

Symmetry-Accelerated Classical Simulation of Clifford-Dominated Circuits

Giulio Camillo,^{1,*} Filipa C.R. Peres^{1,†}, Markus Heinrich^{2,‡} and Juani Bermejo-Vega^{1,3}

¹*Quantum Thermodynamics and Computation Group, Electromagnetism and Matter Physics Department, University of Granada, Granada, Spain*

²*Institute for Theoretical Physics, University of Cologne, Cologne, Germany*

³*Institute Carlos I for Theoretical and Computational Physics, University of Granada, Granada, Spain*



(Received 12 December 2025; accepted 23 April 2026; published 10 June 2026)

Classical simulation of quantum circuits plays a crucial role in validating quantum hardware and delineating the boundaries of quantum advantage. Among the most effective simulation techniques are those based on the stabilizer extent, which quantifies the overhead of representing non-Clifford operations as linear combinations of Clifford unitaries. However, finding optimal decompositions rapidly becomes intractable as it constitutes a superexponentially large optimization problem. In this work, we exploit symmetries in the computation of the stabilizer extent, proving that for real, diagonal, and real-diagonal unitaries, the optimization can be restricted to the corresponding subgroups of the Clifford group without loss of optimality. This “strong symmetry reduction” drastically reduces computational cost, enabling optimal decompositions of unitaries on up to seven qubits using a standard laptop—far beyond previous two-qubit limits. Additionally, we employ a “weak symmetry reduction” method that leverages additional invariances to shrink the search space further. Applying these results, we demonstrate exponential run-time improvements in classical simulations of quantum Fourier transform circuits and measurement-based quantum computations on the Union Jack lattice, as well as insights into the nonstabilizer properties of multicontrolled phase gates and unitaries generating hypergraph states. Our findings establish symmetry exploitation as a powerful route to scale classical simulation techniques and deepen the resource-theoretic understanding of quantum advantage.

DOI: [10.1103/mnzz-481j](https://doi.org/10.1103/mnzz-481j)

I. INTRODUCTION

The central promise of quantum computing is the solution of problems that are out of reach for classical machines. While we thus expect quantum computers to be generally hard to simulate by fully classical means, research into classical simulation algorithms and techniques has proven crucial in developing a better understanding of the precise onset of such a quantum advantage. Indeed, classical simulators allow direct validation of quantum hardware and serve as a benchmark to compare quantum advantage claims against [1–3], while still regularly defying them [4,5]. Improved simulation techniques

thus also raise the bar for better and bigger quantum prototypes that are and will be arising. Moreover, classical simulation provides valuable insights into which quantum resources and phenomena enable the desired quantum advantage. In turn, this can inform the development of more efficient quantum algorithms.

The prototypical example of such a deeper understanding provided by classical simulation is the Gottesman-Knill theorem: It states that even highly entangling quantum circuits—so-called *stabilizer circuits*—are efficiently classically simulable as long as they only involve Clifford gates as well as preparation and measurement in the computational basis [6]. This shows that, besides entanglement, *nonstabilizerness* (also known as *magic*) is a necessary resource for quantum advantage. Strikingly, this resource can be provided by a single non-Clifford gate [7]—often chosen as the *T* gate—or, alternatively, by a continuous supply of *magic states* [8].

The simulation of stabilizer circuits can be extended to arbitrary quantum circuits at the cost of an exponential overhead in the number of nonstabilizer elements (e.g., magic states or non-Clifford gates) and many works have successively developed more and more efficient algorithms

*G.C. and F.C.R.P. contributed equally.

†Contact author: fcperes@onsager.ugr.es

‡Contact author: markus.heinrich@uni-koeln.de

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

to this end [9–14]. These approaches typically decompose a quantum circuit into a linear combination of stabilizer circuits; doing so on the level of the whole circuit is, however, generally computationally intractable. Indeed, previous works have instead been limited to the decomposition of single or two-qubit gates. This makes the runtime of these approaches very sensitive to compilation efficiency, in particular, the ability to minimize nonstabilizer elements in the circuit, but also blind to known submultiplicative effects in nonstabilizer resources: The nonstabilizerness of a circuit is almost always strictly less than the product of its parts. Taken together, these issues can severely limit the applicability of stabilizer-based simulation algorithms.

In this work, we address this problem by exploiting the concept of *symmetry*, providing methods that enable us to exactly and optimally decompose entire nonstabilizer subcircuits (on up to seven qubits). This allows us to show that submultiplicative effects in stabilizer-based simulation are, in fact, drastic: Our methods lead to savings of several orders of magnitude in runtime for paradigmatic circuits.

Concretely, we consider simulation algorithms based on the *stabilizer extent*, which rank among the most successful and best studied approaches [13–17]. By exploiting the symmetry of highly relevant classes of non-Clifford unitaries, we achieve significant reductions in the size of the underlying optimization problem. More specifically, we prove that if a unitary U is real and/or diagonal, we can find an optimal decomposition of U in terms of real and/or diagonal Clifford unitaries only, thus avoiding full optimization over the much larger Clifford group. Strikingly, such a “strong” reduction is not possible for every symmetry. We can, however, exploit any other symmetries of U to perform a “weak” reduction that still leads to a further and strong decrease of the search space.

Importantly, such symmetries are often present in multi-qubit gates or subcircuits that appear as building blocks of quantum algorithms. Examples of this are multicontrolled rotation and phase gates, Toffoli gates, Givens rotations, or the Grover diffusion operator.

We use these findings to accelerate classical simulation based on the sum-over-Cliffords method [13]. By grouping multiple non-Clifford gates and jointly decomposing them, we demonstrate large savings over the simulation of individual gates due to submultiplicative effects [15]. Using our symmetry reduction techniques, we can compute such decompositions for symmetric unitaries supported on at most seven qubits using a consumer laptop, significantly exceeding the previous state of the art, limited to two qubits [15]. We explicitly demonstrate these improvements in the simulation of quantum Fourier transform (QFT) circuits and of (arbitrary, universal) measurement-based quantum computations (MBQCs). To list some examples, a 16-qubit QFT circuit can be simulated two orders of magnitude faster using our results than using prior state-of-the-art stabilizer-based techniques. Additionally, our results

enable an exponential improvement of the classical (weak) simulation of any (universally general) MBQC performed on an n -qubit Union Jack lattice and using only Pauli measurements. From a resource-theoretical perspective, our work allows the in-depth study of the nonstabilizer properties of important families of circuits using the stabilizer extent. Along this line, we make an extensive analysis of the properties of multicontrolled phase gates $C^{n-1}P(\theta)$ as well as of the unitaries generating hypergraph and generalized hypergraph quantum states, thoroughly investigating submultiplicativity effects and also the role played by entanglement in the optimal decompositions.

Finally, let us outline the structure of this paper. Section II gives a brief introduction to the relevant background. In Sec. III, we present our theoretical results on the strong and weak symmetry reduction of the stabilizer extent and provide examples for when strong reduction is possible and when it is not. Section IV is concerned with practical demonstrations of the benefits of our theoretical results: We analyze the nonstabilizerness of multicontrolled phase gates and show how our techniques lead to exponential improvements in the sum-over-Cliffords simulation of QFT circuits, as well as in MBQC using hypergraph states. In Sec. V, we conclude with a discussion of our most important results and give an outlook on the impact of this paper and prospective future research.

II. BACKGROUND

A. Classical simulation and the stabilizer extent

Stabilizer-based classical simulation algorithms stem from the fact that generic circuits can always be expressed over a Clifford + K gate set, where K stands for a single or multiple non-Clifford gate(s). These non-Clifford unitaries are then simulated using stabilizer operations only. Since the latter are by themselves efficiently simulable by Gottesman-Knill, the runtime overhead comes entirely from the simulation of the non-Clifford circuit components.

Most of the non-Clifford gates that are typically considered, for instance, the T , CCZ , and Toffoli gates, can be implemented using a circuit gadget that consumes a resource state—called *magic state* in this context. Since the gadget itself is composed of stabilizer operations, the difficulty in simulating those non-Clifford gates is then given by the simulation of the magic states. To this end, we can express the magic state as a linear combination of stabilizer states $|s\rangle$, $|\psi\rangle = \sum_s x_s |s\rangle$, and use the linearity of the final state’s amplitudes to reduce their computation to a series of stabilizer circuit simulations and suitable preprocessing and postprocessing [11, 13].

An alternative, more general approach is given by the sum-over-Cliffords simulation [13], wherein the non-Clifford unitaries are directly decomposed into a linear combination of Clifford unitaries C , say $U = \sum_C x_C C$.

In perfect analogy with the gadget-based approach, we can then use the Gottesman-Knill theorem to simulate the stabilizer circuits resulting from the decomposition and suitably postprocess the results. This method has several advantages over gadget-based simulation: First, it is applicable to arbitrary non-Clifford unitaries and not only those enabled by magic-state injection; secondly, it is simpler to implement; and, finally, it often involves the manipulation of quantum states with fewer qubits than those in the gadget-based approach, making it faster [13].

In both gadget-based and sum-over-Cliffords simulation, the simulation runtime is proportional to the number of terms in the (nonunique) decomposition of the target states or unitaries. Thus, the simulation works best with a minimal number of terms, also known as the *stabilizer rank* of the target [10,11]. The stabilizer rank can, in principle, be computed by minimizing the so-called ℓ_0 -norm of the coefficient vector $\mathbf{x}$, a problem that is, however, known to be NP-complete [18, Appendix A6]. In practice, even an approximate low-rank decomposition would be sufficient. Bravyi *et al.* [13] show how to probabilistically construct a δ -approximate decomposition with rank $\|\mathbf{x}\|_1^2/\delta^2$ from an exact decomposition with coefficient vector $\mathbf{x}$. Hence, finding decompositions with small ℓ_1 -norm is central to this approach. Importantly, this problem can now be phrased as conic programming and can be solved using standard algorithms. This motivates the definition of the *stabilizer extent* for pure quantum states and matrices as follows.

Definition 1 (Stabilizer extent for pure states [13]). Let STAB_n denote the set of n -qubit pure stabilizer states. The *stabilizer extent* of an arbitrary, n -qubit quantum state $|\psi\rangle$ is defined as

$$\xi(|\psi\rangle) := \min \left\{ \|\mathbf{x}\|_1^2 \mid \mathbf{x} \in \mathbb{C}^{|\text{STAB}_n|} : |\psi\rangle = \sum_{|s\rangle \in \text{STAB}_n} x_s |s\rangle \right\}.$$

We call a decomposition of $|\psi\rangle$ into stabilizer states *optimal* if it is the minimizer of the above optimization problem.

Definition 2 (Stabilizer extent for matrices [13]). Let $\mathcal{C}_n$ be the n -qubit Clifford group. We define the *stabilizer extent* of a matrix $M \in \mathbb{C}^{2^n \times 2^n}$ as

$$\xi(M) := \min \left\{ \|\mathbf{x}\|_1^2 \mid \mathbf{x} \in \mathbb{C}^{|\mathcal{C}_n|} : M = \sum_{C \in \mathcal{C}_n} x_C C \right\}.$$

We again call a decomposition of M optimal if it has $\|\mathbf{x}\|_1^2 = \xi$.

The stabilizer extent has many useful properties that often allow for simplifying or avoiding its computation and finding (sub)-optimal decompositions from known ones. A powerful property of the state version is its multiplicativity under tensor products of few-qubit quantum states [13],

$$\xi \left(\bigotimes_{j=1}^m |\psi_j\rangle \right) = \prod_{j=1}^m \xi(|\psi_j\rangle), \quad (1)$$

where each $|\psi_j\rangle$ is supported on at most three qubits. This result significantly facilitates the runtime estimation of classical simulation in the magic state model, since typically considered magic states are indeed such few-qubit states. Assuming, for instance, that the circuit of interest contains k T gates, the runtime cost of simulating it through gadgetization is given by $\xi(|T\rangle^{\otimes k}) = \xi(|T\rangle)^k = \cos(\pi/8)^{-2k}$ where $|T\rangle := T|+\rangle = (|0\rangle + e^{i\pi/4}|1\rangle)/\sqrt{2}$ [13]. We note that the stabilizer extent is generally submultiplicative for states on more than three qubits [19], thus requiring the direct solution of the defining optimization problem.

Notably, it is not known whether the matrix version of the stabilizer extent is multiplicative under tensor products of few-qubit unitaries. In some cases, we can exploit the following relation between the state and matrix version, which follows directly from the definitions

$$\xi(U|s\rangle) \leq \xi(U). \quad (2)$$

Here, $|s\rangle$ is an arbitrary stabilizer state. An important case is given by a diagonal U and the state $|U\rangle = U|+\rangle^{\otimes n}$. Although these are essentially the same objects, equality in Eq. (2) is only known in certain cases, for instance, if U is a single-qubit Z -rotation, $R_Z(\theta) = e^{-i\frac{\theta}{2}Z}$ (of which the T gate is a special case) or the CCZ gate. In these cases, we can exploit the multiplicativity of the stabilizer extent for states to conclude that $\xi(T^{\otimes k}) = \xi(T)^k = \cos(\pi/8)^{-2k}$ and $\xi(CCZ^{\otimes k}) = \xi(CCZ)^k = (16/9)^k$ [13].

Despite the possible lack of multiplicativity, the stabilizer extent for matrices has several other important properties which we summarize in the following proposition and prove in Appendix A for completeness.

Proposition 1. The stabilizer extent has the following properties:

- (i) $\xi(A) = 1$ for any Clifford unitary $A \in \mathcal{C}_n$.
- (ii) Invariance under multiplication by Clifford unitaries: $\xi(A) = \xi(CA) = \xi(AC') \quad \forall C, C' \in \mathcal{C}_n$.
Monotonicity under stabilizer code projectors Π : $\xi(\Pi A) \leq \xi(A)$, $\xi(A\Pi) \leq \xi(A)$.
- (iii) Submultiplicativity under matrix multiplication and tensor products: $\xi(AB) \leq \xi(A)\xi(B)$, $\xi(A \otimes B) \leq \xi(A)\xi(B)$. Moreover: $\xi(A \otimes B) = \xi(B \otimes A)$.

- (iv) Convexity: $\xi(tA + (1-t)B) \leq t\xi(A) + (1-t)\xi(B)$ for all $t \in [0, 1]$.
 Homogeneity: $\xi(xA) = |x|^2 \xi(A)$ for $x \in \mathbb{C}$.

We note that the stabilizer extent also provides a faithful measure of nonstabilizerness or magic: That is, $\xi(|\psi\rangle) = 1$ if and only if $|\psi\rangle$ is a stabilizer state, and otherwise $\xi(|\psi\rangle) > 1$. The same holds for the matrix version if we restrict to unitaries: $\xi(U) = 1$ iff $U \in \mathcal{C}_n$ and $\xi(U) > 1$ otherwise. For nonunitary matrices, however, ξ can assume values smaller than 1. The monotonicity then makes the stabilizer extent a *magic monotone*, which can be used for lower bounding the cost of (exact) synthesis of certain unitaries [12,20]. Indeed, combining the properties of the stabilizer extent in Proposition 1 leads to the following statement.

Proposition 2 (Exact gate synthesis). For any unitary $U \in \mathcal{U}(2^n)$, there always exists an s so that $\xi(T)^{s-1} < \xi(U) \leq \xi(T)^s$. If U is *exactly synthesizable* in terms of the Clifford + T gate set, then U requires at least s T gates to be synthesized.

We use this proposition to discuss some of our results in Sec. IV. However, since this is by no means central to our work, we leave the proof to Appendix B.

Given a quantum circuit over a Clifford + K gate set, the above-discussed properties also allow us to construct the necessary decompositions for a classical simulation algorithm, the runtime of which can, however, depend highly on the precise construction. One crucial reason for this is the submultiplicativity under matrix multiplication, cf. Proposition 1(iii). As we will also see later in Sec. IV, it can be highly beneficial to “block” non-Clifford gates and try to decompose their matrix product instead of decomposing them one-by-one. For this approach, we can typically only rely on the direct computation of the stabilizer extent by solving the defining optimization problem. Although the latter is a second-order cone program (SOCP) which can be solved in polynomial time in the problem size, the superexponentially growing size of STAB_n and $\mathcal{C}_n$ with n prevents us from doing so in practice. Even small instances require supercomputing resources, cf. Table X in Appendix F. Moreover, the matrix version is generically harder than the state version due to the square increase in the dimension of the underlying vector space and the fact that the Clifford group is superexponentially larger than the size of the set of stabilizer states,

$$\frac{|\mathcal{C}_n|}{|\text{STAB}_n|} = \frac{2^{n^2+2n} \prod_{j=1}^n (4^j - 1)}{2^n \prod_{j=1}^n (2^j + 1)} = \mathcal{O}(2^{n^2}).$$

While impressive progress has been made to speed up the computation of the stabilizer extent for states up to 10

qubits, e.g., by exploiting linear dependencies between stabilizer states [21] or suitably restricting the set of stabilizer states used for optimization [22], the matrix version has received little attention up to now. In this work, we address this gap in the literature and study how symmetries can be exploited to accelerate the computation of the stabilizer extent for unitaries in many important cases.

