A JOURNAL BY SISSA MEDIALAB

JQuant

SUBMITTED: 21st May 2026
REVISED: 1st September 2026
ACCEPTED: 30th September 2026
PUBLISHED: 29th October 2026

Universality of quantum gates in particle and symmetry constrained subspaces

Andreas Stergiou [a,1] **and Nicolas P.D. Sawaya** [b]

[a] *Department of Mathematics, King's College London, Strand, London WC2R 2LS, United Kingdom*

[b] *Azulene Labs, Berkeley, CA 94720, USA*

CONTACTS: andreas.stergiou@kcl.ac.uk, nicolas@azulenelabs.com

ABSTRACT

Simulating physical systems on near-term quantum computers often requires preparing states within constrained subspaces, like those with fixed particle number or spin. We use Lie algebraic techniques to prove that hardware-efficient gates are universal for state preparation in these subspaces. The key mechanism is Pauli Z dressing: commutators of overlapping gates produce Pauli Z operators on shared qubits, acting as spectator projectors that decompose multi-plane rotations into single-plane generators spanning the full $\mathfrak{so}(w)$ algebra, where w is the dimension of the constrained subspace, thereby guaranteeing universality for real state preparation. Adding independent complex phases extends this to $\mathfrak{su}(w)$, enabling arbitrary complex state preparation. We provide a computationally efficient Jacobian criterion for verifying that a circuit can explore any direction on the target manifold from almost any parameter configuration. Our findings are applicable to many problem areas, including Fermi-Hubbard models, Bose-Hubbard models, and molecular electronic structure. We apply our framework to two physical settings: we prove the completeness of the binary encoded multi-level particles ansatz on the conserved-particle-number subspace, and we construct symmetry-preserving circuits for the fuzzy sphere regularisation of the 3D Ising conformal field theory (CFT). For the latter, we variationally prepare the ground and excited states and extract some CFT data.

KEYWORDS: Hamiltonian simulation; Quantum circuits; Quantum criticality / CFT methods

ARXIV: 2605.00979

[1] Contact author.

© The Author(s).
This article is licensed under the terms of the Creative Commons Attribution 4.0 license. All rights for Text and Data Mining, AI training, and similar technologies for commercial purposes, are reserved.

CONTENTS

1 Introduction **1**

2 Lie algebraic framework for constrained state preparation **3**
2.1 Completeness of hardware-efficient ansatze 3
2.2 Conditions for spanning circuits 7

3 Application to bosonic simulation with binary encoding **9**

4 Application to the fuzzy sphere regularisation of the 3D Ising model **10**
4.1 Lowest Landau level physics and fuzzy sphere 10
4.2 State-operator correspondence 12
4.3 Classical extraction via exact diagonalisation 13
4.4 Gates for preserving spin symmetries 14
4.5 Variational quantum simulation 15
4.6 Extracting an OPE coefficient 16
4.7 Scaling with system size 18
4.8 Simulation under a noise model 19

5 Extension to $SU(w)$ with independent complex phases **20**

6 Conclusion **22**

A Commutator identities **23**

B Elementary gate decompositions **25**

References **28**

1 Introduction

Quantum computers promise to provide a pathway to explore previously inaccessible regimes in physics, including strongly-coupled systems and complex many-body dynamics. Due to the limited coherence times and gate fidelities of near-term noisy devices, algorithms typically rely on shallow, parametrised quantum circuits. A central challenge in these variational quantum algorithms (VQAs), such as the variational quantum eigensolver (VQE) [1, 2], is the design of the state preparation ansatz. The ansatz must be expressive enough to reach the desired physical states while remaining hardware-efficient.

In many physical contexts, the relevant Hilbert space is restricted by conservation laws, such as particle number conservation or total spin. Preparing an arbitrary state within such a constrained subspace usually involves complicated unitaries. For example, controlled gates or operations involving long non-local Pauli mappings, such as Jordan-Wigner strings, can prepare exact superpositions. However, their physical implementation demands a prohibitive number of native device gates, drastically increasing circuit depth. Instead, simpler hardware-efficient gates that omit these non-local strings are often employed, relying entirely on adjacent or nearby interactions.

In this work, we use Lie algebraic techniques to establish a mathematical foundation for universal state preparation with hardware-efficient gates within constrained subspaces. A constrained subspace of dimension w, spanned by w computational basis states regarded as coordinate axes of $\mathbb{R}^w$, carries a natural geometric structure: unit-normalised real superpositions live on the sphere S^{w-1}, acted on by

the rotation group $SO(w)$ with Lie algebra $\mathfrak{so}(w)$. We show in Proposition 1 that hardware-efficient 2-qubit "multi-plane" rotation gates, which simultaneously rotate in every plane connecting basis states differing only on the two active qubits, generate the full $\mathfrak{so}(w)$ algebra. The key mechanism, which we term *Pauli Z dressing*, is that commutators of overlapping gates produce Pauli Z operators on shared qubits, which act as spectator projectors and allow the extraction of "single-plane" generators rotating between individual pairs of basis states. Connectivity of the graph of constant-particle-number basis states under pairwise swaps then ensures that these generators span the complete algebra, so circuits composed of such gates can universally prepare any real state within the constrained subspace. When each gate is equipped with an independent complex phase, we show in section 5 that the same spectator mechanism extends to the imaginary hopping generators, promoting the algebra from $\mathfrak{so}(w)$ to the full $\mathfrak{su}(w)$ and enabling universal complex state preparation.

Unlike standard universality proofs for discrete gate sets, which establish dense subgroups via the Solovay-Kitaev theorem [3, 4], our approach exploits the continuous parametrisation of variational gates. Drawing on the Lie algebra rank condition from quantum control theory [5], we show that spanning the Lie algebra guarantees that finite-depth circuits can reach general elements of the relevant connected Lie group exactly. Theorem 1 below provides the essential mathematical foundation for this. Although this continuous approach has been employed previously [6–9], our focus here is on spanning specific constrained subspaces of a Hilbert space.

While Theorem 1 establishes universal reachability in principle, verifying that a specific, finite-depth ansatz spans the target space in practice is a separate challenge. To partly address this, we provide a computationally efficient local criterion, Lemma 1, showing that a minimal circuit of $w-1$ generators has a Jacobian of rank $w-1$ at almost every parameter value — meaning the circuit can explore any direction on S^{w-1} from almost any parameter configuration — provided the Jacobian achieves this rank at a single reference point. Although surjective parametrisations of the target sphere exist, critical points of any circuit map are topologically unavoidable, and for hardware-efficient circuits these can partition the reachable region into disconnected patches. Crossing such boundaries generally requires over-parametrising the circuit.

Our framework applies immediately to any fermionic system with conserved particle number, such as the Fermi-Hubbard model, one of the most widely studied targets for variational quantum simulation [10–12]. We present detailed applications to two settings that involve additional algebraic structure beyond the basic constant-Hamming-weight case. In section 3, we treat the binary encoded multi-level particles ansatz (BEMPA) [13] for bosonic simulation. The BEMPA gate set involves two families of gates ($\hat{A}$ and $\hat{B}$) acting at different binary significance levels; we prove that their commutators activate the same projection mechanism, establishing completeness on the conserved-particle-number subspace (Proposition 2). In section 4, we study the fuzzy sphere regularisation of conformal field theories (CFTs) [14], specifically the 3D Ising model at criticality. Here the physically relevant subspace carries an additional $S_z = 0$ constraint on the total azimuthal angular momentum. Two-qubit gates alone cannot change orbital occupancies while preserving this symmetry, so we consider 4-qubit gates that simultaneously transfer two excitations between orbitals. We then prove (Proposition 3) that the combined 2- and 4-qubit gate set generates the full orthogonal group on the symmetry-constrained subspace, enabling universal state preparation for this model.

The fuzzy sphere Hamiltonian we examine has a finite-dimensional spectrum whose low-lying eigenvalues map, via the state-operator correspondence, to the scaling dimensions of the 3d Ising CFT. Finding the ground and excited states of this Hamiltonian is closely analogous to quantum chemistry problems, and already developed algorithms like VQE and the variational quantum deflation (VQD) algorithm for excited states [15] can be readily applied, as can quantum subspace methods [16]. Because exactly diagonalising this Hamiltonian classically becomes unmanageable for large numbers

of modes, quantum simulation holds immense value. We demonstrate the practical construction of minimal state-preparation circuits tailored to the model's constraints and extract the ground state and two low-lying excited states.

The 3D Ising model at criticality is a particularly attractive target because high-precision data exist through the numerical conformal bootstrap [17]. This provides a rare setting where both strong coupling and reliable independent results coexist, making it an ideal benchmark for quantum simulation: agreement with bootstrap data validates the quantum algorithms, while discrepancies can reveal shortcomings of current devices and guide future improvements.

The remainder of this paper is organised as follows. In section 2 we develop the mathematical framework: we prove the completeness of hardware-efficient gates on constant-Hamming-weight subspaces (Proposition 1) and provide a computationally efficient Jacobian criterion for verifying that a circuit can explore any direction in the target manifold from almost any point (Lemma 1). Section 3 applies this framework to bosonic simulation with binary encoding, establishing the completeness of the BEMPA ansatz (Proposition 2). Section 4 treats the fuzzy sphere regularisation of the 3D Ising CFT: we review the physical setup, construct symmetry-preserving circuits that generate the full orthogonal group on the $S_z = 0$ subspace (Proposition 3), and present variational quantum simulation results for the ground and excited states. We then extract an OPE coefficient from the variationally prepared states, examine how the required circuit size grows with the number of electrons, and determine the gate error rate at which noise begins to dominate the finite-size error. In section 5 we show how independent complex phases extend all preceding results from $\mathfrak{so}(w)$ to $\mathfrak{su}(w)$. We conclude in section 6. Appendix A verifies explicitly the commutator identities on which the three propositions rest, and appendix B collects the elementary gate decompositions used in this work.

2 Lie algebraic framework for constrained state preparation

Efficient state preparation is a critical subroutine in variational quantum algorithms. In many physical systems of interest, such as electronic structure and the fuzzy sphere regularisation discussed later, the target states are constrained by symmetries like particle number conservation or total spin. While controlled operations could be used to prepare generic superpositions within these constrained subspaces [18], they often lead to prohibitively deep circuits on near-term hardware. Instead, hardware-efficient ansatze employing nearest-neighbour or sparsely connected 2-, 3- or 4-qubit gates are preferred.

Here, we provide mathematical results demonstrating how these simpler hardware-efficient gates can achieve universal state preparation within constrained subspaces, drawing on the Lie algebraic properties of the gates. We focus on the subspace of states of constant Hamming weight in this section, and we provide extensions to further cases in the application sections below. We formulate our essential theoretical results as propositions to guide the construction of suitable quantum circuits.

2.1 Completeness of hardware-efficient ansatze

Suppose we want to generate an arbitrary superposition in a subspace of constant Hamming weight n_q using N_q qubits. The dimension of this subspace is

$$w = \binom{N_\mathrm{q}}{n_\mathrm{q}} . \tag{2.1}$$

Let us think of these w basis states as the axes $\mathbf{e}_0, \mathbf{e}_1, \ldots, \mathbf{e}_{w-1}$ of $\mathbb{R}^w$. A unit-normalised, time-reversal symmetric state can be represented by a vector $\mathbf{v}$ in $\mathbb{R}^w$ that can be written as a linear combination of these w states with real coefficients whose squares add up to unity:

$$\mathbf{v} = v_0\mathbf{e}_0 + v_1\mathbf{e}_1 + \cdots + v_{w-1}\mathbf{e}_{w-1} , \qquad v_0^2 + v_1^2 + \cdots + v_{w-1}^2 = 1 . \tag{2.2}$$

In other words, the states we are interested in live on the unit sphere S^{w-1} in $\mathbb{R}^w$. To navigate this subspace, we must apply operations from the special orthogonal group $SO(w)$, which has dimension $w(w-1)/2$.

In this geometric language, the question we're interested in is how to get to a generic location on S^{w-1} starting from a specified location, e.g. $\mathbf{e}_0$. Of course this is possible with the use of $w \times w$ orthogonal matrices with unit determinant (elements of $SO(w)$). However, we would like to achieve this without using gates that grow with w.

A common approach is the use of Givens rotations,

$$G(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{2.3}$$

which are understood to be acting on two qubits only, and we use the ordering $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ for the corresponding column of basis states on which the matrix (2.3) acts. The two qubits on which the Givens rotation is acting need to be specified, for example,

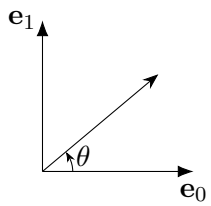
$$G_{01}(\theta)\,|010101\rangle = \cos\theta\,|010101\rangle - \sin\theta\,|100101\rangle\,, \tag{2.4a}$$

$$G_{02}(\theta)\,|010101\rangle = |010101\rangle\,, \tag{2.4b}$$

$$G_{12}(\theta)\,|010101\rangle = \cos\theta\,|010101\rangle + \sin\theta\,|001101\rangle\,, \tag{2.4c}$$

where the first qubit is indexed by 0, the second by 1, and so on.

From a geometric perspective, a Givens rotation can be thought of as a rotation in the plane defined by two axes. For example, if $\mathbf{e}_0$ is the axis representing $|010101\rangle$ and $\mathbf{e}_1$ the axis representing $|001101\rangle$, then (2.4c) represents a rotation by angle θ in the plane defined by $\mathbf{e}_0, \mathbf{e}_1$:



The essential complication is that a Givens rotation acts simultaneously across all possible configurations of the spectator qubits. At some point, this will result in a rotation in more than one plane *simultaneously*, for example if G_{45} is applied on the last of (2.4). One can avoid this with the use of controlled gates, but the increased circuit depth renders them undesirable. As we will see, efficient quantum circuits can be constructed leveraging the algebraic structure of the generators of multi-plane rotations: their commutators produce intermediate Z operators that act as projectors on the spectator qubits, allowing one to isolate single-plane rotations.

Let us start with a fundamental theorem.[1]

Theorem 1. *Let G be a connected Lie group with Lie algebra $\mathfrak{g}$, and let $S \subseteq \mathfrak{g}$ be a subset that generates $\mathfrak{g}$ as a Lie algebra. Furthermore, let $H \leqslant G$ be the subgroup generated by the set $\{\exp(tX) \colon t \in \mathbb{R}, X \in S\}$. Then, $H = G$.*

Proof. We first observe that H is an arcwise connected subgroup of G. By definition, H is generated by elements of the form $\exp(tX)$ for $t \in \mathbb{R}$ and $X \in S$. For any such generator, the map $\gamma : \mathbb{R} \to G$ given

[1] We thank Petr Kravchuk for suggesting the connection to Yamabe's theorem presented in the proof.

by $\gamma(t) = \exp(tX)$ defines a continuous path lying entirely within H that connects the identity element $e = \gamma(0)$ to $\exp(X) = \gamma(1)$. Since arbitrary elements of H are finite products of these generators and group multiplication is continuous, any element $h \in H$ can be connected to the identity e by a continuous path in H. Thus, H is arcwise connected.

By Yamabe's theorem [19], any arcwise connected subgroup of a Lie group is an analytic subgroup. Therefore, H is a Lie subgroup of G in its own right, possessing a well-defined Lie algebra $\mathfrak{h} \subseteq \mathfrak{g}$.

