Qudit gate decomposition for lattice gauge theories on superconducting radio-frequency cavities

Doğa Murat Kürkçüoğlu [1,2,*] Henry Lamm [1,2,†] and Andrea Maestri [3,‡]
[1] *Superconducting and Quantum Materials System Center (SQMS), Batavia, Illinois 60510, USA*
[2] *Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*
[3] *University of Amsterdam, 1090 GL Amsterdam, The Netherlands*

(Received 16 March 2025; accepted 15 July 2026; published 27 August 2026)

In this work we investigate the effect of decomposition basis on primitive qudit gates on superconducting radio-frequency cavity-based quantum computers with applications to lattice gauge theory. Three approaches are tested: SNAP and displacement gates, ECD and single-qubit rotations $R(\theta, \phi)$, and optimal pulse control. For all three decompositions, implementing the necessary sequence of rotations concurrently rather than sequentially can reduce the primitive gate run time. The number of blocks required for the faster ECD and $R_p(\theta)$ is found to scale as $O(d^2)$, while the slower SNAP and displacement set scales sublinearly with d on the permutation gates relevant for lattice gauge theory. For qudits with $d < 10$, the resulting gate times for the two decompositions are similar, though strongly dependent on experimental design choices. Optimal control can outperform both decompositions for small d by a factor of 2 to 12 at the cost of higher classical resources. Lastly, we find that SNAP and displacement are slightly more robust to a simplified noise model.

DOI: 10.1103/bd87-g26f

I. INTRODUCTION

While the advent of quantum computers offers the opportunity to investigate new questions in lattice gauge theories (LGTs) [1–4], at present the gate resources are prohibitively large [5,6]. Due to the large local Hilbert space of gauge theories, d-dimensional qudit-based computers can reduce the resources required [5,7–14]. The potential algorithmic advantage of qudit platforms arises from the increased effective connectivity as native single-qudit $SU(d)$ rotations replace otherwise require nonlocal multiqubit circuits [15–17]. In practice, this may allow for lower individual gate fidelities for the same algorithmic fidelity. Today many promising systems for qudits are being explored including trapped ions [18–20], transmons [21–26], Rydberg arrays [27,28], photonic circuits [29], ultra-cold atomic mixture [30], and superconducting radio frequency (SRF) cavities [31]. In this work we focus on 3D superconducting radio-frequency (SRF) cavities derived from accelerator physics which possess decoherence times on the order of milliseconds [31]. Similar to qubits the possible gate sets are numerous for both native gates [32–35] and fault-tolerant computation [36–46].

Quantum algorithms for lattice gauge theories generally require a set of fundamental group theoretic operations [47]. The identification of primitive subroutines divides the problem of formulating quantum algorithms for LGT into deriving said group-dependent primitives [5,48–53] and group-independent algorithmic design [14,50,54–58]. In addition to the group dependence, the primitive subroutines vary based on the digitization of the gauge degrees of freedom. Some digitize in the representation basis with a maximum representation encoded [11,13,59–74]. Alternatively, the q-deformed formulation obtains a finite-dimensional Hilbert space by changing the symmetry group to be a so-called quantum group [14,50]. Other formulations consider truncations within the bases of gauge-invariant states [75–84]. Further methods begin with different formulations or perform different approximations exist such as light-front quantization [85–87], conformal truncation [88], strong-coupling, and large-N_c expansions [89,90]. Another approach is to formulate an inherently finite-dimensional Hilbert space theory with continuous local gauge symmetry, which is in the same universality class as the original theory. Some methods for this include fuzzy gauge theories which use noncommutative geometry [91] and quantum link models which use rishons and an ancillary dimension [92–102].

In this work we consider gates that arise naturally when digitizing with the discrete subgroup approximation [12,52,58,75,103–116]. The discrete group approximation has advantages over the methods discussed above. It is a finite mapping of group elements to integers that preserves a group structure, therefore reducing the amount of fixed- or floating-point quantum arithmetic and simplifying the primitive gates. This method has its roots in early Euclidean LGT where the discrete group structure allowed for reduction of the classical resources [117–123] and has seen a resurgence in the era of quantum computation [106,107,109,124–126]. While the discrete subgroups are an approximation, they are related to the continuous groups broken by a Higgs mechanism [127–131] with work ongoing to understand this systematically [132].

At present, most work has emphasized qubit devices, but recent demonstrations of multiqudit gates, have increased interest in qudit-based digitization methods [7–9,11–14,53].

*Contact author: dogak@fnal.gov
†Contact author: hlamm@fnal.gov
‡Contact author: andreamaestrimaestri@gmail.com

Given the abundance of platforms and native gates, it would be valuable to explore the relative merits of different gate sets with respect to algorithmic implementation. Qudit-based SRF architectures are one such platform, and we explore prototype gates required for the simulation of LGT within the discrete subgroup approximation on them here. Focus is given to implementations on $d = 4, 6, 8$ states—respectively referred to as *ququart*, *quhexit*, and *quoctit*. We consider both state preparation and gates decomposed in three ways: SNAP and displacement (S + D) gates, the echoed conditional displacement (ECD) and single-qubit rotation gates $R(\theta, \phi)$, and optimal pulse control. Comparisons between the three are made in terms of number of blocks, timescales, and noise resilience.

This paper is organized as follows. A brief review of the field theory motivations and the general protocols we are interested in are found in Sec. II. In Sec. III we briefly summarize the theory behind coupling a cavity to a two-level system to produce a qudit. This is followed in Sec. IV with a discussion about how the interacting Hamiltonian is used to implement digital gates and the three decompositions. Section V is where we describe the methods for determining our gate decompositions, the results of which are found in Sec. VI including an initial study of robustness to noise. We conclude in Sec. VII with summary and future work.

II. LATTICE GAUGE THEORY

The gauge symmetries of particle physics constraint the possible interactions of theories, and thus only a finite set of gauge-group-dependent primitive gates are required for simulation [47,104]. These correspond to operations that can be performed on either one or two elements g, h of the group G stored in registers $|g\rangle, |h\rangle$. For the case of pure gauge theory (theories without the matter), the set of primitive gates are the following: the inverse gate $\mathfrak{U}_{-1}$ which sends a group element to its inverse:

$$\mathfrak{U}_{-1}|g\rangle = |g^{-1}\rangle; \tag{1}$$

the multiplication gate $\mathfrak{U}_{\times}$ that acts on two registers, changing the target to the left product with the control:

$$\mathfrak{U}_{\times}|g\rangle|h\rangle = |g\rangle|gh\rangle; \tag{2}$$

the trace gate $\mathfrak{U}_{\mathrm{Tr}}$ which rotates $|g\rangle$ by a phase dependent on the trace of g in a specified representation:

$$\mathfrak{U}_{\mathrm{Tr}}|g\rangle = e^{i\theta \mathrm{Re}\,\mathrm{Tr}(g)}|g\rangle; \tag{3}$$

and the group Fourier transform $\mathfrak{U}_{FT}$ which acts on a single register with some amplitudes $f(g)$:

$$\mathfrak{U}_F \sum_{g\in G} f(g)|g\rangle = \sum_{\rho\in\hat{G}} \hat{f}(\rho)_{ij}|\rho, i, j\rangle. \tag{4}$$

The second sum is over ρ, the irreducible representations of G; $\hat{f}$ denotes the G Fourier transform of f,

$$\hat{f}(\rho) = \sqrt{\frac{d_\rho}{|G|}} \sum_{g\in G} f(g)\rho(g), \tag{5}$$

where $|G|$ is the group size, d_ρ is the dimensionality of the representation ρ, and f is a function over G.

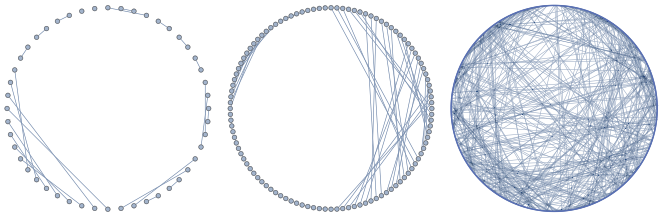

FIG. 1. Permutation graph necessary for implementing $\mathfrak{U}_{-1}$ for different discrete group: (left) the 48 element $\mathbb{BO}$, (center) the 108 element $\Sigma 108$, (right) the 1080 element $\Sigma(1080)$.

Because these operations act on at most two registers, they can be analyzed and optimized more efficiently than working at the algorithmic level, something dramatically demonstrated for $\mathfrak{U}_{FT}$ in a recent work [53]. Further, while $|G|$ (and therefore the optimal qudit dimension d) in the discrete subgroup approximation range from 24 to 1080, groups share structures, which allows us to explore smaller gates while still being informative about preferred gate sets for HEP.

By inspecting Eq. (1), we recognize that $\mathfrak{U}_{-1}$ corresponds to a pair-wise permutation gate $|g\rangle \leftrightarrow |g^{-1}\rangle$ (see Fig. 1 for three example groups). If the qudit $d = |G|$, this gate is formed by a tensor product of at most $|G|/2$ Pauli $X^{(g,h)}$ gates between states $|g\rangle, |h\rangle$.

An example of this for the 24-element $\mathbb{BT}$ mapped to a quicosotetrit ($d = 24$ qudit) is shown in the top of Fig. 2. As such, universal gate sets that can perform multiple $X^{(g,h)}$ concurrently would yield algorithmic improvement. If $d < |G|$, the gate structure generalizes from $X^{(g,h)}$ gates to requiring controlled permutations that couple the qudits. Another important structure exemplified in Fig. 1 is that the maximum distance between $|g\rangle$ and $|g^{-1}\rangle$ is $\sim\frac{1}{2}|G|$, another way to discriminate the efficiency of gate sets.

The only two-register gate, $\mathfrak{U}_{\times}$, corresponds to a controlled permutation gate where a different permutation gate is applied depending on the control qudit state. Writing

$$\mathfrak{U}_{\times} = \sum_{g\in G} |g\rangle\langle g| \otimes P_g \tag{6}$$

makes explicit that each branch P_g is itself a single-qudit permutation of the type studied throughout this work. The two-register gate can therefore be assembled from a CSUM-style controlled-shift primitive between the two qudits together with single-qudit permutations on the target register conditioned on the control. As such, $\mathfrak{U}_{\times}$ has similar implementation requirements to $\mathfrak{U}_{-1}$, and the single-qudit decompositions characterized here are the dominant ingredient. Multimode 3D SRF cavities sharing an ancilla transmon [133] provide a natural hardware substrate for the required two-mode primitive, e.g., through a beamsplitter or controlled-displacement interaction. Further, $\mathfrak{U}_{\mathrm{Tr}}$ is diagonal in state space. Therefore if $d = |G|$, the only elementary gate required is $R_Z^{(g,h)}(\theta)$ which generates a R_Z gate of phase θ between $|g\rangle$ and $|h\rangle$ (see the bottom of Fig. 2).

The final gate, $\mathfrak{U}_{\mathrm{F}}$, can be decomposed for $d = |G|$ into a tensor product of $SU(|\mathcal{C}|)$ rotations $\mathcal{U}_N^{(a,b,\ldots,n)}$ where the size of set $\{\mathcal{C}\}$ is the order of the conjugacy class $\mathcal{C}$ of the group [49]. For $d < |G|$, these must be mapped to controlled $SU(N)$

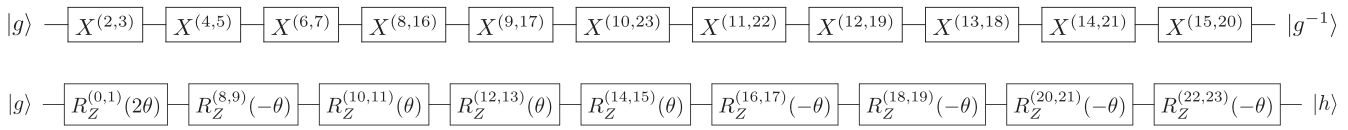

FIG. 2. Implementation of $\mathbb{BT}$ primitive gate assuming a quicosotetrit: (top) $\mathfrak{U}_{-1}$ and (bottom) $\mathfrak{U}_{\mathrm{Tr}}$. Note for the inverse gate the largest distance between a $|g\rangle$ and $|g^{-1}\rangle$ is 13.

gates. Although the spacing between states $a, b, \ldots, n$ could in principle be $\sim|G|$, we find that for the discrete subgroup approximation, $|\mathcal{C}| \lesssim \sqrt{|G|}$ with the maximum $|\mathcal{C}| = 18$. The smallest possible rotation is $SU(2)$, where the Euler angle decomposition can be built from an ZXZ rotation where the superscripts indicate levels that are rotated:

$$\mathcal{U}_2^{(a,b)}(\vec{\theta}) = R_Z^{(a,b)}(\theta_0) R_X^{(a,b)}(\theta_1) R_Z^{(a,b)}(\theta_2). \tag{7}$$

We can combine $\mathcal{U}_N^{(a,b\ldots,n)}$ for the conjugacy class $\mathcal{C}$ into an operator $V_{\mathcal{C}}$. For example, one conjugacy class with order 2 in $\mathbb{BT}$ yields [49]:

$$V_1 = \prod_{a=0}^{11} \mathcal{U}_2^{(2a,2a+1)}(\vec{\theta}). \tag{8}$$

$V_{\mathcal{C}}$ for larger n are constructed recursively with additional Givens rotations [134–137]. $SU(3)$ rotations requires

$$\mathcal{U}_3^{(a,b,c)}(\vec{\theta}) = \mathcal{U}_2^{(a,b)}(\vec{\theta}_0) R_X^{(b,c)}(\theta_1) \mathcal{U}_2^{(a,b)}(\vec{\theta}_2) R_Z^{(b,c)}(\theta_3), \tag{9}$$

while an $SU(6)$ needs from two $\mathcal{U}_3^{(a,b,c)}(\vec{\theta})$ and three $\mathcal{U}_2^{(a,b)}(\vec{\phi})$ (where we suppress the arguments below)

$$\mathcal{U}_6^{(a,b,c,d,e,f)} = \mathcal{U}_3^{(a,b,c)} \mathcal{U}_3^{(d,e,f)} \mathcal{U}_2^{(a,d)} \mathcal{U}_2^{(b,e)} \mathcal{U}_2^{(c,f)}. \tag{10}$$

With these, a $V_{\mathcal{C}}$ with order 6 for $\mathbb{BT}$ is

$$\begin{aligned} V_9 =& \mathcal{U}_6^{(0,9,16,1,8,17)} \mathcal{U}_6^{(2,14,20,3,15,21)} \\ & \times \mathcal{U}_6^{(4,11,23,5,10,22)} \mathcal{U}_6^{(6,12,19,7,13,18)}. \end{aligned} \tag{11}$$

Thus, $\mathfrak{U}_{\mathrm{F}}$ is a tensor product of a few $SU(N < 18)$ rotations between well-separated qudit states.