B. Symmetry and invariance

Central to our work are the notions of *symmetry* and *invariance*. Given a matrix M , we say that it is invariant under a map ϕ if $\phi(M) = M$. In the following, we focus on invertible maps that are linear or antilinear and preserve the Hilbert-Schmidt inner product, that is, $\text{tr}(\phi(A)^\dagger \phi(B)) = \text{tr}(A^\dagger B)$. We call such maps *isometries*. The set of isometries that leave M invariant forms a group, henceforth called the *symmetry group* of M . Given such a group of isometries $\mathcal{G}$, the set of invariant matrices forms a subspace, the so-called *$\mathcal{G}$ -invariant subspace*. Note that if $\mathcal{G}$ contains antilinear isometries, then this is only a real subspace. In the following, we shall focus on finite groups, although the generalization to compact groups is evident. It is an elementary fact that the projection onto the $\mathcal{G}$ -invariant subspace is given by

$$\Pi_{\mathcal{G}}(M) := \frac{1}{|\mathcal{G}|} \sum_{\phi \in \mathcal{G}} \phi(M).$$

Again, if $\mathcal{G}$ contains antilinear isometries, then $\Pi_{\mathcal{G}}$ is only $\mathbb{R}$ -linear, not $\mathbb{C}$ -linear.

The main symmetry groups we consider in this work are the following:

Complex conjugation: Let $\mathcal{K}_n = \{\text{id}, (\cdot)^*\}$ be the group generated by complex conjugation $M \mapsto M^*$ in the computational basis. Its invariant subspace is the real subspace of matrices with real entries, $\mathbb{R}^{2^n \times 2^n}$.

Transposition: Let $\mathcal{T}_n = \{\text{id}, (\cdot)^T\}$ be the group generated by transposition $M \mapsto M^T$ in the computational basis. Its invariant subspace is the complex subspace of symmetric matrices, $\text{Sym}_{2^n}(\mathbb{C})$.

Diagonal Pauli group: Let $\mathcal{Z}_n = \{Z(z) := Z^{z_1} \otimes \cdots \otimes Z^{z_n} \mid z \in \mathbb{Z}_2^n\}$ be the group of diagonal Pauli operators, i.e., the one generated by Pauli Z matrices on every qubit. We can identify $\mathcal{Z}_n$ with the group of quantum channels $M \mapsto Z(z)MZ(z)^\dagger$, the invariant subspace of which is the complex subspace of diagonal matrices.

Permutation group: Let $\mathcal{S}_n$ be the group of qubit permutations generated by all possible SWAP gates. Again, by a slight abuse of notation, we consider the invariant subspace under conjugation with permutations, $M \mapsto \pi M \pi^\dagger$.

We are furthermore interested in the subset of Clifford unitaries $\mathcal{C}_n^{\mathcal{G}}$ that are invariant under a group $\mathcal{G}$. If $\mathcal{G}$ acts

by conjugation, such as for complex conjugation, the diagonal Pauli group, and the permutation group, then $\mathcal{C}_n^{\mathcal{G}}$ is a subgroup of the Clifford group. In contrast, $\mathcal{C}_n^{\mathcal{T}_n}$ is not a subgroup as transposition-invariant Cliffords do not necessarily commute [23]. The subgroups $\mathcal{C}_n^{\mathcal{K}_n}$ and $\mathcal{C}_n^{\mathcal{Z}_n}$ are known as the *real Clifford group* $\mathcal{R}_n$ and *diagonal Clifford group* $\mathcal{D}_n$, respectively. The former is generated by the Pauli Z, CNOT, and Hadamard gates, while the latter is generated by the phase gate S and the controlled-Z gate CZ. We will also consider the group $\mathcal{K}_n \times \mathcal{Z}_n$ generated by the commuting groups $\mathcal{K}_n$ and $\mathcal{Z}_n$. The invariant Clifford subgroup is then the group of real-diagonal Cliffords $\mathcal{RD}_n$, generated by Z and CZ.

III. RESULTS

A. Computation of stabilizer extent under symmetries

In this section, we show that the computation of the stabilizer extent for unitaries that are invariant under certain symmetry groups $\mathcal{G}$ can be reduced to an optimization problem over $\mathcal{C}_n^{\mathcal{G}}$ only, which is typically of a much smaller size than the original one. To this end, we define the $\mathcal{G}$ -symmetric stabilizer extent as follows.

Definition 3 ($\mathcal{G}$ -symmetric stabilizer extent). Take $\mathcal{G}$ a symmetry group and suppose that $\mathcal{C}_n^{\mathcal{G}}$ spans the (real or complex) $\mathcal{G}$ -invariant subspace. Then, the $\mathcal{G}$ -symmetric stabilizer extent of a $\mathcal{G}$ -invariant matrix $M \in \mathbb{C}^{2^n \times 2^n}$ is defined as

$$\xi_{\mathcal{G}}(M) := \min \left\{ \|\mathbf{x}\|_1^2 \mid \mathbf{x} \in \mathbb{C}^{|\mathcal{C}_n^{\mathcal{G}}|} : M = \sum_{C \in \mathcal{C}_n^{\mathcal{G}}} x_C C \right\}.$$

In this work, we focus on the cases where $\mathcal{G} = \mathcal{K}_n, \mathcal{Z}_n, \mathcal{K}_n \times \mathcal{Z}_n$. Then, $\mathcal{C}_n^{\mathcal{G}}$ corresponds to the real, diagonal, and real-diagonal subgroups of the Clifford group, respectively. Since already the respective Pauli subgroups span the $\mathcal{G}$ -invariant subspaces in these cases, $\xi_{\mathcal{G}}$ is well-defined. In the $\mathcal{Z}_n$ and $\mathcal{K}_n \times \mathcal{Z}_n$ case, the matrix U and the Cliffords $C \in \mathcal{C}_n^{\mathcal{G}}$ are real, and we can thus restrict the expansion coefficients $\mathbf{x}$ to real numbers, turning the optimization problem into a linear program (LP).

For $\mathcal{G} = \mathcal{K}_n, \mathcal{Z}_n, \mathcal{K}_n \times \mathcal{Z}_n$, we then show that the stabilizer extent of a $\mathcal{G}$ -invariant unitary coincides with its $\mathcal{G}$ -symmetric stabilizer extent.

Theorem 1 (Strong $\mathcal{G}$ -symmetry reduction). Let $\mathcal{G} \in \{\mathcal{K}_n, \mathcal{Z}_n, \mathcal{K}_n \times \mathcal{Z}_n\}$. For any U such that $U = \Pi_{\mathcal{G}}(U)$, $\xi(U) = \xi_{\mathcal{G}}(U)$.

The proof of Theorem 1 relies on the following intermediate result, the proof of which we defer to Appendix C.

Lemma 1. Let $\mathcal{G} \in \{\mathcal{K}_n, \mathcal{Z}_n, \mathcal{K}_n \times \mathcal{Z}_n\}$, $C \in \mathcal{C}_n$, and $x \in \mathbb{C}$. Then, $\xi_{\mathcal{G}}(\Pi_{\mathcal{G}}(x_C)) \leq |x|^2$.

Proof of Theorem 1. Suppose U has an optimal decomposition $U = \sum_{C \in \mathcal{C}_n} x_C C$, i.e. $\xi(U) = \|\mathbf{x}\|_1^2$. We then have $U = \Pi_{\mathcal{G}}(U) = \sum_{C \in \mathcal{C}_n} \Pi_{\mathcal{G}}(x_C C)$, as $\Pi_{\mathcal{G}}$ is additive (but not necessarily $\mathbb{C}$ -linear). Next, let $\Pi_{\mathcal{G}}(x_C C) = \sum_{S \in \mathcal{C}_n^{\mathcal{G}}} y_S^{(C)} S$ be an optimal decomposition and write

$$U = \sum_{C \in \mathcal{C}_n} \sum_{S \in \mathcal{C}_n^{\mathcal{G}}} y_S^{(C)} S = \sum_{S \in \mathcal{C}_n^{\mathcal{G}}} \left(\sum_{C \in \mathcal{C}_n} y_S^{(C)} \right) S.$$

By Lemma 1, $\|\mathbf{y}^{(C)}\|_1 \leq |x_C|$, thus

$$\begin{aligned} \sqrt{\xi_{\mathcal{G}}(U)} &\leq \sum_{S \in \mathcal{C}_n^{\mathcal{G}}} \left| \sum_{C \in \mathcal{C}_n} y_S^{(C)} \right| \leq \sum_{C \in \mathcal{C}_n} \|\mathbf{y}^{(C)}\|_1 \\ &\leq \|\mathbf{x}\|_1 = \sqrt{\xi(U)}. \end{aligned}$$

Since $\xi(U) \leq \xi_{\mathcal{G}}(U)$ by definition, we have shown equality. ■

In Fig. 1, we see that, by itself, Theorem 1 allows a reduction of orders of magnitude in the number of elements required for the computation of the stabilizer extent, in particular for the cases of diagonal or real-diagonal unitaries. However, the invariant subgroups are still superexponentially large and we find, for instance, $|\mathcal{D}_n| > 10^8$ for $n = 6$, which is already too demanding for the optimizer. Fortunately, this result can be enhanced by leveraging other underlying symmetries of the nonstabilizer unitary of interest.

As an example, let us consider a diagonal unitary U . By Theorem 1, we know it admits an optimal decomposition so that $U = \sum_{D \in \mathcal{D}_n} x_D D$. Let us, in addition, assume that




FIG. 1. Order of the Clifford group C_n and subgroups of interest. The figure presents the size of the Clifford group C_n in solid blue, the real subgroup R_n in dashed orange, the diagonal subgroup D_n in dash-dotted green, and the real-diagonal subgroup in dotted red.

U is also invariant under another group of *linear* isometries $\mathcal{G}$. While Theorem 1 may not apply to this additional symmetry (cf. Sec. III B), we can still exploit the linearity to write

$$\Pi_{\mathcal{G}}(U) = U = \sum_{D \in \mathcal{D}_n} x_D \Pi_{\mathcal{G}}(D).$$

Hence, the optimization can be carried out over the set of projections $\Pi_{\mathcal{G}}(D)$. Coincidences among those can significantly reduce the required search space and thus the size of the optimization problem. An analogous strategy was previously used for other magic monotones [24]. Moreover, in the case of many important symmetry groups, each $\Pi_{\mathcal{G}}(D)$ is a convex combination of Clifford unitaries, allowing us to find not only the optimal value of the stabilizer extent but also a suitable decomposition in terms of Clifford unitaries. We formalize these ideas into the following lemma.

Lemma 2 (Weak $\mathcal{G}$ -symmetry reduction). Let U be an n -qubit unitary whose optimal decomposition can be found over a generic set of matrices, $\mathcal{S}$, so that $U = \sum_{C \in \mathcal{S}} x_C C$. Let $\mathcal{G}$ be a symmetry group of U : $\Pi_{\mathcal{G}}(U) = U$. Then, the stabilizer extent of U can be computed using only the set of different projections of the matrices in the set $\mathcal{S}$.

Unlike Theorem 1, this lemma is completely general. The gain associated with its application is upper-bounded by the number of orbits resulting from the projection of the elements of the original set of matrices. In Appendix D, we illustrate the computation of this upper bound for diagonal unitaries that are also permutation-invariant, i.e., when $\mathcal{G} = \mathcal{S}_n$. In Appendix E, we go one step further and fully enumerate the total number of nonequivalent projections of diagonal Cliffords onto the permutation invariant subspace. This leads to impressive reductions in the size of the search space for the SOCP. As an example, for $n = 6$, we reduce the number of terms from $\sim 1.3 \times 10^8$ to only $\sim 6.4 \times 10^5$, a mere 0.5% of the original number.

In our practical demonstrations, we focus on leveraging Theorem 1 for the case of diagonal and real-diagonal unitaries, since these lead to the strongest reductions as seen by Fig. 1. Additionally, as we will discuss, many important circuits use gates of this form. To conclude, we would like to take this opportunity to point out that our Theorem 1 establishes an important connection between the stabilizer extent of quantum states and the stabilizer extent of diagonal unitaries. Let U be a diagonal unitary. Our result guarantees that the optimal decomposition of U (with respect to the ℓ_1 -norm) can be obtained by decomposing U only in terms of diagonal Clifford unitaries D so that $U = \sum_{D \in \mathcal{D}_n} x_D D$. Let $|U\rangle = U|+\rangle^{\otimes n}$. Then, the quantum state $|U\rangle$ admits a decomposition of the form: $|U\rangle = \sum_{D \in \mathcal{D}_n} x_D |s_D\rangle$, where $|s_D\rangle = D|+\rangle^{\otimes n}$. Since there is a one-to-one correspondence between the diagonal gates

and the set of equatorial stabilizer states, this means that the problem of finding the optimal extent of U can be reduced to the problem of finding the optimal ℓ_1 -norm of a decomposition of $|U\rangle$ using only equatorial stabilizer states. Hence, if the methods proposed in Refs. [21,22] can be modified to find the optimal decomposition of magic states exclusively in terms of equatorial stabilizer states, they might provide a nice way of boosting the results demonstrated below.

B. Is strong symmetry reduction extendable to every symmetry group?

It could be tempting to assume that Theorem 1 is generalizable to include every possible symmetry group. However, for the $\mathcal{G}$ -symmetric stabilizer extent to be well-defined, the subset $\mathcal{C}_n^{\mathcal{G}}$ must span the entire $\mathcal{G}$ -invariant subspace, otherwise the expansion $M = \sum_{C \in \mathcal{C}_n^{\mathcal{G}}} x_C C$ is not guaranteed to exist (cf. Definition 3). The three symmetry groups $\mathcal{G} = \mathcal{K}_n, \mathcal{Z}_n, \mathcal{K}_n \times \mathcal{Z}_n$ for which we proved strong $\mathcal{G}$ -symmetry reduction all meet this condition. For other $\mathcal{G}$, however, this might not be the case and thus our result cannot be applied (see proofs in Appendix C). One example is $\mathcal{C}_n^{\mathcal{S}_n}$, which does not span the permutation-invariant subspace for $n > 3$ [25].

What if $\mathcal{C}_n^{\mathcal{G}}$ does not span the entire $\mathcal{G}$ -invariant subspace, but a specific gate U so that $U = \Pi_{\mathcal{G}}(U)$ can be expanded in terms of that set? In this restricted setting, could a similar result as the one provided by Theorem 1 apply? That is, could we state that $\tilde{\xi}_{\mathcal{G}}(U) = \xi(U)$ (where $\tilde{\xi}_{\mathcal{G}}$ is now restricted to the span of $\mathcal{C}_n^{\mathcal{G}}$)? We answer this question in the negative by providing a specific counterexample. In particular, we consider the important set of two-parameter fermionic simulation (fSim) gates [26], which have the following matrix representation in the computational basis,

$$\text{fSim}(\theta, \phi) := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & -i \sin(\theta) & 0 \\ 0 & -i \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & e^{-i\phi} \end{pmatrix}. \quad (3)$$

It is easily understood from the above representation that these gates enable continuous transformations, with θ controlling the swapping between $|01\rangle$ and $|10\rangle$ and ϕ introducing a phase for $|11\rangle$. Important particular examples of these gates are $CP(\phi) = \text{fSim}(0, -\phi)$ and $\text{iSWAP} = \text{fSim}(-\pi/2, 0)$; $\text{fSim}(\theta, 0)$ is a matchgate [27] and $\text{fSim}(\pi/2, \pi/6)$ was the two-qubit gate implemented in Google's quantum processor in their attempted quantum supremacy experiment [28]. Since these are only two-qubit gates, it is feasible to compute the extent using the entire Clifford group. The results are presented in Fig. 2(a).

Looking at Eq. (3), we can see that the fSim gates are invariant under transposition: $\text{fSim}(\theta, \phi) = \text{fSim}(\theta, \phi)^T$.




FIG. 2. Stabilizer extent of two-qubit $\text{fSim}(\theta, \phi)$ gates as a function of the parameters θ and ϕ . Color map depicting (a) the (true) stabilizer extent, $\xi(\text{fSim}(\theta, \phi))$, (b) the ratio between the minimum ℓ_1 -norm squared obtained only with the transposition-invariant Cliffords and the (true) stabilizer extent, $\tilde{\xi}_{T_2}(\text{fSim}(\theta, \phi)) / \xi(\text{fSim}(\theta, \phi))$, and (c) the ratio between the minimum ℓ_1 -norm squared obtained using the permutation-invariant Cliffords and the (true) stabilizer extent, $\tilde{\xi}_{S_2}(\text{fSim}(\theta, \phi)) / \xi(\text{fSim}(\theta, \phi))$. If the result in Theorem 1 were to hold for this particular example with the transposition and permutation invariances, the ratios depicted in (b) and (c) would have to be 1 for every pair (θ, ϕ) . That is eminently not the case, even though it is possible to expand $\text{fSim}(\theta, \phi)$ in terms of $\mathcal{C}_2^{T_2}$ and $\mathcal{C}_2^{S_2}$.

