

Fault-tolerant resource comparison of qudit and qubit encodings for diagonal quadratic operators

Samuel Godwood^{1,*}, Doğa Murat Kürkçüoğlu², Gabriel N. Perdue², Marina Maneyro¹, and Alessandro Roggero^{3,4}

¹*Department of Physics, University of Liverpool, Liverpool L69 7ZL, United Kingdom*

²*Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

³*Dipartimento di Fisica, University of Trento, Via Sommarive 14, I-38123 Povo, Trento, Italy*

⁴*INFN-TIFPA Trento Institute of Fundamental Physics and Applications, Via Sommarive 14, I-38123 Povo, Trento, Italy*



(Received 27 May 2026; revised 20 July 2026; accepted 23 July 2026; published 12 August 2026)

Finite local Hilbert-space truncations arise naturally in quantum simulations of lattice field theories and motivate qudit encodings, but their fault-tolerant advantage over qubit encodings remains unclear. We compare the non-Clifford cost of implementing quadratic diagonal evolutions, exemplified by $U = e^{-it\phi_z^2}$ in a uniform field-amplitude discretization of a real scalar field, using either one logical d -level qudit or $n_b = \lceil \log_2 d \rceil$ logical qubits. We analyze two standard settings: product-formula simulation and linear combination of unitaries (LCU) per block encoding, taking the resource metric to be the number of non-Clifford gates after synthesis into a discrete logical gate set. Because tight synthesis bounds for general single-qudit rotations are not known, we express the qudit constructions in terms of embedded two-level $SU(2)$ rotations and derive explicit finite- d break-even conditions for their synthesis cost; these serve as compiler targets for when qudit encodings can outperform the qubit baseline. Within the constructive models studied here, product-formula implementations would require an exponentially stronger per-primitive synthesis advantage for qudits to win asymptotically, while in the LCU setting the qubit encoding is asymptotically cheaper in d . Nevertheless, the finite- d threshold analysis identifies low-dimensional regions in which qudits can yield meaningful constant-factor savings, particularly for LCU-based implementations. As a secondary analysis of the LCU construction, we use an idealized negligible-overhead qubit-qudit code-switching model to give an absolute T -count comparison and reinterpret the savings as an allowable per-switch overhead budget.

DOI: [10.1103/832q-bqlx](https://doi.org/10.1103/832q-bqlx)

I. INTRODUCTION

Quantum computers promise first-principles access to strongly correlated quantum dynamics that are difficult to treat classically, with applications spanning high-energy physics and lattice gauge theory, condensed matter and materials, and quantum chemistry [1–6]. In particular, real-time Hamiltonian dynamics and finite-density settings can challenge standard classical methods, including regimes where Monte Carlo approaches suffer from severe sign problems [7–9]. This has motivated a broad program of quantum algorithms for Hamiltonian simulation, spanning product-formula, qubitization and quantum signal processing, and linear combination of unitaries (LCU) [10–16]. Concrete proposals exist for simulating scalar field theories, gauge theories, and related lattice models using digital quantum computers [3,17–23].

A recurring technical feature of these targets is the presence of local degrees of freedom that are naturally infinite-dimensional (bosonic modes) or continuous and

compact (gauge links), requiring a truncation or digitization to a finite local Hilbert space to obtain a finite model [22,24–27]. For scalar fields, one common approach is a field-amplitude (position-basis) digitization on a uniform grid [17,18,28]; for gauge links and rotors, one often imposes representation cutoffs or angular-momentum truncations [22,24–26,29]. Importantly, physically relevant truncations can be modest in some parameter regimes and model families—for example, finite-group and quantum-link formulations, few-qubit Schwinger-model realizations, and qudit gauge-field encodings with local dimensions such as $d = 3$ and $d = 8$ [19,23,25,30–32]—making the structure of the finite- d encoding consequential for end-to-end algorithmic cost.

These observations have motivated sustained interest in qudit-based quantum computation, where the native information carrier is a d -level system rather than a qubit [33–37]. From an algorithmic perspective, qudits can offer a more direct representation of truncated local Hilbert spaces [23,32,33], and in some settings can reduce circuit depth or simplify arithmetic by exploiting generalized Clifford structure [33–36]. More broadly, qudit-based information processing has grown into an active area spanning gate-decomposition techniques, algorithm design, and quantum error correction, as surveyed in recent reviews [33,38]. From an implementation perspective, several leading hardware platforms provide access to well-controlled multilevel structure,

*Contact author: s.godwood@liverpool.ac.uk

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

enabling qudit gates or encoded qudits within a larger physical Hilbert space [39–42].

Prevailing expectations for quantum advantage typically rely on fault-tolerant computation, where error correction and compilation overhead dominate [43–46]. In this regime, resource estimates are often governed by the cost of implementing non-Clifford operations in a discrete logical gate set (e.g., via magic-state distillation) [43,45,47,48]. While there is an extensive literature on qudit algorithms and near-term qudit simulation prospects [23,26,27,32,33,49–51], systematic fault-tolerant resource comparisons that track non-Clifford resources across qudit and qubit encodings remain less standardized, especially at general d .

In this paper, we give a fault-tolerant resource analysis for a representative and widely occurring problem class: onsite quadratic diagonal evolutions, exemplified by $e^{-it\phi_x^2}$ for a digitized scalar field ϕ_x in the field-amplitude basis in position space [3,17,18]. While our primary presentation uses ϕ_x for concreteness, the fault-tolerant qubit-qudit comparison established in this work concerns onsite quadratic diagonal unitaries of the form $e^{-itg(N)}$. Here, g is quadratic in a truncated integer-valued local operator N (e.g., a number operator), as arises in lattice Hamiltonians [24,26,27].

We study two standard simulation paradigms. In Regime 1, we compare a single product-formula (Trotter) step implementing the onsite diagonal evolution using either one d -level qudit acting natively on the d -level truncation, or $n_b = \lceil \log_2 d \rceil$ qubits in the standard binary embedding of the same physical truncation [10–12]. In Regime 2, we compare block-encodings in the LCU framework, where the cost is set by the PREPARE and SELECT oracle implementations [13–16]. Across both regimes, our metric of fault-tolerant cost is the non-Clifford count after synthesis into a discrete logical gate set; synthesis assumptions enter only through the cost of approximating continuous one-parameter primitives [qubit R_z rotations and embedded two-level qudit SU(2) rotations] to a target precision ε [47,48]. Accordingly, the results should be interpreted as a compilation- and resource-model comparison under a chosen fault-tolerant proxy, and not as a complete architecture-level comparison of total logical qudit versus logical qubit overhead. This would additionally require architecture-specific assumptions about code conversion, state injection or distillation, and space-time costs.

Because tight, gate-set-dependent synthesis constants for the required single-qudit embedded SU(2) primitives are not known uniformly in d , we cannot claim closed-form costs for $d > 3$. Instead, we parametrize qudit synthesis by a logarithmic prefactor a and report explicit break-even conditions $a \lesssim a_{\max}^{\text{PF}}(d, \varepsilon)$ in Regime 1 and $a \lesssim a_{\max}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}})$ in Regime 2. $a_{\max}^{\text{PF}}(d, \varepsilon)$ and $a_{\max}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}})$ are break-even compiler targets: they are the largest allowed values of the logarithmic synthesis prefactor a of the relevant embedded two-level qudit primitives for which the qudit construction still uses fewer non-Clifford gates than the stated qubit baseline under our cost model, so a smaller value denotes a more demanding compilation requirement. They are formally defined and analyzed in Secs. III D and IV D 2, respectively.

With this caveat, our contributions are as follows: (1) we derive explicit finite- d break-even synthesis thresholds for the required embedded two-level qudit SU(2) primitives in

TABLE I. End-to-end break-even synthesis prefactor $a_{\max}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}})$ for the fixed-encoding Regime 2 qudit LCU under the coupled per-query error budget, evaluated at $\varepsilon_{\text{sim}} = 10^{-6}$ and $t = 0.1$. Only prime local dimensions d up to 19 are included in the fixed-encoding analysis.

d	$a_{\max}^{\text{LCU}}(d, 0.1, 10^{-6})$
3	2.56
5	1.32
7	0.85
11	0.53
13	0.44
17	0.34
19	0.30

both regimes; (2) we place these thresholds in context with asymptotic results under the explicit models studied here (see Table II): in Regime 1, we show that asymptotic qudit advantage would require an exponentially strong per-primitive synthesis improvement over the qubit baseline, while in Regime 2 we derive asymptotic non-Clifford scalings, for which the qubit implementation is asymptotically cheaper; (3) as a secondary diagnostic, under an idealized free qubit-qudit code-switching assumption in Regime 2, we translate the resulting finite- d savings into a conditional per-switch overhead budget for a hybrid route.

The asymptotic comparisons of contribution (2) are unconditional consequences of the explicit constructions and cost models studied here: within these models, the qubit encoding is asymptotically no more expensive in non-Clifford count in both regimes, and in Regime 2 this holds for any fixed positive synthesis prefactor. By contrast, the finite- d break-even results of contribution (1), expressed through the thresholds $a_{\max}^{\text{PF}}(d, \varepsilon)$ and $a_{\max}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}})$, are conditional compiler-target statements: whether a finite- d qudit advantage is realized depends on the achievable synthesis cost of the embedded two-level qudit primitives, which is not yet known uniformly in d . The code-switching

TABLE II. Leading-order asymptotic scalings for a single onsite quadratic diagonal evolution $e^{-it\phi_x^2}$, comparing the $n_b = \lceil \log_2 d \rceil$ -qubit encoding with the native d -level qudit encoding, in the product-formula (PF, per Trotter step) and LCU (LCU, per block-encoding call) regimes. “Rotations” counts synthesized one-parameter rotations (R_z for qubits, embedded two-level SU(2) for qudits); the qubit LCU baseline uses Toffoli or T gates rather than synthesized rotations. Here ν and ε are errors for one Trotter step and for one block-encoding call, respectively.

Cost	Qubit	Qudit
PF, rotations per step	$\Theta(n_b^2)$	$\Theta(d)$
PF, non-Clifford per step, total step error ν	$\Theta(n_b^2 \log \frac{n_b}{\nu})$	$\Theta(d \log \frac{d}{\nu})$
LCU, rotations per block-encoding call	0	$\Theta(d)$
LCU, non-Clifford per call, total call error ε	$\Theta(n_b + \log \frac{1}{\varepsilon})$	$\Theta(d \log \frac{d}{\varepsilon})$

comparison of contribution (3) is further conditional on an idealized negligible-overhead qubit-qudit code-switching model, and is therefore stated only as an allowable per-switch overhead budget rather than a demonstrated advantage.

The remainder of the work is organized as follows: Section II introduces the truncation model, the qubit and qudit encodings, and the embedded two-level SU(2) primitives used to express single-qudit implementations. Section III develops Regime 1 (product-formula) implementations: Sections III A and III B give the corresponding non-Clifford cost models for qubits and qudits, respectively, Sec. III C studies the asymptotic synthesis behavior, and Sec. III D provides synthesis-parametrized crossover criteria. Section IV develops Regime 2 (LCU per block-encoding) implementations: Sections IV A and IV B give the corresponding constructions and costs, Sec. IV C studies asymptotic behavior, and Secs. IV D 2 and IV D 3 analyze finite- d encoding and hybrid finite- d comparisons, respectively. We discuss our findings in Sec. V, and conclude in Sec. VI.

II. PROBLEM SETUP AND DEFINITIONS

A representative and widely occurring primitive in lattice Hamiltonian simulation is onsite diagonal evolution generated by a quadratic function of a local, digitized degree of freedom. In this work, we use a digitized real scalar field operator ϕ_x at site x as a concrete example, focusing on the implementation of the single-site unitary $e^{-it\phi_x^2}$. We encode ϕ_x using a uniform grid in the field-amplitude basis. This digitization scheme is consistent with the standard formalism for scalar fields in quantum simulation, often referred to as the Jordan–Lee–Preskill (JLP) basis [17,18,52]. Although our focus is strictly on the onsite mass term $\propto \phi_x^2$, we note that the same ideas apply to terms $\propto \pi_x^2$, since π_x is diagonal in the Fourier-conjugate basis. The qudit Fourier transform is Clifford [33], so this basis change does not add non-Clifford cost in the single-qudit setting; by contrast, the n_b -qubit quantum Fourier transform (QFT) is generally non-Clifford and can add overhead in the qubit encoding. Beyond our focus on real scalar fields—important in lattice ϕ^4 theories and dynamical-pion nuclear effective-field theories (EFTs) [53]—operators of this form also arise in lattice gauge theories [27,54,55], condensed matter and quantum optics [56,57], and quantum chemistry [58].

To truncate and discretize ϕ_x , we choose a maximum field value $\phi_{\max}$ and a grid spacing δ_ϕ . Defining

$$M = \frac{\phi_{\max}}{\delta_\phi} \in \mathbb{Z}, \quad d = 2M + 1, \quad (1)$$

we label the basis states $\{|n\rangle\}_{n=0}^{d-1}$. Note that this symmetric grid implies odd d , which we leave implicit throughout unless stated otherwise. This bounded, uniformly spaced field-amplitude truncation is precisely the finite-dimensional local Hilbert-space digitization of the Jordan–Lee–Preskill (JLP) framework for quantum simulation of scalar field theories [17,18,52], to which the encoding used here directly corresponds. The field operator on site x is defined diagonally

in this basis as

$$\phi_x = \sum_{n=0}^{d-1} \lambda_n |n\rangle\langle n|, \quad \lambda_n = -\phi_{\max} + n \delta_\phi. \quad (2)$$

A single d -level qudit naturally represents these $2M + 1$ equally spaced values

$$\{-\phi_{\max}, -\phi_{\max} + \delta_\phi, \dots, +\phi_{\max}\}. \quad (3)$$

For qubit implementations, we use an $n_b = \lceil \log_2 d \rceil$ -qubit register, but the precise embedding depends on the construction being compared. In the product-formula setting of Sec. III, we use the standard binary embedding of the d physical field-amplitude states $\{|n\rangle\}_{n=0}^{d-1}$ into the first d computational basis states of the 2^{n_b} -dimensional register. The digitized field operator is taken to act diagonally, with the remaining $2^{n_b} - d$ computational basis states lying outside the truncation and therefore unused. In the LCU per block-encoding setting of Sec. IV, we instead use the nearest signed-binary register supported by n_b qubits, following the standard projector-LCU baseline of Refs. [59,60]. The qudit and qubit constructions are motivated by the same physical $d = 2M + 1$ target truncation: the qudit construction acts natively on the d -level field-amplitude truncation, while the qubit baseline implements a signed-binary extension of the rounded squared operator that agrees with the physical target on the physical interval $\ell \in \{-M, \dots, M\}$. The Regime-2 finite- d comparisons in this paper are therefore construction-specific comparisons between the native qudit LCU construction and the binary qubit projector-LCU.

To perform an arbitrary SU(d) rotation on a single qudit, one needs a universal basis of operations. This basis can be constructed from a collection of embedded SU(2) rotations that couple different basis states. The operators that generate these fundamental SU(2) subalgebras are defined as generalizations of the Pauli matrices [61]. Let $\{|n\rangle\}_{n=0}^{d-1}$ be the computational basis. For $0 \leq b < c \leq d - 1$, define the embedded two-level generators:

$$X^{(b,c)} = |b\rangle\langle c| + |c\rangle\langle b|, \quad (4)$$

$$Y^{(b,c)} = -i |b\rangle\langle c| + i |c\rangle\langle b|, \quad (5)$$

$$Z^{(b,c)} = |b\rangle\langle b| - |c\rangle\langle c|. \quad (6)$$

We can then define rotations

$$R_G^{(b,c)}(\theta) = \exp\left(-\frac{i\theta}{2} G^{(b,c)}\right), \quad G \in \{X, Y, Z\}, \quad (7)$$

which act only on the $\{|b\rangle, |c\rangle\}$ subspace.

Throughout, unless stated otherwise, we compare a single fault-tolerant d -level qudit with a register of n_b qubits encoding the field operator at site x .

Ideally, we would have a closed-form synthesis bound for approximating the d -dimensional diagonal qudit gate $e^{-it\phi_x^2}$ to accuracy ε using a fixed fault-tolerant qudit gate set. For SU(p) with prime p , a worst-case lower bound for Clifford + T synthesis is given in Ref. [62]. However, although volume and covering arguments provide existential worst-case lower bounds on ε -approximation in SU(p) [62], these bounds do not certify that a particular structured diagonal, such as $e^{-it\phi_x^2}$,

is the worst case. Determining tight cost for a specific target typically requires finding the nearest implementable exactly synthesizable unitary V and the shortest word implementing it, leading to NP-complete decision problems [63]. Extending the single-qutrit ($d = 3$) Clifford + R synthesis analysis of Ref. [64] to arbitrary odd d would require (i) a dimension- d normal form and exact-synthesis framework for a fixed fault-tolerant gate set, and (ii) dimension-specific number-theoretic and algorithmic analysis. Accordingly, deriving a closed-form non-Clifford cost—or proving that $e^{-it\phi_x^2}$ saturates generic lower bounds—for arbitrary d is beyond our scope and is left for future work.