Along with primitive gates, the opportunity for optimizing state preparation in LGT should be considered. In situations where gauge-redundant encodings are necessary, one often prepares a gauge-dependent state and then applies a gauge-symmetrization operator. This gate acts as a projection operator [47,58,104,138]:

$$\begin{aligned} P|U_{ij}\ldots\rangle &= \frac{1}{|G|^N} \left(\int_G \mathrm{d}g_1 \int_G \mathrm{d}g_2 \ldots \right) |g_2 U_{ij} g_1^\dagger \ldots\rangle \\ &= \frac{1}{|G|^N} \int_{G^N} \mathrm{d}g\, \phi(g) |U\rangle. \end{aligned} \tag{12}$$

The structure of this operator is a sum over $\mathfrak{U}_\times$, and thus again a set of controlled permutation matrices. This constraint allows one to restrict the possible state that must be initially prepared, simplifying state preparation optimization. Further, algorithmic choices in state preparation can simplify the possible initial states [115,139,140].

To summarize, the general structures of a LGT primitive gate are tensor products of $X^{(a,b)}$ and $SU(N < |G|)$ rotations between states that are typically well separated. While it could prove experimentally hard to create qudits with $d = |G|$, for smaller d the single-register gates become similar to the controlled permutations required for $\mathfrak{U}_\times$, so one can start with considering how prototype gates on smaller qudits can be compiled.

III. 3D CAVITY QED

The hardware platform we will study here corresponds to an SRF cavity mode (qumode) coupled to a two-level system (qubit). With $\hbar = 1$, a simplistic model of this interaction from circuit QED is the *Jaynes-Cummings model* [141]:

$$H = \omega_c \hat{a}^\dagger \hat{a} + \tfrac{1}{2}\omega_0 \sigma_z + g(\hat{a}^\dagger \sigma_- + \hat{a}\sigma_+), \tag{13}$$

where ω_c is the frequency of the qumode with creation and annihilation operators $\hat{a}, \hat{a}^\dagger$, ω_0 is the frequency of the two-level system which is acted on by the Pauli operators σ_i, and g is a coupling between the qumode and qubit. The regime considered, particularly amenable to hardware, is the so-called dispersive regime, which occurs when $\Delta = |\omega_0 - \omega_c| \gg g$ [142]. In this regime one can approximately diagonalize H using a unitary transformation $D^\dagger H D$ with

$$D = e^{\lambda(\hat{a}^\dagger \sigma_- - \hat{a}\sigma_+)}, \tag{14}$$

where $\lambda = g/\Delta$. By doing so, and expanding to $O(\lambda)$, one obtains

$$H = \omega_c \hat{a}^\dagger \hat{a} + \tfrac{1}{2}\omega_q \sigma_z + \tfrac{1}{2}\chi \hat{a}^\dagger \hat{a} \sigma_z + O(\lambda^2), \tag{15}$$

where $\chi = 2g\lambda = 2g^2/\Delta$ is called the dispersive shift and $\omega_q = \omega_0 + \chi$. With this, the Fock states of the qumode ($|0\rangle, |1\rangle, |2\rangle, \ldots$) can be used as a qudit computational space, and the ancilla qubit provides an interaction to manipulate the qudit.

An easy way to control the system is to add additional fields acting on the qumode and the qubit. A possible choice is to add a single control on both:

$$\begin{aligned} H =& \omega_c \hat{a}^\dagger \hat{a} + \tfrac{1}{2}\omega_q \sigma_z + \chi \hat{a}^\dagger \hat{a} \sigma_z \\ & + \epsilon(t) \hat{a}^\dagger + \epsilon(t)^* \hat{a} + \Omega(t)\sigma_+ + \Omega(t)^* \sigma_-, \end{aligned} \tag{16}$$

which we define as $H = H_0 + H_1$ as a noninteracting $H_0 = \omega_c \hat{a}^\dagger \hat{a} + \frac{1}{2}\omega_q \sigma_z$, and the remaining terms the interacting Hamiltonian H_1.

With this choice of controls and enough time, it is possible to manipulate the system to reach any arbitrary state [143]. It is important to emphasise that for SRF cavities there is not an agreed-upon mode of operation. Therefore, the investigation of different gate sets [32–35] and optimal control [143–151] is key to fully exploit these systems. At a higher level of abstraction, performing efficient compilation of quantum circuits to qudits rather the qubits also requires new techniques [5,152–156].

IV. GATE IMPLEMENTATION

To obtain a theoretical target gate U_t, one manipulates the time-dependent couplings of Eq. (16). All approaches require this optimization at the Hamiltonian level, and in this section we summarize two ways to do so. The first approach relies on only direct optimization of $\epsilon(t)$, $\Omega(t)$, while the other decomposes U_t into a set of simpler gates that have been optimized at the Hamiltonian level. In either case, one considers a quantum state $|\psi(t)\rangle$ evolving according to the Schrödinger equation with a particular time-dependent Hamiltonian $H(t)$. This state can be approximated by dividing t into time steps of size δt:

$$|\psi(t)\rangle \approx \prod_{n=0}^{t/\delta t} e^{-iH(n\delta t)\delta t}|\psi(0)\rangle \equiv V|\psi(0)\rangle, \tag{17}$$

with increasing accuracy as $\delta t \to 0$.

Clearly, $H(t)$ allows transition between quantum states, and thus V correspond to a digital quantum circuit. The goal then is to find a V which implements U_t with high fidelity. It is important to emphasize that V as defined in Eq. (17) makes reference to a state $|\psi(0)\rangle$. It is the case that the state-dependent fidelities vary substantially [157–159], and therefore synthesizing a general unitary is harder than preparing a state from a given $|\psi(0)\rangle$. Proceeding, we will distinguish between gates for *state preparation* which are optimized only for acting on vacuum $|0\rangle$, and general gates which should have high fidelity on any $|\psi(0)\rangle$.

To study these effects, we need to define the infidelity, i.e., a cost function $\mathcal{I}(U_t, V)$ to minimize. Given that the goal is to implement a target gate U_t, a common approach consists of dividing the time into N time steps, and for each of these, the control pulses $\Omega(n\delta t)$, $\epsilon(n\delta t)$ are constant. A main drawback of this approach is the increasing number of time steps—and therefore optimization parameters—required when increasing d of the qudit [35,143]; this is because while the controls can generate the algebra, reaching higher states requires more and more commutators. Another drawback is the pulse discontinuities. By optimizing only the field amplitudes at each time step, we are neglecting the time required to ramp the fields on the physical hardware, which can introduce additional infidelity or must be implemented as expensive constraints on the optimization. Further, increasing N requires more matrix multiplications in the optimization algorithm. Looking at Eq. (17), for a d-dimensional qudit N matrix multiplications are required to compute the evolution of the system, leading to $O(Nd^3)$ operations. Lastly, when optimizing the cost function, unless the optimization method is gradient-free, it is necessary to compute either the gradient or in some cases the Hessian, which further increases the complexity. For N parameters, without recycling computations, the cost function needs to be evaluated $N+1$ times, which leads to an overall complexity of $O(N^2d^3)$.

An alternative approach avoids piecewise constant terms and instead optimize gates in a different basis. Looking at the expression of Eq. (16), a way to generate smooth pulses, making it easier to implement, is to expand them in some polynomial basis [160,161], where the optimal choice can be problem-dependent. Moreover, the number of basis states directly controls how steep the pulse ramping can be. Here we will expand in Chebyshev polynomials which provide an exact nth-order interpolation function for any nth-order polynomial. Approximating instead using the Fourier series requires more terms to reach a comparable accuracy [162]. Rewriting Eq. (16) in the Chebyshev basis of $T_k(t)$ yields

$$H = H_0 + \chi\hat{a}^\dagger\hat{a}\sigma_z + \sum_{k=0}^{N}[(c_k\hat{a}^\dagger + q_k\sigma_+)T_k(t) + \text{H.c.}], \tag{18}$$

where $\{c_k\}$ and $\{q_k\}$ are the coefficients of the qumode and the qubit, respectively. Using this expression for H, it is possible to compute the V. Its magnitude depends on the dimension of the cavity space that is optimized, the duration of the pulses, and the resolution of the hardware on which the pulses are optimized.

A. Interaction picture

Finding an optimal V from Eq. (18) can be made easier by transforming into the interaction frame of reference [163]. Since H_0 is time-independent, the time evolution can be solved analytically:

$$U_0(t) = e^{-iH_0t}.$$

With this expression, it is possible to define a state vector in this new frame of reference:

$$|\psi_I(t)\rangle = U_0(t)|\psi_S(t)\rangle, \tag{19}$$

where $|\psi_S(t)\rangle$ represents the state governed by Schrödinger equation. Using this, it is possible to write

$$i\frac{\partial}{\partial t}|\psi_S(t)\rangle = H|\psi_S(t)\rangle = e^{-iH_0t}\left(H_0 + i\frac{\partial}{\partial t}\right)|\psi_I(t)\rangle$$
$$\Rightarrow i\frac{\partial}{\partial t}|\psi_I(t)\rangle = e^{iH_0t}H_1e^{-iH_0t}|\psi_I(t)\rangle. \tag{20}$$

This can be used to simplified optimization of V by solving Eq. (20) and then transforming back to the initial frame by multiplying by the inverse transformation defined by Eq. (19). The remaining step consists in writing explicitly the term that dictates the evolution of the system:

$$H_{1_I} = e^{iH_0t}H_1e^{-iH_0t}$$

with the resulting operator being

$$H_{1_I} = \chi\hat{a}^\dagger\hat{a}\sigma_z + \sum_{k=0}^{N}[c_kC_k(t)\hat{a}^\dagger + q_kQ_k(t)\sigma_+ + \text{H.c.}], \tag{21}$$

where

$$C_k(t) \equiv T_k(t)e^{i\omega_c t},$$
$$Q_k(t) \equiv T_k(t)e^{i\omega_q t}.$$

While optimizing the pulses in the interaction picture may be easier, the pulses must be converted back to the laboratory frame by multiplying them with the exponential factors that are missing in Eq. (21). This may reduce accuracy, depending on the hardware precision.

To inform the gate synthesis that is developed in this work, it is important to give some realistic near-term properties for 3D cavities. We take for a fiducial decoherence time of

the cavity $T_1 = 10\,\text{ms}$ [164]. With this, we take $\omega_c = 5\,\text{GHz}$, $\omega_q = 3\,\text{GHz}$, and $\chi = 0.3\,\text{MHz}$. Regarding the control pulses, we consider a max time resolution of 1 ns (corresponding to 1 giga-sample/s) for the state preparation, and a longer 10 ns time resolution for the pulses used for gate optimization. An additional constraint is for the drive amplitudes to change by no more than a 1 MHz/ns (to maintain a smooth drive profile) which is consistent with near-term SRF cavities.

B. Gate decomposition in universal sets

One could imagine engineering an optimized pulse for each LGT primitive. This has been show to speed up state preparation [165], although issues with computational expense and error correction may complicate its more general use. Instead, one might decompose each primitive into a limited set of universal gates whose implementation is known at the Hamiltonian level with high fidelity. The price of this is increased runtime of the primitive. While qudits can decrease the circuit depth for LGT [49], improvements are gate-set dependent. Thus, we consider two universal sets: SNAP and displacement and ECD.

1. *SNAP and displacement*

The qudit gate set of SNAP and displacement gates ($S + D$) has been widely discussed in the literature [32–34,152], and we briefly review the key points here. Crucially, the universality of $S + D$ is known [143]. The displacement gate is produced by driving the cavity on resonance with a calibrated pulse of a specific amplitude and length. It can be expressed as

$$D(\alpha) = e^{(\alpha\hat{a}^\dagger - \alpha^*\hat{a})} \tag{22}$$

and generates a coherent state $|\alpha\rangle$ from the vacuum displaced in phase space by α, i.e., $D(\alpha)|0\rangle = |\alpha\rangle$. An important property of $D(\alpha)$ is that its Hermitian conjugate $D^\dagger(\alpha) = D(-\alpha)$. The parameter α of Eq. (22) can be considered real, without loss of generality. Mathematically $D(\alpha)$ represents the Heisenberg group, which may prove useful for compilation. For reasonable χ on GHz SRF cavities with $T_1, T_2 \sim O(1)$ ms, the gate times range from 0.01–0.2 μs depending on α [166].

The SNAP gate applies arbitrary phases to Fock states

$$S(\vec{\theta}) = \sum_{k=0}^{d-1} e^{i\theta_k} |k\rangle\langle k|. \tag{23}$$

This is produced by driving the qubit with π pulses with relative phases, which cause a phase kickback onto the qudit. Importantly, while the qubit is used to implement this gate, upon completion the qubit is unentangled from the qudit. To ensure high fidelity, the driving strength for a SNAP must be taken $\max(Q_k(t)) \ll \chi$, with a good heuristic found to be a pulse length of $\pi/(10\chi)$. While larger χ allows for faster SNAP gates, the need to avoid strong self-Kerr interactions limits its size. Reasonable values for χ for GHz cavities with $T_1, T_2 \sim 1$ ms lead to SNAP gate times between 1 and 50 μs [31,35,164].

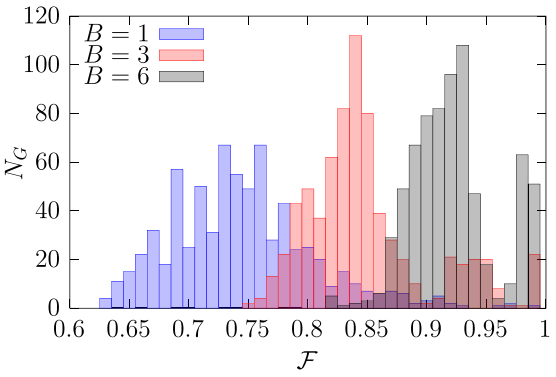

FIG. 3. Number of $d = 6$ permutation gates N_G with given fidelity as a function of B for $S + D$.