Next, we determine the Lie algebra $\mathfrak{h}$. For every $X \in S$, the one-parameter subgroup $t \mapsto \exp(tX)$ lies entirely within H. The tangent vector to this curve at the identity must belong to the Lie algebra of H:

$$\left.\frac{d}{dt}\right|_{t=0} \exp(tX) = X \in \mathfrak{h}\,. \tag{2.5}$$

This implies that $S \subseteq \mathfrak{h}$. By hypothesis, S generates $\mathfrak{g}$ as a Lie algebra. Since $\mathfrak{h}$ is a Lie subalgebra of $\mathfrak{g}$ containing S, it must contain the entire Lie algebra generated by S. Consequently, we obtain $\mathfrak{h} = \mathfrak{g}$.

Finally, since H is a connected Lie subgroup of G and their Lie algebras coincide, H must be the connected component of the identity of G. Because G is assumed to be connected, its identity component is the entire group. Therefore, $H = G$. □

This theorem reduces our problem to showing that a given set of generators is sufficient to generate the desired Lie algebra. For those less familiar with formal group theory, we note that Lie algebra is analogous to quantum operator algebra, and elements of the corresponding Lie group are analogous to exponentials and products of exponentials of the quantum operators.

Definition 1 (Multi-plane generator)**.** *An individual proper rotation gate can be represented as a continuous Lie group element,*

$$G_{ij}(\theta) = \exp\left[i\frac{\theta}{2}(X_iY_j - Y_iX_j)\right]. \tag{2.6}$$

The associated generator is

$$L_{ij} = X_iY_j - Y_iX_j\,. \tag{2.7}$$

The omission of Z-strings means the interacting G-gate applies identical pairwise-swap transitions across all possible configurations of spectator qubits. We call L_{ij} a multi-plane generator because it simultaneously rotates in every 2-dimensional plane connecting pairs of basis states that differ only in the occupations of qubits i and j.

The A-gate [20, 21],

$$A(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & \sin\theta & 0 \\ 0 & \sin\theta & -\cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{2.8}$$

applies the G-gate proper rotation combined with a local reflection matrix that selectively applies a -1 phase to the $|10\rangle$ state,

$$A_{ij}(\theta) = A_{ij}(0)G_{ij}(-\theta)\,, \tag{2.9}$$

where

$$A_{ij}(0) = I - 2\left(\frac{I - Z_i}{2}\right)\left(\frac{I + Z_j}{2}\right). \tag{2.10}$$

Definition 2 (Single-plane generator)**.** *Given two basis states $|x\rangle$ and $|y\rangle$, the single-plane generator is*

$$E_{xy} = |x\rangle\langle y| - |y\rangle\langle x|\,. \tag{2.11}$$

Unlike L_{ij}, the generator E_{xy} rotates in exactly one plane: the single plane spanned by $|x\rangle$ and $|y\rangle$. When $|x\rangle$ and $|y\rangle$ differ on exactly two qubits, i and j, the elementary generator can be expressed explicitly as

$$E_{xy} = \frac{i}{2} L_{ij} \prod_{k \neq i,j} \frac{I + (-1)^{x_k} Z_k}{2}\,, \tag{2.12}$$

where $x_i = 1, x_j = 0$, and x_k is the value of qubit k in $|x\rangle$. The product of projectors pins every spectator qubit to its value in $|x\rangle$ (and $|y\rangle$, since they agree on spectators).

Proposition 1. *Hardware-efficient 2-qubit rotation gates that locally exchange a single excitation (such as the G-gate or the corresponding discrete reflection A-gate), without their non-local Jordan-Wigner strings, generate the full $\mathfrak{so}(w)$ Lie algebra. Consequently, circuits composed of these gates can universally prepare any real state within the constant particle number subspace.*

Proof. To prove that we can generate any transition, we must show that we can isolate a single generator E_{xy} from the multi-plane L_{ij}. This is achieved through the commutators of the proper rotation generators. Throughout, operators such as L_{ij} and Z_k are understood to be acting on the full multi-qubit Hilbert space, with trivial (identity) action on all unlabelled qubits. Commuting two overlapping generators yields

$$[L_{ik}, L_{kj}] = -2i L_{ij} Z_k\,, \tag{2.13}$$

by virtue of $[X, Y] = 2iZ$. The crucial feature here is the emergence of the Pauli Z operator on the shared qubit. Since L_{ij} is a native generator and $L_{ij}Z_k$ arises from the commutator, their linear combinations $L_{ij}(I \pm Z_k)/2$ project spectator qubit k onto a definite value.

Let S be a set of spectators disjoint from $\{i, j\}$. To project on all spectator qubits simultaneously and thereby isolate E_{xy}, we show by induction on $|S|$ that $L_{ij} \prod_{k \in S} Z_k$ lies in the generated algebra for every pair (i, j) and every S. The base cases $|S| = 0$ (native generators) and $|S| = 1$ (equation (2.13)) are established above. For the inductive step, suppose the claim holds for all sets of size less than or equal to s and let $|S| = s + 1$. Pick any $k' \in S$ and set $S' = S \setminus \{k'\}$, so $|S'| = s$. The inductive hypothesis applied to the pair (k', j) yields $L_{k'j} \prod_{l \in S'} Z_l$. Since $S' \cap \{i, k'\} = \emptyset$, every Z_l in the product commutes with the native generator $L_{ik'}$, giving

$$\left[L_{ik'},\ L_{k'j} \prod_{l \in S'} Z_l \right] = [L_{ik'},\ L_{k'j}] \prod_{l \in S'} Z_l = -2i L_{ij} Z_{k'} \prod_{l \in S'} Z_l = -2i L_{ij} \prod_{l \in S} Z_l\,. \tag{2.14}$$

Thus, L_{ij} can be dressed with an arbitrary product of spectator Z operators. Linear combinations of differently dressed generators then yield the full projector product $\prod_{k \neq i,j}(I + (-1)^{x_k} Z_k)/2$ needed to isolate the single-plane generator E_{xy}.

The dressing procedure above directly isolates E_{xy} whenever $|x\rangle$ and $|y\rangle$ differ on exactly two qubits, since such a pair lies in the image of a single multi-plane generator L_{ij}. To obtain E_{xy} for pairs differing on more than two qubits, consider the graph whose vertices are the weight-n_q bitstrings, with edges connecting pairs that differ in exactly two positions. This graph is connected: given any two same-weight bitstrings, one can repeatedly swap a 1 in a position where the strings disagree with a 0 in another disagreeing position, strictly reducing their Hamming distance at each step. Given any path $|x\rangle, |y\rangle, \ldots, |z\rangle$ in this graph, iterated application of the commutator identity

$$[E_{xy}, E_{yz}] = E_{xz} \tag{2.15}$$

constructs the generator for the composite transition. This produces all $w(w-1)/2$ generators of $\mathfrak{so}(w)$.

For the A-gate, the continuous part of the decomposition $A_{ij}(\theta) = A_{ij}(0)\, G_{ij}(-\theta)$ shares the same generator L_{ij}, so the Lie algebra generated by A-gate circuits is identical to that of the corresponding G-gate circuits. □

A related theorem regarding the universality of Hamming-weight preserving gates has recently appeared in [22]. There, the authors analytically evaluate the Lie algebra generated by continuous one-parameter families of quantum gates and prove that for purely continuous hardware-efficient operators to span the full special unitary Lie algebra $\mathfrak{su}(w)$, their generator must inherently include two-body interaction diagonal terms that break free-fermion integrability. Our formalisation in Proposition 1 is fully complementary to their result: while pure G-gates (which represent pure single-particle hopping) are mathematically insufficient to approximate arbitrary complex unitary matrices, as accurately dictated by the conditions established in [22], we establish that they precisely span the special orthogonal algebra $\mathfrak{so}(w)$. This orthogonal universality guarantees that such minimal hardware-efficient circuits can access any arbitrary *real* quantum state within the symmetry-constrained subspace, which constitutes the essential requirement for the variational preparation of ground states of real Hamiltonians.

2.2 Conditions for spanning circuits

For practical scaling on near-term devices, it is desirable to have a robust algebraic technique for confirming if a specialised ansatz spans the required manifold with the minimal number of parameters. The most efficient preparation circuit mandates exactly $w-1$ independent parameters to traverse the manifold. An ansatz that fails to span the target manifold may be unable to reach the true ground state even in principle, so circuit completeness is not merely a convenience but a fundamental computational necessity. Approximating the ground state energy of a generic local Hamiltonian is QMA-complete [23], a complexity class believed to be strictly harder than BQP.

Lemma 1 (Local Spanning Criterion)**.** *Let $|\psi_0\rangle$ be a fixed initial state in the Hamming-weight-n_q subspace of dimension w, and let $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})\,|\psi_0\rangle$ be a parametrised circuit with $p \geqslant w-1$ generators H_k. If the $w \times p$ Jacobian matrix with entries $J_{xk}(\boldsymbol{\theta}) = \partial_{\theta_k} \langle x|\psi(\boldsymbol{\theta})\rangle$, where $\{|x\rangle\}$ are the w computational basis states of the Hamming-weight-n_q subspace, has rank $w-1$ at some reference point $\boldsymbol{\theta}_\mathrm{ref}$, then J has rank $w-1$ at almost every $\boldsymbol{\theta} \in T^p$, where T^p is the p-dimensional torus.*

Proof. Since the circuit is a composition of exponentials of Hermitian operators, the map $\boldsymbol{\theta} \mapsto |\psi(\boldsymbol{\theta})\rangle$ is real-analytic, and hence so is each entry of the $w \times p$ Jacobian matrix $J(\boldsymbol{\theta})$. Differentiating $\langle\psi|\psi\rangle = 1$ shows that the columns of J lie in the tangent space of S^{w-1}, so $\mathrm{rank}(J) \leqslant w-1$. The Jacobian achieves rank $w-1$ if and only if at least one $(w-1)\times(w-1)$ minor of J is non-zero. Each such minor is a real-analytic function of $\boldsymbol{\theta}$. A real-analytic function that is non-zero at a single point can vanish only on a set of Lebesgue measure zero, so a rank-$(w-1)$ Jacobian at $\boldsymbol{\theta}_\mathrm{ref}$ implies rank $w-1$ at almost every $\boldsymbol{\theta} \in T^p$.

When each generator H_k appears at most once, the natural choice is $\boldsymbol{\theta}_\mathrm{ref} = \mathbf{0}$, where all intermediate unitaries collapse to the identity and the entries of J reduce to $J_{xk}(\mathbf{0}) = -i\,\langle x|H_k|\psi_0\rangle$. When a generator appears at multiple depths, evaluation at $\boldsymbol{\theta} = \mathbf{0}$ collapses all instances to the same tangent vector and is inconclusive; one should instead evaluate at a generic $\boldsymbol{\theta}_\mathrm{ref} \neq \mathbf{0}$, at which the non-commuting intermediate unitaries rotate each instance to a distinct tangent vector, and the same real-analytic argument applies. □

The criterion of Lemma 1 provides a powerful, computationally inexpensive classical method to test gate configurations. In the minimal case $p = w-1$, a full-rank Jacobian at a reference point

certifies that the image of the circuit map

$$\begin{aligned} U\colon T^{w-1} &\to S^{w-1} \\ \boldsymbol{\theta} &\mapsto U(\boldsymbol{\theta})|\psi_0\rangle \end{aligned} \tag{2.16}$$

explores a full $(w-1)$-dimensional patch of the target manifold at generic parameter values. This local condition can be used to show, for example, that a circuit composed exclusively of continuous proper G-gates restricted only to adjacent qubits cannot generate the intermediate Z operators necessary to span the subspace of constant Hamming weight. Thus, to achieve universality across the full subspace simply with proper parametrised G-gates, the circuit must include gates acting selectively on non-adjacent qubits.

Corollary 1. *For $2 \leqslant n_\mathrm{q} \leqslant N_\mathrm{q} - 2$, no circuit composed exclusively of continuous proper G-gates on adjacent qubits has a Jacobian of rank $w-1$ for the constant-Hamming-weight subspace of dimension $w = \binom{N_\mathrm{q}}{n_\mathrm{q}}$.*

Proof. Because they act only on neighbouring modes, the adjacent generators $L_{i,i+1} = X_iY_{i+1} - Y_iX_{i+1}$ in the Jordan-Wigner representation are completely free of intermediate Z-strings and thus correspond precisely to quadratic fermionic hopping operators. Because the commutator of any two number-conserving quadratic operators yields another number-conserving quadratic operator, the Lie algebra generated by adjacent-only gates is contained in the Lie algebra of number-conserving free-fermion transformations.

By Theorem 1, the achievable group of unitaries is therefore restricted to the connected Lie subgroup generated by this free-fermion algebra, which is the group of continuous single-particle rotations $SO(N_\mathrm{q})$. A computational basis state of constant Hamming weight represents a single Slater determinant, which corresponds to choosing exactly n_q occupied modes out of the total N_q available single-particle modes. Evolution under the free-fermion group $SO(N_\mathrm{q})$ acts merely as a basis transformation on these underlying modes. Consequently, the n_q-dimensional subspace of occupied states is simply rotated as a rigid block within the overall N_q-dimensional single-particle space. By definition, the geometric space of all such n_q-dimensional subspaces in $\mathbb{R}^{N_\mathrm{q}}$ is the real Grassmannian manifold $\mathrm{Gr}_\mathbb{R}(n_\mathrm{q}, N_\mathrm{q})$. Consequently, the orbit of any computational basis state — and thus the total reachable set of states under adjacent G-gate circuits — is restricted to this Grassmannian, whose manifold dimension is exactly $n_\mathrm{q}(N_\mathrm{q} - n_\mathrm{q})$. Because the generated tangent space is everywhere bounded by this manifold dimension, it cannot span the full tangent space of the unit sphere, provided $n_\mathrm{q}(N_\mathrm{q} - n_\mathrm{q}) < w - 1$.

It remains to verify this inequality for $2 \leqslant n_\mathrm{q} \leqslant N_\mathrm{q} - 2$. By AM-GM, $n_\mathrm{q}(N_\mathrm{q} - n_\mathrm{q}) \leqslant \frac{1}{4}N_\mathrm{q}^2$. The binomial coefficients obey $\binom{N_\mathrm{q}}{n_\mathrm{q}} \geqslant \binom{N_\mathrm{q}}{2} = \frac{1}{2}N_\mathrm{q}(N_\mathrm{q}-1)$ for all $2 \leqslant n_\mathrm{q} \leqslant N_\mathrm{q} - 2$. The bound $\frac{1}{4}N_\mathrm{q}^2 < \frac{1}{2}N_\mathrm{q}(N_\mathrm{q}-1) - 1$ holds for all $N_\mathrm{q} \geqslant 4$ (and the condition $2 \leqslant n_\mathrm{q} \leqslant N_\mathrm{q} - 2$ forces $N_\mathrm{q} \geqslant 4$). Therefore, $n_\mathrm{q}(N_\mathrm{q} - n_\mathrm{q}) \leqslant \frac{1}{4}N_\mathrm{q}^2 < w - 1$. □

The dimensional gap arises because generic states in S^{w-1} are not Slater determinants (e.g., the equal superposition $\frac{1}{\sqrt{2}}(|1100\rangle + |0011\rangle)$ lies strictly outside $\mathrm{Gr}_\mathbb{R}(2,4)$). A direct Pauli computation provides a complementary structural perspective on the restriction of Corollary 1. The commutator of two adjacent generators produces a non-adjacent generator dressed by a Z-string on the shared qubit. Every iterated commutator of adjacent generators has the form $L_{ij}Z_{i+1}\cdots Z_{j-1}$, and since commutators of two such operators again have this form, they close into a Lie subalgebra; this is precisely the free-fermion algebra of Corollary 1. This subalgebra cannot produce the non-contiguous dressed operators $L_{ij}Z_k$ (with $k \notin \{i{+}1, \ldots, j{-}1\}$) required by Proposition 1 to span the full $\mathfrak{so}(w)$.

