additionally, assume that the initial value y_0 is chosen via Lemma A.5 and is kept exact, i.e., $\tilde{y}_0 = y_0$. Then*

$$\left| \frac{1}{\sqrt{g}} - \tilde{y}_n \right| \leq 2^n (\sqrt{2} - 1)^{2^n} + 2^{2-m} \left[\left(\frac{3}{2} \right)^n - 1 \right]. \quad (123)$$

Proof. First note that for $n = 0$, Eq. (123) reproduces Eq. (122) considering that $1/\sqrt{g} \leq 1$. To find a bound on $|1/\sqrt{g} - \tilde{y}_n|$ for $n > 0$, that is the absolute error between the exact value and the value computed by n repetitions of the circuit in Fig. 9, we follow a derivation based on Ref. [176]. First, one notes that $|1/\sqrt{g} - \tilde{y}_n| \leq |1/\sqrt{g} - y_n| + |y_n - \tilde{y}_n|$, and by separately bounding $|1/\sqrt{g} - y_n|$ and $|y_n - \tilde{y}_n|$, the full error can be bounded.

To bound $|1/\sqrt{g} - y_n|$, recall that for the relative error $\delta_k = \frac{y_k - 1/\sqrt{g}}{1/\sqrt{g}}$, Eq. (69) states that $\delta_{k+1} = -\frac{3}{2}\delta_k^2 - \frac{1}{2}\delta_k^3$, which leads to $|\delta_n| \leq 2^n |\delta_0|^{2^n}$ for the relative error on the last output of Newton iterations. The convergence to the exact value $1/\sqrt{g}$ is achieved as long as $|\delta_0| < 1$. This condition is guaranteed using the initial guess from Lemma A.5 since $|\delta_0| \leq \sqrt{2} - 1 < 1$. The bound on $|\delta_n|$ gives:

$$\left| \frac{1}{\sqrt{g}} - y_n \right| = \frac{|\delta_n|}{\sqrt{g}} \leq \frac{2^n (\sqrt{2} - 1)^{2^n}}{\sqrt{g}} \leq 2^n (\sqrt{2} - 1)^{2^n}, \quad (124)$$

where the last inequality holds because $g \geq 1$.

To bound $|y_n - \tilde{y}_n|$, let \mathcal{N}_g be the Newton-iteration function such that $\mathcal{N}_g(y_k) = y_{k+1}$ according to Eq. (119). Denoting ξ_k to be the additive error occurring due to truncating \tilde{y}_k to only m bits of precision, that is

$$\begin{aligned}\tilde{y}_0 &= y_0, \\ \tilde{y}_1 &= \mathcal{N}_g(\tilde{y}_0) + \xi_1, \\ \tilde{y}_2 &= \mathcal{N}_g(\tilde{y}_1) + \xi_2, \\ &\vdots \\ \tilde{y}_n &= \mathcal{N}_g(\tilde{y}_{n-1}) + \xi_n,\end{aligned}$$

one obtains the following relation:

$$|y_k - \tilde{y}_k| \leq |\xi_k| + |y_{k-1} - \tilde{y}_{k-1}| \left| \max_{x \in \mathbb{D}} \mathcal{N}'_g(x) \right|, \quad (125)$$

for some domain \mathbb{D} over which Newton iterations are restricted during the application of each step.

A sufficient domain is $\mathbb{D} = [0, 1/\sqrt{g}]$. This can be seen by noting that: i) if $0 \leq y_k \leq 1/\sqrt{g}$, then $0 \leq y_{k+1} \leq 1/\sqrt{g}$, and ii) $0 \leq y_1 \leq 1/\sqrt{g}$. Combining i) and ii) implies the input to Newton-iteration function \mathcal{N}_g can only be within $[0, 1/\sqrt{g}]$. Number i) follows from the fact that $\mathcal{N}'_g(x) \geq 0$ over \mathbb{D} , implying \mathcal{N}_g takes its minimum at 0 and its maximum at $1/\sqrt{g}$ in this domain. Since $0 \leq y_k \leq 1/\sqrt{g}$ by assumption, then $\mathcal{N}_g(0) \leq \mathcal{N}_g(y_k) \leq \mathcal{N}_g(1/\sqrt{g})$, which implies $0 \leq y_{k+1} \leq 1/\sqrt{g}$. Number ii) follows from the fact that $y_1 - 1/\sqrt{g} = \frac{y_0}{2} (3 - y_0^2 g - \frac{2}{y_0 \sqrt{g}}) \leq 0$ for $\frac{1}{2} \leq y_0^2 g < 2$. This last condition can be shown to follow from the bound on the closeness of the initial guess to the exact function stated in Lemma A.5, so $y_1 \leq 1/\sqrt{g}$ is established. Similarly, since $y_1 = \frac{y_0}{2} (3 - y_0^2 g)$, the initial guess implies $0 \leq y_1$.

With the sufficient domain \mathbb{D} found, it is easy to see that $|\max_{x \in \mathbb{D}} \mathcal{N}'_g(x)| = 3/2$. Furthermore, since $0 \leq y_k \leq 1/\sqrt{g} \leq 1$, each y_k has only one significant bit at the one's place, and the rest of the bits occur after the decimal point. This means that the m -bit-truncated y_k incur an error $\xi \equiv \xi_k = 2^{1-m}$. With these information, Eq. (125) can be further processed:

$$\begin{aligned}|y_n - \tilde{y}_n| &\leq \xi + \frac{3}{2} |y_{n-1} - \tilde{y}_{n-1}| \\ &\leq \xi + \frac{3}{2} \xi + \left(\frac{3}{2}\right)^2 |y_{n-2} - \tilde{y}_{n-2}| \\ &\vdots \\ &\leq \xi + \frac{3}{2} \xi + \left(\frac{3}{2}\right)^2 \xi + \dots + \left(\frac{3}{2}\right)^{n-1} \xi = 2^{2-m} \left[\left(\frac{3}{2}\right)^n - 1 \right].\end{aligned} \quad (126)$$

Ultimately, combining Eqs. (124) and (126) gives a bound on the total absolute error:

$$\left| \frac{1}{\sqrt{g}} - \tilde{y}_n \right| \leq \left| \frac{1}{\sqrt{g}} - y_n \right| + |y_n - \tilde{y}_n| \leq 2^n (\sqrt{2} - 1)^{2^n} + 2^{2-m} \left[\left(\frac{3}{2}\right)^n - 1 \right], \quad (127)$$

as stated. \square

Lemma A.7 (Computing the Schwinger-boson kinetic phase in the fault-tolerant model). *Let η , m , and n be positive integers, with η denoting the size of the $|p\rangle$, $|q\rangle$, and $|p'\rangle$ registers, i.e., $0 \leq p, q, p' \leq 2^\eta - 1$. There exists a quantum circuit for $\mathcal{U}_{\tilde{\mathcal{D}}}$ which maps*

$$|p\rangle |q\rangle |p'\rangle |0\rangle^{\otimes(1+m+\ell+\beta)} \xrightarrow{\mathcal{U}_{\tilde{\mathcal{D}}}} |p\rangle |q\rangle |p+p'+1\rangle |\tilde{\mathcal{D}}(p, q, p')\rangle |\text{junk}\rangle |0\rangle^{\otimes\beta}.$$

Here, the state $|\tilde{\mathcal{D}}(p, q, p')\rangle$ denotes an m -bit fixed-precision approximation to

$$\mathcal{D}(p, q, p') = \sqrt{\frac{pq}{(p+p')(p+p'+1)}}$$

within absolute error $|\delta_{n,m}^{\text{Abs.}}| \leq 2^n (\sqrt{2} - 1)^{2^n} + 2^{2-m} ((3/2)^n - 1)$, where n denotes the number of Newton iterations used to evaluate the function to within this error. Furthermore, the output of each iteration step is truncated to m bits of precision as well. The state $|\text{junk}\rangle$ occupies at most an ℓ -qubit register, where

$$\ell = 3mn + 4\eta n + 15\eta + 3n + 6 + \max(2\eta + 2, m).$$

Routine	T gates	Work bits	Junk bits
$pq : \eta \times \eta$ mult.	$8\eta^2 - 4\eta$	2η	2η
$p + p' : \eta$ -bit adder	4η	η	0
$(pq)(p + p') : 2\eta \times (\eta + 1)$ mult.	$16\eta^2 + 8\eta$	4η	$3\eta + 1$
$p + p' + 1 : (\eta + 1)$ -bit incr.	$4\eta - 4$	$\eta - 2$	0
$(p + p')(pq(p + p' + 1)) : (\eta + 1) \times (3\eta + 1)$ mult.	$24\eta^2 + 20\eta + 4$	$6\eta + 2$	$k = 4\eta + 2$
Initial-value evaluation	$16\eta + 8$	0	$\max(2\eta + 2, m) + 4\eta + 3$
n steps of Newton iteration	$n(160\eta m + 32m^2 + 16\eta + 88m - 8 \max(4\eta + 2, 2m) - 4)$	$12\eta + 9m + 9$	$n(4\eta + 3m + 3)$
$pq \tilde{y}_n : 2\eta \times m$ mult.	$16\eta m - 4 \max(2\eta, m)$	$2 \max(2\eta, m)$	2η
Overall circuit	$160\eta mn + 32m^2 n + 48\eta^2 + 16\eta m + 16\eta n + 88mn - 8n \max(4\eta + 2, 2m) + 48\eta - 4 \max(2\eta, m) - 4n + 8$	$12\eta + 9m + 9$	$3mn + 4\eta n + 15\eta + 3n + 6 + \max(2\eta + 2, m)$

Table 13: Costs associated with the $\mathcal{U}_{\tilde{D}}^{\text{SB}}$ circuit for the Schwinger-boson formulation.

Furthermore, the circuit can be implemented using a number of T gates equal to

$$160\eta mn + 32m^2 n + 48\eta^2 + 16\eta m + 16\eta n + 88mn - 8n \max(4\eta + 2, 2m) + 48\eta - 4 \max(2\eta, m) - 4n + 8.$$

The number of ancillary workspace qubits is given by:

$$\beta = 12\eta + 9m + 9.$$

Additionally, an extra ancilla qubit may be needed for the execution of all the Toffoli gates.

Proof. The size of the workspace that is needed to store all the intermediate values required to approximately compute $\sqrt{pq}/((p + p')(p + p' + 1))$ can be read off from the circuit shown in Fig. 8, and is summarized for each element of the circuit in Table 13. These add up to the value ℓ as stated. This register, denoted as $|\text{junk}\rangle$, is not reset at the end of the $\mathcal{U}_{\tilde{D}}$ evaluation circuit, and hence it will not be usable for later computations. Note that for an input size $k = 4\eta + 2$, the initial-value circuit needs $1 + \lfloor k/2 \rfloor = 2\eta + 2$ qubits to store the value of y_0 . However, if this is smaller than the m -bit register size required to store all the Newton iterates, one should pad the bit representation of y_0 to enlarge it to an m -bit register.

The T-gate and ancillary-qubit cost of the circuit can be evaluated as following. Inspecting the circuit in Fig. 8 from left to right, one first encounters: i) a multiplication circuit for evaluating pq which requires $8\eta^2 - 4\eta$ T gates and 2η ancilla qubits, ii) an addition circuit that evaluates $p + p'$ and requires 4η T gates and η ancilla qubits, iii) a multiplication circuit for evaluating $(pq)(p + p')$ which requires $16\eta^2 + 8\eta$ T gates and 4η ancilla qubits, iv) an incrementer which increases by one the value of $p + p'$ and requires $4\eta - 4$ T gates and $\eta - 2$ ancilla qubits, v) a multiplication circuit for evaluating $(p + p')(pq(p + p' + 1))$ which requires $24\eta^2 + 20\eta + 4$ T gates and $6\eta + 2$ ancilla qubits, vi) a circuit for computing the initial value for the subsequent application Newton's method with the input $pq(p + p')(p + p' + 1)$ with size $k = 4\eta + 2$ which requires $4k$ T gates and no ancilla qubit (beside one for evaluating the Toffoli gates which we take into account later), vii) n applications of Newton-iteration steps with input size k and output size m , which requires in total $n(160\eta m + 32m^2 + 16\eta + 88m - 8 \max(4\eta + 2, 2m) - 4)$ T gates and $2 \max(2\eta, m)$ ancilla qubits, and finally viii) a multiplication circuit for evaluating $pq \tilde{y}_n$ with \tilde{y}_n being the m -bit precision output of the last Newton iteration, which requires $16\eta m - 4 \max(2\eta, m)$ T gates and $2 \max(2\eta, m)$ ancilla qubit.

