

Exact universality of symmetric quantum circuits for arbitrary permutation actions

Abstract

We prove exact universality of symmetric quantum circuits for every finite permutation action on qubits. In the model of arbitrary one-qubit gates and reversible unbounded-fan-in threshold gates, a unitary has a symmetric circuit with clean workspace if and only if it commutes with the given action. Symmetry is required at every layer: each permutation in the given action must admit an extension to all wires that preserves every layer as a set of commuting gates. This answers the arbitrary-group universality question of Castro-Silva, Gur, and Strelchuk (ITCS 2026). For $n \geq 1$ active qubits, the nonuniform construction uses $O(n8^n)$ gates and $O(4^n)$ workspace qubits. The main difficulty is that equivariance of a Pauli expansion makes its coefficients constant on orbits, but does not make its individual terms equivariant. We assign a selector wire to each Pauli word and let the group act on these wires. Symmetric preparation and commuting incidence layers then give a block encoding of the target unitary. Exact oblivious amplification removes the normalization using reflections that tolerate permuted workspace, returning all ancillary qubits to their initial zero state and retaining the target global phase.

Note. This paper was generated entirely by AI, including MiMo, using an automated research pipeline developed by Chenghua Liu and Hanyu Li.

1 Introduction

A unitary can respect a permutation of its qubits while its elementary gate decomposition distinguishes them. The swap of two qubits is a basic example: it commutes with interchanging the two wires, but its usual three-CNOT realization starts with a gate whose control and target are exchanged by that permutation. The first layer is therefore not preserved. Operator equivariance alone does not guarantee the wire automorphisms required of a symmetric circuit.

Castro-Silva, Gur, and Strelchuk formalize this requirement using circuits of arbitrary one-qubit gates and reversible threshold gates, organized into layers of pairwise commuting gates [1]. Every permutation of the active wires in the prescribed group must extend to a permutation of all circuit wires preserving each layer. Workspace qubits may be permuted among themselves. Their universality theorem handles partition groups $S_{n_1} \times \cdots \times S_{n_t}$, acting on disjoint blocks, and they leave universality for arbitrary permutation actions open [1, Theorem 12 and Section 1.5]. We prove that commutation with the action is sufficient in every case, with exact implementation on the full Hilbert space and workspace returned to zero.

The partition-group proof combines simple equivariant unitaries by a symmetric linear-combination construction [1, Lemma 10]. A Pauli expansion supplies simple unitaries for an arbitrary action, but loses the hypothesis that each summand is equivariant. Taking orbit sums does not repair this: even for two interchangeable qubits, the exchanged words $\sigma_1 \otimes I_2$ and $I_2 \otimes \sigma_1$ have a sum with a nontrivial kernel. Orbit sums supply equivariant operators but do not give the unitary summands required by the linear-combination lemma.

We retain the individual Pauli words and give each one a selector wire. Under an active-wire permutation, the selector wires are permuted by the induced action on Pauli words. Equal coefficients

on each orbit then allow a symmetric superposition of one-hot selector states. Selection respects the joint action on the active and selector wires, although selecting an individual word need not respect the active action alone. This is the role of permutable workspace in the proof: it carries the labels needed to apply the individual summands without choosing an orbit representative.

There are two further circuit constraints. First, controlled Pauli words can fail to commute when several selector bits are one. Their behavior on those inputs matters even though the prepared selector is one-hot, because commutation within a layer is an operator identity on the full Hilbert space. We split selection according to the three nonidentity Pauli symbols; each stage uses a commuting CNOT incidence layer between uniform local basis changes. Second, preparation and selection produce a normalized block of the target unitary, rather than the unitary itself. Its normalization is independent of the active input, so exact oblivious amplification applies. The required all-zero workspace test is preserved by every workspace permutation, allowing amplification to use the same wire action as selection.

1.1 Main result

Let X be a finite set of active wires, let $n = |X|$, and write $\mathcal{H}_X = (\mathbb{C}^2)^{\otimes X}$. The symmetric group on X is denoted by S_X . Its tensor-factor permutation representation is

$$R_X(\gamma) \bigotimes_{x \in X} |b_x\rangle_x = \bigotimes_{x \in X} |b_{\gamma^{-1}x}\rangle_x \quad (\gamma \in S_X).$$

We call a unitary U Γ -equivariant if $[U, R_X(\gamma)] = 0$ for all $\gamma \in \Gamma \leq S_X$.

A clean implementation returns every workspace qubit to $|0\rangle$ for every active input, with the target global phase retained. We count each one-qubit or unbounded-fan-in threshold gate once, and write $s(\mathcal{C})$ and $a(\mathcal{C})$ for the number of gates and workspace qubits. Section 2 gives the precise circuit conventions.