Furthermore, in the finite- d compiler-target analysis in Secs. III D and IV D 2, we restrict to prime local dimensions. This restriction reflects the current state of the fault-tolerant qudit-synthesis literature: Clifford + T normal-form and synthesis results and qudit magic-state distillation constructions are most directly developed for prime-dimensional qudits [34,35,37,62,64].

Therefore, for resource estimates, we first express algorithmic costs in terms of continuous one-parameter rotations [e.g., qubit R_z rotations or embedded qudit $SU(2)$ rotations]. In a fault-tolerant setting, these rotations must then be approximately synthesized over a discrete logical gate set to precision ε . Accordingly, the primary resource metric we compare throughout is the resulting number of non-Clifford gates (i.e., the non-Clifford primitives required by the discrete fault-tolerant gate set after synthesis), and in Sec. IV D 3 we specifically compare T counts.

III. REGIME 1: TIME EVOLUTION VIA PRODUCT FORMULAS (TROTTERIZATION)

In this section, we compare the fault-tolerant cost of implementing a single product-formula step of the onsite unitary $e^{-it\phi_x^2}$ in the qubit and qudit encodings. We refer the reader to Ref. [65] for background on product formulas. For fixed formula order and target simulation parameters, the required step count r is encoding independent, so the relative resource comparison reduces to the non-Clifford cost per step at matched accuracy.

A. The qubit implementation

Lemma 1. A single Trotter step of $e^{-it\phi_x^2}$ on a register of $n_b = \lceil \log_2 d \rceil$ qubits can be implemented using at most $[n_b(n_b + 1)]/2$ R_z gates.

Proof. We exploit the binary structure of the register by writing the integer level index as $n = \sum_{m=0}^{n_b-1} 2^m q_m$ with $q_m \in \{0, 1\}$ [53]. Mapping these bits to Pauli operators via $q_m = (1 - Z^{(m)})/2$, the field operator becomes

$$\phi_x = P\mathbb{I} + Q \sum_{m=0}^{n_b-1} 2^m Z^{(m)}, \quad (8)$$

with $P = -\phi_{\max} + \frac{\delta_\phi}{2}(2^{n_b} - 1)$, $Q = -\frac{\delta_\phi}{2}$. Here $Z^{(m)}$ denotes the Pauli Z operator acting on the m th qubit of the register, with $m = 0$ corresponding to the least significant bit. Squaring

the operator yields

$$\begin{aligned} \phi_x^2 &= \left(P\mathbb{I} + Q \sum_{m=0}^{n_b-1} 2^m Z^{(m)} \right)^2 \\ &= P^2\mathbb{I} + 2PQ \sum_{m=0}^{n_b-1} 2^m Z^{(m)} \\ &\quad + Q^2 \sum_{m,m'=0}^{n_b-1} 2^{m+m'} Z^{(m)} Z^{(m')}. \end{aligned} \quad (9)$$

The time-evolution operator $e^{-it\phi_x^2}$ is therefore a product of commuting rotations:

$$\begin{aligned} e^{-it\phi_x^2} &= \exp \left[-it \left(P^2\mathbb{I} + 2PQ \sum_{m=0}^{n_b-1} 2^m Z^{(m)} \right. \right. \\ &\quad \left. \left. + Q^2 \sum_{m,m'=0}^{n_b-1} 2^{m+m'} Z^{(m)} Z^{(m')} \right) \right]. \end{aligned} \quad (10)$$

This requires n_b $R_z(\theta)$ rotations (from the linear terms) and $\frac{n_b(n_b-1)}{2}$ two-qubit ZZ rotations (from the quadratic terms). Since $e^{-i\frac{\theta}{2}(Z \otimes Z)}$ decomposes into two controlled-NOT (CNOT) gates and one $R_z(\theta)$ gate, the total number of $R_z(\theta)$ gates is thus $(n_b^2 + n_b)/2$. ■

B. The qudit implementation

In the qudit computational basis $\{|n\rangle\}_{n=0}^{d-1}$, the digitized field operator ϕ_x is diagonal with eigenvalues λ_n (as defined in Sec. II). It follows that

$$\phi_x^2 = \sum_{n=0}^{d-1} \lambda_n^2 |n\rangle\langle n|, \quad (11)$$

and the corresponding onsite evolution is the diagonal unitary:

$$e^{-it\phi_x^2} = \sum_{n=0}^{d-1} e^{-it\lambda_n^2} |n\rangle\langle n|. \quad (12)$$

Lemma 2. Let $U = \sum_{n=0}^{d-1} e^{-i\beta_n} |n\rangle\langle n|$ be a diagonal unitary on a single d -level qudit. Then U can be implemented up to a global phase using at most $d - 1$ $R_Z^{(k,k+1)}(\theta_k)$ rotations. In particular, taking $\beta_n = t\lambda_n^2$ yields a decomposition of a single Trotter step $e^{-it\phi_x^2}$.

Proof. We first consider an arbitrary diagonal unitary $U = \sum_{n=0}^{d-1} e^{-i\beta_n} |n\rangle\langle n|$ and show how to implement it up to a global phase using adjacent $R_Z^{(k,k+1)}$ rotations; we then specialize to $\beta_n = t\lambda_n^2$. First, each adjacent Z -Givens rotation $R_Z^{(k,k+1)}(\theta_k)$ acts as

$$R_Z^{(k,k+1)}(\theta_k) = e^{-i\theta_k/2} |k\rangle\langle k| + e^{+i\theta_k/2} |k+1\rangle\langle k+1|, \quad (13)$$

where $k = 0, \dots, d-2$. All such rotations are diagonal and therefore commute, so on level $|n\rangle$ of the qudit the net phase is $e^{-i\alpha_n}$, where

$$\alpha_n = \frac{1}{2}\theta_n - \frac{1}{2}\theta_{n-1}, \quad \theta_{-1} = \theta_{d-1} = 0. \quad (14)$$

Since $\sum_{n=0}^{d-1} \alpha_n = \frac{1}{2}(-\theta_{-1} + \theta_{d-1}) = 0$, the product lies in $SU(d)$. The diagonal phases we wish to implement, $(\beta_n)_{n=0}^{d-1}$,

need not satisfy $\sum_n \beta_n = 0$, but they can be realized up to a global phase by writing $e^{-i\beta_n} = e^{-i\bar{\beta}} e^{-i(\beta_n - \bar{\beta})}$ with

$$\bar{\beta} = \frac{1}{d} \sum_{n=0}^{d-1} \beta_n. \quad (15)$$

In our application to $e^{-it\phi_x^2}$ we have $\beta_n = t\lambda_n^2$. Thus we choose the rotation angles as

$$\theta_k = 2 \sum_{n=0}^k (\beta_n - \bar{\beta}) \pmod{4\pi}, \quad (16)$$

so that $\alpha_n = \beta_n - \bar{\beta}$, i.e., the implemented phase equals $e^{-i(\beta_n - \bar{\beta})}$. Hence

$$\prod_{k=0}^{d-2} R_Z^{(k,k+1)}(\theta_k) \quad (17)$$

implements U up to a global phase, so one such decomposition with $d-1$ Givens rotations exists. Specializing to $\beta_n = t\lambda_n^2$ yields $e^{-it\phi_x^2}$ up to a global phase. Moreover, for any target phases $(\alpha_n)_{n=0}^{d-1}$ with $\sum_n \alpha_n = 0$, the linear system in Eq. (14) has the unique solution

$$\theta_k = 2 \sum_{n=0}^k \alpha_n. \quad (18)$$

Thus the angles (θ_k) that realize a given diagonal (up to global phase) are unique modulo 4π . ■

Note that under the symmetric uniform truncation with odd d [see Eq. (2)], the centered partial sums admit the closed form

$$\sum_{n=0}^k (\lambda_n^2 - \mu) = \phi_{\max}^2 \frac{4(k+1)}{3(d-1)^2} \left(k - \frac{d-2}{2} \right) [k - (d-1)], \quad (19)$$

where $\mu = \frac{1}{d} \sum_{n=0}^{d-1} \lambda_n^2$. Since $k \in \{0, \dots, d-2\}$ and $(d-2)/2$ is a half-integer for odd d , this expression never vanishes. Hence $\theta_k \not\equiv 0 \pmod{4\pi}$ for generic t , and therefore $\theta_k \neq 0$.

C. Asymptotic scaling comparison

In this section, we compare the non-Clifford cost of the qudit and qubit implementations. On the qubit side, we consider a fixed fault-tolerant qubit gate set and define $C_{R_z}^{(\text{qb})}(\delta)$ as the non-Clifford synthesis cost of a single-qubit R_z rotation to accuracy δ . On the qudit side, we likewise fix a (dimension- d) fault-tolerant qudit gate set and let $C_{\text{SU}(2)}^{(\text{qd})}(\delta)$ denote the non-Clifford synthesis cost to approximate a single embedded two-level one-parameter $\text{SU}(2)$ rotation to operator-norm error at most δ . As is standard in fault-tolerant settings, neither the qubit R_z rotations nor the qudit embedded $\text{SU}(2)$ rotations are treated as native primitives: in both encodings, these are continuous one-parameter targets that must be compiled into a finite discrete gate set, and our resource metric counts only the resulting non-Clifford gates.

Lemmas 1 and 2 reduce the specific Regime 1 constructions to circuits containing a known number of one-parameter rotations. This allows a gate-set-agnostic asymptotic comparison between the qubit implementation of Lemma 1 and

the qudit implementation obtained from Lemma 2 followed by independent synthesis of the resulting embedded two-level $\text{SU}(2)$ rotations under a uniform per-rotation error allocation. Rather than asserting an explicit $C_{\text{SU}(2)}^{(\text{qd})}$, we ask what relative synthesis performance of those qudit primitives would be required for the qudit construction to use fewer non-Cliffords than the qubit baseline at the same overall precision. Theorem 1 states the resulting requirement.

Theorem 1. Let $U = e^{-it\phi_x^2}$ be a single-site Trotter step on a d -level qudit, and let $n_b = \lceil \log_2 d \rceil$. Consider the qudit implementation obtained by Lemma 2, with each embedded two-level $\text{SU}(2)$ rotation synthesized independently under a uniform per-rotation error allocation. If this qudit decomposition is to asymptotically use fewer non-Cliffords than the qubit implementation of Lemma 1, then the non-Clifford synthesis cost of the required one-parameter $\text{SU}(2)$ rotations must be smaller than that of qubit-optimal R_z synthesis in Clifford + T by a factor

$$\Theta\left(\frac{(\log d)^2}{d}\right) = \Theta\left(\frac{n_b^2}{2^{n_b}}\right), \quad (20)$$

i.e., by an exponentially strong advantage in n_b .

Proof. By Lemma 2, $U = e^{-it\phi_x^2}$ uses $(d-1)$ nontrivial one-parameter embedded $\text{SU}(2)$ rotations; denote this count by

$$L_{\text{qd}}^{(1)}(d) = d - 1. \quad (21)$$

Assume each such $\text{SU}(2)$ rotation is approximated using a fixed discrete gate set, with synthesis cost $C_{\text{SU}(2)}^{(\text{qd})}(\delta)$ to achieve operator-norm error at most δ . Choosing

$$\delta = \frac{\varepsilon}{L_{\text{qd}}^{(1)}(d)} \quad (22)$$

ensures total error at most ε by a standard triangle-inequality bound, so the qudit non-Clifford count satisfies

$$N_{\text{qudit}}^{\text{PF}}(d, \varepsilon) \lesssim L_{\text{qd}}^{(1)}(d) C_{\text{SU}(2)}^{(\text{qd})}\left(\frac{\varepsilon}{L_{\text{qd}}^{(1)}(d)}\right). \quad (23)$$

By Lemma 1, U can be implemented using

$$L_{\text{qb}}^{(1)}(d) = \frac{n_b^2 + n_b}{2} \quad (24)$$

single-qubit R_z rotations on n_b qubits. Allocating per-rotation accuracy

$$\delta_{\text{qb}} = \frac{\varepsilon}{L_{\text{qb}}^{(1)}(d)}, \quad (25)$$

the corresponding qubit non-Clifford count is

$$N_{\text{qubit}}^{\text{PF}}(d, \varepsilon) = L_{\text{qb}}^{(1)}(d) C_{R_z}^{(\text{qb})}\left(\frac{\varepsilon}{L_{\text{qb}}^{(1)}(d)}\right). \quad (26)$$

Under the per-rotation accuracy model, a necessary condition for the qudit implementation of Lemma 2 to use fewer non-Clifford gates than the qubit implementation is

$$N_{\text{qudit}}^{\text{PF}}(d, \varepsilon) \lesssim N_{\text{qubit}}^{\text{PF}}(d, \varepsilon), \quad (27)$$

i.e.,

$$L_{\text{qd}}^{(1)}(d) C_{\text{SU}(2)}^{(\text{qd})} \left(\frac{\varepsilon}{L_{\text{qd}}^{(1)}(d)} \right) \lesssim L_{\text{qb}}^{(1)}(d) C_{R_z}^{(\text{qb})} \left(\frac{\varepsilon}{L_{\text{qb}}^{(1)}(d)} \right). \quad (28)$$

Rearranging yields

$$C_{\text{SU}(2)}^{(\text{qd})} \left(\frac{\varepsilon}{L_{\text{qd}}^{(1)}(d)} \right) \lesssim \frac{L_{\text{qb}}^{(1)}(d)}{L_{\text{qd}}^{(1)}(d)} C_{R_z}^{(\text{qb})} \left(\frac{\varepsilon}{L_{\text{qb}}^{(1)}(d)} \right). \quad (29)$$

Since $L_{\text{qb}}^{(1)}(d) = \Theta((\log d)^2)$, the prefactor satisfies

$$\frac{L_{\text{qb}}^{(1)}(d)}{d-1} = \Theta \left(\frac{(\log d)^2}{d} \right) = \Theta \left(\frac{n_b^2}{2^{n_b}} \right), \quad (30)$$

which is exponentially small in $n_b = \lceil \log_2 d \rceil$. Thus, for fixed ε , the Lemma 2 decomposition plus independent-synthesis route can asymptotically beat the qubit implementation only if embedded SU(2) rotations admit exponentially cheaper synthesis (as a function of n_b) than qubit R_z rotations. ■

For any fixed finite discrete gate set, approximating a continuous one-parameter subgroup to precision δ requires $C_{\text{SU}(2)}^{(\text{qd})}(\delta) = \Omega(\log(1/\delta))$ by standard volume-counting arguments [66]. Moreover, for common fault-tolerant gate sets admitting number-theoretic synthesis, known procedures achieve $C_{\text{prim}}(\delta) = O(\log(1/\delta))$ [47,64], so there is no asymptotic room for an exponentially improving dependence on $n_b = \lceil \log_2 d \rceil$. Consequently, we view the exponential per-rotation advantage demanded by Theorem 1 as incompatible with the standard fault-tolerant gate-synthesis landscape and therefore do not expect the qudit encoding to asymptotically use fewer non-Cliffords than the qubit construction.

D. Finite- d crossover analysis

The SU(2) decomposition still provides a meaningful comparison at the level of the number of continuous one-parameter rotations required. For example, in qutrit Clifford + R synthesis, Ref. [64] focuses on approximating diagonal two-level rotations, since these serve as the basic building blocks for SU(3) synthesis. Furthermore, as discussed above, for any fixed finite discrete gate set, approximating a continuous one-parameter subgroup to precision δ requires $\Omega(\log(1/\delta))$ non-Cliffords by standard volume-counting arguments [66]. Therefore, both qubit rotations and embedded-qudit SU(2) primitives admit logarithmic-in-precision synthesis, $\Omega(\log(1/\delta))$, and we can study the synthesis behavior required of the fault-tolerant qudit implementation as the truncation d grows.