The total T-gate cost of the circuit is the sum of the cost of each subcircuit above and adds up to the value stated. The total auxiliary qubit cost of the circuit is that of the subcircuit which requires the largest workspace, that is the workspace of the Newton iterations. This gives the value of β as stated above. Note that an extra ancilla qubit may be needed throughout the circuit for implementing all Toffoli gates. \square

Lemma A.8 (Computing the LSH kinetic phase in the fault-tolerant model). *Let η , m , n , and n_ℓ be positive integers with $0 \leq n_\ell \leq 2^n - 1$, and n_q and n'_q be single-bit numbers 0 or 1. There exists a quantum circuit for $\mathcal{U}_{\tilde{D}}$ which maps*

$$|n_q\rangle |n'_q\rangle |n_\ell\rangle |0\rangle^{\otimes(1+m+\ell+\beta)} \xrightarrow{\mathcal{U}_{\tilde{D}}} |n_q\rangle |n'_q\rangle |n_\ell + 1 + n_q\rangle |\tilde{D}(n_\ell, n_q, n'_q)\rangle |\text{junk}\rangle |0\rangle^{\otimes\beta}.$$

Routine	T gates	Work bits	Junk bits
$(\eta + 1)$ -bit incr.	$4\eta - 4$	$\eta - 2$	0
Copy $\eta + 1$ bits	0	0	$\eta + 1$
Two controlled $(\eta + 2)$ -bit incrs.	8η	$\eta - 1$	0
$(n_\ell + 1 + n_q)(n_\ell + 1 + n'_q) : (\eta + 1) \times (\eta + 1)$ mult.	$8\eta^2 + 12\eta + 4$	$2\eta + 2$	$2\eta + 2$
Initial-value evaluation	$8\eta + 8$	0	$\max(\eta + 2, m) + 2\eta + 3$
n steps of Newton iteration	$n(80\eta m + 32m^2 + 8\eta + 88m - 16 \max(\eta + 1, m) - 4)$	$6\eta + 9m + 9$	$n(2\eta + 3m + 3)$
$(n_\ell + 1 + n'_q)\tilde{y}_n : (\eta + 1) \times m$ mult.	$8m\eta + 8m - 4 \max(\eta + 1, m)$	$2 \max(\eta + 1, m)$	$\eta + 1$
Overall circuit	$80\eta mn + 32m^2n + 8\eta^2 + 8\eta m + 8\eta n + 88mn - 8n \max(2\eta + 2, 2m) + 32\eta - 4 \max(\eta + 1, m) + 8m - 4n + 8$	$6\eta + 9m + 9$	$2\eta n + 3mn + 6\eta + \max(\eta + 2, m) + 3n + 7$

Table 14: Costs associated with the $\mathcal{Q}_D^{\text{LSH}}$ circuit for the LSH formulation.

Here, the state $|\tilde{\mathcal{D}}(n_\ell, n_q, n'_q)\rangle$ denotes an m -bit fixed-precision approximation to

$$\mathcal{D}(n_\ell, n_q, n'_q) = \sqrt{\frac{n_\ell + 1 + n_q}{n_\ell + 1 + n'_q}}$$

within absolute error $|\delta_{n,m}^{\text{Abs.}}| \leq \sqrt{2}(2^n(\sqrt{2} - 1)^{2^n} + 2^{2-m}((3/2)^n - 1))$, where n denotes the number of Newton iterations used to evaluate the function to within this error. Furthermore, the output of each iteration step is truncated to m bits of precision as well. The state $|\text{junk}\rangle$ occupies at most an ℓ -qubit register, where

$$\ell = 2\eta n + 3mn + 6\eta + \max(\eta + 2, m) + 3n + 7.$$

Furthermore, the circuit can be implemented using a number of T gates equal to

$$80\eta mn + 32m^2n + 8\eta^2 + 8\eta m + 8\eta n + 88mn - 8n \max(2\eta + 2, 2m) + 32\eta - 4 \max(\eta + 1, m) + 8m - 4n + 8.$$

The number of ancillary workspace qubits is given by

$$\beta = 6\eta + 9m + 9.$$

Additionally, an extra ancilla qubit may be needed for the execution of all the Toffoli gates.

Proof. The size of the workspace that is needed to store all the intermediate values required to approximately compute $\sqrt{(n_\ell + 1 + n_q)/(n_\ell + 1 + n'_q)}$ can be read off from the circuit shown in Fig. 11, and is summarized for each element of the circuit in Table 14, adding up to the value ℓ as stated. In particular, the copy operation, the first multiplier, the initial-value circuit, each Newton block, and the final multiplier are the operations that need such workspace. These lead to the state $|\text{junk}\rangle$ that is not reset at the end of the $\mathcal{Q}_D^{\text{LSH}}$ evaluation circuit, and hence it will not be usable for later computations. Note that for the input register size $k = 2\eta + 2$, if $1 + \lfloor k/2 \rfloor = \eta + 2 < m$, the output of the initial-value circuit must be padded to m bits in order to apply Lemma A.4.

The T -gate and ancillary-qubit cost of the circuit can be evaluated as following. Inspecting the circuit in Fig. 11 from left to right, one first encounters: i) an incrementer that is applied to an $(\eta + 1)$ -bit register, with the associated T -gate count $4\eta - 4$ and ancilla-qubit count $\eta - 2$ according to Lemma A.2, a copy operation which can be done with CNOT gates alone, costing no T gates and ancilla qubits, iii) two controlled incrementers each of which can be implemented using an ordinary incrementer on $\eta + 2$ bits (followed by a bit flip on the ‘‘control’’ qubit), costing 4η T gates and $\eta - 1$ ancilla qubits, iv) a multiplier which is applied to two $(\eta + 1)$ -bit integers producing the value $(n_\ell + 1 + n_q)(n_\ell + 1 + n'_q)$ at a cost of $8\eta^2 + 8\eta + 4$ T gates and $2\eta + 2$ ancilla qubits according to Lemma A.3, v) the initial-value circuit that is applied to an argument of size $k = 2\eta + 2$ at a cost of $8\eta + 8$ T gates and no ancilla qubit (beside one reusable one for implementing all Toffoli gates) according to Lemma A.5, vi) n steps of Newton iteration, each with a cost of $80\eta m + 32m^2 + 8\eta + 88m - 16 \max(\eta + 1, m) - 4$ T gates and $6\eta + 9m + 9$ ancilla qubits as implied by Lemma A.4, and finally vii) a multiplier whose inputs are the leading m bits of the last Newton iterate \tilde{y}_n and the $(\eta + 1)$ -bit integer $n_\ell + 1 + n_q$, which costs $8m\eta + 8m - 4 \max(\eta + 1, m)$ T gates and $2 \max(\eta + 1, m)$ ancilla qubits.

The total T-gate cost of the circuit is the sum of the cost of each subcircuit above and adds up to the value stated. The total auxiliary qubit cost of the circuit is that of the subcircuit which requires the largest workspace, that is the workspace of the Newton iterations. This gives the value of β as stated above. Note that an extra ancilla qubit may be needed throughout the circuit for implementing all Toffoli gates. \square

B Deriving the loop-string-hadron hopping subterms

In this appendix, we show that the LSH hopping Hamiltonian written as the two subterms given in Eqs. (50) are equivalent to the original form in Eq. (43) of Ref. [131], that is

$$H_I^{\text{LSH}}(r) = H_I^{(o)}(r) + H_I^{(i)}(r), \quad (128)$$

$$H_I^{(o)}(r) \equiv x \frac{1}{\sqrt{N^L(r)+1}} \mathcal{S}_{\text{out}}^{++}(r) \mathcal{S}_{\text{in}}^{+-}(r+1) \frac{1}{\sqrt{N^R(r+1)+1}} + \text{H.c.}, \quad (129)$$

$$H_I^{(i)}(r) \equiv x \frac{1}{\sqrt{N^L(r)+1}} \mathcal{S}_{\text{out}}^{+-}(r) \mathcal{S}_{\text{in}}^{--}(r+1) \frac{1}{\sqrt{N^R(r+1)+1}} + \text{H.c.}, \quad (130)$$

where $N^{L/R}$ are the flux quantum numbers used throughout the Schwinger-boson and LSH sections and $\mathcal{S}_{\text{out/in}}$ are site-local string operators in the LSH formulation defined in Eqs. (45). One can check, by expanding out the definitions of $N^{L/R}$ and the string operators, that these subterms can be expressed as

$$H_I^{(o)} = x \left[\chi_o^\dagger(r) \Gamma^\dagger(r)^{n_i(r)} \right] \left[\chi_o(r+1) \Gamma^\dagger(r+1)^{1-n_i(r+1)} \right] \quad (131)$$

$$\times \sqrt{\frac{n_\ell(r)+1}{n_\ell(r)+1+n_i(r)}} \sqrt{\frac{n_\ell(r+1)+1+n_i(r+1)}{n_\ell(r+1)+1}} + \text{H.c.}, \quad (132)$$

$$H_I^{(i)} = x \left[\chi_i^\dagger(r+1) \Gamma^\dagger(r+1)^{n_o(r+1)} \right] \left[\chi_i(r) \Gamma^\dagger(r)^{1-n_o(r)} \right] \quad (133)$$

$$\times \sqrt{\frac{n_\ell(r)+1+n_o(r)}{n_\ell(r)+1}} \sqrt{\frac{n_\ell(r+1)+1}{n_\ell(r+1)+1+n_o(r+1)}} + \text{H.c.} \quad (134)$$

The square-root functions in Eqs. (132) and (134) can be simplified using the AGL. For example, consider the square roots in Eq. (132). The AGL on link r can be used to solve for $n_\ell(r+1)$ as

$$n_\ell(r+1) = n_\ell(r) + n_o(r) - n_i(r)n_o(r) - n_i(r+1) + n_i(r+1)n_o(r+1).$$

However, since this will need to be evaluated on the right of $\chi_o^\dagger(r)\chi_o(r+1)$, one can effectively set $n_o(r) \rightarrow 0$ and $n_o(r+1) \rightarrow 1$ in the formula for $n_\ell(r+1)$. Thus,

$$\begin{aligned} n_\ell(r+1) &\rightarrow n_\ell(r) + 0 - 0 - n_i(r+1) + n_i(r+1) = n_\ell(r), \\ &\Rightarrow \sqrt{\frac{n_\ell(r)+1}{n_\ell(r)+1+n_i(r)}} \sqrt{\frac{n_\ell(r+1)+1+n_i(r+1)}{n_\ell(r+1)+1}} \rightarrow \sqrt{\frac{n_\ell(r)+1+n_i(r+1)}{n_\ell(r)+1+n_i(r)}}. \end{aligned} \quad (135)$$

An analogous calculation can be done to simplify Eq. (132) by first solving for $n_\ell(r)$ using the AGL and then substituting this into the square-root functions in Eq. (132). The result is

$$\sqrt{\frac{n_\ell(r)+1+n_o(r)}{n_\ell(r)+1}} \sqrt{\frac{n_\ell(r+1)+1}{n_\ell(r+1)+1+n_o(r+1)}} \rightarrow \sqrt{\frac{n_\ell(r+1)+1+n_o(r)}{n_\ell(r+1)+1+n_o(r+1)}}.$$

To summarize, using the AGL, one can equivalently use

$$H_I^{(o)}(r) \equiv x \left[\chi_o^\dagger(r) \Gamma^\dagger(r)^{n_i(r)} \right] \left[\chi_o(r+1) \Gamma^\dagger(r+1)^{1-n_i(r+1)} \right] \sqrt{\frac{n_\ell(r+1)+1+n_i(r+1)}{n_\ell(r+1)+1+n_i(r)}} + \text{H.c.}, \quad (136)$$

$$H_I^{(i)}(r) \equiv x \left[\chi_i^\dagger(r+1) \Gamma^\dagger(r+1)^{n_o(r+1)} \right] \left[\chi_i(r) \Gamma^\dagger(r)^{1-n_o(r)} \right] \sqrt{\frac{n_\ell(r)+1+n_o(r)}{n_\ell(r)+1+n_o(r+1)}} + \text{H.c.}, \quad (137)$$

from which H_I^{LSH} in Eq. (50) is reproduced.

C Second-order commutator bounds

In this appendix, we compute the second-order Trotter error bound of Ref. [104]:

$$\begin{aligned} \|V(\theta) - e^{-i\theta H}\| &\leq \frac{\theta^3}{24} \sum_{\gamma_1=1}^{\Upsilon} \left\| \left[H_{\gamma_1}, \left[H_{\gamma_1}, \sum_{\gamma_2=\gamma_1+1}^{\Upsilon} H_{\gamma_2} \right] \right] \right\| \\ &\quad + \frac{\theta^3}{12} \sum_{\gamma_1=1}^{\Upsilon} \left\| \left[\sum_{\gamma_3=\gamma_1+1}^{\Upsilon} H_{\gamma_3}, \left[\sum_{\gamma_2=\gamma_1+1}^{\Upsilon} H_{\gamma_2}, H_{\gamma_1} \right] \right] \right\|. \end{aligned} \quad (138)$$

for both the Schwinger-boson and the LSH formulations. The ordered operators, H_γ , are chosen as

$$\left\{ \dots, H_M(r), H_E(r), \{H_I^{(j)}(r)\}_{j=1}^\nu, H_M(r+1), H_E(r+1), \{H_I^{(j)}(r+1)\}_{j=1}^\nu, \dots \right\}, \quad (139)$$

for both formulations, where $\nu = 8$ for the Schwinger-boson formulation and $\nu = 2$ for the LSH formulation.

The first step one may take is to move the summations to the outside of the operator norms, making use of the triangle inequality:

$$\begin{aligned} \|V(\theta) - e^{-i\theta H}\| &\leq \frac{\theta^3}{24} \sum_{\gamma_1=1}^{\Upsilon} \sum_{\gamma_2=\gamma_1+1}^{\Upsilon} \left\| \left[H_{\gamma_1}, \left[H_{\gamma_1}, H_{\gamma_2} \right] \right] \right\| \\ &\quad + \frac{\theta^3}{12} \sum_{\gamma_1=1}^{\Upsilon} \sum_{\gamma_2=\gamma_1+1}^{\Upsilon} \sum_{\gamma_3=\gamma_1+1}^{\Upsilon} \left\| \left[H_{\gamma_3}, \left[H_{\gamma_2}, H_{\gamma_1} \right] \right] \right\|. \end{aligned} \quad (140)$$

This is because we generally do not see cancellations take place by keeping the sums inside the commutator brackets, and whenever a norm of the form $\|A+B\|$ is encountered, a better upper bound on this than $\|A\| + \|B\|$ is often not known.