It is worth noting that the specific A-gate circuits proposed by Gard et al. [21], which use repeating blocks of strictly adjacent A-gates, appear to bypass this limitation but only if the input state is

of the Néel type (i.e. $|0101\ldots\rangle$). Discrete reflections $A_{ij}(0)$ in adjacent A-gates, when applied to the alternating occupation pattern of the Néel state, appear to break the free-fermion integrability that would otherwise constrain adjacent-only continuous rotations. While adjacent A-gates are not generically universal for arbitrary state-to-state mappings across the full algebra, they might be sufficient to construct complete state-preparation ansatze when anchored to the correct starting point, as suggested in [21]. A rigorous proof of this statement, however, remains elusive.

While Lemma 1 ensures that a minimal circuit can explore a local neighbourhood of states, it does not guarantee that every possible state in the subspace can be reached. Any smooth mapping from the circuit's periodic parameters to the target sphere must encounter "critical points" where the circuit loses its ability to move in certain directions. This is due to a fundamental topological difference: the parameter space (a torus) has "holes" or cycles that the sphere lacks, making it impossible for any smooth map between them to have full-rank Jacobian everywhere. In practice, hardware-efficient circuits are even more restricted because each gate acts on several pairs of states simultaneously. This structure often partitions the reachable states into separate "islands" or patches, meaning that a single minimal circuit may be "trapped" within one region and unable to reach states beyond certain topological boundaries.

For variational algorithms such as VQE, which target low-lying excitations near a known reference state, local coverage of a single patch is typically sufficient. Achieving strict global reachability of arbitrary superpositions, however, generally requires over-parametrising the circuit, for instance by repeating the minimal gate sequence, to supply the extra degrees of freedom needed to cross these boundaries. The construction of [21] appears to achieve a surjective map from T^w to S^{w-1} precisely by introducing one additional parameter beyond the minimal count.

3 Application to bosonic simulation with binary encoding

The mathematical framework developed in section 2 for dynamically generating the full orthogonal group on a constrained subspace has a variety of extensions. One of them is the simulation of bosonic systems with a conserved particle number. Several considerations related to the quantum simulation of bosonic and phononic problems have been previously studied [24–30].

Recently, the binary encoded multi-level particles ansatz (BEMPA) [13] was proposed to simulate such systems using standard binary encoding for the bosonic modes. In this encoding, the available states of a bosonic mode are represented by a qubit register where the k-th qubit corresponds to a weight of 2^{k-1}. To enforce global particle conservation, the total sum of these weights across all modes must be constant.

The BEMPA circuit relies on two fundamental building blocks, whose generators we denote $\hat{G}_A$ and $\hat{G}_B$ (using hats to distinguish them from the G-gates of section 2). The $\hat{A}$ gate acts on two qubits of the same significance 2^{k-1} (belonging to different bosonic modes) and performs an exact rotation in the subspace spanned by $|01\rangle$ and $|10\rangle$. Its generator is $\hat{G}_A = i(|01\rangle\langle 10| - |10\rangle\langle 01|)$, which in Pauli form reads

$$\hat{G}_A = \tfrac{1}{2}(X \otimes Y - Y \otimes X)\,. \tag{3.1}$$

The $\hat{B}$ gate acts on three qubits, transferring amplitude between two qubits of significance 2^{k-1} and one qubit of significance 2^k; its generator is $\hat{G}_B = i(|001\rangle\langle 110| - |110\rangle\langle 001|)$. The corresponding Pauli string representation is given by

$$\hat{G}_B = \tfrac{1}{4}(X \otimes X \otimes Y - Y \otimes Y \otimes Y - X \otimes Y \otimes X - Y \otimes X \otimes X)\,. \tag{3.2}$$

Crucially, just as the G gate is applied directly to the qubit registers without intervening non-local Jordan-Wigner Z-strings, so too are the BEMPA $\hat{A}$ and $\hat{B}$ gates.

Proposition 2. *The BEMPA $\hat{A}$ and $\hat{B}$ gates, applied directly without non-local Z-strings, generate the full $\mathfrak{so}(w_{\text{boson}})$ Lie algebra on the conserved-particle-number subspace of dimension w_{boson}. Consequently, circuits composed of these gates can universally prepare any real state within this subspace.*

Proof. The argument follows the projection mechanism of Proposition 1. A single application of a bare BEMPA $\hat{A}$ or $\hat{B}$ gate acts symmetrically across all possible configurations of the spectator qubits. To isolate individual rotations, the projection mechanism requires both a bare generator L and its Z-dressed counterpart $L\,Z_k$ for some spectator qubit k, so that their linear combination yields the projector $(I \pm Z_k)/2$. The key identity that realises this for BEMPA is the commutator of a bare 2-body $\hat{G}_A$ gate with a bare 3-body $\hat{G}_B$ gate sharing one qubit,

$$[\hat{G}_{A;ik},\ \hat{G}_{B;kjl}] = -i\hat{G}_{B;ijl}Z_k\,, \tag{3.3}$$

which is verified term by term from the Pauli decompositions above in appendix A.

Since the bare $\hat{G}_{B;ijl}$ is also available as a native gate, the pair $\hat{G}_{B;ijl}$ and $\hat{G}_{B;ijl}Z_k$ enables the projection mechanism on qubit k. The same mechanism extends to every spectator qubit: for any spectator qubit k' at significance level $2^{k'-1}$, there exists at least one $\hat{G}_A$ generator pairing qubits of that significance across two bosonic modes, and at least one $\hat{G}_B$ generator bridging that level to the next. This is guaranteed by the structure of the binary encoding, in which every significance level hosts one qubit per bosonic mode, ensuring that the overlap conditions for (3.3) are always satisfiable. The inductive dressing argument of Proposition 1 then carries through: since the Z factors produced by (3.3) act on qubits disjoint from the active sites of subsequent commutators, iterating the procedure builds up an arbitrary product of spectator Z operators. Linear combinations of differently dressed generators yield the full projector product that isolates the single-plane generator $E_{xy} = |x\rangle\langle y| - |y\rangle\langle x|$ for each pair of basis states connected by a single $\hat{A}$ or $\hat{B}$ transition.

To obtain E_{xy} for pairs not connected by a single transition, we use the connectivity of the configuration graph. The $\hat{A}$ gates redistribute amplitude within each significance level and the $\hat{B}$ gates transfer amplitude between adjacent levels. Their combined action connects all valid binary configurations that share the same total particle number, since any redistribution of particle counts among the bosonic modes can be achieved by a sequence of such moves. Iterated application of the commutator identity $[E_{xy}, E_{yz}] = E_{xz}$ along paths in this graph then constructs all $w_{\text{boson}}(w_{\text{boson}}-1)/2$ generators of $\mathfrak{so}(w_{\text{boson}})$. □

4 Application to the fuzzy sphere regularisation of the 3D Ising model

We now turn our attention to the physical context that strongly motivated the preceding algebraic developments: the non-perturbative simulation of conformal field theories on quantum computers. We consider the non-commutative fuzzy sphere geometry as a regularisation of the 3D Ising critical point [14]. This approach is particularly advantageous because, unlike lattice regularisations, it preserves the rotation symmetry that is an essential ingredient of any CFT in three dimensions.

4.1 Lowest Landau level physics and fuzzy sphere

Consider electrons (charged particles) on a sphere with a magnetic monopole of charge $4\pi s$, $2s \in \mathbb{Z}$, at its centre. The purpose of the monopole is to create a radial magnetic field. If we neglect electron-electron interactions, then due to their interaction with the magnetic field the electrons will occupy Landau levels [31]. As it turns out, the interaction Hamiltonian is given in terms of the $SU(2)$ Casimir operator and, consequently, its eigenvalue problem is reduced to obtaining the irreducible

representations of $SU(2)$. For electrons in the lowest Landau level (LLL) in the presence of the monopole, such a representation is furnished by the functions[2]

$$\Phi_m(\theta,\phi) = (-1)^m \sqrt{\frac{(2s+1)!}{4\pi(s-m)!(s+m)!}}\, e^{im\phi} \cos^{s-m}\left(\frac{\theta}{2}\right) \sin^{s+m}\left(\frac{\theta}{2}\right), \qquad m = -s, -s+1, \ldots, s, \tag{4.1}$$

where θ, ϕ are coordinates on the sphere. These are obtained as the $l = s$ case of ${}_sY_{lm}(\theta,\phi)$ spin-weighted or monopole spherical harmonics [32–34]. There are obviously $2s+1$ orbitals in the LLL, so $N = 2s+1$ electrons will fill the LLL.

The Landau orbitals (4.1) are in one-to-one correspondence with states on the fuzzy sphere non-commutative geometry [35, 36]. Here the fuzzy sphere of radius R is a fuzzy manifold whose coordinates $\hat{x}_i, i = 1,2,3$, satisfy $SU(2)$ commutation relations,

$$[\hat{x}_i, \hat{x}_j] = i\frac{R}{s}\varepsilon_{ijk}\hat{x}_k\,, \tag{4.2}$$

where ε_{ijk} is the fully antisymmetric Levi-Civita symbol with $\varepsilon_{123} = 1$. Thus, fuzzy sphere geometry and LLL physics are equivalent.

To realise a phase transition in the same universality class as the 3D Ising model, which must involve the breaking of a $\mathbb{Z}_2$ global symmetry, let us assign a $\mathbb{Z}_2$ isospin quantum number to each electronic mode, which we will denote by an up or down arrow. This effectively doubles the number of electrons that can fit in the LLL, since now there are two distinct electron states that can occupy each Landau orbital. With the additional isospin quantum number, $N = 2s+1$ electrons provide a half-filling of the LLL.

So far the electrons we have considered have been interacting only with the background magnetic field of the monopole. Let us now add electron-electron interactions and assume that all electrons remain in the LLL. To represent the interactions, let us consider second-quantisation creation and annihilation operators $c^\dagger_{m\uparrow}, c^\dagger_{m\downarrow}$ and $c_{m\uparrow}, c_{m\downarrow}$, respectively, for $m = -s, -s+1, \ldots, s$.[3] For brevity of notation we may define the 2-vector $\mathbf{c}_m = (c_{m\uparrow} \quad c_{m\downarrow})^T$, in terms of which we construct the Hamiltonian

$$H = \sum_{m_1,\ldots,m_4=-s}^{s} V_{m_1,m_2,m_3,m_4}\left[(\mathbf{c}^\dagger_{m_1}\mathbf{c}_{m_4})(\mathbf{c}^\dagger_{m_2}\mathbf{c}_{m_3}) - (\mathbf{c}^\dagger_{m_1}\sigma^z\mathbf{c}_{m_4})(\mathbf{c}^\dagger_{m_2}\sigma^z\mathbf{c}_{m_3})\right] - h\sum_{m=-s}^{s}\mathbf{c}^\dagger_m\sigma^x\mathbf{c}_m, \tag{4.3}$$

where σ^x and σ^z are the first and third Pauli matrices, respectively, h is a transverse magnetic field and the interaction potential is given by

$$V_{m_1,m_2,m_3,m_4} = \delta_{m_1+m_2,m_3+m_4}\sum_{l=0}^{2s} V_l(4s-2l+1)\begin{pmatrix} s & s & 2s-l \\ m_1 & m_2 & -m_1-m_2\end{pmatrix}\times\begin{pmatrix} s & s & 2s-l \\ m_3 & m_4 & -m_3-m_4\end{pmatrix}, \tag{4.4}$$

where V_l are coefficients that depend on the specific type of electron-electron interactions one uses and $\left(\begin{smallmatrix} j_1 & j_2 & j_3 \\ m_1 & m_2 & m_3\end{smallmatrix}\right)$ is the Wigner $3j$ symbol. One can easily check that H preserves particle number and is invariant under the $\mathbb{Z}_2$ transformation

$$c_{m\uparrow} \leftrightarrow c_{m\downarrow}\,. \tag{4.5}$$

By virtue of being placed on a sphere, it also enjoys an $SU(2)$ rotation symmetry, with $\mathbf{c}_m, m = -s, -s+1, \ldots, s$, forming a spin-$s$ representation of $SU(2)$.

[2] $\Phi_m(\theta,\phi), m = -s, -s+1, \ldots, s$, furnish a spin-$s$ representation of $SU(2)$.

[3] As an example, the operator $c^\dagger_{s\uparrow}$ creates an electron in a state with S_z eigenvalue s and $\mathbb{Z}_2$ isospin $\uparrow$ in the LLL.

Following [14], we choose

$$V_0 \neq 0\,, \qquad V_1 = 1\,, \qquad V_l = 0 \text{ for } l > 1\,, \tag{4.6}$$

which describe ultra-local density-density interactions on the physical sphere. One can then see that by tuning the couplings V_0 and h the system will move between different phases. For $h = 0$ there will be two degenerate ground states, namely

$$|\Psi_\uparrow\rangle = \prod_{m=-s}^{s} c^\dagger_{m\uparrow}\,|0\rangle\,, \qquad |\Psi_\downarrow\rangle = \prod_{m=-s}^{s} c^\dagger_{m\downarrow}\,|0\rangle\,, \tag{4.7}$$

corresponding to a phase in which the $\mathbb{Z}_2$ symmetry is spontaneously broken. For $h \gg V_0$ the transverse magnetic field term in (4.3) will dominate, so that there will be a single ground state,

$$|\Psi_x\rangle = \prod_{m=-s}^{s} (c^\dagger_{m\uparrow} + c^\dagger_{m\downarrow})\,|0\rangle\,, \tag{4.8}$$

corresponding to a phase with unbroken $\mathbb{Z}_2$ symmetry. One then expects that there will be a phase transition between these regimes lying in the Ising universality class. This is attained on a line in the V_0-h plane. The point

$$V_0 = 4.75\,, \qquad h = 6.32 \tag{4.9}$$

lies on that critical line and yields surprisingly small finite-size effects [14].

4.2 State-operator correspondence

Let us briefly describe how conformal field theories allow the extraction of physical observables from the fuzzy sphere construction. The interested reader can find more details in standard references such as [37].

It is widely believed that the 3D Ising model at criticality is a CFT. CFTs are quantum field theories that are invariant under the conformal group, whose generators include translations P_μ, rotations $M_{\mu\nu}$, dilatations D, and special conformal transformations K_μ. The fundamental degrees of freedom of a CFT are captured by operators appearing in correlation functions. Any operator $\mathcal{O}$ has a well-defined scaling dimension $\Delta_\mathcal{O}$, which is revealed by the action of D on it and which governs its 2-point correlation function:

$$\langle \mathcal{O}(x)\mathcal{O}(0)\rangle = \frac{1}{|x|^{2\Delta_\mathcal{O}}}\,. \tag{4.10}$$

Computing $\Delta_\mathcal{O}$ for the operators of a given CFT is a highly non-trivial problem, especially in strongly-coupled cases like the 3D Ising model.