We can thus attempt to compute the optimal decomposition of $\text{fSim}(\theta, \phi)$ into symmetric Clifford gates $\mathcal{C}_2^{T_2}$ and compare the associated squared ℓ_1 -norm $\tilde{\xi}_{T_2}(\text{fSim}(\theta, \phi))$ to the stabilizer extent $\xi(\text{fSim}(\theta, \phi))$. Our numerical studies demonstrate that, while for some parameter ranges $\tilde{\xi}_{T_2}(\text{fSim}(\theta, \phi))$ corresponds to the true stabilizer extent, for the most part $\tilde{\xi}_{T_2}(\text{fSim}(\theta, \phi)) > \xi(\text{fSim}(\theta, \phi))$. That is, we fail to find the correct value of the stabilizer extent if we restrict to expansions over the set $\mathcal{C}_n^{T_2}$. Figure 2(b) provides a clear visualization of the results.

The fSim gates are also invariant under qubit permutations. We work similarly to the case of transposition invariance and consider decompositions over the permutation-invariant Clifford subgroup $\mathcal{C}_2^{S_2}$ with associated stabilizer extent $\tilde{\xi}_{S_2}(\text{fSim}(\theta, \phi))$. In this case, the results are even worse than the ones obtained with the transposition-invariant Clifford gates; in particular, while the optimization problem could always be solved, the agreement between ξ and $\tilde{\xi}_{S_2}$ is restricted to a smaller parameter region and the ratio $\tilde{\xi}_{S_2}/\xi$ achieves much larger values [see Fig. 2(c)].

IV. APPLICATIONS

A. Details of the practical implementation

As previously explained, this work aims to improve the solution of a particular SOCP. To this end, and before the optimization itself, we have to build the set of elements to be considered in the main constraint given to the optimizer. To generate the group of Clifford unitaries, we used the algorithm provided by Koenig and Smolin [29]. On the other hand, to generate the diagonal and real-diagonal subgroups, we constructed our own algorithm [30]. Both of these algorithms exploit a particular emergent structure from the isomorphism between $\mathcal{C}_n/\mathcal{P}_n$ and $\text{Sp}(2n)$, where

$\text{Sp}(2n)$ refers to the symplectic group of $2n \times 2n$ matrices (see, for example, Ref. [31]).

In our code, Lemma 2 is leveraged to reduce the size of the search space of the optimization problem for diagonal and real-diagonal unitaries that are also permutation-invariant. This is done by directly filtering the elements of the starting set of diagonal and real-diagonal Cliffords. Finally, to solve the optimization problem itself, we employed the third-party software Gurobi [32] on a laptop with an Intel(R) Core(TM) i5-155U processor and 32GB of RAM.

We verified our code by computing the stabilizer extent of $T^{\otimes k}$ for $k \in \{1, \dots, 6\}$, CS , $CS^{\otimes 2}$, CCZ , and $CCZ^{\otimes 2}$ and comparing the results against the known theoretical values by Bravyi *et al.* [13]. All of the instances yielded the correct outcomes. In the following sections, we present results for other important unitaries.

Our code is openly available on GitHub [30].

B. Characterization of multicontrolled phase gates

We start by calculating the stabilizer extent of the family of multicontrolled phase gates on n qubits, denoted $C^{n-1}P(\theta)$ and defined as

$$C^{n-1}P(\theta) = \text{diag}(1, 1, \dots, 1, e^{i\theta}) \in \mathbb{C}^{2^n}. \quad (4)$$

The results obtained for the stabilizer extent of this family of gates are presented in Fig. 3. We see that, as the number of (control) qubits increases, the distance between the curves starts to diminish. Moreover, the position of the absolute maxima progressively moves toward $\theta = \pi$ and the local minimum present at that position seems to become less pronounced. We conjecture that, in the limit $n \rightarrow \infty$, the local minimum disappears and the curve has a single maximum at $\theta_{\max} = \pi$.




FIG. 3. Stabilizer extent of the multicontrolled- $P(\theta)$ gates as a function of θ . The figure shows the stabilizer extent of the $C^{n-1}P(\theta)$ gates as defined in Eq. (4) for $n \in \{1, \dots, 6\}$. The background of the plot is divided into regions colored in different shades of blue, highlighting values of the stabilizer extent between $\xi(T^{s-1})$ and $\xi(T^s)$ for $s \in \{1, 7\}$ and allow us to immediately identify the minimum number of T gates needed to synthesize the different gates, as per Proposition 2.

In Table I, we present the explicit values of the stabilizer extent for the $C^{n-1}Z$ and $C^{n-1}S$ gates. We note that, up to $n = 4$, $\xi(C^{n-1}Z) < \xi(C^{n-1}S)$, but for $n \geq 5$ this relation is reversed. This is consistent with the previous conjecture that in the limit $n \rightarrow \infty$ the maximum occurs at $\theta_{\max} = \pi$. Additionally, the table also presents the maximum extent value associated with the gate family $C^{n-1}P(\theta)$ for each value of n and the corresponding rotation angle $\theta_{\max}$.

We use this set of gates to make interesting studies concerning the properties of submultiplicativity of the stabilizer extent under tensor product and composition. We go further by analyzing the role of entanglement in the decompositions of this important family of unitaries and conclude with comments on the cost of synthesizing some of these gates.

TABLE I. Stabilizer extent, ξ , for relevant gates. We present the value of the stabilizer extent for the (a) $C^{n-1}Z$ gates, (b) $C^{n-1}S$ gates, and (c) gates with maximal stabilizer extent for each n . The latter results were obtained by interpolation of the values used to construct the curves in Fig. 3.

(a)			(b)			(c)			
n	Gate	ξ	n	Gate	ξ	n	Gate	$\theta_{\max}$	ξ
1	Z	1	1	S	1	1	T	$\pi/4$	1.171 57
2	CZ	1	2	CS	1.6	2	CS	$\pi/2$	1.6
3	CCZ	1.777 78	3	C^2S	2.050 00	3	$C^2P(\theta_{\max})$	0.6476π	2.132 63
4	C^3Z	2.250 00	4	C^3S	2.312 50	4	$C^3P(\theta_{\max})$	0.7048π	2.500 00
5	C^4Z	2.506 94	5	C^4S	2.453 13	5	$C^4P(\theta_{\max})$	0.7288π	2.707 92
6	C^5Z	2.640 63	6	C^5S	2.525 78	6	$C^5P(\theta_{\max})$	0.7397π	2.817 74
7	C^6Z	2.708 77	7	C^6S	—	7	$C^6P(\theta_{\max})$	—	—

1. (Sub)Multiplicativity of the stabilizer extent under tensor product and composition

We numerically tested the (sub)multiplicativity of the stabilizer extent under tensor products and multiplicativity holds for all the tested instances. Namely, up to five tensor products of the same single-qubit $P(\theta)$, as well as tensor products of $CP(\theta)$ with other $CP(\theta)$ or the $C^2P(\theta)$ gate, and also of $P(\theta)$ and $C^3P(\theta)$, for $\theta \in [0, 2\pi]$; that is to say, we tested the multiplicativity of the curves appearing in Fig. 3. More generally, we also considered tensor products with elements having different angles randomly selected within that interval, including $CP(\theta_1)$ with $CP(\theta_2)$, $CP(\theta_1)$ with $C^2P(\theta_2)$, $P(\theta_1)$ with $C^3P(\theta_2)$, and $P(\theta_1)$ with multiple other phase gates. As we found no counter-example, the question of whether multiplicativity of the stabilizer extent holds for one-, two-, and three-qubit unitaries (as it does for quantum states) remains open.

Now, a particular subset of the family of multicontrolled phase gates, the $CP(2\pi/2^j)$ gates, is used in the textbook formulation of the QFT subroutine. We focus on these gates to analyze the submultiplicativity of the stabilizer extent under composition. In particular, we consider unitary blocks of the form

$$U_k = \prod_{j=2}^{k+1} CP_{1,j}(2\pi/2^j),$$

where the subscript indicates that the two-qubit gate is controlled in the first qubit and targets the j th qubit. Hence, each U_k is a $(k+1)$ -qubit unitary with k controlled-phase gates.

The results are depicted in Fig. 4. We consider different values of k ranging from 2 to 12. The blue circles represent the values obtained by computing the stabilizer extent of each two-qubit gate in U_k and multiplying the results. Contrastingly, the yellow shapes (squares, diamonds, and octagons) correspond to the values obtained by splitting each U_k into unitary blocks of five qubits (the largest blocks whose stabilizer extent can be computed using our




FIG. 4. Submultiplicativity of the controlled-phase gates in the quantum Fourier transform subroutine. The blue circles depict the values obtained by multiplying the stabilizer extent of the individual $\{CP_{1,j}(2\pi/2^j)\}_{j=2}^{k+1}$ gates. Each circle corresponds to the value obtained for the largest unitary block, U_k , in a QFT subroutine with $(k+1)$ qubits. In yellow, we present the results obtained when splitting each U_k into unitary blocks of the maximum size. Up to $k=4$, the results correspond to the true value of the stabilizer extent. For $5 \leq k \leq 8$, each U_k is split into two blocks whose stabilizer extent is multiplied. Finally, for $k \geq 9$, each U_k is split into three blocks.

code), calculating the stabilizer extent for each block, and multiplying the results. The figure perfectly captures the strict submultiplicativity of the stabilizer extent.

These results will be useful in Sec. IV C, when we analyze the remarkable improvements that our results provide to the simulation of QFT circuits using the naive sum-over-Cliffords approach.

2. The role of entanglement

We now study how entanglement impacts the value of the (squared) ℓ_1 -norm of stabilizer decompositions of $C^{n-1}Z$ gates. Because these gates are real and diagonal, Theorem 1 guarantees that an optimal decomposition can be found in terms of real-diagonal Clifford unitaries. Every $C \in \mathcal{RD}_n$ can be written as

$$C = \prod_{j=1}^n Z_j^{a_j} \prod_{k>j}^n CZ_{jk}^{b_{jk}},$$

where $a_j \in \{0,1\}$ determines whether a Z gate is acting on the j th qubit and $b_{jk} \in \{0,1\}$ establishes whether a CZ gate is being applied between qubits j and k . It is therefore easy to see that the number of CZ gates in a given n -qubit real-diagonal Clifford unitary is bounded between 0 and $n(n-1)/2$.

To gain a deeper understanding of the entangling structures underpinning the optimal decompositions of $C^{n-1}Z$ gates, we break down the real-diagonal Clifford unitaries into sets defined by the number of CZ gates. We denote these sets $\mathcal{C}_i$, where the subscript i denotes the number of CZ gates of the real-diagonal Clifford unitaries in the set. For example, $\prod_{j=1}^n Z_j^{a_j} \in \mathcal{C}_0$, every $\prod_{j=1}^n Z_j^{a_j} CZ_{kl}$ with $l > k$ is in $\mathcal{C}_1$, every $\prod_{j=1}^n Z_j^{a_j} CZ_{kl} CZ_{mp}$ with $l > k, p > m$, and $(k, l) \neq (m, p)$ is in $\mathcal{C}_2$, and so on. After doing this separation of $\mathcal{RD}_n$ into subsets $\{\mathcal{C}_i\}_i$, we can investigate how the square of the ℓ_1 -norm of the best stabilizer decomposition of $C^{n-1}Z$ gates changes when we narrow the search space for the LP from the whole $\mathcal{RD}_n$ group to specific collections of $\mathcal{C}_i$'s. In doing this additional (heuristic) reduction of the search space, we are not guaranteed to find the absolute optimal decomposition and are likely to uncover only a suboptimal one. Remarkably, our results show that, for $C^{n-1}Z$ gates with $n > 3$, the optimal decomposition can be found within a restricted collection of $\mathcal{C}_i$ sets, highlighting that while some entangling structures are strictly necessary, others are not.

In Fig. 5, we show how our figure of merit changes as we increase the number of sets $\mathcal{C}_i$ that are being used during the optimization. As we can see, for $3 < n \leq 7$, we do not need to consider all $\mathcal{C}_i$ in order to achieve the optimal solution: $\xi(C^{n-1}Z)$. We conjecture that the same will hold also for $n > 7$, suggesting that further heuristic reductions




FIG. 5. The role of entanglement in the computation of $\xi(C^{n-1}Z)$. We show the minimum value found for the value of the ℓ_1 -norm squared of the decomposition of $C^{n-1}Z$ gates in terms of real-diagonal Cliffords. This value is plotted against the number of CZ sets, $\mathcal{C}_i$, used during the optimization. Note that given a fixed number of sets, several different combinations are possible; we tested all of these combinations and present only the minimum value in each case. Note that, for $n \leq 6$, the value of the stabilizer extent (i.e., the absolute minimum ℓ_1 -norm squared) is achieved using only four $\mathcal{C}_i$ sets. On the other hand, for $n = 7$, five such sets are needed.

TABLE II. Minimum ℓ_1 -norm squared of the decomposition of the CCZ gate, using different combinations of CZ sets. For $n = 3$, diagonal-real Clifford operators can be separated into four CZ sets, $\{\mathcal{C}_i\}_{i=0}^3$, where i fixes the number of CZ gates of the unitaries in each set. We present the results from the largest to the smallest best possible value of the square of the ℓ_1 -norm. In this case, the absolute optimal value is reached only when all three sets $\{\mathcal{C}_i\}_i$ are used.

$\mathcal{C}_i$ sets used during optimization	Min. ℓ_1 -norm squared
$\mathcal{C}_0$ or $\mathcal{C}_3$	6.250 00
$\mathcal{C}_1$ or $\mathcal{C}_2$	4.000 00
$\mathcal{C}_0 \cup \mathcal{C}_2$ or $\mathcal{C}_1 \cup \mathcal{C}_3$	4.000 00
$\mathcal{C}_0 \cup \mathcal{C}_3$	4.000 00
$\mathcal{C}_0 \cup \mathcal{C}_1$ or $\mathcal{C}_2 \cup \mathcal{C}_3$	2.560 00
$\mathcal{C}_1 \cup \mathcal{C}_2$	2.250 00
any $\mathcal{C}_i \cup \mathcal{C}_j \cup \mathcal{C}_k$ with $i \neq j \neq k \neq i$	2.250 00
$\bigcup_{i=0}^3 \mathcal{C}_i$	1.777 77 (optimal)

to the search space might be possible, potentially enabling the computation of the stabilizer extent for even larger unitaries. A finer look into Fig. 5 reveals that, for $n \leq 6$, four different $\mathcal{C}_i$ sets suffice to achieve the optimal decomposition of $C^{n-1}Z$ gates. On the other hand, for $n = 7$, five different $\mathcal{C}_i$ sets are needed. Nevertheless, this is still a large reduction when we consider the total of 22 different sets $\{\mathcal{C}_i\}_{i=0}^{21}$ that exist for that number of qubits.

We complement the analysis above with the information in Table II, which explicitly presents the smallest ℓ_1 -norm squared obtained when different $\mathcal{C}_i$ sets are used to decompose the CCZ gate.

We conclude this discussion by pointing out that, since the $C^{n-1}Z$ gates are not only real-diagonal but also permutation-invariant, we also leveraged Lemma 2 to further reduce the memory requirements. For details, the interested reader is referred to Appendices D and E.

3. Statements concerning gate synthesis

Combining Proposition 2 with the enhanced ability to calculate the stabilizer extent of unitaries, we improve our capacity to lower bound the number of T gates needed to synthesize a given unitary. Figure 3 allows us to immediately identify what is the minimum number of T gates that are needed for synthesizing each gate in the family of multicontrolled phase gates. For instance, we see that CS requires at least 3 T gates, the CCZ at least 4, and the CCS requires at least 5. These results are consistent with the ones obtained in Ref. [12] using the robustness of magic. However, in that work, the authors limit their analysis to three-qubit gates of the third level of the Clifford hierarchy. Here, we easily extend this analysis to a larger set of gates. For instance, we note that the C^3Z , C^3S , C^4Z , and C^4S gates all require at least 6 T magic gates to be synthesized.

Usually, gates with smaller rotation angles require more T gates to be synthesized. This is not reflected by our results, indicating that our lower bound is not tight and that the stabilizer extent might not be the best monotone for drawing conclusions concerning the cost of circuit synthesis. Even so, it provides a good estimate for many gates, as seen by the aforementioned examples of the CS and CCZ gates, which can be exactly synthesized with 3 and 4 T gates, respectively.

C. Speeding up the classical simulation of quantum Fourier transform circuits

The QFT is a fundamental subroutine in many important quantum algorithms, including quantum phase estimation, Shor's factoring algorithm, and quantum support vector machines. While QFT has been shown to be efficiently classically simulable on low-entangled inputs [33], this is no longer the case when embedded as a component within these algorithms. Here, we investigate how the simulation of the QFT block (potentially embedded within an arbitrary quantum algorithm) is significantly enhanced by our results.