The second-order Trotter-error-bound calculation has now been turned into an exercise in evaluating various double commutators of simulatable Hamiltonian terms. One first identifies non-vanishing commutators in the two sums, noting that many of the commutators are identically zero due to the commutativity of operators from distant sites/links. In the double sum, $\sum_{\gamma_1} \sum_{\gamma_2 > \gamma_1}$, the five non-vanishing operator forms are

$$\mathcal{O}_1 = \left[H_M(r), [H_M(r), H_I^{(j)}(r)] \right], \quad (141)$$

$$\mathcal{O}_2 = \left[H_E(r), [H_E(r), H_I^{(j)}(r)] \right], \quad (142)$$

$$\mathcal{O}_3 = \left[H_I^{(j)}(r), [H_I^{(j)}(r), H_I^{(k)}(r)] \right] \quad (k > j), \quad (143)$$

$$\mathcal{O}_4 = \left[H_I^{(j)}(r), [H_I^{(j)}(r), H_M(r+1)] \right], \quad (144)$$

$$\mathcal{O}_5 = \left[H_I^{(j)}(r), [H_I^{(j)}(r), H_I^{(k)}(r+1)] \right]. \quad (145)$$

$$(146)$$

In the triple sum, $\sum_{\gamma_1} \sum_{\gamma_2 > \gamma_1} \sum_{\gamma_3 > \gamma_1}$, the nineteen non-vanishing operator forms are

$$\mathcal{O}_6 = \left[H_E(r), [H_I^{(j)}(r), H_M(r)] \right], \quad (147)$$

$$\mathcal{O}_7 = \left[H_I^{(k)}(r), [H_I^{(j)}(r), H_M(r)] \right], \quad (148)$$

$$\mathcal{O}_8 = \left[H_M(r+1), [H_M(r), H_I^{(j)}(r)] \right], \quad (149)$$

$$\mathcal{O}_9 = \left[H_I^{(k)}(r+1), [H_I^{(j)}(r), H_M(r)] \right], \quad (150)$$

$$(151)$$

$$\mathcal{O}_{10} = \left[H_I^{(k)}(r), [H_I^{(j)}(r), H_E(r)] \right], \quad (152)$$

$$\mathcal{O}_{11} = \left[H_M(r+1), [H_I^{(j)}(r), H_E(r)] \right], \quad (153)$$

$$\mathcal{O}_{12} = \left[H_I^{(k)}(r+1), [H_I^{(j)}(r), H_E(r)] \right], \quad (154)$$

(155)

$$\mathcal{O}_{13} = \left[H_I^{(l)}(r), [H_I^{(k)}(r), H_I^{(j)}(r)] \right] \quad (k > j, l > j), \quad (156)$$

$$\mathcal{O}_{14} = \left[H_M(r+1), [H_I^{(k)}(r), H_I^{(j)}(r)] \right] \quad (k > j), \quad (157)$$

$$\mathcal{O}_{15} = \left[H_I^{(l)}(r+1), [H_I^{(k)}(r), H_I^{(j)}(r)] \right] \quad (k > j), \quad (158)$$

(159)

$$\mathcal{O}_{16} = \left[H_I^{(k)}(r), [H_M(r+1), H_I^{(j)}(r)] \right] \quad (k > j), \quad (160)$$

$$\mathcal{O}_{17} = \left[H_M(r+1), [H_M(r+1), H_I^{(j)}(r)] \right], \quad (161)$$

$$\mathcal{O}_{18} = \left[H_I^{(k)}(r+1), [H_M(r+1), H_I^{(j)}(r)] \right] \quad (k > j), \quad (162)$$

(163)

$$\mathcal{O}_{19} = \left[H_I^{(l)}(r), [H_I^{(k)}(r+1), H_I^{(j)}(r)] \right] \quad (l > j), \quad (164)$$

$$\mathcal{O}_{20} = \left[H_M(r+1), [H_I^{(k)}(r+1), H_I^{(j)}(r)] \right], \quad (165)$$

$$\mathcal{O}_{21} = \left[H_E(r+1), [H_I^{(k)}(r+1), H_I^{(j)}(r)] \right], \quad (166)$$

$$\mathcal{O}_{22} = \left[H_I^{(l)}(r+1), [H_I^{(k)}(r+1), H_I^{(j)}(r)] \right], \quad (167)$$

$$\mathcal{O}_{23} = \left[H_M(r+2), [H_I^{(k)}(r+1), H_I^{(j)}(r)] \right], \quad (168)$$

$$\mathcal{O}_{24} = \left[H_I^{(l)}(r+2), [H_I^{(k)}(r+1), H_I^{(j)}(r)] \right]. \quad (169)$$

Strictly speaking, some of the above operator forms are not relevant at the boundaries of the lattice. For simplicity, the Trotterization error bound will be slightly inflated by ignoring the OBCs and acting as if every one of these forms appears L times.

Turning now to the evaluations, some of the commutators are evaluated (or partially evaluated) by first noting that each hopping term in either formulations has the form

$$H_I^{(j)}(r) = x(K^{(j)}(r) + K^{(j)\dagger}(r)), \quad (170)$$

where $K^{(j)}$ ($K^{(j)\dagger}$) can be read from the hopping Hamiltonians in Eq. (50). They contain both bosonic and fermionic creation and/or annihilation operators. Then one can show that generally

$$[\psi^\dagger(r)\psi(r), H_I^{(j)}(r)] = \pm x (K^{(j)\dagger}(r) - K^{(j)}(r)), \quad (171a)$$

$$[\psi^\dagger(r+1)\psi(r+1), H_I^{(j)}(r)] = \mp x (K^{(j)\dagger}(r) - K^{(j)}(r)), \quad (171b)$$

where the \pm and \mp are correlated and depend on the specific $H_I^{(j)}$ being considered. Note that the commutators in Eqs. (171) hold whether or not the $K^{(j)}$ has a cutoff built in. Furthermore, the electric Hamiltonian at each link in either formulation can be written as $J^L(r)(J^L(r) + 1)$ where $J^L = N^L/2$ in either formulation. Then one can show that

$$[J^L(r)(J^L(r) + 1), H_I^{(j)}(r)] = \left(J^L(r) + \frac{1}{4} \right) x K^{(j)\dagger}(r) - x K^{(j)}(r) \left(J^L(r) + \frac{1}{4} \right), \quad (172)$$

for both formulations. Following whatever simplifications are available, one may eventually resort to using derivatives of the triangle inequality of the form

$$\|[A, B]\| \leq 2\|A\|\|B\|, \quad (173)$$

$$\|[A, [B, C]]\| \leq 4\|A\|\|B\|\|C\|. \quad (174)$$

In this work we do not find useful simplifications to commutators of the form $[H_I^{(j)}(r), H_I^{(k)}(r')]$ to apply before invoking one of Eqs. (173) or (174).

In summary, the norm evaluation will always boil down to some combinations of $\|H_M(r)\|$, $\|J^L(r) + 1/4\|$, and/or $\|xK^{(j)\dagger} \pm xK^{(j)}\|$. While $\|H_M(r)\| \leq \mu$ for both formulations, the bounds on $\|J^L(r) + 1/4\|$ and

$\|xK^{(j)\dagger} \pm xK^{(j)}\|$ are formulation dependent, which we evaluate in the respective subsections below. Nonetheless, a general intermediate bound can be still placed on $\|xK^{(j)\dagger} \pm xK^{(j)}\|$ as follows.

One first notes that $K^{(j)}$ contains at least one fermionic operator, e.g., ψ_x^\dagger . Let $|\Psi\rangle = c_0 |\phi_0\rangle + c_1 |\phi_1\rangle$ denote an eigenvector of $K^{(j)} + K^{(j)\dagger}$ with the largest eigenvalue, where $\psi_x^\dagger \psi_x |\phi_b\rangle = b |\phi_b\rangle$ for $b = 0, 1$, $\langle \phi_{b'} | \phi_b \rangle = \delta_{b' b}$, and hence $|c_0|^2 + |c_1|^2 = 1$. Now,

$$\begin{aligned} \|K^{(j)} \pm K^{(j)\dagger}\| &= \left| \langle K^{(j)} \pm K^{(j)\dagger} | \Psi \rangle \right| \\ &= \left| c_0 K^{(j)} |\phi_0\rangle \pm c_1 K^{(j)\dagger} |\phi_1\rangle \right| \\ &= \sqrt{|c_0|^2 \langle \phi_0 | K^{(j)\dagger} K^{(j)} | \phi_0 \rangle + |c_1|^2 \langle \phi_1 | K^{(j)} K^{(j)\dagger} | \phi_1 \rangle} \\ &\leq \sqrt{(|c_0|^2 + |c_1|^2) \max \left(\langle \phi_0 | K^{(j)\dagger} K^{(j)} | \phi_0 \rangle, \langle \phi_1 | K^{(j)} K^{(j)\dagger} | \phi_1 \rangle \right)} \\ &= \sqrt{\max \left(\langle \phi_0 | K^{(j)\dagger} K^{(j)} | \phi_0 \rangle, \langle \phi_1 | K^{(j)} K^{(j)\dagger} | \phi_1 \rangle \right)} \\ &= \max \left(\sqrt{\langle \phi_0 | K^{(j)\dagger} K^{(j)} | \phi_0 \rangle}, \sqrt{\langle \phi_1 | K^{(j)} K^{(j)\dagger} | \phi_1 \rangle} \right) \\ &\leq \max(\|K^{(j)}\|, \|K^{(j)\dagger}\|). \end{aligned}$$

From this, it immediately follows that

$$\|xK^{(j)\dagger} \pm xK^{(j)}\| \leq \max(\|xK^{(j)}\|, \|xK^{(j)\dagger}\|). \quad (175)$$

The bound on $\|K^{(j)}\|$ and $\|K^{(j)\dagger}\|$ will be discussed in the following for each formulation.

C.0.1 Schwinger-boson commutator bounds

Recall the generalized Schwinger-boson hopping term as written in Eq. (59),

$$H_I^{\text{SB}(j)} = \pm x \psi_x^\dagger \psi_y \lambda_p^- \lambda_q^- \mathcal{D}^{\text{SB}}(p, q, p') + \text{H.c.}$$

Above, one can identify

$$K^{(j)} = \pm \psi_x^\dagger \psi_y \lambda_p^- \lambda_q^- \mathcal{D}^{\text{SB}}(p, q, p'), \quad (176a)$$

$$K^{(j)\dagger} = \pm \psi_y^\dagger \psi_x \lambda_p^+ \lambda_q^+ \mathcal{D}^{\text{SB}}(p+1, q+1, p'), \quad (176b)$$

where $\mathcal{D}^{\text{SB}}(p, q, p') = \sqrt{pq/((p+p')(p+p'+1))}$. Given Eq. (175), the next step is the evaluation of $\max(\|xK^{(j)}\|, \|xK^{(j)\dagger}\|)$ which boils down to the bounds on $\|\mathcal{D}^{\text{SB}}\|$. It is easy to see that $\|\mathcal{D}^{\text{SB}}(p, q, p')\| \leq 1$. This is because the AGL implies $q \leq p + p'$, and therefore,

$$\frac{pq}{(p+p')(p+p'+1)} \leq \frac{p}{p+p'+1} \leq \frac{p}{p+1} < 1.$$

A similar argument can be made to show that $\|\mathcal{D}^{\text{SB}}(p+1, q+1, p')\| \leq 1$ as well. In either case, the norm upper bound immediately follows upon reinserting the factor of x ,¹⁸

$$\|xK^{(j)\dagger} \pm xK^{(j)}\| \leq x. \quad (177)$$

As for the commutators involving $H_E^{\text{SB}}(r)$, one notes that for the Schwinger-boson formulation,

$$J^L(r) = \frac{1}{2}(N_1^L(r) + N_2^L(r)), \quad (178)$$

hence

$$\|J^L(r) + \frac{1}{4}\| \leq \Lambda + \frac{1}{4}. \quad (179)$$

¹⁸Note that in our convention, $x \geq 0$.