Assume we wish to implement the onsite unitary to accuracy ε . Starting from Eq. (29), we model the cost of a fault-tolerant continuous one-parameter embedded SU(2) rotation as

$$C_{\text{SU}(2)}^{(\text{qd})} \approx a \log_2(1/\delta), \quad (31)$$

where $a > 0$ is a constant prefactor that captures gate-set and compilation overheads. For the qubit baseline, we use the single-qubit R_z synthesis model [67],

$$C_{R_z}^{(\text{qb})}(\delta) = 0.57 \log_2(1/\delta) + 8.83. \quad (32)$$

Substituting $\delta = \varepsilon/L_{\text{qd}}^{(1)}(d)$ for the qudit decomposition and $\delta = \varepsilon/L_{\text{qb}}^{(1)}(d)$ for the qubit construction yields the approximate break-even condition

$$\begin{aligned} L_{\text{qd}}^{(1)}(d) a \log_2 \left(\frac{L_{\text{qd}}^{(1)}(d)}{\varepsilon} \right) \\ \lesssim L_{\text{qb}}^{(1)}(d) \left[0.57 \log_2 \left(\frac{L_{\text{qb}}^{(1)}(d)}{\varepsilon} \right) + 8.83 \right]. \end{aligned} \quad (33)$$

Rearranging gives an explicit requirement on the qudit synthesis prefactor,

$$a \lesssim a_{\text{max}}^{\text{PF}}(d, \varepsilon) = \frac{L_{\text{qb}}^{(1)}(d) [0.57 \log_2(L_{\text{qb}}^{(1)}(d)/\varepsilon) + 8.83]}{L_{\text{qd}}^{(1)}(d) \log_2(L_{\text{qd}}^{(1)}(d)/\varepsilon)}. \quad (34)$$

For $\varepsilon = 10^{-6}$ (a representative value), and restricting to prime local dimensions in the fixed-encoding comparison, this yields

$$\begin{aligned} a_{\text{max}}^{\text{PF}}(3, 10^{-6}) &\approx 1.51, \\ a_{\text{max}}^{\text{PF}}(5, 10^{-6}) &\approx 1.48, \\ a_{\text{max}}^{\text{PF}}(7, 10^{-6}) &\approx 0.96. \end{aligned} \quad (35)$$

The quantity $a_{\text{max}}^{\text{PF}}(d, \varepsilon)$ should be read as a compiler target: if an embedded two-level qudit SU(2) rotation can be synthesized with cost $a < a_{\text{max}}^{\text{PF}}(d, \varepsilon)$, then the qudit implementation uses fewer non-Clifford gates than the qubit baseline.

It is also useful to compare this compiler target with the effective prefactor that would reproduce the qubit R_z synthesis cost at the same primitive precision required of the qudit rotations,

$$a_{R_z}^{\text{PF}}(d, \varepsilon) = \frac{0.57 \log_2(L_{\text{qd}}^{(1)}(d)/\varepsilon) + 8.83}{\log_2(L_{\text{qd}}^{(1)}(d)/\varepsilon)}. \quad (36)$$

When $a_{\text{max}}^{\text{PF}}(d, \varepsilon) > a_{R_z}^{\text{PF}}(d, \varepsilon)$, the qudit construction can tolerate embedded-SU(2) synthesis that is worse than qubit R_z synthesis at the same primitive precision. When $a_{\text{max}}^{\text{PF}}(d, \varepsilon) < a_{R_z}^{\text{PF}}(d, \varepsilon)$, a genuine per-primitive qudit synthesis advantage is required.

For $\varepsilon = 10^{-6}$, we find $a_{\text{max}}^{\text{PF}}(d, \varepsilon) > a_{R_z}^{\text{PF}}(d, \varepsilon)$ only for $d = 3$ and $d = 5$ among the computed prime dimensions. Determining whether these compiler targets are attainable for $d \geq 3$ requires explicit synthesis estimates for the relevant embedded primitives.

IV. REGIME 2: TIME EVOLUTION VIA LINEAR COMBINATION OF UNITARIES (LCU)

In the block-encoding LCU framework, the motivation for using qudits is stronger than in the continuous-rotation setting: when we express the same diagonal operator as a linear combination of unitaries, some terms that are non-Clifford on qubits can become Clifford on qudits. Prominent examples include the d -level incrementer and the quantum Fourier transform, both of which are Clifford operations on a single qudit [33]. Our goal is to determine whether, at fixed physical precision, increasing the local dimension (i.e., using a d -level

qudit instead of $\lceil \log_2 d \rceil$ qubits) can yield a constant-factor or asymptotic reduction in non-Clifford resources.

We consider two comparisons within Regime 2: a fully fixed-encoding qudit-versus-qubit comparison, and a secondary idealized hybrid code-switching comparison. In both comparisons, the qubit baseline is the standard nearest-register signed-binary projector-LCU construction of Refs. [59,60] on $n_b = \lceil \log_2 d \rceil$ qubits, which implements the signed-binary extension of the rounded squared operator that agrees with the physical $d = 2M + 1$ target on $\ell \in \{-M, \dots, M\}$. The main distinction between the two comparisons is that the fixed-encoding model asks what native fault-tolerant qudit synthesis performance would be required for an advantage, while the hybrid model obtains an absolute qubit Clifford + T comparison by allowing non-Clifford subroutines to be implemented after idealized qubit-qudit code-switching. The finite- d Regime 2 results below are therefore construction-specific comparisons between the native qudit LCU construction and this standard qubit baseline, and not optimal lower bounds over all possible qubit encodings of the truncation.

In the fixed-encoding comparison, the qudit implementation remains entirely in the qudit encoding, and we derive the corresponding break-even condition $a \lesssim a_{\max}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}})$. In the hybrid comparison, an $n_b = \lceil \log_2 d \rceil$ -qubit index register is allowed to code-switch to a single logical d -level qudit and back, giving an idealized absolute T -count comparison. Here, code-switching is modeled as an isometry from the valid index subspace to $\mathbb{C}^d$, followed by the inverse isometry after the qudit operation. The invalid computational states $|d\rangle, \dots, |2^{n_b} - 1\rangle$ are assumed to be unpopulated in the ideal circuit; the physical cost of implementing the isometry is not included, except through the overhead-budget analysis in Sec. IV D 3. For the purposes of this section, we focus on the hybrid qudit encoding, with details of the fixed-encoding implementation given in Appendix B.

Given a Hamiltonian H_i expressed as a linear combination of unitaries,

$$H_i = \sum_{j=0}^{M_i-1} h_{ij} U_{ij}, \quad \lambda_i = \sum_{j=0}^{M_i-1} |h_{ij}|, \quad (37)$$

one can define a preparation unitary PREP_i and a selection unitary SELECT_i by

$$\text{PREP}_i : |0\rangle \mapsto \sum_{j=0}^{M_i-1} \sqrt{\frac{|h_{ij}|}{\lambda_i}} |j\rangle, \quad (38)$$

$$\text{SELECT}_i = \sum_{j=0}^{M_i-1} |j\rangle\langle j| \otimes U_{ij}. \quad (39)$$

Here the first register (prepared by PREP_i) is the term-selection register, which we call the index register; SELECT_i acts on both this register and on the system register. Projecting the index register onto $|0\rangle$ then yields a block-encoding of H_i [13,68]:

$$\langle 0 | \text{PREP}_i^\dagger \text{SELECT}_i \text{PREP}_i | 0 \rangle = \frac{H_i}{\lambda_i}. \quad (40)$$

This block-encoding can then be used as input to standard Hamiltonian-simulation primitives (e.g., qubitization or quantum signal processing and its generalization via quantum singular value transformation) to implement real-time evolution $e^{-iH_i t}$, and more general functions of H_i , with polylogarithmic dependence on target precision. For a fault-tolerant resource comparison in Regime 2, we focus on the dominant oracle costs of PREP_i and SELECT_i , since these set the query complexity and, in our constructions, drive the leading non-Clifford gate counts. For background on block-encodings and their use in Hamiltonian simulation, see Refs. [13,15,16].

A. The qubit baseline

The qubit baseline studied here is the standard nearest-register signed-binary projector-LCU construction of Refs. [59,60]. The physical target is the same $d = 2M + 1$ field-amplitude truncation as in the qudit construction; the qubit baseline implements the nearest signed-binary extension of the rounded squared operator on

$$n_b = \lceil \log_2 d \rceil \quad (41)$$

qubits. Starting from the uniform field operator in Eq. (2), we define the centered field label

$$\ell = n - M. \quad (42)$$

We label a computational basis state of the n_b -qubit register by a sign bit $\ell_{\text{sign}} \in \{0, 1\}$ and magnitude bits $\ell_{n_b-2}, \dots, \ell_0 \in \{0, 1\}$,

$$|\ell\rangle = |\ell_{\text{sign}}, \ell_{n_b-2}, \dots, \ell_0\rangle, \quad (43)$$

and assign the integer field label

$$\ell = (-1)^{\ell_{\text{sign}}} \sum_{r=0}^{n_b-2} 2^r \ell_r, \quad |\ell| = \sum_{r=0}^{n_b-2} 2^r \ell_r, \quad (44)$$

so that

$$\ell \in \{-(2^{n_b-1} - 1), \dots, 2^{n_b-1} - 1\}. \quad (45)$$

Note that the two computational strings with all magnitude bits zero (sign bit $\ell_{\text{sign}} = 0$ or $\ell_{\text{sign}} = 1$) are both assigned the field label $\ell = 0$. The redundant negative-zero string is therefore not excluded; it is a second computational basis representative of $\ell = 0$. Because the squared operator below depends only on the magnitude bits (ℓ_r), the sign bit is irrelevant whenever the magnitude is zero, and the negative-zero string is an invariant zero-eigenvalue state of the squared operator.

Rather than restricting to the physical interval $\{-M, \dots, M\}$ and enforcing a physical-subspace projector, we work directly with the full n_b -qubit register. With a slight abuse of notation, we denote by ϕ_x^2 the corresponding rounded diagonal operator defined on every computational basis state of the register,

$$\phi_x^2 = \delta_\phi^2 \sum_{\ell_{\text{sign}}, \ell_{n_b-2}, \dots, \ell_0 \in \{0, 1\}} \ell^2 |\ell\rangle\langle\ell|, \quad (46)$$

where ℓ on the right-hand side is the integer label assigned above. By construction, this operator agrees with the

physical target on the physical interval $\ell \in \{-M, \dots, M\}$ and is the standard signed-binary extension to the full register used throughout Refs. [59,60]. Equivalently, indexing by the integer label,

$$\phi_x^2 = \delta_\phi^2 \sum_{\ell=-(2^{n_b-1}-1)}^{2^{n_b-1}-1} \ell^2 |\ell\rangle\langle\ell|, \quad (47)$$

with the understanding that $\ell = 0$ is realized by both sign-bit values. Since the sign bit drops out upon squaring, we have

$$\ell^2 = \left(\sum_{r=0}^{n_b-2} 2^r \ell_r \right)^2 = \sum_{r,s=0}^{n_b-2} 2^{r+s} \ell_r \ell_s. \quad (48)$$

Accordingly, for each pair (r, s) we define the diagonal bit-projector on the full n_b -qubit register,

$$P_{r,s} = \sum_{\ell_{\text{sign}}, \ell_{n_b-2}, \dots, \ell_0 \in \{0,1\}} \ell_r \ell_s |\ell\rangle\langle\ell|, \quad (49)$$

which depends only on the magnitude bits and is therefore insensitive to the negative-zero redundancy. Because each bit ℓ_r is either 0 or 1, the product $\ell_r \ell_s$ is equal to 1 exactly when the r th and s th nonsign bits are both equal to 1, and is otherwise 0. Hence,

$$P_{r,s} |\ell\rangle = \begin{cases} |\ell\rangle, & \ell_r = \ell_s = 1 \\ 0, & \text{otherwise,} \end{cases} \quad (50)$$

so $P_{r,s}$ is the projector onto the span of basis states whose selected nonsign bits are both equal to 1. Substituting the bit expansion of ℓ^2 into the diagonal operator gives

$$\phi_x^2 = \delta_\phi^2 \sum_{r,s=0}^{n_b-2} 2^{r+s} P_{r,s}. \quad (51)$$

Since the operators $P_{r,s}$ are projectors rather than unitaries, we give a block encoding of ϕ_x^2 using the projector-based LCU methods of Refs. [59,60]. The PREP unitary acts on the index and ancilla registers and does not move the system register out of any computational-basis subspace, while the SELECT unitary is diagonal in the signed-binary computational basis [it conditions on the index pair (r, s) and applies a bit-test phase on the system]; both are therefore well-defined on every computational basis state of the n_b -qubit Hilbert space.

Lemma 3. The operator ϕ_x^2 admits a block encoding with normalization factor

$$\alpha = \delta_\phi^2 (2^{n_b-1} - 1)^2 \quad (52)$$

on an n_b -qubit system register, using two index registers of $\lceil \log_2(n_b - 1) \rceil$ qubits each and one additional ancilla qubit. Equivalently, there exist unitaries PREP and SELECT such that

$$(|0\rangle \otimes I) \text{PREP}^\dagger \text{SELECT PREP} (|0\rangle \otimes I) = \frac{\phi_x^2}{\alpha}, \quad (53)$$

where $|0\rangle$ denotes the joint initial state of the two index registers and the ancilla qubit. Moreover, the T count for one block-encoding call is

$$T_{\phi_x^2}^{\text{LCU}}(\varepsilon) = 32b_r + 24n_b - 116, \quad (54)$$

where

$$b_r = \left\lceil \frac{1}{2} \log_2 \left(\frac{9\pi^2}{2\varepsilon} \right) \right\rceil. \quad (55)$$

Proof. Following Refs. [59,60], we implement the projectors $P_{r,s}$ within a block encoding by introducing a single ancilla qubit b prepared in the state

$$|+\rangle_b = \frac{|0\rangle_b + |1\rangle_b}{\sqrt{2}}, \quad (56)$$

together with the weighted index-register state

$$|\psi\rangle_{rs} = \frac{1}{\sqrt{\Lambda}} \sum_{r,s=0}^{n_b-2} 2^{(r+s)/2} |r, s\rangle, \quad (57)$$

where

$$\Lambda = \sum_{r,s=0}^{n_b-2} 2^{r+s}. \quad (58)$$

Let PREP be the unitary defined by

$$\text{PREP} |0\rangle_{rs} |0\rangle_b = |\psi\rangle_{rs} |+\rangle_b. \quad (59)$$

For each pair (r, s) , define a unitary $\text{SELECT}_{r,s}$ acting on the ancilla and signed-binary field register by

$$\text{SELECT}_{r,s} |b\rangle |\ell\rangle = (-1)^{b(\ell_r \ell_s \oplus 1)} |b\rangle |\ell\rangle. \quad (60)$$

Equivalently, conditioned on the chosen pair (r, s) , the circuit queries the r th and s th nonsign bits of the signed-binary label and applies a Z to the ancilla unless both selected bits are equal to 1. Therefore, for every basis state $|\ell\rangle$,

$$(|+\rangle_b \otimes I) \text{SELECT}_{r,s} (|+\rangle_b \otimes |\ell\rangle) = \ell_r \ell_s |\ell\rangle, \quad (61)$$

and hence

$$(|+\rangle_b \otimes I) \text{SELECT}_{r,s} (|+\rangle_b \otimes I) = P_{r,s}. \quad (62)$$

Thus the desired projector appears as the $|+\rangle_b$ -to- $|+\rangle_b$ matrix element of a unitary.

The normalization is

$$\Lambda = \sum_{r,s=0}^{n_b-2} 2^{r+s} = \left(\sum_{r=0}^{n_b-2} 2^r \right)^2 = (2^{n_b-1} - 1)^2. \quad (63)$$

Next define the full selection unitary

$$\text{SELECT} = \sum_{r,s=0}^{n_b-2} |r, s\rangle\langle r, s| \otimes \text{SELECT}_{r,s}, \quad (64)$$

which applies the corresponding bit-test unitary conditioned on the index pair (r, s) . Sandwiching this operator between the weighted index state and the ancilla state gives

$$(|\psi\rangle_{rs} \langle +|_b \otimes I) \text{SELECT} (|\psi\rangle_{rs} |+\rangle_b \otimes I) = \frac{1}{\Lambda} \sum_{r,s=0}^{n_b-2} 2^{r+s} P_{r,s}. \quad (65)$$

Using the projector decomposition above,

$$\phi_x^2 = \delta_\phi^2 \sum_{r,s=0}^{n_b-2} 2^{r+s} P_{r,s}, \quad (66)$$

we obtain

$$\langle \psi |_{rs} \langle + |_b \otimes I \rangle \text{SELECT} (| \psi \rangle_{rs} | + \rangle_b \otimes I) = \frac{\phi_x^2}{\delta_\phi^2 \Lambda}. \quad (67)$$

Therefore, this construction gives an exact qubit block encoding of ϕ_x^2 with normalization factor

$$\alpha = \delta_\phi^2 \Lambda = \delta_\phi^2 (2^{n_b-1} - 1)^2. \quad (68)$$

Equivalently,

$$(|0\rangle_{rs} \langle 0|_b \otimes I) \text{PREP}^\dagger \text{SELECTPREP} (|0\rangle_{rs} |0\rangle_b \otimes I) = \frac{\phi_x^2}{\alpha}, \quad (69)$$

which is the claimed exact block encoding at the oracle level.

