

Quantum circuits for $SU(3)$ lattice gauge theory

Praveen Balaji,¹ Cianán Conefrey-Shinozaki¹, Patrick Draper^{1,*}, Jason K. Elhaderi¹, Drishti Gupta¹,
Luis Hidalgo¹, Andrew Lytle¹, and Enrico Rinaldi²

¹*Department of Physics, University of Illinois, Urbana, Illinois 61801, USA*

²*Quantinuum K.K., Otemachi Financial City Grand Cube 3F, 1-9-2 Otemachi, Chiyoda-ku, Tokyo, Japan*



(Received 11 May 2025; accepted 26 August 2025; published 24 September 2025; corrected 7 November 2025)

Lattice gauge theories in varying dimensions, lattice volumes, and truncations offer a rich family of targets for Hamiltonian simulation on quantum devices. In return, formulating quantum simulations can provide new ways of thinking about the quantum structure of gauge theories. In this work, we consider pure $SU(3)$ gauge theory in two and three spatial dimensions in a streamlined version of the electric basis. We use a formulation of the theory that balances locality of the Hamiltonian and size of the gauge-invariant state space, and we classically pre-compute dictionaries of plaquette operator matrix elements for use in circuit construction. We build circuits for simulating time evolution on arbitrary lattice volumes, spanning circuits suitable for Noisy Intermediate-Scale Quantum era hardware to future fault-tolerant devices. Relative to spin models, time evolution in lattice gauge theories involves more complex local unitaries, and the Hilbert space of all quantum registers may have large unphysical subspaces. Based on these features, we develop general, volume-scalable tools for optimizing circuit depth, including pruning and fusion algorithms for collections of large multicontrolled unitaries. We describe scalings of quantum resources needed to simulate larger circuits and some directions for future algorithmic development.

DOI: [10.1103/k8f6-yft8](https://doi.org/10.1103/k8f6-yft8)

I. INTRODUCTION

Rapid advances in quantum computing technology have led to renewed interest in the Hamiltonian simulation of fundamental physics. The classical simulation of many quantum systems is believed to be intractable, due in part to the size of the systems' state spaces, which may require exponential classical resources to store and manipulate. Quantum computers, utilizing native quantum degrees of freedom, offer an exciting alternative.

Gauge theories underpin the Standard Model of particle physics and have long been studied by theoretical, experimental, and computational means. The Hamiltonian formulation of lattice gauge theories was first developed in [1]. In the intervening decades, much of the computational work focused instead on the Lagrangian formulation and numerical path integral evaluation. In cases where the Euclidean action is real and bounded from below, the path integral weight can be interpreted as a probability, making these theories amenable to (classical) Monte Carlo techniques. The most well-known and widely studied example is lattice

quantum chromodynamics (QCD), where high-performance classical simulations have successfully calculated hadronic masses [2], form factors [3], and other equilibrium properties to high precision [4,5]. Despite these advances, many properties of gauge theories have remained difficult or intractable to study classically, including general scattering processes, nonequilibrium dynamics, and physics at large baryon chemical potential or topological θ angle. Quantum computing opens up a new avenue by which lattice gauge theories may be studied in the Hamiltonian formulation [6–8]. In light of recent technological progress, the problem of mapping gauge theory degrees of freedom and the Hamiltonian to those of quantum hardware platforms has been of great interest [9–23].

A first step toward placing the Standard Model on a quantum computer is to develop real-time simulations of QCD or one of its simpler relatives, including pure $SU(2)$ or $SU(3)$ gauge theory, or QCD in lower dimensions. Several approaches to the quantum simulation of $SU(3)$ gauge theory have been explored; see, e.g., [9,24–37]. In the future, some hybridization of the techniques developed in these studies may provide the earliest exhibitions of “quantum advantage for particle physics.”

In this work, we develop circuits and tools for simulating pure $SU(3)$ gauge theory on general-purpose qubit-based quantum computers. Although the underlying technology used to implement qubits can vary (trapped ions, photonic, Rydberg atoms, superconducting transmons, etc.), the

*Contact author: pdraper@illinois.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

circuit-level specifications of time evolution operators that we develop here can be deployed on any of these devices.

We work in a reduced version of the electric basis, in which the basic degrees of freedom are gauge fluxes on links, supplemented by singlet information encoded on vertices. Because of the latter, it is a generalization of what is known as the local multiplet basis in the literature. The plaquette-local Hamiltonian is computed classically using the representation theory of $SU(3)$, ensuring that only physical transitions that satisfy the Gauss law constraint are present at the noiseless level. By specifying the transitions locally, with a local Hilbert space that grows with the size of the electric truncation, the overall quantum resources (qubits and gates) scales with the size/volume of the system. Our specification is complete, taking into account subtleties associated with different multiplicities of gauge-invariant states that can arise at each vertex.

We have developed a modular code that generates Hamiltonian evolution circuits for the $SU(3)$ gauge theory, with a specified truncation to render the Hilbert space finite-dimensional, and in 1, 2, or 3 spatial dimensions, with lattice volume and truncations specified by the user. Our circuits are based on Givens rotations that drive transitions between any two physical bitstrings. We have developed and implemented various optimizations to reduce overall circuit depths, including strategies that remove controls from rotation gates which are redundant on the physical subspace of the simulation Hilbert space. These optimizations typically reduce circuit depths by a factor of 5–10. The code is being assembled into a Python package which we plan on open-sourcing in the near future, and will include additional tools supporting the downstream analysis of simulation results such as the reconstruction of physical lattice states, the computation of simple electric-basis observables from shots data, and options for performing rudimentary error mitigation.

The remainder of this article is organized as follows. In Sec. II we construct the reduced electric basis, which balances locality of the lattice Hilbert space while preserving gauge invariance. We provide a “master formula” for the matrix elements of the plaquette operator in Eqs. (15) and (16), and describe its efficient classical precomputation. In Sec. III we build quantum circuits using a general approach that applies to any lattice and truncation. We describe several optimizations and provide resource estimates for a variety of small lattices in Table IV. In Sec. IV we perform classical state vector simulations for small lattices in $d = 3/2$, $d = 2$, and $d = 3$ ($d = 3/2$ refers to a line of plaquettes) with different boundary conditions. We study effects of varying the coupling, irrep truncation, number of Trotter steps, and approximation on magnetic transitions. For the 2×2 lattice in $d = 2$, we also perform noisy simulations using the Quantinuum System Model H2 emulator [38]. For the $2 \times 2 \times 2$ lattice in $d = 3$, we perform a simple tensor network analysis using matrix

product states. Finally, in Sec. V we conclude with a summary of some directions for future study.

II. THEORETICAL SETUP

Our target theory is given by the Kogut-Susskind (KS) Hamiltonian, describing $SU(N)$ gauge theory discretized on a rectilinear spatial lattice

$$H = H_E + H_M$$

$$= \frac{1}{2a} \left[g^2 \sum_{\ell \in \mathcal{L}} E_\ell^2 + \frac{1}{g^2} \sum_{P \in \mathcal{P}} 2N - \square_P - \square_P^\dagger \right]. \quad (1)$$

g is the gauge field coupling, and a is the lattice spacing. Throughout this work, we set $a = 1$. $\mathcal{L}$ denotes the set of all links and $\mathcal{P}$ the set of all plaquettes of the lattice. We work in a version of the electric basis, in which the electric operator E_ℓ^2 is diagonal. The magnetic Hamiltonian on plaquette P is defined in terms of fundamental-representation link operators U^f circumscribing P

$$\square_P \equiv \text{Tr}(U^f(s, \vec{e}_i) U^f(s + \vec{e}_i, \vec{e}_j) U^f(s + \vec{e}_j, \vec{e}_i)^\dagger U^f(s, \vec{e}_j)^\dagger). \quad (2)$$

Here s denotes a lattice site at a corner of P and $\vec{e}_{i,j}$ are orthonormal lattice vectors. The trace is taken in the defining representation. Of fundamental importance will be the classical precomputation of matrix elements of $\square_P + \square_P^\dagger$ in the gauge-invariant basis. First we describe the construction of this basis.

A. Reduced electric basis states

We begin in the Hilbert space associated with the standard electric basis [24]. (Subsequently we will refer to the standard basis as just the *electric basis*, to be distinguished from the “gauge-invariant electric basis” and “reduced electric basis” defined below.) The whole-lattice Hilbert space is constructed as a tensor product of Hilbert spaces associated with each link. For a single link $\ell(s, \vec{e}_i)$, the electric basis states $|R, \lambda_-, \lambda_+\rangle$ are labeled by an $SU(N)$ irrep R and two multi-indices $\lambda_{+/-}$. The latter represent two states in the irrep R , associated with the two ends of the link ℓ , respectively. In $SU(3)$, for example, an irrep can be labeled by whole numbers $R = (p, q)$, and the value of the quadratic Casimir operator is given by

$$E^2 |R, \lambda_-, \lambda_+\rangle = \frac{1}{3} (p^2 + q^2 + pq + 3(p + q)) |R, \lambda_-, \lambda_+\rangle. \quad (3)$$

It is notationally convenient to split the link state into a product of two “half-link” states (this terminology was also used in [18]). For each link $\ell(s, \vec{e}_i)$ in the lattice, we introduce the half-link states $|R_\ell, \lambda_{\ell^\pm}\rangle$. The “positive” and “negative” half-links are defined according to the unit vector $\vec{e}_i$; the positive half-link ℓ^+ is the end of the link

toward which $\vec{e}_i$ points, and the negative half-link ℓ^- is the end toward which $-\vec{e}_i$ points. Electric basis states of ℓ are then written in terms of half-link states as

$$|R_\ell, \lambda_{\ell^-}, \lambda_{\ell^+}\rangle \equiv |R_\ell, \lambda_{\ell^-}\rangle \otimes |R_\ell, \lambda_{\ell^+}\rangle. \quad (4)$$

This notation will be useful later to simplify the expression for the plaquette operator matrix elements.

Even for a single link, the dimension of the Hilbert space is infinite. A finite-dimensional Hilbert space is obtained by truncating the electric basis to irreps of low-lying electric energy density $\propto E^2$. We denote the truncation to r -index tensor irreps by T_r . For instance, T_1 corresponds to the truncation to irreps $\{1, 3, \bar{3}\}$ in $SU(3)$.

This construction, detailed in [24], has the significant advantage that it expresses the lattice Hilbert space as a tensor product of local Hilbert spaces. Since the Hamiltonian is also local at the level of plaquettes, the construction has good “in-principle” scalability properties with the lattice volume. Its primary disadvantage is that the dimension of the local Hilbert space grows rapidly with irrep truncation, and in a way not particularly naturally adapted to qubits, such that the plaquette operator becomes rather complicated to implement. Furthermore most of the lattice Hilbert space is unphysical, since the Gauss’ law has not been imposed.