Equipped with these bounds, one can proceed to evaluating a bound on the double commutators. When the double commutators involve only one $H_I^{(j)}$ operator, Eqs. (171) help one to ‘fully’ evaluate the norms as

$$\|\mathcal{O}_1\| = \|[H_M(r), [H_M(r), H_I^{(j)}(r)]]\| \leq x\mu^2, \quad (180)$$

$$\|\mathcal{O}_2\| = \|[H_E(r), [H_E(r), H_I^{(j)}(r)]]\| \leq x \left(\Lambda + \frac{1}{4} \right)^2, \quad (181)$$

$$\|\mathcal{O}_6\| = \|[H_E(r), [H_I^{(j)}(r), H_M(r)]]\| \leq x \left(\Lambda + \frac{1}{4} \right) \mu, \quad (182)$$

$$\|\mathcal{O}_8\| = \|[H_M(r+1), [H_I^{(j)}(r), H_M(r)]]\| \leq x\mu^2, \quad (183)$$

$$\|\mathcal{O}_{11}\| = \|[H_M(r+1), [H_I^{(j)}(r), H_E(r)]]\| \leq x \left(\Lambda + \frac{1}{4} \right) \mu, \quad (184)$$

$$\|\mathcal{O}_{17}\| = \|[H_M(r+1), [H_M(r+1), H_I^{(j)}(r)]]\| \leq x\mu^2, \quad (185)$$

without ever appealing to Eqs. (173) or (174). When the ‘inner’ commutator and ‘outer’ commutator each contain a $H_I^{(j)}$, one can use the bounds obtained on $K^{(j)}$ and J^L , in combination with Eq. (173), to obtain

$$\|\mathcal{O}_4\| = \|[H_I^{(j)}(r), [H_I^{(j)}(r), H_M(r+1)]]\| \leq 2x^2\mu, \quad (186)$$

$$\|\mathcal{O}_7\| = \|[H_I^{(k)}(r), [H_I^{(j)}(r), H_M(r)]]\| \leq 2x^2\mu, \quad (187)$$

$$\|\mathcal{O}_9\| = \|[H_I^{(k)}(r+1), [H_I^{(j)}(r), H_M(r)]]\| \leq 2x^2\mu, \quad (188)$$

$$\|\mathcal{O}_{10}\| = \|[H_I^{(k)}(r), [H_I^{(j)}(r), H_E(r)]]\| \leq 2x^2 \left(\Lambda + \frac{1}{4} \right), \quad (189)$$

$$\|\mathcal{O}_{12}\| = \|[H_I^{(k)}(r+1), [H_I^{(j)}(r), H_E(r)]]\| \leq 2x^2 \left(\Lambda + \frac{1}{4} \right), \quad (190)$$

$$\|\mathcal{O}_{16}\| = \|[H_I^{(k)}(r), [H_M(r+1), H_I^{(j)}(r)]]\| \leq 2x^2\mu \quad (k > j), \quad (191)$$

$$\|\mathcal{O}_{18}\| = \|[H_I^{(k)}(r+1), [H_M(r+1), H_I^{(j)}(r)]]\| \leq 2x^2\mu \quad (k > j). \quad (192)$$

When the ‘inner’ commutator involves two $H_I^{(j)}$ terms, one usually appeals directly to Eq. (174), obtaining

$$\|\mathcal{O}_3\| = \|[H_I^{(j)}(r), [H_I^{(j)}(r), H_I^{(k)}(r)]]\| \leq 4x^3 \quad (k > j), \quad (193)$$

$$\|\mathcal{O}_5\| = \|[H_I^{(j)}(r), [H_I^{(j)}(r), H_I^{(k)}(r+1)]]\| \leq 4x^3, \quad (194)$$

$$\|\mathcal{O}_{13}\| = \|[H_I^{(l)}(r), [H_I^{(k)}(r), H_I^{(j)}(r)]]\| \leq 4x^3 \quad (k > j, l > j), \quad (195)$$

$$\|\mathcal{O}_{14}\| = \|[H_M(r+1), [H_I^{(k)}(r), H_I^{(j)}(r)]]\| \leq 4x^2\mu \quad (k > j), \quad (196)$$

$$\|\mathcal{O}_{15}\| = \|[H_I^{(l)}(r+1), [H_I^{(k)}(r), H_I^{(j)}(r)]]\| \leq 4x^3 \quad (k > j), \quad (197)$$

$$\|\mathcal{O}_{19}\| = \|[H_I^{(l)}(r), [H_I^{(k)}(r+1), H_I^{(j)}(r)]]\| \leq 4x^3 \quad (l > j), \quad (198)$$

$$\|\mathcal{O}_{20}\| = \|[H_M(r+1), [H_I^{(k)}(r+1), H_I^{(j)}(r)]]\| \leq 4x^2\mu, \quad (199)$$

$$\|\mathcal{O}_{22}\| = \|[H_I^{(l)}(r+1), [H_I^{(k)}(r+1), H_I^{(j)}(r)]]\| \leq 4x^3, \quad (200)$$

$$\|\mathcal{O}_{24}\| = \|[H_I^{(l)}(r+2), [H_I^{(k)}(r+1), H_I^{(j)}(r)]]\| \leq 4x^3. \quad (201)$$

In the cases of \mathcal{O}_{21} and \mathcal{O}_{23} , however, it is observed that the Jacobi identity could be used in combination with Eq. (173) to derive

$$\|\mathcal{O}_{21}\| = \|[H_E(r+1), [H_I^{(k)}(r+1), H_I^{(j)}(r)]]\| \leq 2x^2 \left(\Lambda + \frac{1}{4} \right), \quad (202)$$

$$\|\mathcal{O}_{23}\| = \|[H_M(r+2), [H_I^{(k)}(r+1), H_I^{(j)}(r)]]\| \leq 2x^2\mu. \quad (203)$$

With upper bounds derived for all commutator norms, the only remaining step is to add them all up with appropriate weights. Every operator form involves at least one $H_I^{(j)}$ operator, and hence belongs to one or more

Operator form	Coefficient	\mathcal{O}_i multiplicity	\mathcal{O}_i bound (SB)	\mathcal{O}_i bound (LSH)
\mathcal{O}_1	$\frac{1}{24}$	ν	$\mu^2 x$	$\sqrt{2}\mu^2 x$
\mathcal{O}_2	$\frac{1}{24}$	ν	$(\Lambda + \frac{1}{4})^2 x$	$\sqrt{2}(\frac{\Lambda}{2} + \frac{3}{4})^2 x$
\mathcal{O}_3	$\frac{1}{24}$	$\frac{1}{2}(\nu - 1)\nu$	$4x^3$	$8\sqrt{2}x^3$
\mathcal{O}_4	$\frac{1}{24}$	ν	$2\mu x^2$	$4\mu x^2$
\mathcal{O}_5	$\frac{1}{24}$	ν^2	$4x^3$	$8\sqrt{2}x^3$
\mathcal{O}_6	$\frac{1}{12}$	ν	$(\Lambda + \frac{1}{4})\mu x$	$\sqrt{2}(\frac{\Lambda}{2} + \frac{3}{4})\mu x$
\mathcal{O}_7	$\frac{1}{12}$	ν^2	$2\mu x^2$	$4\mu x^2$
\mathcal{O}_8	$\frac{1}{12}$	ν	$\mu^2 x$	$\sqrt{2}\mu^2 x$
\mathcal{O}_9	$\frac{1}{12}$	ν^2	$2\mu x^2$	$4\mu x^2$
\mathcal{O}_{10}	$\frac{1}{12}$	ν^2	$2(\Lambda + \frac{1}{4})x^2$	$4(\frac{\Lambda}{2} + \frac{3}{4})x^2$
\mathcal{O}_{11}	$\frac{1}{12}$	ν	$(\Lambda + \frac{1}{4})\mu x$	$\sqrt{2}(\frac{\Lambda}{2} + \frac{3}{4})\mu x$
\mathcal{O}_{12}	$\frac{1}{12}$	ν^2	$2(\Lambda + \frac{1}{4})x^2$	$4(\frac{\Lambda}{2} + \frac{3}{4})x^2$
\mathcal{O}_{13}	$\frac{1}{12}$	$\frac{1}{6}(\nu - 1)\nu(2\nu - 1)$	$4x^3$	$8\sqrt{2}x^3$
\mathcal{O}_{14}	$\frac{1}{12}$	$\frac{1}{2}(\nu - 1)\nu$	$4\mu x^2$	$8\mu x^2$
\mathcal{O}_{15}	$\frac{1}{12}$	$\frac{1}{2}(\nu - 1)\nu^2$	$4x^3$	$8\sqrt{2}x^3$
\mathcal{O}_{16}	$\frac{1}{12}$	$\frac{1}{2}(\nu - 1)\nu$	$2\mu x^2$	$4\mu x^2$
\mathcal{O}_{17}	$\frac{1}{12}$	ν	$\mu^2 x$	$\sqrt{2}\mu^2 x$
\mathcal{O}_{18}	$\frac{1}{12}$	$\frac{1}{2}(\nu - 1)\nu$	$2\mu x^2$	$4\mu x^2$
\mathcal{O}_{19}	$\frac{1}{12}$	$\frac{1}{2}(\nu - 1)\nu^2$	$4x^3$	$8\sqrt{2}x^3$
\mathcal{O}_{20}	$\frac{1}{12}$	ν^2	$4\mu x^2$	$8\mu x^2$
\mathcal{O}_{21}	$\frac{1}{12}$	ν^2	$2(\Lambda + \frac{1}{4})x^2$	$4(\frac{\Lambda}{2} + \frac{3}{4})x^2$
\mathcal{O}_{22}	$\frac{1}{12}$	ν^3	$4x^3$	$8\sqrt{2}x^3$
\mathcal{O}_{23}	$\frac{1}{12}$	ν^2	$2\mu x^2$	$4\mu x^2$
\mathcal{O}_{24}	$\frac{1}{12}$	ν^3	$4x^3$	$8\sqrt{2}x^3$

Table 15: Summary of non-zero commutator upper bounds, including the pure number coefficients from Eq. (140) and multiplicities of each operator form.

sums over the ν distinct subterms, leading to a certain multiplicity associated with each \mathcal{O}_i . These multiplicities are recorded in Table 15. When the $\|\mathcal{O}_i\|$ are summed up with appropriate coefficients and multiplicities, the final result is

$$\|V^{\text{SB}}(\theta) - e^{-i\theta H^{\text{SB}}}\| \leq L\theta^3 \rho^{\text{SB}}(x, \Lambda, \mu), \quad (204)$$

with

$$\rho^{\text{SB}}(x, \Lambda, \mu) \equiv \frac{1658x^3}{3} + 32\Lambda x^2 + \frac{218\mu x^2}{3} + 8x^2 + \frac{\Lambda^2 x}{3} + \frac{4\Lambda\mu x}{3} + \frac{\Lambda x}{6} + \frac{5\mu^2 x}{3} + \frac{\mu x}{3} + \frac{x}{48}. \quad (205)$$

C.0.2 Loop-string-hadron commutator bounds

The derivation of the second-order Trotter error for the LSH formulation closely follows that of the Schwinger-boson formulation. One difference concerns the norms $\|xK^{(j)} \pm xK^{(j)\dagger}\|$. Specifically, one of the generalized LSH hopping subterms is as written in Eq. (81):

$$H_I^{\text{LSH}(j)} = x \chi_{y'}^\dagger \chi_{x'} (\Gamma_q^\dagger)^{n_y} (\Gamma_p^\dagger)^{1-n_x} \mathcal{D}^{\text{LSH}}(p, n_x, n_y) + \text{H.c.}$$

Above, one identifies

$$K^{(j)\dagger} = \chi_{y'}^\dagger \chi_{x'} (\Gamma_q^\dagger)^{n_y} (\Gamma_p^\dagger)^{1-n_x} \mathcal{D}^{\text{LSH}}(p, n_x, n_y), \quad (206a)$$

$$K^{(j)} = \chi_{x'}^\dagger \chi_{y'} (\Gamma_q)^{n_y} (\Gamma_p)^{1-n_x} \mathcal{D}^{\text{LSH}}(p-1+n_x, n_y, n_x), \quad (206b)$$

where $\mathcal{D}^{\text{LSH}}(p, n, n') = \sqrt{(p+1+n)/(p+1+n')}$. The norm bound reduces to the evaluation of $\|\mathcal{D}^{\text{LSH}}\|$ as before since $\|xK^{(j)\dagger} \pm xK^{(j)}\| \leq \max(\|xK^{(j)}\|, \|xK^{(j)\dagger}\|)$. It is easy to see that $\|\mathcal{D}^{\text{LSH}}(p, n, n')\| \leq \sqrt{2}$. This

is because the numerator is larger when $n = 1$ and the denominator is smaller when $n' = 0$. The expression $\sqrt{(p+2)/(p+1)}$ is then largest for $p = 0$, giving the norm of $\sqrt{2}$ as was claimed. Similar reasoning can be used to deduce that $\|\mathcal{D}^{\text{LSH}}(p-1+n', n, n')\| \leq \sqrt{2}$ as well. In either case, one arrives at

$$\|xK^{(j)\dagger} \pm xK^{(j)}\| \leq \sqrt{2} x.$$

As for the commutators involving $H_E^{\text{LSH}}(r)$, one notes that for the LSH formulation, 5

$$J^L(r) = \frac{1}{2}(n_\ell(r) + n_o(r)(1 - n_i(r))), \quad (207)$$

hence

$$\|J^L(r) + \frac{1}{4}\| \leq \frac{\Lambda}{2} + \frac{3}{4}. \quad (208)$$

Fortunately, the forms of non-vanishing commutators in the LSH formulation are the same as in the Schwinger-boson formulation. That is, there are twenty four \mathcal{O}_i expressions from before, although the operator forms of the Hamiltonian terms are different. As the evaluations proceed the same way, rather than repeating the explanations, we simply state what differences arise in the norm evaluations:

1. One should use $\|J^L + 1/4\| \leq \Lambda/2 + 3/4$ for the LSH formulation instead of $\Lambda + 1/4$ for the Schwinger-boson formulation.
2. Wherever a factor of x is acquired in the Schwinger-boson results, one now picks up an additional factor of $\sqrt{2}$.
3. The \mathcal{O}_i have different multiplicities arising from $\nu = 2$ for the LSH formulation, as opposed to $\nu = 8$ for the Schwinger-boson formulation.

Table 15 summarizes the LSH commutator upper bounds. When the $\|\mathcal{O}_i\|$ are summed up with appropriate coefficients and multiplicities, the final result is

$$\|V^{\text{LSH}}(\theta) - e^{-i\theta H^{\text{LSH}}}\| \leq L\theta^3 \rho^{\text{LSH}}(x, \Lambda, \mu), \quad (209)$$

with

$$\rho^{\text{LSH}}(x, \Lambda, \mu) \equiv \frac{47\sqrt{2}x^3}{3} + 2\Lambda x^2 + \frac{25\mu x^2}{3} + 3x^2 + \frac{\Lambda^2 x}{24\sqrt{2}} + \frac{\Lambda\mu x}{3\sqrt{2}} + \frac{\Lambda x}{8\sqrt{2}} + \frac{5\mu^2 x}{6\sqrt{2}} + \frac{\mu x}{2\sqrt{2}} + \frac{3x}{32\sqrt{2}}. \quad (210)$$

D Gate-cost tables

In this appendix, we tabulate the gate counts for near- and far-term approaches to the Trotterized time evolution as detailed in Sec. 3.3, over a range of possible simulation parameters. In the near-term case, we envisage small lattices (no more than 20 staggered sites), harsh electric truncations (not exceeding four qubits per bosonic DOF), total spectral-norm-error bound no smaller than 5%, and bare couplings that are not particularly weak ($x \leq 1$). In the far-term case, we envisage large lattices (up to 1000 staggered sites), modestly larger bosonic registers (up to eight-qubit registers), per-cent and per-mille spectral-norm-error tolerances, as well as weaker bare couplings ($x \geq 1$).