For the fault-tolerant cost, we specialize the kinetic-term construction of Refs. [59,60] to a single signed-binary field register. One PREP oracle for the weighted bit-pair state has Toffoli count

$$T_{\text{prep}}^{(\text{Toff})} = 4b_r + 2n_b - 16, \quad b_r = \left\lceil \frac{1}{2} \log_2 \left(\frac{9\pi^2}{2\varepsilon} \right) \right\rceil, \quad (70)$$

and $\text{PREP}^\dagger$ has the same cost. The selector consists only of the two bit queries and the final controlled phase, giving

$$T_{\text{sel}}^{(\text{Toff})} = 2(n_b - 1), \quad T_{\text{sel}}^{(T)} = 20. \quad (71)$$

Using the convention 1 Toffoli = $4T$ [69], the full block-encoding call

$$W = (\text{PREP}^\dagger \otimes I) \text{SELECT} (\text{PREP} \otimes I) \quad (72)$$

therefore has total T count of

$$T_{\phi_x^2}^{\text{LCU}}(\varepsilon) = 4(T_{\text{prep}}^{(\text{Toff})} + T_{\text{prep}}^{(\text{Toff})} + T_{\text{sel}}^{(\text{Toff})}) + T_{\text{sel}}^{(T)} \quad (73)$$

$$= 32b_r + 24n_b - 116. \quad (74)$$

■

B. The qudit implementation

For the qudit case, we first define the generalized Pauli Z on a d -level qudit [33]

$$Z_d = \sum_{j=0}^{d-1} \omega^j |j\rangle \langle j|, \quad (75)$$

where $\omega = e^{2\pi i/d}$, which belongs to the qudit Clifford group.

1. Oracle definitions

Lemma 4. The squared field operator ϕ_x^2 admits a unique expansion in the generalized Pauli basis as

$$\phi_x^2 = \sum_{r=0}^{d-1} \beta_r Z_d^r, \quad (76)$$

with

$$\begin{aligned} \beta_0 &= \phi_{\text{max}}^2 \frac{d+1}{3(d-1)}, \\ \beta_r &= \frac{2\phi_{\text{max}}^2}{(d-1)^2} e^{i\pi r/d} \frac{\cos(\pi r/d)}{\sin^2(\pi r/d)}, \quad r = 1, \dots, d-1. \end{aligned} \quad (77)$$

This decomposition is irreducible in the sense that $\beta_r \neq 0$ for all $r \in \{1, \dots, d-1\}$, so ϕ_x^2 has full support on the nontrivial generalized Pauli operators $\{Z_d^r\}_{r=1}^{d-1}$.

Proof. The set $\{Z_d^r\}_{r=0}^{d-1}$ forms a complete orthogonal basis for the d -dimensional vector space of diagonal operators on $\mathbb{C}^d$ [70]. Consequently, any diagonal operator admits a unique expansion

$$\phi_x^2 = \sum_{r=0}^{d-1} \beta_r Z_d^r, \quad (78)$$

with eigenvalues related by

$$\lambda_n^2 = \sum_{r=0}^{d-1} \beta_r \omega^{rn}, \quad (79)$$

i.e., the eigenvalue sequence $(\lambda_n^2)_{n=0}^{d-1}$ is the inverse discrete Fourier transform of $(\beta_r)_{r=0}^{d-1}$. Therefore, the coefficients β_r are given by the discrete Fourier transform

$$\beta_r = \frac{1}{d} \sum_{n=0}^{d-1} \lambda_n^2 \omega^{-rn}. \quad (80)$$

For the uniformly spaced grid $\lambda_n = -\phi_{\text{max}} + \frac{2\phi_{\text{max}}}{d-1}n$, a direct evaluation of this sum yields

$$\beta_0 = \phi_{\text{max}}^2 \frac{d+1}{3(d-1)}, \quad (81)$$

and, for $r \neq 0$,

$$\beta_r = \frac{2\phi_{\text{max}}^2}{(d-1)^2} e^{i\pi r/d} \frac{\cos(\pi r/d)}{\sin^2(\pi r/d)}. \quad (82)$$

Since $\lambda_n^2 \in \mathbb{R}$ for all n , the Fourier coefficients obey $\beta_{d-r} = \beta_r^*$, which ensures Hermiticity of ϕ_x^2 .

Finally, for d odd and $r \in \{1, \dots, d-1\}$ we have $\sin(\pi r/d) \neq 0$ and $\cos(\pi r/d) \neq 0$, so the real prefactor is nonzero. Hence $\beta_r \neq 0$ for all $r \neq 0$, implying that ϕ_x^2 has full support on $\{Z_d^r\}_{r=1}^{d-1}$ and cannot be represented with fewer nontrivial generalized Pauli terms. ■

Note that both the “full support” and the associated minimality statements are understood relative to the generalized-Pauli basis expansion used here, rather than claiming optimality over all possible LCU decompositions or over alternative operator bases, which could in principle yield different representations.

Definition 1. The LCU oracles for ϕ_x^2 are defined as

$$\text{PREP}_{\phi_x^2} : |0\rangle \mapsto \sum_{r=1}^{d-1} \sqrt{\frac{|\beta_r|}{\Lambda}} |r\rangle, \quad (83)$$

with normalization constant $\Lambda = \sum_{r=1}^{d-1} |\beta_r|$, and

$$\text{SELECT}_{\phi_x^2} = |0\rangle \langle 0| \otimes I + \left(\sum_{r=1}^{d-1} e^{i\theta_r} |r\rangle \langle r| \otimes Z_d^r \right). \quad (84)$$

Proof. Note that the coefficients β_r are in general complex. Following the standard LCU convention (see, e.g., Ref. [71]), we absorb the phase of each coefficient into the corresponding unitary and work with non-negative real weights. Therefore,

we define

$$\tilde{U}_r = e^{i\theta_r} Z_d^r, \quad \theta_r = \arg \beta_r. \quad (85)$$

Note that the operator can be decomposed as $\phi_x^2 = \beta_0 \mathbb{I} + \phi_x^{2'}$. Thus it is sufficient to implement $e^{-i\phi_x^{2'}}$, since $(e^{-i\phi_x^{2'}})$ is equal to $(e^{-i\phi_x^2})$ up to a global phase [68]. With some abuse of notation, we drop the identity term but continue to label the shifted operator by (ϕ_x^2) , i.e.,

$$\phi_x^2 = \sum_{r=1}^{d-1} |\beta_r| \tilde{U}_r. \quad (86)$$

Thus we take an index register over $r = 0, \dots, d-1$, i.e., a single d -level qudit (or $\lceil \log_2(d) \rceil$ qubits). Using the standard LCU convention and defining $\Lambda = \sum_{r=1}^{d-1} |\beta_r|$, we obtain

$$\text{PREP}_{\phi_x^2} : |0\rangle \mapsto \sum_{r=1}^{d-1} \sqrt{\frac{|\beta_r|}{\Lambda}} |r\rangle \quad (87)$$

and

$$\begin{aligned} \text{SELECT}_{\phi_x^2} &= |0\rangle\langle 0| \otimes I + \sum_{r=1}^{d-1} |r\rangle\langle r| \otimes \tilde{U}_r \\ &= |0\rangle\langle 0| \otimes I + \left(\sum_{r=1}^{d-1} e^{i\theta_r} |r\rangle\langle r| \otimes Z_d^r \right), \end{aligned} \quad (88)$$

where the diagonal term acts on the index register. ■

2. The SELECT Oracle

With the index and system registers encoded as d -level qudits, the entangling part of SELECT (84),

$$\sum_{r=0}^{d-1} |r\rangle\langle r| \otimes Z_d^r, \quad (89)$$

is the (Clifford) qudit generalized controlled-Z gate [35,72]. The only non-Clifford cost in SELECT therefore arises from the diagonal phases that accompany the LCU coefficients, i.e.,

$$D = |0\rangle\langle 0| + \sum_{r=1}^{d-1} e^{i\theta_r} |r\rangle\langle r|. \quad (90)$$

For the SELECT cost analysis below, we adopt the idealized hybrid model in which code-switching between the d -level index qudit and an n_b -qubit index register is treated as free. The overhead of this conversion is discussed in Sec. IV D 3.

Corollary 1. The diagonal phase operator

$$D = \sum_{r=0}^{d-1} e^{i\theta_r} |r\rangle\langle r| \quad (91)$$

can be implemented using $4n_b$ T gates, $n_b R_z(\theta)$ rotations, and $n_b + 1$ ancillas on an $n_b = \lceil \log_2 d \rceil$ qubit index register.

Proof. In Appendix A, we show that the Fourier coefficients of ϕ_x^2 admit the decomposition $\beta_r = c_r e^{i\pi r/d}$ with $c_r \in \mathbb{R}$, so that

$$D = \sum_{r=0}^{d-1} \text{sgn}(c_r) e^{i\pi r/d} |r\rangle\langle r| = D_{\text{sign}} D_{\text{clock}}. \quad (92)$$

Here $D_{\text{clock}} = \exp(i\frac{\pi}{d}N)$, where N is the number operator on the index register, and D_{sign} applies a ± 1 phase depending on r . Using a binary expansion of r and a standard reversible comparator circuit, we show that D_{clock} can be implemented with $n_b R_z(\theta)$ rotations (Lemma A2), and D_{sign} can be implemented with $4n_b$ T gates and $n_b + 1$ ancillas (Lemma A3). ■

3. The PREP Oracle

Lemma 5. The PREP oracle for the qudit LCU decomposition,

$$\text{PREP}_{\phi_x^2} : |0\rangle \mapsto \sum_{r=1}^{d-1} \sqrt{\frac{|\beta_r|}{\Lambda}} |r\rangle, \quad (93)$$

requires at most $2^{n_b} - 1$ single-qubit R_z rotations, up to Cliffords, when the index is encoded in an $n_b = \lceil \log_2 d \rceil$ qubit register.

Proof. By Lemma 4, ϕ_x^2 has full support on the nontrivial generalized Pauli operators $\{Z_d^r\}_{r=1}^{d-1}$. Thus the coefficients $|\beta_r|/\Lambda$ define a generic $(d-1)$ -component state on the subspace spanned by $\{|r\rangle\}_{r=1}^{d-1}$. Viewing the index register as an $n_b = \lceil \log_2 d \rceil$ -qubit system, we thus upper-bound $\text{PREP}_{\phi_x^2}$ as a generic real state-preparation task on that subspace.

A standard state-preparation construction uses uniformly controlled rotations [73]. In the circuit decomposition for state preparation, the “phase equalization” step is performed by a cascade of uniformly controlled z -rotations Ξ_z , and the “amplitude shaping” step is performed by a cascade of uniformly controlled y -rotations Ξ_y . This two-stage structure is given explicitly as a product over qubit levels in Eq. (7) of [73]. Because the coefficients $\frac{|\beta_r|}{\Lambda}$ are real (hence no relative phases are required), the phase-equalization stage is trivial: $\Xi_z = I$. Thus it suffices to implement $\Xi_y^\dagger$, i.e., a cascade of uniformly controlled R_y rotations.

Fix a target qubit $j \in \{1, \dots, n_b\}$. The corresponding gate in the cascade is a $(j-1)$ -fold uniformly controlled y -rotation $F_j^{j-1}(y, \cdot)$ (in the notation of Ref. [73]). Note that a k -fold uniformly controlled one-parameter rotation decomposes into exactly 2^k one-qubit elementary rotations [plus $O(k)$ CNOTs] [74]. Therefore, the total number of one-qubit rotations appearing in the amplitude cascade is

$$\sum_{j=1}^{n_b} 2^{j-1} = 2^{n_b} - 1. \quad (94)$$

Finally, each single-qubit $R_y(\theta)$ can be implemented as one $R_z(\theta)$ conjugated by Clifford gates, so the non-Clifford cost is $2^{n_b} - 1$ single-qubit R_z rotations. ■

C. Asymptotic scaling comparison

We first analyze the hybrid code-switching model in a common qubit R_z -rotation and T -count metric, which gives a clean asymptotic comparison; the fully qudit fixed-encoding analysis is deferred to Sec. IV D 2. This choice does not affect the asymptotic conclusion: the fixed-encoding block encoding still requires $\Theta(d)$ embedded two-level $\text{SU}(2)$ rotations per call (Appendix B), matching the $\Theta(d)$ rotation count in the code-switching construction.

Under the code-switching assumption, our non-Clifford metric is explicitly qubit T gates. For a single application of PREP and SELECT, via Corollary 1 and Lemma 5, the qudit implementation has $(2^{n_b} - 1) + n_b$ single-qubit R_z gates, and

$$N_{\text{nc}}^{(\text{qd})}(d) = 4n_b \quad (95)$$

T gates. By contrast, for the qubit implementation in Sec. IV A, the full Regime 2 per-call block-encoding cost is

$$N_{\text{qubit}}^{\text{LCU}}(d, \varepsilon) = 32 b_r(\varepsilon) + 24n_b - 116, \quad (96)$$

where

$$b_r(\varepsilon) = \left\lceil \frac{1}{2} \log_2 \left(\frac{9\pi^2}{2\varepsilon} \right) \right\rceil. \quad (97)$$

Thus, in the qubit Regime 2 baseline there is no synthesized R_z contribution.

A single block-encoding call corresponds to $W = (\text{PREP}^\dagger \otimes I) \text{SELECT} (\text{PREP} \otimes I)$, so the total number of synthesized R_z gates per qudit call is

$$L_{R_z, \text{qd}}^{(2)}(d) = 2(2^{n_b} - 1) + n_b. \quad (98)$$

Allocating the full-target precision ε uniformly across these synthesized qudit-side R_z gates gives

$$\delta^{(\text{qd})}(d, \varepsilon) = \frac{\varepsilon}{L_{R_z, \text{qd}}^{(2)}(d)}. \quad (99)$$

Using the single-qubit synthesis model

$$C_{R_z}^{(\text{qb})}(\delta) = 0.57 \log_2 \left(\frac{1}{\delta} \right) + 8.83, \quad (100)$$

the total Regime 2 non-Clifford cost models become

$$N_{\text{qubit}}^{\text{LCU}}(d, \varepsilon) = 32 b_r(\varepsilon) + 24n_b - 116, \quad (101)$$

and

$$N_{\text{qudit}}^{\text{LCU}}(d, \varepsilon) = L_{R_z, \text{qd}}^{(2)}(d) \left[0.57 \log_2 \left(\frac{1}{\delta^{(\text{qd})}(d, \varepsilon)} \right) + 8.83 \right] + N_{\text{nc}}^{(\text{qd})}(d). \quad (102)$$

Theorem 2. For fixed ε , the cost models above imply that the qudit construction uses asymptotically more T gates, i.e.,

$$\frac{N_{\text{qubit}}^{\text{LCU}}(d, \varepsilon)}{N_{\text{qudit}}^{\text{LCU}}(d, \varepsilon)} \rightarrow 0 \quad (d \rightarrow \infty). \quad (103)$$

Proof. We have $N_{\text{qubit}}^{\text{LCU}}(d, \varepsilon) = \Theta(n_b)$ for fixed ε , since $b_r(\varepsilon)$ is independent of d and the qubit per-call cost is $32 b_r(\varepsilon) + 24n_b - 116$. By contrast, $L_{R_z, \text{qd}}^{(2)}(d) = \Theta(2^{n_b})$, and hence

$$N_{\text{qudit}}^{\text{LCU}}(d, \varepsilon) = \Theta(2^{n_b} n_b) \quad (104)$$

for fixed ε because $\log_2(1/\delta^{(\text{qd})}) = \log_2(L_{R_z, \text{qd}}^{(2)}/\varepsilon) = \Theta(n_b)$. Therefore, $N_{\text{qubit}}^{\text{LCU}}(d, \varepsilon)/N_{\text{qudit}}^{\text{LCU}}(d, \varepsilon) \rightarrow 0$, as claimed. ■

Note that in the $d \rightarrow \infty$ limit at fixed ϕ_{max} and fixed $t, \varepsilon_{\text{sim}}$, both LCU normalizations $\alpha_{\text{qb}}(d)$ and $\alpha_{\text{qd}}(d)$ remain $O(1)$ in d (the qudit normalization converges to a finite constant; the qubit normalization is bounded by $4\phi_{\text{max}}^2$ but does not approach a single limit because of the dependence on $n_b = \lceil \log_2 d \rceil$). Hence the query count $Q = O(\alpha t + \log(1/\varepsilon_{\text{sim}}))$ is $O(1)$ in d for both encodings. Consequently, incorporating the

query complexity as detailed in Sec. IV D 1 does not change the d -asymptotic separation established in Theorem 2: the asymptotic ordering is determined entirely by the per-call non-Clifford cost growth.