We use the phase-free Hermitian Pauli basis $\mathcal{P}_X = \{I_2, \sigma_1, \sigma_2, \sigma_3\}^{\otimes X}$, where

$$\sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Theorem 1.1 (Universality for finite permutation actions). *Let X be finite and $\Gamma \leq S_X$. A unitary on $\mathcal{H}_X$ has an exact clean-workspace Γ -symmetric circuit if and only if it commutes with $R_X(\Gamma)$.*

For $n = |X| \geq 1$, write $U = \sum_{P \in \mathcal{P}_X} \alpha_P P$, and put

$$J = \{P \in \mathcal{P}_X : \alpha_P \neq 0\}, \quad M = |J|, \quad \lambda = \sum_{P \in J} |\alpha_P|.$$

If U is Γ -equivariant, a realizing circuit can be chosen with

$$\begin{aligned} a(\mathcal{C}) &\leq M + 5 \leq 4^n + 5, \\ s(\mathcal{C}) &= O(nM\lambda). \end{aligned} \tag{1}$$

In particular, since $\lambda \leq \sqrt{M}$, the gate count is $O(nM^{3/2})$ and hence $O(n8^n)$. These bounds hold both when equality gates are counted as single gates and when they are expanded into threshold and one-qubit gates. If $X = \emptyset$, one fixed clean workspace qubit and one one-qubit gate suffice.

The theorem is nonuniform: the one-qubit parameters depend on the exact Pauli coefficients, and the size bound does not include their classical computation or representation. In the worst case

the circuit has exponential size. The gate model also counts an unbounded-fan-in threshold test as one gate, independently of its number of control wires. These are the resources under which the universality statement holds.

Several related notions of symmetry impose different constraints. Representation-adapted layers and their learning-theoretic properties are studied in equivariant quantum neural networks [4, 5]. Symmetry-adapted projections for many-body simulation are studied by Bastidas et al. [6]. Our concern is the realizability of an arbitrary equivariant unitary by elementary gates whose individual layers admit the required wire automorphisms, rather than the construction of symmetry-adapted operators alone.

Universality also depends on the available interactions. Kazi, Larocca, and Cerezo identify relative-phase restrictions for fully permutation-equivariant generators of bounded bodyness [7]. Deliyannis and Marvian prove realizability of all fully permutation-invariant unitaries, up to global phase, using global fields and a Tavis–Cummings interaction with a bosonic mode returned to its initial state [8, Theorem 1]. Local symmetric-gate models can exhibit both nonuniversality and its removal by ancillary resources [9]. In our model, symmetry may permute the elementary gates within a commuting layer, with the ancillary wires participating in the same action. The proof below constructs those layer automorphisms explicitly. Numbered citations to [1] refer to its full arXiv version 2.

2 Circuit symmetry and selector preparation

We first fix the circuit conventions, then construct the selector state used in the Pauli expansion. Besides arbitrary gates in $U(2)$, the elementary gates are reversible threshold gates $\text{Th}_t^{S,h}$, where the control set S does not contain the head wire h . On a computational-basis state, this gate toggles h if and only if $\sum_{u \in S} b_u \geq t$, leaving all other bits unchanged. We also write $\text{Eq}_t^{S,h}$ for the corresponding equality test $\sum_{u \in S} b_u = t$. For $0 \leq t \leq |S|$,

$$\text{Eq}_t^{S,h} = \text{Th}_t^{S,h} \text{Th}_{t+1}^{S,h}, \quad (2)$$

where $\text{Th}_0^{S,h}$ is the one-qubit bit flip and $\text{Th}_{|S|+1}^{S,h} = I$. Thus an equality gate abbreviates at most two allowed gates and uses no additional qubit. The control sets are unbounded in size; the gate count does not count control incidences separately. Throughout, operator products act from right to left.

Definition 2.1 (Symmetric circuits and clean implementation). A circuit on disjoint active and workspace wire sets X and W is a sequence of layers, each consisting of pairwise commuting elementary gates. It is Γ -*symmetric* if, for every $\gamma \in \Gamma$, there is a permutation $\tilde{\gamma} \in S_{X \sqcup W}$ whose restriction to X is γ , such that relabeling wires by $\tilde{\gamma}$ maps each layer to itself as a set of labeled gates. In particular, gate parameters, control sets, and head wires are respected under this relabeling.

The circuit $\mathcal{C}$ *exactly implements* U *with clean workspace* if

$$\mathcal{C}(|\psi\rangle_X |0^W\rangle) = (U|\psi\rangle)_X |0^W\rangle \quad \text{for every } |\psi\rangle \in \mathcal{H}_X. \quad (3)$$

Its size $s(\mathcal{C})$ is its number of gates, and its workspace size is $a(\mathcal{C}) = |W|$.