For a comparison between the formulations, one can compute the relative CNOT cost of the LSH formulation to that of the Schwinger-boson formulation. Using the data from Table 16, one observes the relative CNOT cost of the LSH to Schwinger bosons to be most strongly influenced by η , and to a lesser degree by x . For example, the reduction is about 18.4 times at $(\eta, x) = (2, 0.1)$, 24.5 times at $(\eta, x) = (2, 1)$, 317 times at $(\eta, x) = (4, 0.1)$, and 380 times at $(\eta, x) = (4, 1)$. In a similar vein, using Table 17, one observes that the relative T-gate cost of LSH to Schwinger bosons to be most strongly influenced by x , with the reduction being about 21-fold at $x = 1$, and 24-fold at $x = 10$.

References

- [1] Richard P. Feynman. ‘‘Simulating physics with computers’’. *Int. J. Theor. Phys.* **21**, 467–488 (1982). 1
- [2] Seth Lloyd. ‘‘Universal quantum simulators’’. *Science* **273**, 1073–1078 (1996). 1
- [3] John Preskill. ‘‘Quantum computing in the NISQ era and beyond’’. *Quantum* **2**, 79 (2018). arXiv:1801.00862. 1

m/g	Δ_{Trot}	x	L	η	t/a_s	Schwinger bosons			LSH		
						Qubits	Min. s	Min. CNOTs	Qubits	Min. s	Min. CNOTs
1	10%	0.1	10	2	1	92	186	4.8613×10^6	40	63	2.63088×10^5
1	5%	0.1	10	2	1	92	262	6.84763×10^6	40	89	3.71664×10^5
1	10%	1	10	2	1	92	102	2.66587×10^6	40	26	1.08576×10^5
1	5%	1	10	2	1	92	144	3.76358×10^6	40	37	1.54512×10^5
1	10%	0.1	10	4	1	164	433	5.21403×10^8	60	136	1.64261×10^6
1	5%	0.1	10	4	1	164	613	7.38153×10^8	60	193	2.33105×10^6
1	10%	1	10	4	1	164	129	1.55337×10^8	60	34	4.10652×10^5
1	5%	1	10	4	1	164	182	2.19158×10^8	60	48	5.79744×10^5
1	10%	0.1	20	2	1	192	262	1.44561×10^7	80	89	7.84624×10^5
1	5%	0.1	20	2	1	192	371	2.04703×10^7	80	126	1.11082×10^6
1	10%	1	20	2	1	192	144	7.94534×10^6	80	37	3.26192×10^5
1	5%	1	20	2	1	192	203	1.12007×10^7	80	52	4.58432×10^5
1	10%	0.1	20	4	1	344	613	1.55832×10^9	120	193	4.92111×10^6
1	5%	0.1	20	4	1	344	866	2.20148×10^9	120	272	6.93546×10^6
1	10%	1	20	4	1	344	182	4.62667×10^8	120	48	1.2239×10^6
1	5%	1	20	4	1	344	257	6.53326×10^8	120	68	1.73386×10^6
1	10%	0.1	10	2	5	92	2072	5.41538×10^7	40	702	2.93155×10^6
1	5%	0.1	10	2	5	92	2929	7.65523×10^7	40	993	4.14677×10^6
1	10%	1	10	2	5	92	1133	2.96121×10^7	40	288	1.20269×10^6
1	5%	1	10	2	5	92	1602	4.18699×10^7	40	407	1.69963×10^6
1	10%	0.1	10	4	5	164	4841	5.82936×10^9	60	1519	1.83465×10^7
1	5%	0.1	10	4	5	164	6846	8.24371×10^9	60	2149	2.59556×10^7
1	10%	1	10	4	5	164	1432	1.72436×10^9	60	375	4.52925×10^6
1	5%	1	10	4	5	164	2024	2.43723×10^9	60	531	6.41342×10^6
1	10%	0.1	20	2	5	192	2929	1.61611×10^8	80	993	8.75429×10^6
1	5%	0.1	20	2	5	192	4143	2.28594×10^8	80	1404	1.23777×10^7
1	10%	1	20	2	5	192	1602	8.8392×10^7	80	407	3.58811×10^6
1	5%	1	20	2	5	192	2266	1.25029×10^8	80	575	5.0692×10^6
1	10%	0.1	20	4	5	344	6846	1.74034×10^{10}	120	2149	5.47952×10^7
1	5%	0.1	20	4	5	344	9682	2.46128×10^{10}	120	3038	7.74629×10^7
1	10%	1	20	4	5	344	2024	5.14526×10^9	120	531	1.35394×10^7
1	5%	1	20	4	5	344	2863	7.2781×10^9	120	750	1.91235×10^7

Table 16: Near-term simulation costs as a function of Hamiltonian parameters ($m/g, x, L$, and η), evolution time ($t/a_s = 2xT$), and desired bound on the controlled sources of error (Δ_{Trot}). Qubit counts are the register size of the lattice and exclude possible ancilla qubits (which are insignificant in the near-term circuits cost). Other tabulated costs are the minimal required number of second-order Trotter steps (based on the decomposition of Eq. (95b)) and the associated CNOT-gate count (for the naive circuitization approach based on the full Pauli decomposition of diagonal phase functions) in the zero-noise limit.

- [4] Iulia M Georgescu, Sahel Ashhab, and Franco Nori. “Quantum simulation”. *Reviews of Modern Physics* **86**, 153 (2014). [arXiv:1308.6253](#). 1
- [5] Dave Wecker, Matthew B Hastings, Nathan Wiebe, Bryan K Clark, Chetan Nayak, and Matthias Troyer. “Solving strongly correlated electron models on a quantum computer”. *Physical Review A* **92**, 062318 (2015). [arXiv:1506.05135](#). 1
- [6] Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C Benjamin, and Xiao Yuan. “Quantum computational chemistry”. *Reviews of Modern Physics* **92**, 015003 (2020). [arXiv:1808.10402](#).
- [7] Yudong Cao, Jonathan Romero, Jonathan P Olson, Matthias Degroote, Peter D Johnson, Mária Kieferová, Ian D Kivlichan, Tim Menke, Borja Peropadre, Nicolas PD Sawaya, et al. “Quantum chemistry in the age of quantum computing”. *Chemical Reviews* **119**, 10856–10915 (2019). [arXiv:1812.09976](#).
- [8] Ryan Babbush, Nathan Wiebe, Jarrod McClean, James McClain, Hartmut Neven, and Garnet Kin-Lic Chan. “Low-depth quantum simulation of materials”. *Physical Review X* **8**, 011044 (2018). [arXiv:1706.00023](#).
- [9] Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet Kin-Lic Chan. “Quantum algorithms for

m/g	x	η	L	t/a_s	Δ	$\alpha_{\text{Trot.}}$	$\alpha_{\text{Newt.}}$	Schwinger bosons		LSH	
								Qubits	T gates	Qubits	T gates
1	1	4	100	1	0.01	90%	9%	2626	8.19713×10^{11}	1319	3.91817×10^{10}
1	1	4	100	1	0.001	90%	9%	2704	3.09951×10^{12}	1397	1.5172×10^{11}
1	10	4	100	1	0.01	90%	9%	2626	5.7715×10^{11}	1319	2.31055×10^{10}
1	10	4	100	1	0.001	90%	9%	2704	2.18285×10^{12}	1397	8.94211×10^{10}
1	1	8	100	1	0.01	90%	9%	4398	5.79224×10^{12}	1807	2.72735×10^{11}
1	1	8	100	1	0.001	90%	9%	4476	2.1482×10^{13}	1885	1.03709×10^{12}
1	10	8	100	1	0.01	90%	9%	4398	1.33288×10^{12}	1807	5.40046×10^{10}
1	10	8	100	1	0.001	90%	9%	4476	4.94102×10^{12}	1885	2.05007×10^{11}
1	1	4	1000	1	0.01	90%	9%	18904	3.12769×10^{13}	6797	1.53099×10^{12}
1	1	4	1000	1	0.001	90%	9%	19008	1.22564×10^{14}	6875	5.81562×10^{12}
1	10	4	1000	1	0.01	90%	9%	18904	2.2027×10^{13}	6797	9.02342×10^{11}
1	10	4	1000	1	0.001	90%	9%	19008	8.63137×10^{13}	6875	3.42924×10^{12}
1	1	8	1000	1	0.01	90%	9%	35076	2.16773×10^{14}	10885	1.04652×10^{13}
1	1	8	1000	1	0.001	90%	9%	35180	8.30098×10^{14}	10963	3.91414×10^{13}
1	10	8	1000	1	0.01	90%	9%	35076	4.98595×10^{13}	10885	2.06871×10^{12}
1	10	8	1000	1	0.001	90%	9%	35180	1.9092×10^{14}	10963	7.73624×10^{12}
1	1	4	100	10	0.01	90%	9%	2704	3.0993×10^{13}	1397	1.51643×10^{12}
1	1	4	100	10	0.001	90%	9%	2808	1.2146×10^{14}	1475	5.76229×10^{12}
1	10	4	100	10	0.01	90%	9%	2704	2.18258×10^{13}	1397	8.94083×10^{11}
1	10	4	100	10	0.001	90%	9%	2808	8.55326×10^{13}	1475	3.39741×10^{12}
1	1	8	100	10	0.01	90%	9%	4476	2.14816×10^{14}	1885	1.03705×10^{13}
1	1	8	100	10	0.001	90%	9%	4580	8.22615×10^{14}	1963	3.87886×10^{13}
1	10	8	100	10	0.01	90%	9%	4476	4.94053×10^{13}	1885	2.04958×10^{12}
1	10	8	100	10	0.001	90%	9%	4580	1.89192×10^{14}	1963	7.66615×10^{12}
1	1	4	1000	10	0.01	90%	9%	19008	1.22564×10^{15}	6875	5.81468×10^{13}
1	1	4	1000	10	0.001	90%	9%	19086	4.48657×10^{15}	6979	2.29217×10^{14}
1	10	4	1000	10	0.01	90%	9%	19008	8.63103×10^{14}	6875	3.4283×10^{13}
1	10	4	1000	10	0.001	90%	9%	19086	3.15948×10^{15}	6979	1.35149×10^{14}
1	1	8	1000	10	0.01	90%	9%	35180	8.30094×10^{15}	10963	3.91412×10^{14}
1	1	8	1000	10	0.001	90%	9%	35258	2.99214×10^{16}	11067	1.5154×10^{15}
1	10	8	1000	10	0.01	90%	9%	35180	1.90912×10^{15}	10963	7.73585×10^{13}
1	10	8	1000	10	0.001	90%	9%	35258	6.88164×10^{15}	11067	2.99513×10^{14}

Table 17: Far-term simulation costs as a function of Hamiltonian parameters ($m/g, x, L$, and $\Lambda = 2^n - 1$), evolution time (t/a_s), and desired bound on the controlled sources of error (Δ). Qubit counts are the sum of qubits needed to represent lattice DOFs and ancilla qubits used for implementing the time evolution.

quantum chemistry and quantum materials science”. *Chemical Reviews* **120**, 12685–12717 (2020). [arXiv:2001.03685](https://arxiv.org/abs/2001.03685).

- [10] Vera von Burg, Guang Hao Low, Thomas Häner, Damian S Steiger, Markus Reiher, Martin Roetteler, and Matthias Troyer. “Quantum computing enhanced computational catalysis”. *Physical Review Research* **3**, 033055 (2021). [arXiv:2007.14460](https://arxiv.org/abs/2007.14460).
- [11] He Ma, Marco Govoni, and Giulia Galli. “Quantum simulations of materials on near-term quantum computers”. *npj Computat. Mater.* **6**, 85 (2020). 1
- [12] Matthew Dietrich, David Hertzog, Martin J. Savage, et al. “Nuclear Physics and Quantum Information Science: Report by the NSAC QIS Subcommittee”. Technical Report NSAC-QIS-2019. NSF & DOE Office of Science (2019). url: https://science.osti.gov/-/media/np/pdf/Reports/NSAC_QIS_Report.pdf. 1
- [13] Christian W. Bauer et al. “Quantum Simulation for High-Energy Physics”. *PRX Quantum* **4**, 027001 (2023). [arXiv:2204.03381](https://arxiv.org/abs/2204.03381). 1
- [14] Simon Catterall et al. “Report of the snowmass 2021 theory frontier topical group on quantum information science”. In *Snowmass 2021*. (2022). [arXiv:2209.14839](https://arxiv.org/abs/2209.14839).
- [15] Travis S. Humble, Gabriel N. Perdue, and Martin J. Savage. “Snowmass computational frontier: Topical group report on quantum computing” (2022). [arXiv:2209.06786](https://arxiv.org/abs/2209.06786). 1