One crucial feature of conformal transformations is that they can map flat space $\mathbb{R}^3$ to the "cylinder" $S^2 \times \mathbb{R}$. Under this mapping, the generator of dilatations D on $\mathbb{R}^3$ becomes the generator of time translations on $S^2 \times \mathbb{R}$, i.e. the Hamiltonian on S^2. Operators on $\mathbb{R}^3$ give rise to states on S^2 and vice versa; this is the state-operator correspondence. The scaling dimensions of the operators are therefore identical to the eigenvalues of the S^2 Hamiltonian.

The fuzzy sphere construction is based on the idea that H of (4.3) with the parameter choices (4.6) and (4.9) is precisely such a Hamiltonian. Extracting its energy spectrum therefore gives direct access to the low-lying scaling dimensions of the 3D Ising CFT. Moreover, the identification of eigenstates enables the determination of 3-point function coefficients, which constitute additional CFT observables. As explained in [38, 39], this is achieved by computing ratios of appropriate amplitudes and becomes precise in the limit of infinite system size. We carry this out explicitly for $f_{\sigma\sigma\epsilon}$ in section 4.6, using the states prepared variationally in section 4.5.

Operator	Bootstrap	$N = 4$	Errors	ℓ	$\mathbb{Z}_2$
$\mathbb{1}$	0	0	N/A	0	Even
σ	0.518	0.53014	2.3%	0	Odd
ϵ	1.413	1.38210	2.2%	0	Even
$\partial_\mu\sigma$	1.518	1.52202	0.26%	1	Odd
$\partial_\mu\epsilon$	2.413	2.33705	3.1%	1	Even
$\Box\sigma$	2.518	2.42728	3.6%	0	Odd
$\partial_\mu\partial_\nu\sigma$	2.518	2.42839	3.6%	2	Odd
$\partial_\mu\partial_\nu\partial_\rho\sigma$	3.518	2.84791	19%	3	Odd
$T_{\mu\nu}$	3	3	N/A	2	Even
$\partial_\mu\partial_\nu\epsilon$	3.413	3.12561	8.4%	2	Even
$\partial_\mu\Box\sigma$	3.518	3.29166	6.4%	1	Odd
$\Box\epsilon$	3.413	3.57740	4.8%	0	Even
$\partial_\rho T_{\mu\nu}$	4	3.66295	8.4%	3	Even
ϵ'	3.830	4.01886	4.9%	0	Even
$\varepsilon_\mu{}^{\rho\lambda}\partial_\lambda T_{\rho\nu}$	4	4.05373	1.3%	2	Even
$\sigma_{\mu\nu}$	4.18	4.24153	1.5%	2	Odd
$\sigma_{\mu\nu\rho}$	4.638	4.61755	0.44%	3	Odd
$\partial_\mu\partial_\nu T_{\rho\lambda}$	5	4.85625	2.9%	4	Even

Table 1. Low-lying spectrum of the 3D Ising CFT as obtained with ED from the fuzzy-sphere regularisation method for $N = 4$ electrons, at the couplings (4.6) and (4.9). The full obtained spectrum is presented and conformal bootstrap results [42] are also included for comparison. The naming of operators follows standard notation in the physics literature. The spin ℓ and $\mathbb{Z}_2$ charge of each level are obtained as expectation values of the angular momentum Casimir and of the $\mathbb{Z}_2$ operator (4.5) in the corresponding eigenvector.

4.3 Classical extraction via exact diagonalisation

The second-quantisation Hamiltonian (4.3) can be conveniently constructed using `OpenFermion` [40], which provides methods for performing computations with fermionic creation and annihilation operators directly. Since H preserves particle number and total azimuthal angular momentum S_z, it is block-diagonal in these quantum numbers. Working at half-filling and restricting to the $S_z = 0$ sector reduces the Hilbert space dimension considerably.[4] This sector is sufficient because every spin-ℓ multiplet in the spectrum contains an $\ell_z = 0$ component that appears in this sector.

To find the spectrum we can proceed with exact diagonalisation (ED) for this small system size. To extract operator dimensions, the ground-state energy is subtracted from all energies, so that the ground state corresponds to the identity operator with scaling dimension zero. The shifted spectrum is then rescaled by an overall factor, fixed by demanding that the leading $\mathbb{Z}_2$-even state with spin two has energy exactly 3, as required for the energy-momentum tensor, $T_{\mu\nu}$, of the 3D Ising CFT. Table 1 shows the resulting low-lying spectrum for $N = 4$ electrons, compared with conformal bootstrap results [42].

The primary limitation of this classical approach is the rapid growth of the Hilbert space dimension with system size. While the $N = 4$ case is easily tractable, the dimension grows quickly: for example, the $N = 40$ case requires a Hilbert space of dimension 2,944,055,592 (for the $S_z = 0$ sector). On a quantum computer, by contrast, we would need only 80 qubits (two per fermionic mode). This provides strong motivation for exploring quantum simulation of this problem, as it holds the promise of reaching system sizes far beyond the reach of any classical method. Cases requiring even more

[4] The dimension of the Hilbert space sufficient for our considerations is given by sequence A125809 of the Online Encyclopedia of Integer Sequences [41].

fermionic modes, such as the Heisenberg model recently studied in [43–45] and others [46–58], would also benefit greatly from a realisation on a quantum computer.

4.4 Gates for preserving spin symmetries

For the most efficient implementation of the fuzzy sphere model, states must not only be half-filled but also maintain a total azimuthal angular momentum $S_z = 0$. This requires unitary operations to commute with the S_z operator. Symmetry constraints such as these are commonly encountered in a variety of problems.

The total azimuthal angular momentum on the fuzzy sphere is the sum of the magnetic quantum numbers m of all occupied orbitals:

$$S_z = \sum_k m_k \frac{I - Z_k}{2} \,. \tag{4.11}$$

For any parametrised continuous rotation to unconditionally preserve the $S_z = 0$ constraint, its Lie algebra generator must commute with the S_z operator.

If we consider the 2-qubit G-gate, we have

$$[L_{ij}, S_z] = i(m_i - m_j)(X_iX_j + Y_iY_j) \,. \tag{4.12}$$

This commutator vanishes if and only if $m_i = m_j$. For our encoding, this restricts $G^{(2)}$ gates to swapping the two isospins ($\uparrow$ and $\downarrow$) of the *same* m orbital. Such operations cannot change overall orbital occupancies and are thus individually insufficient to explore the full $S_z = 0$ subspace.

To achieve transitions that change orbital occupancies while simultaneously preserving S_z, we must introduce 4-qubit interaction gates, denoted $G^{(4)}$, which act on an ordered quartet of qubits (i, j, k, l) and generate the transition $|1100\rangle \leftrightarrow |0011\rangle$. This operation preserves S_z provided $m_i + m_j = m_k + m_l$. The continuous parametrised gate is given by

$$G^{(4)}_{ijkl}(\theta) = \exp\left[i\frac{\theta}{8}L_{ijkl}\right] \,, \tag{4.13}$$

with the explicit generator

$$\begin{aligned} L_{ijkl} = {} & X_iX_jX_kY_l + X_iX_jY_kX_l - X_iY_jX_kX_l - Y_iX_jX_kX_l \\ & + X_iY_jY_kY_l + Y_iX_jY_kY_l - Y_iY_jX_kY_l - Y_iY_jY_kX_l \,. \end{aligned} \tag{4.14}$$

Analogous to the 2-qubit gate, the improper reflection form of the rotation is

$$A^{(4)}_{ijkl}(\theta) = A^{(4)}_{ijkl}(0)G^{(4)}_{ijkl}(-\theta) \,, \tag{4.15}$$

where the fixed reflection component

$$A^{(4)}_{ijkl}(0) = I - 2\left(\frac{I - Z_i}{2}\right)\left(\frac{I - Z_j}{2}\right)\left(\frac{I + Z_k}{2}\right)\left(\frac{I + Z_l}{2}\right) \tag{4.16}$$

operates as the identity on all states except $|1100\rangle$, to which it applies a -1 phase. For clarity of presentation, we will use $G^{(2)}$ for (2.3) and $A^{(2)}$ for (2.8) from now on.

Proposition 3. *Circuits built from 2-qubit $G^{(2)}$ and 4-qubit $G^{(4)}$ gates, applied directly without non-local Z-strings and chosen to respect the strict symmetry constraint $S_z = 0$, dynamically generate the full continuous orthogonal group on that symmetry-constrained target subspace.*

Proof. The argument proceeds as in Proposition 1, but now both $G^{(2)}$ and $G^{(4)}$ generators contribute. Evaluating nested commutators of overlapping gates produces intermediate Pauli Z operators on shared spectator qubits. In particular, commuting overlapping generators yields Z-strings on the shared sites, just as commuting L_{ik} with L_{kj} produces Z_k in the 2-qubit case. For instance, commuting a 2-body and a 4-body generator sharing one qubit gives

$$[L_{ij}, L_{jklm}] = -2iL_{iklm}Z_j\,, \tag{4.17}$$

while commuting two 4-body generators that overlap on two qubits gives

$$[L_{ijkl}, L_{klmn}] = -4iL_{ijmn}(Z_k + Z_l)\,. \tag{4.18}$$

Both identities are derived explicitly in appendix A. They provide the base case for the inductive dressing of Proposition 1: since the Z factors act on qubits disjoint from the active sites of subsequent commutators, the same induction builds up an arbitrary product of spectator Z operators. Linear combinations of differently dressed generators then yield the full projector product that isolates the single-plane generator E_{xy} for each pair of basis states connected by a single allowed transition.

To obtain E_{xy} for pairs not directly connected, we use the connectivity of the configuration graph on the $S_z = 0$ subspace. The $m_i = m_j$ transitions (from $G^{(2)}$ gates) freely exchange the isospin of any single orbital without altering orbital occupancies, while the $m_i + m_j = m_k + m_l$ transitions (from $G^{(4)}$ gates) transfer a pair of particles between any two pairs of orbitals with equal total magnetic quantum number. Since repeated application of these two types of moves can rearrange the particle occupancies into any valid $S_z = 0$ pattern, the graph is fully connected. Iterated application of the commutator identity $[E_{xy}, E_{yz}] = E_{xz}$ along paths in this graph therefore constructs all generators of $\mathfrak{so}(w_0)$, where w_0 is the dimension of the symmetry-constrained subspace. □

4.5 Variational quantum simulation

The Python code used in this subsection and below can be found in a dedicated GitHub repository, andstergiou/fuzzy-ising-vqe, which is archived on Zenodo [59].

For the $N = 4$ case we have $s = 3/2$, with the magnetic quantum number taking the values $m = -\frac{3}{2}, -\frac{1}{2}, \frac{1}{2}, \frac{3}{2}$. Each value of m accommodates two isospin states, giving $N_q = 8$ qubits in total. Restricting to the half-filled $S_z = 0$ subspace leaves $w_0 = 18$ allowed basis states.

Using Proposition 3 and Lemma 1, we find that a circuit with exactly $w_0 - 1 = 17$ independently parametrised gates has a full-rank Jacobian at a reference point, establishing that the Jacobian has rank $w_0 - 1$ at almost every parameter value. This local criterion does not, however, certify that the circuit map $f\colon T^{w_0-1} \to S^{w_0-1}$ is globally surjective. Numerical optimisation over the circuit parameters (minimising $\|f(\boldsymbol{\theta}) - \mathbf{y}\|^2$ for random targets $\mathbf{y} \in S^{w_0-1}$) reveals that minimal 17-parameter circuits generically miss a small fraction of the sphere: the best such circuit reached 47 out of 50 random targets, with unreachable targets plateauing at residual costs $\sim 10^{-4}$, clearly separated from the $\sim 10^{-15}$ costs of reachable targets. Adding two extra parameters closes this gap: the 19-parameter G-gate circuit shown in figure 1 achieves 100/100 random-target reachability across multiple independent seeds, with maximum cost $\sim 10^{-14}$, providing strong numerical evidence for global surjectivity. Even though it is not necessarily minimal or unique, we will use the circuit of figure 1 for our VQE tests. Its resource cost is given in appendix B.

Since our system involves fermionic creation and annihilation operators, we use OpenFermion [40] to simulate it. After a Jordan-Wigner transformation, the resulting Hamiltonian is imported into PennyLane [60], and the circuit in figure 1 is constructed using PennyLane's SingleExcitation and DoubleExcitation rotations. Our setup exploits all symmetries of the problem. Besides the

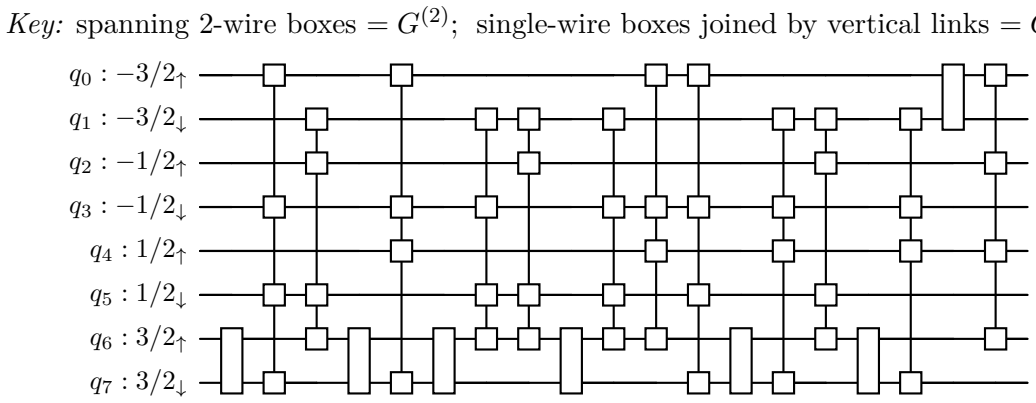


Figure 1. A 19-parameter G-gate circuit that is numerically surjective on the $S_z = 0$ subspace for $N = 4$, starting from the input state $|01010101\rangle$. The first 17 gates form a minimal spanning circuit of $G^{(2)}$ gates on the (q_6, q_7) isospin pair interleaved with various $G^{(4)}$ gates; the two additional gates $G^{(2)}_{01}$ and $G^{(4)}_{0624}$ close the local-to-global gap, so that numerical optimisation reaches every tested random target on S^{17}.

optimised circuit depth obtained, the benefits of using symmetries for convergence of VQE can also be considerable [61].

For our variational quantum simulation of the fuzzy sphere Hamiltonian for $N = 4$ electrons we use a classical state vector simulator (`lightning.qubit` backend of `PennyLane`). For the ground state, the cost function is the energy expectation value,

$$C_0(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|H|\psi(\boldsymbol{\theta})\rangle\,. \tag{4.19}$$

Optimisation proceeds in two phases: a coarse phase using the `Adam` adaptive gradient method [62] (step size 0.1, up to 500 iterations) to rapidly navigate the energy landscape and escape local minima, followed by a fine phase using gradient descent (step size 0.02) for precise convergence. Starting from a random initialisation, the VQE converges to the ground-state energy to within 1.2×10^{-7} in absolute value.