For this discussion, we deliberately focus on the runtime savings in simulating the QFT circuit with sum-over-Clifford simulators that stem from our symmetry reduction techniques. This allows us to illustrate how our methods can be employed in practice to speed up stabilizer-based simulators. The actual simulation runtime depends on other factors as well, for instance, what input states are used or how the QFT is embedded in a larger circuit. This is true for any simulation method (e.g., tensor networks) and determines which one should be used in practice.

For simplicity, we start by considering the five-qubit quantum circuit implementing QFT as seen in Fig. 6. The unitary blocks U_1 to U_4 can be decomposed into sums of Cliffords using our Theorem 1 for diagonal unitaries. The stabilizer extent values obtained for each of these blocks (and previously depicted by the yellow squares in Fig. 4) are $\xi(U_1) = 1.6$, $\xi(U_2) = 1.835\,66$, $\xi(U_3) = 1.958\,57$, and $\xi(U_4) = 2.034\,80$. This can be used to quantify the cost of simulating such a unitary block: $\prod_{i=1}^4 \xi(U_i) \approx 11.705$. Alternatively, up to now, one would be limited to computing the decomposition of each controlled-phase gate individually. In that case, the cost of simulation would be determined by the product of the values of the first four blue points in Fig. 4: $\prod_{j=2}^5 \xi(CP(2\pi/2^j))^{6-j} \approx 32.015$. Hence, owing to the strong submultiplicativity of the stabilizer extent of controlled-phase gates, our results enable speeding up the simulation of this quantum circuit roughly 2.7 times.

Although this may appear small, the advantage of our method becomes increasingly significant as the number of qubits increases. For instance, for the eight-qubit QFT circuit, our results enable the use of a (suboptimal)




FIG. 6. Example of a five-qubit circuit implementing the quantum Fourier transform. Each unitary block U_j has j controlled phase gates. The stabilizer extent of various of these blocks can be seen in Fig. 4.

decomposition with a total one-norm squared of 125.25. Contrastingly, decomposing each non-Clifford controlled-phase gate individually leads to a decomposition with a one-norm squared of 1264.6, a value that is approximately 10 times larger. Similarly, for the sixteen-qubit QFT circuit, using the commonly used approximation where any controlled-phase with a rotation angle smaller than $2\pi/2^{13}$ is ignored [34], our results enable a decomposition with a one-norm squared of 9.2×10^4 . On the other hand, under the same approximation, the individual decomposition of each controlled-phase gate leads to a one-norm squared of 3.0×10^7 . Hence, our approach speeds up the simulation of such a circuit roughly 322 times.

Motivated by fault-tolerant computation, quantum circuits are often rewritten into the Clifford + T gate set. We could thus compare the cost of classically simulating the QFT circuits natively using our technique, to that of simulating the equivalent Clifford + T circuits. Based on the results from Refs. [34,35], we expect the five-qubit QFT circuit to be approximated by a Clifford+ T quantum circuits with 60–100 T gates. This corresponds to an associated stabilizer extent between 1.3×10^4 and 7.5×10^6 ; hence, our approach allows the direct classical simulation of the five-qubit QFT circuit in Fig. 6 between 1.1×10^3 and 6.4×10^5 times faster than one can simulate the corresponding Clifford+ T quantum circuit. Analogously, Nam *et al.* [34] report that their algorithm can synthesize the eight-qubit and sixteen-qubit QFT circuits, respectively, with 303 T gates and 1162 T gates. This leads to $\xi(T^{\otimes 303}) = 6.9 \times 10^{20}$ and $\xi(T^{\otimes 1162}) = 8.1 \times 10^{79}$, values that are, respectively, 5.4×10^{18} and 8.8×10^{74} times larger than the values of the one-norm squared of the decompositions obtained leveraging our strong symmetry result.

The task of synthesizing a given unitary into a specific gate set, viz., the Clifford+ T gate set, is incredibly hard. Hence, it is reasonable to assume that better algorithms than the one presented by Nam *et al.* could be built. Even still, for the cost of classically simulating these circuits to compete with that of simulating them natively, leveraging

our approach, the five-, eight-, and sixteen-qubit QFT circuits would have to be synthesized with T counts of 16, 31, and 73, respectively. It seems highly unlikely that such values can be achieved. This highlights an important point, notably that the cost of classically simulating a given quantum computation is sensitive to the set of operations used to express it.

D. Hypergraph quantum states

We now turn our attention to circuits generating a special type of quantum states known as hypergraph quantum states [36]. Formally, these states are defined as follows.

Definition 4 (Hypergraph quantum states [36]). Let $H = (V, E)$ be an undirected hypergraph where V is the set of vertices and E the set of hyperedges. Then, the corresponding n -qubit *hypergraph quantum state* is

$$|\Psi_H\rangle := \prod_{k=1}^n \prod_{e \in E^{(k)}} C^{k-1} Z_e |+\rangle^{\otimes n},$$

where $E^{(k)} \subset E$ denotes the set of hyperedges of order k , i.e., the set of hyperedges connecting k qubits. For $k = 1$, $E^{(1)} = V$.

Put into words, given an undirected hypergraph H , the corresponding hypergraph state $|\Psi_H\rangle$ is generated by first associating with each vertex the stabilizer state $|+\rangle$. Then, for every hyperedge e , a multicontrolled- Z gate $C^{k-1} Z_e$ is applied to the qubits connected by e . We note that graph states are a special type of hypergraph state where all hyperedges are of order $k = 2$.

Hypergraph states are relevant in MBQC [37–39], quantum error correction [40,41], condensed-matter physics [38,41,42], quantum metrology [43,44], and quantum foundations [40,43,45]. Moreover, these states are intimately connected to important quantum algorithms like the Deutsch-Jozsa and Grover algorithms [36]. Recently, complete categories of four-qubit hypergraph states were

experimentally prepared and verified in a reprogrammable silicon-photon quantum chip [46].

Additionally, the recent interest in hypergraph quantum states has also unveiled important connections between the resource theory of magic, Boolean function analysis, coding theory, and quantum complexity of phases of matter [38,47]. Still, while other nonclassical properties of hypergraph states have been thoroughly studied [40,43,45,48,49], the characterization of nonstabilizerness remains fairly unexplored. Reference [47] made a generic study of many-body quantum magic that included interesting results on the nonstabilizerness of the Union Jack state—a special hypergraph state known to be universal for MBQC [38]. More recently, Ref. [50] studied the nonstabilizerness of hypergraph states via stabilizer Rényi entropies.

1. Characterization of k -uniform hypergraph states

Here, we leverage our results to compute the stabilizer extent of the unitaries generating five-qubit hypergraph quantum states. Because the number of hypergraphs of five vertices is extremely large, we focus on the study of k -uniform hypergraphs, that is, hypergraphs whose hyperedges have fixed order k . Since when $k = 2$, there is no magic, we start with $k = 3$. The number of nonequivalent three-uniform hypergraphs on five vertices is 33. Interestingly, the unitaries that generate the corresponding hypergraph quantum states separate themselves into only two distinct values of the stabilizer extent: $\xi = 1.77778$ and $\xi = 2.56000$. In Fig. 7, we depict four examples of three-uniform hypergraphs on five vertices, the circuit generating the corresponding hypergraph state from the state $|+\rangle^{\otimes 5}$, and the associated stabilizer extent. We also highlight how hypergraph states with the same number of hyperedges may exhibit different values of the stabilizer extent of the underlying unitary [cf. Figs. 7(c) and 7(d) and Table III].

Figure 8 depicts all five nonequivalent four-uniform hypergraphs on five qubits, as well as the circuit that generates the corresponding hypergraph quantum state from state $|+\rangle^{\otimes 5}$ and the associated value of the stabilizer extent. In this case, there are three distinct values for this quantity: $\xi = 2.25000$, $\xi = 2.56000$, and $\xi = 2.77778$.

2. Speeding up the classical simulation of (universal) measurement-based quantum computations

The Union Jack lattice depicted in Fig. 9 is known to be universal for MBQC with Pauli measurements [38]. Thus, any quantum computation can be written as a circuit on n qubits initialized in the state $|+\rangle^{\otimes n}$, followed by the application of the magic and entangling unitary that generates the Union Jack state, and Pauli measurements at the end. Such a computation can be simulated using the sum-over-Cliffords approach. If we were to use




FIG. 7. Example of four nonequivalent three-uniform hypergraphs on five vertices. In this case, there are 33 inequivalent hypergraph states but only two distinct values for the stabilizer extent of the unitaries generating them. For each of these two values, we depict only two examples: [(a) and (c)] $\xi = 1.77778$ and [(b) and (d)] $\xi = 2.56000$. The state depicted in (c) will be important later when we discuss classical simulation of measurement-based quantum computations. Note that (c) and (d) have the same number of hyperedges, but different values of the stabilizer extent.

the CCZ decomposition used by Bravyi *et al.* [13], the simulation time would grow as $\mathcal{O}(2^{1.6601n})$. Alternatively, our result allows us to decompose each yellow cell in Fig. 9 directly, which leads to an associated simulation runtime of $\mathcal{O}(2^{0.4150n})$. This corresponds to an exponential speed up of $\mathcal{O}(2^{1.2451n})$ to the classical simulation of

TABLE III. Distribution of the 33 nonequivalent three-uniform hypergraph states of five qubits by number of hyperedges and value of the stabilizer extent of the generating unitary. The first column fixes the number of hyperedges and the second indicates how many of the nonequivalent hypergraphs have that number of hyperedges. The last two columns indicate how the hypergraphs distribute themselves across the two different values of the stabilizer extent of the generating unitary.

No. of hyperedges	No. of nonequivalent hypergraphs	$\xi = 1.77778$	$\xi = 2.56000$
1	1	1	0
2	2	1	1
3	4	2	2
4	6	2	4
5	6	1	5
6	6	1	5
7	4	1	3
8	2	1	1
9	1	0	1
10	1	0	1




FIG. 8. All five nonequivalent four-uniform hypergraphs on five vertices. We see that there are only three distinct values for the stabilizer extent for the unitaries generating the corresponding hypergraph states: $\xi = 2.25000$ when we have one (a) and two (b) $C^{(3)}Z$ gates, $\xi = 2.56000$ when all $C^{(3)}Z$ are present (c), and $\xi = 2.77778$ when there are three (d) and four (e) $C^{(3)}Z$ gates.

any MBQC written in terms of the Union Jack state with Pauli measurements.

From Ref. [47], it follows immediately that $\xi(|\Psi_{\text{UJ}}\rangle) < 2^{0.46n}$, where $|\Psi_{\text{UJ}}\rangle$ denotes the Union Jack hypergraph quantum state. Hence, our results also provide a better

upper bound for the stabilizer extent of the Union Jack state

$$\xi(|\Psi_{\text{UJ}}\rangle) < \xi(U_{\text{UJ}}) = \mathcal{O}(2^{0.4150n}),$$

with $|\Psi_{\text{UJ}}\rangle = U_{\text{UJ}}|+\rangle^{\otimes n}$.




FIG. 9. Union Jack lattice. The red shaded area represents the unit cell of the lattice. The yellow shaded square represents the three-uniform five-qubit hypergraph state depicted in Fig. 7(c), which can also be used to generate the entire Union Jack lattice.

E. Generalized hypergraph quantum states

As the name suggests, generalized hypergraph quantum states are a generalization of hypergraph states where the entangling gates are allowed to be multicontrolled phase gates $C^{k-1}P(\theta_k)$ with arbitrary angles, instead of restricting them to be multicontrolled Z gates ($\theta_k = \pi \forall k$). Moreover, the angles θ_k may differ for each order k . Formally, we define them as follows.

Definition 5 (Generalized hypergraph quantum states). Let $H = (V, E)$ be an undirected hypergraph, with V its set of vertices and E its set of hyperedges. For each integer $1 \leq k \leq |V|$, set a $(k-1)$ -controlled phase gate $C^{k-1}P(\theta_k)$, with fix angle θ_k . Then, the corresponding n -qubit generalized hypergraph quantum state is

$$|\Psi_{\text{GH}}\rangle := \prod_{k=1}^n \prod_{e \in E^{(k)}} C^{k-1}P_e(\theta_k) |+\rangle^{\otimes n},$$

where $E^{(k)} \subset E$ denotes the set of hyperedges of order k . For $k = 1$, $E^{(1)} = V$.




FIG. 10. A five-qubit generalized hypergraph circuit. For each integer $1 \leq k \leq 5$, we set the multicontrolled phase gates with maximal stabilizer extent, as given by Table I(c). We then construct unitary blocks $U_{\max}^{E(k)}$ associated with k -uniform hypergraphs and intended to have large (ideally maximal) stabilizer extent values. The associated results are: $\xi(T^{\otimes 5}) = 2.20723$, $\xi(U_{\max}^{E(2)}) = 3.76471$, $\xi(U_{\max}^{E(3)}) = 3.34751$, $\xi(U_{\max}^{E(4)}) = 3.18877$, and $\xi(U_{\max}^{E(5)}) = 2.70792$. We compose these blocks as depicted in the figure. The stabilizer extent of the whole circuit is $\xi = 3.59638$.

Here, motivated by our previous results, we compute the stabilizer extent of the unitaries generating certain generalized hypergraph states. Specifically, we consider five-qubit states and, for each value of $1 \leq k \leq 5$, we fix the multicontrolled phase gate to be the one with maximum value of the stabilizer extent (as seen in Table I(c)). Then, for $k = 1$, the unitary with maximum stabilizer extent is $T^{\otimes 5}$, with a value of $\xi(T^{\otimes 5}) = 2.20723$. For $k = 2$, we took all two-uniform nonequivalent graphs and computed the extent of the associated unitary obtained by associating with each edge the CS gate (see Appendix G for the complete set of results). We then selected the unitary with the maximum stabilizer extent, which is denoted $U_{\max}^{E(2)}$ and has a stabilizer extent of $\xi(U_{\max}^{E(2)}) = 3.76471$. For $k = 3$ and $k = 4$, we took the three- and four-uniform hypergraphs in Figs. 7 and 8 that were associated with the largest value of the stabilizer extent when using multicontrolled Z gates, and replaced those gates by the $C^2P(0.6476\pi)$ and $C^3P(0.7048\pi)$ gates, respectively. We denote the resulting unitaries, respectively, $U_{\max}^{E(3)}$ and $U_{\max}^{E(4)}$ and their stabilizer extents are $\xi(U_{\max}^{E(3)}) = 3.34751$ and $\xi(U_{\max}^{E(4)}) = 3.18877$. As expected, these results are larger than the ones obtained with C^2Z and C^3Z . Finally, for $k = 5$, we straightforwardly chose $U_{\max}^{E(5)} = C^4P(0.7288\pi)$ for which $\xi(U_{\max}^{E(5)}) = 2.70792$ (cf. Table I(c)).

The unitaries $U_{\max}^{E(k)}$ can be used to construct a quantum circuit as depicted in Fig. 10. Given how it was constructed, one might be tempted to conjecture that such a circuit would have a very large stabilizer extent or even the largest possible stabilizer extent. However, that is emphatically not the case, as the circuit depicted has $\xi = 3.59638 < \xi(U_{\max}^{E(2)})$. Going even further, we study how the stabilizer extent changes under different compositions of the unitary blocks $U_{\max}^{E(k)}$. The results can be seen in Table IV. We note that some of the combinations have a stabilizer extent lower than that of the whole circuit, while

others have a larger value. Remarkably, none of the tested compositions achieves a stabilizer extent larger than that of $U_{\max}^{E(2)}$. As another example, we note that $T^{\otimes 5}$ and $U_{\max}^{E(5)}$ are two blocks which are permutation invariant. However, while removing the former causes a meaningful reduction in the value of the stabilizer extent (from 3.59638 to 3.08596), removing the latter makes the value of the stabilizer increase (to 3.60744).

Another interesting point to consider is that the composition $T^{\otimes 5} \cdot U_{\max}^{E(2)}$ is an instance of the so-called *instantaneous quantum polynomial* (IQP) circuits capable of demonstrating quantum advantage. As shown in Ref. [51], efficient classical simulation (in the weak sense) is impossible for a class of IQP circuits composed only of $\{T, CS\}$ or $\{Z, CZ, CCZ\}$ (see also Ref. [52]). This is interesting, as the aforementioned composition is not the one with the maximum value of the stabilizer extent. In fact, of all the compositions presented in Table IV, this is the one with the smallest stabilizer extent, making it the most easily simulable by the sum-over-Cliffords approach. These observations seem to connect to the work by Liu and

TABLE IV. Stabilizer extent of five-qubit unitaries obtain by composing unitaries $U_{\max}^{E(k)}$. The stabilizer extent of each individual unitary can be found in the text.