D. Finite- d crossover analysis

1. Required number of queries

Given an LCU decomposition

$$\hat{H} = \sum_{j=1}^L \alpha_j \hat{U}_j, \quad \alpha = \sum_{j=1}^L |\alpha_j|, \quad (105)$$

qubitization-based simulation of e^{-iHt} to error ε_{sim} uses

$$Q(\alpha; t, \varepsilon_{\text{sim}}) = O(\alpha t + \log(1/\varepsilon_{\text{sim}})) \quad (106)$$

queries to a block-encoding of $\hat{H}$ [15], where $\alpha = \sum_j |\alpha_j|$ is the LCU coefficient 1-norm. A single query consists of one application of the block-encoding unitary

$$W = (\text{PREP}^\dagger \otimes I) \text{SELECT} (\text{PREP} \otimes I). \quad (107)$$

Thus the end-to-end non-Clifford cost is obtained by multiplying the per-call block-encoding cost by the corresponding query count.

We denote by $\alpha_{\text{qb}}(d)$ and $\alpha_{\text{qd}}(d)$ the LCU normalizations entering the qubit and qudit query counts, respectively. For the qubit projector construction of Sec. IV A, Lemma 3 gives

$$\alpha_{\text{qb}}(d) = (\delta_\phi)^2 (2^{n_b-1} - 1)^2. \quad (108)$$

For the qudit generalized-Pauli expansion in Definition 1, we discard the identity contribution $\beta_0 I$, since it contributes only a global phase to the time evolution. The corresponding shifted-operator normalization is therefore

$$\alpha_{\text{qd}}(d) \equiv \Lambda = \sum_{r=1}^{d-1} |\beta_r|, \quad (109)$$

with β_r as in Eq. (82).

Following Ref. [75], we define the encoding-dependent query counts

$$\begin{aligned} Q_{\text{qb}}(d; t, \varepsilon_{\text{sim}}) &= Q(\alpha_{\text{qb}}(d); t, \varepsilon_{\text{sim}}), \\ Q_{\text{qd}}(d; t, \varepsilon_{\text{sim}}) &= Q(\alpha_{\text{qd}}(d); t, \varepsilon_{\text{sim}}). \end{aligned} \quad (110)$$

The qubit end-to-end baseline is therefore

$$T_{\text{tot}}^{\text{qb}}(d) = Q_{\text{qb}}(d; t, \varepsilon_{\text{sim}}) N_{\text{qubit}}^{\text{LCU}}(d, \varepsilon_{\text{BE}}^{\text{qb}}(d)), \quad (111)$$

where $N_{\text{qubit}}^{\text{LCU}}$ is the non-Clifford cost of one qubit block-encoding call. For any specified qudit implementation with per-call cost $N_{\text{qudit}}^{\text{LCU}}$, the analogous total cost is

$$T_{\text{tot}}^{\text{qd}}(d) = Q_{\text{qd}}(d; t, \varepsilon_{\text{sim}}) N_{\text{qudit}}^{\text{LCU}}(d, \varepsilon_{\text{BE}}^{\text{qd}}(d)). \quad (112)$$

In the fixed-encoding comparison below, Eq. (111) gives the qubit baseline against which the qudit synthesis threshold is defined. In the code-switching comparison, Eqs. (111) and (112) are used directly to form the total-cost ratio.

The per-call block-encoding accuracy is coupled to the target simulation error using a per-query error budget. By a standard telescoping bound in operator norm, Q queries each

accurate to ε_{BE} contribute at most $Q \varepsilon_{\text{BE}}$ additional error. It is therefore sufficient to take

$$\varepsilon_{\text{BE}}^{\text{qb}}(d) = \frac{\varepsilon_{\text{sim}}}{Q_{\text{qb}}(d; t, \varepsilon_{\text{sim}})}, \quad \varepsilon_{\text{BE}}^{\text{qd}}(d) = \frac{\varepsilon_{\text{sim}}}{Q_{\text{qd}}(d; t, \varepsilon_{\text{sim}})}. \quad (113)$$

For the code-switching model, the total-cost ratio is then

$$\mathcal{R}_{\text{tot}}(d) = \frac{T_{\text{tot}}^{\text{qb}}(d)}{T_{\text{tot}}^{\text{qd}}(d)}, \quad (114)$$

where $\mathcal{R}_{\text{tot}}(d) > 1$ indicates a lower total non-Clifford cost for the qudit-based construction.

For the numerical crossover estimates below, we instantiate Eq. (106) using the continuous proxy

$$Q(\alpha; t, \varepsilon_{\text{sim}}) = \alpha t + \log_2(1/\varepsilon_{\text{sim}}), \quad (115)$$

rather than an exact integer query count. Since both encodings use the same qubitization primitive, algorithm-dependent $O(1)$ prefactors affect Q_{qb} and Q_{qd} in the same way. They therefore mainly rescale absolute costs and, through the choice $\varepsilon_{\text{BE}} = \varepsilon_{\text{sim}}/Q$, enter the per-query synthesis budgets only logarithmically. The finite- d crossover thresholds reported below should consequently be interpreted up to small $O(1)$ shifts near break-even.

The proxy in Eq. (115) separates two regimes. In the precision-dominated regime, $\log_2(1/\varepsilon_{\text{sim}})$ is the larger contribution to the query count; in the time-dominated regime, $\alpha(d)t$ is larger. For $\varepsilon_{\text{sim}} = 10^{-6}$, we have $\log_2(1/\varepsilon_{\text{sim}}) \approx 20$. Over the odd dimensions considered here, $3 \leq d \leq 1000$, one finds $\max_d \alpha_{\text{qb}}(d) \times 0.1 \approx 0.40$ and $\max_d \alpha_{\text{qd}}(d) \times 0.1 \approx 0.07$, so both query counts are precision-dominated at $t = 0.1$. By contrast, $\min_d \alpha_{\text{qb}}(d) \times 3000 \approx 3 \times 10^3$ and $\min_d \alpha_{\text{qd}}(d) \times 3000 \approx 2 \times 10^3$, so both query counts are time dominated at $t = 3000$.

Accordingly, we use $t = 0.1$ and $t = 3000$ as representative precision- and time-dominated points, respectively. Throughout this finite- d analysis, we fix $\phi_{\text{max}} = 1$ and vary d by refining the uniform grid spacing $\delta_\phi = 2\phi_{\text{max}}/(d-1)$. These values of t are not intended as application-specific runtimes; they are chosen to isolate the two limiting behaviors of the query model.

2. Fixed-encoding model

We first study the fixed-encoding model and ask how favorable the qudit synthesis prefactor must be for a fully qudit implementation to match or beat the qubit baseline in total non-Clifford cost. As in Sec. III D, we consider prime d for the qudit construction, model the synthesis cost of an embedded two-level SU(2) rotation as

$$C_{\text{prim}}(\delta) \approx a \log_2(1/\delta), \quad (116)$$

and derive a break-even prefactor $a_{\text{max}}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}})$ such that $a < a_{\text{max}}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}})$ implies the qudit construction is no more expensive end-to-end.

In the fixed-encoding construction, the non-Clifford parts of the qudit block-encoding call are expressed in terms of embedded two-level SU(2) rotations. In Appendix B we show that the non-Clifford part of SELECT can be synthesized, up to global phase, using at most $(d-1)$ embedded two-level SU(2) rotations (Corollary B1), while PREP and PREP[†] each

require $(d-1)$ such rotations (Lemma B1). Therefore, for the prime local dimensions considered in the fixed-encoding analysis, the total number of synthesized two-level rotations per block-encoding call is

$$L_{\text{qd}}^{(2)}(d) = 3d - 3. \quad (117)$$

Using the per-query budget from Sec. IV D 1 and allocating it uniformly across these rotations, we take

$$\delta = \frac{\varepsilon_{\text{BE}}^{\text{qd}}(d)}{L_{\text{qd}}^{(2)}(d)}, \quad (118)$$

so that each two-level rotation is synthesized to error at most δ . This gives the per-call qudit non-Clifford estimate

$$N_{\text{qudit}}^{\text{LCU}}(d, \varepsilon_{\text{BE}}^{\text{qd}}(d)) \approx L_{\text{qd}}^{(2)}(d) a \log_2 \left(\frac{L_{\text{qd}}^{(2)}(d)}{\varepsilon_{\text{BE}}^{\text{qd}}(d)} \right). \quad (119)$$

To compare against the qubit Regime 2 baseline of Eq. (111), we define the largest synthesis prefactor a for which the fixed-encoding qudit construction is no more expensive in total cost as

$$a_{\text{max}}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}}) = \frac{T_{\text{tot}}^{\text{qb}}(d)}{Q_{\text{qd}}(d) L_{\text{qd}}^{(2)}(d) \log_2(L_{\text{qd}}^{(2)}(d)/\varepsilon_{\text{BE}}^{\text{qd}}(d))}, \quad (120)$$

where $T_{\text{tot}}^{\text{qb}}(d)$, $Q_{\text{qd}}(d)$, and $\varepsilon_{\text{BE}}^{\text{qd}}(d)$ are defined in Sec. IV D 1. Thus, if $a < a_{\text{max}}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}})$, the fixed-encoding qudit implementation uses fewer non-Cliffords than the qubit encoding within this cost model.

As in the product-formula analysis, it is useful to compare this compiler target with the effective logarithmic prefactor that would reproduce qubit R_z synthesis at the same primitive precision. For Regime 2 fixed-encoding, the relevant primitive precision is $\delta = \varepsilon_{\text{BE}}^{\text{qd}}(d)/L_{\text{qd}}^{(2)}(d)$, so we define

$$a_{R_z}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}}) = \frac{0.57 \log_2(L_{\text{qd}}^{(2)}(d)/\varepsilon_{\text{BE}}^{\text{qd}}(d)) + 8.83}{\log_2(L_{\text{qd}}^{(2)}(d)/\varepsilon_{\text{BE}}^{\text{qd}}(d))}. \quad (121)$$

The comparison

$$a_{\text{max}}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}}) > a_{R_z}^{\text{LCU}}(d, t, \varepsilon_{\text{sim}}) \quad (122)$$

indicates that the fixed-encoding qudit construction can tolerate embedded-SU(2) synthesis at least as costly as qubit R_z synthesis at the same primitive precision. Conversely, if $a_{\text{max}}^{\text{LCU}} < a_{R_z}^{\text{LCU}}$, then a genuine per-primitive qudit synthesis advantage is required.

In the precision-dominated regime, with $\varepsilon_{\text{sim}} = 10^{-6}$ and $t = 0.1$, Table I reports representative values of $a_{\text{max}}^{\text{LCU}}(d, 0.1, 10^{-6})$. The condition $a_{\text{max}}^{\text{LCU}} > a_{R_z}^{\text{LCU}}$ holds only for $d = 3$ and $d = 5$. At $d = 7$ the two quantities are already nearly equal, with $a_{\text{max}}^{\text{LCU}}/a_{R_z}^{\text{LCU}} \approx 0.97$, and for larger prime dimensions the fixed-encoding qudit construction requires an increasingly favorable embedded-SU(2) synthesis prefactor.

In the time-dominated regime, with $\varepsilon_{\text{sim}} = 10^{-6}$ and $t = 3000$, the relative query normalizations make the fixed-encoding comparison more favorable, as shown in Fig. 1. With the same-precision reference prefactor $a_{R_z}^{\text{LCU}}$, the favorable prime dimensions are those for which $a_{\text{max}}^{\text{LCU}} > a_{R_z}^{\text{LCU}}$:

$$d \in \{3, 5, 7, 11, 13, 17, 19\}. \quad (123)$$





FIG. 1. Break-even synthesis prefactor $a_{\max}^{\text{LCU}}(d, 3000, 10^{-6})$ for the fixed-encoding Regime 2 implementation under the coupled per-query error budget. The dashed curve marks the same-precision qubit- R_z effective prefactor $a_{R_z}^{\text{LCU}}(d, 3000, 10^{-6})$ from Eq. (121). Values of $a_{\max}^{\text{LCU}}$ above this reference prefactor indicate that the fully qudit implementation can match or outperform the qubit baseline in total non-Clifford cost without requiring an embedded two-level qudit synthesis prefactor smaller than the corresponding qubit- R_z effective prefactor. Only prime local dimensions d are included in the fixed-encoding analysis.

The largest value in this regime occurs at $d = 5$, where $a_{\max}^{\text{LCU}}(5, 3000, 10^{-6}) = 4.794611$ and $a_{R_z}^{\text{LCU}}(5, 3000, 10^{-6}) = 0.825901$. The last favorable prime dimension is $d = 19$, where $a_{\max}^{\text{LCU}}(19, 3000, 10^{-6}) = 1.339724$ and $a_{R_z}^{\text{LCU}}(19, 3000, 10^{-6}) = 0.810783$. For $d \geq 23$ among the prime dimensions considered, the break-even prefactor lies below the same-precision qubit- R_z reference prefactor.

3. Idealized code-switching model

We now turn to the code-switching model. Using the end-to-end proxies $T_{\text{tot}}^{\text{qb/qd}}(d)$ and the ratio $\mathcal{R}_{\text{tot}}(d)$ defined



in Sec. IV D 1, we consider the idealized hybrid setting in which the n_b -qubit index register can be code-switched to a single d -level qudit (and back) to implement the controlled- Z_d interaction.

Figure 2 shows the resulting total T counts $T_{\text{tot}}^{\text{qb}}(d)$ and $T_{\text{tot}}^{\text{qd}}(d)$, together with the corresponding ratio $\mathcal{R}_{\text{tot}}(d)$ at $\varepsilon_{\text{sim}} = 10^{-6}$ and $t = 0.1$. For these parameters, the qudit construction is cheaper only for $d = 3$ and $d = 5$, while the qubit construction is cheaper for $d \geq 7$. The crossover occurs between $d = 5$ and $d = 7$: $\mathcal{R}_{\text{tot}}(5) = 1.006205$ and $\mathcal{R}_{\text{tot}}(7) = 0.999963$. The piecewise-constant structure in d is a consequence of the dependence on $n_b = \lceil \log_2 d \rceil$. The largest relative qudit advantage occurs at $d = 3$, where $\mathcal{R}_{\text{tot}}(3) = 2.033787$, and the largest absolute savings also occur at $d = 3$, where the qudit construction saves $T_{\text{tot}}^{\text{qb}}(3) - T_{\text{tot}}^{\text{qd}}(3) \approx 4.20 \times 10^3$ T gates.

Figure 3 shows the resulting total T counts, $T_{\text{tot}}^{\text{qb}}(d)$ and $T_{\text{tot}}^{\text{qd}}(d)$, together with the corresponding ratio $\mathcal{R}_{\text{tot}}(d)$ at $\varepsilon_{\text{sim}} = 10^{-6}$ and $t = 3000$ (time-dominated regime). For these parameters, the qudit construction is cheaper for $d \in [3, 13] \cup [17, 21]$, while the qubit construction is cheaper outside these windows. The final crossover occurs between $d = 21$ and $d = 23$: $\mathcal{R}_{\text{tot}}(21) = 1.062653$, whereas $\mathcal{R}_{\text{tot}}(23) = 0.835319$. As in the precision-dominated comparison, the piecewise structure in d originates from the dependence on $n_b = \lceil \log_2 d \rceil$. The largest relative qudit advantage occurs at $d = 5$, where $\mathcal{R}_{\text{tot}}(5) = 3.959978$, while the largest absolute savings occur at $d = 9$, where the qudit construction saves $T_{\text{tot}}^{\text{qb}}(9) - T_{\text{tot}}^{\text{qd}}(9) \approx 3.65 \times 10^6$ T gates.

We can also interpret the qudit advantage as a code-switching budget. Let

$$\Delta_{\text{tot}}(d) = T_{\text{tot}}^{\text{qb}}(d) - T_{\text{tot}}^{\text{qd}}(d), \quad (124)$$

and suppose the implementation requires k directional code-switches per query. In the hybrid construction considered here, the n_b -qubit index register is used for PREP (and the diagonal R_z -synthesis term in SELECT); it is then converted to a single d -level qudit to implement the controlled- Z_d




FIG. 2. Non-Clifford resource comparison for the Regime 2 LCU per block-encoding implementation under the code-switching model, for $\varepsilon_{\text{sim}} = 10^{-6}$ and $t = 0.1$. Panel (a) shows the ratio $\mathcal{R}_{\text{tot}}(d) = T_{\text{tot}}^{\text{qb}}(d)/T_{\text{tot}}^{\text{qd}}(d)$ as a function of the local dimension d . Panel (b) shows the corresponding total T -count estimates $T_{\text{tot}}^{\text{qb}}(d)$ and $T_{\text{tot}}^{\text{qd}}(d)$.