For the excited states we use the variational quantum deflation (VQD) method [15], augmenting the cost function with penalty terms that penalise overlaps with the previously found eigenstates:

$$C_k(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|H|\psi(\boldsymbol{\theta})\rangle + \sum_{j<k} \beta_j |\langle\psi(\boldsymbol{\theta})|\psi_j\rangle|^2\,, \tag{4.20}$$

where $|\psi_j\rangle$ is the j-th eigenstate already prepared and $\beta_j > 0$ is a penalty strength chosen to lift the corresponding eigenvalue above the target. We use $\beta_0 = 10$ for the first excited state and $\beta_0 = 30$, $\beta_1 = 20$ for the second excited state; each VQD step employs the same two-phase `Adam` (1000 iterations, step size 0.1) then gradient-descent (step size 0.01) strategy.

The convergence curves are displayed in figure 2, and the final energies are compared with ED in table 2. All three eigenvalues are recovered to within 4×10^{-7} in absolute energy. Moreover, the overlap matrix $|\langle\psi_k(\boldsymbol{\theta}^*)|\phi_j\rangle|^2$, where $|\phi_j\rangle$ are the exact eigenvectors, equals the identity matrix to ten significant figures, confirming that the circuit faithfully prepares each target eigenstate.

4.6 Extracting an OPE coefficient

The energy spectrum is not the only CFT observable accessible from the fuzzy sphere. As anticipated in section 4.2, 3-point function coefficients follow from ratios of matrix elements of a local operator [38, 39]. It is worth verifying that the variational pipeline reaches these too, since they probe the prepared eigenvectors rather than merely the eigenvalues, and are therefore a considerably more stringent test of the ansatz.

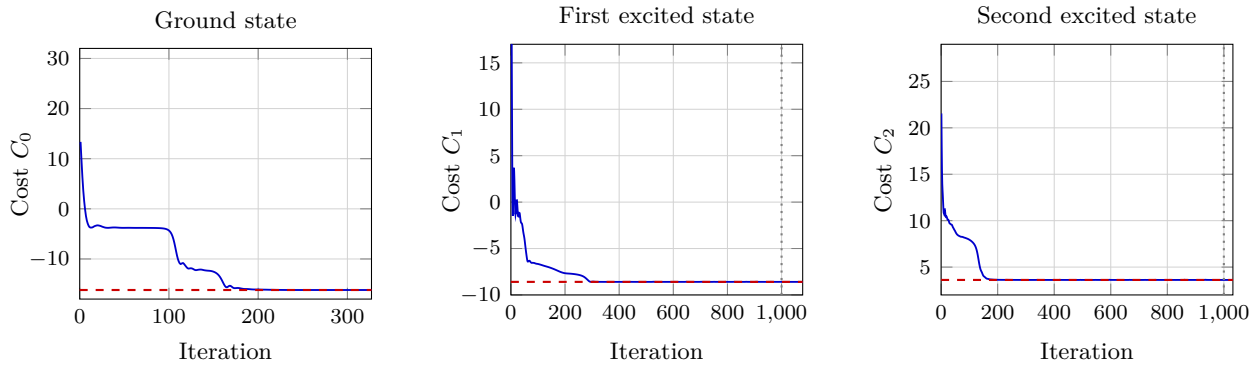


Figure 2. VQE/VQD convergence for the three lowest eigenstates of the fuzzy sphere Hamiltonian ($N = 4$). Each panel shows the cost function at every iteration (blue curve), with the exact eigenvalue from ED indicated by the red dashed line. For the excited states (centre and right panels), the dotted vertical line marks the transition from the `Adam` phase to the gradient-descent phase.

State	ED	VQE/VQD
Ground state	−16.18995794	−16.18995782
First excited state	−8.59299820	−8.59299785
Second excited state	3.61550790	3.61550806

Table 2. Energies of the three lowest eigenstates of the fuzzy sphere Hamiltonian ($N = 4$, $V_0 = 4.75$, $V_1 = 1$, $h = 6.32$) from ED and from the VQE/VQD algorithm using the circuit of figure 1.

The construction is particularly simple at the north pole of the sphere. From (4.1) the Landau orbital $\Phi_m(\theta, \phi)$ carries a factor $\sin^{s+m}(\theta/2)$, which vanishes at $\theta = 0$ unless $m = -s$. Any local bilinear evaluated at the north pole therefore collapses onto the single orbital $m = -s$. Choosing the $\mathbb{Z}_2$-odd channel gives

$$\Lambda_\sigma \propto \mathbf{c}^\dagger_{-s}\sigma^z\mathbf{c}_{-s} = n_{-s\uparrow} - n_{-s\downarrow}\,, \tag{4.21}$$

which commutes with both the particle number and S_z, and which is diagonal in the computational basis. Thus, it can be read off from a single Z-basis measurement of the two qubits encoding the $m = -s$ orbital.

Expanding the lattice operator in CFT operators, $\Lambda_\sigma = c_\sigma\sigma + \ldots$, the state-operator correspondence gives $\langle\mathbb{1}|\,\Lambda_\sigma\,|\sigma\rangle = c_\sigma$ from the 2-point normalisation and $\langle\epsilon|\,\Lambda_\sigma\,|\sigma\rangle = c_\sigma f_{\sigma\sigma\epsilon}$ from the 3-point function. The unknown overlap c_σ cancels in the ratio,

$$f_{\sigma\sigma\epsilon} = \frac{\langle\epsilon|\,\Lambda_\sigma\,|\sigma\rangle}{\langle\mathbb{1}|\,\Lambda_\sigma\,|\sigma\rangle}\,, \tag{4.22}$$

and, conveniently, the three states appearing here are exactly the ground state and the two excited states already prepared in table 2: $|\mathbb{1}\rangle$ is the ground state, $|\sigma\rangle$ the first excited state and $|\epsilon\rangle$ the second, as identified by their spin and $\mathbb{Z}_2$ quantum numbers in table 1. No further variational optimisation is needed.

Evaluating (4.22) on the VQE/VQD states of table 2 gives

$$f_{\sigma\sigma\epsilon}\big|_{\text{VQE/VQD}} = 0.96141\,, \qquad f_{\sigma\sigma\epsilon}\big|_{\text{ED}} = 0.96144\,, \tag{4.23}$$

the two agreeing to four significant figures. The most accurate determination of this quantity is given by the conformal bootstrap value $f_{\sigma\sigma\epsilon} = 1.0518537(41)$ [63]. The residual 8.6% discrepancy is a finite-size effect of the $N = 4$ truncation, not a failure of the variational preparation: the ED value computed in

N	Qubits	$w_0 - 1$	p (10^{-2})	p (10^{-4})	p (10^{-8})
2	4	3	2	3	3
4	8	17	12	17	19
6	12	119	42	116	141

Table 3. Number of parameters p required for a layered S_z-preserving circuit to reach the exact ground state to infidelity below 10^{-2}, 10^{-4} and 10^{-8}. The corresponding CNOT counts are $96, 142, 170$ for $N = 4$ and $444, 1336, 1686$ for $N = 6$. The value $p = 19$ recovered at $N = 4$ coincides with the parameter count of the circuit of figure 1, which was constructed independently.

the same 18-dimensional subspace carries the identical deviation. It is larger than the corresponding errors on the low-lying scaling dimensions in table 1, which is expected, since OPE coefficients are known to converge more slowly with system size than the spectrum [38]. The point of the exercise is that our symmetry-constrained variational method reproduces the exact subspace answer for an observable built from the eigenvectors, and does so using only measurements the circuit already supports.

4.7 Scaling with system size

The results above are for $N = 4$ alone, and a single system size cannot distinguish polynomial from exponential growth of the gate count. Since the near-term relevance of the construction depends on how that count grows, we now examine $N = 2, 4, 6$, for which the $S_z = 0$ half-filled subspace has dimensions $w_0 = 4, 18, 120$ on $4, 8, 12$ qubits respectively.

We first verify that the gate set of section 4 remains adequate as N grows. The configuration graph on the $S_z = 0$ subspace, with edges given by the allowed $G^{(2)}$ and $G^{(4)}$ transitions, is connected for every case we examined, including $N = 8$ (where $w_0 = 1064$), confirming the connectivity step in the proof of Proposition 3. We note in passing that connectivity of this graph is necessary but not sufficient: for $N = 2$ the two $G^{(2)}$ gates alone already connect all four basis states, yet they generate only states that factorise across the two orbitals and so cannot reach the entangled ground state. This is an instance of the same dimensional gap identified in Corollary 1, here between the 2-dimensional set of reachable factorised states and the 3-sphere on which the sector lives.

To measure the required circuit size we use a layered ansatz of fixed structure. One layer consists of all N $G^{(2)}$ gates together with a set of $G^{(4)}$ gates added until the configuration graph is connected; the layer is then repeated, and truncated at the required length, to give a circuit with p parameters. For each p we minimise the infidelity $1 - |\langle \psi_{\mathrm{gs}} | \psi(\boldsymbol{\theta}) \rangle|^2$ with the exact ground state from many random starts, and record the smallest p at which a chosen infidelity is reached. Optimising the infidelity directly, rather than the energy, isolates the expressivity of the circuit from the trainability of the VQE landscape; we found that greedy operator-selection schemes stall in local minima well before exhausting the expressivity of the ansatz, so the fixed-structure multi-start protocol gives a more faithful measure. Table 3 collects the results.

As we see from these results, for essentially exact preparation the required parameter count tracks the dimension of the subspace: $p(10^{-8}) = 3, 19, 141$ against $w_0 - 1 = 3, 17, 119$. Since w_0 grows exponentially with N, so does the cost of exact ground state preparation with this gate set. Our completeness theorems guarantee that the target is reachable at finite depth, but they carry no promise that the depth is small, and the data show that it is not. The picture is more favourable if only moderate accuracy is demanded: $p(10^{-2})$ grows from 12 to 42 between $N = 4$ and $N = 6$, a factor of 3.5, while w_0 grows by a factor of 6.7. Approximate preparation thus scales appreciably better than the subspace dimension, though still steeply, and it is in this regime that variational preparation on near-term devices is plausible.

The over-parametrisation noted in section 4.5 is a separate question, since it concerns reaching *arbitrary* states in the sector rather than one particular ground state, and it is the one on which the topological obstruction of section 2.2 bears. Repeating the random-target test used there, the smallest layered circuit that reaches every member of a set of random targets to infidelity below 10^{-8} has

$$p = 4,\ 19,\ 124 \qquad \text{for } N = 2,\ 4,\ 6\,, \tag{4.24}$$

against the minimal counts $w_0 - 1 = 3, 17, 119$. At $N = 4$ this independently reproduces the two extra parameters found in section 4.5. The over-parametrisation therefore persists and grows slowly in absolute terms, by $1, 2, 5$ parameters, but shrinks steadily relative to the minimal count, from 33% to 12% to 4%. The topological obstruction does not disappear as N increases, but neither does it become a leading contribution to the circuit size; the dominant cost is the minimal count $w_0 - 1$ itself.

It is worth recording that the physics improves with N even as the circuits become more expensive. Identifying the energy-momentum tensor and normalising on it as in section 4.3, the $N = 6$ system gives $\Delta_\sigma = 0.52180$ and $\Delta_\epsilon = 1.38471$, against 0.53014 and 1.38210 at $N = 4$ and the bootstrap values 0.518 and 1.413. The error on Δ_σ falls from 2.3% to 0.73%.

4.8 Simulation under a noise model

All results so far are for noiseless state vector simulation, which leaves open the question of what gate fidelity would actually be required to extract the spectrum on hardware. To answer this question, we repeat the $N = 4$ calculation with the state represented as a density matrix rather than a state vector, as the noise channels map the pure ansatz state to a mixed one. The circuit of figure 1 is decomposed into its 182 CNOT gates and 196 single-qubit rotations, and we apply a two-qubit depolarising channel of rate γ after every CNOT together with a single-qubit depolarising channel of rate $\gamma/10$ after every single-qubit gate. The VQE and VQD optimisations are re-run in full at each value of γ, rather than reusing the noiseless parameters. Scaling dimensions are obtained with the $T_{\mu\nu}$ normalisation held fixed at its exact value, which isolates the effect of the noise on the measured gaps.

Because the ansatz preserves both particle number and S_z exactly, any population that the noise moves out of the half-filled $S_z = 0$ sector is detectable in a computational-basis measurement and can be discarded. We therefore also report the symmetry-post-selected estimator, in which the density matrix is projected back onto the physical sector and renormalised before the energies are read off. This costs nothing beyond the measurements already performed, and is available precisely because the circuit is symmetry-preserving by construction.

The results are shown in figure 3. The noise-induced error grows essentially linearly in γ throughout the regime in which the calculation remains meaningful, as expected when the dominant effect is the accumulation of independent errors over a fixed number of gates. Reading off the crossing with the finite-size error gives

$$\gamma^* \approx 2.0 \times 10^{-4} \quad \text{(raw)}\,, \qquad \gamma^* \approx 6.5 \times 10^{-4} \quad \text{(symmetry post-selected)}\,, \tag{4.25}$$

so that below these rates the $N = 4$ truncation, rather than the hardware, limits the accuracy of the extracted spectrum. We should be clear that this is a demanding target: it is roughly an order of magnitude beyond typical current two-qubit gate error rates, and it is a direct consequence of the 182 CNOT gates in the circuit. Symmetry post-selection improves things by a factor of about three in the tolerable error rate, and by a factor of between three and four in the error at fixed γ throughout the regime where it is useful, which is a concrete practical benefit of insisting that the ansatz preserve the symmetry exactly rather than approximately.

The method degrades rapidly above $\gamma = 10^{-3}$, where the weight remaining in the physical sector has fallen to 0.85 and the raw Δ_σ is underestimated by 12%, although post-selection still recovers it

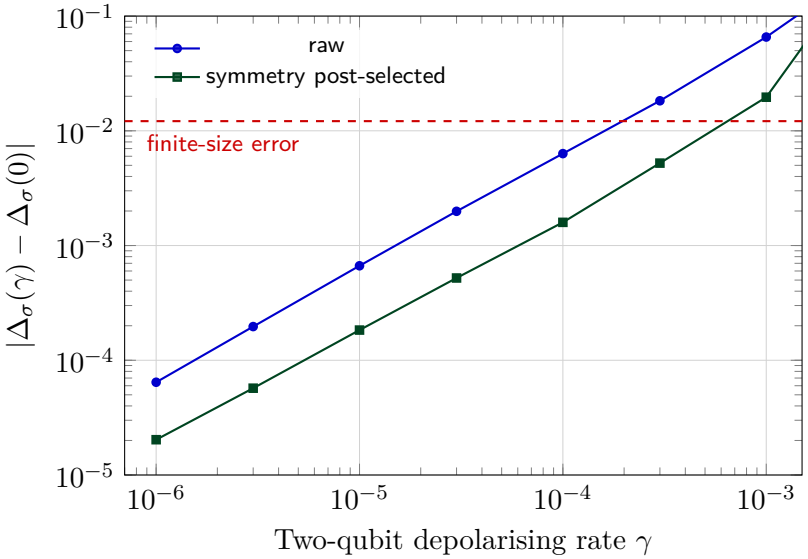


Figure 3. Noise-induced error on the extracted scaling dimension Δ_σ as a function of the two-qubit depolarising rate γ, for $N = 4$. The dashed line is the finite-size error $|\Delta_\sigma(0) - \Delta_\sigma^{\text{bootstrap}}| = 1.21 \times 10^{-2}$. The noise-induced error grows essentially linearly in γ and crosses the finite-size error at $\gamma \approx 2.0 \times 10^{-4}$ for the raw estimator, and at $\gamma \approx 6.5 \times 10^{-4}$ after symmetry post-selection.

to within 3.7%. Beyond the range plotted the breakdown is complete: by $\gamma = 3 \times 10^{-3}$ the weight is 0.63, the extracted spectrum has lost all resemblance to the CFT, and post-selection no longer helps at all, since the surviving population in the physical sector is itself close to uniform. We note that re-optimising the parameters at each γ, rather than reusing the noiseless ones, does not improve the extracted dimensions: to leading order the noise acts as a global depolarising channel, which rescales all energies by a common factor and adds a $\boldsymbol{\theta}$-independent constant, leaving the optimum where it was and contracting all gaps uniformly.