Composition	ξ
$T^{\otimes 5} \cdot U_{\max}^{E(2)}$	3.08388
$T^{\otimes 5} \cdot U_{\max}^{E(2)} \cdot U_{\max}^{E(3)}$	3.37165
$T^{\otimes 5} \cdot U_{\max}^{E(2)} \cdot U_{\max}^{E(3)} \cdot U_{\max}^{E(4)}$	3.60744
$T^{\otimes 5} \cdot U_{\max}^{E(2)} \cdot U_{\max}^{E(3)} \cdot U_{\max}^{E(4)} \cdot U_{\max}^{E(5)}$	3.59638
$U_{\max}^{E(2)} \cdot U_{\max}^{E(3)}$	3.23195
$U_{\max}^{E(2)} \cdot U_{\max}^{E(3)} \cdot U_{\max}^{E(4)}$	3.15329
$U_{\max}^{E(2)} \cdot U_{\max}^{E(3)} \cdot U_{\max}^{E(4)} \cdot U_{\max}^{E(5)}$	3.08596

Winter [47], who pointed out that more magic does not necessarily equal more computational power.

V. DISCUSSION AND OUTLOOK

In this paper, we exploit symmetry reductions to accelerate state-of-the-art classical algorithms for simulating Clifford-dominated circuits. Our approach yields more efficient methods for decomposing multiqubit, non-Clifford gates that are invariant under certain symmetry groups. This allowed us to find (previously out of reach) optimal expansions of several multiqubit quantum gates as sums of Clifford unitaries. These were then leveraged to demonstrate drastic improvements to the efficiency of sum-over-Clifford simulation for many cases of practical interest. In particular, we showed that our method leads to remarkable speedups in stabilizer-extent-based classical algorithms for simulating QFTs, IQP circuits, and MBQCs based on hypergraph states.

Beyond this, our method should be of further general interest in quantum computing problems where sequences of symmetric quantum gates appear naturally because of, e.g., operational constraints or existing problem structures. For instance, sequences of diagonal and real multiqubit gates have been proposed in applications such as quantum compiling and quantum approximate optimization [53,54]. Commuting (diagonal) gates are also native in quantum platforms such as ion traps and Rydberg atoms, and complex-sequences thereof are employed in quantum simulation of long-ranged quantum Ising models [55–61]. Real and symmetric gates further appear commonly in fault tolerant quantum architectures, as transversal or fault-tolerant gates, e.g., while doing error correction with surface codes [33,62,63].

Our work motivates further research into the classification of the mathematical symmetry groups that admit what we called *strong symmetry reduction* (Theorem 1), which takes us to a new optimization search space of reduced size where further *weak symmetry reductions* are feasible to implement (Lemma 2). In particular, we leave the following open question: What is the underlying structure or physical property behind strong symmetry reduction? In quantum computation, group-theoretic methods, such as representation theory or group cohomology, have already been successful in the development of stabilizer-based classical simulation methods [64–66], the identification of computational phases of matter [67–69], and in the study of quantum resources such as quantum contextuality [70–72]. Developing further methods to understand symmetry in quantum circuits could yield new insights into the nature of quantum speedups, as well as better algorithms for classical simulation and hybrid quantum-classical computation [73].

A potential avenue for improving our methods would be studying the effect of approximation errors in the sum-over-Cliffords expansion. In situations where, even after leveraging strong symmetry reduction, the dimensionality of the convex problem remains unmanageable, we could try to approximate the expansion of the target gate using a smaller set of Clifford operators, trading computational cost for an error. Along a similar line, developing a heuristic algorithm with convergence guarantees reminiscent of the techniques proposed in Ref. [22] should also lead to relevant enhancements of our results. Lastly, as mentioned in Sec. III A, weak symmetry reduction is applicable to any symmetry group with linear isometries. However, the cost of implementing this type of filtering has only been investigated in a reduced number of cases. Further improvements of our method can be achieved by designing better algorithms for performing the symmetry projection of a Clifford unitary.

A salient advantage of our symmetry reduction methods is their enhanced ability to find Clifford decompositions of sequences of quantum gates invariant under certain symmetry groups. This includes examples of important gates such as circuits of T , CS , and CCZ gates. We observe that directly decomposing such sequences has a remarkably positive impact on the cost of classical simulation (e.g. cf. Fig. 4). Underlying this advantage is the strict submultiplicativity of the stabilizer extent for these sequences—a fundamental feature of this monotone, which has not been analyzed in detail earlier. This signals out the stabilizer extent as a potential quantifier to study the nonclassical aspects of deep universal quantum circuits. Can the stabilizer extent capture quantum features that power quantum computation? How does this monotone relate to other notions of nonclassicality that have been proposed for sequences of quantum transformations, such as, e.g., sequential contextuality [74,75], transformation contextuality [76,77], or the memory cost of quantum measurement sequences [78,79]? An answer to these questions would lead to a deeper understanding of the foundations of quantum computational advantage.

ACKNOWLEDGMENTS

We thank David Gross, Hammam Qassim, D. Luis Manuel Máñez Espina, Rhea Alexander-Turner, and José M. Martín for profitable research discussions. We acknowledge support from Ayuda Consolidación CNS2023-145392 (MICIU/AEI/10.13039/501100011033, NextGenerationEU/PRTR); Ramón y Cajal RYC2022-036209-I (MICIU/AEI/10.13039/501100011033, ESF+); EU HORIZON RIA FoQaCiA GA 101070558; Project Generation of Knowledge PID 2024-162155OB-I00 (MICIU/AEI/10.13039/501100011033, ERDF/EU); Project FEDER C-EXP-256-UGR23 (Consejería de Universidad, Investigación e Innovación y UE Programa FEDER

Andalucía 2021-2027). M.H. acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—Cluster of Excellence Matter and Light for Quantum Computing (ML4Q) EXC 2004/1-390534769.

DATA AVAILABILITY

The code used to produce the numerical results is publicly available on Github [30].

APPENDIX A: PROPERTIES OF THE STABILIZER EXTENT FOR MATRICES

In this Appendix, we prove the properties of the stabilizer extent for matrices as stated in Proposition 1.

Proof of Proposition 1. We will prove the properties of the stabilizer extent for matrices one by one.

(i) Let $A = \sum_{C \in \mathcal{C}_n} x_C C$ be a unitary operator $A \in \mathcal{U}(2^n)$ so that $AA^\dagger = A^\dagger A = \mathbb{1}$. This implies that

$$|\text{tr}[A^\dagger A]| = 2^n \iff 2^n = \left| \text{tr} \left[\sum_{C, C' \in \mathcal{C}_n} x_C x_C^* C^\dagger C' \right] \right|.$$

Analyzing the right-hand side, we note that

$$\begin{aligned} \left| \text{tr} \left[\sum_{C, C' \in \mathcal{C}_n} x_C x_C^* C^\dagger C' \right] \right| &= \left| \sum_{C, C' \in \mathcal{C}_n} x_C x_C^* \text{tr}[C^\dagger C'] \right| \\ &\leq 2^n \sum_{C, C' \in \mathcal{C}_n} |x_C| |x_C^*| = 2^n \|\mathbf{x}\|_1^2, \end{aligned}$$

which implies that $\|\mathbf{x}\|_1^2 \geq 1$. This means that any decomposition of a unitary operator must have an ℓ_1 -norm squared greater than or equal to one, thus, $\xi(A) \geq 1 \forall A \in \mathcal{U}(2^n)$. When A is not only a unitary but a Clifford unitary, the trivial decomposition of A (written as itself) achieves the minimum possible value of 1, thus $\xi(A) = 1 \forall A \in \mathcal{C}_n$. From the arguments above, it is not hard to see that for any non-Clifford unitary A , $\xi(A) > 1$. The extent as defined in Definition 2 can only be smaller than 1 for nonunitary matrices.

(ii) Let A be a matrix with an optimal decomposition $A = \sum_{C \in \mathcal{C}_n} x_C C$ meaning that $\xi(A) = \|\mathbf{x}\|_1^2$. Suppose that A is left-multiplied by a Clifford unitary $K \in \mathcal{C}_n$ so that

$$B = KA = \sum_{C \in \mathcal{C}_n} x_C KC = \sum_{C' \in \mathcal{C}_n} x_{K^\dagger C'} C'.$$

By the definition of the stabilizer extent, the decomposition obtained above for B must obey $\xi(B) \leq \|\mathbf{x}\|_1^2 = \xi(A)$. Next, suppose that the optimal decomposition of B is

$B = \sum_{C \in \mathcal{C}_n} y_C C$ so that, by definition, $\xi(B) = \|\mathbf{y}\|_1^2$. Since $B = KA$, then $A = K^\dagger B$ which implies that

$$A = \sum_{C \in \mathcal{C}_n} y_C K^\dagger C = \sum_{C' \in \mathcal{C}_n} y_{KC'} C'.$$

Again, by the definition of the stabilizer extent, this decomposition implies that $\xi(A) \leq \|\mathbf{y}\|_1^2 = \xi(B)$. Thus, we obtain that the stabilizer extent of A and the stabilizer extent of $B = KA$ are the same. Similar calculations follow for right multiplication. Thus, the stabilizer extent for matrices is preserved under left and right multiplication by Clifford unitaries.

To prove monotonicity under stabilizer code projectors, consider an abelian subgroup of the Pauli group $\mathcal{G} \subset \mathcal{P}_n$ such that $|\mathcal{G}| = 2^r$. A projector onto the codespace defined by $\mathcal{G}$ has the form $\Pi_{\mathcal{G}} = 2^{-r} \sum_{P \in \mathcal{G}} P$. Then, we write

$$\begin{aligned} \xi(\Pi_{\mathcal{G}} A) &= \xi \left(2^{-r} \sum_{P \in \mathcal{G}} PA \right) \leq 2^{-r} \sum_{P \in \mathcal{G}} \xi(PA) \\ &= 2^{-r} \sum_{P \in \mathcal{G}} \xi(A) = \xi(A), \end{aligned}$$

where we leveraged convexity (see below) and invariance under Cliffords. A similar reasoning allows to prove that $\xi(A \Pi_{\mathcal{G}}) \leq \xi(A)$.

(iii) To prove the submultiplicativity of the stabilizer extent under matrix multiplication consider two matrices $A, B \in \mathbb{C}^{2^n \times 2^n}$ with optimal decompositions $A = \sum_{C \in \mathcal{C}_n} x_C C$ and $B = \sum_{C' \in \mathcal{C}_n} y_{C'} C'$ so that $\xi(A) = \|\mathbf{x}\|_1^2$ and $\xi(B) = \|\mathbf{y}\|_1^2$. Next, we take the product of these two matrices

$$\begin{aligned} AB &= \sum_{C, C' \in \mathcal{C}_n} x_C y_{C'} C C' = \sum_{K \in \mathcal{C}_n} \left(\sum_{C \in \mathcal{C}_n} x_C y_{C^\dagger K} \right) K \\ &= \sum_{K \in \mathcal{C}_n} z_K K. \end{aligned}$$

By definition $\xi(AB) \leq \|\mathbf{z}\|_1^2$ and we can upperbound $\|\mathbf{z}\|_1$ as follows:

$$\begin{aligned} \|\mathbf{z}\|_1 &= \sum_{K \in \mathcal{C}_n} \left| \sum_{C \in \mathcal{C}_n} x_C y_{C^\dagger K} \right| \leq \sum_{K, C \in \mathcal{C}_n} |x_C y_{C^\dagger K}| \\ &= \|\mathbf{x}\|_1 \|\mathbf{y}\|_1. \end{aligned}$$

Hence, $\xi(AB) \leq \xi(A) \xi(B)$.

Submultiplicativity with respect to tensor products is fairly straightforward to show. As before, let us consider two matrices $A \in \mathbb{C}^{2^n \times 2^n}$, $B \in \mathbb{C}^{2^m \times 2^m}$ with optimal decompositions $A = \sum_{C \in \mathcal{C}_n} x_C C$ and $B = \sum_{C' \in \mathcal{C}_m} y_{C'} C'$ so

that $\xi(A) = \|\mathbf{x}\|_1^2$ and $\xi(B) = \|\mathbf{y}\|_1^2$. Considering their tensor product, we note that

$$A \otimes B = \sum_{C \in \mathcal{C}_n} \sum_{C' \in \mathcal{C}_{n'}} x_C y_{C'} (C \otimes C').$$

By definition $\xi(A \otimes B) \leq \left(\sum_{C \in \mathcal{C}_n} \sum_{C' \in \mathcal{C}_{n'}} |x_C y_{C'}| \right)^2 = \xi(A) \xi(B)$.

Showing that $\xi(A \otimes B) = \xi(B \otimes A)$ uses property (ii). Namely, there is always a suitable swapping operation, which we denote $\text{SWAP} \in \mathcal{C}_{n+m}$, so that $\text{SWAP}(A \otimes B) \text{SWAP} = (B \otimes A)$. Hence, by property (ii), $\xi(A \otimes B) = \xi(\text{SWAP}(A \otimes B) \text{SWAP}) = \xi(B \otimes A)$.

(iv) To demonstrate convexity of the stabilizer extent for matrices, consider again two matrices $A, B \in \mathbb{C}^{2^n \times 2^n}$ with optimal decompositions $A = \sum_{C \in \mathcal{C}_n} x_C C$ and $B = \sum_{C' \in \mathcal{C}_n} y_{C'} C'$. Then, for any $t \in [0, 1]$, $tA + (1-t)B = \sum_{C \in \mathcal{C}_n} (tx_C + (1-t)y_C) C$ which implies that $\xi(tA + (1-t)B) \leq \|t\mathbf{x} + (1-t)\mathbf{y}\|_1^2 \leq t\xi(A) + (1-t)\xi(B)$ by the convexity of the norm and the square function. Homogeneity follows straightforwardly from Definition 2 and the homogeneity of the ℓ_1 -norm.

This concludes the proof of Proposition 1. $\blacksquare$

We call the reader's attention to the proof of property (i): If we were to restrict the definition of the stabilizer extent to unitaries, this would be a faithful measure in the sense that $\xi(A) = 1$ iff $A \in \mathcal{C}_n$ and $\xi(A) > 1$ otherwise. As it stands, Definition 2 applies to any arbitrary matrix and the stabilizer extent can therefore be smaller than 1; this is seen explicitly in Lemma 1 and its proof in Appendix C.

APPENDIX B: LOWER BOUNDS ON GATE SYNTHESIS

Here, we prove Proposition 2, which explicitly states how the stabilizer extent can be used to lower bound the cost of exactly synthesizing certain unitaries.

Proof of Proposition 2. By assumption, $s \in \mathbb{N}$ is the largest positive integer fulfilling the lower bound $\xi(T)^{s-1} < \xi(U)$, which implies the upper bound $\xi(U) \leq \xi(T)^s$. Suppose wlog that the unitary gate of interest can be written as a quantum circuit of form $U = C_m L_m C_{m-1} L_{m-1} \cdots C_1 L_1 C_0$, where $\{C_j\}$ are Clifford unitaries and $\{L_j\}$ are non-Clifford layers containing T gates acting on distinct qubits. Let t be the total number of T gates in this circuit. Then, the submultiplicativity of the extent implies $\xi(U) \leq \xi(T)^t$. Putting this together with our lower bound, it follows that $t > s - 1$. This holds for any quantum circuit of Clifford+ T gates that can exactly synthesize U . Hence, U cannot possibly be synthesized with $s - 1$ (or fewer) T gates.

Let us consider that the optimal gate synthesis of U in terms of the Clifford + T gate set requires exactly

$t_{\min}$ T gates. Then, assigning these gates into the non-stabilizer layers and using Proposition 1 leads to $\xi(U) = \xi(C_m L_m C_{m-1} L_{m-1} \cdots C_1 L_1 C_0) \leq \xi(T)^{t_{\min}}$ and, necessarily, $t_{\min} \geq s$. Thus, the number of T gates needed to synthesize U is at least s . $\blacksquare$

It is well known that magic monotones can be used to lower bound the cost of (exact) synthesis of certain unitaries [12,20]. For instance, a similar result was given in Ref. [12] using the robustness of magic, $\mathcal{R}$. In particular, let $|U\rangle = U|+\rangle^{\otimes n}$; if $\mathcal{R}(|T\rangle^{\otimes(s-1)}) < \mathcal{R}(|U\rangle) < \mathcal{R}(|T\rangle^{\otimes s})$, then a minimum number of s T gates are needed to (exactly) synthesize U . However, this result is limited by the submultiplicativity of the robustness of magic, which makes it hard to determine $\mathcal{R}(|T\rangle^{\otimes s})$ for values of s beyond 10. Using the stabilizer extent has no such limitation, since $\xi(T^{\otimes s}) = \xi(T)^s$ (recall Sec. II A) and, therefore, the upper and lower bounds depend uniquely on $\xi(T) = \cos^{-2}(\pi/8) \approx 1.1716$.

The result in Proposition 2 can also be extended to other gate sets of the form Clifford+ M provided that M is a non-Clifford gate on at most three qubits which allows for a magic state $|M\rangle$ such that $\xi(M) = \xi(|M\rangle)$ (recall discussion in Sec. II A).