The symmetry condition is Definition 12 of [1]; cleanliness is the additional implementation requirement in (3). A lift $\tilde{\gamma}$ necessarily preserves W setwise, but it need not fix its wires individually. The definition asks for a lift separately for each γ , without requiring a homomorphism. Our construction supplies a single explicit group action on all wires. This also ensures that the same lifts work for every subcircuit we concatenate.

Lemma 2.2. *Every unitary implemented by a clean Γ -symmetric circuit is Γ -equivariant.*

Proof. Fix γ and an associated lift $\tilde{\gamma}$. Relabeling permutes the gates in each layer, and these gates commute. Hence the circuit unitary commutes with $R_{X \sqcup W}(\tilde{\gamma})$. Since the lift preserves X and W separately and fixes $|0^W\rangle$, applying this commutation identity to $|\psi\rangle|0^W\rangle$ and using (3) gives $UR_X(\gamma)|\psi\rangle = R_X(\gamma)U|\psi\rangle$. $\square$

For the converse, the Pauli coefficients determine both the selector amplitudes and the phases to be applied during selection. Assume $n \geq 1$. Orthogonality of the Hermitian Pauli basis gives

$$U = \sum_{P \in \mathcal{P}_X} \alpha_P P, \quad \alpha_P = 2^{-n} \text{Tr}(PU). \quad (4)$$

For $\gamma \in \Gamma$, define $\gamma P = R_X(\gamma)PR_X(\gamma)^\dagger$. This action only permutes tensor factors; it introduces no signs.

Lemma 2.3. *If U is Γ -equivariant, then $\alpha_{\gamma P} = \alpha_P$ for every $P \in \mathcal{P}_X$ and $\gamma \in \Gamma$. Moreover, $\sum_P |\alpha_P|^2 = 1$.*

Proof. Cyclicity of trace and commutation with U give

$$\begin{aligned} \alpha_{\gamma P} &= 2^{-n} \text{Tr}(R_X(\gamma)PR_X(\gamma)^\dagger U) \\ &= 2^{-n} \text{Tr}(PR_X(\gamma)^\dagger UR_X(\gamma)) = \alpha_P. \end{aligned}$$

Since $2^{-n} \text{Tr}(P^\dagger Q) = \delta_{P,Q}$, expanding $2^{-n} \text{Tr}(U^\dagger U)$ in this basis yields $\sum_P |\alpha_P|^2 = 2^{-n} \text{Tr}(I) = 1$. $\square$

It follows that J is Γ -stable and

$$1 \leq \lambda \leq \sqrt{M}. \quad (5)$$

Indeed, the lower bound is $\|\alpha\|_1 \geq \|\alpha\|_2 = 1$, and the upper bound is Cauchy–Schwarz on the M nonzero coefficients. Define

$$d_P = \sqrt{\frac{|\alpha_P|}{\lambda}}, \quad \omega_P = \frac{\alpha_P}{|\alpha_P|} \quad (P \in J).$$

Both parameters are constant on each orbit, and $\sum_{P \in J} d_P^2 = 1$.

Introduce selector wires $Q = \{q_P : P \in J\}$, with group action $q_P \mapsto q_{\gamma P}$. Let $|e_P\rangle$ be the computational-basis state whose only nonzero bit is on q_P . We will prepare

$$|\chi\rangle = \sum_{P \in J} d_P |e_P\rangle. \quad (6)$$

Each distinct Pauli word has one selector wire, irrespective of its stabilizer. Thus there is no multiplicity from indexing by group elements.

The amplitudes in (6) have no preferred order within an orbit, so we prepare them through simultaneous local rotations and a Hamming-weight test. A product of one-qubit states with amplitude ratios d_P already has the desired relative amplitudes on weight one. The probability of that subspace lies between e^{-1} and $1/2$, regardless of M . We lower it to $\sin^2(\pi/10)$ with one auxiliary rotation, making two amplification steps exact. We record the reflection identity first; its isometry formulation will later amplify every active input at once.