A spatial gauge transformation at lattice site s is generated by

$$G^a(s) = \sum_{i=1}^3 [E_+^a(s - e_i, e_i) - E_-^a(s, e_i)]. \quad (5)$$

Here E_+^a and E_-^a are the electric field operators generating the left and right action of the gauge group on the link operators, and the λ ’s in Eq. (3) and (4) are the eigenvalues of the $E_{+,-}^a$ in the Cartan subalgebra. Gauss’ law can then be enforced by making physical states satisfy $G^a(s)|\text{phys}\rangle = 0$, and it is equivalent to the requirement that the half-link states of $|\text{phys}\rangle$ meeting at site s form an $SU(N)$ singlet. Thus, a basis of *physical* lattice states may be obtained from the electric basis by contracting the link states meeting at each vertex s with an $SU(N)$ -invariant tensor. Explicitly, a single physical electric basis state $|\omega\rangle$ may be expanded in terms of electric basis states¹ using generalized Clebsch-Gordan coefficients (CGCs)

¹Under a gauge transformation Ω , the positive and negative half-link are transformed as $|R_\ell, \lambda_{\ell^+}\rangle \rightarrow D_{\lambda\rho}^f(\Omega)|R_\ell, \rho_{\ell^+}\rangle$ and $|R_\ell, \lambda_{\ell^-}\rangle \rightarrow D_{\rho\lambda}^f(\Omega^\dagger)|R_\ell, \rho_{\ell^-}\rangle$ for unitary representation matrices $D^f(\Omega)$ (see, e.g., Refs. [12,39]). The relation between the matrix elements $D_{\rho\lambda}^f(\Omega^\dagger)$ and $D_{\lambda\rho}^f(\Omega)$ depends on sign conventions, generalizing the Condon-Shortley phases. In general $D_{\rho\lambda}^f(\Omega^\dagger) = D_{\lambda\rho}^{f*}(\Omega) = \phi(\lambda)\phi(\rho)D_{\tilde{\lambda}\tilde{\rho}}^f(\Omega)$, where $\tilde{f}$ and $\tilde{\lambda}$ are defined in the text above and the ϕ ’s are the conventional phases. We discuss them more thoroughly around Eqs. (12) and (13). This distinction between the basis vectors of barred irreps and naive conjugation is the origin of the phases in Eq. (6).

$$|\omega\rangle = \bigotimes_{s \in S} \sum_{\lambda} \left(\prod_{\ell \in \ell^-(s)} \phi(\lambda) \right) \langle 1, \Gamma_s | \bigotimes_{\ell \in \ell^-(s)} (\bar{R}, \tilde{\lambda}) \bigotimes_{\ell \in \ell^+(s)} (R, \lambda) \rangle \bigotimes_{\ell \in \ell^+(s)} (R, \lambda) \rangle. \quad (6)$$

Here $\ell(s)$ denotes the set of half-links meeting at site s , and $\ell^\pm(s)$ are the subsets of positive/negative half-links at the site. $\bar{R}$ is the representation conjugate to R ; if R is a real representation, then $\bar{R}$ is the same as R . Apart from possible phase conventions $\phi(\lambda)$ which will be revisited below (see footnote), $\tilde{\lambda}$ is the basis state of $\bar{R}$ obtained by negating the weights of the state λ of R . The CGC $\langle 1, \Gamma_s | \bigotimes_{\ell \in \ell^-(s)} (\bar{R}, \tilde{\lambda}) \bigotimes_{\ell \in \ell^+(s)} (R, \lambda) \rangle$ is the Γ^{th} invariant tensor in the corresponding product of irreps. The allowed invariant tensors at each vertex depend only on the link irreps meeting at the vertex, but in general, they are not unique; for example, in $SU(3)$ there are two singlets in $8 \otimes 8 \otimes 8$ and 145 in $8 \otimes 8 \otimes 8 \otimes 8 \otimes 8 \otimes 8$. We discuss the enumeration of singlets further in Sec. II C.

Thus, a complete specification of a gauge invariant basis state is given by an assignment of irrep to every link along with a consistent choice of singlet at every vertex. In particular, we no longer need separate basis states for each irrep internal state (e.g. red, green, blue for $R = 3$). The *gauge-invariant electric basis* spans the physical Hilbert space, which is much smaller than the Hilbert space spanned by the electric basis.

The fact that the consistent vertex singlets are dependent on the link irreps is inconvenient for locality and for implementing the magnetic time evolution operator. Consistent assignments of irreps to every link and singlets to every vertex is essentially a global problem. For this reason, it is advantageous to relax the “consistency” requirement, and simply define a *reduced electric basis* of states by an assignment of irrep to every link and a choice of “singlet multiplicity index” at every vertex. In this construction, the Hilbert space is built again from a tensor product of local Hilbert spaces. However, basis states are included regardless of whether the tensor product of the link irreps contains even one singlet, or of whether the singlet multiplicity index is in the range of the number of singlets in the tensor product. Thus some unphysical states are reintroduced.

The idea of using link irreps to define the Hilbert space, rather than the full electric basis of [24], was developed for quantum simulations of gauge theories in [9,40,41], and is often referred to as the “local multiplet basis” in the literature. The reduced electric basis defined above extends the local multiplet basis to include site singlet multiplicity information, which is generically essential to obtain a complete description of the gauge invariant Hilbert space and obtain meaningful matrix elements. Some previous works, for example [9], do not have to assign separate singlet multiplicity degrees of freedom because the truncation is sufficiently strong that no multiplicity arises.

In $SU(3)$, with irreps truncated to $\{1, 3, \bar{3}\}$ and working with a line of plaquettes, the singlets associated with any collection of three irreps is unique. This can be extended to higher dimensions by working with lattices with only triangular vertices (see, e.g., [42]), but it still depends on the irrep truncation: once two-index tensor irreps are included, the singlets are no longer uniquely identified by the link irreps.

The nice locality properties of the reduced electric basis is manifest. Let us illustrate the price—some unphysical basis states—with some examples. In $SU(3)$ in three spatial dimensions with link irreps drawn from $\{1, 3, \bar{3}\}$, the largest number of distinct invariant tensors (6) is found in $3 \otimes 3 \otimes 3 \otimes \bar{3} \otimes \bar{3} \otimes \bar{3}$, so the singlet multiplicity index must run from one to six. However, other link states may contain fewer singlets. For example, $3 \otimes 3 \otimes 3 \otimes 1 \otimes 1 \otimes 1$ contains only one singlet, so when the links are in these irreps there are five choices of singlet multiplicity index at the vertex that are inconsistent with the links. Furthermore $3 \otimes 3 \otimes 3 \otimes 3 \otimes \bar{3} \otimes \bar{3}$ contains zero singlets, so the corresponding link irreps are inconsistent by themselves. Thus the reduced electric basis expands the Hilbert space to include unphysical states, but the Hilbert space dimension is still considerably lower than in the full electric basis, and the Hamiltonian is local.²

To summarize, a basis state $|\omega\rangle$ of the reduced electric basis is given by

$$|\omega\rangle = \bigotimes_{s \in \mathcal{S}, \ell \in \mathcal{L}} |R_\ell\rangle |\Gamma_s\rangle \quad (7)$$

where

- (i) the link irreps R_ℓ are drawn independently from the irrep truncation, and
- (ii) the site singlet multiplicity index Γ_s runs from one to the maximum number of independent singlets that appear in any tensor product of $2d$ elements of the irrep truncation. (We may also impose an additional truncation on the collection of singlets; see Sec. II D.)

Note that the reduced electric basis includes all the basis states of the gauge-invariant electric basis exactly once.

B. Matrix elements of the plaquette operator

We seek to classically precompute matrix elements of the gauge-invariant plaquette operator between gauge-invariant basis states. This involves the action of the operator both on

²It is not clear what impact these unphysical states will have on fidelity in simulations on Noisy Intermediate-Scale Quantum (NISQ) era devices, apart from making it worse. This must be tested on small systems, as below in Fig. 15, in order to develop mitigation strategies. In the fault-tolerant era, the hope (shared for essentially any scalable lattice gauge theory simulation) is that error correction will be sufficient to prevent wandering off into unphysical sectors.

the link irreps and the site singlets. We begin by discussing the two types of Clebsch-Gordan coefficients that appear in the final master formula for these matrix elements.

The first type of CGC is associated with the site singlets. In d spatial dimensions, at most $2d$ half-links meet at each site—in $d = 3$, for example, three negative half-links and three positive half-links meet.³ Let $\ell(s) = \ell^-(s) \cup \ell^+(s)$ denote the set of all negative and positive half-links meeting at site s , respectively. Fixing the irreps $R_{\ell(s)}$, we define a tensor product space

$$\bigotimes_{\ell \in \ell^+(s)} R_\ell \bigotimes_{\ell \in \ell^-(s)} \bar{R}_\ell \quad (8)$$

associated with half-links meeting at s . It is important that we impose an arbitrary but fixed ordering of the half-links meeting at each site, which we call an *F-order*. This ordering is implicit in the notation $\bigotimes_{\ell \in \ell^\pm(s)}$. It is necessary to avoid permutation ambiguities when the same irrep appears on multiple links meeting at a site: in general, the states appearing in an irrep direct sum decomposition of a tensor product transform nontrivially under permutations.

A standard electric basis for this space is

$$\bigotimes_{\ell \in \ell^+(s)} |R_\ell, \lambda_\ell\rangle \bigotimes_{\ell \in \ell^-(s)} |\bar{R}_\ell, \tilde{\lambda}_\ell\rangle. \quad (9)$$

Gauge-invariant states correspond to singlets in the Hilbert space spanned by (9). The CGCs for these singlets were introduced already in (6) and are denoted

$$\langle 1, \Gamma_s | \bigotimes_{\ell^-} \bar{R}_\ell \tilde{\lambda}_{\ell^-} \bigotimes_{\ell^+} R_\ell \lambda_{\ell^+} \rangle \quad (10)$$

where Γ_s labels the Γ^{th} singlet living in the space (9). The half-link states appear in the tensor product in a chosen *F-order*, which means that each (R, λ) state appearing in a CGC can be attributed to a specific half-link on the lattice. (See Fig. 1.)

The second type of CGC is associated with the link irreps. A link operator in representation r acts on the electric basis states of the link as

$$\begin{aligned} U_{\sigma\sigma'}^r |R_\ell, \lambda_{\ell^-}, \lambda_{\ell^+}\rangle \\ = \sum_{\Omega_\ell, \Lambda_{\ell^-}, \Lambda_{\ell^+}} \sqrt{\frac{\dim(R_\ell)}{\dim(\Omega_\ell)}} \langle \Omega_\ell, \Lambda_{\ell^-} | R_\ell, \lambda_{\ell^-}; r, \sigma \rangle \\ \times \langle \Omega_\ell, \Lambda_{\ell^+} | R_\ell, \lambda_{\ell^+}; r, \sigma' \rangle | \Omega_\ell, \Lambda_{\ell^-}, \Lambda_{\ell^+} \rangle. \end{aligned} \quad (11)$$

In this expression, objects like $\langle \Omega_\ell, \Lambda_{\ell^-} | R_\ell, \lambda_{\ell^-}; r, \sigma \rangle$ denote CGCs.

The gauge invariant plaquette operator appearing in the KS Hamiltonian was defined in Eq. (2) in terms of the fundamental representation link operators. It is convenient

³There will be fewer half-links meeting at sites on the boundary of nonperiodic lattices.




FIG. 1. An F -ordering is a link ordering at each vertex which is used to uniquely specify invariant tensors associated with singlet states. At each vertex, the precise order in which the link irreps appear in the “tensor product side” of the CGC is given by the F -ordering.

to convert $(U^f)^\dagger$ into $U^{\bar{f}}$. As mentioned previously, this relation can involve conventional signs

$$[(U^f)^\dagger]_{\sigma\sigma'} = (U^f_{\sigma'\sigma})^* = \phi(\sigma)\phi(\sigma')U^{\bar{f}}_{\sigma'\sigma}. \quad (12)$$

Here, $\phi(\sigma) = \pm 1$ relates the basis states associated with naïvely conjugating f to the actual basis states we use for $\bar{f}$. This map involves a (conventional) reversal of orientation of some basis vectors. Since the invariant tensor on $R \otimes R^*$ is just the identity matrix, the ϕ are related to the singlet in $R \otimes \bar{R}$ by

$$\begin{aligned} \phi(\lambda) &= \phi(\tilde{\lambda}) = \text{sign}(\langle 1 | R, \lambda; \bar{R}, \tilde{\lambda} \rangle) \\ &= \text{sign}(\langle 1 | \bar{R}, \tilde{\lambda}; R, \lambda \rangle). \end{aligned} \quad (13)$$

Note that $\phi(\lambda)\phi(\lambda) = 1$.