APPENDIX C: PROOF OF LEMMA 1

Lemma 1 is essential to the proof of our main result. In this appendix, we give a proof thereof, separated into three lemmas covering the three different subgroups. The arguments use that stabilizer states have the following well-known form in the computational basis [80]:

$$|\psi\rangle = \frac{1}{\sqrt{|A|}} \sum_{\mathbf{x} \in A} i^{\mathbf{a} \cdot \mathbf{x}} (-1)^{q(\mathbf{x}) + \mathbf{b} \cdot \mathbf{x}} |\mathbf{x}\rangle, \quad (\text{C1})$$

where $A \subset \mathbb{Z}_2^n$ is a suitable affine subspace, $\mathbf{a} \in \mathbb{Z}_2^n$, and q is a quadratic function. We note that the exponent of i is meant to be evaluated modulo two. Every stabilizer state has the form Eq. (C1), and vice versa, any state of the form Eq. (C1) is a stabilizer state.

The matrix entries of a Clifford unitary can be deduced from Eq. (C1) by referring to its Choi state $C \otimes \mathbb{1} |\phi^+\rangle$. Since $|\phi^+\rangle := 2^{-n/2} \sum_{\mathbf{x} \in \mathbb{Z}_2^n} |\mathbf{x}\mathbf{x}\rangle$ is a stabilizer state, so is $C \otimes \mathbb{1} |\phi^+\rangle$, and we can write

$$C \otimes \mathbb{1} |\phi^+\rangle = \frac{1}{\sqrt{|W|}} \sum_{\mathbf{w} \in W} i^{\mathbf{a} \cdot \mathbf{w}} (-1)^{q(\mathbf{w}) + \mathbf{b} \cdot \mathbf{w}} |\mathbf{w}\rangle,$$

where $W \subset \mathbb{Z}_2^{2n}$ is a suitable affine subspace of dimension $n + s$ with $0 \leq s \leq n$ [81] (s is the dimension of the support of every $C|\mathbf{x}\rangle$). Writing $\mathbf{w} = (\mathbf{x}, \mathbf{y})$, we then find the matrix

entries of C to be

$$\begin{aligned} \langle \mathbf{x} | C | \mathbf{y} \rangle &= 2^{n/2} \langle \mathbf{x} \mathbf{y} | C \otimes \mathbb{1} | \phi^+ \rangle \\ &= \frac{1}{2^{s/2}} i^{\mathbf{a} \cdot (\mathbf{x}, \mathbf{y})} (-1)^{q(\mathbf{x}, \mathbf{y}) + \mathbf{b} \cdot (\mathbf{x}, \mathbf{y})} 1_W(\mathbf{x}, \mathbf{y}). \end{aligned} \quad (\text{C2})$$

Here $1_W(w)$ is the indicator function on W which is one if $w \in W$ and zero else. Note that not every $2n$ -qubit stabilizer state corresponds to a Clifford unitary; only the maximally entangled ones do. By explicitly taking partial traces, it is straightforward to verify that this restricts the affine subspaces $W \in \mathbb{Z}_2^{2n}$ and the quadratic forms q , but not the linear part given by $\mathbf{a}$ and $\mathbf{b}$. In particular, the matrix C' obtained by setting $\mathbf{a} = 0$ is another valid Clifford unitary.

Finally, we note that the form of stabilizer states, and also Clifford unitaries, in the computational basis is only defined up to a global phase. In the standard definition of the Clifford group as generated by Hadamard, CNOT, and phase gate, its center is $\mathbb{Z}_8$, and thus for any $C \in \mathcal{C}_n$, we also have $\omega_8 C \in \mathcal{C}_n$ for any eighth root of unity ω_8 . However, for the stabilizer extent, this global phase does not matter: we can simply group the coefficients associated with the same Clifford unitary up to a phase. Concretely, let $\bar{\mathcal{C}}_n$ be the subset given by the phase convention Eq. (C2) and then write $U = \sum_{C \in \bar{\mathcal{C}}_n} (\sum_{k=0}^7 x_C^{(k)} \omega_8^k) C$. This procedure can only decrease the ℓ_1 -norm and thus any optimal decomposition can be chosen to only contain Cliffords from $\bar{\mathcal{C}}_n$.

Lemma C1. For all Clifford unitaries $C \in \mathcal{C}_n$, we have $\xi_{\mathcal{Z}_n}(\Pi_{\mathcal{Z}_n}(C)) \leq 1$.

Note that $\Pi_{\mathcal{Z}_n}$ is $\mathbb{C}$ -linear, hence we do not need to explicitly consider a complex factor $x \in \mathbb{C}$.

Proof. For C diagonal Clifford, the result is trivially true, as for $\Pi_{\mathcal{Z}_n}(C) = 0$. Hence, let us assume that C is nondiagonal and that $d := \Pi_{\mathcal{Z}_n}(C) \neq 0$.

By Eq. (C2), the nonvanishing entries of d then have the form

$$d_{\mathbf{x}\mathbf{x}} = 2^{-s/2} i^{\mathbf{a}' \cdot \mathbf{x}} (-1)^{q'(\mathbf{x}) + \mathbf{b}' \cdot \mathbf{x}}, \quad (\text{C3})$$

where $\mathbf{a}', \mathbf{b}' \in \mathbb{Z}_2^n$ and q' are obtained from above in the obvious way. Next, we show that d can be extended to a diagonal Clifford unitary. Let $S \subset \mathbb{Z}_2^n$ be the support of $\mathbf{x} \mapsto d_{\mathbf{x}\mathbf{x}}$ and let us define another diagonal matrix κ on the complement of S by extending the linear and quadratic forms in Eq. (C3) (in some way, this may not be unique). Then, the diagonal matrix $D = 2^{s/2}(d + \kappa)$ is by construction Clifford because its diagonal is proportional to a (full-support) stabilizer state [80,81].

The support S of d is essentially the intersection of the affine support of $C \otimes \mathbb{1} | \phi^+ \rangle$ with the diagonal $\Delta =$

$\{(\mathbf{x}, \mathbf{x}) \mid \mathbf{x} \in \mathbb{Z}_2^n\}$, and thus also an affine subspace. We can hence write it as $S = V + \mathbf{t}$ for a linear subspace V and a vector $\mathbf{t}$. $V^\perp$ is given as the subspace of vectors that are orthogonal to V with respect to the binary dot product. Then, $G := \{(-1)^{\mathbf{t} \cdot \mathbf{z}} Z(\mathbf{z}) \mid \mathbf{z} \in V^\perp\}$ is an Abelian group which stabilizes d ,

$$(-1)^{\mathbf{t} \cdot \mathbf{z}} Z(\mathbf{z}) \cdot d = \sum_{\mathbf{x} \in S} (-1)^{\mathbf{z} \cdot (\mathbf{x} + \mathbf{t})} d_{\mathbf{x}\mathbf{x}} | \mathbf{x} \rangle \langle \mathbf{x} | = d,$$

where the last step follows from $\mathbf{x} + \mathbf{t} \in V$ for all $\mathbf{x} \in S$. The average over G annihilates κ ,

$$\begin{aligned} \frac{1}{|V^\perp|} \sum_{\mathbf{z} \in V^\perp} (-1)^{\mathbf{t} \cdot \mathbf{z}} Z(\mathbf{z}) \cdot \kappa &= \frac{1}{|V^\perp|} \sum_{\mathbf{x} \in S^c} \sum_{\mathbf{z} \in V^\perp} (-1)^{\mathbf{z} \cdot (\mathbf{x} + \mathbf{t})} \kappa_{\mathbf{x}\mathbf{x}} | \mathbf{x} \rangle \langle \mathbf{x} | \\ &= \frac{1}{|V^\perp|} \sum_{\mathbf{x} \in S^c} \underbrace{\sum_{\mathbf{z} \in V^\perp} (-1)^{\mathbf{z} \cdot (\mathbf{x} + \mathbf{t})}}_{=|V^\perp| \delta(\mathbf{x} + \mathbf{t} \in V)} \kappa_{\mathbf{x}\mathbf{x}} | \mathbf{x} \rangle \langle \mathbf{x} | = 0. \end{aligned}$$

Hence, we get the desired decomposition of d in terms of diagonal Cliffords,

$$\begin{aligned} \frac{2^{-s/2}}{|V^\perp|} \sum_{\mathbf{z} \in V^\perp} (-1)^{\mathbf{t} \cdot \mathbf{z}} Z(\mathbf{z}) \cdot D &= \frac{1}{|V^\perp|} \sum_{\mathbf{z} \in V^\perp} (-1)^{\mathbf{a} \cdot \mathbf{z}} Z(\mathbf{z}) \cdot (d + \kappa) = d. \end{aligned}$$

This shows that $\xi(d) = \xi(\Pi_{\mathcal{Z}_n}(C)) \leq 2^{-s} \leq 1$ as claimed. In fact, this is also optimal since we find for any decomposition $d = \sum_i c_i D^{(i)}$ into diagonal Cliffords $D^{(i)}$,

$$2^{-s/2} = |d_{\mathbf{x}\mathbf{x}}| \leq \sum_i |c_i| \underbrace{|D_{\mathbf{x}\mathbf{x}}^{(i)}|}_{=1} = \|c\|_1. \quad \blacksquare$$

Lemma C2. For all Clifford unitaries $C \in \mathcal{C}_n$ and $x \in \mathbb{C}$, we have $\xi_{\mathcal{K}_n}(\Pi_{\mathcal{K}_n}(xC)) \leq |x|^2$.

Proof. We will show the following: If the phase of x is not a second or fourth root of unity, then $\Pi_{\mathcal{K}_n}(xC) = 0$ and the statement is trivially true. Otherwise, we construct a decomposition

$$\Pi_{\mathcal{K}_n}(xC) = \alpha |x| \frac{1}{|G|} \sum_{P \in G} P C', \quad (\text{C4})$$

where $0 \leq \alpha \leq 1$, C' is a real Clifford unitary, and the P 's are real Pauli operators from a suitable stabilizer group G . From this, the claim $\xi_{\mathcal{K}_n}(xC) \leq \alpha^2 |x|^2 \leq |x|^2$ follows directly.

For now, let us assume that $x = 1$ and show $\xi_{\mathcal{K}_n}(\Pi_{\mathcal{K}_n}(C)) \leq 1$. As we will see, it is pretty straightforward to modify our argument to cover the general case. By Eq. (C2), the matrix entries of C are of the form

$$\langle \mathbf{x} | C | \mathbf{y} \rangle = \frac{1}{2^{s/2}} i^{\mathbf{a} \cdot (\mathbf{x}, \mathbf{y})} (-1)^{q(\mathbf{x}, \mathbf{y}) + \mathbf{b} \cdot (\mathbf{x}, \mathbf{y})} \mathbf{1}_W(\mathbf{x}, \mathbf{y}).$$

Recall that $\mathbf{a}, \mathbf{b}$, and q are suitable linear and quadratic forms on a $(n + s)$ -dimensional affine subspace $W \subset \mathbb{Z}_2^{2n}$. The matrix $\eta = \Pi_{\mathcal{K}_n}(C)$ then has the following entries:

$$\eta_{\mathbf{x}, \mathbf{y}} = 2^{-s/2} \mathbf{1}_W(\mathbf{x}, \mathbf{y}) \begin{cases} (-1)^{q(\mathbf{x}, \mathbf{y}) + \mathbf{b} \cdot (\mathbf{x}, \mathbf{y})} & \text{if } \mathbf{a} \cdot (\mathbf{x}, \mathbf{y}) = 0, \\ 0 & \text{otherwise.} \end{cases}$$

We define another real matrix κ by

$$\kappa_{\mathbf{x}, \mathbf{y}} := 2^{-s/2} \mathbf{1}_W(\mathbf{x}, \mathbf{y}) \begin{cases} (-1)^{q(\mathbf{x}, \mathbf{y}) + \mathbf{b} \cdot (\mathbf{x}, \mathbf{y})} & \text{if } \mathbf{a} \cdot (\mathbf{x}, \mathbf{y}) \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

By construction, we have $C = \eta + i\kappa$. Moreover, if $C_{\mathbf{y}}$ denote the columns of C and similarly $\eta_{\mathbf{y}}$ and $\kappa_{\mathbf{y}}$ those of η and κ , respectively, then it is easy to check that

$$\frac{\mathbb{1} + (-1)^{\mathbf{v} \cdot \mathbf{y}} Z(\mathbf{u})}{2} |C_{\mathbf{y}}\rangle = |\eta_{\mathbf{y}}\rangle$$

and

$$\frac{\mathbb{1} - (-1)^{\mathbf{v} \cdot \mathbf{y}} Z(\mathbf{u})}{2} |C_{\mathbf{y}}\rangle = i |\kappa_{\mathbf{y}}\rangle.$$

Here, $\mathbf{a} = (\mathbf{u}, \mathbf{v})$. This shows that the columns of η and κ are (potentially subnormalized) stabilizer states, and we can deduce their stabilizer groups using standard arguments (cf. Ref. [82, Sec. 4.4]). Note that the columns of C form a *stabilizer basis*, namely $|C_{\mathbf{y}}\rangle = C |\mathbf{y}\rangle$, thus their stabilizer groups $G_{\mathbf{y}}$ are the same up to signs.

Case 1: Suppose $Z(\mathbf{u})$ commutes with all stabilizers, i.e. it is contained in all $G_{\mathbf{y}}$ up to a sign that we can write as $(-1)^{\mathbf{m} \cdot \mathbf{y} + c}$. Then, $C_{\mathbf{y}}$ is either left invariant by $\frac{1}{2}(\mathbb{1} + (-1)^{\mathbf{v} \cdot \mathbf{y} + c} Z(\mathbf{u}))$ or is in its kernel, depending on whether $\mathbf{m} \cdot \mathbf{y} = \mathbf{v} \cdot \mathbf{y}$ or not.

- (1a) Assume that $\mathbf{m} = \mathbf{v}$. Then all columns of C are left invariant and either $C = \eta$ or $C = i\kappa$, depending on $c = 0$ or not. In the first case, C is real and we can take $C = C'$, $\alpha = 1$, $G = \{\mathbb{1}\}$. In the second case, C is purely imaginary, $\Pi_{\mathcal{K}_n}(C) = 0$, and we can take $\alpha = 0$.
- (1b) Assume that $\mathbf{m} \neq \mathbf{v}$ and for simplicity $c = 0$. Then, exactly half of the columns of C are left invariant and become the columns of η and the other half become the columns of $i\kappa$. Note that the $|\mathbf{y}\rangle$ with $(\mathbf{m} + \mathbf{v}) \cdot \mathbf{y} = 0$ span a stabilizer code which is stabilized by $Z(\mathbf{m} + \mathbf{v})$. Thus, the same is true for the

corresponding $|C_{\mathbf{y}}\rangle$ and the columns of η are stabilized by $G := \{\mathbb{1}, CZ(\mathbf{m} + \mathbf{v})C^\dagger\}$. Note that since η is real, G has to be a real stabilizer group. By construction, the columns of κ are (-1) -eigenvectors of $CZ(\mathbf{m} + \mathbf{v})C^\dagger$ and thus

$$\frac{1}{|G|} \sum_{h \in G} h(\eta + \kappa) = \eta = \Pi_{\mathcal{K}_n}(C).$$

Note that $\eta + \kappa = C'$ where C' is a real matrix with entries

$$C'_{\mathbf{x}, \mathbf{y}} = 2^{-s/2} (-1)^{q(\mathbf{x}, \mathbf{y}) + \mathbf{b} \cdot (\mathbf{x}, \mathbf{y})} \mathbf{1}_W(\mathbf{x}, \mathbf{y}).$$

This is effectively the Clifford C with $\mathbf{a} = 0$, thus C' defines a valid Clifford unitary with real entries by the discussion after Eq. (C2). Note that for $c = 1$ the roles of η and κ are simply reversed, which does not change the result. Setting $\alpha = 1$, this concludes the construction of Eq. (C4) for case 1.