FIG. 3. Non-Clifford resource comparison for the Regime 2 LCU per block-encoding implementation under the code-switching model, for $\varepsilon_{\text{sim}} = 10^{-6}$ and $t = 3000$. Panel (a) shows the ratio $\mathcal{R}_{\text{tot}}(d) = T_{\text{tot}}^{\text{qb}}(d)/T_{\text{tot}}^{\text{qd}}(d)$ as a function of the local dimension d . Panel (b) shows the corresponding total T -count estimates $T_{\text{tot}}^{\text{qb}}(d)$ and $T_{\text{tot}}^{\text{qd}}(d)$. The piecewise-constant structure and repeated crossover windows arise from the binary embedding.

interaction, and finally converted back to qubits, so $k = 2$. If each switch incurs non-Clifford overhead T_{cs} , the added overhead is $Q_{\text{qd}}(d)k T_{\text{cs}}$, where $Q_{\text{qd}}(d)$ is the qudit query count from Sec. IV D 1. Thus, qudits remain advantageous provided

$$T_{\text{cs}} < \frac{\Delta_{\text{tot}}(d)}{Q_{\text{qd}}(d)k}. \quad (125)$$

For $\varepsilon_{\text{sim}} = 10^{-6}$ and $t = 0.1$ with $k = 2$, this yields representative budgets $T_{\text{cs}} \approx 1.05 \times 10^2$ (for $d = 3$) and $T_{\text{cs}} \approx 1.35 \times 10^0$ (for $d = 5$); the budget is already negative by $d = 7$, reflecting the fact that the total-cost crossover has already occurred.

For $\varepsilon_{\text{sim}} = 10^{-6}$ and $t = 3000$, this yields representative budgets $T_{\text{cs}} \approx 2.87 \times 10^2$ (for $d = 3$), $T_{\text{cs}} \approx 7.42 \times 10^2$ (for $d = 5$), $T_{\text{cs}} \approx 8.97 \times 10^2$ (for $d = 9$), and $T_{\text{cs}} \approx 6.65 \times 10^2$ (for $d = 17$), while near the final crossover the budget drops to $T_{\text{cs}} \approx 6.34 \times 10^1$ at $d = 21$ and becomes negative by $d = 23$. We discuss the implications of such a comparison in Sec. V.

V. DISCUSSION

We compare fault-tolerant qudit and qubit encodings for implementing the onsite diagonal operator $e^{-it\phi_x^2}$ arising from a d -level field-amplitude truncation at matched accuracy. The central question is not whether qudits win asymptotically for this operator class, but whether finite- d windows exist in which they can reduce non-Clifford cost. Our results show that such windows do exist, but their size and relevance depend strongly on the simulation regime, the achievable synthesis cost of embedded two-level qudit primitives, and, in hybrid constructions, the overhead of code-switching.

Table II summarizes the leading-order non-Clifford scalings established below for both regimes and encodings, within the constructive cost models analyzed here.

In Regime 1, the binary qubit encoding has an asymptotic advantage for structural reasons: the binary expansion of the onsite quadratic phase yields only one- and two-body

Z and ZZ terms, and hence only $\Theta((\log d)^2)$ synthesized R_z rotations per Trotter step. By contrast, the single-qudit implementation analyzed in Sec. III B uses exactly $d - 1$ nontrivial embedded two-level $R_Z^{(k,k+1)}(\theta_k)$ rotations for the symmetric truncation. Under the common logarithmic-precision synthesis model, this means that for the qudit route to asymptotically beat the qubit route, the per-rotation synthesis prefactor for the embedded $\text{SU}(2)$ rotations would need to be smaller than that of a qubit R_z by a factor $\Theta((\log d)^2/d)$, i.e., an exponentially strong improvement in $n_b = \lceil \log_2 d \rceil$. Since known fault-tolerant synthesis procedures already achieve $O(\log(1/\delta))$ scaling for common discrete gate sets [47,64,76,77], we do not expect such an exponential per-rotation advantage to arise. Likewise, phase kickback does not evade the basic issue: after coherently computing n^2 , one must still synthesize a diagonal phase function of that value.

The finite- d Regime 1 crossover analysis points in the same direction. Modeling each embedded qudit $\text{SU}(2)$ primitive as costing $\approx a \log_2(1/\delta)$ non-Cliffords and comparing against the qubit R_z synthesis model $0.57 \log_2(1/\delta) + 8.83$ yields the break-even threshold $a_{\text{max}}^{\text{PF}}(d, \varepsilon)$ and the qubit- R_z effective prefactor $a_{R_z}^{\text{PF}}(d, \varepsilon)$ (Sec. III D). For representative $\varepsilon = 10^{-6}$, we find $a_{\text{max}}^{\text{PF}}(3, 10^{-6}) \approx 1.51$ and $a_{\text{max}}^{\text{PF}}(5, 10^{-6}) \approx 1.48$, while $a_{\text{max}}^{\text{PF}}(7, 10^{-6}) \approx 0.96$ already falls below $a_{R_z}^{\text{PF}}$. The condition $a_{\text{max}}^{\text{PF}} > a_{R_z}^{\text{PF}}$ holds only for $d = 3$ and $d = 5$ among prime dimensions at this precision. We therefore view Regime 1 as allowing, at most, a narrow low- d opportunity for qudits, rather than a robust route to scalable savings.

In Regime 2, the comparison is between the native qudit LCU construction and the standard signed-binary qubit projector-LCU baseline of Refs. [59,60], both targeting the same physical $d = 2M + 1$ field-amplitude truncation. The qudit construction makes SELECT particularly cheap: when both index and system are encoded as qudits, the entangling part $\sum_r |r\rangle\langle r| \otimes Z'_d$ is a generalized controlled- Z and hence Clifford, leaving only a simple diagonal on the index register

as the non-Clifford component. However, the qubit baseline still has an asymptotic advantage overall, because in the generalized-Pauli expansion on a d -level system, ϕ_x^2 has irreducible LCU length $d - 1$, so the corresponding PREP oracle amounts to generic $(d - 1)$ -dimensional state preparation requiring $d - 1$ embedded two-level rotations (Lemma B1). By contrast, the qubit baseline's per-call cost grows only as $\Theta(n_b)$ (Lemma 3). Thus, while SELECT is especially favorable in the qudit encoding, the full oracle cost is not asymptotically improved (Theorem 2). As elsewhere in Regime 2, this is a construction-specific statement about the standard signed-binary qubit projector-LCU baseline; alternative qubit LCU constructions could change the asymptotics.

We therefore analyze finite- d behavior in an end-to-end sense by weighting per-block-encoding costs by the qubitization query complexity (Sec. IV D 2). In the fixed-encoding setting, the precision-dominated regime ($\varepsilon_{\text{sim}} = 10^{-6}$, $t = 0.1$) is relatively restrictive: as shown in Table I, the condition $a_{\text{max}}^{\text{LCU}} > a_{R_z}^{\text{LCU}}$ holds only for $d = 3$ ($a_{\text{max}}^{\text{LCU}} = 2.56$) and $d = 5$ ($a_{\text{max}}^{\text{LCU}} = 1.32$), while $a_{\text{max}}^{\text{LCU}}$ falls below the qubit- R_z reference prefactor for $d \geq 7$.

In the time-dominated regime ($\varepsilon_{\text{sim}} = 10^{-6}$, $t = 3000$), the picture is more favorable to qudits: $a_{\text{max}}^{\text{LCU}}(d, 3000, 10^{-6})$ exhibits multiple threshold crossings as a function of d , reflecting the piecewise dependence on the binary embedding through $n_b = \lceil \log_2 d \rceil$ (Fig. 1). The condition $a_{\text{max}}^{\text{LCU}} > a_{R_z}^{\text{LCU}}$ holds for all prime dimensions $d \in \{3, 5, 7, 11, 13, 17, 19\}$, with the largest value $a_{\text{max}}^{\text{LCU}}(5, 3000, 10^{-6}) = 4.794611$ and the last favorable prime $d = 19$, where $a_{\text{max}}^{\text{LCU}}(19, 3000, 10^{-6}) = 1.339724$ compared with $a_{R_z}^{\text{LCU}}(19, 3000, 10^{-6}) = 0.810783$. For $d \geq 23$ among prime dimensions, the break-even prefactor lies below the qubit- R_z reference. These values are not uniform in d , so any practical advantage remains truncation-dependent.

For context, existing qutrit synthesis results report effective logarithmic prefactors around 3.16 to 3.24 (in R count) for relevant diagonal or two-level primitives [64,76]. Because the fixed-encoding break-even thresholds $a_{\text{max}}^{\text{LCU}}$ are defined relative to the qubit T-gate synthesis model, a direct numerical comparison with R -count data requires a gate-set conversion not available in general; nevertheless, the favorable $d = 3$ window in the time-dominated regime has a break-even prefactor substantially above the qubit reference $a_{R_z}^{\text{LCU}} \approx 0.83$, suggesting that concrete qutrit synthesis data could fall within the achievable range. Furthermore, for odd prime dimensions, Ref. [62] gives an asymptotic worst-case lower bound for approximating arbitrary $\text{SU}(p)$ unitaries by single-qudit Clifford + T circuits. This bound sets a floor on achievable synthesis cost and provides context for whether the break-even thresholds identified here are plausibly reachable, but a definitive comparison requires dimension-specific synthesis results.

We then consider the code-switching setting, which we emphasize is entirely idealized: determining the fault-tolerant overhead of converting between a logical qudit and a logical qubit register remains an open problem. In the precision-dominated regime ($t = 0.1$, $\varepsilon_{\text{sim}} = 10^{-6}$, Fig. 2), the qudit construction is cheaper only for $d = 3$ and $d = 5$, with the crossover occurring between $d = 5$ and $d = 7$:

$\mathcal{R}_{\text{tot}}(5) = 1.006205$ while $\mathcal{R}_{\text{tot}}(7) = 0.999963$. The largest relative advantage occurs at $d = 3$, where $\mathcal{R}_{\text{tot}}(3) = 2.033787$, and the largest absolute savings also occur at $d = 3$, where the qudit construction saves approximately 4.20×10^3 T gates.

In the time-dominated regime ($t = 3000$, $\varepsilon_{\text{sim}} = 10^{-6}$, Fig. 3), the end-to-end comparison exhibits multiple crossings tied to the binary embedding, with the qudit construction cheaper for $d \in [3, 13] \cup [17, 21]$ and the qubit encoding cheaper outside these windows. The final crossover occurs between $d = 21$ and $d = 23$: $\mathcal{R}_{\text{tot}}(21) = 1.062653$ while $\mathcal{R}_{\text{tot}}(23) = 0.835319$. The largest relative advantage occurs at $d = 5$, where $\mathcal{R}_{\text{tot}}(5) = 3.959978$, while the largest absolute savings occur at $d = 9$, where the qudit construction saves approximately 3.65×10^6 T gates.

Equivalently, these savings define a loose break-even budget for code-switching. Writing $\Delta_{\text{tot}}(d) = T_{\text{tot}}^{\text{qb}}(d) - T_{\text{tot}}^{\text{qd}}(d)$ and assuming $k = 2$ directional code switches per query, qudit implementations remain advantageous provided the additional non-Clifford overhead per switch satisfies Eq. (125). In the precision-dominated regime, representative budgets are $T_{\text{cs}} \approx 1.05 \times 10^2$ T gates per switch at $d = 3$ and $T_{\text{cs}} \approx 1.35$ at $d = 5$; the budget is already negative at $d = 7$. In the time-dominated regime, budgets reach up to $O(10^3)$ T gates per switch near the most favorable dimensions (e.g., $T_{\text{cs}} \approx 8.97 \times 10^2$ at $d = 9$), declining toward the final crossover ($T_{\text{cs}} \approx 6.34 \times 10^1$ at $d = 21$, negative for $d \geq 23$).

However, comparing T_{cs} with the actual fault-tolerant cost of code-switching in this setting is highly nontrivial. The overhead of converting between a fault-tolerant d -level qudit encoding and a fault-tolerant qubit register is strongly architecture- and code-dependent, and a concrete implementation with explicit space-time accounting remains an open problem. Relatedly, Ref. [78] introduced qubit fusion and qudit fission, effecting a conversion between two qubits and a single $d = 4$ qudit via stabilizer circuits and the consumption of a nonstabilizer resource state $|F\rangle$. In that framework, a natural route to a T -equivalent estimate is to compare the cost of supplying high-fidelity $|F\rangle$ states with that of supplying high-fidelity qubit magic states. However, the absolute and relative costs of $|F\rangle$ versus $|T\rangle$ depend sensitively on the available distillation protocols and the underlying fault-tolerant architecture. Consequently, a full end-to-end resource estimate for code switching—including state injection, distillation, and the associated space-time overhead of the conversion circuitry—is beyond the scope of this work.

The small- d window of potential advantage identified in Regime 1 and in the precision-dominated Regime 2 comparison is the most accessible regime for near-term qudit synthesis efforts. The time-dominated Regime 2 model extends this window further: in the fixed-encoding comparison, the break-even condition $a_{\text{max}}^{\text{LCU}} > a_{R_z}^{\text{LCU}}$ holds for prime dimensions up to $d = 19$; in the code-switching comparison, the favorable range extends to $d \leq 21$. Reference [18] argues that roughly four to seven qubits per site (corresponding to $d = 2^4$ to 2^7) are required before digitization error in ϕ_x is reduced below typical Trotter error. This scale is at or above the end of the favorable windows identified here: only the lower end of this range ($d \approx 16$ to 19) overlaps with the time-dominated break-

even region, while larger truncations ($d \geq 23$) lie outside the favorable window in both models. Furthermore, diagonal evolutions generated by quadratic onsite or link operators of the same form arise naturally in other settings, where the physically relevant Hilbert-space cutoffs are often explicitly small integers. For example, for gauge links in $(1 + 1)$ -dimensional $[(1 + 1)\text{D}]$ scalar quantum electrodynamics (QED), a truncation $d \leq 13$ appears sufficient [61,79]. For pure $U(1)$ gauge theory in $(2 + 1)\text{D}$, Ref. [80] finds per-mille plaquette accuracy with only $d = 7$ states per lattice site; and for two-dimensional (2D) lattice QED plaquette or gauge-link simulations, Refs. [81,82] explicitly work with minimal-to-improved gauge-field discretizations at $d = 3$ and $d = 5$, with convergence trends indicating that these modest d already capture the relevant low-energy physics. These examples further motivate investigations of fault-tolerant qudits for $d > 3$.

Finally, we comment on how far these conclusions may extend beyond the specific operator ϕ_x^2 . For a fixed-degree polynomial diagonal term $h(N)$ in a truncated integer-valued operator N , the Regime 1 asymptotic lesson is expected to persist under the same constructive route: the binary qubit expansion yields a number of Z -type rotations that grows polynomially in $n_b = \lceil \log_2 d \rceil$, whereas the direct embedded-rotation qudit diagonal implementation of Lemma 2 generally uses $O(d)$ two-level rotations. The quadratic case is therefore not uniquely responsible for the asymptotic ordering. The finite- d crossover thresholds, by contrast, are operator-dependent, so the specific break-even values reported here would need to be recomputed for higher-order potentials. For generic, unstructured diagonal operators we claim no universal conclusion: the comparison can become sensitive to the choice of basis, oracle construction, and synthesis model. Conversely, some lattice-gauge-theory settings, for example $\mathbb{Z}_N$ lattice gauge theories, carry qudit structure that may structurally favor qudit-based constructions: recent work exploiting prime- N $\mathbb{Z}_N$ lattice-gauge-theory structure within a qudit stabilizer-code framework is one example [83]. We therefore regard the present quadratic analysis as a controlled benchmark for qudit-versus-qubit non-Clifford costs, rather than a universal statement about all diagonal Hamiltonians.

VI. CONCLUSIONS

Overall, these results suggest that fault-tolerant qudits should be viewed not as providing a generic asymptotic advantage for this class of quadratic diagonal evolutions, but rather as a potentially useful finite- d resource when the truncation, simulation regime, and achievable synthesis overhead are favorable. Within the constructions studied here, the LCU per block-encoding (Regime 2) comparison is the more promising setting: the fixed-encoding analysis identifies break-even conditions at low prime dimensions (favorable up to $d = 19$ in the time-dominated regime), while the idealized code-switching analysis shows that meaningful constant-factor savings can arise at selected truncations up to $d \lesssim 21$, subject to architecture-dependent conversion costs. The Regime 2 comparison is, in particular, a construction-specific comparison between the native qudit LCU construction and the standard signed-binary qubit projector-LCU baseline of

Refs. [59,60]; alternative qubit LCU constructions could change the comparison in either direction.