We are now ready to construct our master formula. Label the sites and links around a plaquette P counter-clockwise as $s_1, \dots, s_4$ and $\ell_1, \dots, \ell_4$. The set of plaquette sites and links are $\mathcal{S}(P)$ and $\mathcal{L}(P)$. At each site, there are at most four links that are not part of the plaquette, but they must be known to calculate singlet CGCs. Following [9,41] we will refer to these “control” or “external” links. The collection of control links at site s_i is $\vec{C}_{s_i}$. (See Fig. 1.) Let $\langle \omega_{n,m} |$ denote two gauge-invariant electric basis states. (As discussed above, these are also a subset of the reduced electric basis states; we require matrix elements only for the former.) $\langle \omega_n | \square | \omega_m \rangle$ can then be found by acting on $|\omega_m\rangle$ with the four link operators, then contracting with $\langle \omega_n |$, collapsing any sums over half-link states that are not modified (any half-links not part of the plaquette).

Carrying out this procedure we arrive at the master formula for the plaquette operator matrix elements in the reduced electric basis. The expression is somewhat lengthy, so to preserve readability of the text we have placed it in the Appendix, Eq. (A1). Here we provide a compact version, as follows. First, we can take a sign convention for the CGCs such that

$$\begin{aligned} \langle (A, a) \otimes (f, \sigma) | B, b \rangle \\ = \phi(a)\phi(\sigma)\phi(b) \langle (\bar{A}, \bar{a}) \otimes (\bar{f}, \bar{\sigma}) | \bar{B}, \bar{b} \rangle \end{aligned} \quad (14)$$

holds. (A similar relation from interchanging f and $\bar{f}$ would also hold.) Second, we define the site factors

$$\begin{aligned} \left\langle \begin{array}{ccc} R & G & S \\ f & \vec{C} & \bar{f} \\ R' & G' & S' \end{array} \right\rangle &= \sum_{\sigma=1}^{\dim(f)} \sum_{r=1}^{\dim(R)} \sum_{r'=1}^{\dim(R')} \sum_{s=1}^{\dim(S)} \sum_{s'=1}^{\dim(S')} \sum_{\vec{c} \in \vec{C}} \phi(\sigma) \\ &\times \langle R', r' | (R, r) \otimes (f, \sigma) \rangle \langle S', s' | (S, s) \otimes (\bar{f}, \bar{\sigma}) \rangle \\ &\times \langle 1, G | (R, r) \otimes (S, s) \otimes (\vec{C}, \vec{c}) \rangle_F \\ &\times \langle 1, G' | (R', r') \otimes (S', s') \otimes (\vec{C}, \vec{c}) \rangle_F. \end{aligned} \quad (15)$$

The subscript F on each singlet CGC reminds us that the irreps in the set $\{R, S, \vec{C}\}$ should be placed into the lattice F -order, as described above.

These site factors are somewhat similar to the 9-R symbols that appear in the matrix element formulas derived in [9]. However, here their form is fixed for each site (f always appears on the left column and $\bar{f}$ always appears on the right column), and more significantly, singlet

multiplicity indices are not summed over. See Fig. 2 for a sketch of how different singlet multiplicities appear in $\square$ matrix elements. These site factors also apply to rectilinear lattices in general spatial dimensions; the only dimension dependence is in the number of irreps appearing in $\vec{C}$.

Using (14) and (15), and suppressing the obvious Kronecker deltas, the master formula (A1) can be written in the compact form

$$\begin{aligned}
\langle \omega_n | \square | \omega_m \rangle = & \left(\prod_{k=1}^4 \sqrt{\frac{\dim(R_{\ell_k}^m)}{\dim(R_{\ell_k}^n)}} \right) \left\langle \left\langle \begin{matrix} \bar{R}_{\ell_4}^m & \Gamma_{s_1}^m & \bar{R}_{\ell_1}^m \\ f & \vec{C}_{s_1} & \bar{f} \\ \bar{R}_{\ell_4}^n & \Gamma_{s_1}^n & \bar{R}_{\ell_1}^n \end{matrix} \right\rangle \right\rangle \left\langle \left\langle \begin{matrix} R_{\ell_1}^m & \Gamma_{s_2}^m & \bar{R}_{\ell_2}^m \\ f & \vec{C}_{s_2} & \bar{f} \\ R_{\ell_1}^n & \Gamma_{s_2}^n & \bar{R}_{\ell_2}^n \end{matrix} \right\rangle \right\rangle \\
& \times \left\langle \left\langle \begin{matrix} R_{\ell_2}^m & \Gamma_{s_3}^m & \bar{R}_{\ell_3}^m \\ f & \vec{C}_{s_3} & \bar{f} \\ R_{\ell_2}^n & \Gamma_{s_3}^n & \bar{R}_{\ell_3}^n \end{matrix} \right\rangle \right\rangle \left\langle \left\langle \begin{matrix} \bar{R}_{\ell_3}^m & \Gamma_{s_4}^m & R_{\ell_4}^m \\ f & \vec{C}_{s_4} & \bar{f} \\ \bar{R}_{\ell_3}^n & \Gamma_{s_4}^n & R_{\ell_4}^n \end{matrix} \right\rangle \right\rangle
\end{aligned} \quad (16)$$

The factorization of these matrix elements into site factors is enormously useful for classical precomputation of the matrix elements. Classical precomputation begins by calculating all nonvanishing site factors, enumerating all different ways that Eq. (15) can be populated with irreps and singlets given a truncation. Because of the F -ordering imposed on the singlet CGCs, site factors corresponding to the same site but to plaquettes lying in different planes of the lattice need not be equal. Furthermore, if open boundary conditions are used, then there may be different site factors at different “ends” of the lattice.

From a list of nonzero site factors, four at a time can then be “glued” together by making sure the columns “match.” More precisely, note that where the site factors about each other in Eq. (16), the adjacent columns refer to the same active link. However, one column holds the active link irreps, while the adjacent column in the neighboring site factor holds their conjugates. Therefore, two site factors can initially be glued together if the S and S' irreps of the first site factor, in the notation of Eq. (15), are conjugate to the R and R' irreps of the second. Then the rest of the site factors can be glued in this manner, forming the entire plaquette matrix element. Extra care should be taken when periodic boundary conditions constrain certain control links of separate sites to be in the same state. In a 2×2 lattice, for instance, $\vec{C}_{s_1}$ and $\vec{C}_{s_2}$ will have one link in common. This process will be repeated for plaquettes lying on different

planes of the lattice (and ends of the lattice, if open boundary conditions are used.)

Apart from the plaquette matrix elements, it is useful to generate sets of *physical plaquette states*—an assignment of irrep to each of the active and control links of a single plaquette—such that each of the four sites supports a singlet, and valid assignment of singlet multiplicity index to the sites. This can be done sequentially as follows. For s_1 , generate all possible singlets allowed by the irrep truncation, assigning irreps to specific links. For s_2 (s_3), connect all singlets present on s_1 (s_2) with another singlet that has an ℓ_1 (ℓ_2) irrep conjugate to the ℓ_1 (ℓ_2) irrep on s_1 (s_2). Then for s_4 , a similar connection must be done taking into account the ℓ_3 irrep of s_3 and the ℓ_4 irrep of s_1 . Throughout this process, boundary conditions may constrain the control link irreps present on each singlet.

We see in Table I a rapid growth in the number of $\square$ matrix elements in $SU(3)$ with dimension and irrep truncation. Fortunately, the plaquette operator matrix elements are very sparse, and we will provide some evidence that only a small subset of these transitions are needed to accurately capture dynamics.

C. Numerical computation of Clebsch-Gordan coefficients

The site factors (15) appearing in the master formula (16) for the physical state matrix elements of the magnetic Hamiltonian are composed of sums of $SU(N)$ Clebsch-Gordan coefficients, for example, quantities of the form

$$C_{M_1, M_2, \dots}^{M, \alpha} = \langle M, \alpha | M_1 \otimes M_2 \otimes \dots \rangle \quad (17)$$

TABLE I. The general number of $\square$ matrix elements in different lattices and truncations with periodic boundary conditions in $SU(3)$. T_1 is the irrep truncation $\{1, 3, \bar{3}\}$, and T_2 is the irrep truncation $\{1, 3, \bar{3}, 6, \bar{6}, 8\}$.

d	T_1	T_2
3/2	81	28858
2	19329	$\sim 1.48 \times 10^{10}$
3	$\sim 2.07 \times 10^{10}$	$\gtrsim 10^{20}$




FIG. 2. A representation of the action of $\square$ on a particular plaquette. Observe that initially all the vertex singlets are 0 or unoccupied because the irreps at each vertex only decompose into one singlet irrep. However $3 \otimes 3 \otimes \bar{3} \otimes \bar{3}$ contains two singlets in its decomposition so the operation of $\square$ will yield two plaquette states, one with nontrivial multiplicity at the top-right vertex. The amplitude for the transition to either state is given by the master formula and will in general be distinct.

where each of $M, M_1, M_2, \dots$ denotes both an irrep and a state within that irrep, and α indexes the outer multiplicity of M in the direct sum decomposition. To find the CGCs, we use the method described in [43] for products of two irreps and generalize it recursively to products of arbitrary numbers of irreps. We briefly describe the algorithm here, and refer to [43] for details.

- (1) The tensor product of irreps from the input are decomposed into a direct sum of irreps using the Little-Richardson rule, which involves constructing and composing the Gelfand-Tsetlin (GT) patterns of the tensored irreps to obtain new GT patterns associated with irreps in the direct sum decomposition. Reference [43] offers an iterative algorithm based on this rule that is utilized in our code.
- (2) For each irrep in the decomposition, the CGCs of the highest weight state are found. Simultaneous eigenstates of $N-1$ commuting matrices $(J_z^{(l)}, 1 \leq l \leq N-1)$ can be raised or lowered by generalized ladder operators, $J_+^{(l)}$ and $J_-^{(l)}$. The highest weight state H is defined as the state $|H, \alpha\rangle$ such that

$$J_+^{(l)} |H, \alpha\rangle = \sum_{M_1, M_2} C_{M_1, M_2}^{H, \alpha} (J_+^{S_1} \otimes \mathbb{I}^{S_2} + \mathbb{I}^{S_1} \otimes J_+^{S_2}) \times |M \otimes M'\rangle = 0 \quad (18)$$

Here $S_{1,2}$ denote the irreps to which the states $M_{1,2}$ belong. Since $\{|M_1 \otimes M_2\rangle\}$ forms an orthonormal basis, Eq. (18) gives a solvable system of linear equations for the CGCs. The number of independent solutions for the highest weight CGCs for a given irrep corresponds to its outer multiplicity in the decomposition. Since the quantities we calculate differentiate between copies of the same irrep, we construct an orthonormal basis of solutions.

- (3) Using the above $J_-^{(l)}$ lowering operators, the CGCs for the lower-weight states are constructed from the known CGCs of the higher-weight states. With the selection rules of the lowering and raising operators, states of the same weight have the same “parent” states that give rise to superpositions of those states once lowered. It then becomes a matter of solving linear sets of equations for their CGCs, as in the second step.

We have implemented this algorithm as a supplementary Python package to our primary simulation code base and plan to make it open-source in the future.

D. Truncations

Thus far the formalism described applies to general $SU(N)$ gauge group, apart from explicit examples given for $N=3$. Going forward we restrict attention to $SU(3)$.

We employ several types of truncations in order to map the full lattice $SU(3)$ theory to approximations of it that are more tractable for present quantum hardware.

For notational convenience we define T_i to refer to the i -index irrep truncation; e.g., $T_1 = \{1, 3, \bar{3}\}$ (one-index tensors), $T_2 = \{1, 3, \bar{3}, 6, \bar{6}, 8\}$ (two-index tensors), etc. These truncations can be viewed as bounding the electric energy density on a link, cf. (3).