Case 2: Let us now assume that $Z(\mathbf{u})$ does not commute with all stabilizers. Note that this implies that the norm of $|\eta_{\mathbf{y}}\rangle$ and $|\kappa_{\mathbf{y}}\rangle$ is $1/\sqrt{2}$, independent of other details. Moreover, we can find Paulis $g_1, \dots, g_n$ (possibly with a $\pm$ sign) such that $Z(\mathbf{u})$ anticommutes with g_1 , it commutes with $g_2, \dots, g_n$, and $G_{\mathbf{y}} = \langle (-1)^{\mathbf{t}_1 \cdot \mathbf{y}} g_1, \dots, (-1)^{\mathbf{t}_n \cdot \mathbf{y}} g_n \rangle$. Here, $\mathbf{t}_i \in \mathbb{Z}_2^n$ is uniquely determined by g_i . The stabilizer groups of the columns of η and κ are then given by $G_{\mathbf{y}}^\eta = \langle (-1)^{\mathbf{v} \cdot \mathbf{y}} Z(\mathbf{u}), (-1)^{\mathbf{t}_2 \cdot \mathbf{y}} g_2, \dots, (-1)^{\mathbf{t}_n \cdot \mathbf{y}} g_n \rangle$ and $G_{\mathbf{y}}^\kappa = \langle (-1)^{\mathbf{v} \cdot \mathbf{y}} Z(\mathbf{u}), (-1)^{\mathbf{t}_2 \cdot \mathbf{y}} g_2, \dots, (-1)^{\mathbf{t}_n \cdot \mathbf{y}} g_n \rangle$, respectively. Note that since $\eta_{\mathbf{y}}$ and $\kappa_{\mathbf{y}}$ are proportional to real stabilizer states, their stabilizer groups are real, too. We now have to further distinguish two subcases:

- (2a) If $\mathbf{v}, \mathbf{t}_2, \dots, \mathbf{t}_n$ are linearly independent, then all columns of η are mutually orthogonal (and so are those of κ). But then, $\Pi_{\mathcal{K}_n}(C) = \eta = C'/\sqrt{2}$ for a real Clifford unitary C' and Eq. (C4) holds with $\alpha = 1/\sqrt{2}$ and $G = \{\mathbb{1}\}$.
- (2b) If $\mathbf{v}, \mathbf{t}_2, \dots, \mathbf{t}_n$ are linearly dependent, we can write $\mathbf{v} + \sum_{i=2}^n \lambda_i \mathbf{t}_i = 0$ for $\lambda_i \in \mathbb{Z}_2$. Thus

$$\begin{aligned} G_{\mathbf{y}} &\ni (-1)^{\mathbf{v} \cdot \mathbf{y}} Z(\mathbf{u}) \prod_{i=2}^n [(-1)^{\mathbf{t}_i \cdot \mathbf{y}} g_i]^{\lambda_i} \\ &= (-1)^{(\mathbf{v} + \sum_{i=2}^n \lambda_i \mathbf{t}_i) \cdot \mathbf{y}} Z(\mathbf{u}) \prod_{i=2}^n g_i^{\lambda_i} \\ &= Z(\mathbf{u}) \prod_{i=2}^n g_i^{\lambda_i} =: g \end{aligned}$$

is an element contained in every stabilizer group $G_{\mathbf{y}}$. The real stabilizer group $G := \langle g \rangle$ then simultaneously stabilizes all columns of η . Since every

column of κ is by construction a (-1) -eigenvector of g , the average over G annihilates κ . Following the argument in case 1b, we can thus again choose the real Clifford $C' = \eta + \kappa$ and $\alpha = 1$, which concludes the construction of Eq. (C4) for case 2.

Finally, we argue for a general $x = |x|e^{i\varphi}$. Since the matrix entries of C have a phase that is a multiple of $\pi/4$, φ has to be a multiple of $\pi/4$ itself, otherwise none of the matrix entries would be real and the projection would be identically zero, $\Pi_{\mathcal{K}_n}(xC) = 0$, and the claim trivially holds. Hence, we can assume that $x = |x|(-1)^{k/l}i^l$ where $k, l \in \mathbb{Z}_2$. Since the projection $\Pi_{\mathcal{K}_n}$ is $\mathbb{R}$ -linear, we have $\Pi_{\mathcal{K}_n}(|x|(-1)^{k/l}i^l C) = |x|(-1)^k \Pi_{\mathcal{K}_n}(i^l C)$. The multiplication by i simply exchanges the role of η and κ . However, all of our arguments apply to both of them, implying a decomposition of the form

$$\Pi_{\mathcal{K}_n}(i^l C) = \alpha \frac{1}{|G|} \sum_{P \in G} PC',$$

where the definition of α , G , and C' are as in above cases, possibly with η and κ exchanged. The general case now follows by multiplying this equation with $|x|(-1)^k$ and noting that we can absorb a potential minus sign in the definition of C' . ■

Lemma C3. For all Clifford unitaries $C \in \mathcal{C}_n$, we have $\xi_{\mathcal{Z}_n, \mathcal{K}_n}(\Pi_{\mathcal{Z}_n, \mathcal{K}_n}(xC)) \leq |x|^2$.

Proof. Note that $\mathcal{Z}_n$ and $\mathcal{K}_n$ commute, thus $\Pi_{\mathcal{Z}_n, \mathcal{K}_n} = \Pi_{\mathcal{K}_n} \Pi_{\mathcal{Z}_n}$ and we can modify the proof of Lemma C1 as follows. Taking the real part of $d = \Pi_{\mathcal{Z}_n}(C)$ results in real-diagonal matrix d_r with coefficients as in Eq. (C3), but with $\mathbf{a}' = 0$ and the affine support $S_r = S \cap \{\mathbf{x} \mid \mathbf{a}' \cdot \mathbf{x} = 0\}$. Now continue as in the proof of Lemma C1 and note that $D_r = 2^{s/2}(d_r + \kappa_r)$ is a real diagonal matrix, so are the elements of G . ■

APPENDIX D: COMPUTING THE CARDINALITY OF $\mathcal{D}_n/\mathcal{S}_n$ AND $\mathcal{R}\mathcal{D}_n/\mathcal{S}_n$

The quotient set $\mathcal{D}_n/\mathcal{S}_n$ denotes the set of all orbits from the action of $\mathcal{S}_n$ on $\mathcal{D}_n$ by conjugation. Its cardinality can be computed through a direct use of Burnside's theorem, i.e., by averaging the number of fixed points in $\mathcal{D}_n$ by the action of each permutation in $\mathcal{S}_n$. That is to say,

$$|\mathcal{D}_n/\mathcal{S}_n| = \frac{1}{n!} \sum_{\sigma \in \mathcal{S}_n} |\mathcal{D}_n^\sigma|,$$

where $\mathcal{D}_n^\sigma$ denotes the set of elements of $\mathcal{D}_n$ fixed by σ .

To calculate this number, we start with some considerations about $\mathcal{S}_n$. It is well known that a permutation σ can be uniquely decomposed into disjoint cyclic components,

such that its effect can be understood by the way it cycles collections of qubits. A cycle of length k will contain k qubits, say with labels $\{i_1, i_2, \dots, i_k\}$, and σ will act on it with the shifts $i_1 \rightarrow i_2, i_2 \rightarrow i_3, \dots, i_k \rightarrow i_1$, if the cycle appears in the decomposition of such a permutation. Each σ will then correspond to a different partition of the n qubit labels. However, as we will see, for the purposes of counting the number of elements in $\mathcal{D}_n^\sigma$, it only matters how many cycles m_k of length k appear on the decomposition of σ . Let $\lambda \vdash n$ (i.e. λ partition of n) denote the collection of those permutations with the same number $m_k \forall k$. The previous sum can then be rewritten as

$$|\mathcal{D}_n/\mathcal{S}_n| = \frac{1}{n!} \sum_{\lambda \vdash n} \frac{n!}{\prod_k k^{m_k} m_k!} |\mathcal{D}_n^{\sigma \in \lambda}|,$$

where the factor $n!/\prod_k k^{m_k} m_k!$ counts the number of permutations with same partition λ . [Starting from $n!$ possible permutations, we divide by $m_k!$ $\forall k$ to account for the permutation of similar cycles and by $k^{m_k} \forall k$ to account for the possible k (cyclic) relabels within each cycle of length k .]

We now compute $|\mathcal{D}_n^{\sigma \in \lambda}|$. Notice that any diagonal Clifford unitary $D \in \mathcal{D}_n$ can be written as

$$D = \prod_{j=1}^n S_j^{a_j} \prod_{k>j}^n CZ_{jk}^{b_{jk}}, \quad (\text{D1})$$

where $a_j \in \{0, 1, 2, 3\}$ and $b_{jk} \in \{0, 1\}$ establishes the presence or absence of a CZ gate between qubits j and k . This highlights that $\mathcal{D}_n$ is isomorphic to $\mathbb{Z}_4^n \times \mathbb{Z}_2^{\binom{n}{2}}$. Now, for a permutation $\sigma \in \lambda$ to fix D , a_j must be constant within each cycle of σ and b_{jk} must be constant within an orbit of pairs $\{j, k\}$ under σ . The first condition tells us that, if σ has $l = \sum_k m_k$ cycles, then there are 4^l possible choices for the coefficients a_j . The second condition informs us that there is a factor of 2 for every orbit of pairs $\{j, k\}$. The number of such orbits within a given cycle of length k is $k/2 - 1$ if k is odd and $k/2$ if it is even. On the other hand, between cycles of lengths k_1 and k_2 , the number of orbits is equal to the *greatest common divisor* (gcd) between them. Hence, denoting by $f(\sigma)$ the number of these orbits, we finally have that

$$|\mathcal{D}_n^{\sigma \in \lambda}| = 4^{\sum_k m_k} 2^{f(\sigma)},$$

with

$$f(\sigma) = \sum_k m_k \left\lfloor \frac{k}{2} \right\rfloor + k \binom{m_k}{2} + \sum_{k_1 < k_2} m_{k_1} m_{k_2} \gcd(k_1, k_2).$$

Finally, this allows us to write that

$$|\mathcal{D}_n/\mathcal{S}_n| = \sum_{\lambda \vdash n} \frac{4^{\sum_k m_k} 2^{f(\sigma \in \lambda)}}{\prod_k k^{m_k} m_k!}.$$

TABLE V. Cardinality of $\mathcal{D}_n/\mathcal{S}_n$ and $\mathcal{RD}_n/\mathcal{S}_n$ for small values of n .

n	$ \mathcal{D}_n/\mathcal{S}_n $	$ \mathcal{RD}_n/\mathcal{S}_n $
1	4	2
2	20	4
3	128	16
4	1616	117
5	32 768	1215
6	1 516 480	19 305
7	164 003 840	438 900
8	25 434 312 192	14 373 432

To conclude, we note that the same reasoning could be followed to find the cardinality of the quotient set $\mathcal{RD}_n/\mathcal{S}_n$, the difference being the reduction of possible phases to 2 instead of 4. Table V shows the cardinality of the quotient sets $\mathcal{D}_n/\mathcal{S}_n$ and $\mathcal{RD}_n/\mathcal{S}_n$ for some values of n .

APPENDIX E: ADDITIONAL REDUCTIONS VIA PERMUTATION INVARIANCE

Appendix D showed how to compute the number of orbits due to the action by the group of qubit permutations $\mathcal{S}_n$ on the group of diagonal and real-diagonal Clifford operators. This gives us an upper bound for the number of required terms in decompositions leveraging Lemma 2. However, in practice, several of these terms can be eliminated once we notice that diagonals that are equivalent by conjugation through some permutation yield the same $\mathcal{S}_n$ -projection.

Looking at Eq. (D1), we observe that conjugation by a permutation preserves the number of CZ gates. Thus, we can restrict ourselves to searching for nonequivalent unitaries within categories described by the number of CZ gates that are present on them. Table VI shows the total number of diagonal Cliffords present in each such category before any elimination has taken place. From a permutation-invariance perspective, nonequivalent diagonal Cliffords within a CZ category can be found by noticing that the problem boils down to finding nonisomorphic (under vertex-permutation) simple graphs with the same

number of edges. The results are presented in Table VII. To push things even further, an additional, internal elimination can be performed within the resulting elements accounted in Table VII, when all possible phase configurations are taken into account. The final number of terms that have to be considered when leveraging Lemma 2 for the particular case of permutation-invariant diagonal unitaries are shown in Table VIII. A direct comparison with the values presented in Table V shows us that, in this case, the upper bound is rather weak.

For small n (typically up to 10), the underlying graph problem is relatively easy and quick to solve numerically, and even ready-to-go datasets can be found on the internet. In our case, we have used the software *SageMath* [83] to yield the graphs. As a final comment, notice that we have only removed redundancies within each individual category. Because we can have equal projections across different categories, the values presented in Table VIII can still be further reduced, if necessary.

APPENDIX F: PRACTICAL CONSIDERATIONS ON CODE IMPLEMENTATION

Memory is the main bottleneck associated with the optimization problem under consideration. This is because, as mentioned in Sec. IIA, the size of the search space scales (super)exponentially with n (see Fig. 1). To give the reader an impression of the practical demands associated with the calculation of the stabilizer extent, in Table X one can check the amount of memory required to *formulate* the optimization problem for computing $\xi(C^{n-1}Z)$, using different sets of Clifford operators, namely, the entire Clifford group, the group of diagonal Clifford unitaries, and the group of real-diagonal Cliffords. This clearly highlights the beneficial impact of Theorem 1 memory-wise. Additionally, we also understand that the time it takes to *solve* the problem is not a concern.

Given this memory bottleneck, it is important to optimize its use. For the particular case of diagonal Clifford operators, by noticing that the nonzero entries of the matrices always have real and/or imaginary parts of $\{\pm 1\}$, we can use the most economic data type available, namely

TABLE VI. Number of diagonal Clifford operators for different numbers of qubits arranged into categories defined by the CZ count. The numbers in each row are filled up to the existing category for that n . The underlined bold values refer to the category with the largest number of elements for each n . The shadowed values highlight the symmetric nature of the numbers in each row; specifically, counting the number of nonisomorphic, simple graphs with k edges (for given n) is equivalent to counting nonisomorphic, simple graphs without k edges from the maximum (for that n). The last column shows the total for each n and corresponds to $|\mathcal{D}_n|$.

n	0 CZs	1 CZ	2 CZs	3 CZs	4 CZs	5 CZs	6 CZs	7 CZs	...	Total
2	16	16	—	—	—	—	—	—	—	32
3	64	192	192	64	—	—	—	—	—	512
4	256	1536	3840	5120	3840	1536	256	—	—	16 384
5	1024	10 240	46 080	122 880	215 040	258 048	215 040	122 880	...	1 048 576
6	4096	61 440	430 080	1 863 680	5 591 040	12 300 288	20 500 480	26 357 760	...	134 217 728

TABLE VII. Number of diagonal Clifford operators for different numbers of qubits arranged into categories defined by the CZ count obtained by eliminating connectivity redundancies. The notation follows that of Table VI. The last column presents the fraction of the original set of diagonal Clifford unitaries that survives the first filtering stage.

n	0 CZs	1 CZ	2 CZs	3 CZs	4 CZs	5 CZs	6CZs	7 CZs	...	Total	Fraction
2	16	16	—	—	—	—	—	—	—	32	100%
3	64	64	64	64	—	—	—	—	—	256	50%
4	256	256	512	768	512	256	256	—	—	2816	17.2%
5	1024	1024	2048	4096	6144	6144	6144	4096	...	34 816	3.3%
6	4096	4096	8192	20 480	36 864	61 440	86 016	98 304	...	638 976	0.48%

TABLE VIII. Number of different S_n -projections in each category. The notation follows that of Table VI. The last column informs the fraction of the original set of diagonal Clifford operators that survives the complete filtering procedure described in the text. We see that the reductions obtained are very impressive. If needed, these results can be improved further by considering equivalences across different categories.

n	0 CZs	1 CZ	2 CZs	3 CZs	4 CZs	5 CZs	6CZs	7 CZs	...	Total	Fraction
2	10	10	—	—	—	—	—	—	—	20	62.5%
3	20	40	40	20	—	—	—	—	—	120	23.4%
4	35	100	215	296	215	100	35	—	—	996	6.08%
5	56	200	620	1464	2384	2760	2384	1464	...	12 208	1.16%
6	84	350	1350	4380	11 226	22 530	35 734	45 106	...	241 520	0.180%

TABLE IX. Memory requirements when computing $\xi(T^{\otimes n})$ using only diagonal Clifford operators as a function of the number of qubits. The first row presents the memory needed to store a matrix whose columns are the $|\mathcal{D}_n|$ diagonal Clifford unitaries encoded as vectors using the data type *int8*. This matrix, denoted A_n , is used in the formulation of the optimization problem for diagonal unitaries. The values in this first row were calculated by hand. The second row indicates the memory needed by Gurobi to precisely formulate the entire optimization problem, as given by the size of the output *.mps* file. Finally, the last row tells us the amount of extra memory needed while actually solving the problem, i.e., the incremented value of RAM due to resolution, which has been estimated by standard *memory-profiling* for Python. Empty spaces indicate instances for which we were unable to obtain results. * indicates a value that was obtained using the Proteus supercomputer.

n	3	4	5	6	7	8	9
Integer A_n	8.19 kB	0.52 MB	67.1 MB	17.2 GB	8.80 TB	9.01 PB	18.4 EB
Problem definition	443 kB	25.0 MB	3.07 GB	475 GB*	—	—	—
Problem resolution	2.0 MB	12.5 MB	1.18 GB	—	—	—	—

TABLE X. Comparison of the computational resources required to compute $\xi(C^{n-1}Z)$ when using different decomposition sets. The presented time to solve is given by Gurobi, while the memory values are the size of the *.mps* files underlying each optimization problem. The results for $\mathcal{D}_n$ and $\mathcal{RD}_n$ employ both the strong and weak reductions, that is, Theorem 1 and Lemma 2, respectively. The value marked with * identifies an instance for which the problem was not solved because the optimizer was unable to allocate the required amount of RAM.

n	2	3	4	5	6	7
Memory, $\mathcal{C}_n$	11.5 MB	88.6 GB*	—	—	—	—
Memory, $\mathcal{D}_n$	12 kB	165 kB	2.394 MB	48.582 MB	2.69 GB	—
Memory, $\mathcal{RD}_n$	3 kB	12 kB	66 kB	686 kB	14.136 MB	572 MB
Time, $\mathcal{C}_n$	300 ms	—	—	—	—	—
Time, $\mathcal{D}_n$	10 ms	20 ms	30 ms	320 ms	10.55 s	—
Time, $\mathcal{RD}_n$	10 ms	10 ms	10 ms	30 ms	120 ms	4.32 s




FIG. 11. Selection of simple graphs of five vertices and the stabilizer extent of the corresponding unitaries generated by associating a qubit to each vertex and applying a CS gate to each pair of connected vertices. Each different value of ξ among all nonequivalent graphs defines what we call a family. Here, we depict a single representative for each family, choosing the graph with fewer edges. Note that by applying these unitaries to the state $|+\rangle^{\otimes 5}$, we obtain instances of two-uniform generalized hypergraph states of five qubits.

int8, to store them. The same observation holds for real-diagonal Cliffords whose nonzero entries are always ± 1 .