We therefore interpret the break-even thresholds derived here as concrete compiler targets rather than as a blanket demonstration of qudit superiority. A more definitive comparison will require explicit gate-set-dependent synthesis costs for diagonal $SU(d)$ operations, or equivalently for the embedded two-level qudit $SU(2)$ primitives used throughout, together with architecture-level resource estimates for code conversion, state injection and distillation, and associated space-time overheads. It would also be valuable to extend this analysis to other operator classes for which the structural simplifications available to binary qubit encodings for diagonal operators do not apply.

ACKNOWLEDGMENTS

We thank Joseph Carlson, Mino Kabirnezhad, Andy C. Y. Li, and Constantinos Andreopoulos for valuable conversations, as well as Cameron Cook for helpful comments on the draft. S.G. acknowledges an STFC Ph.D. studentship at the LIV.INNO Centre for Doctoral Training “Innovation in Data Intensive Science.” G.N.P. and D.M.K. were partially supported for this work by the DOE/HEP QuantISED program grant “HEP Machine Learning and Optimization Go Quantum,” Identification No. 0000240323. M.M. acknowledges the support of Agencia Nacional de Investigación e Innovación, project POS_EXT_2023_1_175687, and LIV.INNO. This work was produced 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. The United States Government retains and the publisher, by accepting the work for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this work, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan [84]. Fermilab Report No. FERMILAB-PUB-26-0287-ETD.

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.

APPENDIX A: SELECT-ORACLE COST DETAILS

Lemma A1. The diagonal phase operator

$$D = \sum_{r=0}^{d-1} e^{i\theta_r} |r\rangle\langle r| \quad (\text{A1})$$

with $\theta_0 = 0$ can be written as

$$D = \sum_{r=0}^{d-1} \text{sgn}(c_r) e^{i\pi r/d} |r\rangle\langle r| = D_{\text{sign}} D_{\text{clock}}, \quad (\text{A2})$$

where $c_r \in \mathbb{R}$ and

$$D_{\text{clock}} = \sum_{r=0}^{d-1} e^{i\pi r/d} |r\rangle\langle r|, \\ D_{\text{sign}} = |0\rangle\langle 0| + \sum_{r=1}^{d-1} \text{sgn}(c_r) |r\rangle\langle r|. \quad (\text{A3})$$

Proof. The operator ϕ_x^2 is diagonal with real eigenvalues λ_n^2 , so its expansion in the generalized Pauli basis is

$$\phi_x^2 = \sum_{r=0}^{d-1} \beta_r Z_d^r, \quad (\text{A4})$$

with Fourier coefficients

$$\beta_r = \frac{1}{d} \sum_{n=0}^{d-1} \lambda_n^2 \omega^{-rn}, \quad \omega = e^{2\pi i/d}. \quad (\text{A5})$$

For the uniform symmetric truncation used here, the explicit evaluation of this sum [Eq. (77)] gives, for $r \neq 0$,

$$\beta_r = \frac{2\phi_{\text{max}}^2}{(d-1)^2} e^{i\pi r/d} \frac{\cos(\pi r/d)}{\sin^2(\pi r/d)} = c_r e^{i\pi r/d}, \quad (\text{A6})$$

where

$$c_r = \frac{2\phi_{\text{max}}^2}{(d-1)^2} \frac{\cos(\pi r/d)}{\sin^2(\pi r/d)} \in \mathbb{R}. \quad (\text{A7})$$

Thus each β_r has a fixed phase $e^{i\pi r/d}$ and a real amplitude c_r .

Defining the phases of D by $e^{i\theta_r} = \beta_r/|\beta_r|$ and using $\text{sgn}(c_r) = c_r/|c_r|$, we obtain

$$e^{i\theta_r} = \frac{\beta_r}{|\beta_r|} = \frac{c_r e^{i\pi r/d}}{|c_r|} = \text{sgn}(c_r) e^{i\pi r/d}, \quad (\text{A8})$$

so

$$D = \sum_{r=0}^{d-1} e^{i\theta_r} |r\rangle\langle r| = D_{\text{sign}} D_{\text{clock}}, \quad (\text{A9})$$

as claimed. $\blacksquare$

Lemma A2. The operator $D_{\text{clock}} = \exp(i\frac{\pi}{d}N)$ can be implemented using $O(n_b)$ single-qubit Z rotations on the index register, i.e., $O(n_b)$ non-Clifford gates.

Proof. We first write the integer r in binary as

$$r = \sum_{m=0}^{n_b-1} 2^m q_m, \quad q_m \in \{0, 1\}, \quad (\text{A10})$$

where q_m is the bit stored on qubit m of the index register. In terms of Pauli operators,

$$q_m = \frac{1 - Z^{(m)}}{2}, \quad (\text{A11})$$

where $Z^{(m)}$ is the Pauli- Z on qubit m . Substituting this into the definition of the number operator, we obtain

$$N = \sum_{m=0}^{n_b-1} 2^m q_m = \frac{1}{2} \sum_{m=0}^{n_b-1} 2^m - \frac{1}{2} \sum_{m=0}^{n_b-1} 2^m Z^{(m)}. \quad (\text{A12})$$

It follows that

$$D_{\text{clock}} = \exp\left(i\frac{\pi}{d}N\right) \\ = \exp\left(i\frac{\pi}{2d} \sum_{m=0}^{n_b-1} 2^m\right) \prod_{m=0}^{n_b-1} \exp\left(-i\frac{\pi}{2d} 2^m Z^{(m)}\right). \quad (\text{A13})$$

The prefactor

$$\exp\left(i\frac{\pi}{2d} \sum_{m=0}^{n_b-1} 2^m\right) \quad (\text{A14})$$

is a global phase and can be ignored. Hence D_{clock} decomposes as a product of n_b single-qubit Z rotations on the index register, one on each qubit m with angle $-\frac{\pi}{2d} 2^m$. In particular, the cost of implementing D_{clock} is just n_b single-qubit rotations. $\blacksquare$

Lemma A3. The diagonal operator

$$D_{\text{sign}} = |0\rangle\langle 0| + \sum_{r=1}^{d-1} \text{sgn}(c_r) |r\rangle\langle r| \quad (\text{A15})$$

can be implemented using $4n_b$ T gates, together with n_b scratch ancillas and a single flag ancilla.

Proof. From the explicit form of the Fourier coefficients of ϕ_x^2 , we have, for $r \neq 0$,

$$\beta_r = c_r e^{i\pi r/d}, \quad c_r = K \frac{\cos(\pi r/d)}{\sin^2(\pi r/d)}, \quad K \in \mathbb{R}, \quad (\text{A16})$$

so $c_r \in \mathbb{R}$ for all r . For odd d , we have $0 < \frac{\pi r}{d} < \pi$ for $1 \leq r \leq d-1$, so $\sin(\pi r/d) > 0$ and the sign of c_r is determined by $\cos(\pi r/d)$ alone. Since

$$\cos(\pi r/d) > 0 \quad \text{for } 1 \leq r \leq \frac{d-1}{2},$$

$$\cos(\pi r/d) < 0 \quad \text{for } \frac{d+1}{2} \leq r \leq d-1, \quad (\text{A17})$$

it follows that

$$c_r > 0 \quad \text{for } 1 \leq r \leq \frac{d-1}{2},$$

$$c_r < 0 \quad \text{for } \frac{d+1}{2} \leq r \leq d-1 \quad (\text{A18})$$

and the pattern is antisymmetric $c_{d-r} = -c_r$. In particular, $\text{sgn}(c_r)$ is a single-threshold function of the integer label r with threshold $T = (d+1)/2$.

We implement D_{sign} using a standard reversible comparator acting on the n_b -qubit index register and a single ancilla. Let U_{cmp} be a comparator unitary that compares the value r on the index register with the fixed classical threshold T and write the result into an ancilla qubit $|a\rangle$ initialized in $|0\rangle$:

$$U_{\text{cmp}} : |r\rangle|0\rangle \mapsto |r\rangle|f(r)\rangle, \quad (\text{A19})$$

where

$$f(r) = \begin{cases} 0, & r < T \\ 1, & r \geq T. \end{cases} \quad (\text{A20})$$

Starting from $|r\rangle|0\rangle$, we perform the following three-step sequence:

$$\begin{aligned} |r\rangle|0\rangle &\xrightarrow{U_{\text{cmp}}} |r\rangle|f(r)\rangle \xrightarrow{Z_a} (-1)^{f(r)}|r\rangle|f(r)\rangle \\ &\xrightarrow{U_{\text{cmp}}^\dagger} (-1)^{f(r)}|r\rangle|0\rangle. \end{aligned} \quad (\text{A21})$$

After this sequence, the ancilla is returned to $|0\rangle$, and the net action on the index register is

$$|r\rangle \mapsto (-1)^{f(r)}|r\rangle = \begin{cases} +|r\rangle, & r < T \\ -|r\rangle, & r \geq T, \end{cases} \quad (\text{A22})$$

which is precisely the desired diagonal

$$\sum_{r=1}^{d-1} (-1)^{f(r)}|r\rangle\langle r| = \sum_{r=1}^{d-1} \text{sgn}(c_r)|r\rangle\langle r|. \quad (\text{A23})$$

In compact form we can write

$$D_{\text{sign}} = U_{\text{cmp}}^\dagger (I \otimes Z_a) U_{\text{cmp}}. \quad (\text{A24})$$

The cost of this implementation is dominated by the comparator. Reference [68] [Appendix C 1(a), Eqs. (C22)–(C23)] presents an n_b -bit comparator and shows that an optimized variant CMP' can be implemented using n_b logical ANDs, hence a T count of $4n_b$ (up to Clifford gates), while $\text{CMP}'^\dagger$ is implemented using measurements and classically conditioned Clifford gates. Therefore, the phase-kickback construction $D_{\text{sign}} = U_{\text{cmp}}^\dagger (I \otimes Z_a) U_{\text{cmp}}$ has a T count of $4n_b$, using $O(n_b)$ scratch qubits for carries and one flag ancilla. ■

APPENDIX B: REGIME-2 CROSSOVER DETAILS

Corollary B1. Define D as in Corollary 1. Then there exist angles $\vartheta_0, \dots, \vartheta_{d-2}$ (unique modulo 4π) and a global phase $\gamma \in \mathbb{R}$ (unique modulo 2π) such that

$$D = e^{i\gamma} \prod_{k=0}^{d-2} R_Z^{(k,k+1)}(\vartheta_k). \quad (\text{B1})$$

Consequently, D can be implemented using at most $d-1$ adjacent embedded $\text{SU}(2)$ rotations.

Proof. Adjoin the trivial phase on $|0\rangle$ and write

$$D = \sum_{n=0}^{d-1} e^{i\theta_n} |n\rangle\langle n|, \quad \theta_0 = 0, \quad (\text{B2})$$

where for $r \geq 1$ the phases are those defining D from Eq. (A8). With Eq. (A18) determining $\text{sgn}(c_r)$, we have

$$\theta_r = \begin{cases} \frac{\pi r}{d}, & \text{sgn}(c_r) = +1 \\ \frac{\pi r}{d} + \pi, & \text{sgn}(c_r) = -1, \end{cases} \quad (\text{B3})$$

so that $\theta_r \in [0, 2\pi)$ and $e^{i\theta_r} = \beta_r/|\beta_r|$. Define the global phase

$$\gamma = \frac{1}{d} \sum_{n=0}^{d-1} \theta_n, \quad \alpha_n = \theta_n - \gamma \quad (\text{B4})$$

so that $\sum_{n=0}^{d-1} \alpha_n = 0$, and

$$D = e^{i\gamma} \sum_{n=0}^{d-1} e^{i\alpha_n} |n\rangle\langle n|. \quad (\text{B5})$$

With this form, we can invoke Lemma 2: for any diagonal phases $(\alpha_n)_{n=0}^{d-1}$ with $\sum_n \alpha_n = 0$, there exists a unique choice of adjacent angles

$$\vartheta_k = -2 \sum_{n=0}^k \alpha_n, \quad k = 0, 1, \dots, d-2 \quad (\text{B6})$$

such that

$$\prod_{k=0}^{d-2} R_Z^{(k,k+1)}(\vartheta_k) = \sum_{n=0}^{d-1} e^{i\alpha_n} |n\rangle\langle n|. \quad (\text{B7})$$

Therefore,

$$D = e^{i\gamma} \prod_{k=0}^{d-2} R_Z^{(k,k+1)}(\vartheta_k), \quad (\text{B8})$$

proving existence; uniqueness (mod 4π) follows from the uniqueness statement in Lemma 2. In our case, by Eq. (A8), we have

$$e^{i\theta_r} = \frac{\beta_r}{|\beta_r|} = \text{sgn}(c_r) e^{i\pi r/d}, \quad r = 1, \dots, d-1, \quad (\text{B9})$$

where c_r is determined by Eq. (A18). Substituting $\alpha_n = \theta_n - \gamma$ into Eq. (B6) yields

$$\vartheta_k = -2 \left(\sum_{n=0}^k \theta_n - (k+1)\gamma \right), \quad k = 0, 1, \dots, d-2. \quad (\text{B10})$$

Evaluating these sums gives the explicit closed form

$$\vartheta_k = \begin{cases} \frac{\pi}{d}(k+1)(4m-k), & 0 \leq k \leq m \\ \frac{\pi}{d}(k+1)(4m-k) - 2\pi(k-m), & m < k \leq d-2, \end{cases} \quad (\text{B11})$$

where $m = (d-1)/2$. ■

Note that $R_Z^{(k,k+1)}(\vartheta_k)$ is trivial if and only if $\vartheta_k \equiv 0 \pmod{4\pi}$. In the numerical range considered in Sec. IV D 2, the exact value of $s(d)$ is evaluated directly from the closed form in Eq. (B11). For odd $d \leq 513$, one finds the decomposition uses $d-1$, $d-2$, or $d-4$ adjacent embedded rotations. Therefore, Corollary B1 gives the uniform upper bound $d-1$.

Lemma B1. Define the state

$$|\psi_\beta\rangle = \sum_{r=1}^{d-1} \sqrt{\frac{|\beta_r|}{\Lambda}} |r\rangle, \quad \Lambda = \sum_{r=1}^{d-1} |\beta_r| \quad (\text{B12})$$

from Eq. (93). Then there exists a product of exactly $d-1$ embedded two-level $\text{SU}(2)$ rotations that prepare $|\psi_\beta\rangle$ from $|0\rangle$, i.e.,

$$\text{PREP} |0\rangle = |\psi_\beta\rangle, \quad (\text{B13})$$

with the explicit decomposition

$$\text{PREP} = \prod_{r=d-1}^1 R_Y^{(0,r)}(\theta_r). \quad (\text{B14})$$

Proof. Write

$$a_r = \sqrt{\frac{|\beta_r|}{\Lambda}}, \quad r = 1, \dots, d-1, \quad (\text{B15})$$

so that $a_r \geq 0$ and $\sum_{r=1}^{d-1} a_r^2 = 1$. For each $r \in \{1, \dots, d-1\}$, let

$$U_r = R_Y^{(0,r)}(\theta_r) \quad (\text{B16})$$

denote the embedded two-level Y -rotation acting only on $\text{span}\{|0\rangle, |r\rangle\}$. In that two-dimensional subspace it acts as

$$\begin{aligned} |0\rangle &\mapsto \cos\left(\frac{\theta_r}{2}\right)|0\rangle + \sin\left(\frac{\theta_r}{2}\right)|r\rangle, \\ |r\rangle &\mapsto -\sin\left(\frac{\theta_r}{2}\right)|0\rangle + \cos\left(\frac{\theta_r}{2}\right)|r\rangle, \end{aligned} \quad (\text{B17})$$

and it acts trivially on all basis states $|k\rangle$ with $k \neq 0, r$. We apply these rotations sequentially to $|0\rangle$, i.e.,

$$|\psi_r\rangle = U_r U_{r-1} \cdots U_1 |0\rangle. \quad (\text{B18})$$

Since U_r only mixes the amplitudes on $|0\rangle$ and $|r\rangle$, we can track the amplitudes explicitly. After r steps the state has the form

$$\begin{aligned} |\psi_r\rangle &= \left(\prod_{k=1}^r \cos\left(\frac{\theta_k}{2}\right)\right) |0\rangle \\ &+ \sum_{j=1}^r \left(\sin\left(\frac{\theta_j}{2}\right) \prod_{k=1}^{j-1} \cos\left(\frac{\theta_k}{2}\right)\right) |j\rangle. \end{aligned} \quad (\text{B19})$$

Indeed, at step r the operator U_r multiplies the current amplitude on $|0\rangle$ by $\cos(\theta_r/2)$, and creates a new amplitude on $|r\rangle$ equal to $\sin(\theta_r/2)$ times the current amplitude on $|0\rangle$, while leaving all previously prepared levels $|1\rangle, \dots, |r-1\rangle$ unchanged.