In the T_i , no truncation of the site singlets is assumed beyond those imposed automatically by the link irrep truncation. In Table II, we report the numbers of unique singlet states per site in various dimensions and truncations. However, we will also make use of a further truncation T'_1 , where the irreps are truncated at $\{1, 3, \bar{3}\}$ and the site singlets are limited so that only two link irreps entering the site are allowed to be nontrivial at a time. This truncation can also be viewed as cutting off the electric energy density at a site.

These approximations to the full theory can be systematically improved in the obvious way. Their main value is to limit the size of the local Hilbert space associated with a plaquette, and thus, in turn, to limit the complexity of magnetic time evolution. However, these truncations may also fail to lower gate counts to a sufficient degree on present hardware. In this case it is useful to impose a further cutoff on the individual transitions induced by the magnetic Hamiltonian. We discuss a truncation scheme for the magnetic matrix elements further in Sec. IV.

E. Magnetic matrix element distributions

The plaquette operator matrix elements between physical plaquette states populate distributions. In Figs. 3 and 4 we show a sampling of these distributions. In the $d=3$ case, the precise matrix elements vary between the three different planes in which a plaquette can lie.⁴ However, the distributions are essentially the same and we show only one representative.

III. QUANTUM CIRCUITS

The time evolution operator for the Kogut-Susskind Hamiltonian is $U(t) = e^{-iHt}$. However, it is difficult and inefficient to create an exact quantum circuit that implements $U(t)$ on a digital quantum computer. Product formulas allow an arbitrary n -qubit unitary to be approximated by a product of unitaries. In this work we use the first-order Lie-Trotter approximation. Dividing the time t into N steps of size $\Delta t = t/N$, we have $e^{-iHt} = e^{-iH\Delta t} \dots e^{-iH\Delta t} + \mathcal{O}(\Delta t^2)$. Each $e^{-iH\Delta t}$ is as difficult to simulate exactly as e^{-iHt} , but the short time window Δt allows further Trotterization. We decompose

⁴This is because the F -ordering of links needed to define the basis renders the basis noninvariant under discrete rotations. Physical quantities like the expectation value of the electric energy on links will, however, exhibit discrete rotation invariance, and this serves as a nontrivial consistency check on the implementation.




FIG. 3. Distributions of magnitudes of matrix elements of the plaquette operator for the “line of plaquettes” lattice, with T_1 truncation (left) and T_2 truncation (right).




FIG. 4. Distributions of magnitudes of matrix elements of the plaquette operator for the $d = 2$ (left) and $d = 3$ (right) lattices with T_1 truncation.

$e^{-iH\Delta t} = \prod_h e^{-ih\Delta t} + \mathcal{O}(\Delta t^2)$, where each of the unitaries $e^{-ih\Delta t}$ is easier to simulate.

For the decomposition of each timestep, we use

$$e^{-iH\Delta t} \approx \prod_k e^{-iH_{B,k}\Delta t} \prod_j e^{-iH_{E,j}\Delta t}. \quad (19)$$

Here each k corresponds to a single Givens rotation on a single plaquette induced by the local magnetic Hamiltonian $\square_P + \square_P^\dagger$. Each j corresponds to a term in the Pauli string decomposition of the local electric Hamiltonian on a single link. These decompositions and their circuitizations will be described below.

A. Registers

First, we encode the degrees of freedom in quantum registers. There is a register dedicated to each of the N_s sites and N_ℓ links of the lattice. Link registers encode the irrep states of each link and site registers encode the singlet states of each site in the reduced electric basis.

We adopt a dense binary encoding of the singlet and irrep states. For example, with 3- and 2-qubit registers for sites and links, respectively, “000” encodes some specific singlet state and “10” encodes some specific irrep state. In general the site registers contain $n_s = \lceil \log_2(N_s) \rceil$ qubits and the link registers contain $n_r = \lceil \log_2(N_r) \rceil$ qubits, where N_s (N_r) is the number of singlets (irreps) of the truncation. Thus, the total number of qubits required is $N_s n_s + N_\ell n_r$.

We use a particular organization of the quantum registers based on the locations of the plaquettes on a lattice. Each plaquette is identified with a site; multiple plaquettes are associated with each site for $d > 2$. For a given site s , choose two unit vectors e_i and e_j with $i < j$. Then the

plaquette is defined by $s_1 = s$, $\ell_1 = \ell(s, e_i)$, and $\ell_4 = \ell(s, e_j)$. This identification avoids double-counting plaquettes as we sweep over the lattice. Unique site and link registers are then created sequentially, per plaquette, in the order $s_1 s_2 s_3 s_4 \ell_1 \ell_2 \ell_3 \ell_4$ as we iterate over all sites. Because the same sites and links may appear in adjacent plaquettes, only the first occurrence of a site/link in this process requires a quantum register. This means that plaquettes which are nearby in the lattice will typically share several link and vertex registers once encoded.

The precise encoding of the singlet and irrep states in register bitstrings is somewhat arbitrary. We choose an irrep encoding based on a lexicographic ordering on the i-weights of the underlying irreps. In the T_1 truncation, $(0, 0, 0) < (1, 0, 0) < (1, 1, 0)$, so $(0, 0, 0) := “00,”$ $(1, 0, 0) := “01,”$ and $(1, 1, 0) := “10.”$ Because $N_r = 3$ in T_1 , we require $n_r = 2$ qubits per encoded bitstring. The singlet states at each lattice site are completely specified by the irrep values on all the connected half-links along with a zero-indexed positive integer specifying a multiplicity of that singlet. We encode this multiplicity index in an additional register associated with each lattice site. If in a particular geometry and irrep truncation the maximal value which occurs for the zero-indexed multiplicity is n , then, each site register will require $\lceil \log_2(n + 1) \rceil$ qubits. In cases where $n = 0$ such as in the T'_1 truncation, no site registers are needed. This encoding generalizes straightforwardly to other truncations and dimensions.

The encoding is used to translate Hamiltonian matrix elements from transitions between singlets and irreps to transitions between bitstrings encoding either single-link or single-plaquette data. The operators $H_{E,j}$ of the electric Hamiltonian are diagonal in the electric basis and remain

diagonal in the computational basis after the encoding. Each one-plaquette term $H_{B,k}$ in the magnetic Hamiltonian acts (nondiagonally) on the states of four site registers, four “active” link registers, and $4n$ total “control” link registers, $|s_1\rangle|s_2\rangle|s_3\rangle|s_4\rangle|\ell_1\rangle|\ell_2\rangle|\ell_3\rangle|\ell_4\rangle|c_{1,1}\rangle\cdots|c_{4,n}\rangle$.⁵ The number of control link registers per site n depends on the size of the lattice and its boundary conditions, taking a maximal value of $n = 2(d-1)$ in d dimensions. The bitstring of the $8 + 4n$ combined registers specifies an electric basis state of a plaquette. Thus the terms in the magnetic Hamiltonian may be thought of as transitions between two particular plaquette bitstrings. The full-lattice Trotter step circuit is then the product of matrix exponentials of these terms.

Finally, the full-lattice Trotter step circuit is repeatedly appended to obtain the time evolution circuit $U(t)$. This circuit is run on a classical simulator or quantum device and all qubits are measured. Inverting the encoding above, the measured bitstrings can be translated back to singlet and irrep states on the lattice. This allows, for example, for the calculation of transition probabilities between electric basis states, as well as chromoelectric expectation values and equal-time correlation functions. We will also employ simple classical postprocessing error mitigation to deal with shots falling into unphysical bins.

B. Rotation circuits

The full-lattice Trotter step circuit is a product of $e^{-iH_{E,j}\Delta t}$ circuits and $e^{-iH_{B,k}\Delta t}$ circuits. Because each link register contains $\mathcal{O}(1)$ qubits, we use Pauli string decomposition to construct $e^{-iH_{E,j}\Delta t}$.⁶ However, the set of registers for a plaquette may contain tens of qubits, and $e^{-iH_{B,k}\Delta t}$ is better constructed with multicontrolled unitary (MCU) gates. In this section we describe our implementation of these circuits. We use the notation $\pi^{0,1}$ for projectors on the $|0\rangle$ and $|1\rangle$ states of a qubit, and $\sigma^\pm$ for the raising and lowering operators on a qubit.

1. Electric evolution

Since $H_{E,j}$ is diagonal in the electric basis, it can be decomposed into a sum of Pauli strings containing only I and Z . One way to think of the electric Hamiltonian at a link is in terms of projectors onto each irrep in the truncation.

For example, consider the term $H_{E,j} = \frac{g^2 E_j^2}{2a} \pi^0 \pi^1$. Here

$$\pi^0 = \frac{I+Z}{2} \quad \pi^1 = \frac{I-Z}{2}. \quad (20)$$

⁵The notation $|c_{i,j}\rangle$ denotes the j th control link connected to site s_i , with some arbitrary convention set to fix the ordering of controls at a particular site.

⁶For higher truncations, involving more qubits per link register, it is advantageous to implement a multicontrolled phase gate instead of Pauli decomposition.


This operator projects onto the irrep encoded by “10” with electric Casimir E_j^2 . Its Pauli decomposition is

$$H_{E,j} = \frac{g^2 E_j^2}{8a} (II - IZ + ZI - ZZ) \quad (21)$$

The identity string II is discarded because it amounts to a global phase. All I - Z strings commute, so we may write

$$e^{-iH_{E,j}\Delta t} = e^{-i\frac{\theta}{2}IZ} e^{-i\frac{\theta}{2}ZZ} e^{-i\frac{\theta}{2}ZI} \quad (22)$$

The rotation angles θ are all $\pm \frac{g^2 E_j^2}{4a} \Delta t$. Strings that contain one Z are simply an $R_Z(\theta)$ gate on that qubit. Strings that contain more than one Z require a parity-counting scheme, e.g.,



$$e^{-i\frac{\theta}{2}ZZZ} = \quad (23)$$

2. Magnetic evolution

To Trotterize the magnetic timestep at a plaquette, we expand the magnetic plaquette Hamiltonian in a sum of two-state transitions. In this decomposition, each operator $H_{B,k}$ is proportional to a (symmetrized) string of projector and ladder operators, e.g.,

$$\pi^1 \sigma^+ \sigma^- \pi^0 + \text{H.c.} \equiv |1100\rangle\langle 1010| + \text{H.c.} \quad (24)$$

In this form it is clear that each term in the decomposition encodes a transition between two plaquette states P_i and P_f . Thus each factor $e^{-iH_{B,k}\Delta t}$ in the Trotterization is a Givens rotation⁷ with a rotation angle proportional to $\langle P_f | \square + \square^\dagger | P_i \rangle$.