Lemma 2.4 (Two-reflection identity). *Let Π be an orthogonal projector on a finite-dimensional Hilbert space, and let S, G be isometries into that space from the same Hilbert space $\mathcal{K}$. Suppose that $\Pi S = \sin \vartheta G$, where $0 < \vartheta < \pi/2$. Put*

$$B = \frac{S - \sin \vartheta G}{\cos \vartheta}, \quad R_G = I - 2\Pi, \quad R_S = I - 2SS^\dagger.$$

Then B is an isometry orthogonal to G , and for every integer $j \geq 0$,

$$(R_S R_G)^j S = (-1)^j (\sin((2j+1)\vartheta)G + \cos((2j+1)\vartheta)B). \quad (7)$$

Proof. Applying Π again gives $\Pi G = G$. Hence $G^\dagger S = \sin \vartheta I$, and substitution into the definition of B gives $G^\dagger B = 0$, $B^\dagger B = I$, and $\Pi B = 0$. In the orthogonal coordinates given by G and B , the starting isometry is represented by $\begin{pmatrix} \sin \vartheta \\ \cos \vartheta \end{pmatrix} \otimes I_{\mathcal{K}}$, and

$$R_S R_G = - \begin{pmatrix} \cos(2\vartheta) & \sin(2\vartheta) \\ -\sin(2\vartheta) & \cos(2\vartheta) \end{pmatrix} \otimes I_{\mathcal{K}}$$

on their joint range. Multiplying this matrix sends $\begin{pmatrix} \sin t \\ \cos t \end{pmatrix}$ to $-\begin{pmatrix} \sin(t+2\vartheta) \\ \cos(t+2\vartheta) \end{pmatrix}$. Iteration proves (7). $\square$

We use the rotation convention

$$R_y(t) = \begin{pmatrix} \cos(t/2) & -\sin(t/2) \\ \sin(t/2) & \cos(t/2) \end{pmatrix}, \quad |-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}.$$

In particular, $R_y(-\pi/2)|0\rangle = |-\rangle$ and $R_y(\pi/2)|-\rangle = |0\rangle$. A reversible predicate test with head in $|-\rangle$ leaves that head unchanged and contributes a minus sign exactly when the predicate holds. We use this phase-kickback implementation for all reflections below.

Lemma 2.5 (Symmetric one-hot preparation). *There is a circuit B on Q and three additional wires such that*

$$B|0^Q\rangle|0^3\rangle = |\chi\rangle|0^3\rangle, \quad s(B) = O(M).$$

Every permutation of the selector wires within each Γ -orbit, extended by fixing the three additional wires, preserves every layer of B .

Proof. Apply $R_y(2 \arctan d_P)$ to each q_P , and denote the resulting one-layer circuit by A_0 . Then

$$A_0|0^Q\rangle = c_0 \bigotimes_{P \in J} (|0\rangle + d_P|1\rangle), \quad c_0 = \prod_{P \in J} (1 + d_P^2)^{-1/2}.$$

Let P_1 project onto the weight-one subspace of Q . We have $P_1 A_0|0^Q\rangle = c_0|\chi\rangle$, so the success probability is $p_0 = c_0^2$. The inequalities

$$2 = 1 + \sum_P d_P^2 \leq \prod_P (1 + d_P^2) \leq \exp\left(\sum_P d_P^2\right) = e$$

give $e^{-1} \leq p_0 \leq 1/2$.

The weight-one probability need not correspond to an exact integer number of Grover rotations. The adjustment wire r corrects this, while f stores a predicate and z supplies phase kickback. Set $\vartheta_0 = \pi/10$, and choose $\beta_0 \in (0, \pi)$ by

$$\sin(\beta_0/2) = \frac{\sin \vartheta_0}{\sqrt{p_0}}. \quad (8)$$

This is possible because $0 < \sin(\pi/10) < 1/2 < e^{-1/2} \leq \sqrt{p_0}$. Let T_0 apply A_0 on Q and $R_y(\beta_0)$ on r , acting trivially on f, z . On Q, f, r , define

$$|s_0\rangle = T_0|0^Q, 0_f, 0_r\rangle, \quad |g_0\rangle = |\chi\rangle|0_f\rangle|1_r\rangle.$$

The projector onto weight one on Q , zero on f , and one on r sends $|s_0\rangle$ to $\sin \vartheta_0 |g_0\rangle$.

Here are explicit gates for the two required reflections. Put

$$D = \text{Eq}_1^{Q,f}, \quad F_G = D \text{Eq}_2^{\{f,r\},z} D, \quad F_0 = \text{Eq}_0^{Q \cup \{f,r\},z}. \quad (9)$$

On the subspace where $f = 0$ and $z = |-\rangle$, the rightmost D computes the weight-one predicate into f , the middle gate contributes a minus sign precisely when this predicate and $r = 1$ both hold, and the leftmost $D = D^\dagger$ returns f to zero. Thus F_G is the good reflection on this invariant subspace. With $z = |-\rangle$, the gate F_0 reflects about $|0^Q, 0_f, 0_r\rangle$; consequently, $T_0 F_0 T_0^\dagger$ reflects about $|s_0\rangle$. Both reflections preserve the subspace $f = 0$.