5 Extension to $SU(w)$ with independent complex phases

The results of section 2 generate the real orthogonal group $SO(w)$, the natural setting for states with purely real amplitudes. Many physical applications, e.g. systems lacking time-reversal symmetry and complex ground-state preparation, require the full special unitary group $SU(w)$. We now show that the Pauli Z dressing mechanism extends directly to $SU(w)$ once the gates carry independent complex phases.

Replacing the real Givens rotation $G(\theta)$ with a parametrised $SU(2)$ rotation on the active $\{|01\rangle, |10\rangle\}$ subspace introduces an additional phase ϕ:

$$G(\theta, \phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -e^{i\phi}\sin\theta & 0 \\ 0 & e^{-i\phi}\sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{5.1}$$

In exponential form,

$$G_{ij}(\theta, \phi) = \exp\left[i\frac{\theta}{2}\left(\cos\phi\, L^a_{ij} - \sin\phi\, L^s_{ij} \right) \right], \tag{5.2}$$

where $L^a_{ij} \equiv L_{ij} = X_iY_j - Y_iX_j$ is the real (antisymmetric) hopping generator defined earlier, and $L^s_{ij} = X_iX_j + Y_iY_j$ is the corresponding imaginary (symmetric) hopping generator. When ϕ varies independently across different gates in the circuit, the continuous generators span both hopping directions.

The proof of Proposition 1 rests on the overlapping commutator (2.13), whose Z spectator on the shared qubit enables the isolation of single-plane generators via dressing. We now show that the same spectator structure emerges for every combination of L^s and L^a. Using $[X, Y] = 2iZ$ and the tensor product identity $[A_iB_j, C_jD_k] = A_i[B_j, C_j]D_k$, we compute

$$[L^a_{ik}, L^a_{kj}] = -2iL^a_{ij}Z_k\,, \tag{5.3a}$$
$$[L^s_{ik}, L^s_{kj}] = 2iL^a_{ij}Z_k\,, \tag{5.3b}$$
$$[L^a_{ik}, L^s_{kj}] = -2iL^s_{ij}Z_k\,. \tag{5.3c}$$

The first line is the commutator used in the proof of Proposition 1; the remaining two extend it. Every commutator produces a Pauli Z on the shared qubit, regardless of whether the constituent generators are real or imaginary. Commuting two generators of the same type (aa or ss) yields a spectator-dressed real generator L^aZ, while commuting generators of different types (sa or as) yields a spectator-dressed imaginary generator L^sZ.

Because the spectator Z emerges uniformly in all cases, the Lie algebraic generation carries through unchanged with L^s in place of L^a. Concretely, isolating single-plane rotations as in the $\mathfrak{so}(w)$ case applied to L^a yields the single-plane antisymmetric generators $E^a_{xy} = |x\rangle\langle y| - |y\rangle\langle x|$; the same operations applied to L^s isolate single-plane symmetric generators $E^s_{xy} = i(|x\rangle\langle y| + |y\rangle\langle x|)$. The algebra $\mathfrak{su}(w)$ requires three types of generators for each pair (x, y): the antisymmetric E^a_{xy}, the symmetric E^s_{xy}, and a diagonal generator $E^d_{xy} = i(|x\rangle\langle x| - |y\rangle\langle y|)$. The first two are thus provided structurally by L^a and L^s respectively, and the third follows from their commutator:

$$[E^a_{xy}, E^s_{xy}] = 2E^d_{xy}\,. \tag{5.4}$$

Therefore, the diagonal generators require no additional circuit elements. The set $\{E^a_{xy}, E^s_{xy}, E^d_{xy}\}$ over all pairs (x, y) therefore spans the full $\mathfrak{su}(w)$ Lie algebra on the target subspace.

This result connects to the criteria of Yan et al. [22], whose Theorem 1 proves that hardware-efficient generators (without non-local Jordan-Wigner Z-strings) can span $\mathfrak{su}(w)$ on the constant Hamming-weight subspace only if they break free-fermion integrability via a two-body diagonal Z_iZ_j component. At the level of a single link, the same-site commutator

$$[L^a_{ij}, L^s_{ij}] = 4i(Z_i - Z_j) \tag{5.5}$$

is single-body and remains within the free-fermion algebra. This means that the complex phase alone does not break integrability. In our framework, it is the Pauli Z dressing that does so: dressing a hopping term with Z-string projectors produces effective generators containing multi-body diagonal components Z_iZ_j and higher, exactly the terms demanded by the Yan et al. condition. The independent phase ϕ then supplies the second hopping direction within this already non-integrable structure.

The same logic applies to further symmetry-constrained sectors. For the $S_z = 0$ subspace of section 4, recall that the 4-qubit gates $G^{(4)}_{ijkl}(\theta)$ have a real generator L^a_{ijkl} driving $|1100\rangle\langle 0011| - |0011\rangle\langle 1100|$. An independent phase ϕ now adds the imaginary generator L^s_{ijkl}. The overlapping commutators of these 4-qubit generators produce spectator Z-strings by exactly the same mechanism as the 2-qubit case, so the dressing procedure of Proposition 3 isolates both real and imaginary single-plane generators across the connected graph of the $S_z = 0$ subspace, completing the $\mathfrak{su}(w_0)$ algebra. As a consistency check with the Yan et al. criterion, the same-site commutator $[L^s_{ijkl}, L^a_{ijkl}]$ is proportional to $|1100\rangle\langle 1100| - |0011\rangle\langle 0011|$, which expands into single-body (Z_i) and three-body ($Z_iZ_jZ_k$) Pauli terms, the latter manifestly breaking free-fermion integrability.

6 Conclusion

We have established a Lie algebraic framework for proving that hardware-efficient quantum gates, applied without non-local Jordan-Wigner strings, generate the full orthogonal group on constrained subspaces of fixed particle number and symmetry. Central to this is the mechanism of Pauli Z dressing: the emergence of Pauli Z operators through commutators of overlapping generators, which act as spectator projectors and allow multi-plane rotations to be decomposed into elementary single-plane generators E_{xy}. We formalised this in three propositions: Proposition 1 for the constant Hamming weight subspace, Proposition 2 for the BEMPA ansatz in bosonic systems with conserved particle number, and Proposition 3 for further symmetry-constrained sectors such as $S_z = 0$. Because these propositions rely on continuously parametrised generators spanning the relevant Lie algebras, they guarantee exact reachability across the target connected Lie group. This distinguishes our completeness results from traditional quantum gate universality proofs, which rely on discrete gate sets and only guarantee asymptotic density. The local Lemma 1 provides a computationally efficient sufficient criterion for establishing that a given minimal circuit can explore any direction on the target manifold from a generic parameter configuration, though it does not settle global reachability of every point in the manifold.

We applied this framework to two physical settings. For the BEMPA ansatz, Proposition 2 establishes that the $\hat{A}$ and $\hat{B}$ gates generate the full algebra for bosonic systems with conserved particle number. For the fuzzy sphere regularisation of the 3D Ising CFT, we constructed explicit circuits with 19 parameters that span the full relevant subspace for $N = 4$ electrons, with numerical evidence for global surjectivity, and demonstrated agreement of the extracted spectrum between ED and a VQE simulation. Beyond the spectrum, the same variationally prepared states yield the OPE coefficient $f_{\sigma\sigma\epsilon}$ from a ratio of amplitudes requiring no additional circuitry. Extending the study to $N = 2$ and $N = 6$ shows that the parameter count needed for exact ground state preparation tracks the dimension w_0 of the constrained subspace, and therefore grows exponentially with N, so that our completeness results should not be read as a promise of shallow circuits; approximate preparation is appreciably cheaper. Under a depolarising noise model the extracted scaling dimensions are limited by the hardware rather than by finite size once the two-qubit error rate exceeds $\gamma^* \approx 2 \times 10^{-4}$, which symmetry post-selection relaxes by roughly a factor of three.

We further showed in section 5 that the framework extends from $SO(w)$ to the full special unitary group $SU(w)$ once each gate carries an independent complex phase. The key observation is that the overlapping commutators producing spectator Z operators, the mechanism underlying all three propositions, emerge identically for the imaginary hopping generators $L^s_{ij} = X_iX_j + Y_iY_j$ as for the real generators $L^a_{ij} = X_iY_j - Y_iX_j$. The Pauli Z dressing therefore isolates both real and imaginary single-plane generators for every pair of basis states, spanning the full $\mathfrak{su}(w)$. This is consistent with the criteria of Yan et al. [22]: the integrability-breaking diagonal components their theorem demands arise from the Pauli Z dressing itself, while the independent phases supply the additional hopping directions.

We suspect that Pauli Z dressing is, in fact, a universal structural feature of hardware-efficient ansatze rather than a property specific to the systems studied here. The mechanism relies on a single algebraic ingredient: commutators of overlapping multi-qubit generators produce Z operators on shared qubits, which then serve as spectator projectors. This ingredient is available whenever the gate set contains generators with overlapping support on at least one qubit, a condition satisfied by essentially any non-trivial hardware-efficient circuit topology. It would be interesting to determine whether there exist natural constrained subspaces for which the Z dressing mechanism fails to generate the full algebra, or whether completeness is instead guaranteed for any connected gate graph acting on a connected configuration space. Such a classification could provide a unifying criterion for the expressibility of hardware-efficient ansatze across many fields of study. For example, establishing

exactly what set of independent complex phases enables $\mathfrak{su}(w)$ completeness across different symmetry-constrained physical models, and classifying their minimal gate sets and circuit topologies, remains a valuable target for future work.

More broadly, the systematic construction of minimal circuits, i.e. those with the minimum number of independent parameters, that are locally surjective onto a given target manifold remains an open problem. While the Jacobian criterion of Lemma 1 can verify specific ansatze locally, it offers no guidance on how to find one. A constructive procedure that, given a target subspace and a set of available gates, produces a minimal-depth spanning circuit, or establishes that no such circuit of a given topology exists, would significantly accelerate the deployment of variational algorithms on near-term hardware.

On the applications side, the fuzzy sphere construction is particularly promising for larger system sizes, where classical methods become impractical but the qubit count remains modest. At these sizes, noise-aware circuit optimisation and error mitigation strategies become essential. The noise analysis of section 4.8 suggests that the symmetry structure is itself a resource in this respect, since exact conservation makes leakage out of the physical sector detectable at no additional measurement cost; the algebraic framework developed here could inform the design of circuits that are simultaneously minimal in depth and robust to hardware errors.

Acknowledgments

AS is grateful to Jake Belton and Nadav Drukker for useful discussions and collaboration in the initial stages of this project, and to Petr Kravchuk for insightful conversations. Claude Code and Gemini were both used in this work for coding tests and claim verification. They also aided in the writing of this paper. AS is supported by the Royal Society under grant URF\R1\211417 and by STFC under grant ST/X000753/1. NS acknowledges use of resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC award NERSC DDR-ERCAP0038154.

Appendix

A Commutator identities

The proofs of Propositions 1, 2 and 3 all rest on commutators of overlapping generators producing Pauli Z operators on the shared qubits. In this appendix we verify these identities explicitly. They are all instances of a single elementary statement about Pauli strings, which also explains why a Z is produced in every case.

Lemma 2 (Overlap lemma)**.** *Let*

$$P = \tilde{P} \prod_{k \in S} p_k \,, \qquad Q = \tilde{Q} \prod_{k \in S} q_k \,, \tag{A.1}$$

be Pauli strings whose supports intersect exactly in the set S, with $p_k, q_k \in \{X, Y\}$ and with $\tilde{P}$, $\tilde{Q}$ supported on disjoint sets of qubits, both disjoint from S. Let $D = \{k \in S \colon p_k \neq q_k\}$, $d = |D|$, and $\eta_k = +1$ if $(p_k, q_k) = (X, Y)$ and $\eta_k = -1$ if $(p_k, q_k) = (Y, X)$. Then

$$[P, Q] = (i^d - (-i)^d) \Big(\prod_{k \in D} \eta_k \Big) \tilde{P}\tilde{Q} \prod_{k \in D} Z_k \,. \tag{A.2}$$

In particular, $[P, Q] = 0$ whenever d is even, and $[P, Q] = 2i\eta_k \tilde{P}\tilde{Q}Z_k$ when $D = \{k\}$.

Proof. Since $\tilde{P}$ and $\tilde{Q}$ have disjoint supports they commute with each other and with every p_k, q_k, so $PQ = \tilde{P}\tilde{Q}\prod_{k\in S} p_k q_k$ and $QP = \tilde{P}\tilde{Q}\prod_{k\in S} q_k p_k$. For $k \notin D$ we have $p_k q_k = q_k p_k = I$, while for $k \in D$ the relations $XY = iZ$ and $YX = -iZ$ give $p_k q_k = i\eta_k Z_k$ and $q_k p_k = -i\eta_k Z_k$. Hence $PQ = i^d(\prod_{k\in D}\eta_k)\tilde{P}\tilde{Q}\prod_{k\in D} Z_k$ and QP is the same expression with $i \to -i$, which yields (A.2). □

The Z operators are therefore not an accident of any particular gate set: two hardware-efficient generators built from X and Y only, overlapping on an odd number of qubits with mismatched Paulis there, *necessarily* produce a Z on precisely those qubits. Equivalently, writing $\sigma^{\pm} = \frac{1}{2}(X \pm iY)$, so that $L_{ij} = 2i(|01\rangle\langle 10| - |10\rangle\langle 01|)_{ij}$, the shared qubit is acted on by $\sigma^+\sigma^- = \frac{1}{2}(I + Z)$ or $\sigma^-\sigma^+ = \frac{1}{2}(I - Z)$ depending on the order of the product, and the commutator retains exactly the Z part.