- [16] Tim Byrnes and Yoshihisa Yamamoto. “Simulating lattice gauge theories on a quantum computer”. *Phys. Rev. A* **73**, 022328 (2006). [arXiv:quant-ph/0510027](#). 1
- [17] Stephen P. Jordan, Keith S.M. Lee, and John Preskill. “Quantum Algorithms for Quantum Field Theories”. *Science* **336**, 1130–1133 (2012). [arXiv:1111.3633](#). 1
- [18] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. “Quantum Computation of Scattering in Scalar Quantum Field Theories”. *Quant. Inf. Comput.* **14**, 1014–1080 (2014). [arXiv:1112.4833](#).
- [19] Erez Zohar and Benni Reznik. “Confinement and lattice QED electric flux-tubes simulated with ultracold atoms”. *Phys. Rev. Lett.* **107**, 275301 (2011). [arXiv:1108.1562](#).
- [20] L. Tagliacozzo, A. Celi, A. Zamora, and M. Lewenstein. “Optical Abelian Lattice Gauge Theories”. *Annals Phys.* **330**, 160–191 (2013). [arXiv:1205.0496](#).
- [21] D. Banerjee, M. Dalmonte, M. Muller, E. Rico, P. Stebler, U.-J. Wiese, and P. Zoller. “Atomic Quantum Simulation of Dynamical Gauge Fields coupled to Fermionic Matter: From String Breaking to Evolution after a Quench”. *Phys. Rev. Lett.* **109**, 175302 (2012). [arXiv:1205.6366](#).
- [22] Erez Zohar, J. Ignacio Cirac, and Benni Reznik. “Cold-Atom Quantum Simulator for $SU(2)$ Yang-Mills Lattice Gauge Theory”. *Phys. Rev. Lett.* **110**, 125304 (2013). [arXiv:1211.2241](#).
- [23] Erez Zohar, J. Ignacio Cirac, and Benni Reznik. “Quantum simulations of gauge theories with ultracold atoms: local gauge invariance from angular momentum conservation”. *Phys. Rev. A* **88**, 023617 (2013). [arXiv:1303.5040](#).
- [24] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. “Quantum Algorithms for Fermionic Quantum Field Theories” (2014). [arXiv:1404.7115](#).
- [25] Erez Zohar and Michele Burrello. “Formulation of lattice gauge theories for quantum simulations”. *Phys. Rev. D* **91**, 054506 (2015). [arXiv:1409.3085](#). 3.1
- [26] Kevin Marshall, Raphael Pooser, George Siopsis, and Christian Weedbrook. “Quantum simulation of quantum field theory using continuous variables”. *Phys. Rev. A* **92**, 063825 (2015). [arXiv:1503.08121](#).
- [27] A. Mezzacapo, E. Rico, C. Sabin, I.L. Egusquiza, L. Lamata, and E. Solano. “Non-Abelian $SU(2)$ Lattice Gauge Theories in Superconducting Circuits”. *Phys. Rev. Lett.* **115**, 240502 (2015). [arXiv:1505.04720](#).
- [28] E.A. Martinez et al. “Real-time dynamics of lattice gauge theories with a few-qubit quantum computer”. *Nature* **534**, 516–519 (2016). [arXiv:1605.04570](#).
- [29] Erez Zohar, Alessandro Farace, Benni Reznik, and J. Ignacio Cirac. “Digital quantum simulation of \mathbb{Z}_2 lattice gauge theories with dynamical fermionic matter”. *Phys. Rev. Lett.* **118**, 070501 (2017). [arXiv:1607.03656](#).
- [30] Erez Zohar, Alessandro Farace, Benni Reznik, and J. Ignacio Cirac. “Digital lattice gauge theories”. *Phys. Rev. A* **95**, 023604 (2017). [arXiv:1607.08121](#).
- [31] Ali Hamed Moosavian and Stephen Jordan. “Faster Quantum Algorithm to Simulate Fermionic Quantum Field Theory”. *Phys. Rev. A* **98**, 012332 (2018). [arXiv:1711.04006](#).
- [32] T.V. Zache, F. Hebenstreit, F. Jendrzejewski, M.K. Oberthaler, J. Berges, and P. Hauke. “Quantum simulation of lattice gauge theories using Wilson fermions”. *Sci. Technol.* **3**, 034010 (2018). [arXiv:1802.06704](#).
- [33] Frederik Görg, Kilian Sandholzer, Joaquín Minguzzi, Rémi Desbuquois, Michael Messer, and Tilman Esslinger. “Realization of density-dependent Peierls phases to engineer quantized gauge fields coupled to ultracold matter”. *Nature Phys.* **15**, 1161–1167 (2019). [arXiv:1812.05895](#).
- [34] Christian Schweizer, Fabian Grusdt, Moritz Berngruber, Luca Barbiero, Eugene Demler, Nathan Goldman, Immanuel Bloch, and Monika Aidelsburger. “Floquet approach to \mathbb{Z}_2 lattice gauge theories with ultracold atoms in optical lattices”. *Nature Physics* **15**, 1168–1173 (2019). [arXiv:1901.07103](#).
- [35] N. Klco, E.F. Dumitrescu, A.J. McCaskey, T.D. Morris, R.C. Pooser, M. Sanz, E. Solano, P. Lougovski, and M.J. Savage. “Quantum-classical computation of Schwinger model dynamics using quantum computers”. *Phys. Rev. A* **98**, 032331 (2018). [arXiv:1803.03326](#).
- [36] Hsuan-Hao Lu et al. “Simulations of Subatomic Many-Body Physics on a Quantum Frequency Processor”. *Phys. Rev. A* **100**, 012320 (2019). [arXiv:1810.03959](#). 1
- [37] Arpan Bhattacharyya, Arvind Shekar, and Aninda Sinha. “Circuit complexity in interacting QFTs and RG flows”. *JHEP* **10**, 140 (2018). [arXiv:1808.03105](#).
- [38] Jesse R. Stryker. “Oracles for Gauss’s law on digital quantum computers”. *Phys. Rev. A* **99**, 042301 (2019). [arXiv:1812.01617](#). 1, 3.1
- [39] Indrakshi Raychowdhury and Jesse R. Stryker. “Solving Gauss’s Law on Digital Quantum Computers with Loop-String-Hadron Digitization”. *Phys. Rev. Res.* **2**, 033039 (2020). [arXiv:1812.07554](#). 1
- [40] Di Luo, Jiayu Shen, Michael Highman, Bryan K. Clark, Brian DeMarco, Aida X. El-Khadra, and Bryce Gadway. “Framework for simulating gauge theories with dipolar spin systems”. *Phys. Rev. A* **102**, 032617 (2020). [arXiv:1912.11488](#).

- [41] Federica M. Surace, Paolo P. Mazza, Giuliano Giudici, Alessio Lerose, Andrea Gambassi, and Marcello Dalmonte. “Lattice gauge theories and string dynamics in Rydberg atom quantum simulators”. *Phys. Rev. X* **10**, 021041 (2020). [arXiv:1902.09551](#).
- [42] Alexander Mil, Torsten V. Zache, Apoorva Hegde, Andy Xia, Rohit P. Bhatt, Markus K. Oberthaler, Philipp Hauke, Jürgen Berges, and Fred Jendrzejewski. “A scalable realization of local U(1) gauge invariance in cold atomic mixtures”. *Science* **367**, 1128–1130 (2020). [arXiv:1909.07641](#).
- [43] Natalie Klco, Jesse R. Stryker, and Martin J. Savage. “SU(2) non-Abelian gauge field theory in one dimension on digital quantum computers”. *Phys. Rev. D* **101**, 074512 (2020). [arXiv:1908.06935](#). 3
- [44] Natalie Klco and Martin J. Savage. “Digitization of scalar fields for quantum computing”. *Phys. Rev. A* **99**, 052335 (2019). [arXiv:1808.10378](#).
- [45] Christian W. Bauer, Wibe A. de Jong, Benjamin Nachman, and Davide Provasoli. “Quantum Algorithm for High Energy Physics Simulations”. *Phys. Rev. Lett.* **126**, 062001 (2021). [arXiv:1904.03196](#). 1
- [46] Zohreh Davoudi, Mohammad Hafezi, Christopher Monroe, Guido Pagano, Alireza Seif, and Andrew Shaw. “Towards analog quantum simulations of lattice gauge theories with trapped ions”. *Phys. Rev. Res.* **2**, 023015 (2020). [arXiv:1908.03210](#).
- [47] Natalie Klco and Martin J. Savage. “Systematically Localizable Operators for Quantum Simulations of Quantum Field Theories”. *Phys. Rev. A* **102**, 012619 (2020). [arXiv:1912.03577](#).
- [48] Henry Lamm, Scott Lawrence, and Yukari Yamauchi. “Parton physics on a quantum computer”. *Phys. Rev. Res.* **2**, 013272 (2020). [arXiv:1908.10439](#).
- [49] Niklas Mueller, Andrey Tarasov, and Raju Venugopalan. “Deeply inelastic scattering structure functions on a hybrid quantum computer”. *Phys. Rev. D* **102**, 016007 (2020). [arXiv:1908.07051](#).
- [50] Henry Lamm, Scott Lawrence, and Yukari Yamauchi. “General Methods for Digital Quantum Simulation of Gauge Theories”. *Phys. Rev. D* **100**, 034518 (2019). [arXiv:1903.08807](#). 4
- [51] Andrei Alexandru, Paulo F. Bedaque, Siddhartha Harmalkar, Henry Lamm, Scott Lawrence, and Neill C. Warrington. “Gluon Field Digitization for Quantum Computers”. *Phys. Rev. D* **100**, 114501 (2019). [arXiv:1906.11213](#). 4
- [52] Natalie Klco and Martin J. Savage. “Fixed-point quantum circuits for quantum field theories”. *Phys. Rev. A* **102**, 052422 (2020). [arXiv:2002.02018](#).
- [53] Bing Yang, Hui Sun, Robert Ott, Han-Yi Wang, Torsten V. Zache, Jad C. Halimeh, Zhen-Sheng Yuan, Philipp Hauke, and Jian-Wei Pan. “Observation of gauge invariance in a 71-site Bose–Hubbard quantum simulator”. *Nature* **587**, 392–396 (2020). [arXiv:2003.08945](#).
- [54] Alexander F. Shaw, Pavel Lougovski, Jesse R. Stryker, and Nathan Wiebe. “Quantum Algorithms for Simulating the Lattice Schwinger Model”. *Quantum* **4**, 306 (2020). [arXiv:2002.11146](#). 1, 2, 2.1, 2.2.1, 3, 3.2.1, 11, 3.2.2, 4, 14
- [55] Bipasha Chakraborty, Masazumi Honda, Taku Izubuchi, Yuta Kikuchi, and Akio Tomiya. “Classically emulated digital quantum simulation of the Schwinger model with a topological term via adiabatic state preparation”. *Phys. Rev. D* **105**, 094503 (2022). [arXiv:2001.00485](#).
- [56] Junyu Liu and Yuan Xin. “Quantum simulation of quantum field theories as quantum chemistry”. *JHEP* **12**, 011 (2020). [arXiv:2004.13234](#).
- [57] Michael Kreshchuk, William M. Kirby, Gary Goldstein, Hugo Beauchemin, and Peter J. Love. “Quantum simulation of quantum field theory in the light-front formulation”. *Phys. Rev. A* **105**, 032418 (2022). [arXiv:2002.04016](#).
- [58] Jan F. Haase, Luca Dellantonio, Alessio Celi, Danny Paulson, Angus Kan, Karl Jansen, and Christine A. Muschik. “A resource efficient approach for quantum and classical simulations of gauge theories in particle physics”. *Quantum* **5**, 393 (2021). [arXiv:2006.14160](#). 2.1
- [59] Danny Paulson et al. “Towards simulating 2D effects in lattice gauge theories on a quantum computer”. *PRX Quantum* **2**, 030334 (2021). [arXiv:2008.09252](#).
- [60] Raka Dasgupta and Indrakshi Raychowdhury. “Cold-atom quantum simulator for string and hadron dynamics in non-Abelian lattice gauge theory”. *Phys. Rev. A* **105**, 023322 (2022). [arXiv:2009.13969](#).
- [61] Simon V. Mathis, Guglielmo Mazzola, and Ivano Tavernelli. “Toward scalable simulations of lattice gauge theories on quantum computers”. *Phys. Rev. D* **102**, 094501 (2020). [arXiv:2005.10271](#).
- [62] Yasar Y. Atas, Jinglei Zhang, Randy Lewis, Amin Jahanpour, Jan F. Haase, and Christine A. Muschik. “SU(2) hadrons on a quantum computer via a variational approach”. *Nature Commun.* **12**, 6499 (2021). [arXiv:2102.08920](#). 3
- [63] Sarmed A Rahman, Randy Lewis, Emanuele Mendicelli, and Sarah Powell. “SU(2) lattice gauge theory on a quantum annealer”. *Phys. Rev. D* **104**, 034501 (2021). [arXiv:2103.08661](#). 3