With these, any single-qudit gate can be approximately decomposed by [33]

$$U_t = \left(\prod_{k=0}^{B} D(\alpha_k) S(\vec{\theta}_k)\right) D(\alpha_{B+1}), \tag{24}$$

where B is the number of blocks used. For a d-state qudit, an B block decomposition has $(d+1)B + 1$ free parameters. Since an arbitrary qudit state can be obtained by an $SU(d)$ rotation with $d^2 - 1$ parameters, a naive estimate for the decomposition of any single qudit state is at most $B = \frac{d^2-2}{d+1} \sim O(d)$ [167]. This represented an algorithmic improvement over the $O(d^2)$ scaling of a $\log_2(d)$ qubit device with access to only one- and two-qubit gates [17]. The authors of [33] empirically found that a broad set of qudit gates (with $d \leqslant 10$) can be decomposed according to Eq. (24) with only three to four blocks and that the fidelity $\mathcal{F} \approx 1 - e^{-6.49(B/d)^{1.91}}$.

Following the LGT insight discussed in Sec. II that permutation matrices represent a sizable fraction of the primitives necessary, we investigated whether these gates represent a special subset of all qudit gates in terms of their fidelity under $S + D$ decomposition. For all $d!$ permutation gates with $d = [3, 6]$ we performed 10 decompositions with $B = [1, 6]$, where the parameters for the decompositions were randomly initialized. From these data, the best fit to the functional form advocated in [33] was $\mathcal{F}_{\text{perm}} \approx 1 - e^{-3.8(5)(B/d)^{0.5(1)}}$. This dependence shows that for values of $B/d \lesssim 1$, permutation gates can be decomposed with higher fidelity than general gates. In Fig. 3 we present the distribution of permutation gate fidelities for $d = 6$ as a function of B. From the figure, one observes a clear bimodal distribution in $\mathcal{F}$ for $B = 3, 6$.Beyond finding that permutation gates are easier to decompose than general $SU(d)$ matrices, one could ask if specific permutation decompositions are better, especially in light of the bimodal distribution. Given the freedom in encoding of LGT, such algorithmic input could be used similar to Gray codes on qubits to reduce noise or gates [168–171]. As such, we investigated the dependence of $\mathcal{F}$ for $d = 6$ permutation gates on two distance metrics used for distinguishing between strings when we write the permutations as tuples, for example,

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \to (1, 0, 2), \tag{25}$$

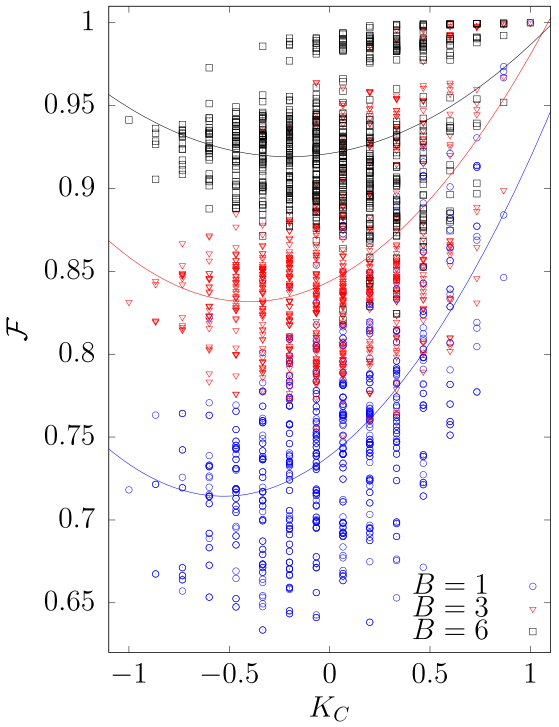

FIG. 4. Fidelity of $d = 6$ permutation gates vs K_C and B. Quadratic best fit lines are shown to guide the eye.

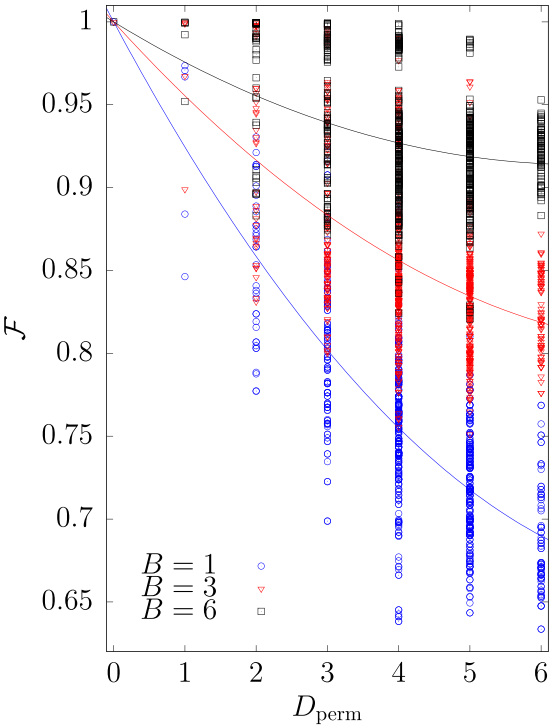

FIG. 5. Fidelity of $d = 6$ permutation gates vs D_{perm} for different block sizes B. Quadratic best fit lines with $\mathcal{F}(0) = 1$ are shown to guide the eye.

where the ith entry j of the tuple corresponds to permuting the state $|j\rangle \to |i\rangle$. The first is the *Kendall rank correlation coefficient* K_C which is used as a statistic measure between two sets of data $x = \{x_0 \dots x_n\}$ and $y = \{y_0 \dots y_n\}$ of their relative monotonicity [172]. It is defined in terms of the number of pairs (x_i, y_i) and (x_j, y_j) that are concordant, i.e., for $i < j$ either $x_i < x_j$ and $y_i < y_j$ or $x_i > x_j$ and $y_i > y_j$. If we call this number n_c, and the remaining discordant pairs n_d, then K_C is

$$K_C = \frac{n_c - n_d}{n_c + n_d}, \tag{26}$$

where K_C can range from 1 when the two sets are identical to -1 when the two sets are exactly reverse. We take $x = (0, 1, \dots, d) = \mathbb{1}_d$, and thus K_C is measured from it. The results are found in Fig. 4. From these it is clear that K_C does capture some of the dependence of $\mathcal{F}$ on permutation, with higher K_C typically having larger $\mathcal{F}$. In particular we see that the bimodal distribution of Fig. 3 is correlated with positive K_C. Investigating further, one observes that $\mathcal{F}(K_C)$ is nonmonotonic, reaching a minimum at $K_C = 0$ with a large variance, before rising slightly toward $K_C \to -1$.

Another metric for comparing sets, with its origin in spell checking, is the *Damerau-Levenshtein distance* [173,174] D_{perm} which here counts only the number of transposition of two adjacent characters required to change one word into another. Figure 5 presents $\mathcal{F}$ as a function of D_{perm}. We do observe decreasing $\mathcal{F}$ with D_{perm}. From the best fit lines in Fig. 5, we see that $1 - \mathcal{F} \approx [10^{-3}, 10^{-2}]D_{\text{perm}}$ with some weak quadratic term. This allows us to compare two strategies for implementing a permutation: (1) decomposing the full target permutation directly into a single $S + D$ circuit with block count B, versus (2) implementing it as a sequence of $N = D_{\text{perm}}$ adjacent single transpositions $X^{(a,a+1)}$, each decomposed separately into $S + D$ at smaller B with per-gate fidelity $\mathcal{F}_{\text{ind}}$. Equating the worst-case direct fidelity to the product fidelity of the sequenced version gives, in the best case for the single gates,

$$1 - 10^{-2} D_{\text{perm}} = \mathcal{F}_{\text{ind}}^{D_{\text{perm}}} \Rightarrow \mathcal{F}_{\text{ind}} > 0.990; \tag{27}$$

i.e.. each individual transposition would have to be decomposed to fidelity >0.990 just to match the direct decomposition. Since this is a stringent requirement and direct $S + D$ decomposition already reaches it with modest $B < D_{\text{perm}}$, decomposing the full permutation directly is always more efficient in this regime.

Taken together, we conclude that permutation gates represent an important subset of gates under $S + D$ decomposition with higher than average fidelity. Further, specific classes of permutation gates appear to have still better scaling and thus could be useful in optimizing digitization of LGT.

2. ECD decomposition

One limitation of $S + D$ is that operations require time comparable to $2\pi/\chi$. This is a limiting factor in the weak-dispersive regime because decoherence will reduce fidelity. To overcome this challenge, another universal set has been proposed which combines an echoed conditional displacement (ECD) gate with single-qubit rotations [35]:

$$U_t = \left(\prod_{k=0}^{B} R(\theta_k, \phi_k) \text{ECD}(\alpha_k) \right) R(\theta_{B+1}, \phi_{B+1}), \tag{28}$$

where the ECD gate corresponds to applying D gates conditioned on the state of the qubit,

$$\text{ECD}(\alpha) = D(\alpha/2)|e\rangle\langle g| + D(-\alpha/2)|g\rangle\langle e|.$$

To enable universality, this is coupled to unselective qubit rotation:

$$R(\theta,\phi) = \exp\left(-\frac{i\theta}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi)\right).$$

In constrast to S + D, only the ECD gate acts directly on the qudit. In this scheme $R(\theta,\phi)$ acts on the qubit changing the relative displacement. As such, the qubit and qudit may still be entangled at the end of a block. This leads to greater sensitivity of the qudit to noise from the qubit. Modern transmon qubits can perform $R(\theta,\phi)$ in $O(0.01\ \mu\text{s})$. ECD(α) can be implemented in $T \propto \frac{\alpha}{\chi\alpha_0}$ where α_0 is the magnitude of the intermediate oscillator displacement. For GHz cavities and reasonable values of $\alpha_0 \sim 30$ this would correspond to 0.1–0.5 μs [31,35].

Counting parameters, B blocks have $3B+2$. Since the qubit may not be decoupled, the general state is given by an $SU(2d)$ rotation with $4d^2-1$ parameters. This naively estimates that $\frac{4}{3}d^2 - 1$ blocks are needed–$O(d^2)$ scaling, which is parametrically worse that $S+D$. As will be seen below, this larger B can be compensated by the faster gate times.

One should also compare to state-of-the-art qubit implementations, which also scale asymptotically as $O(d^2)$. An $SU(2^n)$ unitary where $d = 2^n$ requires [17] a number of sequential cnot gates,

$$N_{\text{CNOT}} = \tfrac{21}{16}d^2 - d - \tfrac{3}{4}d\log_2(d), \tag{29}$$

which suggests a comparable circuit depth to ECD. Further, the theoretical lower bound is $N_{\text{CNOT}} = \frac{1}{4}[d^2 - 1 - 3\log_2(d)]$ [175] while the lower bound on ECD is unknown. Importantly, modern transmon-based platforms can perform cnot gates in ∼0.5 μs [176], similar to the higher estimates for near-term ECD gates. Taken together, these results suggest that the algorithmic advantage of qudit-based platforms with ECD gate sets over qubit devices depends upon acheiving gate times below 0.5 μs and clearer decomposition estimates.

V. METHODS

For optimization, we have used `LBFGS` as our optimization algorithm [177]. For preparing $|\Psi\rangle$, one solves the pulse optimization of Eq. (17) using the parameters in Eq. (18) with a fixed $|\psi(0)\rangle$. For general gate preparation, the more complicated optimization procedure motivates using the interaction picture of Eq. (21). Further, given the larger number of optimization parameters, we found efficiency gains from first optimizing pulses with a small number of coefficients, and then growing the bases while using as initial conditions for the next iterations the previous results.[1] In order to compare the optimized gates with the target, it is necessary to define a cost function. In the literature, there is a number considered,

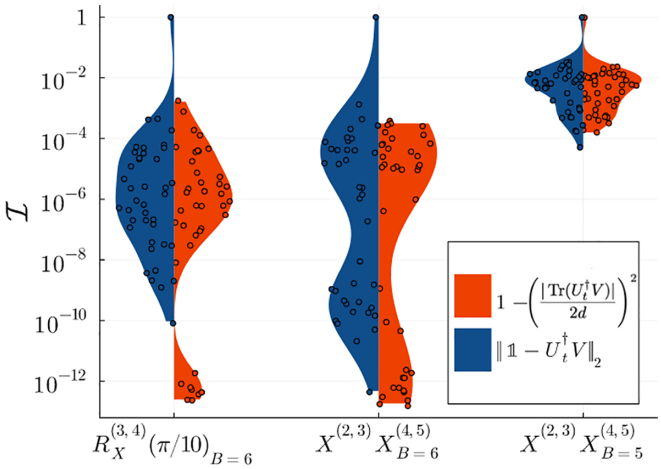

FIG. 6. Distribution of $\mathcal{I}$ of two cost functions for three gates on a quhexit using $S+D$ decomposition, after 30 optimizations and using different numbers of blocks B.

typically related to p norms [32,33,35,49,143,145,147,149]:

$$||A||_p = \sup_{x\neq 0}\frac{||Ax||_p}{||x||_p}, \tag{30}$$

which use the p vector norm,

$$||x||_p = \left(\sum_i |x_i|^p\right)^{1/p}. \tag{31}$$

If the goal is preparing $|\Psi\rangle$, the optimization problem can be restricted to only a subspace of the full system. Here we follow prior work and use as a cost function the infidelity:

$$\mathcal{I} = 1 - |\langle\Psi|V|\psi(0)\rangle|^2. \tag{32}$$

Typically, $|\psi(0)\rangle = |0\rangle$, but we will investigate a limited set of other initial states as a way to understand robustness of state preparation to noise.

In case where we wish to optimize V to approximate a gate U_t, one must account for working in the interacting picture by multiplying the target by the inverse of the transformation $U_0(T)$ defined in Eq. (19). With this, the most common cost function to minimize is

$$\mathcal{I} = 1 - \left(\frac{|\text{Tr}(U_t^\dagger U_0(T)V)|}{2d}\right)^2, \tag{33}$$

where $2d$ is the dimension of the combined cavity and qubit space. The idea behind this measure is that if $V \approx U_t$, using the properties of the unitary matrices we have $\mathcal{I} \longrightarrow 0$. Other options to measure the distance between two matrices could be

$$\mathcal{I} = ||\mathbb{1} - U_t^\dagger U_0(T)V||_2. \tag{34}$$

which still relies on unitarity, and the norm of the difference between the matrices:

$$\mathcal{I} = ||U_t - U_0(T)V||_2. \tag{35}$$

VI. NUMERICAL RESULTS

In this section we numerically optimize decompositions of a set of prototypical gates for HEP. As a first step, we

[1]The entire parallelized codebase is written in `Julia` and can be found at [178].