The 2-qubit identity (2.13). With $L_{ik} = X_iY_k - Y_iX_k$ and $L_{kj} = X_kY_j - Y_kX_j$, the overlap is $S = \{k\}$. Of the four pairs of terms, the two with matching Pauli on k drop out by Lemma 2. The surviving pairs are (X_iY_k, X_kY_j), with $(p_k, q_k) = (Y, X)$ and hence $\eta_k = -1$, and $(-Y_iX_k, -Y_kX_j)$, with $(p_k, q_k) = (X, Y)$ and hence $\eta_k = +1$. They contribute $-2iX_iY_jZ_k$ and $(-1)(-1)(+1)2iY_iX_jZ_k$ respectively, so that

$$[L_{ik}, L_{kj}] = -2iX_iY_jZ_k + 2iY_iX_jZ_k = -2i(X_iY_j - Y_iX_j)Z_k = -2iL_{ij}Z_k\,. \tag{A.3}$$

The BEMPA identity (3.3). Here

$$\hat{G}_{A;ik} = \tfrac{1}{2}(X_iY_k - Y_iX_k)\,, \qquad \hat{G}_{B;kjl} = \tfrac{1}{4}(X_kX_jY_l - Y_kY_jY_l - X_kY_jX_l - Y_kX_jX_l)\,, \tag{A.4}$$

where in $\hat{G}_B$ the shared qubit k carries the first Pauli of each string. The two operators overlap only on k, so Lemma 2 applies to each of the $2 \times 4 = 8$ pairs of terms. Four pairs have matching Pauli on k ($d = 0$) and vanish; the surviving four are

$\hat{G}_{A;ik}$ term	$\hat{G}_{B;kjl}$ term	(p_k, q_k)	contribution
$+\frac{1}{2}X_iY_k$	$+\frac{1}{4}X_kX_jY_l$	(Y, X)	$-\frac{i}{4}X_iX_jY_lZ_k$
$+\frac{1}{2}X_iY_k$	$-\frac{1}{4}X_kY_jX_l$	(Y, X)	$+\frac{i}{4}X_iY_jX_lZ_k$
$-\frac{1}{2}Y_iX_k$	$-\frac{1}{4}Y_kY_jY_l$	(X, Y)	$+\frac{i}{4}Y_iY_jY_lZ_k$
$-\frac{1}{2}Y_iX_k$	$-\frac{1}{4}Y_kX_jX_l$	(X, Y)	$+\frac{i}{4}Y_iX_jX_lZ_k$

Summing the four contributions,

$$[\hat{G}_{A;ik}, \hat{G}_{B;kjl}] = -\tfrac{i}{4}(X_iX_jY_l - Y_iY_jY_l - X_iY_jX_l - Y_iX_jX_l)Z_k = -i\hat{G}_{B;ijl}Z_k\,, \tag{A.5}$$

which is (3.3). Note that the four surviving terms reassemble precisely into $\hat{G}_B$ with k replaced by i: the sign η_k supplied by Lemma 2 is exactly compensated by the relative sign carried by the corresponding term of $\hat{G}_A$, so the internal sign pattern of $\hat{G}_B$ is transported unchanged.

Two overlapping $\hat{B}$ gates generate the 4-body generators of section 4 by the same mechanism,

$$[\hat{G}_{B;ijk}, \hat{G}_{B;lmk}] = \tfrac{i}{8}L_{ijlm}Z_k\,, \tag{A.6}$$

where the two strings share only the highest-significance qubit k.

The identities of Proposition 3. For the mixed 2-body/4-body commutator the overlap is again a single qubit, $S = \{j\}$, carrying the first Pauli of L_{jklm}. Splitting L_{jklm} according to that Pauli,

$$\begin{aligned} L_{jklm} = {} & X_j(X_kX_lY_m + X_kY_lX_m - Y_kX_lX_m + Y_kY_lY_m) \\ & + Y_j(-X_kX_lX_m + X_kY_lY_m - Y_kX_lY_m - Y_kY_lX_m)\,, \end{aligned} \tag{A.7}$$

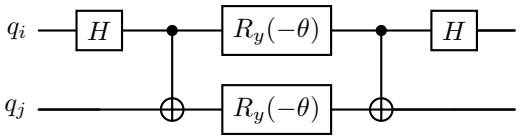


Figure 4. Decomposition of the 2-qubit Givens rotation $G^{(2)}_{ij}(\theta)$ into two CNOT gates and single-qubit rotations.

the term X_iY_j of L_{ij} pairs only with the X_j block ($\eta = -1$) and $-Y_iX_j$ only with the Y_j block ($\eta = +1$). The eight surviving contributions are

$$\begin{aligned}[L_{ij}, L_{jklm}] = -2iZ_j \big[&X_iX_kX_lY_m + X_iX_kY_lX_m - X_iY_kX_lX_m - Y_iX_kX_lX_m \\ &+ X_iY_kY_lY_m + Y_iX_kY_lY_m - Y_iY_kX_lY_m - Y_iY_kY_lX_m\big]\,,\end{aligned} \tag{A.8}$$

and the bracket is exactly L_{iklm}, so $[L_{ij}, L_{jklm}] = -2iL_{iklm}Z_j$.

For two 4-body generators overlapping on two qubits, $S = \{k, l\}$, the even-d part of Lemma 2 does the essential work: every pair of terms that disagrees on *both* shared qubits has $d = 2$ and therefore commutes, and pairs agreeing on both have $d = 0$ and also commute. Only pairs disagreeing on exactly one shared qubit survive, producing a single Z on that qubit. Collecting the terms that disagree on k alone and on l alone gives the two symmetric families

$$[L_{ijkl}, L_{klmn}] = -4iL_{ijmn}Z_k - 4iL_{ijmn}Z_l = -4iL_{ijmn}(Z_k + Z_l)\,, \tag{A.9}$$

as quoted in the proof of Proposition 3.

B Elementary gate decompositions

In this appendix, we provide decompositions of the gates used in this work into elementary single-qubit rotations and CNOT gates. Throughout, $R_k(\alpha) = \exp(-i\alpha\sigma_k/2)$ for $k = x, y, z$, and H denotes the Hadamard gate. Note the half-angle convention: $R_y(\alpha)|0\rangle = \cos(\alpha/2)|0\rangle + \sin(\alpha/2)|1\rangle$, so the argument of R_y must be 2θ to achieve a net rotation by θ in a 2-dimensional subspace.

2-qubit Givens rotation $G^{(2)}_{ij}(\theta)$. The Givens rotation acts as a rotation by θ in the $\{|01\rangle, |10\rangle\}$ subspace while leaving $|00\rangle$ and $|11\rangle$ invariant. This is essentially the `SingleExcitation` rotation of `PennyLane`. The Hadamard on q_i followed by a CNOT computes the parity of the two qubits into q_j; the two $R_y(-\theta)$ rotations then act with opposite effective signs in the even- and odd-parity sectors, producing a net rotation only in the active subspace; see figure 4.

2-qubit A-gate $A^{(2)}_{ij}(\theta)$. Since $A(\theta) = A(0)\,G(-\theta)$ and the discrete reflection matrix $A(0) = \mathrm{diag}(1, 1, -1, 1)$ flips the sign of the determinant in the active subspace, it is impossible to decompose $A^{(2)}_{ij}(\theta)$ using only two CNOTs and single-qubit $SU(2)$ rotations. The reflection $A(0)$ admits a single-CNOT implementation as $(I \otimes R_y(\pi/2))\,\mathrm{CNOT}\,(I \otimes R_y(-\pi/2))$, which combined with $G^{(2)}(-\theta)$ yields a 3-CNOT decomposition; see figure 5.

BEMPA $\hat{B}$ gate. The 3-qubit $\hat{B}$ gate rotates in the $\{|001\rangle, |110\rangle\}$ subspace. Its decomposition follows the same compute-rotate-uncompute pattern as $G^{(4)}$, adapted to three qubits; see figure 6.

4-qubit Givens rotation $G^{(4)}_{ijkl}(\theta)$. The 4-qubit gate rotates in the $\{|1100\rangle, |0011\rangle\}$ subspace. It is essentially the `DoubleExcitation` rotation of `PennyLane`. The strategy is to use a cascade of CNOT gates to map the Hamming distance of the active basis states from four down to one, apply a multi-controlled rotation, and uncompute; see figure 7.

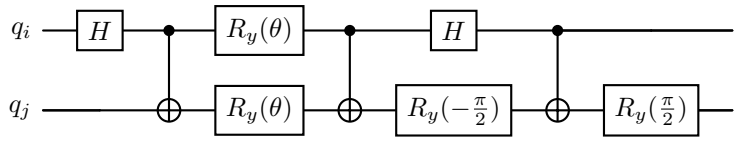


Figure 5. Minimal exact decomposition of the 2-qubit A-gate $A_{ij}^{(2)}(\theta)$ into three CNOT gates. The first two CNOTs implement $G(-\theta)$ as in figure 4 (note the sign change from $R_y(-\theta)$ to $R_y(\theta)$, which reverses the rotation direction), and the third implements the reflection $A(0)$.

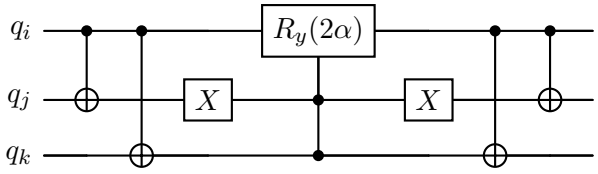


Figure 6. Decomposition of the BEMPA $\hat{B}$ gate. The CNOT cascade controlled by q_i maps $|001\rangle$ and $|110\rangle$ to states differing only in q_i. The X gate on q_j isolates these two states so that only they trigger the $C^2R_y(2\alpha)$ rotation on q_i.

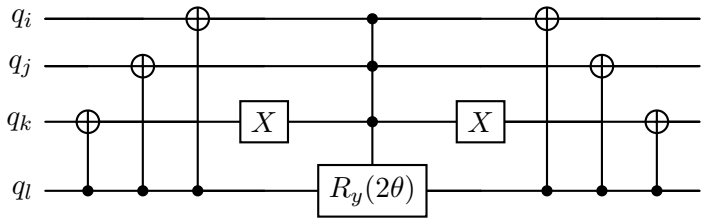


Figure 7. Decomposition of the 4-qubit Givens rotation $G_{ijkl}^{(4)}(\theta)$ acting on the $\{|1100\rangle, |0011\rangle\}$ subspace. The CNOT cascade controlled by q_l maps $|1100\rangle \mapsto |1110\rangle$ and $|0011\rangle \mapsto |1111\rangle$, which differ only in q_l. The X gate on q_k ensures these are the only states with $q_i = q_j = q_k = 1$, so a single $C^3R_y(2\theta)$ implements the rotation. The cascade is then reversed to uncompute. The decomposition of the multi-controlled R_y into elementary gates is given in figure 9.

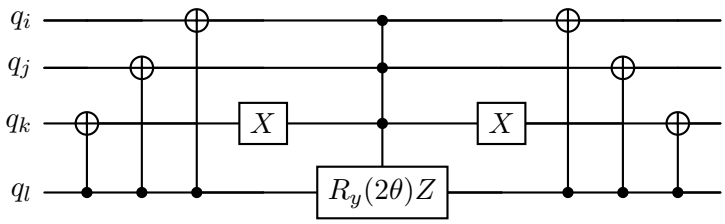


Figure 8. Decomposition of the 4-qubit A-gate $A_{ijkl}^{(4)}(\theta)$. The CNOT cascade is identical to figure 7. The multi-controlled gate on q_l applies $C^3(R_y(2\theta)\,Z)$ combined with a $C^3(-I)$ phase correction to implement the reflection $A^{(4)}(0)$ together with $G^{(4)}(-\theta)$. In practice, the overall $C^3(-I)$ can be absorbed into the decomposition of the multi-controlled gate.

4-qubit A-gate $A_{ijkl}^{(4)}(\theta)$. Since $A^{(4)}(\theta) = A^{(4)}(0)\,G^{(4)}(-\theta)$ and $A^{(4)}(0)$ applies a -1 phase solely to $|1100\rangle$, the circuit reuses the compute-rotate-uncompute structure of $G^{(4)}$ but replaces the multi-controlled rotation. In the mapped basis, the two active states $|1110\rangle$ (from $|1100\rangle$) and $|1111\rangle$ (from $|0011\rangle$) are distinguished by q_l. The combined rotation and reflection require the single-qubit gate $R_y(2\theta)\,Z$ on q_l (which acts as $\left[\begin{smallmatrix}\cos\theta & \sin\theta \\ \sin\theta & -\cos\theta\end{smallmatrix}\right]$), with an additional $C^3(-I)$ phase to account for the $A^{(4)}(0)$ reflection; see figure 8.

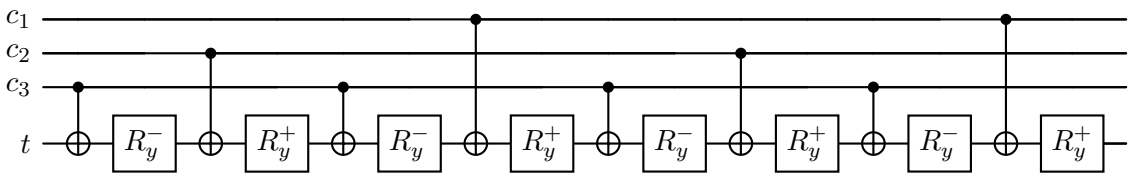


Figure 9. Decomposition of $C^3R_y(\alpha)$ into 8 CNOT gates and 8 single-qubit rotations $R_y^{\pm} \equiv R_y(\pm\alpha/8)$ of alternating sign. The CNOT controls follow the Gray code pattern $c_3, c_2, c_3, c_1, c_3, c_2, c_3, c_1$, where c_3 toggles most frequently. The same pattern with two controls and 4 CNOTs gives $C^2R_y(\alpha)$ with $R_y(\pm\alpha/4)$.

Gate	Qubits	CNOTs	Depth
$G_{ij}^{(2)}(\theta)$	2	2	5
$A_{ij}^{(2)}(\theta)$	2	3	7
$\hat{B}(\alpha)$	3	8	12
$G_{ijkl}^{(4)}(\theta)$	4	14	22
$A_{ijkl}^{(4)}(\theta)$	4	14	22

Table 4. Total circuit resources for each gate decomposition after expanding multi-controlled rotations into elementary gates. The CNOT counts include both the CNOT cascade and the decomposed C^mR_y gate (2^m CNOTs). The depth accounts for parallelisation of gates on disjoint qubits.

Decomposition of multi-controlled rotations. The multi-controlled $C^mR_y(\alpha)$ gates appearing in the circuits above can be decomposed into 2^m CNOT gates and 2^m single-qubit R_y rotations using the uniformly controlled rotation technique [64]. The construction interleaves CNOT gates — whose controls cycle through the m control qubits in Gray code order — with $R_y(\pm\alpha/2^m)$ rotations of alternating sign on the target qubit; see figure 9. More generally, the same technique applies to any multi-controlled single-qubit gate $C^m(U)$ with arbitrary $U \in U(2)$, using 2^m CNOTs with general $e^{i\alpha}R_z \cdot R_y \cdot R_z$ blocks between them instead of bare R_y rotations [64]. For $A^{(4)}$, the combined gate $C^3(-R_y(2\theta)\,Z)$, which absorbs the $C^3(-I)$ phase correction of figure 8, has $U = -R_y(2\theta)\,Z \notin SU(2)$ (since $\det U = -1$), but is decomposed with the same $2^3 = 8$ CNOT count; the non-unit determinant is accommodated by a global phase factor $e^{i\alpha}$ in the single-qubit blocks between the CNOTs.

Summary of circuit resources. The *circuit depth* is the number of sequential time steps when gates acting on disjoint qubits are executed in parallel. Table 4 summarises the total resource counts after decomposing all multi-controlled gates as in figure 9.