- [64] Zohreh Davoudi, Norbert M. Linke, and Guido Pagano. “Toward simulating quantum field theories with controlled phonon-ion dynamics: A hybrid analog-digital approach”. *Phys. Rev. Res.* **3**, 043072 (2021). [arXiv:2104.09346](#). 4
- [65] João Barata, Niklas Mueller, Andrey Tarasov, and Raju Venugopalan. “Single-particle digitization strategy for quantum computation of a ϕ^4 scalar field theory”. *Phys. Rev. A* **103**, 042410 (2021). [arXiv:2012.00020](#).
- [66] Wibe A. de Jong, Kyle Lee, James Mulligan, Mateusz Płoskoń, Felix Ringer, and Xiaojun Yao. “Quantum simulation of nonequilibrium dynamics and thermalization in the Schwinger model”. *Phys. Rev. D* **106**, 054508 (2022). [arXiv:2106.08394](#).
- [67] Anthony N. Ciavarella and Ivan A. Chernyshev. “Preparation of the SU(3) lattice Yang-Mills vacuum with variational quantum methods”. *Phys. Rev. D* **105**, 074504 (2022). [arXiv:2112.09083](#).
- [68] 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”. *Phys. Rev. D* **103**, 094501 (2021). [arXiv:2101.10227](#). 1, 3.1, 4
- [69] Angus Kan and Yunseong Nam. “Lattice Quantum Chromodynamics and Electrodynamics on a Universal Quantum Computer” (2021). [arXiv:2107.12769](#). 1, 2.1, 3, 3.2, 15, 4
- [70] Thomas D. Cohen, Henry Lamm, Scott Lawrence, and Yukari Yamauchi. “Quantum algorithms for transport coefficients in gauge theories”. *Phys. Rev. D* **104**, 094514 (2021). [arXiv:2104.02024](#).
- [71] Bárbara Andrade, Zohreh Davoudi, Tobias Graß, Mohammad Hafezi, Guido Pagano, and Alireza Seif. “Engineering an effective three-spin Hamiltonian in trapped-ion systems for applications in quantum simulation”. *Quantum Sci. Technol.* **7**, 034001 (2022). [arXiv:2108.01022](#).
- [72] M. Sohaib Alam, Stuart Hadfield, Henry Lamm, and Andy C. Y. Li. “Primitive quantum gates for dihedral gauge theories”. *Phys. Rev. D* **105**, 114501 (2022). [arXiv:2108.13305](#).
- [73] Nhung H. Nguyen, Minh C. Tran, Yingyue Zhu, Alaina M. Green, C. Huerta Alderete, Zohreh Davoudi, and Norbert M. Linke. “Digital Quantum Simulation of the Schwinger Model and Symmetry Protection with Trapped Ions”. *PRX Quantum* **3**, 020324 (2022). [arXiv:2112.14262](#).
- [74] Jinglei Zhang, Ryan Ferguson, Stefan Kühn, Jan F. Haase, C. M. Wilson, Karl Jansen, and Christine A. Muschik. “Simulating gauge theories with variational quantum eigensolvers in superconducting microwave cavities”. *Quantum* **7**, 1148 (2023). [arXiv:2108.08248](#).
- [75] Masazumi Honda, Etsuko Itou, Yuta Kikuchi, Lento Nagano, and Takuya Okuda. “Classically emulated digital quantum simulation for screening and confinement in the Schwinger model with a topological term”. *Phys. Rev. D* **105**, 014504 (2022). [arXiv:2105.03276](#).
- [76] Zhao-Yu Zhou, Guo-Xian Su, Jad C. Halimeh, Robert Ott, Hui Sun, Philipp Hauke, Bing Yang, Zhen-Sheng Yuan, Jürgen Berges, and Jian-Wei Pan. “Thermalization dynamics of a gauge theory on a quantum simulator”. *Science* **377**, 311–314 (2022). [arXiv:2107.13563](#).
- [77] Daniel González-Cuadra, Torsten V. Zache, Jose Carrasco, Barbara Kraus, and Peter Zoller. “Hardware Efficient Quantum Simulation of Non-Abelian Gauge Theories with Qudits on Rydberg Platforms”. *Phys. Rev. Lett.* **129**, 160501 (2022). [arXiv:2203.15541](#).
- [78] Jesse Osborne, Ian P. McCulloch, Bing Yang, Philipp Hauke, and Jad C. Halimeh. “Large-Scale 2+1D U(1) Gauge Theory with Dynamical Matter in a Cold-Atom Quantum Simulator” (2022). [arXiv:2211.01380](#).
- [79] Zohreh Davoudi, Niklas Mueller, and Connor Powers. “Towards Quantum Computing Phase Diagrams of Gauge Theories with Thermal Pure Quantum States”. *Phys. Rev. Lett.* **131**, 081901 (2023). [arXiv:2208.13112](#).
- [80] Niklas Mueller, Joseph A. Carolan, Andrew Connelly, Zohreh Davoudi, Eugene F. Dumitrescu, and Kübra Yeter-Aydeniz. “Quantum Computation of Dynamical Quantum Phase Transitions and Entanglement Tomography in a Lattice Gauge Theory”. *PRX Quantum* **4**, 030323 (2023). [arXiv:2210.03089](#).
- [81] Edison M. Murairi, Michael J. Cervia, Hersh Kumar, Paulo F. Bedaque, and Andrei Alexandru. “How many quantum gates do gauge theories require?”. *Phys. Rev. D* **106**, 094504 (2022). [arXiv:2208.11789](#).
- [82] Roland C. Farrell, Ivan A. Chernyshev, Sarah J. M. Powell, Nikita A. Zemlevskiy, Marc Illa, and Martin J. Savage. “Preparations for quantum simulations of quantum chromodynamics in 1+1 dimensions. I. Axial gauge”. *Phys. Rev. D* **107**, 054512 (2023). [arXiv:2207.01731](#).
- [83] Roland C. Farrell, Ivan A. Chernyshev, Sarah J. M. Powell, Nikita A. Zemlevskiy, Marc Illa, and Martin J. Savage. “Preparations for quantum simulations of quantum chromodynamics in 1+1 dimensions. II. Single-baryon β -decay in real time”. *Phys. Rev. D* **107**, 054513 (2023). [arXiv:2209.10781](#).
- [84] Giuseppe Clemente, Arianna Crippa, and Karl Jansen. “Strategies for the determination of the running coupling of (2+1)-dimensional QED with quantum computing”. *Phys. Rev. D* **106**, 114511 (2022). [arXiv:2206.12454](#).

- [85] Guy Pardo, Tomer Greenberg, Aryeh Fortinsky, Nadav Katz, and Erez Zohar. “Resource-efficient quantum simulation of lattice gauge theories in arbitrary dimensions: Solving for Gauss’s law and fermion elimination”. *Phys. Rev. Res.* **5**, 023077 (2023). [arXiv:2206.00685](#).
- [86] M.C. Banuls et al. “Simulating Lattice Gauge Theories within Quantum Technologies”. *Eur. Phys. J. D* **74**, 165 (2020). [arXiv:1911.00003](#).
- [87] Natalie Klco, Alessandro Roggero, and Martin J. Savage. “Standard model physics and the digital quantum revolution: thoughts about the interface”. *Rept. Prog. Phys.* **85**, 064301 (2022). [arXiv:2107.04769](#).
- [88] Erez Zohar. “Quantum simulation of lattice gauge theories in more than one space dimension—requirements, challenges and methods”. *Phil. Trans. A. Math. Phys. Eng. Sci.* **380**, 20210069 (2021). [arXiv:2106.04609](#). 1
- [89] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski. “Cloud Quantum Computing of an Atomic Nucleus”. *Phys. Rev. Lett.* **120**, 210501 (2018). [arXiv:1801.03897](#). 1
- [90] Omar Shehab, Kevin A. Landsman, Yunseong Nam, Daiwei Zhu, Norbert M. Linke, Matthew J. Keesan, Raphael C. Pooser, and Christopher R. Monroe. “Toward convergence of effective field theory simulations on digital quantum computers”. *Phys. Rev. A* **100**, 062319 (2019). [arXiv:1904.04338](#).
- [91] Alessandro Roggero and Joseph Carlson. “Dynamic linear response quantum algorithm”. *Phys. Rev. C* **100**, 034610 (2019). [arXiv:1804.01505](#).
- [92] Alessandro Roggero, Andy C. Y. Li, Joseph Carlson, Rajan Gupta, and Gabriel N. Perdue. “Quantum Computing for Neutrino-Nucleus Scattering”. *Phys. Rev. D* **101**, 074038 (2020). [arXiv:1911.06368](#).
- [93] Weijie Du, James P. Vary, Xingbo Zhao, and Wei Zuo. “Quantum simulation of nuclear inelastic scattering”. *Phys. Rev. A* **104**, 012611 (2021). [arXiv:2006.01369](#).
- [94] Weijie Du, James P. Vary, Xingbo Zhao, and Wei Zuo. “Ab initio nuclear structure via quantum adiabatic algorithm” (2021). [arXiv:2105.08910](#).
- [95] Alessandro Roggero, Chenyi Gu, Alessandro Baroni, and Thomas Papenbrock. “Preparation of excited states for nuclear dynamics on a quantum computer”. *Phys. Rev. C* **102**, 064624 (2020). [arXiv:2009.13485](#).
- [96] Eric T. Holland, Kyle A. Wendt, Konstantinos Kravvaris, Xian Wu, W. Erich Ormand, Jonathan L DuBois, Sofia Quaglioni, and Francesco Pederiva. “Optimal Control for the Quantum Simulation of Nuclear Dynamics”. *Phys. Rev. A* **101**, 062307 (2020). [arXiv:1908.08222](#).
- [97] Dmitri E. Kharzeev and Yuta Kikuchi. “Real-time chiral dynamics from a digital quantum simulation”. *Phys. Rev. Res.* **2**, 023342 (2020). [arXiv:2001.00698](#).
- [98] Michael Kreshchuk, Shaoyang Jia, William M. Kirby, Gary Goldstein, James P. Vary, and Peter J. Love. “Simulating Hadronic Physics on NISQ devices using Basis Light-Front Quantization”. *Phys. Rev. A* **103**, 062601 (2021). [arXiv:2011.13443](#).
- [99] Khadeejah Bepari, Sarah Malik, Michael Spannowsky, and Simon Williams. “Towards a quantum computing algorithm for helicity amplitudes and parton showers”. *Phys. Rev. D* **103**, 076020 (2021). [arXiv:2010.00046](#).
- [100] Christian W. Bauer, Marat Freytsis, and Benjamin Nachman. “Simulating Collider Physics on Quantum Computers Using Effective Field Theories”. *Phys. Rev. Lett.* **127**, 212001 (2021). [arXiv:2102.05044](#). 1
- [101] Andrew M Childs and Yuan Su. “Nearly optimal lattice simulation by product formulas”. *Physical review letters* **123**, 050503 (2019). [arXiv:1901.00564](#). 1, 1
- [102] Masuo Suzuki. “General theory of fractal path integrals with applications to many-body theories and statistical physics”. *Journal of Mathematical Physics* **32**, 400–407 (1991). 1
- [103] Nathan Wiebe, Dominic Berry, Peter Hoyer, and Barry C Sanders. “Higher order decompositions of ordered operator exponentials”. *Journal of Physics A: Mathematical and Theoretical* **43**, 065203 (2010). [arXiv:0812.0562](#).
- [104] Andrew M Childs, Yuan Su, Minh C Tran, Nathan Wiebe, and Shuchen Zhu. “Theory of Trotter error with commutator scaling”. *Physical Review X* **11**, 011020 (2021). [arXiv:1912.08854](#). 1, 3.3, C
- [105] Andrew M Childs and Nathan Wiebe. “Hamiltonian simulation using linear combinations of unitary operations”. *Quantum Information and Computation* **12**, 901–921 (2012). [arXiv:1202.5822](#). 1
- [106] Dominic W Berry, Andrew M Childs, Richard Cleve, Robin Kothari, and Rolando D Somma. “Simulating Hamiltonian dynamics with a truncated Taylor series”. *Physical Review Letters* **114**, 090502 (2015). [arXiv:1412.4687](#). 1
- [107] Guang Hao Low and Isaac L. Chuang. “Optimal Hamiltonian Simulation by Quantum Signal Processing”. *Phys. Rev. Lett.* **118**, 010501 (2017). [arXiv:1606.02685](#). 1
- [108] Guang Hao Low and Isaac L Chuang. “Hamiltonian simulation by qubitization”. *Quantum* **3**, 163 (2019). [arXiv:1610.06546](#). 1