Lemma 2.4, applied within that subspace with $\mathcal{K} = \mathbb{C}$, now gives

$$(T_0 F_0 T_0^\dagger F_G)^2 (|s_0\rangle |-\rangle_z) = |g_0\rangle |-\rangle_z,$$

since $5\vartheta_0 = \pi/2$ and $(-1)^2 = 1$. In particular, an explicit preparation circuit is

$$B = \sigma_{1,r} R_y(\pi/2)_z (T_0 F_0 T_0^\dagger F_G)^2 T_0 R_y(-\pi/2)_z. \quad (10)$$

The final bit flip resets r , the rotation resets z , and the computed predicate on f has already been uncomputed. Since the two reflection steps contribute phase $(-1)^2 = 1$, this is the exact preparation identity, including its phase.

To verify the circuit symmetry, implement the rotations on Q as one layer, each fixed-wire rotation as its own layer, and each equality gate in (9) as its own layer. The parameters d_P are equal within each orbit, so the layers of A_0 and $A_0^\dagger$ are preserved under all orbitwise selector permutations. Every equality control set either contains all of Q and some fixed wires or consists entirely of fixed wires. Its head is fixed. These permutations therefore preserve every layer of (10). There are five occurrences of A_0 or $A_0^\dagger$ and only a constant number of other gates, giving $s(B) = O(M)$. Expanding the equality gates by (2), using a separate layer for each factor, preserves the same symmetries and count up to a constant factor. $\square$

This is an explicit instance of exact amplitude amplification [2, Section 2.1], with the symmetry preserved as in [1, Lemma 6]. Here the initial product rotations and predicate tests supply the entire state-preparation circuit, together with its action on the workspace wires.

3 A symmetric Pauli block encoding

A prepared one-hot selector asks for exactly one Pauli word to be applied. The circuit implementing this request must still be a unitary on every selector input. In particular, controlled Pauli words cannot all be put in one commuting layer: two anticommuting words still anticommute on the branch where both of their controls are one. Ordering these whole-word gates arbitrarily would distinguish labels that the group may exchange. We instead order the three Pauli symbols, which are fixed by wire permutations, and put every incidence of a given symbol in the same commuting stage.

For $s \in \{1, 2, 3\}$, let L_s be the unitary of the incidence layer

$$\mathcal{L}_s = \{\text{CNOT}(q_P \longrightarrow x) : P \in J, x \in X, P_x = \sigma_s\}. \quad (11)$$

Here P_x denotes the factor of P at wire x . All controls lie in Q and all targets in X , so no target is a control of another gate. Gates with a common target apply the same bit flip, and gates with distinct targets also commute. Thus $\mathcal{L}_s$ is a legal layer. Each CNOT is a threshold gate with one control.

Let H be the Hadamard gate and $S = \text{diag}(1, i)$. The identities $H\sigma_1 H = \sigma_3$ and $S\sigma_1 S^\dagger = \sigma_2$ give controlled-symbol stages

$$C_1 = L_1, \quad C_3 = H^{\otimes X} L_3 H^{\otimes X}, \quad C_2 = S^{\otimes X} L_2 (S^\dagger)^{\otimes X}.$$

Each uniform tensor product is a layer of identical one-qubit gates on X . Define the selector phase layer and the selection circuit by

$$\Delta = \prod_{P \in J} \text{diag}(1, \omega_P)_{q_P}, \quad \text{SELECT} = C_2 C_3 C_1 \Delta. \quad (12)$$

For a one-hot selector, only the incidences belonging to its selected word are active. The three stages then act on disjoint positions in that word, so their order introduces no Pauli phase. It follows that

$$\text{SELECT}(|\psi\rangle_X |e_P\rangle_Q) = \omega_P(P|\psi\rangle)_X |e_P\rangle_Q. \quad (13)$$

Equation (13) specifies the action needed for the block encoding. The same elementary circuit defines a unitary on all other selector inputs, and every incidence layer remains a legal commuting layer.

For every $\gamma \in \Gamma$, use the lift

$$x \mapsto \gamma x, \quad q_P \mapsto q_{\gamma P}, \quad \text{all three preparation wires fixed.} \quad (14)$$

The phase layer is preserved because $\omega_{\gamma P} = \omega_P$. The identity $(\gamma P)_{\gamma x} = P_x$ shows that each incidence layer is preserved. The uniform basis-change layers are preserved as well. Lemma 2.5 supplies precisely the same lift for B , and also for $B^\dagger$. More generally, a lift preserving a circuit's layers preserves its inverse layers: taking adjoints preserves commutativity and reverses only the order of the layers.