An additional observation is that the memory needed by the optimizer to actually *solve* the problem is larger than

that needed to simply *formulate* the problem. This is seen explicitly in Table IX, where we present the use of memory when computing $\xi(T^{\otimes n})$ with $\mathcal{D}_n$ as search space. The values presented in this case are without leveraging our

weak symmetry reduction results. They become much better when Lemma 2 is exploited (as can be understood by a direct comparison with Table X).

We conclude by remarking that, when we want to use the weak symmetry reduction lemma, time may become the main bottleneck, not because of the optimization problem itself, but due to the intermediate step of filtering vectors with a given symmetry (as done for presenting the results in Appendix E). Fortunately, this kind of filtering procedure benefits immensely from parallelization. For our particular case, parallelization (and compilation) implemented through the use of *Numba* [84] was sufficient, even though there is room to make our subroutines faster.

APPENDIX G: NONEQUIVALENT GRAPH STATES OF FIVE VERTICES

We considered all nonequivalent, simple graphs of five qubits and, for each, calculated the stabilizer extent of the unitary generated by associating with each edge the *CS* gate. From this thorough analysis, ten different values of the stabilizer extent emerged. Figure 11 summarizes the results, depicting the simplest graph instance associated with each of the ten values obtained.

-
- [1] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
 - [2] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, *Nature* **618**, 500 (2023).
 - [3] A. Morvan *et al.*, Phase transitions in random circuit sampling, *Nature* **634**, 328 (2024).
 - [4] T. Begušić, J. Gray, and G. K.-L. Chan, Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance, *Sci. Adv.* **10**, eadk4321 (2024).
 - [5] J. Tindall, M. Fishman, E. M. Stoudenmire, and D. Sels, Efficient tensor network simulation of IBM's Eagle kicked ising experiment, *PRX Quantum* **5**, 010308 (2024).
 - [6] D. Gottesman, Stabilizer codes and quantum error correction, PhD Thesis, California Institute of Technology, 1997.
 - [7] G. Nebe, E. M. Rains, and N. J. A. Sloane, The invariants of the Clifford groups, *Des. Codes Cryptogr.* **24**, 99 (2001).
 - [8] S. Bravyi and A. Kitaev, Universal quantum computation with ideal Clifford gates and noisy ancillas, *Phys. Rev. A* **71**, 022316 (2005).
 - [9] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
 - [10] S. Bravyi, G. Smith, and J. A. Smolin, Trading classical and quantum computational resources, *Phys. Rev. X* **6**, 021043 (2016).
 - [11] S. Bravyi and D. Gosset, Improved classical simulation of quantum circuits dominated by Clifford gates, *Phys. Rev. Lett.* **116**, 250501 (2016).
 - [12] M. Howard and E. Campbell, Application of a resource theory for magic states to fault-tolerant quantum computing, *Phys. Rev. Lett.* **118**, 090501 (2017).
 - [13] S. Bravyi, D. Browne, P. Calpin, E. Campbell, D. Gosset, and M. Howard, Simulation of quantum circuits by low-rank stabilizer decompositions, *Quantum* **3**, 181 (2019).
 - [14] J. R. Seddon, B. Regula, H. Pashayan, Y. Ouyang, and E. T. Campbell, Quantifying quantum speedups: Improved classical simulation from tighter magic monotones, *PRX Quantum* **2**, 010345 (2021).
 - [15] H. Qassim, J. J. Wallman, and J. Emerson, Clifford recompilation for faster classical simulation of quantum circuits, *Quantum* **3**, 170 (2019).
 - [16] H. Pashayan, O. Reardon-Smith, K. Korzekwa, and S. D. Bartlett, Fast estimation of outcome probabilities for quantum circuits, *PRX Quantum* **3**, 020361 (2022).
 - [17] H. Qassim, H. Pashayan, and D. Gosset, Improved upper bounds on the stabilizer rank of magic states, *Quantum* **5**, 606 (2021).
 - [18] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (W. H. Freeman and Company, San Francisco, CA, 1979).
 - [19] A. Heimendahl, F. Montealegre-Mora, F. Vallentin, and D. Gross, Stabilizer extent is not multiplicative, *Quantum* **5**, 400 (2021).
 - [20] M. Beverland, E. Campbell, M. Howard, and V. Kliuchnikov, Lower bounds on the non-Clifford resources for quantum computations, *Quantum Sci. Technol.* **5**, 035009 (2020).
 - [21] N. de Silva, M. Yin, and S. Strelchuk, Bases for optimising stabiliser decompositions of quantum states, *Quantum Sci. Technol.* **9**, 045004 (2024).
 - [22] H. Hamaguchi, K. Hamada, N. Marumo, and N. Yoshioka, Faster computation of nonstabilizerness, *Phys. Rev. Appl.* **23**, 014069 (2025).
 - [23] For instance, Hadamard and *CZ* are symmetric matrices, but $(H_1 CZ)^T = CZ^T H_1^T = CZ H_1 \neq H_1 CZ$.
 - [24] M. Heinrich and D. Gross, Robustness of magic and symmetries of the stabiliser polytope, *Quantum* **3**, 132 (2019).
 - [25] The Clifford group is not a 4-design, thus the set of operators commuting with all $C^{\otimes 4}$ is larger than the span of S_4 . If $C_n^{S_n}$ span the entire permutation-invariant subspace, there would be operators in the commutant of the unitary group that are not linear combinations of permutations, contradicting Schur-Weyl duality.
 - [26] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Quantum simulation of electronic structure with linear depth and connectivity, *Phys. Rev. Lett.* **120**, 110501 (2018).
 - [27] L. G. Valiant, Quantum circuits that can be simulated classically in polynomial time, *SIAM J. Comput.* **31**, 1229 (2002).
 - [28] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019b).
 - [29] R. Koenig and J. A. Smolin, How to efficiently select an arbitrary Clifford group element, *J. Math. Phys.* **55**, 122202 (2014).

- [30] G. Camillo, Symmetric Stabilizer Extent for Unitaries, 2025, https://github.com/Giuhcs/Symmetric_Stabilizer_Extent_for_Unitaries.
- [31] N. Rengaswamy, R. Calderbank, H. D. Pfister, and S. Kadhe, in *2018 IEEE International Symposium on Information Theory (ISIT)* (IEEE, Vail, CO, USA, 2018), pp. 791–795.
- [32] Gurobi Optimization, LLC, Gurobi Optimizer Reference Manual, 2024, <https://www.gurobi.com>.
- [33] D. E. Browne, Efficient classical simulation of the quantum Fourier transform, *New J. Phys.* **9**, 146 (2007).
- [34] Y. Nam, Y. Su, and D. Maslov, Approximate quantum Fourier transform with $O(n \log(n))$ T gates, *npj Quantum Inf.* **6**, 26 (2020).
- [35] V. Gheorghiu, M. Mosca, and P. Mukhopadhyay, T-count and T-depth of any multi-qubit unitary, *npj Quantum Inf.* **8**, 141 (2022).
- [36] M. Rossi, M. Huber, D. Bruß, and C. Macchiavello, Quantum hypergraph states, *New J. Phys.* **15**, 113022 (2013).
- [37] R. Raussendorf and H. J. Briegel, A one-way quantum computer, *Phys. Rev. Lett.* **86**, 5188 (2001).
- [38] J. Miller and A. Miyake, Hierarchy of universal entanglement in 2D measurement-based quantum computation, *npj Quantum Inf.* **2**, 16036 (2016).
- [39] Y. Takeuchi, T. Morimae, and M. Hayashi, Quantum computational universality of hypergraph states with Pauli-X and Z basis measurements, *Sci. Rep.* **9**, 13585 (2019).
- [40] D. W. Lyons, N. P. Gibbons, M. A. Peters, D. J. Upchurch, S. N. Walck, and E. W. Wertz, Local Pauli stabilizers of symmetric hypergraph states, *J. Phys. A: Math. Theor.* **50**, 245303 (2017).
- [41] G. Zhu, T. Jochym-O'Connor, and A. Dua, Topological order, quantum codes, and quantum computation on fractal geometries, *PRX Quantum* **3**, 030338 (2022).
- [42] S. Lu, X. Gao, and L.-M. Duan, Efficient representation of topologically ordered states with restricted Boltzmann machines, *Phys. Rev. B* **99**, 155136 (2019).
- [43] M. Gachechiladze, C. Budroni, and O. Gühne, Extreme violation of local realism in quantum hypergraph states, *Phys. Rev. Lett.* **116**, 070401 (2016).
- [44] N. Shettell and D. Markham, Graph states as a resource for quantum metrology, *Phys. Rev. Lett.* **124**, 110502 (2020).
- [45] O. Gühne, M. Cuquet, F. E. S. Steinhoff, T. Moroder, M. Rossi, D. Bruß, B. Kraus, and C. Macchiavello, Entanglement and nonclassical properties of hypergraph states, *J. Phys. A: Math. Theor.* **47**, 335303 (2014).
- [46] J. Huang, X. Li, X. Chen, C. Zhai, Y. Zheng, Y. Chi, Y. Li, Q. He, Q. Gong, and J. Wang, Demonstration of hypergraph-state quantum information processing, *Nat. Commun.* **15**, 2601 (2024).
- [47] Z.-W. Liu and A. Winter, Many-body quantum magic, *PRX Quantum* **3**, 020333 (2022).
- [48] Y. Zhou and A. Hamma, Entanglement of random hypergraph states, *Phys. Rev. A* **106**, 012410 (2022).
- [49] R. Sarkar, S. Dutta, S. Banerjee, and P. K. Panigrahi, Phase squeezing of quantum hypergraph states, *J. Phys. B: At., Mol. Opt. Phys.* **54**, 135501 (2021).
- [50] J. Chen, Y. Yan, and Y. Zhou, Magic of quantum hypergraph states, *Quantum* **8**, 1351 (2024).
- [51] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Average-case complexity versus approximate simulation of commuting quantum computations, *Phys. Rev. Lett.* **117**, 080501 (2016).
- [52] X. Ni and M. van den Nest, Commuting quantum circuits: Efficiently classical simulations versus hardness results, *Quantum Inf. Comput.* **13**, 54 (2013).
- [53] V. Vandrale, Quantum binary field multiplication with subquadratic Toffoli gate count and low space-time cost, [arXiv:2501.16136](https://arxiv.org/abs/2501.16136) [quant-ph].
- [54] F. Liu, T. Fan, G. Shu, C. Zhang, W. Wang, X. Qi, X. Zhu, H. Yang, and Y. Fei, Applications of the CCZS gate in quantum circuit synthesis, *Quantum Sci. Technol.* **10**, 035057 (2025).
- [55] P. Schauss, Quantum simulation of transverse Ising models with Rydberg atoms, *Quantum Sci. Technol.* **3**, 023001 (2018).
- [56] B.-W. Li, Y.-K. Wu, Q.-X. Mei, R. Yao, W.-Q. Lian, M.-L. Cai, Y. Wang, B.-X. Qi, L. Yao, L. He, Z.-C. Zhou, and L.-M. Duan, Probing critical behavior of long-range transverse-field Ising model through quantum Kibble-Zurek mechanism, *PRX Quantum* **4**, 010302 (2023).
- [57] M. L. Baez, M. Goihl, J. Haferkamp, J. Bermejo-Vega, M. Gluza, and J. Eisert, Dynamical structure factors of dynamical quantum simulators, *Proc. Natl. Acad. Sci.* **117**, 26123 (2020).
- [58] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, *Nature* **551**, 601 (2017).
- [59] R. Islam, E. E. Edwards, K. Kim, S. Korenblit, C. Noh, H. Carmichael, G.-D. Lin, L.-M. Duan, C.-C. Joseph Wang, J. K. Freericks, and C. Monroe, Onset of a quantum phase transition with a trapped ion quantum simulator, *Nat. Commun.* **2**, 377 (2011).
- [60] J. G. Bohnet, B. C. Sawyer, J. W. Britton, M. L. Wall, A. M. Rey, M. Foss-Feig, and J. J. Bollinger, Quantum spin dynamics and entanglement generation with hundreds of trapped ions, *Science* **352**, 1297 (2016).
- [61] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, *Nature* **551**, 579 (2017).
- [62] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A: At., Mol. Opt. Phys.* **86**, 032324 (2012).
- [63] D. Horsman, A. G. Fowler, S. Devitt, and R. V. Meter, Surface code quantum computing by lattice surgery, *New J. Phys.* **14**, 123011 (2012).
- [64] J. Bermejo-Vega and M. Van Den Nest, Classical simulations of Abelian-group normalizer circuits with intermediate measurements, *Quantum Inf. Comput.* **14**, 181 (2014).
- [65] J. Bermejo-Vega, C. Y.-Y. Lin, and M. Van Den Nest, Normalizer circuits and a Gottesman-Knill theorem for

- infinite-dimensional systems, *Quantum Inf. Comput.* **16**, 361 (2016).
- [66] J. Bermejo-Vega, Normalizer circuits and quantum computation, PhD Thesis, Technische Universität München, 2016.
- [67] D. V. Else, I. Schwarz, S. D. Bartlett, and A. C. Doherty, Symmetry-protected phases for measurement-based quantum computation, *Phys. Rev. Lett.* **108**, 240505 (2012).
- [68] R. Raussendorf, C. Okay, D.-S. Wang, D. T. Stephen, and H. P. Nautrup, Computationally universal phase of quantum matter, *Phys. Rev. Lett.* **122**, 090501 (2019).
- [69] D. T. Stephen, H. P. Nautrup, J. Bermejo-Vega, J. Eisert, and R. Raussendorf, Subsystem symmetries, quantum cellular automata, and computational phases of quantum matter, *Quantum* **3**, 142 (2019).
- [70] S. Abramsky and A. Brandenburger, The sheaf-theoretic structure of non-locality and contextuality, *New J. Phys.* **13**, 113036 (2011).
- [71] S. Abramsky, R. S. Barbosa, K. Kishida, R. Lal, and S. Mansfield, in *24th EACSL Annual Conference on Computer Science Logic (CSL 2015)*, *Leibniz International Proceedings in Informatics (LIPIcs)* (Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2015), Vol. 41, pp. 211–228.
- [72] R. Raussendorf, C. Okay, M. Zurel, and P. Feldmann, The role of cohomology in quantum computation with magic states, *Quantum* **7**, 979 (2023).
- [73] F. C. R. Peres and E. F. Galvão, Quantum circuit compilation and hybrid computation using Pauli-based computation, *Quantum* **7**, 1126 (2023).
- [74] S. Mansfield and E. Kashefi, Quantum advantage from sequential-transformation contextuality, *Phys. Rev. Lett.* **121**, 230401 (2018).
- [75] P.-E. Emeriau, M. Howard, and S. Mansfield, Quantum advantage in information retrieval, *PRX Quantum* **3**, 020307 (2022).
- [76] R. W. Spekkens, Contextuality for preparations, transformations, and Unsharp Measurements, *Physical Review A* **71**, 052108 (2005).
- [77] D. Schmid, Contextual advantage for state discrimination, *Phys. Rev. X* **8**, 011015 (2018).
- [78] C. Budroni, G. Fagundes, and M. Kleinmann, Memory cost of temporal correlations, *New J. Phys.* **21**, 093018 (2019).
- [79] C. Budroni, Contextuality, memory cost and non-classicality for sequential measurements, *Philos. Trans. R. Soc., A* **377**, 20190141 (2019).
- [80] D. Gross and M. Van den Nest, The LU-LC conjecture, diagonal local operations and quadratic forms over GF(2), *Quantum Inf. Comput.* **8**, 263 (2008).
- [81] J. Dehaene and B. De Moor, Clifford group, stabilizer states, and linear and quadratic operations over GF(2), *Phys. Rev. A* **68**, 042318 (2003).
- [82] M. Heinrich, On stabiliser techniques and their application to simulation and certification of quantum devices, Ph.D. thesis, Universität zu Köln, Cologne, 2021.
- [83] W. Stein *et al.*, Sage Mathematics Software (Version x.y.z), The Sage Development Team, <http://www.sagemath.org>.
- [84] S. K. Lam, A. Pitrou, and S. Seibert, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM '15* (Association for Computing Machinery, New York, NY, USA, 2015).