To construct the circuit for $e^{-iH_{B,k}\Delta t}$, we diagonalize $H_{B,k}$ using CX gates. To do this, we observe that terms in the magnetic Hamiltonian can be grouped into families of mutually commuting operators. A convenient choice is given by the *LP families* (L for “ladder” and P for “projector”). An *LP* family is designated by a choice of L or P for each qubit; e.g., $LLPLPPL\dots$, etc. On n qubits

⁷A Givens rotation rotates two components of a vector into each other. $e^{-iH_{B,k}\Delta t}$ rotates two computational basis states, corresponding to a complex Givens rotation on the computational Hilbert space. This is simply the most obvious way to implement a unitary given by an exponential of a matrix that has no particular structure. This approach is often used to implement magnetic evolution in an electric basis, see, e.g., [9,44]. An early adopter of Givens rotations was Jacobi, who used them in an iterative algorithm to diagonalize matrices [45].

there are 2^n commuting operators per LP family. Each magnetic Hamiltonian term can be grouped into the LP family corresponding to its arrangement of ladder and projector operators. For instance, $\pi^1 \sigma^- \pi^1 + \text{H.c.}$ belongs to the PLP family. Then, for each LP family, there is a circuit V that transforms all $H_{B,k}$ in the family into operators comprised of $n - 1$ projectors and a rotation on a single *pivot qubit* q' .⁸

For each $H_{B,k}$, we fix a convention where the pivot qubit q' is identified as the first qubit acted on by $\sigma^\pm$ (i.e. the first L in the LP family), when reading $H_{B,k}$ from left to right. The circuit V can then be read off of the LP family in the following way. For each qubit after q' , append a CX gate controlled on q' and targeted to each other qubit on which there is an L . These entangling CX gates cause the rotation of the pivot to propagate to all other L -qubits. Then $H_{B,k} \rightarrow V\tilde{H}_{B,k}V^\dagger$, where $\tilde{H}_{B,k}$ is a string of projectors and an X gate on q' (e.g. $\pi^1 X \pi^1 \pi^0$). Thus $e^{-iH_{B,k}\Delta t} = V e^{-i\tilde{H}_{B,k}\Delta t} V^\dagger$, where $e^{-i\tilde{H}_{B,k}\Delta t}$ is a multicontrolled




FIG. 5. Gives rotation circuit for magnetic evolution. The CX gates transform the q_2 ladder operator into a projector. Then $e^{-i\frac{\theta}{2}\pi^1 X \pi^1 \pi^0}$ is a multicontrolled $R_X(\theta)$ gate. Note that the magnetic Hamiltonian term shown here belongs to the $PLLP$ family and q_2 is q' .

$R_X(\theta)$ gate. An example is shown in Fig. 5. The rotation angle, $\theta = \frac{2}{2ag^2} \langle P_f | \square + \square^\dagger | P_i \rangle \Delta t$, is classically precomputed. Because multicontrolled $R_X(\theta)$ gates are not native on most current quantum hardware, it is useful to decompose them into multicontrolled X gates, which can then be decomposed into CX s:

(25)

C. Resources and optimization

The bulk of the gate depth in the Trotterized circuit comes from multicontrolled CX gates (C^nX) which generate the multicontrolled rotations (MCR). While native implementations of multicontrolled unitaries are a focus of research interest [46,47], it is currently more feasible to decompose them into single- and double-qubit unitaries. Of particular interest are the double-unitary CX or $CNOT$ gate and the non-Clifford unitary T , and as the number of matrix elements, plaquettes, and qubits per plaquette register grow, the CX and T gate costs grow as well. In an effort to reduce the depth of $e^{-iH\Delta t}$, we perform several optimizations.

1. Gray code ordering

In both electric and magnetic evolution circuits, there are CX gates that surround a rotation. It is possible to reduce these CX s by ordering the local Hamiltonian operators in a Gray code [48] fashion described below.

In the electric evolution, each $H_{E,j}$ contains a sum of Pauli strings. It is always possible to reorder these Pauli

⁸ V and q' depend on the LP family, but this dependence is suppressed in the notation.

strings such that their sequence returns a Gray code. In such a sequence, neighboring terms differ by one Pauli operator. For example, (IZ, ZZ, ZI) gives a Gray code sequence because IZ (ZI) differs from ZZ only on the left (right) operator. Another example of a Gray code sequence is $(IIZ, IZZ, IZI, ZZI, ZZZ, ZIZ, ZII)$. In the matrix exponential, $e^{-iH_{E,j}\Delta t}$, a Gray code ordering will ensure that at most $1CX$ is necessary between each $R_Z(\theta)$ gate.

Generally, the magnetic Hamiltonian will contain terms from many different LP families. The CX s present in each $e^{-iH_{B,k}\Delta t}$ appear in relation to the L operators of their respective LP family. This suggests an optimization done by ordering the LP families in Gray code. For example, (PLL, PLP, LLP, LPP) is a Gray code sequence of LP families. Of course, the number of LP families in the magnetic Hamiltonian can be much smaller than 2^n , but this is an optimal ordering of the families.

2. Control pruning

One of the simplest lattices is a single plaquette with open boundary conditions. In the T_1 irrep truncation, this lattice requires $n_r = 2$ qubits per link and no site qubits. Therefore, encoding this lattice requires 8 qubits, and the magnetic evolution circuit naïvely contains C^7X gates.

However, there are only three physical plaquette states. The 7 controls on the C^7X are ensuring that only two out of the $2^8 = 256$ possible basis states in the computational Hilbert space are rotated into each other, even though 253 of these states will never appear in a noiseless simulation. Typically, many of the 7 controls per C^7X can be removed, or “pruned,” without changing the action of the unitary operator on the physical subspace. Pruning in the classical circuit preparation stage can substantially lower the CX and T gate cost of the magnetic evolution, once the MCUs are decomposed into one and two qubit gates.

The control pruning algorithm requires knowledge of the set $\mathcal{P}$ of gauge-invariant plaquette states present in the given truncation scheme, the bitstrings P_i and P_f involved in the Givens rotation, and the LP family corresponding to the rotation.⁹ The output is the set of qubits Q in the Givens rotation that must be controlled on. The algorithm is as follows:

- (1) Choose a representative P from the Givens rotation bitstrings (either P_i or P_f .) Identify the pivot qubit q' associated with P (see Sec. III B 2 for the definition of the pivot.) Because the $R_X(\theta)$ gate of the Givens rotation is applied to q' , q' cannot be in Q , but any other qubit associated with the plaquette may appear in Q .
- (2) Using the diagonalizing circuit V described in Sec. III B 2, translate $\mathcal{P}$ into $\tilde{\mathcal{P}}$, the set of physical plaquette states after applying V . For example, if the family is $LLPP$, then V takes $1100 \rightarrow 1000$, $0000 \rightarrow 0000$, etc. Similarly, relabel the representative $P \rightarrow \tilde{P}$. (This step is necessary because the multicontrolled $R_X(\theta)$ in a Givens rotation circuit operates on the states in $\tilde{\mathcal{P}}$ rather than $\mathcal{P}$.)
- (3) Iterate once through $\tilde{\mathcal{P}}$, identifying any bitstrings (excluding the pivot qubit q') with Hamming distance one from $\tilde{P}$. These qubits are a subset of the necessary control qubits; without them, these bitstrings would be erroneously rotated by the $MCRX$. Collect these qubits into the set Q , and eliminate all bitstrings from $\tilde{\mathcal{P}}$ that differ from $\tilde{P}$ at any qubits in Q : these strings are vetoed by controls on the Q .¹⁰
- (4) Iterate through the remaining bitstrings of $\tilde{\mathcal{P}}$. For each bitstring, populate the list of qubits (excluding the pivot) at which the bitstring differs from $\tilde{P}$. At the end of this iteration, find the qubit q that appeared most frequently in these lists. (If there are multiple, choose one.) Append q to Q , and eliminate from $\tilde{\mathcal{P}}$ all bitstrings that differ from $\tilde{P}$ at q .
- (5) Repeat step 4 until $\tilde{\mathcal{P}}$ only contains $\tilde{P}_i$ and $\tilde{P}_f$.

⁹This data is automatically collected during the classical construction of $\square$ matrix elements.

¹⁰While this step is not strictly necessary, it reduces the number of loops required in the final steps of the algorithm.




FIG. 6. Example of a fusion of multicontrolled unitaries.

The algorithm terminates with the set Q , the set of control qubits necessary to distinguish P_i and P_f from all other states in $\mathcal{P}$. The number of controls scales logarithmically with the number of plaquette states: $|Q| \approx \lceil \log_2(|\mathcal{P}|) \rceil$.

3. Control fusion

A final optimization of the multicontrolled unitaries is possible for a given LP family. If there are multiple terms $H_{B,k}$ of the magnetic Hamiltonian that correspond to the same LP family and they possess the same $\langle P_f | \square + \square^\dagger | P_i \rangle$ coefficient, then it may be possible to fuse these terms in powers of two. This is due to the property shown in Fig. 6.

4. Quantum resource estimates

Here we provide rough estimates for the quantum resources required to run simulations, in terms of CX and T gate counts.

- (i) For a C^nX gate with n controls, we quote the result of [49], which found a decomposition in terms of $12n - 32$ CX gates and $16n - 48$ T gates for $n \geq 6$, without ancillas. With ancillas, we quote the result of [50], which found a decomposition in terms of $6n - 18$ CX gates and $8n - 25$ T gates.
- (ii) If there are N_M terms in the magnetic Hamiltonian, then there are $2N_M$ C^nX gates in a Trotter circuit without optimization.¹¹ The CXs between each term add no more than $2nN_M$ CXs. Neglecting the electric evolution which contributes a small number of CX and further optimizations we obtain the estimates tabulated in Table III.
- (iii) Let $N_\square$ be the number of physical plaquette states in the simulation. Control pruning causes each C^nX to be controlled on order $\lceil \log_2(N_\square) \rceil$ qubits. If control pruning is the only optimization used, then n in the points above can be replaced by $\lceil \log_2(N_\square) \rceil$.
- (iv) If control fusion and Gray code orderings are used, then the resources can vary further. Control fusion reduces both N_M and n . Gray code orderings linearly reduce the CX cost.

¹¹Note that depending on implementation, one may be able to decompose the multicontrolled rotation directly into CX rather than into two MCX [50].


TABLE II. The number of unique singlets found at a site.

d	T_1	T_2
3/2	4	16
2	6	66
3	28	1646

TABLE III. General gate cost estimates per plaquette per Trotter step with no optimizations. $n + 1$ is the number of qubits needed to encode a plaquette state.

Gate	CX	T
W/o ancillas	$2N_M(13n - 32)$	$2N_M(16n - 48)$
W/ ancillas	$2N_M(7n - 18)$	$2N_M(8n - 25)$

Table IV summarizes gate counts and resource requirements for a variety of small lattices, based on explicit circuit construction, and illustrates the improvements afforded by the pruning and fusion optimizations. Simulation results for several of these lattices are reported in the next section. As seen also in Fig. 7, using ancilla in the MCX decomposition reduces gate depth by about half for any lattice. Control pruning has a larger impact for lattices with a larger control scheme, with almost a five-fold reduction for the $2 \times 2 \times 2$ lattice. The effect of pruning can be seen in the decrease of the number of control qubits in Table IV. Observe that

FIG. 7. Reduction of gate depth for CX count (solid) and T count (dashed) for three of the lattices described in Table IV.

Gray-ordering and control fusion impact lattices with a large number of transitions relative to the size of their control scheme, but have negligible impact for lattices with strong truncations on the number of physical states. The effects of control fusion can be seen in the decrease of MCX after optimization.

IV. SIMULATIONS

In this section we report results for noiseless simulations and noisy emulations of Quantumum systems in $d = 3/2$, $d = 2$, and $d = 3$. Throughout, we begin in the electric

TABLE IV. Gate cost comparison between unoptimized and optimized circuits. “Trans(itions)” refers to the number of MCR per trotter step per plaquette (equivalently, it is the number of nonzero matrix elements of $\square$). “Ctrl” is the average number of controls per MCX . “Unopt” means that each MCR is directly transpiled into CX and T and then combined. “Opt” means that each MCR has been pruned, then Gray-ordered, then fused, then implemented using ancilla qubits and the v-chain method [51]. n -plaq are $d = 3/2$ lattices of length n with PBC (Fig. 8). $n \times n$ are $d = 2$ lattices of size n^2 with PBC (Fig. 12). 1-cube is a $d = 3$ lattice with open boundary conditions (Fig. 16). $n \times n \times n$ are lattices of size n^3 with PBC. All-to-all connectivity is assumed in CX counts. Table made with [52].