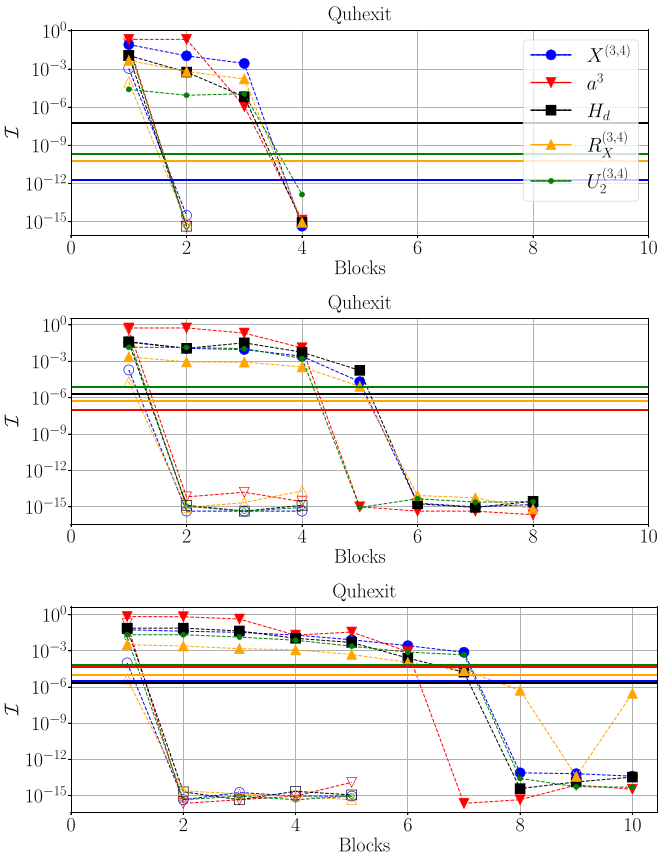

FIG. 7. Infidelity $\mathcal{I}$ vs block count B for $S+D$ (open points) and ECD (closed points) decompositions of the state preparations defined in the text, for the ququart ($d=4$, top panel), the quhexit ($d=6$, middle panel), and the quoctit ($d=8$, bottom panel). The horizontal lines correspond to the reference values of the pulse-optimization infidelities at the times listed in Table I.

compared the three candidate cost functions defined in Eqs. (33), (35), and (34) in order to decide which one to use in the subsequent analysis. The test cases were $S+D$ decompositions on a quhexit of two gates: the double-X gate $X^{(2,3)}X^{(4,5)}$ (at two different block sizes) and the $R_X^{(2,3)}(\pi/10)$ gate. The optimization based on the matrix-norm cost function of Eq. (35) showed poor, slow convergence, so we restrict the in-depth analysis to the remaining two cost functions, whose resulting infidelity distributions are presented in Fig. 6. In all cases the displayed infidelity is the trace-distance metric of

TABLE I. $\mathcal{I}$ from pulse optimization of state preparation on different qudits.

d	T [μs]	X	R_X	U_2	$\|0\rangle \to \|3\rangle$	$\|0\rangle \to \|H\rangle$
4	0.1	2×10^{-12}	6×10^{-11}	2×10^{-10}	2×10^{-10}	6×10^{-8}
6	0.2	2×10^{-6}	5×10^{-7}	8×10^{-6}	1×10^{-7}	2×10^{-6}
8	0.5	3×10^{-6}	9×10^{-6}	7×10^{-5}	5×10^{-5}	2×10^{-6}

Eq. (33), evaluated on the gates returned by optimizing each of the two cost functions.

One observes from these results that varying cost function leads to different decompositions, with distinct distributions related to how the matrix elements are weighted. A more sophisticated two-sided t-test for unequal variances confirms this. Putting everything together, the trace distance defined in Eq. (33) generally produced lower $\mathcal{I}$. Therefore this metric will be used in what follows.

We now study preparation of the following states $|\Psi\rangle$ on all three qudits and using all three decompositions:

(1) Fock state 3: $a^3|0\rangle = |3\rangle$, which is a representative of an electric basis state

(2) d-Hadamard state: $H_d|0\rangle = \frac{1}{\sqrt{d}}\sum_i |i\rangle$, which can be taken as a prototype of the weak-coupling gauge-invariant vacuum state in the magnetic basis

(3) $X^{(3,4)}$, $R_X^{(3,4)}(\pi/10)$, $U_2^{(3,4)}(\frac{\pi}{5}, \frac{2\pi}{15}, \frac{\pi}{10})$ applied to randomly generated state: $|\psi_{\text{ran}}\rangle \to U_{\text{target}}|\psi_{\text{ran}}\rangle$, correspond to prototypes of general HEP states that are prepared on a potentially noisy device.

For each decomposition, $\mathcal{I}$ as a function of B are presented in Fig. 7. From these $|\Psi\rangle$, the $S+D$ decompositions appear independent of qudit size d, always requiring a block size $B_{S+D}=2$ to reach $\mathcal{I} < 10^{-15}$. In contrast reaching $\mathcal{I} < 10^{-15}$ requires $B_{\text{ECD}} \approx d$. Thus these $|\Psi\rangle$ confirms the expectations of $B_{\text{ECD}} > B_{S+D}$ and the scaling from naive parameter counting. Further, while ECD requires more blocks, the potential 2–50× faster time per block can mitigate this so the ultimate preference depends on experimental hardware. These results can be compared to pulse optimization (PO) at total pulse durations chosen to be of order a single ECD block, but scaling with the qudit dimension: $T_{\text{PO}}^{d=4} = 0.1$ μs, $T_{\text{PO}}^{d=6} = 0.2$ μs, and $T_{\text{PO}}^{d=8} = 0.5$ μs. Example pulses for the quhexit are shown in Fig. 8. While these pulses do not start and end at zero power, we found that imposing smooth ramping had only a small effect on $\mathcal{I}$ but greatly increased convergence time. From these results we observe that the PO infidelity $\mathcal{I}_{\text{PO}}$ has a clear

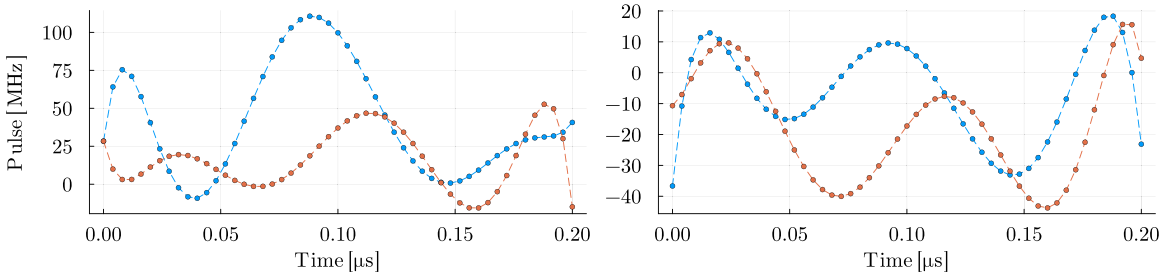

FIG. 8. Pulse optimization for the state prep of $H_6|0\rangle = \frac{1}{\sqrt{6}}\sum_i |i\rangle$ for a quhexit.

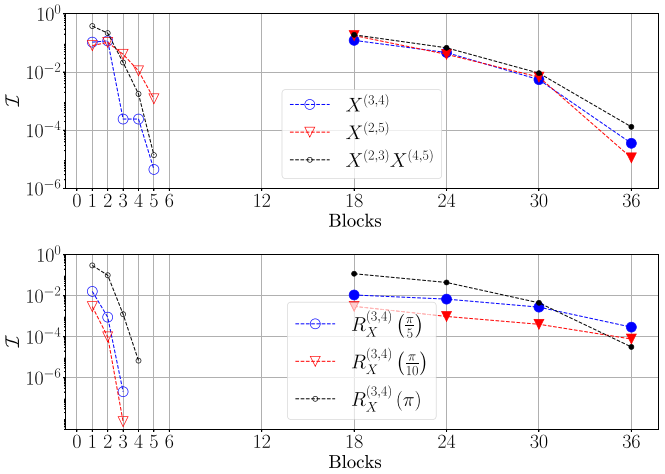

FIG. 9. $\mathcal{I}$ vs blocks for S + D and ECD for quhexit gates.

dependence on d, ranging from $\mathcal{I}_{\mathrm{PO}} \lesssim 10^{-8}$ for the ququart up to $\mathcal{I}_{\mathrm{PO}} \sim 10^{-5}$ for the quoctit (see Table I).

Moving to gate decomposition, we consider the following gates which are prototypical of HEP primitives:

(1) $X^{(3,4)}$, $X^{(2,5)}$, and $X^{(2,3)}X^{(4,5)}$ which can be related to $\mathfrak{U}_{-1}$ and $\mathfrak{U}_{\times}$

(2) $R_X^{(3,4)}(\theta)$ with $\theta = \pi, \pi/5, \pi/10$ which model $\mathfrak{U}_{\mathrm{Tr}}$ and $\mathfrak{U}_{\mathrm{F}}$.

For these gates, we consider their decomposition on quhexits with $S+D$ and ECD (see Fig. 9). In contrast to state preparation, $\mathcal{I}$ depends strongly upon B. We observe little difference in scaling for $X^{(3,4)}$ and $X^{(2,3)}X^{(4,5)}$, but $X^{(2,5)}$—with its larger distance between states—appears to converge more slowly for $S+D$. Notice how $\mathcal{I}$ is not strongly affected by applying multiple X gates, confirming the $S+D$ observation and suggesting it also holds for ECD decompositions. For the $R_X^{(3,4)}(\theta)$ gates, $\mathcal{I}$ increases with θ for both decompositions. Further, at fixed B and for the small θ investigated here, R_X has lower $\mathcal{I}$ than X gates. In all cases $B_{S+D} \ll B_{\mathrm{ECD}}$ with the relative factor being $\sim 6 = d$ which comports with the scaling from parameter counting.

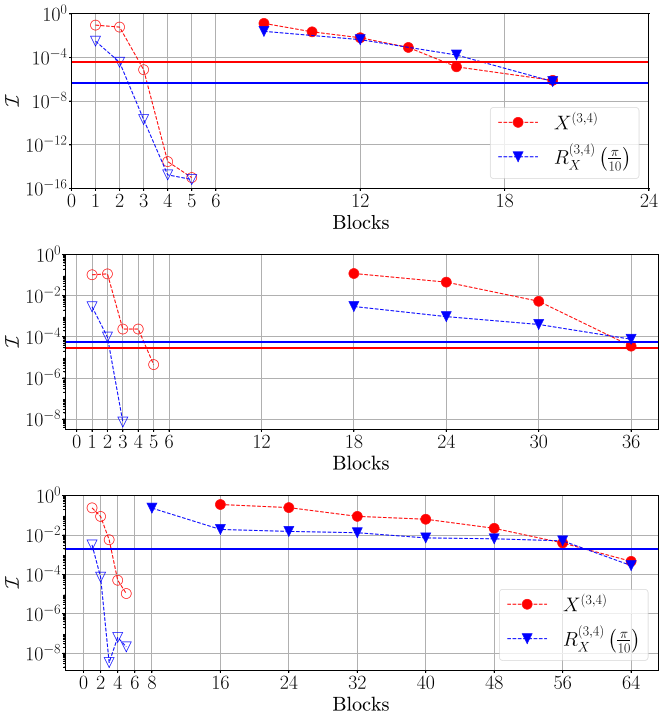

FIG. 11. $\mathcal{I}$ vs blocks for $S+D$ and ECD decomposition of an $X^{(3,4)}$ and $R_x^{(3,4)}(\pi/10)$ gate, for ququart, quhexit, and quoctit from top to bottom. The horizontal lines correspond to the reference values of the PO infidelities.

The state-dependency analysis above considered sequences of at most two transpositions on a quhexit. To verify that the advantage of decomposing combined permutations as a single $S+D$ circuit extends to longer sequences and to higher qudit dimension, we performed a systematic comparison at $d = 8$ for the four-disjoint-transposition operator $X^{(0,1)}X^{(2,3)}X^{(4,5)}X^{(6,7)}$. Three decomposition strategies are optimized side by side: (1) four single transpositions decomposed individually and concatenated, (2) two interleaved disjoint pairs $X^{(0,1)}X^{(4,5)}$ and $X^{(2,3)}X^{(6,7)}$ each decomposed separately, and (3) the entire four-transposition operator

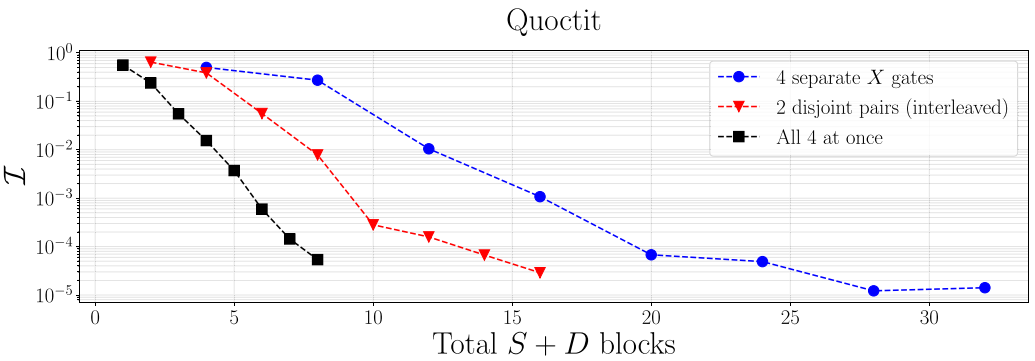

FIG. 10. Cumulative infidelity $\mathcal{I}$ vs the total number of $S+D$ blocks used to assemble the four-disjoint-transposition operator $X^{(0,1)}X^{(2,3)}X^{(4,5)}X^{(6,7)}$ on a quoctit ($d = 8$). Three decomposition strategies are compared: four single transpositions decomposed individually (blue circles, total $= 4B$), two interleaved disjoint pairs decomposed separately (red triangles, total $= 2B$), and the full four-transposition operator decomposed as a single $S+D$ circuit (black squares, total $= B$). Each point is the best of 20 random LBFGS initializations, and cumulative infidelity assumes multiplicative composition $1 - \prod_i (1 - \mathcal{I}_i)$ across the suboperators.