Let W_0 contain the M selectors and the three preparation wires. Extend SELECT trivially to the latter wires, and put

$$A = (I_X \otimes B^\dagger) \text{SELECT} (I_X \otimes B). \quad (15)$$

This is a Γ -symmetric circuit under (14), which in fact defines a homomorphism on all its wires. Its all-zero workspace block is

$$\begin{aligned} & (I_X \otimes \langle 0^{W_0} |) A (I_X \otimes |0^{W_0}\rangle) \\ &= (I_X \otimes \langle \chi |) \text{SELECT} (I_X \otimes |\chi\rangle) \\ &= \sum_{P \in J} d_P^2 \omega_P P = \frac{1}{\lambda} \sum_{P \in J} \alpha_P P = \frac{U}{\lambda}. \end{aligned} \quad (16)$$

The first equality only uses the action of B on the all-zero state and its adjoint. Thus the unspecified action of the preparation circuit on other inputs does not enter the block identity. The normalization λ leaves a nonzero-workspace component when $\lambda > 1$; its norm is determined by unitarity, as follows.

In that case, with $\Pi_0 = I_X \otimes |0^{W_0}\rangle\langle 0^{W_0}|$, unitarity of A and U gives, for each unit vector $|\psi\rangle$,

$$A(|\psi\rangle|0^{W_0}\rangle) = \lambda^{-1} U|\psi\rangle|0^{W_0}\rangle + \sqrt{1 - \lambda^{-2}} |\Phi_\psi\rangle, \quad (17)$$

where

$$|\Phi_\psi\rangle = \frac{(I - \Pi_0)A(|\psi\rangle|0^{W_0}\rangle)}{\sqrt{1 - \lambda^{-2}}}$$

is a unit vector in $\ker \Pi_0$. In particular, the success probability λ^{-2} is independent of the input.

There are at most nM CNOT incidences across the three symbol stages, M selector phases, and $4n$ basis-change gates. Together with Lemma 2.5, this yields

$$|W_0| = M + 3, \quad s(A) = O(nM). \quad (18)$$

4 From block encoding to exact universality

The block encoding succeeds with the same probability on every active input. This allows amplification without a reflection about the unknown input state. This is the input-independent setting of oblivious amplitude amplification used for linear combinations of unitaries [3, Section I]. The symmetric oblivious-amplification statement in [1, Lemma 7] assumes workspace fixed pointwise, whereas our selector wires are permuted. The proof only needs a reflection about the whole zero-workspace subspace. Its equality test uses the workspace as an unordered control set, so it remains symmetric under precisely those permutations. We give the construction with its exact phase to obtain literal equality with the target unitary.

Lemma 4.1 (Exact symmetric oblivious amplification). *Let A be a Γ -symmetric circuit on active wires X and workspace W , and let V be a unitary on $\mathcal{H}_X$. Suppose that, for a specified $p \in (0, 1)$,*

$$(I_X \otimes \langle 0^W |) A (I_X \otimes | 0^W \rangle) = \sqrt{p} V. \quad (19)$$

Then V has an exact clean-workspace Γ -symmetric circuit using $|W| + 2$ workspace qubits and $O((s(A) + 1)/\sqrt{p})$ gates. Every lift witnessing the symmetry of A extends to a lift for this circuit by fixing the two new wires.

Proof. Add an adjustment wire r and a phase-kickback wire z , both fixed under every lift. Set

$$K = \left\lceil \frac{\pi}{4\sqrt{p}} \right\rceil, \quad \vartheta = \frac{\pi}{4K + 2}, \quad \cos(\zeta/2) = \frac{\sin \vartheta}{\sqrt{p}}. \quad (20)$$

The angle $\zeta \in (0, \pi)$ exists uniquely because

$$0 < \frac{\sin \vartheta}{\sqrt{p}} < \frac{\vartheta}{\sqrt{p}} < \frac{\pi}{4K\sqrt{p}} \leq 1.$$

Let T apply A on $X \sqcup W$ and $R_y(\zeta)$ on r , and act trivially on z . For the moment omit the last wire, and define the zero-workspace embedding

$$E : \mathcal{H}_X \longrightarrow \mathcal{H}_{X \sqcup W \sqcup \{r\}}, \quad E|\psi\rangle = |\psi\rangle |0^W\rangle |0\rangle_r.$$

Put $\Pi = EE^\dagger$, $S = TE$, and $G = EV$. The maps S and G are isometries because A and V are unitary. The block identity (19) gives

$$E^\dagger TE = \sqrt{p} \cos(\zeta/2) V = \sin \vartheta V, \quad \Pi S = \sin \vartheta G.$$

The scalar $\sin \vartheta$ is the same on all of $\mathcal{H}_X$. Consequently the two isometric copies of $\mathcal{H}_X$ rotate through a common angle, and Lemma 2.4 applies simultaneously to every input, with domain $\mathcal{K} = \mathcal{H}_X$.