	Lattice	T_i	Trans	Qubits	MCX	Ctrl	CX	T
Unopt opt	2-plaq	T_1	27	12	108	11	1.20×10^4	1.38×10^4
				15	108	4	2.80×10^3	3.10×10^3
	3-plaq	T_1	81	18	486	15	7.92×10^4	9.33×10^4
				23	486	4	1.25×10^4	1.40×10^4
	5-plaq	T_1	81	30	810	15	1.32×10^5	1.56×10^5
				35	810	4	2.05×10^4	2.30×10^4
	2-plaq	T_2	3798	22	15192	21	3.66×10^6	4.38×10^6
				33	14912	9	8.10×10^5	9.40×10^5
	2×2	T_1	835	20	6680	19	1.44×10^6	1.71×10^6
				27	5104	7	2.20×10^5	2.50×10^5
	3×3	T'_1	83	36	1494	23	3.99×10^5	4.78×10^5
				46	1494	8	7.25×10^4	8.40×10^4
	1-cube	T_1	81	24	972	15	1.58×10^5	1.87×10^5
				29	972	4	2.45×10^4	2.75×10^4
	$2 \times 2 \times 2$	T'_1	127	48	6096	31	2.26×10^6	2.73×10^6
				66	6096	8	2.90×10^5	3.40×10^5
	$3 \times 3 \times 3$	T'_1	707	162	114534	39	5.44×10^7	6.60×10^7
				180	114534	9	6.19×10^6	7.22×10^6

vacuum state and measure the expectation value of the electric field $\langle E^2 \rangle$ averaged over the lattice.

Due to the rapid growth in the number of physical plaquette states and magnetic matrix elements with dimension and truncation, it is essential to understand which matrix elements contribute significantly to a given simulation. Matrix elements which contribute insignificantly to a noiseless simulation still contribute noise on a real device, and cost computation time. Therefore, a main focus in this section will be to develop assessments of significance (“scores”) for the precomputed matrix elements, and test additional truncations based on these scores.

A. $d = 3/2$

The “line of plaquettes” lattice, in the local basis, with T_1 truncation, two plaquettes, and periodic boundary conditions along the line, was studied in [9]. The lattice is shown in Fig. 8. Figure 9 compares the Trotter evolution of the electric vacuum in the T_1 and T_2 truncations on this lattice at couplings $g = 1$ and $g = 2$.¹² We observe that including two-index tensor irreps has a much smaller impact at stronger coupling, reflecting larger gaps.¹³

On this lattice, with the T_1 irrep truncation, the plaquette operator can induce 27 different transitions between physical plaquette states. To assess the relative importance of each transition, we assign a score

$$\text{Score}(i \leftrightarrow j, t) = \max(|\langle j | (\square + \square^\dagger) | i \rangle| P(i, t) dt, |\langle j | (\square + \square^\dagger) | i \rangle| P(j, t) dt). \quad (26)$$

Here $P(i, t)$ is the probability that a given plaquette is in physical state i at time t . Thus the score measures the importance of an $i \leftrightarrow j$ transition at a given time, weighing both the amplitude of the transition and the probability of being in one of the states in the transition. To get an accurate sample of the significant transitions involved in time evolution up to some time T , we sample scores at several values of $0 < t < T$ and take the maximum score for each matrix element over the sample.

Figure 10 shows the distribution of (nonzero) matrix element scores, for $g = 1$ and $g = 2$ and $T = 3$, sampling scores over five equally spaced times. Each bar represents a matrix element. Notably, only 9 out of 27 matrix elements actually participate; others describe transitions between physical states that cannot be reached by time evolution

¹²We have checked that the matrix elements used for the T_1 truncation in $d = 3/2$ agree with those of [9]. The left-hand plot of Fig. 9 agrees with Figs. 12 (14) of [9], after rescaling by 3 (6) to match the whole-lattice electric energy (whole-lattice $\langle E^2 \rangle$).

¹³Recent work [53] has studied a general phenomenon of “fragmentation” of the lattice gauge theory Hilbert space, where transitions between sectors of significantly different electric flux are effectively decoupled at larger g . This is a distinct phenomenon with similar effects.




FIG. 8. The 2-plaquette lattice with PBCs.

from the electric vacuum. Among the 9, most are numerically insignificant, and the gap between scores grows with g . Figure 11 shows the effect of truncating transitions based on score. In this case, keeping only the first 3 matrix elements—transitions away from the vacuum and a transition involving the first set of excited states—still gives an excellent approximation to the evolution with all matrix elements.

We emphasize that matrix element truncation based on the score introduced above is not a scalable algorithm, since it takes input from classical simulations that are limited to $\lesssim 40$ qubits. It is also dependent on the initial state and duration of evolution, and may have different performance for different observables. (For observables $\mathcal{O}$ that are diagonal in the electric basis, a variant of the algorithm is to weight the scores by $\mathcal{O}(i)$ or $\mathcal{O}(j)$ for an $i \leftrightarrow j$ transition.) We introduce the score cut only to illustrate the point that not all matrix elements contribute equally importantly to the evolution of a given state. Exploiting this point is likely to be essential in scaling simulations to larger irrep truncations and dimensions, where the number of physical plaquette states and transitions grows rapidly. A more elegant and scalable version of the algorithm would attempt to construct a score without requiring full classical simulations. We will not construct such an algorithm here.¹⁴ Below, we will assess the same scoring and truncation algorithm in $d = 2$, while in $d = 3$ we will revert to a simpler truncation on matrix element magnitudes alone.

B. $d = 2$

We turn now to the simplest 2D lattice, 2×2 plaquettes with periodic boundary conditions. The lattice is shown in Fig. 12.

We apply the plaquette matrix element scoring procedure described above for $d = 3/2$. The left-hand panel of Fig. 13 shows the distribution of scores for $T = 3$ again taken over 5 samples. Similar to $d = 3/2$, out of the 835 matrix elements describing the evolution of 2×2 -plaquette, only 247 describe transitions that can be reached from time evolution from the vacuum. Again we observe clear hierarchies among the matrix elements. Figure 14 compares the Trotter evolution of the electric vacuum in the T_1 truncations at $g = 1$. Trotter error for $N_t = 1$ becomes significant at times of order one in lattice units. For earlier

¹⁴Previous works have examined cuts on matrix element magnitudes alone, see, e.g., [54]. Ideally, a more scalable algorithm will also account for the “cost” of getting to the states i or j , as the score does.




FIG. 9. Time evolution of the link-averaged electric Casimir operator, starting from the electric vacuum on the $d = 3/2$ two-plaquette lattice. Different curves correspond to different numbers of Trotter steps N_t and irrep truncations T_1 and T_2 . Left: $g = 1$. Right: $g = 2$. “Exact” curves refer to results obtained by exact diagonalization. For larger coupling, the impact of the irrep truncation is suppressed.




FIG. 10. Matrix element scores for the $d = 3/2$ two-plaquette lattice with T_1 irrep truncation, and $g = 1$ (left) and $g = 2$ (right).




FIG. 11. Time evolution of the link-averaged electric Casimir operator, starting from the electric vacuum on the $d = 3/2$ two-plaquette lattice in the T_1 irrep truncation. Different curves correspond to different numbers of Trotter steps N_t and truncations on the magnetic matrix element scores shown in Fig. 10; only the elements with the B largest scores are kept in the truncation. Left: $g = 1$. Right: $g = 2$.




FIG. 12. The 2×2 -plaquette lattice with PBCs.

times, a truncation to the transitions with the $B = 5$ largest scores gives a good match to the exact evolution. For $N_t = 5$, reasonable agreement with the exact evolution is obtained when retaining on order of tens of the dominant transitions.

The scoring algorithm has a time duration parameter T . Scores are determined by sampling over a set of times up to T . The relative insensitivity of the scores to T is illustrated




FIG. 13. Left: matrix element scores for the 2×2 -plaquette lattice with $g = 1$ and T_1 irrep truncation. Elements are ordered by descending score. Right: matrix element scores for the same lattice and parameters, with two different time ranges ($T = 1$, blue, and $T = 10$, orange) used in the sampling. Here the matrix element ordering is arbitrary but fixed between the two values of T . Substantial overlap (purple) indicates that the scoring is largely insensitive to the sampling duration.




FIG. 14. Time evolution of the link-averaged electric Casimir operator, starting from the electric vacuum on the 2×2 -plaquette lattice in the T_1 irrep truncation. Different curves correspond to different numbers of Trotter steps N_t and truncations on the magnetic matrix element scores shown in Fig. 13; only the elements with the B largest scores are kept in the truncation.

in the right-hand panel of Fig. 13. We see that only a few of the scores evolve appreciably as the duration is increased by an order of magnitude.

1. Noisy emulation

The circuits developed in this work are best suited for quantum platforms with substantial qubit connectivity: although the Hamiltonian is local in the sense of coupling nearest-neighbor links and plaquettes, the plaquette Hilbert space can span tens of qubits, which are all coupled together. In Fig. 15 we report results of noisy simulations of the 2×2 PBC lattice of Fig. 12, with T_1 irrep truncation and $B = 3$ transition truncation, using the Quantinuum System Model H2 emulator [38]. H2 is a 56 qubit trapped ion device with all-to-all connectivity. The emulator allows scaling the noise levels and we present results with all noise sources uniformly dialed to 1% and 10%. As a mild form of error mitigation, we veto shots that go outside the physical link Hilbert space (with irrep truncation T_1 , the two-qubit 11 bitstring is not assigned an irrep), and then to veto shots that fail the Gauss law at some vertex. At 1% the noise is insignificant for this circuit. Above 10% of percent, the noise is not effectively mitigated by the simple mitigation strategy.




FIG. 15. Quantum H2 emulator results for the evolution of the electric vacuum, $g = 1$, and one Trotter step on the 2×2 lattice with PBCs and the T_1 irrep truncation. Ancillas are used in the decomposition of multicontrolled gates to reduce two-qubit gate counts. We apply a rudimentary form of error mitigation, vetoing shots that correspond to unphysical states. Dashed curves correspond to different scale factors applied to the noise. Error bars on these results reflect shot noise. (200 shots were used for the first four “1%” points, increased to 1000 shots for the final point to reduce noise. 1000 shots were used for all “10%” points.) The black curve corresponds to the noiseless simulation.

C. $d=3$

Although we are able to build circuits for larger lattices in $d = 3$, our ability to execute noiseless state vector simulations is limited to $n_{\text{qubit}} \lesssim 40$. A simple lattice where we can perform simulations is the single cube with open boundary conditions, shown in Fig. 16.

In this case, we opt for a simpler truncation of magnetic transitions, cutting on the magnitude of matrix elements alone, rather than the score defined in Eq. (26). Results for $g = 1$, T_1 truncation are shown in the top panel of Fig. 17. We also examine the T'_1 singlet truncation in the bottom panel of Fig. 17. We find that for this lattice the full T'_1 simulation (17 transitions) well approximates the full T_1 truncation (81 transitions).




FIG. 16. The cube lattice with OBCs.




FIG. 17. Link-averaged electric Casimir for the cube lattice with OBC in T_1 (top) and T'_1 (bottom). Different curves correspond to varying the number of Trotter steps and the cutoff on magnetic matrix elements (N_t , cutoff). Top: a cutoff of 0 means all 81 magnetic matrix elements were used. Cutoffs of (0.3,0.5,1) correspond to (71,27,3) matrix elements, respectively. Bottom: a cutoff of 0 means all 17 magnetic matrix elements were used. Cutoffs of (0.3,1) correspond to (15,3) matrix elements, respectively.




FIG. 18. Vacuum-persistence probability (left) and link-averaged electric Casimir (right) for the $2 \times 2 \times 2$ (24-plaquette) lattice with PBCs and T'_1 truncation, computed using matrix product state simulation, one Trotter step, with truncation threshold set to $\epsilon = 10^{-6}$ (red) and 10^{-8} (blue, dashed). The curves are essentially coincident.