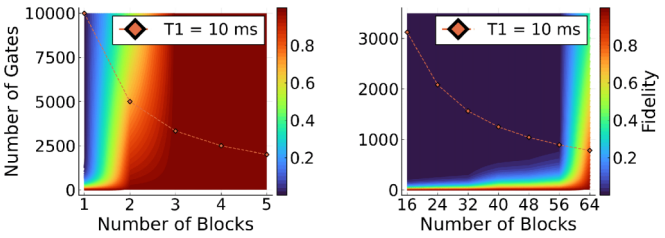

FIG. 12. $R_X^{(3,4)}(\pi/10)$ gate decomposition on quoctit using (left) $S+D$ and (right) ECD. The T_1 is plotted as a reference for the maximum number of gates given a certain amount of blocks, assuming fixed gate times of $T_{\text{SNAP}} = 1\,\mu\text{s}$ and $T_{\text{ECD}} = 200\,\text{ns}$.

decomposed as a single $S+D$ circuit. The cumulative infidelity over the assembled gate is plotted against the total number of $S+D$ blocks used in Fig. 10. Across the entire range of cumulative infidelity, decomposing all four transpositions at once requires fewer total blocks by a factor of 3 to 4 relative to running four individually decomposed transpositions, and by roughly a factor of 2 relative to the two-pair split. At a target $\mathcal{I} = 10^{-4}$, the all-at-once decomposition reaches the target with ∼7 blocks, the two-pair split with ∼14, and the four-single split with ∼24. The combined-decomposition advantage therefore persists at $d = 8$ and at sequence length 4. Structurally, this is what one would expect: a fixed-structure k-disjoint-transposition operator acts nontrivially on a $2k$-dimensional subspace independent of d, while the $S+D$ parameter count grows as $(d+1)B$, so the same combined-decomposition advantage is expected to extend to higher qudit dimensions as long as $2k \lesssim d$.

After looking at the dependence of $\mathcal{I}$ on the two decomposition methods, we perform an analysis of two gates $X^{(3,4)}$ and $R_X^{(3,4)}(\pi/10)$ for $d = 4, 6, 8$ qudits, which is presented in Fig. 11. If we take as a fiducial $\mathcal{I} = 10^{-4}$, we find that nearly constant $B_{S+D} \sim 2-3$ for all d, while $B_{\text{ECD}} = 16, 36, 64$ for $d = 4, 6, 8$ respectively. This again suggests B_{S+D} scales sublinearly with d. In contrast, ECD scales consistently with $B_{\text{ECD}} \sim d^2$.

Taking these results, one can estimate the break-even dimension d_b when the decompositions require the same amount of time by setting

$$B_{\text{ECD}} T_{\text{ECD}} = B_{S+D} T_{S+D}. \tag{36}$$

TABLE II. $\mathcal{I}_{\text{PO}}$ of two gates obtained with pulse optimization for qudits. For all cases, the time interval was divided into 50 time steps and the best optimization of 10 trials was taken. The final truncation order of the Chebysev series is denoted by o_C.

d	T [μs]	o_C	X	R_x
4	0.5	18	4×10^{-5}	5×10^{-7}
6	0.5	30	3×10^{-5}	6×10^{-5}
8	2.0	50	2×10^{-3}	2×10^{-3}

With the empirically observed constant B_{S+D} and relative factor of ∼1 between the decompositions, one finds

$$\frac{T_{S+D}}{T_{\text{ECD}}} = d_b^2 \Rightarrow d_b = \sqrt{\frac{[1, 50]\,\mu\text{s}}{[0.1, 0.5]\,\mu\text{s}}} = [2, 22]; \tag{37}$$

if instead, one takes the scaling suggested from parameter counting $B_{S+D} = d$, the range is $d_b = [2, 500]$.

These heuristics suggest while neither gate decomposition is decisively better, hardware-obtainable T_{S+D} and T_{ECD} would resolve it.

This logic can also be used to estimate the number of HEP primitive gates could be simulated within $T_1 = 10$ ms of the qudit with fixed $\mathcal{I}$ versus B. In Fig. 12 is a fiducial result for the $R_X^{(3,4)}(\pi/10)$ gate decomposition on a quoctit with fixed gate times of $T_{S+D} = 1\,\mu\text{s}$ and $T_{\text{ECD}} = 0.2\,\mu\text{s}$. Together these suggest such devices could reasonably achieve circuit depth of thousands of gates with total $\mathcal{I} < 1\%$. Given the resources estimates for $SU(2)$ in [49], this would be sufficient for quantum utility in toy models.

Finally, we consider optimal control. The total pulse duration was set to $T_{\text{PO}} = 0.5\,\mu\text{s}$ for the ququart and the quhexit, in order to be competitive with the single-block gate-set decompositions, but had to be extended to $T_{\text{PO}} = 2\,\mu\text{s}$ for the quoctit in order to reach $\mathcal{I} < 10^{-2}$. An example of the resulting pulses is shown in Fig. 13. The lowest infidelity from 10 trials for all three qudit dimensions are shown in Table II. An additional optimization that proved useful for the quoctit was to start with a smaller basis set, truncated at the 32th order. Then every 500 iterations, new terms were added to the decomposition. The new terms were added in three different batches consisting of 8, 5, and 5.

For the case of the quhexit, the infidelity from pulse optimization $\mathcal{I}_{\text{PO}} \lesssim 10^{-5}$ was close to that obtained for a five-block $S+D$ and a 36-block ECD decomposition, allowing a

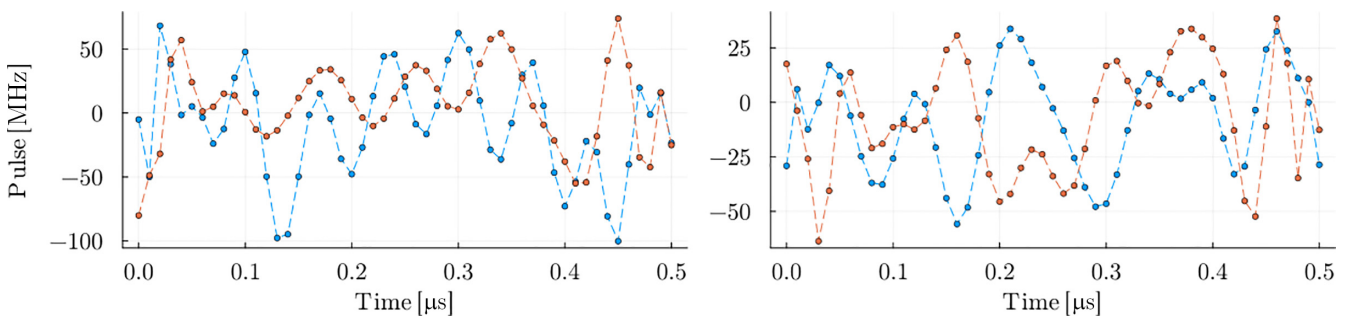

FIG. 13. Example X gate pulses for a quhexit, with a control on the pulses every 10 ns.

TABLE III. Total gate times T for $X^{(3,4)}$ and $R_X^{(3,4)}(\pi/10)$ on the quhexit ($d = 6$) and quoctit ($d = 8$), for pulse optimization (PO), $S + D$, and ECD decompositions, at the matched infidelities described in the text. The ranges for $S + D$ and ECD correspond to the per-block timing ranges $T_{S+D} \in [1, 50]$ μs and $T_{\text{ECD}} \in [0.1, 0.5]$ μs given in Sec. IV.

d	$\mathcal{I}$	T_{PO} [μs]	T_{S+D} [μs]	T_{ECD} [μs]
6	$\lesssim 10^{-5}$	0.5	[5,250]	[3.6,18]
8	$\lesssim 10^{-3}$	2.0	[4,200]	[6.4,32]

direct comparison of the resulting gate times for $X^{(3,4)}$ and $R_X^{(3,4)}(\pi/10)$. For the quoctit, the best pulse-optimization infidelity was $\mathcal{I}_{\text{PO}} \lesssim 10^{-3}$, which was matched by a four-block $S + D$ and a 64-block ECD decomposition. The corresponding total gate times for the two qudits and the three approaches are summarized in Table III.

For the quhexit, pulse optimization at fixed infidelity yields at least a factor of ∼10 reduction in gate time relative to either gate-set decomposition. For the quoctit, the improvement is a more modest factor of ∼2, which may indicate a diminishing advantage at higher dimension, or a reflection of the increased difficulty of classical optimization at high dimension.

While these results demonstrate it is possible to decompose gates of interest to a universal set, the decomposition generically has poor scaling $O(d^n)$ with $n \geqslant 3$ [152,156,167] that without mitigation will become prohibitive for large d. It is therefore worth investigating the relations between optimized parameters to seek heuristics to guide the optimization. Since $S + D$ requires fewer blocks than ECD, we will restrict our analysis to this decomposition and study the five-block versions of $X^{(3,4)}$, $X^{(2,3)}X^{(4,5)}$, $R_X^{(3,4)}(\pi/10)$ gates acting on a quhexit. The optimization is performed 200 times for each gate. The distributions of $\mathcal{I}$ are shown in Fig. 14, where one observes a clear ordering of $\mathcal{I}$ from smallest to largest of $R_X^{(3,4)}(\pi/10)$, $X^{(3,4)}$, $X^{(2,3)}X^{(4,5)}$. This supports the intuition that the distance from $\mathbb{1}$ is a proxy for difficulty.

Probing further, one can study the distribution of $S + D$ parameters across the 200 optimizations. We define $P(\alpha)$ as the empirical probability density of optimal displacement

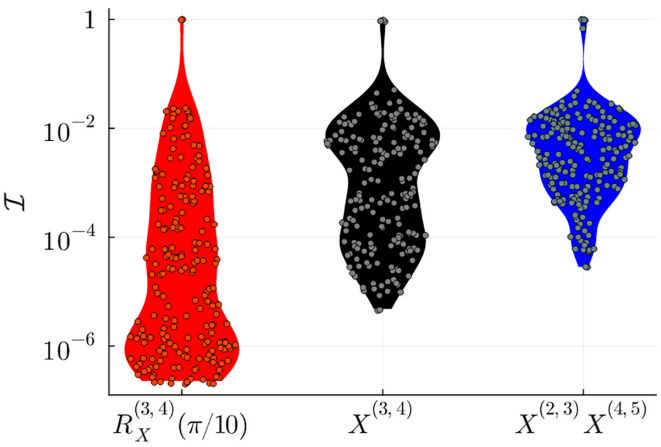

FIG. 14. Distribution of infidelities for five-block $S + D$ decomposition for different quhexit gates starting from the same initial parameters.

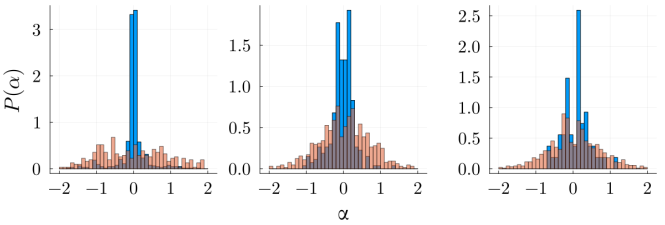

FIG. 15. Empirical distribution $P(\alpha)$ of the optimal displacement amplitudes α_k, pooled over all blocks and over all 200 optimization runs for each gate, for five-block $S + D$ decompositions on the quhexit. Orange histograms correspond to runs with $\mathcal{I} > 10^{-4}$ and blue histograms to runs with $\mathcal{I} \leqslant 10^{-4}$. From left to right: $R_X^{(3,4)}(\pi/10)$, $X^{(3,4)}$, and $X^{(2,3)}X^{(4,5)}$.

amplitudes α_k, obtained by pooling all α_k values from every block of every optimization run for a given gate, separated by the resulting infidelity, and normalizing the histogram. The analogous quantity $P(\theta)$ for the SNAP angles is found to be approximately uniform on $[-\pi, \pi]$ regardless of $\mathcal{I}$. In contrast, $P(\alpha)$ clearly distinguishes between good and poor optimizations, as seen in Fig. 15: optimizations with $\mathcal{I} \leqslant 10^{-4}$ yield $P(\alpha)$ sharply peaked near $\alpha = 0$, whereas optimizations with $\mathcal{I} > 10^{-4}$ produce a much broader $P(\alpha)$. This separation can be exploited to restrict the search parameter space in future optimizations.

Another route to accelerating the optimization is to exploit correlations between the optimal parameters. This is further motivated by the empirical observation that, for many gates, the optimal single-block $S + D$ decomposition ($B_{S+D} = 1$) satisfies $\alpha_1 = -\alpha_0$. We searched for correlations across optimizations of $X^{(3,4)}$. Figure 16 shows the correlation matrix corresponding to high and low infidelities. The correlation between two decomposition parameters $X, Y \in \{\alpha_k, \theta_k\}$ were computed using

$$\rho(X, Y) = \frac{1}{\sigma_X \sigma_Y (n-1)} \sum_{i=1}^{n} (X_i - \bar{X})(Y_i - \bar{Y}),$$

where σ_X and $\bar{X}$ are the standard deviation and mean of X. From the analysis of these correlations, we found different distributions when comparing low- and high-infidelity decompositions, when $\mathcal{I} > 10^{-5}$ the parameters are highly symmetric and with large correlation (or anticorrelation). In contrast, for $\mathcal{I} \leqslant 10^{-5}$ the parameters display decreasing correlation with gates farther away in the decomposition, i.e., gates in the same block are more correlated. Additionally,

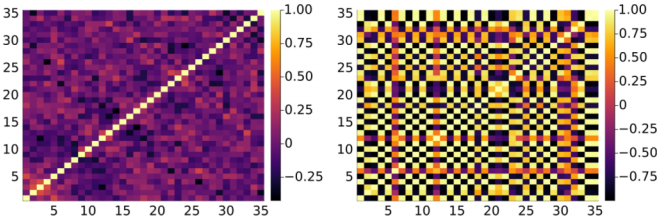

FIG. 16. Correlation matrix across parameters for $S + D$ decompositions of $X^{(3,4)}$ gate for a 5-block quhexit optimizations with (left) $\mathcal{I} \leqslant 10^{-5}$ and (right) $\mathcal{I} > 10^{-5}$.