- [109] Shantanav Chakraborty, András Gilyén, and Stacey Jeffery. “The power of block-encoded matrix powers: improved regression techniques via faster Hamiltonian simulation”. *Leibniz International Proceedings in Informatics (LIPIcs)* **132**, 33:1–33:14 (2019). [arXiv:1804.01973](#). 1
- [110] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. “Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics”. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. Page 193–204. New York, NY, USA (2019). Association for Computing Machinery. [arXiv:1806.01838](#). 1
- [111] Amir Kalev and Itay Hen. “Quantum algorithm for simulating hamiltonian dynamics with an off-diagonal series expansion”. *Quantum* **5**, 426 (2021). [arXiv:2006.02539](#). 1
- [112] Abhishek Rajput, Alessandro Roggero, and Nathan Wiebe. “Hybridized Methods for Quantum Simulation in the Interaction Picture”. *Quantum* **6**, 780 (2022). [arXiv:2109.03308](#). 1, 4
- [113] Torin F. Stetina, Anthony Ciavarella, Xiaosong Li, and Nathan Wiebe. “Simulating Effective QED on Quantum Computers”. *Quantum* **6**, 622 (2022). [arXiv:2101.00111](#). 1
- [114] Johann Ostmeier. “Optimised Trotter decompositions for classical and quantum computing”. *J. Phys. A* **56**, 285303 (2023). [arXiv:2211.02691](#). 1
- [115] Peter W Shor. “Fault-tolerant quantum computation”. In *Proceedings of 37th Conference on Foundations of Computer Science*. Pages 56–65. IEEE (1996). [arXiv:quant-ph/9605011](#). 1
- [116] Jesse R. Stryker. “Shearing approach to gauge invariant Trotterization” (2021). [arXiv:2105.11548](#). 1, 2.1
- [117] Andrew M Childs and Wim Van Dam. “Quantum algorithms for algebraic problems”. *Reviews of Modern Physics* **82**, 1 (2010). [arXiv:0812.0380](#). 1
- [118] Thomas Häner, Martin Roetteler, and Krysta M. Svore. “Optimizing Quantum Circuits for Arithmetic” (2018). [arXiv:1805.12445](#). 1, 16, A
- [119] Thomas Haener, Mathias Soeken, Martin Roetteler, and Krysta M Svore. “Quantum circuits for floating-point arithmetic”. In *International Conference on Reversible Computation*. Pages 162–174. Springer (2018). [arXiv:1807.02023](#). 1
- [120] Ian D Kivlichan, Nathan Wiebe, Ryan Babbush, and Alán Aspuru-Guzik. “Bounding the costs of quantum simulation of many-body physics in real space”. *Journal of Physics A: Mathematical and Theoretical* **50**, 305301 (2017). [arXiv:1608.05696](#). 1
- [121] Yuan Su, Dominic W. Berry, Nathan Wiebe, Nicholas Rubin, and Ryan Babbush. “Fault-Tolerant Quantum Simulations of Chemistry in First Quantization”. *PRX Quantum* **2**, 040332 (2021). [arXiv:2105.12767](#).
- [122] Ryan Babbush, Dominic W Berry, Ian D Kivlichan, Annie Y Wei, Peter J Love, and Alán Aspuru-Guzik. “Exponentially more precise quantum simulation of fermions in second quantization”. *New Journal of Physics* **18**, 033032 (2016). [arXiv:1506.01020](#).
- [123] Poul Jorgensen. “Second quantization-based methods in quantum chemistry”. Elsevier. (2012).
- [124] Nikolaž Moll, Andreas Fuhrer, Peter Staar, and Ivano Tavernelli. “Optimizing qubit resources for quantum chemistry simulations in second quantization on a quantum computer”. *Journal of Physics A: Mathematical and Theoretical* **49**, 295301 (2016). [arXiv:1510.04048](#).
- [125] Ryan Babbush, Dominic W Berry, Yuval R Sanders, Ian D Kivlichan, Artur Scherer, Annie Y Wei, Peter J Love, and Alán Aspuru-Guzik. “Exponentially more precise quantum simulation of fermions in the configuration interaction representation”. *Quantum Science and Technology* **3**, 015006 (2017). [arXiv:1506.01029](#). 1
- [126] John B. Kogut and Leonard Susskind. “Hamiltonian Formulation of Wilson’s Lattice Gauge Theories”. *Phys. Rev. D* **11**, 395–408 (1975). 1, 2.1, 3, 3.1
- [127] J. Schwinger. “On angular momentum”. *Technical report*. Harvard University (1952). 1
- [128] Manu Mathur. “Harmonic oscillator prepotentials in SU(2) lattice gauge theory”. *J. Phys. A* **38**, 10015–10026 (2005). [arXiv:hep-lat/0403029](#).
- [129] Ramesh Anishetty, Manu Mathur, and Indrakshi Raychowdhury. “Irreducible SU(3) Schwinger Bosons”. *J. Math. Phys.* **50**, 053503 (2009). [arXiv:0901.0644](#). 3
- [130] Manu Mathur, Indrakshi Raychowdhury, and Ramesh Anishetty. “SU(N) Irreducible Schwinger Bosons”. *J. Math. Phys.* **51**, 093504 (2010). [arXiv:1003.5487](#). 1
- [131] Indrakshi Raychowdhury and Jesse R. Stryker. “Loop, String, and Hadron Dynamics in SU(2) Hamiltonian Lattice Gauge Theories”. *Phys. Rev. D* **101**, 114502 (2020). [arXiv:1912.06133](#). 1, 3.1.2, 3.1.2, B
- [132] Zohreh Davoudi, Indrakshi Raychowdhury, and Andrew Shaw. “Search for efficient formulations for Hamiltonian simulation of non-Abelian lattice gauge theories”. *Phys. Rev. D* **104**, 074505 (2021). [arXiv:2009.11802](#). 1, 3, 3.1, 3.1, 4
- [133] Jad C. Halimeh, Haifeng Lang, Julius Mildenberger, Zhang Jiang, and Philipp Hauke. “Gauge-Symmetry Protection Using Single-Body Terms”. *PRX Quantum* **2**, 040311 (2021). [arXiv:2007.00668](#). 1

- [134] Minh C. Tran, Yuan Su, Daniel Carney, and Jacob M. Taylor. “Faster Digital Quantum Simulation by Symmetry Protection”. *Phys. Rev. X. Quantum.* **2**, 010323 (2021). [arXiv:2006.16248](#).
- [135] Valentin Kasper, Torsten V. Zache, Fred Jendrzejewski, Maciej Lewenstein, and Erez Zohar. “Non-Abelian gauge invariance from dynamical decoupling”. *Phys. Rev. D* **107**, 014506 (2023). [arXiv:2012.08620](#).
- [136] Henry Lamm, Scott Lawrence, and Yukari Yamauchi. “Suppressing Coherent Gauge Drift in Quantum Simulations” (2020). [arXiv:2005.12688](#).
- [137] Jad C. Halimeh, Haifeng Lang, and Philipp Hauke. “Gauge protection in non-abelian lattice gauge theories”. *New J. Phys.* **24**, 033015 (2022). [arXiv:2106.09032](#). 1
- [138] Saurabh V. Kadam, Indrakshi Raychowdhury, and Jesse R. Stryker. “Loop-string-hadron formulation of an SU(3) gauge theory with dynamical quarks”. *Phys. Rev. D* **107**, 094513 (2023). [arXiv:2212.04490](#). 1
- [139] Yuan Su, Hsin-Yuan Huang, and Earl T. Campbell. “Nearly tight Trotterization of interacting electrons”. *Quantum* **5**, 495 (2021). [arXiv:2012.09194](#). 2, 4
- [140] Burak Şahinoğlu and Rolando D. Somma. “Hamiltonian simulation in the low-energy subspace”. *npj Quantum Inf.* **7**, 119 (2021). [arXiv:2006.02660](#).
- [141] Changhao Yi and Elizabeth Crosson. “Spectral analysis of product formulas for quantum simulation”. *npj Quantum Information* **8**, 37 (2022). [arXiv:2102.12655](#). 2, 4
- [142] Wikipedia contributors. “Logic synthesis — Wikipedia, the free encyclopedia” (2013). [Online; accessed Dec-2022]. 2
- [143] Boris Golubov, Aleksandr Efimov, and Valentin Skvortsov. “Walsh series and transforms: theory and applications”. *Volume 64*. Springer Science & Business Media. (2012). 2.2.1
- [144] Rao K Yarlagadda and John E Hershey. “Hadamard matrix analysis and synthesis: with applications to communications and signal/image processing”. *Volume 383*. Springer Science & Business Media. (2012). 2.2.1
- [145] Jonathan Welch, Daniel Greenbaum, Sarah Mostame, and Alan Aspuru-Guzik. “Efficient quantum circuits for diagonal unitaries without ancillas”. *New Journal of Physics* **16**, 033040 (2014). [arXiv:1306.3991](#). 2.2.1, 3.2.1
- [146] Christopher Kane, Dorota M. Grabowska, Benjamin Nachman, and Christian W. Bauer. “Efficient quantum implementation of 2+1 U(1) lattice gauge theories with Gauss law constraints” (2022). [arXiv:2211.10497](#). 2.2.1, 3.2.1
- [147] Manu Mathur and T. P. Sreeraj. “Lattice Gauge Theories and Spin Models”. *Phys. Rev. D* **94**, 085029 (2016). [arXiv:1604.00315](#). 3, 3.1
- [148] Manu Mathur and Atul Rathor. “Exact duality and local dynamics in SU(N) lattice gauge theory”. *Phys. Rev. D* **107**, 074504 (2023). [arXiv:2109.00992](#). 3.1
- [149] N. E. Ligterink, N. R. Walet, and R. F. Bishop. “Towards a many body treatment of Hamiltonian lattice SU(N) gauge theory”. *Annals Phys.* **284**, 215–262 (2000). [arXiv:hep-lat/0001028](#).
- [150] Pietro Silvi, Enrique Rico, Marcello Dalmonte, Ferdinand Tschirsich, and Simone Montangero. “Finite-density phase diagram of a (1+1)-d non-abelian lattice gauge theory with tensor networks”. *Quantum* **1**, 9 (2017). [arXiv:1606.05510](#).
- [151] R. Brower, S. Chandrasekharan, and U. J. Wiese. “QCD as a quantum link model”. *Phys. Rev. D* **60**, 094502 (1999). [arXiv:hep-th/9704106](#). 3
- [152] Stefan Kühn, J. Ignacio Cirac, and Mari Carmen Bañuls. “Non-Abelian string breaking phenomena with Matrix Product States”. *JHEP* **07**, 130 (2015). [arXiv:1505.04441](#). 3
- [153] Mari Carmen Bañuls, Krzysztof Cichy, J. Ignacio Cirac, Karl Jansen, and Stefan Kühn. “Efficient basis formulation for 1+1 dimensional SU(2) lattice gauge theory: Spectral calculations with matrix product states”. *Phys. Rev. X* **7**, 041046 (2017). [arXiv:1707.06434](#).
- [154] P. Sala, T. Shi, S. Kühn, M. C. Bañuls, E. Demler, and J. I. Cirac. “Variational study of U(1) and SU(2) lattice gauge theories with Gaussian states in 1+1 dimensions”. *Phys. Rev. D* **98**, 034505 (2018). [arXiv:1805.05190](#). 3
- [155] C. J. Hamer, Wei-hong Zheng, and J. Oitmaa. “Series expansions for the massive Schwinger model in Hamiltonian lattice theory”. *Phys. Rev. D* **56**, 55–67 (1997). [arXiv:hep-lat/9701015](#). 3.1
- [156] Yu Tong, Victor V. Albert, Jarrod R. McClean, John Preskill, and Yuan Su. “Provably accurate simulation of gauge theories and bosonic systems”. *Quantum* **6**, 816 (2022). [arXiv:2110.06942](#). 3.1
- [157] Frank Gray. “Pulse code communication”. U.S. Patent No. 2,632,058 (1953). 12
- [158] Stephen S Bullock and Igor L Markov. “Smaller circuits for arbitrary n-qubit diagonal computations”. *Quantum Information and Computation* **4**, 027–047 (2004). [arXiv:quant-ph/0303039](#). 13
- [159] Eyal Kushilevitz and Yishay Mansour. “Learning decision trees using the fourier spectrum”. In Proceedings of the twenty-third annual ACM symposium on Theory of computing. Pages 455–464. (1991). 14

- [160] Alex Bocharov, Martin Roetteler, and Krysta M Svore. “Efficient synthesis of universal repeat-until-success quantum circuits”. *Physical Review Letters* **114**, 080502 (2015). [arXiv:1404.5320](#). 3.2.1, 3.2.1, 3.2.1, 3.2.2, 3.2.2, 3.2.2
- [161] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John Smolin, and Harald Weinfurter. “Elementary gates for quantum computation”. *Phys. Rev. A* **52**, 3457 (1995). [arXiv:quant-ph/9503016](#). 3.2.2
- [162] Yong He, Ming-Xing Luo, E. Zhang, Hong-Ke Wang, and Xiao-Feng Wang. “Decompositions of n-qubit toffoli gates with linear circuit complexity”. *International Journal of Theoretical Physics* **56**, 2350–2361 (2017). 3.2.2
- [163] Z. Davoudi and J. R. Styker. “On the quantum-computing cost of lattice quantum chromodynamics”. work in progress (2023). 4
- [164] Daniel C. Hackett, Kiel Howe, Ciaran Hughes, William Jay, Ethan T. Neil, and James N. Simone. “Digitizing Gauge Fields: Lattice Monte Carlo Results for Future Quantum Computers”. *Phys. Rev. A* **99**, 062341 (2019). [arXiv:1811.03629](#). 4
- [165] Tobias Hartung, Timo Jakobs, Karl Jansen, Johann Ostmeyer, and Carsten Urbach. “Digitising SU(2) gauge fields and the freezing transition”. *Eur. Phys. J. C* **82**, 237 (2022). [arXiv:2201.09625](#). 4
- [166] Andrew M Childs, Dmitri Maslov, Yunseong Nam, Neil J Ross, and Yuan Su. “Toward the first quantum simulation with quantum speedup”. *Proceedings of the National Academy of Sciences* **115**, 9456–9461 (2018). [arXiv:1711.10980](#). 4
- [167] Dong An, Di Fang, and Lin Lin. “Time-dependent unbounded Hamiltonian simulation with vector norm scaling”. *Quantum* **5**, 459 (2021). [arXiv:2012.13105](#). 4
- [168] Qi Zhao, You Zhou, Alexander F. Shaw, Tongyang Li, and Andrew M. Childs. “Hamiltonian Simulation with Random Inputs”. *Phys. Rev. Lett.* **129**, 270502 (2022). [arXiv:2111.04773](#). 4
- [169] Marcela Carena, Henry Lamm, Ying-Ying Li, and Wanqiang Liu. “Lattice renormalization of quantum simulations”. *Phys. Rev. D* **104**, 094519 (2021). [arXiv:2107.01166](#). 4
- [170] Anthony Ciavarella. “Algorithm for quantum computation of particle decays”. *Phys. Rev. D* **102**, 094505 (2020). [arXiv:2007.04447](#).
- [171] Raúl A. Briceño, Juan V. Guerrero, Maxwell T. Hansen, and Alexandru M. Sturzu. “Role of boundary conditions in quantum computations of scattering observables”. *Phys. Rev. D* **103**, 014506 (2021). [arXiv:2007.01155](#). 4
- [172] Michael A Nielsen and Isaac Chuang. “Quantum computation and quantum information”. Cambridge University Press. (2002). A
- [173] Craig Gidney. “Halving the cost of quantum addition”. *Quantum* **2**, 74 (2018). [arXiv:1709.06648](#). A, 15, A, A, A
- [174] Cody Jones. “Low-overhead constructions for the fault-tolerant toffoli gate”. *Physical Review A* **87**, 022328 (2013). [arXiv:1212.5069](#). A
- [175] Steven A. Cuccaro, Thomas G. Draper, Samuel A. Kutin, and David Petrie Moulton. “A new quantum ripple-carry addition circuit” (2004). [arXiv:quant-ph/0410184](#). A
- [176] Mihir K Bhaskar, Stuart Hadfield, Anargyros Papageorgiou, and Iasonas Petras. “Quantum algorithms and circuits for scientific computing”. *Quantum Information and Computation* **16**, 0197–0236 (2016). [arXiv:1511.08253](#). A