1. Tensor networks

Higher n_{qubit} can be accessed in classical simulations using tensor network approximations. We perform preliminary analysis of a $2 \times 2 \times 2$ (24 plaquette) lattice with PBCs and a maximum-element truncation of the magnetic transitions, using the out-of-the-box matrix product state (MPS) method included in Qiskit. This is a 48 qubit lattice, out of range of statevector simulation. Figure 18 shows the electric Casimir at $g = 1$ for a single Trotter step and TN thresholds of 10^{-6} and 10^{-8} . We observed convergence for threshold $\lesssim 10^{-6}$.

The 48 qubit implementation of the circuit does not use ancillas in the MCU decomposition. Consequently it is a deep circuit, well out of range of current NISQ hardware. With ancillas, the depth is reduced, but the qubit count (66) exceeds that of current devices with all-to-all connectivity. It represents an interesting target for next-generation hardware. This lattice at higher N_t , and the $3 \times 3 \times 3$ lattice, may be useful benchmarks for comparing classical TN methods to quantum simulations. Further tensor treatments of the circuits constructed here will be presented elsewhere.

V. FUTURE DIRECTIONS

We close with a brief summary of directions for future development.

On the physics front, it is necessary to incorporate quarks and a θ term; to develop ground state preparation methods; to implement measurement of expectation values of plaquette operators; and to study the continuum limit. For example, unlike the pure gauge analysis here, inclusion of quarks will require allowing gauge links at sites to fuse into nontrivial irreps. Matrix elements of localized operators in the Hamiltonian can still be classically precomputed using generalized Clebsch-Gordan coefficients. Similarly, minimizing the resource requirements to approach the continuum limit will require the addition of improvement terms [55] (for a quantum simulation study, see [56].) In the approach taken here, incorporating at least classical

improvement terms amounts to computing master formulas for the matrix elements of some additional operators.

Minimizing the quantum resource requirements for physically relevant simulations will require further advances in circuit compression and development of new approximation schemes. We have begun preliminary explorations of some algorithms; for example, binning magnetic rotations by magnitude and rebasing the states so that the rotations in a bin can be fused into a smaller set of multicontrolled unitaries. This will be reported on elsewhere. We have also described here some crude methods for assessing “which matrix elements really matter” in a given simulation. In our examples the answer is generally “much less than all of them,” but clearly there is much more to be understood about this structure and how to determine it *a priori*.

There are several natural directions for extending the functionality of our modular circuit generation code. These include extending circuit-generation capability from $SU(3)$ to $SU(N)$ utilizing the methods outlined in Sec. II C, integrating the precomputation of matrix elements into the simulation pipeline to support dynamic circuit generation in arbitrary lattice dimensionalities and truncations, and the implementation of higher-order Trotterization in simulation circuits.

Finally, as is well-known but abundantly clear in the circuits we have constructed, qubit platforms with one- and two-qubit gates furnish sufficient but quite unnatural architectures for simulating lattice gauge theories. Qudit systems, or qubits with native multicontrolled gates, could offer significant reductions in circuit depth for these kinds of simulation problems. It would be interesting to sharpen the precise hardware advances that would be most impactful.

ACKNOWLEDGMENTS

We thank Aida El-Khadra for helpful conversations and collaboration in the early stages of this work. This work was supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics Quantum Information Science Enabled Discovery (QuantISED) program. We gratefully acknowledge BlueQubit [57] for enabling classical state vector simulations presented in this work. We also gratefully acknowledge the use of Quantinuum H1-1 and H2-1 emulators [38]. This work used the Delta [58] system at the National Center for Supercomputing Applications through allocation PHY230137 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants No. 2138259, No. 2138286, No. 2138307, No. 2137603, and No. 2138296. The Delta advanced computing and data resource is supported by the National Science Foundation (Award No. OAC 2005572) and the State of Illinois. Delta is a joint effort of the University of Illinois Urbana-Champaign and its National Center for Supercomputing Applications.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX: MASTER FORMULA

$$\begin{aligned}
\langle \omega_n | \square \{P(\vec{s}, \vec{e}_i, \vec{e}_j)\} | \omega_m \rangle &= \sqrt{\frac{\dim(R_{\ell_0}^m) \dim(R_{\ell_1}^m) \dim(R_{\ell_2}^m) \dim(R_{\ell_3}^m)}{\dim(R_{\ell_0}^n) \dim(R_{\ell_1}^n) \dim(R_{\ell_2}^n) \dim(R_{\ell_3}^n)}} \prod_{s \notin \mathcal{S}(P)} \delta_{\Gamma_s^n, \Gamma_s^m} \prod_{\ell \notin \mathcal{L}(P)} \delta_{R_\ell^n, R_\ell^m} \\
&\times \sum_{\sigma_1=1}^{\dim(f)} \sum_{\lambda_{\ell_1}^m=1}^{\dim(R_{\ell_1}^m)} \sum_{\lambda_{\ell_1}^n=1}^{\dim(R_{\ell_1}^n)} \sum_{\lambda_{\ell_4}^m=1}^{\dim(R_{\ell_4}^m)} \sum_{\lambda_{\ell_4}^n=1}^{\dim(R_{\ell_4}^n)} \sum_{\vec{\chi} \in \vec{\mathcal{C}}_{s_1}} \phi(\sigma_1) \phi(\lambda_{\ell_1}^m) \phi(\lambda_{\ell_1}^n) \phi(\lambda_{\ell_4}^m) \phi(\lambda_{\ell_4}^n) \\
&\times \langle R_{\ell_1}^n, \lambda_{\ell_1}^n | (R_{\ell_1}^m, \lambda_{\ell_1}^m) \otimes (f, \sigma_1) \rangle \langle R_{\ell_4}^n, \lambda_{\ell_4}^n | (R_{\ell_4}^m, \lambda_{\ell_4}^m) \otimes (\bar{f}, \bar{\sigma}_1) \rangle \\
&\times \langle 1, \Gamma_{s_1}^m | (\bar{R}_{\ell_1}^m, \tilde{\lambda}_{\ell_1}^m) \otimes (\bar{R}_{\ell_4}^m, \tilde{\lambda}_{\ell_4}^m) \otimes (\vec{\mathcal{C}}_{s_1}, \vec{\chi}) \rangle \langle 1, \Gamma_{s_1}^n | (\bar{R}_{\ell_1}^n, \tilde{\lambda}_{\ell_1}^n) \otimes (\bar{R}_{\ell_4}^n, \tilde{\lambda}_{\ell_4}^n) \otimes (\vec{\mathcal{C}}_{s_1}, \vec{\chi}) \rangle \\
&\times \sum_{\sigma_2=1}^{\dim(f)} \sum_{\lambda_{\ell_1}^m=1}^{\dim(R_{\ell_1}^m)} \sum_{\lambda_{\ell_1}^n=1}^{\dim(R_{\ell_1}^n)} \sum_{\lambda_{\ell_2}^m=1}^{\dim(R_{\ell_2}^m)} \sum_{\lambda_{\ell_2}^n=1}^{\dim(R_{\ell_2}^n)} \sum_{\vec{\chi} \in \vec{\mathcal{C}}_{s_2}} \phi(\lambda_{\ell_2}^m) \phi(\lambda_{\ell_2}^n) \\
&\times \langle R_{\ell_1}^n, \lambda_{\ell_1}^n | (R_{\ell_1}^m, \lambda_{\ell_1}^m) \otimes (f, \sigma_2) \rangle \langle R_{\ell_2}^n, \lambda_{\ell_2}^n | (R_{\ell_2}^m, \lambda_{\ell_2}^m) \otimes (f, \sigma_2) \rangle \\
&\times \langle 1, \Gamma_{s_2}^m | (R_{\ell_1}^m, \lambda_{\ell_1}^m) \otimes (\bar{R}_{\ell_2}^m, \tilde{\lambda}_{\ell_2}^m) \otimes (\vec{\mathcal{C}}_{s_2}, \vec{\chi}) \rangle \langle 1, \Gamma_{s_2}^n | (R_{\ell_1}^n, \lambda_{\ell_1}^n) \otimes (\bar{R}_{\ell_2}^n, \tilde{\lambda}_{\ell_2}^n) \otimes (\vec{\mathcal{C}}_{s_2}, \vec{\chi}) \rangle
\end{aligned}$$

$$\begin{aligned}
& \times \sum_{\sigma_f=1}^{\dim(f)} \sum_{\lambda_{\ell_3^+}^m=1}^{\dim(R_{\ell_3}^m)} \sum_{\lambda_{\ell_3^+}^n=1}^{\dim(R_{\ell_3}^n)} \sum_{\lambda_{\ell_2^+}^m=1}^{\dim(R_{\ell_2}^m)} \sum_{\lambda_{\ell_2^+}^n=1}^{\dim(R_{\ell_2}^n)} \sum_{\vec{\chi} \in \vec{C}_{s_3}} \phi(\sigma_3) \\
& \times \langle R_{\ell_3}^n, \lambda_{\ell_3^+}^n | (R_{\ell_3}^m, \lambda_{\ell_3^+}^m) \otimes (\tilde{f}, \tilde{\sigma}_3) \rangle \langle R_{\ell_2}^n, \lambda_{\ell_2^+}^n | (R_{\ell_2}^m, \lambda_{\ell_2^+}^m) \otimes (f, \sigma_3) \rangle \\
& \times \langle 1, \Gamma_{s_3}^m | (R_{\ell_3}^m, \lambda_{\ell_3^+}^m) \otimes (R_{\ell_2}^m, \lambda_{\ell_2^+}^m) \otimes (\vec{C}_{s_3}, \vec{\chi}) \rangle \langle 1, \Gamma_{s_3}^n | (R_{\ell_3}^n, \lambda_{\ell_3^+}^n) \otimes (R_{\ell_2}^n, \lambda_{\ell_2^+}^n) \otimes (\vec{C}_{s_3}, \vec{\chi}) \rangle \\
& \times \sum_{\sigma_4=1}^{\dim(f)} \sum_{\lambda_{\ell_3^+}^m=1}^{\dim(R_{\ell_3}^m)} \sum_{\lambda_{\ell_3^+}^n=1}^{\dim(R_{\ell_3}^n)} \sum_{\lambda_{\ell_4^+}^m=1}^{\dim(R_{\ell_4}^m)} \sum_{\lambda_{\ell_4^+}^n=1}^{\dim(R_{\ell_4}^n)} \sum_{\vec{\chi} \in \vec{C}_{s_4}} \phi(\lambda_{\ell_3^+}^m) \phi(\lambda_{\ell_3^+}^n) \\
& \times \langle R_{\ell_3}^n, \lambda_{\ell_3^+}^n | (R_{\ell_3}^m, \lambda_{\ell_3^+}^m) \otimes (\tilde{f}, \tilde{\sigma}_4) \rangle \langle R_{\ell_4}^n, \lambda_{\ell_4^+}^n | (R_{\ell_4}^m, \lambda_{\ell_4^+}^m) \otimes (\tilde{f}, \tilde{\sigma}_4) \rangle \\
& \times \langle 1, \Gamma_{s_4}^m | (\bar{R}_{\ell_3}^m, \tilde{\lambda}_{\ell_3^+}^m) \otimes (R_{\ell_4}^m, \lambda_{\ell_4^+}^m) \otimes (\vec{C}_{s_4}, \vec{\chi}) \rangle \langle 1, \Gamma_{s_4}^n | (\bar{R}_{\ell_3}^n, \tilde{\lambda}_{\ell_3^+}^n) \otimes (R_{\ell_4}^n, \lambda_{\ell_4^+}^n) \otimes (\vec{C}_{s_4}, \vec{\chi}) \rangle. \quad (A1)
\end{aligned}$$