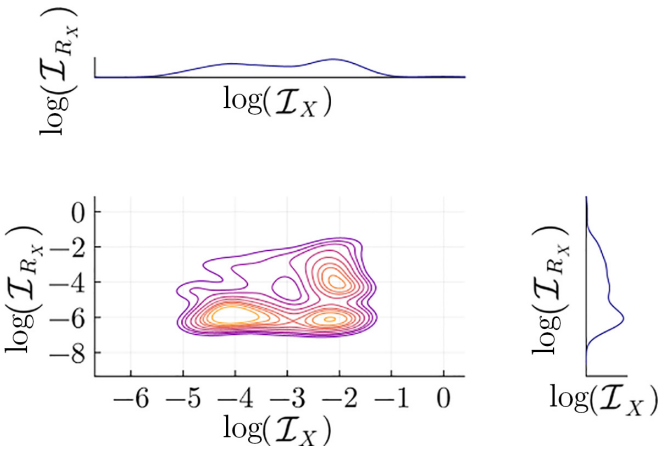

FIG. 17. Contours of the infidelities reached from same initial parameters for a R_X and a X gate.

there is some clear, weak anticorrelation structure suggesting new Ansätze for decomposition. Moreover, we found that the same initial conditions gave either good results for two different gates, or bad results for both. This is presented in Fig. 17 and suggests that there might be only a subspace of the parameters. However, this hypothesis requires more analytical work should be explored in the future.

A final consideration is the robustness to quantum noise. As a first step toward understanding this, we studied how $\mathcal{I}$ degrades when the parameters of the gate sets and pulse optimization were perturbed by a Gaussian noise. This basic noise model could be taken to approximate the stochastic fluctuations in the pulses driving the cavity and qubit. In this model the parameters are varied:

$$\alpha \to \alpha(x) = \alpha + \frac{1}{\beta|\alpha|\sqrt{2\pi}} \exp^{-\frac{1}{2}(\frac{x}{\beta|\alpha|})^2}, \tag{38}$$

where α is the parameter of the decomposition, and β is the square root of the variance.

Taking the noiseless $X^{(3,4)}$ gates on the quhexit with $\mathcal{I} \approx 10^{-5}$: five-block $S+D$ with $\mathcal{I} = 3.97 \times 10^{-5}$, 36-block ECD with $\mathcal{I} = 3.57 \times 10^{-5}$, and 30th-order pulse optimization for 0.5 μs with $\mathcal{I} = 3.10 \times 10^{-5}$; 50 random sample perturbations for each β are computed. The plot showing the results is Fig. 18. One observes that for $\beta < 10^{-4}$, the effect of noise is negligible. At larger β, pulse optimization and ECD are

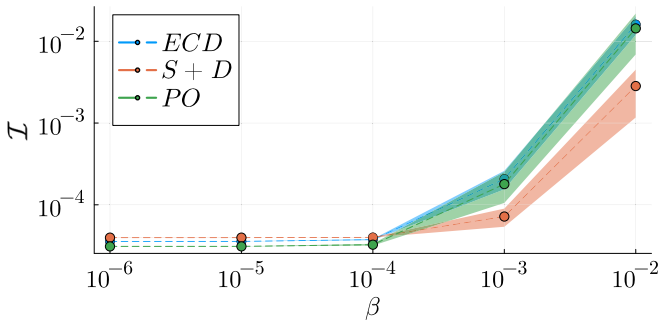

FIG. 18. Stability analysis for all the decompositions for a X gate on a quhexit.

found to demonstrate similar sized infidelities from the noise, while $S+D$ is found to have $\mathcal{I}$ about half the size of the other two.

VII. CONCLUSION

In this work we compared the three methods for implementing prototypical primitive gates for simulating lattice gauge theories on a 3D SRF qudit architecture: SNAP and displacement gates, ECD and single-qubit rotations, and optimal pulse control. We observed that lattice gauge theory primitives have common properties that distinguish them from general $SU(d)$ circuits, and thus the choice of decomposition could greatly change the algorithmic reach of devices. We found evidence that for subsets of permutation gates the number of SNAP and displacement blocks required to approximate them scales sublinearly with qudit dimension—in contrast to the general case of linear scaling—while the ECD block depth quadratically. In all cases, we found that decomposing a sequence of $X^{(a,b)}$ gates rather than individual reduces the total number of blocks required and consequently the total time for fixed infidelity. With the gate sets, we found that $\mathcal{I} < 10^{-5}$ is possible for $B_{S+D} = d$ and $B_{\text{ECD}} = d^2$. In addition, pulse decomposition with modest classical resources can outperform in run time for $d < 8$, while becoming comparable to gate decomposition for quoctits. Beyond run time, the $S+D$ decomposition was found to be more robust to Gaussian perturbations.

From the analysis of the $S+D$ decomposition, we found that good gate decompositions can be distinguished by symmetries, correlations, and the distribution of parameters. Further, we found correlations between initial conditions, suggesting a subspace useful as initial conditions.

Looking forward to the higher dimensions of practical interest for HEP (such as $d = 32$, which arises naturally from faithful representations of the discrete subgroups of $SU(3)$ and from recent fault-tolerant qudit code constructions [179]), the binding constraint is unlikely to be classical optimization. The optimization of Eq. (17) via LBFGS reached $\mathcal{I} \lesssim 10^{-5}$ at $d = 8$ with modest computational resources, including the iterative basis-growth scheme described in Sec. VI. The parameter count $O(Nd^2)$ and the per-evaluation cost $O(Nd^3)$ both extrapolate cleanly to the dimensions of interest on a single GPU node. Likewise, hardware coherence already supports the relevant block depths: recent ultracoherent SRF cavities reach $T_1 \sim 20$ ms and prepare Fock states up to $N = 20$ at over 95% fidelity [133], which leaves ample budget for $B \sim d^2$ ECD circuits at $T_{\text{ECD}} \sim 0.5$ μs. The dominant constraint instead becomes the precision of the time-domain controls. At $d \gtrsim 16$, drive-amplitude calibration, dispersive-shift uniformity $\chi(n)$ at high Fock occupation, and the higher-order self- and cross-Kerr corrections to the dispersive Hamiltonian, all of which our work treats as small, become non-negligible.

Building on these single-register results, a near-term proof of principle for non-Abelian lattice gauge theory on this hardware is within reach. As written in Eq. (6), the multiplication gate factors as a sum of single-qudit permutations P_g controlled by the state of the first register. Three ingredients therefore suffice to assemble $\mathfrak{U}_\times$: (1) the single-qudit primitives $\mathfrak{U}_{-1}$, $\mathfrak{U}_{\text{Tr}}$, and $\mathfrak{U}_F$ characterized in this work, applied

per cavity mode; (2) a CSUM-style two-mode primitive coupling adjacent cavity modes through a shared transmon, of the type already demonstrated in two-mode SRF platforms [133]; and (3) the single-qudit permutations of (1) applied to the target register conditioned on the control state. With these, a single-plaquette pure-gauge simulation for a small non-Abelian discrete subgroup (such as $\mathbb{Q}_8$ at $d = 8$ or $\mathbb{BT}$ at $d = 24$) becomes realizable, and the same architecture extends naturally to multilink simulations by tiling additional cavity modes. This provides a concrete path from the per-gate characterizations of the present work to a proof-of-principle lattice gauge theory simulation on SRF hardware.

From this initial work, we found that algorithmic advantage and hardware-specific gate times can be used to decisively choose between different native gate sets on qudit devices for the restricted class of gates needed for high energy physics. In the future, more in-depth investigation of the scaling of the decompositions should be performed. Further, it would be of invaluable help to try to implement some of the decompositions on the actual hardware, in order to better understand and develop a noise model and how the infidelities are affected.

ACKNOWLEDGMENTS

The authors would like to thank Tanay Roy, Nicholas Bornman, Andy Li, and Judah Unmuth-Yockey for helpful discussions. This work was supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Superconducting Quantum Materials and Systems Center (SQMS), under Contract No. 89243024CSC000002. Fermilab is operated by Fermi Forward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

DATA AVAILABILITY

There are no publicly available research data or software supporting this paper. Requests for further information or data should be sent to the authors.

[1] N. Klco, A. Roggero, and M. J. Savage, Standard Model physics and the digital quantum revolution: Thoughts about the interface, Rep. Prog. Phys. **85**, 064301 (2022).
[2] M. C. Bañuls *et al.*, Simulating lattice gauge theories within quantum technologies, Eur. Phys. J. D **74**, 165 (2020).
[3] C. W. Bauer *et al.*, Quantum simulation for high-energy physics, PRX Quantum **4**, 027001 (2023).
[4] A. Di Meglio *et al.*, Quantum computing for high-energy physics: State of the art and challenges. Summary of the QC4HEP Working Group, PRX Quantum **5**, 037001 (2024).
[5] 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).
[6] M. Rhodes, M. Kreshchuk, and S. Pathak, Exponential improvements in the simulation of lattice gauge theories using near-optimal techniques, PRX Quantum **5**, 040347 (2024).
[7] E. Gustafson *et al.*, Large scale multi-node simulations of $\mathbb{Z}_2$ gauge theory quantum circuits using Google Cloud Platform, arXiv:2110.07482.
[8] E. J. Gustafson, Prospects for simulating a qudit-based model of (1+1)D scalar QED, Phys. Rev. D **103**, 114505 (2021).
[9] P. P. Popov, M. Meth, M. Lewestein, P. Hauke, M. Ringbauer, E. Zohar, and V. Kasper, Variational quantum simulation of U(1) lattice gauge theories with qudit systems, Phys. Rev. Res. **6**, 013202 (2024).
[10] D. M. Kurkcuoglu, M. S. Alam, J. A. Job, A. C. Y. Li, A. Macridin, G. N. Perdue, and S. Providence, Quantum simulation of ϕ^4 theories in qudit systems, arXiv:2108.13357.
[11] G. Calajò, G. Magnifico, C. Edmunds, M. Ringbauer, S. Montangero, and P. Silvi, Digital quantum simulation of a (1+1)D SU(2) lattice gauge theory with ion qudits, PRX Quantum **5**, 040309 (2024).
[12] D. González-Cuadra, T. V. Zache, J. Carrasco, B. Kraus, and P. Zoller, Hardware efficient quantum simulation of non-Abelian gauge theories with qudits on Rydberg platforms, Phys. Rev. Lett. **129** 160501 (2022).
[13] M. Illa, C. E. P. Robin, and M. J. Savage, Qu8its for quantum simulations of lattice quantum chromodynamics, Phys. Rev. D **110**, 014507 (2024).
[14] T. V. Zache, D. González-Cuadra, and P. Zoller, Fermion-qudit quantum processors for simulating lattice gauge theories with matter, Quantum **7**, 1140 (2023).
[15] D. Janković, J.-G. Hartmann, M. Ruben, and P.-A. Hervieux, Noisy qudit vs multiple qubits: Conditions on gate efficiency for enhancing fidelity, npj Quantum Inf. **10**, 59 (2024).
[16] A. S. Nikolaeva, E. O. Kiktenko, and A. K. Fedorov, Efficient realization of quantum algorithms with qudits, EPJ Quantum Technol. **11**, 43 (2024).
[17] M. B. Mansky, S. L. n. Castillo, V. R. Puigvert, and C. Linnhoff-Popien, Near-optimal quantum circuit construction via Cartan decomposition, Phys. Rev. A **108**, 052607 (2023).
[18] M. Ringbauer, M. Meth, L. Postler, R. Stricker, R. Blatt, P. Schindler, and T. Monz, A universal qudit quantum processor with trapped ions, Nat. Phys. **18**, 1053 (2022).
[19] A. S. Nikolaeva *et al.*, Scalable improvement of the generalized Toffoli gate realization using trapped-ion-based qutrits, Phys. Rev. Lett. **135** 060601 (2025).
[20] P. J. Low, B. White, and C. Senko, Control and readout of a 13-level trapped ion qudit, arXiv:2306.03340.
[21] N. Goss *et al.*, High-fidelity qutrit entangling gates for superconducting circuits, Nat. Commun. **13**, 7481 (2022) [Erratum: Nat. Commun. **14**, 4256 (2023)].
[22] K. Luo *et al.*, Experimental realization of two qutrits gate with tunable coupling in superconducting circuits, Phys. Rev. Lett. **130**, 030603 (2023).
[23] N. Goss, S. Ferracin, A. Hashim, A. Carignan-Dugas, J. M. Kreikebaum, R. K. Naik, D. I. Santiago, and I. Siddiqi, Extending the computational reach of a superconducting qutrit processor, arXiv:2305.16507.
[24] S. Cao, M. Bakr, G. Campanaro, S. D. Fasciati, J. Wills, D. Lall, B. Shteynas, V. Chidambaram, I. Rungger, and P. Leek, Emulating two qubits with a four-level transmon qudit