Initialize z in $|-\rangle$ and set

$$F = \text{Eq}_0^{W \cup \{r\}, z}. \quad (21)$$

On this invariant z -subspace, F implements $I - 2\Pi$, while $TFT^\dagger$ implements $I - 2\Pi T^\dagger = I - 2SS^\dagger$. Consequently, (7) and $(2K + 1)\vartheta = \pi/2$ imply

$$(TFT^\dagger F)^K S = (-1)^K G,$$

where the unchanged factor $|- \rangle_z$ is suppressed. The good output has W and r exactly zero. Apply $R_y(\pi/2)$ to z to return it to zero as well. Finally, let Ω_K be the identity if K is even and the one-qubit gate $-I_2$ on r if K is odd. An explicit circuit for literal V , including its global phase, is therefore

$$\mathcal{C}_V = \Omega_K R_y(\pi/2)_z (TFT^\dagger F)^K T R_y(-\pi/2)_z. \quad (22)$$

The scalar correction does not disturb a clean qubit and uses no additional wire. No unknown input-dependent reflection is required.

It remains to check the syntax. Fix a lift for A . Its restriction to W is a permutation, so it preserves the whole control set $W \cup \{r\}$ in (21); the head z is fixed. Every fixed-wire gate and every layer of A or $A^\dagger$ is also preserved. Implement T by the layers of A , followed by a separate layer for its rotation on r . The circuit in (22) then has the asserted layer automorphisms. The same check applies to each factor when F is expanded using (2).

There are $2K + 1$ calls to A or $A^\dagger$, together with $O(K + 1)$ other gates. Since $K = O(p^{-1/2})$, this gives the claimed size. Only r and z have been added to W . $\square$

Applying the lemma to the Pauli block encoding completes the construction. The unit-normalization case can be handled directly, so no singular amplification parameters are needed.

Proof of Theorem 1.1. Necessity is Lemma 2.2. For sufficiency, first assume $n \geq 1$ and use the block circuit A from (15). If $\lambda > 1$, apply Lemma 4.1 with $V = U$ and $p = \lambda^{-2}$. Its hypotheses follow from (16), and the lifts (14) fix the two new wires. The resulting circuit exactly implements U and returns all workspace to zero, including the selector and both sets of amplification wires.

If $\lambda = 1$, amplification is unnecessary. For a unit vector $|\psi\rangle$, the zero-workspace component of $A(|\psi\rangle|0^{W_0}\rangle)$ is $U|\psi\rangle|0^{W_0}\rangle$, already of norm one. As A is unitary, its orthogonal component has norm zero. Thus A itself is a clean implementation. In this case, the identity

$$\left(\sum_P |\alpha_P| \right)^2 = \sum_P |\alpha_P|^2 + \sum_{P \neq Q} |\alpha_P| |\alpha_Q|$$

also shows that $M = 1$; the target is a scalar of modulus one times a single Pauli word.

For $\lambda > 1$, the amplification count is $K = \lceil \pi\lambda/4 \rceil = O(\lambda)$. Combining (18) with Lemma 4.1, and using $n, M \geq 1$, gives $s(\mathcal{C}) = O(nM\lambda)$. The same bound holds when $\lambda = 1$. There are $M + 3$ block-encoding workspace qubits and at most two more for outer amplification. Equation (5) and $M \leq 4^n$ now give every bound in (1). All equality gates used here occur in separate layers. Their threshold expansions retain these symmetries and add no workspace, so the bounds apply to the stated elementary gate set.

Finally, if $X = \emptyset$, then $\mathcal{H}_X \cong \mathbb{C}$ and $U = e^{i\varphi}$ for some φ . The allowed one-qubit gate $e^{i\varphi} I_2$, applied to one fixed workspace wire in $|0\rangle$, implements this scalar exactly and leaves that wire clean. $\square$

Corollary 4.2. *Let a finite group G act on X through a homomorphism $\rho : G \rightarrow S_X$, not necessarily injectively. Every unitary commuting with $R_X(\rho(G))$ has an exact clean circuit symmetric under this action, with the bounds of Theorem 1.1.*

Proof. Apply the theorem to $\Gamma = \text{im } \rho$, and compose its explicit action on the circuit wires with ρ . The kernel of ρ acts trivially on every wire. $\square$

The selector cost is controlled by the number of Pauli words, not the number of their orbits. Replacing an orbit by a single selector label would no longer specify which of its individual words should be applied; the corresponding orbit sum need not be unitary. Thus orbit-constant coefficients explain symmetry of the construction but do not, by themselves, compress it. Appendix A records a different route through symmetric Choi-state preparation, with that preparation made an explicit hypothesis.