We now choose the angles θ_r so that the amplitude on $|r\rangle$ equals the desired value a_r . From (B19), we require

$$a_r = \left(\prod_{k=1}^{r-1} \cos\left(\frac{\theta_k}{2}\right)\right) \sin\left(\frac{\theta_r}{2}\right), \quad (\text{B20})$$

or equivalently,

$$\sin\left(\frac{\theta_r}{2}\right) = \frac{a_r}{\prod_{k=1}^{r-1} \cos\left(\frac{\theta_k}{2}\right)}. \quad (\text{B21})$$

This gives the explicit update rule for θ_r . This recursion is well-defined because the denominator is exactly the current amplitude of $|0\rangle$. Moreover, since

$$\left(\prod_{k=1}^{r-1} \cos^2\left(\frac{\theta_k}{2}\right)\right) = 1 - \sum_{j=1}^{r-1} a_j^2 \geq a_r^2, \quad (\text{B22})$$

the ratio in (B21) lies in $[0,1]$, so θ_r always exists. With this choice, (B19) implies that after $d-1$ steps the amplitudes satisfy

$$|\psi_{d-1}\rangle = \left(\prod_{k=1}^{d-1} \cos\left(\frac{\theta_k}{2}\right)\right) |0\rangle + \sum_{r=1}^{d-1} a_r |r\rangle. \quad (\text{B23})$$

Finally, since $\sum_{r=1}^{d-1} a_r^2 = 1$, the leftover amplitude on $|0\rangle$ must vanish, i.e.,

$$\prod_{k=1}^{d-1} \cos\left(\frac{\theta_k}{2}\right) = 0, \quad (\text{B24})$$

and therefore

$$\left(\prod_{r=d-1}^1 R_Y^{(0,r)}(\theta_r)\right) |0\rangle = \sum_{r=1}^{d-1} a_r |r\rangle = |\psi_\beta\rangle, \quad (\text{B25})$$

which proves (B14). As $\beta_r \neq 0$ by Lemma 4, we have $a_r > 0$ for all $r = 1, \dots, d-1$, so each level $|r\rangle$ must receive nonzero amplitude. Hence the construction requires (and uses) exactly $d-1$ nontrivial embedded two-level rotations. ■

- [1] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [2] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [3] J. Preskill, Simulating quantum field theory with a quantum computer, *PoS LATTICE2018*, 024 (2018).
- [4] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, *Rev. Mod. Phys.* **86**, 153 (2014).
- [5] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, Quantum chemistry in the age of quantum computing, *Chem. Rev. (Washington, DC)* **119**, 10856 (2019).
- [6] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, *npj Comput. Mater.* **6**, 85 (2020).
- [7] M. Troyer and U.-J. Wiese, Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations, *Phys. Rev. Lett.* **94**, 170201 (2005).
- [8] C. Gatteringer and K. Langfeld, Approaches to the sign problem in lattice field theory, *Int. J. Mod. Phys. A* **31**, 1643007 (2016).

- [9] K. Nagata, Finite-density lattice QCD and sign problem: Current status and open problems, *Prog. Part. Nucl. Phys.* **127**, 103991 (2022).
- [10] H. F. Trotter, On the product of semi-groups of operators, *Proc. Am. Math. Soc.* **10**, 545 (1959).
- [11] M. Suzuki, Fractal decomposition of exponential operators with applications to many-body theories and Monte Carlo simulations, *Phys. Lett. A* **146**, 319 (1990).
- [12] A. M. Childs and Y. Su, Nearly optimal lattice simulation by product formulas, *Phys. Rev. Lett.* **123**, 050503 (2019).
- [13] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, *Quantum Inf. Comput.* **12**, 901 (2012).
- [14] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating Hamiltonian dynamics with a truncated Taylor series, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [15] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
- [16] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st*

- Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)* (Association for Computing Machinery, New York, NY, USA, 2019), pp. 193–204.
- [17] S. P. Jordan, K. S. M. Lee, and J. Preskill, Quantum algorithms for quantum field theories, *Science* **336**, 1130 (2012).
 - [18] N. Klco and M. J. Savage, Digitization of scalar fields for quantum computing, *Phys. Rev. A* **99**, 052335 (2019).
 - [19] E. Zohar, A. Farace, B. Reznik, and J. I. Cirac, Digital lattice gauge theories, *Phys. Rev. A* **95**, 023604 (2017).
 - [20] 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).
 - [21] N. Klco, J. R. Stryker, and M. J. Savage, SU(2) non-Abelian gauge field theory in one dimension on digital quantum computers, *Phys. Rev. D* **101**, 074512 (2020).
 - [22] Y. Ji, H. Lamm, and S. Zhu, Gluon field digitization via group space decimation for quantum computers, *Phys. Rev. D* **102**, 114513 (2020).
 - [23] M. Illa, C. E. P. Robin, and M. J. Savage, Qu8its for quantum simulations of lattice quantum chromodynamics, *Phys. Rev. D* **110**, 014507 (2024).
 - [24] J. B. Kogut and L. Susskind, Hamiltonian formulation of Wilson's lattice gauge theories, *Phys. Rev. D* **11**, 395 (1975).
 - [25] S. Chandrasekharan and U.-J. Wiese, Quantum link models: A discrete approach to gauge theories, *Nucl. Phys. B* **492**, 455 (1997).
 - [26] 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).
 - [27] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero, C. A. Muschik, B. Reznik, E. Rico, L. Tagliacozzo, K. Van Acoleyen, F. Verstraete, U.-J. Wiese, M. Wingate, J. Zakrzewski, and P. Zoller, Simulating lattice gauge theories within quantum technologies, *Eur. Phys. J. D* **74**, 165 (2020).
 - [28] A. Macridin, P. Spentzouris, J. Amundson, and R. Harnik, Electron-phonon systems on a universal quantum computer, *Phys. Rev. Lett.* **121**, 110504 (2018).
 - [29] 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).
 - [30] E. A. Martinez, C. A. Muschik, P. Schindler, D. Nigg, A. Erhard, M. Heyl, P. Hauke, M. Dalmonte, T. Monz, P. Zoller, and R. Blatt, Real-time dynamics of lattice gauge theories with a few-qubit quantum computer, *Nature (London)* **534**, 516 (2016).
 - [31] F. M. Surace, P. P. Mazza, G. Giudici, A. Lerose, A. Gambassi, and M. Dalmonte, Lattice gauge theories and string dynamics in Rydberg atom quantum simulators, *Phys. Rev. X* **10**, 021041 (2020).
 - [32] 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).
 - [33] Y. Wang, Z. Hu, B. C. Sanders, and S. Kais, Qudits and high-dimensional quantum computing, *Front. Phys.* **8**, 589504 (2020).
 - [34] D. Gottesman, Fault-tolerant quantum computation with higher-dimensional systems, in *Quantum Computing and Quantum Communications*, edited by C. P. Williams, Lecture Notes in Computer Science, Vol. 1509 (Springer, Berlin, Heidelberg, 1998), pp. 302–313.
 - [35] E. Hostens, J. Dehaene, and B. De Moor, Stabilizer states and Clifford operations for systems of arbitrary dimensions, and modular arithmetic, *Phys. Rev. A* **71**, 042315 (2005).
 - [36] J. M. Farinholt, An ideal characterization of the Clifford operators, *J. Phys. A: Math. Theor.* **47**, 305303 (2014).
 - [37] E. T. Campbell, H. Anwar, and D. E. Browne, Magic-state distillation in all prime dimensions using quantum Reed-Muller codes, *Phys. Rev. X* **2**, 041021 (2012).
 - [38] E. O. Kiktenko, A. S. Nikolaeva, and A. K. Fedorov, *Colloquium*: Qudits for decomposing multiqubit gates and realizing quantum algorithms, *Rev. Mod. Phys.* **97**, 021003 (2025).
 - [39] 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).
 - [40] M. Neeley, M. Ansmann, R. C. Bialczak, M. Hofheinz, E. Lucero, A. D. O'Connell, D. Sank, H. Wang, J. Wenner, A. N. Cleland, and J. M. Martinis, Emulation of a quantum spin with a superconducting phase qudit, *Science* **325**, 722 (2009).
 - [41] L. B. Nguyen, N. Goss, K. Siva, Y. Kim, E. Younis, B. Qing, A. Hashim, D. I. Santiago, and I. Siddiqi, Empowering a qudit-based quantum processor by traversing the dual bosonic ladder, *Nat. Commun.* **15**, 7117 (2024).
 - [42] M. Erhard, M. Krenn, and A. Zeilinger, Advances in high-dimensional quantum entanglement, *Nat. Rev. Phys.* **2**, 365 (2020).
 - [43] S. Bravyi and A. Kitaev, Universal quantum computation with ideal Clifford gates and noisy ancillas, *Phys. Rev. A* **71**, 022316 (2005).
 - [44] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
 - [45] D. Litinski, A game of surface codes: Large-scale quantum computing with lattice surgery, *Quantum* **3**, 128 (2019).
 - [46] A. Katabarwa, K. Gratsea, A. Caesura, and P. D. Johnson, Early fault-tolerant quantum computing, *PRX Quantum* **5**, 020101 (2024).
 - [47] N. J. Ross and P. Selinger, Optimal ancilla-free Clifford + T approximation of Z rotations, *Quantum Inf. Comput.* **16**, 901 (2016).
 - [48] V. Kliuchnikov, D. Maslov, and M. Mosca, Fast and efficient exact synthesis of single-qubit unitaries generated by Clifford and T gates, *Quantum Inf. Comput.* **13**, 607 (2013).
 - [49] D. M. Kurkcuoglu, H. Lamm, and A. Maestri, Qudit gate decomposition dependence for lattice gauge theories, [arXiv:2410.16414](https://arxiv.org/abs/2410.16414).
 - [50] R. Joshi, M. Meth, J. C. Louw, J. J. Osborne, K. Mato, M. Ringbauer, and J. C. Halimeh, Efficient qudit circuit for quench dynamics of $2 + 1D$ quantum link electrodynamics, [arXiv:2507.12589](https://arxiv.org/abs/2507.12589).
 - [51] R. Joshi, J. C. Louw, M. Meth, J. J. Osborne, K. Mato, G.-X. Su, M. Ringbauer, and J. C. Halimeh, Probing hadron scattering in lattice gauge theories on qudit quantum computers, [arXiv:2507.12614](https://arxiv.org/abs/2507.12614).

- [52] S. P. Jordan, K. S. M. Lee, and J. Preskill, Quantum computation of scattering in scalar quantum field theories, *Quantum Inf. Comput.* **14**, 1014 (2014).
- [53] J. D. Watson, J. Bringewatt, A. F. Shaw, A. M. Childs, A. V. Gorshkov, and Z. Davoudi, Quantum algorithms for simulating nuclear effective field theories, [arXiv:2312.05344](https://arxiv.org/abs/2312.05344).
- [54] T. Byrnes and Y. Yamamoto, Simulating lattice gauge theories on a quantum computer, *Phys. Rev. A* **73**, 022328 (2006).
- [55] A. F. Shaw, P. Lougovski, J. R. Stryker, and N. Wiebe, Quantum algorithms for simulating the lattice Schwinger model, *Quantum* **4**, 306 (2020).
- [56] W. C. Chung, J. de Hond, J. Xiang, E. Cruz-Colón, and W. Ketterle, Tunable single-ion anisotropy in spin-1 models realized with ultracold atoms, *Phys. Rev. Lett.* **126**, 163203 (2021).
- [57] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, Cold bosonic atoms in optical lattices, *Phys. Rev. Lett.* **81**, 3108 (1998).
- [58] S. McArdle, A. Mayorov, X. Shan, S. Benjamin, and X. Yuan, Digital quantum simulation of molecular vibrations, *Chem. Sci.* **10**, 5725 (2019).
- [59] Y. Su, D. W. Berry, N. Wiebe, N. Rubin, and R. Babbush, Fault-tolerant quantum simulations of chemistry in first quantization, *PRX Quantum* **2**, 040332 (2021).
- [60] L. Spagnoli, C. Lissoni, and A. Roggero, Quantum simulation of nuclear dynamics in first quantization, [arXiv:2507.22814](https://arxiv.org/abs/2507.22814).
- [61] E. J. Gustafson, Prospects for simulating a qudit-based model of $(1 + 1)$ D scalar QED, *Phys. Rev. D* **103**, 114505 (2021).
- [62] S. Prakash, A. R. Kalra, and A. Jain, A normal form for single-qudit Clifford + T operators, *Quantum Inf. Process.* **20**, 341 (2021).
- [63] K. L. Manders and L. Adleman, NP-complete decision problems for binary quadratics, *J. Comput. Syst. Sci.* **16**, 168 (1978).
- [64] E. J. Gustafson, H. Lamm, D. Liu, E. M. Murairi, and S. Zhu, Synthesis of single-qutrit circuits from Clifford + R gates, *Phys. Rev. A* **112**, 062414 (2025).
- [65] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of Trotter error with commutator scaling, *Phys. Rev. X* **11**, 011020 (2021).
- [66] A. W. Harrow, B. Recht, and I. L. Chuang, Efficient discrete approximations of quantum gates, *J. Math. Phys.* **43**, 4445 (2002).
- [67] V. Kliuchnikov, K. Lauter, R. Minko, A. Paetznick, and C. Petit, Shorter quantum circuits via single-qubit gate approximation, *Quantum* **7**, 1208 (2023).
- [68] A. Hardy, P. Mukhopadhyay, M. S. Alam, R. Konik, L. Hormozi, E. Rieffel, S. Hadfield, J. A. Barata, R. Venugopalan, D. E. Kharzeev, and N. Wiebe, Scattering processes from quantum simulation algorithms for scalar field theories, *PRX Quantum* **7**, 010343 (2026).
- [69] C. Jones, Low-overhead constructions for the fault-tolerant Toffoli gate, *Phys. Rev. A* **87**, 022328 (2013).
- [70] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, Cambridge, 2018).
- [71] S. Chakraborty, Implementing any linear combination of unitaries on intermediate-term quantum computers, *Quantum* **8**, 1496 (2024).
- [72] M. Karácsony, L. Oroszlány, and Z. Zimborás, Efficient qudit based scheme for photonic quantum computing, *SciPost Phys. Core* **7**, 032 (2024).
- [73] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, Transformation of quantum states using uniformly controlled rotations, *Quantum Inf. Comput.* **5**, 467 (2005).
- [74] V. Bergholm, J. J. Vartiainen, M. Möttönen, and M. M. Salomaa, Quantum circuits with uniformly controlled one-qubit gates, *Phys. Rev. A* **71**, 052330 (2005).
- [75] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Encoding electronic spectra in quantum circuits with linear T complexity, *Phys. Rev. X* **8**, 041015 (2018).
- [76] A. Bocharov, A note on optimality of quantum circuits over metaplectic basis, *Quantum Inf. Comput.* **18**, 0001 (2018).
- [77] P. Selinger, Efficient Clifford + T approximation of single-qubit operators, *Quantum Inf. Comput.* **15**, 159 (2015).
- [78] J. E. Moussa, Quantum circuits for qubit fusion, *Quantum Inf. Comput.* **16**, 1113 (2016).
- [79] J. Zhang, Y. Meurice, and S.-W. Tsai, Truncation effects in the charge representation of the $O(2)$ model, *Phys. Rev. B* **103**, 245137 (2021).
- [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, *Phys. Rev. D* **107**, L031503 (2023).
- [81] D. Paulson, L. Dellantonio, J. F. Haase, A. Celi, A. Kan, A. Jena, C. Kokail, R. van Bijnen, K. Jansen, P. Zoller, and C. A. Muschik, Simulating 2D effects in lattice gauge theories on a quantum computer, *PRX Quantum* **2**, 030334 (2021).
- [82] M. Meth, J. F. Haase, J. Zhang, C. Edmunds, L. Postler, A. Steiner, A. J. Jena, L. Dellantonio, R. Blatt, P. Zoller, T. Monz, P. Schindler, C. Muschik, and M. Ringbauer, Simulating two-dimensional lattice gauge theories on a qudit quantum computer, *Nat. Phys.* **21**, 570 (2025).
- [83] L. Spagnoli, A. Roggero, and N. Wiebe, Qudit stabiliser codes for $\mathbb{Z}_n$ lattice gauge theories with matter, [arXiv:2602.20661](https://arxiv.org/abs/2602.20661).
- [84] [http://energy.gov/downloads/doe-public-access-plan](https://energy.gov/downloads/doe-public-access-plan).