-
- [1] J. B. Kogut and L. Susskind, Hamiltonian formulation of Wilson's lattice gauge theories, *Phys. Rev. D* **11**, 395 (1975).
- [2] J. Bulava, R. Briceño, W. Detmold, M. Döring, R. G. Edwards, A. Francis, F. Knechtli, R. Lewis, S. Prelovsek, S. M. Ryan, A. Rusetsky, S. R. Sharpe, A. Szczepaniak, C. E. Thomas, M. L. Wagman, and M. Wagner, Hadron spectroscopy with lattice QCD, [arXiv:2203.03230](#).
- [3] A. S. Meyer, A. Walker-Loud, and C. Wilkinson, Status of lattice QCD determination of nucleon form factors and their relevance for the few-GeV neutrino program, *Annu. Rev. Nucl. Part. Sci.* **72**, 205 (2022).
- [4] A. S. Kronfeld *et al.* (USQCD Collaboration), Lattice QCD and particle physics, [arXiv:2207.07641](#).
- [5] Y. Aoki *et al.* (Flavour Lattice Averaging Group (FLAG) Collaboration), FLAG review 2024, [arXiv:2411.04268](#).
- [6] C. W. Bauer, Z. Davoudi, N. Klco, and M. J. Savage, Quantum simulation of fundamental particles and forces, *Nat. Rev. Phys.* **5**, 420 (2023).
- [7] C. W. Bauer *et al.*, Quantum simulation for high-energy physics, *PRX Quantum* **4**, 027001 (2023).
- [8] A. Di Meglio *et al.*, Quantum computing for high-energy physics: State of the art and challenges, *PRX Quantum* **5**, 037001 (2024).
- [9] A. Ciavarella, N. Klco, and M. 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).
- [10] M. S. Alam, S. Hadfield, H. Lamm, and A. C. Y. Li (SQMS Collaboration), Primitive quantum gates for dihedral gauge theories, *Phys. Rev. D* **105**, 114501 (2022).
- [11] Z. Davoudi, A. F. Shaw, and J. R. Stryker, General quantum algorithms for Hamiltonian simulation with applications to a non-Abelian lattice gauge theory, *Quantum* **7**, 1213 (2023).
- [12] C. W. Bauer, I. D'Andrea, M. Freytsis, and D. M. Grabowska, A new basis for Hamiltonian SU(2) simulations, *Phys. Rev. D* **109**, 074501 (2024).
- [13] J. Zhang, S. W. Tsai, and Y. Meurice, Critical behavior of lattice gauge theory Rydberg simulators from effective Hamiltonians, *Phys. Rev. D* **110**, 034513 (2024).
- [14] Z. Davoudi, C.-C. Hsieh, and S. V. Kadam, Scattering wave packets of hadrons in gauge theories: Preparation on a quantum computer, *Quantum* **8**, 1520 (2024).
- [15] D. M. Grabowska, C. F. Kane, and C. W. Bauer, A fully gauge-fixed SU(2) Hamiltonian for quantum simulations, *Phys. Rev. D* **111**, 114516 (2025).
- [16] H. Lamm, Y.-Y. Li, J. Shu, Y.-L. Wang, and B. Xu, Block encodings of discrete subgroups on a quantum computer, *Phys. Rev. D* **110**, 054505 (2024).
- [17] B. Assi and H. Lamm, Digitization and subduction of SU(N) gauge theories, *Phys. Rev. D* **110**, 074511 (2024).
- [18] I. M. Burbano and C. W. Bauer, Gauge loop-string-hadron formulation on general graphs and applications to fully gauge fixed Hamiltonian lattice gauge theory, [arXiv:2409.13812](#).
- [19] D. M. Kırkcıoğlu, H. Lamm, and A. Maestri, Qudit gate decomposition dependence for lattice gauge theories, [arXiv:2410.16414](#).
- [20] T. Jakobs, M. Garofalo, T. Hartung, K. Jansen, J. Ostmeier, S. Romiti, and C. Urbach, Dynamics in Hamiltonian lattice gauge theory: Approaching the continuum limit with partitionings of SU(2), [arXiv:2503.03397](#).
- [21] V. Ale, N. M. Bauer, R. G. Jha, F. Ringer, and G. Siopsis, Quantum computation of SU(2) lattice gauge theory with continuous variables, *J. High Energy Phys.* **06** (2025) 084.
- [22] J. C. Halimeh, M. Hanada, S. Matsuura, F. Nori, E. Rinaldi, and A. Schäfer, A universal framework for the quantum simulation of Yang-Mills theory, [arXiv:2411.13161](#).
- [23] P. Fontana, M. M. Rianza, and A. Celi, An efficient finite-resource formulation of non-Abelian lattice gauge theories beyond one dimension, [arXiv:2409.04441](#).
- [24] T. Byrnes and Y. Yamamoto, Simulating lattice gauge theories on a quantum computer, *Phys. Rev. A* **73**, 022328 (2006).

- [25] S. V. Kadam, I. Raychowdhury, and J. R. Stryker, Loop-string-hadron formulation of an SU(3) gauge theory with dynamical quarks, *Phys. Rev. D* **107**, 094513 (2023).
- [26] A. N. Ciavarella, Quantum simulation of lattice QCD with improved Hamiltonians, *Phys. Rev. D* **108**, 094513 (2023).
- [27] R. C. Farrell, I. A. Chernyshev, S. J. M. Powell, N. A. Zemlevskiy, M. Illa, and M. J. Savage, Preparations for quantum simulations of quantum chromodynamics in $1+1$ dimensions. I. Axial gauge, *Phys. Rev. D* **107**, 054512 (2023).
- [28] Y. Y. Atas, J. F. Haase, J. Zhang, V. Wei, S. M.-L. Pfaendler, R. Lewis, and C. A. Muschik, Simulating one-dimensional quantum chromodynamics on a quantum computer: Real-time evolutions of tetra- and pentaquarks, *Phys. Rev. Res.* **5**, 033184 (2023).
- [29] A. Kan and Y. Nam, Lattice quantum chromodynamics and electrodynamics on a universal quantum computer, *arXiv:2107.12769*.
- [30] M. Fromm, O. Philipsen, W. Unger, and C. Winterowd, Quantum gate sets for lattice QCD in the strong-coupling limit: $n_f = 1$, *Eur. Phys. J. Quantum Technol.* **11**, 24 (2024).
- [31] E. J. Gustafson, Y. Ji, H. Lamm, E. M. Murairi, S. O. Perez, and S. Zhu, Primitive quantum gates for an SU(3) discrete subgroup: $\Sigma(36 \times 3)$, *Phys. Rev. D* **110**, 034515 (2024).
- [32] E. J. Gustafson, Y. Ji, H. Lamm, E. M. Murairi, S. O. Perez, and S. Zhu, Primitive quantum gates for an SU(3) discrete subgroup: $\Sigma(36 \times 3)$, *Phys. Rev. D* **110**, 034515 (2024).
- [33] A. N. Ciavarella and C. W. Bauer, Quantum simulation of SU(3) lattice Yang-Mills theory at leading order in large- N_c expansion, *Phys. Rev. Lett.* **133**, 111901 (2024).
- [34] S. V. Kadam, A. Naskar, I. Raychowdhury, and J. R. Stryker, Loop-string-hadron approach to SU(3) lattice Yang-Mills theory: Gauge invariant Hilbert space of a trivalent vertex, *arXiv:2407.19181*.
- [35] G. Bergner, M. Hanada, E. Rinaldi, and A. Schäfer, Toward QCD on quantum computer: Orbifold lattice approach, *J. High Energy Phys.* **05** (2024) 234.
- [36] A. T. Than *et al.*, The phase diagram of quantum chromodynamics in one dimension on a quantum computer, *arXiv:2501.00579*.
- [37] S. V. Kadam, A. Naskar, I. Raychowdhury, and J. R. Stryker, Loop-string-hadron approach to SU(3) lattice Yang-Mills theory: Gauge invariant Hilbert space of a trivalent vertex, *Phys. Rev. D* **111**, 074516 (2025).
- [38] Quantinuum, <https://www.quantinuum.com/> (2025).
- [39] E. Zohar and M. Burrello, Formulation of lattice gauge theories for quantum simulations, *Phys. Rev. D* **91**, 054506 (2015).
- [40] M. C. Bañuls, K. Cichy, J. I. Cirac, K. Jansen, and S. 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).
- [41] N. Klco, J. R. Stryker, and M. J. Savage, SU(2) non-Abelian gauge field theory in one dimension on digital quantum computers, *Phys. Rev. D* **101**, 074512 (2020).
- [42] A. H. Z. Kavaki and R. Lewis, From square plaquettes to triamond lattices for SU(2) gauge theory, *Commun. Phys.* **7**, 208 (2024).
- [43] A. Alex, M. Kalus, A. Huckleberry, and J. von Delft, A numerical algorithm for the explicit calculation of SU(N) and SL(N,C) Clebsch-Gordan coefficients, *J. Math. Phys. (N.Y.)* **52**, 023507 (2011).
- [44] A. Ciavarella, N. Klco, and M. J. Savage, Some conceptual aspects of operator design for quantum simulations of non-Abelian lattice gauge theories, *arXiv:2203.11988*.
- [45] C. Jacobi, Über ein leichtes verfahren die in der theorie der säcularstörungen vorkommenden gleichungen numerisch aufzulösen, *J. Reine Angew. Math.* **1846**, 51 (1846).
- [46] Z. Jia, W. Huie, L. Li, W. K. C. Sun, X. Hu, Aakash, H. Kogan, A. Karve, J. Y. Lee, and J. P. Covey, An architecture for two-qubit encoding in neutral ytterbium-171 atoms, *npj Quantum Inf.* **10**, 106 (2024).
- [47] M. Li, F.-Q. Guo, Z. Jin, L.-L. Yan, E.-J. Liang, and S.-L. Su, Multiple-qubit controlled unitary quantum gate for Rydberg atoms using shortcut to adiabaticity and optimized geometric quantum operations, *Phys. Rev. A* **103**, 062607 (2021).
- [48] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, England, 2010).
- [49] B. Zindorf and S. Bose, Efficient implementation of multi-controlled quantum gates, *arXiv:2404.02279*.
- [50] D. Maslov, Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization, *Phys. Rev. A* **93**, 022311 (2016).
- [51] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Phys. Rev. A* **52**, 3457 (1995).
- [52] Tables generator, https://www.tablesgenerator.com/latex_tables (accessed: 2025-02-11).
- [53] A. N. Ciavarella, C. W. Bauer, and J. C. Halimeh, Generic Hilbert space fragmentation in Kogut-Susskind lattice gauge theories, *arXiv:2502.03533*.
- [54] C. Kane, D. M. Grabowska, B. Nachman, and C. W. Bauer, Efficient quantum implementation of $2+1$ U(1) lattice gauge theories with Gauss law constraints, *arXiv:2211.10497*.
- [55] G. P. Lepage and P. B. Mackenzie, On the viability of lattice perturbation theory, *Phys. Rev. D* **48**, 2250 (1993).
- [56] M. Carena, H. Lamm, Y.-Y. Li, and W. Liu, Improved Hamiltonians for quantum simulations of gauge theories, *Phys. Rev. Lett.* **129**, 051601 (2022).
- [57] BlueQubit, <https://www.bluequbit.io/> (2025).
- [58] W. Gropp, T. Boerner, B. Bode, G. Bauer, and C. Stirm, Delta: Balancing GPU performance with advanced system interfaces, Technical Report, 2023, <https://hdl.handle.net/2142/117179>.

Correction: Missing lines in the quantum circuits represented in Eq. (23), Eq. (25), Fig. 5, and Fig. 6 have been restored.