A A conditional construction from a Choi state

An exact symmetric preparation of the Choi state of U would also suffice for exact symmetric synthesis of U . The reduction below records the required workspace action and normalization. It is conditional on that state-preparation circuit and independent of the Pauli construction.

Let $n \geq 1$, let $d = 2^n$, and let D, A, B be three copies of the wire set X . Define

$$|\Phi_d\rangle_{AB} = d^{-1/2} \sum_{b \in \{0,1\}^X} |b\rangle_A |b\rangle_B, \quad |J_U\rangle_{AB} = (I_A \otimes U_B) |\Phi_d\rangle_{AB}.$$

The diagonal action of Γ on $A \sqcup B$ fixes $|\Phi_d\rangle$: it permutes the summation labels in both factors in the same way. Since U commutes with this action on B , the diagonal action also fixes $|J_U\rangle$.

Suppose a circuit C_J , symmetric under that diagonal action, prepares $|J_U\rangle_{AB}$ from zero using additional workspace T_J , returned to zero. Each of its chosen lifts is assumed to restrict to the prescribed action on both A and B . Let F_{DA} be Bell preparation on corresponding pairs: a uniform Hadamard layer on D , followed by the layer of CNOTs from D to A . Then $F_{DA}|0^D, 0^A\rangle = |\Phi_d\rangle_{DA}$. Define

$$A_J = \text{SWAP}_{DB} F_{DA}^\dagger C_J,$$

with every factor extended by the identity on unused wires. The register swap is three layers of CNOTs on corresponding pairs. Extending each lift of C_J by the same action on D preserves every Bell-unpreparation and swap layer. Thus A_J is symmetric with active register D and workspace $A \sqcup B \sqcup T_J$.

For $|\psi\rangle = \sum_b \psi_b |b\rangle$, direct contraction gives

$$(\langle \Phi_d |_{DA} \otimes I_B)(|\psi\rangle_D \otimes |J_U\rangle_{AB}) = \frac{1}{d} \sum_b \psi_b U |b\rangle_B = \frac{1}{d} U |\psi\rangle_B. \quad (23)$$

Before the swap, the all-zero component on $D \sqcup A$ is therefore $|0^D, 0^A\rangle U |\psi\rangle_B / d$; T_J is already zero. The swap moves the output to D , so

$$(I_D \otimes \langle 0^{A \sqcup B \sqcup T_J} |) A_J (I_D \otimes |0^{A \sqcup B \sqcup T_J}\rangle) = \frac{U}{d}.$$

Lemma 4.1 applies with $p = d^{-2}$. The resulting resource estimates depend on $s(C_J)$ and $a(C_J)$. An independently efficient symmetric Choi-state preparation would therefore give an alternative synthesis route; the bounds in (1) follow from the explicit Pauli selector instead.

References

- [1] D. Castro-Silva, T. Gur, and S. Strelchuk. Symmetric quantum computation. In *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, volume 362 of *LIPIcs*, article 35, 2026. doi:10.4230/LIPIcs.ITCS.2026.35. Full version: arXiv:2501.01214v2, 2025.
- [2] G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002. doi:10.1090/conm/305/05215.
- [3] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma. Simulating Hamiltonian dynamics with a truncated Taylor series. *Physical Review Letters*, 114:090502, 2015. doi:10.1103/PhysRevLett.114.090502.
- [4] Q. T. Nguyen, L. Schatzki, P. Braccia, M. Ragone, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo. Theory for equivariant quantum neural networks. *PRX Quantum*, 5:020328, 2024. doi:10.1103/PRXQuantum.5.020328.
- [5] L. Schatzki, M. Larocca, Q. T. Nguyen, F. Sauvage, and M. Cerezo. Theoretical guarantees for permutation-equivariant quantum neural networks. *npj Quantum Information*, 10:12, 2024. doi:10.1038/s41534-024-00804-1.
- [6] V. M. Bastidas, N. Fitzpatrick, K. J. Joven, Z. M. Rossi, S. Islam, T. Van Voorhis, I. L. Chuang, and Y. Liu. Unification of finite symmetries in the simulation of many-body systems on quantum computers. *Physical Review A*, 111:052433, 2025. doi:10.1103/PhysRevA.111.052433.
- [7] S. Kazi, M. Larocca, and M. Cerezo. On the universality of S_n -equivariant k -body gates. *New Journal of Physics*, 26:053030, 2024. doi:10.1088/1367-2630/ad4819.
- [8] P. Deliyannis and I. Marvian. Permutation-invariant N -body gates via the Tavis–Cummings interaction. arXiv:2506.03453, 2025; version 3, 2026.
- [9] I. Marvian. Restrictions on realizable unitary operations imposed by symmetry and locality. *Nature Physics*, 18:283–289, 2022. doi:10.1038/s41567-021-01464-0.