for variational quantum algorithms, Quantum Sci. Technol. **9**, 035003 (2024).
[25] L. M. Seifert, Z. Li, T. Roy, D. I. Schuster, F. T. Chong, and J. M. Baker, Exploring ququart computation on a transmon using optimal control, Phys. Rev. A **108**, 062609 (2023).
[26] Z. Wang, R. W. Parker, E. Champion, and M. S. Blok, Systematic study of high E_J/E_C transmon qudits up to $d = 12$, Phys. Rev. Appl. **23**, 034046 (2025).
[27] A. Kruckenhauser, R. van Bijnen, T. V. Zache, M. Di Liberto, and P. Zoller, High-dimensional SO(4)-symmetric Rydberg manifolds for quantum simulation, Quantum Sci. Technol. **8**, 015020 (2023).
[28] S. R. Cohen and J. D. Thompson, Quantum computing with circular Rydberg atoms, PRX Quantum **2**, 030322 (2021).
[29] Y. Chi *et al.*, A programmable qudit-based quantum processor, Nat. Commun. **13**, 1166 (2022).
[30] V. Kasper, D. González-Cuadra, A. Hegde, A. Xia, A. Dauphin, F. Huber, E. Tiemann, M. Lewenstein, F. Jendrzejewski, and P. Hauke, Universal quantum computation and quantum error correction with ultracold atomic mixtures, Quantum Sci. Technol. **7**, 015008 (2022).
[31] T. Roy, T. Kim, A. Romanenko, and A. Grassellino, Qudit-based quantum computing with SRF cavities at Fermilab, PoS **LATTICE2023**, 127 (2024).
[32] R. W. Heeres, B. Vlastakis, E. Holland, S. Krastanov, V. V. Albert, L. Frunzio, L. Jiang, and R. J. Schoelkopf, Cavity state manipulation using photon-number selective phase gates, Phys. Rev. Lett. **115**, 137002 (2015).
[33] T. Fösel, S. Krastanov, F. Marquardt, and L. Jiang, Efficient cavity control with SNAP gates, arXiv:2004.14256.
[34] M. Kudra, M. Kervinen, I. Strandberg, S. Ahmed, M. Scigliuzzo, A. Osman, D. P. Lozano, M. O. Tholén, R. Borgani, D. B. Haviland, *et al.*, Robust preparation of Wigner-negative states with optimized SNAP-displacement sequences, PRX Quantum **3**, 030301 (2022).
[35] A. Eickbusch, V. Sivak, A. Z. Ding, S. S. Elder, S. R. Jha, J. Venkatraman, B. Royer, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, Fast universal control of an oscillator with weak dispersive coupling to a qubit, Nat. Phys. **18**, 1464 (2022).
[36] G. K. Brennen, S. S. Bullock, and D. P. O'Leary, Efficient circuits for exact-universal computation with qudits, Quant. Inf. Comput. **6**, 436 (2006).
[37] M. Howard and J. Vala, Qudit versions of the qubit $\pi/8$ gate, Phys. Rev. A **86**, 022316 (2012).
[38] A. Bocharov, A note on optimality of quantum circuits over metaplectic basis, Quant. Inf. Comput. **18**, 1 (2018).
[39] L. Yeh and J. van de Wetering, Constructing all qutrit controlled Clifford+T gates in Clifford+T, in *Reversible Computation (RC 2022)*, Lecture Notes in Computer Science, Vol. 13354 (Springer, Cham, 2022), pp. 28–50.
[40] Y. Nakajima, Y. Kawano, H. Sekigawa, M. Nakanishi, S. Yamashita, and Y. Nakashima, Synthesis of quantum circuits for d-level systems by using cosine-sine, Quant. Inf. Comput. **9**, 423 (2009).
[41] S. Prakash, A. R. Kalra, and A. Jain, A normal form for single-qudit Clifford+T operators, Quant. Inf. Proc. **20**, 341 (2021).
[42] A. R. Kalra, D. Valluri, and M. Mosca, Synthesis and arithmetic of single qutrit circuits, Quantum **9**, 1647 (2023).
[43] R. I. Booth and T. Carette, Complete ZX-calculi for the stabiliser fragment in odd prime dimensions, in *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 241 (Schloss Dagstuhl Leibniz-Zentrum für Informatik, 2022), pp. 24:1–24:15.
[44] A. N. Glaudell, N. J. Ross, J. van de Wetering, and L. Yeh, Qutrit metaplectic gates are a subset of Clifford+T, in *17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 232 (SchlossDagstuhlLeibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany, 2022), pp. 12:1–12:15.
[45] A. N. Glaudell, N. J. Ross, J. van de Wetering, and L. Yeh, Exact synthesis of multiqutrit Clifford-cyclotomic circuits, EPTCS **406**, 44 (2024).
[46] S. Prakash, A. Jain, B. Kapur, and S. Seth, Normal form for single-qutrit Clifford $+T$ operators and synthesis of single-qutrit gates, Phys. Rev. A **98**, 032304 (2018).
[47] H. Lamm, S. Lawrence, and Y. Yamauchi (NuQS Collaboration), General methods for digital quantum simulation of gauge theories, Phys. Rev. D **100**, 034518 (2019).
[48] 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).
[49] E. J. Gustafson, H. Lamm, F. Lovelace, and D. Musk, Primitive quantum gates for an $SU(2)$ discrete subgroup: Binary tetrahedral, Phys. Rev. D **106**, 114501 (2022).
[50] T. V. Zache, D. González-Cuadra, and P. Zoller, Quantum and classical spin-network algorithms for q-deformed Kogut-Susskind gauge theories, Phys. Rev. Lett. **131**, 171902 (2023).
[51] E. J. Gustafson, H. Lamm, and F. Lovelace, Primitive quantum gates for an $SU(2)$ discrete subgroup: Binary octahedral, Phys. Rev. D **109**, 054503 (2024).
[52] H. Lamm, Y.-Y. Li, J. Shu, Y.-L. Wang, and B. Xu, Block encodings of discrete subgroups on quantum computer, arXiv:2405.12890.
[53] E. M. Murairi, M. Sohaib Alam, H. Lamm, S. Hadfield, and E. Gustafson, Highly-efficient quantum Fourier transformations for some nonabelian groups, arXiv:2408.00075.
[54] T. D. Cohen, H. Lamm, S. Lawrence, and Y. Yamauchi (NuQS Collaboration), Quantum algorithms for transport coefficients in gauge theories, Phys. Rev. D **104**, 094514 (2021).
[55] M. Carena, H. Lamm, Y.-Y. Li, and W. Liu, Improved Hamiltonians for quantum simulations of gauge theories, Phys. Rev. Lett. **129**, 051601 (2022).
[56] E. J. Gustafson, Stout smearing on a quantum computer, arXiv:2211.05607.
[57] E. Gustafson and R. Van de Water, Improved fermion Hamiltonians for quantum simulation, PoS **LATTICE2023**, 215 (2024).
[58] M. Carena, H. Lamm, Y.-Y. Li, and W. Liu, Quantum error thresholds for gauge-redundant digitizations of lattice field theories, arXiv:2402.16780.
[59] E. Zohar, J. I. Cirac, and B. Reznik, Cold-atom quantum simulator for SU(2) Yang-Mills lattice gauge theory, Phys. Rev. Lett. **110**, 125304 (2013).
[60] E. Zohar, J. I. Cirac, and B. Reznik, Simulating compact quantum electrodynamics with ultracold atoms: Probing con-

finement and nonperturbative effects, Phys. Rev. Lett. **109**, 125302 (2012).
[61] E. Zohar, J. I. Cirac, and B. Reznik, Quantum simulations of gauge theories with ultracold atoms: Local gauge invariance from angular momentum conservation, Phys. Rev. A **88**, 023617 (2013).
[62] E. Zohar, J. I. Cirac, and B. Reznik, Quantum simulations of lattice gauge theories using ultracold atoms in optical lattices, Rep. Prog. Phys. **79**, 014401 (2016).
[63] A. Bazavov, Y. Meurice, S.-W. Tsai, J. Unmuth-Yockey, and J. Zhang, Gauge-invariant implementation of the Abelian Higgs model on optical lattices, Phys. Rev. D **92**, 076003 (2015).
[64] J. Zhang, J. Unmuth-Yockey, J. Zeiher, A. Bazavov, S. W. Tsai, and Y. Meurice, Quantum simulation of the universal features of the Polyakov loop, Phys. Rev. Lett. **121**, 223201 (2018).
[65] J. Unmuth-Yockey, J. Zhang, A. Bazavov, Y. Meurice, and S.-W. Tsai, Universal features of the Abelian Polyakov loop in 1+1 dimensions, Phys. Rev. D **98**, 094511 (2018).
[66] N. Klco, M. J. Savage, and J. R. Stryker, SU(2) non-Abelian gauge field theory in one dimension on digital quantum computers, Phys. Rev. D **101**, 074512 (2020).
[67] R. C. Farrell, M. Illa, A. N. Ciavarella, and M. J. Savage, Scalable circuits for preparing ground states on digital quantum computers: The Schwinger model vacuum on 100 qubits, PRX Quantum **5**, 020315 (2024).
[68] R. C. Farrell, M. Illa, A. N. Ciavarella, and M. J. Savage, Quantum simulations of hadron dynamics in the Schwinger model using 112 qubits, Phys. Rev. D **109**, 114510 (2024).
[69] A. Ciavarella, N. Klco, and M. J. Savage, A trailhead for quantum simulation of SU(3) Yang-Mills lattice gauge theory in the local multiplet basis, Phys. Rev. D **103** 094501 (2021).
[70] A. Bazavov, S. Catterall, R. G. Jha, and J. Unmuth-Yockey, Tensor renormalization group study of the non-Abelian Higgs model in two dimensions, Phys. Rev. D **99**, 114507 (2019).
[71] A. J. Buser, T. Bhattacharya, L. Cincio, and R. Gupta, Quantum simulation of the qubit-regularized $O(3)$-sigma model, Phys. Rev. D **102** 114514 (2020).
[72] T. Bhattacharya, A. J. Buser, S. Chandrasekharan, R. Gupta, and H. Singh, Qubit regularization of asymptotic freedom, Phys. Rev. Lett. **126** 172001 (2021).
[73] A. H. Z. Kavaki and R. Lewis, From square plaquettes to triamond lattices for SU(2) gauge theory, Commun. Phys. **7**, 208 (2024).
[74] E. M. Murairi, M. J. Cervia, H. Kumar, P. F. Bedaque, and A. Alexandru, How many quantum gates do gauge theories require? Phys. Rev. D **106** 094504 (2022).
[75] 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).
[76] I. Raychowdhury and J. R. Stryker, Solving Gauss's law on digital quantum computers with loop-string-hadron digitization, Phys. Rev. Res. **2**, 033039 (2020).
[77] S. V. Kadam, Theoretical developments in lattice gauge theory for applications in double-beta decay processes and quantum simulation, Ph.D. thesis, University of Maryland, College Park, 2023.
[78] Z. Davoudi, I. Raychowdhury, and A. Shaw, Search for efficient formulations for Hamiltonian simulation of non-Abelian lattice gauge theories, Phys. Rev. D **104** 074505 (2021).
[79] E. Mathew and I. Raychowdhury, Protecting local and global symmetries in simulating (1+1)D non-Abelian gauge theories, Phys. Rev. D **106**, 054510 (2022).
[80] C. W. Bauer and D. M. Grabowska, Efficient representation for simulating U(1) gauge theories on digital quantum computers at all values of the coupling, arXiv:2111.08015.
[81] D. M. Grabowska, C. Kane, B. Nachman, and C. W. Bauer, Overcoming exponential scaling with system size in Trotter-Suzuki implementations of constrained Hamiltonians: 2+1 U(1) lattice gauge theories, arXiv:2208.03333.
[82] D. M. Grabowska, C. F. Kane, and C. W. Bauer, A fully gauge-fixed SU(2) Hamiltonian for quantum simulations, arXiv:2409.10610.
[83] T. Li, Quantum simulations of quantum electrodynamics in Coulomb gauge, arXiv:2406.01204.
[84] P. Fontana, M. M. Riaza, and A. Celi, An efficient finite-resource formulation of non-Abelian lattice gauge theories beyond one dimension, Phys. Rev. X **15** 031065 (2025.
[85] M. Kreshchuk, W. M. Kirby, G. Goldstein, H. Beauchemin, and P. J. Love, Quantum simulation of quantum field theory in the light-front formulation, Phys. Rev. A **105** 032418 (2022).
[86] M. Kreshchuk, S. Jia, W. M. Kirby, G. Goldstein, J. P. Vary, and P. J. Love, Simulating hadronic physics on NISQ devices using basis light-front quantization, Phys. Rev. A **103** 062601 (2021).
[87] M. Kreshchuk, S. Jia, W. M. Kirby, G. Goldstein, J. P. Vary, and P. J. Love, Light-front field theory on current quantum computers, Entropy **23** 597 (2020).
[88] J. Liu and Y. Xin, Quantum simulation of quantum field theories as quantum chemistry, J. High Energ. Phys. 12 (2020) 011.
[89] M. Fromm, O. Philipsen, W. Unger, and C. Winterowd, Quantum gate sets for lattice QCD in the strong coupling limit: $N_f = 1$, EPJ Quantum Technol. **11**, 24 (2024).
[90] A. N. Ciavarella and C. W. Bauer, Quantum simulation of SU(3) lattice Yang Mills theory at leading order in large N, arXiv:2402.10265.
[91] A. Alexandru, P. F. Bedaque, A. Carosso, M. J. Cervia, E. M. Murairi, and A. Sheng, Fuzzy gauge theory for quantum computers, Phys. Rev. D **109**, 094502 (2024).
[92] R. Brower, S. Chandrasekharan, and U. J. Wiese, QCD as a quantum link model, Phys. Rev. D **60**, 094502 (1999).
[93] H. Singh, Qubit $O(N)$ nonlinear sigma models, arXiv:1911.12353.
[94] H. Singh and S. Chandrasekharan, Qubit regularization of the $O(3)$ sigma model, Phys. Rev. D **100**, 054505 (2019).
[95] U.-J. Wiese, Towards quantum simulating QCD, Nucl. Phys. A **931**, 246 (2014).
[96] R. C. Brower, D. Berenstein, and H. Kawai, Lattice gauge theory for a quantum computer, PoS **LATTICE2019**, 112 (2019).
[97] S. V. Mathis, G. Mazzola, and I. Tavernelli, Toward scalable simulations of lattice gauge theories on quantum computers, Phys. Rev. D **102**, 094501 (2020).
[98] J. C. Halimeh, R. Ott, I. P. McCulloch, B. Yang, and P. Hauke, Robustness of gauge-invariant dynamics against defects in ultracold-atom gauge theories, Phys. Rev. Res. **2**, 033361 (2020).

[99] T. Budde, M. K. Marinković, and J. C. P. Barros, Quantum many-body scars for arbitrary integer spin in 2 + 1D Abelian gauge theories, arXiv:2403.08892.

[100] J. Osborne, I. P. McCulloch, and J. C. Halimeh, Quantum many-body scarring in 2 + 1D gauge theories with dynamical matter, arXiv:2403.08858.

[101] J. Osborne, B. Yang, I. P. McCulloch, P. Hauke, and J. C. Halimeh, Spin-$SU(1)$ quantum link models with dynamical matter on a quantum simulator, arXiv:2305.06368.

[102] D. Luo, J. Shen, M. Highman, B. K. Clark, B. DeMarco, A. X. El-Khadra, and B. Gadway, A framework for simulating gauge theories with dipolar spin systems, Phys. Rev. A **102** 032617 (2020).