Cost of the circuit of figure 1. Assembling the 7 $G^{(2)}$ and 12 $G^{(4)}$ gates of figure 1 from the decompositions above gives 182 CNOT gates in total, with a CNOT-depth of 165 when gates on disjoint qubits are executed in parallel. The count is dominated by the $G^{(4)}$ gates, at 14 CNOTs each; these are unavoidable here, since by the argument of section 4 no 2-qubit gate can change orbital occupancies while preserving S_z. We note that the 14-CNOT decomposition of $G^{(4)}$ used here follows the transparent uniformly-controlled-rotation route of figure 9 and is not optimised. Throughout this work "hardware-efficient" should be read as parameter-economical while remaining exactly symmetry-preserving, rather than as depth-minimal.

Funding Statement. This paper has been funded by the following grants/awards. Royal Society: `URF\R1\211417`; Science and Technology Facilities Council: `ST/X000753/1`.

Data Availability Statement. My article has no associated data or the data will not be deposited.

Code Availability Statement. Code available at https://github.com/andstergiou/fuzzy-ising-vqe.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License (CC-BY4.0), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

[1] A. Peruzzo et al., *A variational eigenvalue solver on a photonic quantum processor*, *Nature Commun.* **5** (2014) 4213 [arXiv:1304.3061] [INSPIRE].

[2] J. Tilly et al., *The Variational Quantum Eigensolver: A review of methods and best practices*, *Phys. Rept.* **986** (2022) 1 [arXiv:2111.05176] [INSPIRE].

[3] A.Y. Kitaev, *Quantum computations: algorithms and error correction*, *Russ. Math. Surv.* **52** (1997) 1191.

[4] C.M. Dawson and M.A. Nielsen, *The Solovay-Kitaev algorithm*, *Quant. Inf. Comput.* **6** (2006) 081 [quant-ph/0505030] [INSPIRE].

[5] V. Jurdjevic and H.J. Sussmann, *Control systems on Lie groups*, *J. Diff. Eq.* **12** (1972) 313.

[6] D.P. DiVincenzo, *Two-bit gates are universal for quantum computation*, *Phys. Rev. A* **51** (1995) 1015 [cond-mat/9407022] [INSPIRE].

[7] S. Lloyd, *Almost Any Quantum Logic Gate is Universal*, *Phys. Rev. Lett.* **75** (1995) 346 [INSPIRE].

[8] A. Barenco, *A Universal two bit gate for quantum computation*, *Proc. Roy. Soc. Lond. A* **449** (1995) 679 [quant-ph/9505016] [INSPIRE].

[9] D. Deutsch, A. Barenco and A. Ekert, *Universality in quantum computation*, *Proc. Roy. Soc. Lond. A* **449** (1995) 669 [quant-ph/9505018] [INSPIRE].

[10] D. Wecker et al., *Solving strongly correlated electron models on a quantum computer*, *Phys. Rev. A* **92** (2015) 062318 [arXiv:1506.05135] [INSPIRE].

[11] Z. Jiang et al., *Quantum algorithms to simulate many-body physics of correlated fermions*, *Phys. Rev. Applied* **9** (2018) 044036 [arXiv:1711.05395] [INSPIRE].

[12] C. Cade, L. Mineh, A. Montanaro and S. Stanisic, *Strategies for solving the Fermi-Hubbard model on near-term quantum computers*, *Phys. Rev. B* **102** (2020) 235122 [arXiv:1912.06007] [INSPIRE].

[13] S. Bahrami and N. Sawaya, *Particle-conserving quantum circuit ansatz with applications in variational simulation of bosonic systems*, arXiv:2402.18768 [INSPIRE].

[14] W. Zhu et al., *Uncovering Conformal Symmetry in the 3D Ising Transition: State-Operator Correspondence from a Quantum Fuzzy Sphere Regularization*, *Phys. Rev. X* **13** (2023) 021009 [arXiv:2210.13482] [INSPIRE].

[15] O. Higgott, D. Wang and S. Brierley, *Variational Quantum Computation of Excited States*, *Quantum* **3** (2019) 156 [arXiv:1805.08138] [INSPIRE].

[16] K. Klymko et al., *Real-Time Evolution for Ultracompact Hamiltonian Eigenstates on Quantum Hardware*, *PRX Quantum* **3** (2022) 020323 [arXiv:2103.08563] [INSPIRE].

[17] D. Poland, S. Rychkov and A. Vichi, *The Conformal Bootstrap: Theory, Numerical Techniques, and Applications*, *Rev. Mod. Phys.* **91** (2019) 015002 [arXiv:1805.04405] [INSPIRE].

[18] J.M. Arrazola et al., *Universal quantum circuits for quantum chemistry*, *Quantum* **6** (2022) 742 [arXiv:2106.13839] [INSPIRE].

[19] H. Yamabe, *On an arcwise connected subgroup of a Lie group*, *Osaka Math. J.* **2** (1950) 13.

[20] P.K. Barkoutsos et al., *Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions*, *Phys. Rev. A* **98** (2018) 022322 [arXiv:1805.04340] [INSPIRE].

[21] B.T. Gard et al., *Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm*, *npj Quantum Inf.* **6** (2020) 10 [arXiv:1904.10910] [INSPIRE].

[22] G. Yan et al., *Universal 2-Local Symmetry-Preserving Quantum Neural Networks for Fermionic Systems*, arXiv:2412.04825 [INSPIRE].

[23] J. Kempe, A. Kitaev and O. Regev, *The Complexity of the Local Hamiltonian Problem*, *SIAM J. Comput.* **35** (2006) 1070 [quant-ph/0406180] [INSPIRE].

[24] R. Somma, G. Ortiz, E. Knill and J. Gubernatis, *Quantum Simulations of Physics Problems*, *Int. J. Quant. Inf.* **01** (2003) 417 [quant-ph/0304063] [INSPIRE].

[25] S. McArdle et al., *Digital quantum simulation of molecular vibrations*, *Chem. Sci.* **10** (2019) 5725 [arXiv:1811.04069] [INSPIRE].

[26] P.J. Ollitrault, A. Baiardi, M. Reiher and I. Tavernelli, *Hardware efficient quantum algorithms for vibrational structure calculations*, *Chem. Sci.* **11** (2020) 6842 [arXiv:2003.12578] [INSPIRE].

[27] N.P.D. Sawaya, G.G. Guerreschi and A. Holmes, *On connectivity-dependent resource requirements for digital quantum simulation of d-level particles*, in the proceedings of the *2020 IEEE International Conference on Quantum Computing and Engineering*, (2020) [DOI:10.1109/QCE49297.2020.00031] [arXiv:2005.13070] [INSPIRE].

[28] H. Jnane et al., *Analog quantum simulation of non-Condon effects in molecular spectroscopy*, *ACS Photonics* **8** (2021) 7 [arXiv:2011.05553] [INSPIRE].

[29] A.T. Schmitz, N.P.D. Sawaya, S. Johri and A.Y. Matsuura, *Graph optimization perspective for low-depth Trotter-Suzuki decomposition*, *Phys. Rev. A* **109** (2024) 042418 [arXiv:2103.08602] [INSPIRE].

[30] D. Trenev et al., *Refining resource estimation for the quantum computation of vibrational molecular spectra through Trotter error analysis*, *Quantum* **9** (2025) 1630 [arXiv:2311.03719] [INSPIRE].

[31] F.D.M. Haldane, *Fractional quantization of the Hall effect: A hierarchy of incompressible quantum fluid states*, *Phys. Rev. Lett.* **51** (1983) 605 [INSPIRE].

[32] T.T. Wu and C.N. Yang, *Dirac Monopole Without Strings: Monopole Harmonics*, *Nucl. Phys. B* **107** (1976) 365 [INSPIRE].

[33] E.T. Newman and R. Penrose, *Note on the Bondi-Metzner-Sachs group*, *J. Math. Phys.* **7** (1966) 863 [INSPIRE].

[34] T. Dray, *The Relationship Between Monopole Harmonics and Spin Weighted Spherical Harmonics*, *J. Math. Phys.* **26** (1985) 1030 [INSPIRE].

[35] J. Madore, *The Fuzzy sphere*, *Class. Quant. Grav.* **9** (1992) 69 [INSPIRE].

[36] K. Hasebe, *Hopf Maps, Lowest Landau Level, and Fuzzy Spheres*, *SIGMA* **6** (2010) 071 [arXiv:1009.1192] [INSPIRE].

[37] P. Di Francesco, P. Mathieu and D. Senechal, *Conformal Field Theory*, Springer-Verlag, New York (1997) [DOI:10.1007/978-1-4612-2256-9] [INSPIRE].

[38] L. Hu, Y.-C. He and W. Zhu, *Operator Product Expansion Coefficients of the 3D Ising Criticality via Quantum Fuzzy Spheres*, *Phys. Rev. Lett.* **131** (2023) 031601 [arXiv:2303.08844] [INSPIRE].

[39] C. Han, L. Hu, W. Zhu and Y.-C. He, *Conformal four-point correlators of the three-dimensional Ising transition via the quantum fuzzy sphere*, *Phys. Rev. B* **108** (2023) 235123 [arXiv:2306.04681] [INSPIRE].

[40] J.R. McClean et al., *OpenFermion: the electronic structure package for quantum computers*, *Quantum Sci. Technol.* **5** (2020) 034014 [arXiv:1710.07629] [INSPIRE].

[41] OEIS Foundation Inc., *A125809 entry, The On-Line Encyclopedia of Integer Sequences*, https://oeis.org/search?q=A125809.

[42] D. Simmons-Duffin, *The Lightcone Bootstrap and the Spectrum of the 3d Ising CFT*, *JHEP* **03** (2017) 086 [arXiv:1612.08471] [INSPIRE].

[43] C. Han, L. Hu and W. Zhu, *Conformal operator content of the Wilson-Fisher transition on fuzzy sphere bilayers*, *Phys. Rev. B* **110** (2024) 115113 [arXiv:2312.04047] [INSPIRE].

[44] A. Dey, L. Herviou, C. Mudry and A.M. Läuchli, *Conformal Data for the O(3) Wilson-Fisher Conformal Field Theory from Fuzzy Sphere Realization of the Quantum Rotor Model*, *Phys. Rev. Lett.* **136** (2026) 206504 [arXiv:2510.09755] [INSPIRE].

[45] W. Guo, Z. Zhou, T.-C. Wei and Y.-C. He, *O(N) free-scalar and Wilson-Fisher conformal field theories on the fuzzy sphere*, *Phys. Rev. B* **114** (2026) 065102 [arXiv:2512.02234] [INSPIRE].

[46] Z. Zhou, L. Hu, W. Zhu and Y.-C. He, *SO(5) Deconfined Phase Transition under the Fuzzy-Sphere Microscope: Approximate Conformal Symmetry, Pseudo-Criticality, and Operator Spectrum*, *Phys. Rev. X* **14** (2024) 021044 [arXiv:2306.16435] [INSPIRE].

[47] Z. Zhou and Y.-C. He, *3D Conformal Field Theories with Sp(N) Global Symmetry on a Fuzzy Sphere*, *Phys. Rev. Lett.* **135** (2025) 026504 [arXiv:2410.00087] [INSPIRE].

[48] R. Fan, J. Dong and A. Vishwanath, *Simulating the non-unitary Yang-Lee conformal field theory on the fuzzy sphere*, *JHEP* **09** (2026) 041 [arXiv:2505.06342] [INSPIRE].

[49] E. Arguello Cruz, I.R. Klebanov, G. Tarnopolsky and Y. Xin, *Yang-Lee Quantum Criticality in Various Dimensions*, *Phys. Rev. X* **16** (2026) 011022 [arXiv:2505.06369] [INSPIRE].

[50] J. Elias Miró and O. Delouche, *Flowing from the Ising model on the fuzzy sphere to the 3D Lee-Yang CFT*, *JHEP* **10** (2025) 037 [arXiv:2505.07655] [INSPIRE].

[51] Y.-C. He, *Free real scalar CFT on fuzzy sphere: spectrum, algebra and wavefunction ansatz*, arXiv:2506.14904 [INSPIRE].

[52] J. Taylor, C. Voinea, Z. Papić and R. Fan, *Conformal Scalar Field Theory from Ising Tricriticality on the Fuzzy Sphere*, *Phys. Rev. Lett.* **136** (2026) 056503 [arXiv:2506.22539] [INSPIRE].

[53] Z. Zhou, C. Wang and Y.-C. He, *Chern-Simons-Matter Conformal Field Theory on the Fuzzy Sphere: Higgs Transition of the Kalmeyer-Laughlin Chiral Spin Liquid*, *Phys. Rev. Lett.* **137** (2026) 106504 [arXiv:2507.19580] [INSPIRE].

[54] C. Voinea, W. Zhu, N. Regnault and Z. Papić, *Critical Majorana Fermion at a Topological Quantum Hall Bilayer Transition*, *Phys. Rev. Lett.* **136** (2026) 076601 [arXiv:2509.08036] [INSPIRE].

[55] Y. Tang et al., *Emergence of 3D Superconformal Ising Criticality on the Fuzzy Sphere*, *Phys. Rev. Lett.* **137** (2026) 026502 [arXiv:2512.25054] [INSPIRE].

[56] E. Huffman, Z. Zhou, Y.-C. He and J.S. Hofmann, *Generalizing deconfined criticality to 3D N-flavor SU(2) quantum chromodynamics on the fuzzy sphere*, *Phys. Rev. D* **114** (2026) 034503 [arXiv:2602.11255] [INSPIRE].

[57] A. Dey et al., *Conformal Data for the $O(2)$ Wilson-Fisher CFT in $(2+1)$-Dimensional Spacetime from Exact Diagonalization and Matrix Product States on the Fuzzy Sphere*, arXiv:2604.18705 [INSPIRE].

[58] A. Stergiou, *Quantum rotors on the fuzzy sphere and the cubic conformal field theory*, *Phys. Rev. B* **114** (2026) 225103 [arXiv:2604.24840] [INSPIRE].

[59] A. Stergiou, *Fuzzy Sphere Ising VQE*, Zenoto (2026), DOI:10.5281/zenodo.20343458.

[60] V. Bergholm et al., *PennyLane: Automatic differentiation of hybrid quantum-classical computations*, arXiv:1811.04968 [INSPIRE].

[61] I. Miháliková, J. Carlson, D. Neill and I. Stetcu, *State preparation and symmetries*, arXiv:2510.06702 [INSPIRE].

[62] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, in the proceedings of the *International Conference on Learning Representations*, San Diego, U.S.A., 07–09 May 2015 [arXiv:1412.6980] [INSPIRE].

[63] F. Kos, D. Poland, D. Simmons-Duffin and A. Vichi, *Precision Islands in the Ising and $O(N)$ Models*, *JHEP* **08** (2016) 036 [arXiv:1603.04436] [INSPIRE].

[64] M. Möttönen, J.J. Vartiainen, V. Bergholm and M.M. Salomaa, *Quantum Circuits for General Multiqubit Gates*, *Phys. Rev. Lett.* **93** (2004) 130502 [quant-ph/0404089] [INSPIRE].

[65] Andreas Stergiou, *fuzzy-ising-vqe*, https://github.com/andstergiou/fuzzy-ising-vqe.