[103] E. Zohar and M. Burrello, Formulation of lattice gauge theories for quantum simulations, Phys. Rev. D **91**, 054506 (2015).

[104] E. Zohar, A. Farace, B. Reznik, and J. I. Cirac, Digital lattice gauge theories, Phys. Rev. A **95**, 023604 (2017).

[105] J. Bender, E. Zohar, A. Farace, and J. I. Cirac, Digital quantum simulation of lattice gauge theories in three spatial dimensions, New J. Phys. **20**, 093001 (2018).

[106] D. C. Hackett, K. Howe, C. Hughes, W. Jay, E. T. Neil, and J. N. Simone, Digitizing gauge fields: Lattice Monte Carlo results for future quantum computers, Phys. Rev. A **99**, 062341 (2019).

[107] A. Alexandru, P. F. Bedaque, S. Harmalkar, H. Lamm, S. Lawrence, and N. C. Warrington (NuQS Collaboration), Gluon field digitization for quantum computers, Phys. Rev. D **100**, 114501 (2019).

[108] A. Yamamoto, Real-time simulation of (2+1)-dimensional lattice gauge theory on qubits, Prog. Theor. Exp. Phys. **2021**, 013B06 (2021).

[109] Y. Ji, H. Lamm, and S. Zhu (NuQS Collaboration), Gluon field digitization via group space decimation for quantum computers, Phys. Rev. D **102**, 114513 (2020).

[110] J. F. Haase, L. Dellantonio, A. Celi, D. Paulson, A. Kan, K. Jansen, and C. A. Muschik, A resource efficient approach for quantum and classical simulations of gauge theories in particle physics, Quantum **5**, 393 (2021).

[111] M. Carena, H. Lamm, Y.-Y. Li, and W. Liu, Lattice renormalization of quantum simulations, Phys. Rev. D **104** 094519 (2021).

[112] T. Armon, S. Ashkenazi, G. García-Moreno, A. González-Tudela, and E. Zohar, Photon-mediated stroboscopic quantum simulation of a $\mathbb{Z}_2$ lattice gauge theory, Phys. Rev. Lett. **127**, 250501 (2021).

[113] C. Charles, E. J. Gustafson, E. Hardt, F. Herren, N. Hogan, H. Lamm, S. Starecheski, R. S. Van de Water, and M. L. Wagman, Simulating $\mathbb{Z}_2$ lattice gauge theory on a quantum computer, arXiv:2305.02361.

[114] R. Irmejs, M. C. Banuls, and J. I. Cirac, Quantum simulation of $\mathbb{Z}_2$ lattice gauge theory with minimal resources, Phys. Rev. D **108**, 074503 (2023).

[115] E. J. Gustafson and H. Lamm, Toward quantum simulations of $\mathbb{Z}_2$ gauge theory without state preparation, Phys. Rev. D **103**, 054507 (2021).

[116] T. Hartung, T. Jakobs, K. Jansen, J. Ostmeyer, and C. Urbach, Digitising SU(2) gauge fields and the freezing transition, Eur. Phys. J. C **82**, 237 (2022).

[117] M. Creutz, L. Jacobs, and C. Rebbi, Monte Carlo study of Abelian lattice gauge theories, Phys. Rev. D **20**, 1915 (1979).

[118] M. Creutz and M. Okawa, Generalized actions in $Z(p)$ lattice gauge theory, Nucl. Phys. B **220**, 149 (1983).

[119] G. Bhanot and C. Rebbi, Monte Carlo simulations of lattice models with finite subgroups of SU(3) as gauge groups, Phys. Rev. D **24**, 3319 (1981).

[120] D. Petcher and D. H. Weingarten, Monte Carlo calculations and a model of the phase structure for gauge theories on discrete subgroups of SU(2), Phys. Rev. D **22**, 2465 (1980).

[121] G. Bhanot, SU(3) lattice gauge theory in four-dimensions with a modified Wilson action, Phys. Lett. B **108**, 337 (1982).

[122] D. H. Weingarten and D. N. Petcher, Monte Carlo integration for lattice gauge theories with fermions, Phys. Lett. B **99**, 333 (1981).

[123] D. Weingarten, Monte Carlo evaluation of hadron masses in lattice gauge theories with fermions, Phys. Lett. B **109**, 57 (1982).

[124] Y. Ji, H. Lamm, and S. Zhu, Gluon digitization via character expansion for quantum computers, Phys. Rev. D **107** 114503 (2023).

[125] A. Alexandru, P. F. Bedaque, R. Brett, and H. Lamm, Spectrum of digitized QCD: Glueballs in a S(1080) gauge theory, Phys. Rev. D **105**, 114508 (2022).

[126] M. Carena, E. J. Gustafson, H. Lamm, Y.-Y. Li, and W. Liu, Gauge theory couplings on anisotropic lattices, Phys. Rev. D **106**, 114504 (2022).

[127] J. B. Kogut, 1/n expansions and the phase diagram of discrete lattice gauge theories with matter fields, Phys. Rev. D **21**, 2316 (1980).

[128] J. Romers, Discrete gauge theories in two spatial dimensions, Ph.D. thesis, Master's thesis, Universiteit van Amsterdam, 2007.

[129] E. H. Fradkin and S. H. Shenker, Phase diagrams of lattice gauge theories with Higgs fields, Phys. Rev. D **19**, 3682 (1979).

[130] D. Harlow and H. Ooguri, Symmetries in quantum field theory and quantum gravity, arXiv:1810.05338.

[131] D. Horn, M. Weinstein, and S. Yankielowicz, Hamiltonian approach to $Z(N)$ lattice gauge theories, Phys. Rev. D **19**, 3715 (1979).

[132] B. Assi and H. Lamm, Digitization and subduction of $SU(N)$ gauge theories, Phys. Rev. D **110**, 074511 (2024).

[133] T. Kim *et al.*, Ultracoherent superconducting cavity-based multiqudit platform with error-resilient control, arXiv:2506.03286.

[134] T. Tilma and E. C. G. Sudarshan, Generalized Euler angle parametrization for $SU(N)$, J. Phys. A: Math. Gen. **35**, 10467 (2002).

[135] S. Bertini, S. L. Cacciatori, and B. L. Cerchiai, On the Euler angles for SU(N), J. Math. Phys. **47**, 043510 (2006).

[136] T. Tilma and E. C. G. Sudarshan, Generalized Euler angle parameterization for $U(N)$ with applications to $SU(N)$ coset volume measures, J. Geom. Phys. **52**, 263 (2004).

[137] F. S. Khan and M. Perkowski, Synthesis of multi-qudit hybrid and d-valued quantum logic circuits by decomposition, Theor. Comput. Sci. **367**, 336 (2006).

[138] J. R. Stryker, Oracles for Gauss's law on digital quantum computers, Phys. Rev. A **99**, 042301 (2019).

[139] S. Harmalkar, H. Lamm, and S. Lawrence (NuQS Collaboration), Quantum simulation of field theories without state preparation, arXiv:2001.11490 [hep-lat].

[140] J. Saroni, H. Lamm, P. P. Orth, and T. Iadecola, Reconstructing thermal quantum quench dynamics from pure states, Phys. Rev. B **108**, 134301 (2023).
[141] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, Circuit quantum electrodynamics, Rev. Mod. Phys. **93**, 025005 (2021).
[142] M. Boissonneault, J. M. Gambetta, and A. Blais, Dispersive regime of circuit QED: Photon-dependent qubit dephasing and relaxation rates, Phys. Rev. A **79**, 013819 (2009).
[143] S. Krastanov, V. V. Albert, C. Shen, C.-L. Zou, R. W. Heeres, B. Vlastakis, R. J. Schoelkopf, and L. Jiang, Universal control of an oscillator with dispersive coupling to a qubit, Phys. Rev. A **92**, 040303(R) (2015).
[144] F. W. Strauch, All-resonant control of superconducting resonators, Phys. Rev. Lett. **109**, 210501 (2012).
[145] C. K. Law and J. H. Eberly, Arbitrary control of a quantum electromagnetic field, Phys. Rev. Lett. **76**, 1055 (1996).
[146] T. F. Roque, A. A. Clerk, and H. Ribeiro, Engineering fast high-fidelity quantum operations with constrained interactions, arXiv:2003.12096.
[147] X. Fei, L. T. Brady, J. Larson, S. Leyffer, and S. Shen, Binary control pulse optimization for quantum systems, Quantum **7**, 892 (2023).
[148] L. M. Seifert, J. Chadwick, A. Litteken, F. T. Chong, and J. M. Baker, Time-efficient qudit gates through incremental pulse re-seeding, in *2022 IEEE International Conference on Quantum Computing and Engineering* (IEEE, Piscataway, NJ, 2022), pp. 304–313.
[149] P. Luchi, F. Turro, V. Amitrano, F. Pederiva, X. Wu, K. Wendt, J. L. Dubois, and S. Quaglioni, Control optimization for parametric Hamiltonians by pulse reconstruction, Eur. Phys. J. A **59**, 196 (2023).
[150] A. Reineri, S. Zorzetti, T. Roy, and X. You, Exploration of superconducting multi-mode cavity architectures for quantum computing, arXiv:2308.11740.
[151] X. You *et al.*, Crosstalk-robust quantum control in multimode bosonic systems, Phys. Rev. Appl. **22**, 044072 (2024).
[152] J. Job, Efficient, direct compilation of SU(N) operations into SNAP & displacement gates, arXiv:2307.11900.
[153] O. Ogunkoya, J. Kim, B. Peng, A. B. Özgüler, and Y. Alexeev, Qutrit circuits and algebraic relations: A pathway to efficient spin-1 Hamiltonian simulation, Phys. Rev. A **109**, 012426 (2024).
[154] P. Roy, J. van de Wetering, and L. Yeh, The qudit ZH-calculus: Generalised Toffoli+Hadamard and universality, EPTCS **384**, 142 (2023).
[155] K. Mato, M. Ringbauer, S. Hillmich, and R. Wille, Compilation of entangling gates for high-dimensional quantum systems, in *Proceedings of the 28th Asia and South Pacific Design Automation Conference, ASPDAC'23* (ACM Press, New York, NY, 2023), pp. 202–208.
[156] O. Ogunkoya, K. Morris, and D. M. Kürkçüoglu, Investigating parameter trainability in the SNAP-displacement protocol of a qudit system, arXiv:2309.14942.
[157] B. Şahinoğlu and R. D. Somma, Hamiltonian simulation in the low energy subspace, npj Quantum Inf **7**, 119 (2020).
[158] T. Hatomura, State-dependent error bound for digital quantum simulation of driven systems, Phys. Rev. A **105**, L050601 (2022).
[159] D. Burgarth, N. Galke, A. Hahn, and L. van Luijk, State-dependent Trotter limits and their approximations, Phys. Rev. A **107**, L040201 (2023).
[160] K. Kormann, S. Holmgren, and H. Karlsson, A Fourier-coefficient based solution of an optimal control problem in quantum chemistry, J. Optim Theory Appl. **147**, 491 (2010).
[161] A. Pagano, M. M. Müller, T. Calarco, S. Montangero, and P. Rembold, The role of bases in quantum optimal control, Phys. Rev. A **110** 062608 (2024).
[162] N. Karjanto, Properties of Chebyshev polynomials, arXiv:2002.01342.
[163] C. Gerry and P. Knight, *Introductory Quantum Optics* (Cambridge University Press, Cambridge, 2004).
[164] M. S. Alam *et al.*, Quantum computing hardware for HEP algorithms and sensing, arXiv:2204.08605.
[165] J. Y. Araz, S. Bhowmick, M. Grau, T. J. McEntire, and F. Ringer, State preparation of lattice field theories using quantum optimal control, Phys. Rev. D **111**, 034506 (2025).
[166] C. J. Axline, L. D. Burkhart, W. Pfaff, M. Zhang, K. Chou, P. Campagne-Ibarcq, P. Reinhold, L. Frunzio, S. Girvin, L. Jiang, *et al.*, On-demand quantum state transfer and entanglement between remote microwave cavity memories, Nat. Phys. **14**, 705 (2018).
[167] T. Fösel, S. Krastanov, F. Marquardt, and L. Jiang, Efficient cavity control with SNAP gates, arXiv:2004.14256.
[168] J. J. Vartiainen, M. Möttönen, and M. M. Salomaa, Efficient decomposition of quantum gates, Phys. Rev. Lett. **92**, 177902 (2004).
[169] N. P. D. Sawaya, T. Menke, T. H. Kyaw, S. Johri, A. Aspuru-Guzik, and G. G. Guerreschi, Resource-efficient digital quantum simulation of d-level systems for photonic, vibrational, and spin-s Hamiltonians, npj Quantum Inf. **6**, 49 (2020).
[170] O. Di Matteo, A. McCoy, P. Gysbers, T. Miyagi, R. M. Woloshyn, and P. Navrátil, Improving Hamiltonian encodings with the Gray code, Phys. Rev. A **103**, 042405 (2021).
[171] C. C. Chang, K. S. McElvain, E. Rrapaj, and Y. Wu, Improving Schrödinger equation implementations with Gray code for adiabatic quantum computers, PRX Quantum **3**, 020356 (2022).
[172] M. G. Kendall, A new measure of rank correlation, Biometrika **30**, 81 (1938).
[173] F. J. Damerau, A technique for computer detection and correction of spelling errors, Commun. ACM **7**, 171 (1964).
[174] V. Levenshtein, Binary codes capable of correcting deletions, insertions, and reversals, in Proceedings of the Soviet physics doklady (1966).
[175] V. V. Shende, I. L. Markov, and S. S. Bullock, Minimal universal two-qubit controlled-not-based circuits, Phys. Rev. A **69**, 062321 (2004).
[176] D. C. McKay, I. Hincks, E. J. Pritchett, M. Carroll, L. C. G. Govia, and S. T. Merkel, Benchmarking quantum processor performance at scale, arXiv:2311.05933.
[177] D. C. Liu and J. Nocedal, On the limited memory BFGS method for large scale optimization, Math. Program. **45**, 503 (1989).
[178] https://github.com/AndreaMaestri18/Opt3DQalgs.
[179] M. J. Cervia, H. Lamm, D. Liu, E. M. Murairi, and S. Zhu, Magic state distillation using asymptotically good codes on qudits, arXiv:2512.21874.