University of Trento
Department of Physics
PhD Thesis

Beyond Qubits: Quantum Optimization and Lattice Gauge Theories in a qudit framework

Supervisors
Prof. Philipp Hans Juergen Hauke
Dr. Rer. Nat. Sebastian Schmitt

PhD Student
Alberto Bottarelli

Acknowledgments

Abstract

Quantum computing holds the promise of tackling problems that are intractable for classical machines, yet the path from current noisy intermediate-scale quantum (NISQ) devices to fault-tolerant processors demands both better algorithms and a deeper understanding of how to exploit the full structure of quantum hardware. A key direction in this effort is the move *beyond qubits*: exploiting higher-dimensional local Hilbert spaces — *qudits* — to encode information, enforce constraints, and simulate complex physical systems more naturally and efficiently.

This thesis investigates two interconnected research fronts in which the qudit framework provides concrete benefits: *quantum optimization* and the *quantum simulation of lattice gauge theories* (LGTs). On the optimization side, we show how a single ancilla qudit can enforce inequality constraints in the Quantum Approximate Optimization Algorithm (QAOA) without inflating the solution space, and how counterdiabatic driving, enhanced by symmetry reduction, improves variational performance for qudit-based algorithms. We additionally study the noise robustness of variational algorithms on real hardware, and investigate how nonstabilizerness builds up and relates to optimization performance across both qubit and qutrit protocols. On the simulation side, we analyze the ground-state structure and phase transitions of non-Abelian lattice gauge theories with dihedral symmetry, characterizing confinement and string breaking in D_N models, and probing the quantum complexity of ground states across $\mathbb{Z}_N$, D_3, and SU(2) gauge theories.

Together, these results clarify where and how the qudit framework provides algorithmic and representational advantages over the standard qubit paradigm, and contribute to the broader effort of developing practical quantum simulation and optimization tools for near-term hardware.

Contents

Introduction

Starting from Feynman's revolutionary idea [1], quantum mechanics has introduced a radically different computational paradigm that has drastically evolved in the last two decades.

Since the early proposals by Feynman and Deutsch [2] and the discovery of exponentially fast algorithms for structured problems [3, 4], the prospect of a *quantum advantage* — the ability of quantum processors to solve problems that are intractable for any classical computer — has driven an entire field of research. Today, programmable quantum devices based on superconducting circuits [5], trapped ions [6, 7], Rydberg atoms [8], and photonic platforms [9] have demonstrated coherent control over tens to hundreds of quantum degrees of freedom, placing us in the noisy intermediate-scale quantum (NISQ) era [10]. In this regime, hardware is too noisy for deep fault-tolerant circuits yet too large for exact classical simulation, creating a unique opportunity to explore quantum algorithms on real devices and to identify which features of quantum mechanics are most valuable for practical computation.

The standard unit of quantum information — the qubit — is by definition a two-level system. However, many physical platforms natively provide access to local Hilbert spaces of dimension $d > 2$. Trapped-ion processors, for instance, can coherently control up to eight internal levels of a single ion, implementing universal single- and two-qudit gates with high fidelity [6, 11]. Extending the computational framework from qubits to *qudits* is not merely a matter of hardware capability: it provides genuine algorithmic advantages. A d-dimensional local degree of freedom can encode more information per physical subsystem, naturally represent finite groups without truncation overhead, and implement certain constraint structures without expanding the solution space. These features make qudits particularly attractive for two application domains that form the core of this thesis: *quantum optimization* and the *quantum simulation of lattice gauge theories*.

Quantum optimization aims to use quantum hardware to solve combinatorial problems that are hard for classical computers [12, 13, 14]. The Quantum Approximate Optimization Algorithm (QAOA) [12] is a leading candidate for near-term applications: it prepares a parametrised quantum state by alternating between a cost Hamiltonian and a mixing Hamiltonian, and a classical outer loop tunes the variational parameters to minimise the expected cost. Despite its appeal, QAOA faces significant practical challenges: real-world problems are often subject to inequality constraints that are costly to encode in the standard binary framework; the onset of Barren Plateaus (BPs) — the exponential vanishing of cost-function gradients with system size [15, 16] — can make classical optimization of the variational parameters intractable; and the impact of hardware noise on the trainability of these algorithms remains an active area of investigation.

Lattice gauge theories (LGTs) are among the most important models in theoretical physics, underpinning our understanding of the fundamental forces and of strongly correlated quantum matter, from quantum chromodynamics (QCD) to topological phases [17, 18]. Classical simulation of non-Abelian LGTs is severely limited by the sign problem and the

exponential growth of the Hilbert space, making quantum simulation a highly sought-after alternative [19, 20, 21, 22]. Qudits are natural building blocks for LGT simulators: the gauge degrees of freedom living on lattice links transform under finite-dimensional representations of group algebras, and qudits can encode these representations more compactly than binary encodings [23]. Beyond the encoding question, understanding the ground-state structure of discrete non-Abelian gauge theories — confinement, string breaking, phase transitions — is a problem of intrinsic physical interest and a necessary step toward quantum simulations of more complex theories such as QCD.

Together, these two directions illustrate how expanding the native dimensionality of quantum hardware can simultaneously broaden the class of tractable problems and deepen our physical understanding of complex quantum systems, placing qudits at the frontier of both near-term and long-term quantum computing research

This thesis is organised as follows.

Ch.1 provides the theoretical background. It introduces the qubit and qudit formalisms, including generalized Pauli operators, the Clifford group, and trapped-ion qudit hardware. It then reviews classical and quantum optimization, covering variational quantum eigensolvers, QAOA, and the problem of Barren Plateaus. Finally, it introduces lattice gauge theories, from the Kogut–Susskind Hamiltonian and its symmetries to the specific models studied in the rest of the thesis.

Ch.2 addresses the problem of enforcing inequality constraints in QAOA using qudits, based on [24]. Two complementary strategies are developed and benchmarked: a qudit generalization of slack variables, and a direct penalty approach that uses a single ancilla qudit to imprint phases on infeasible states without enlarging the solution space. The direct penalty approach is shown to be particularly advantageous when many constraints are present simultaneously, as it keeps the fraction of feasible solutions constant regardless of the number of constraints.

Ch.3 develops a symmetry-enhanced counterdiabatic approach for qudit QAOA, based on [25]. Counterdiabatic driving suppresses non-adiabatic transitions during the variational evolution; exploiting the symmetries of the problem reduces the parameter count of the variational ansatz from $\mathcal{O}(N^2)$ to a problem-specific symmetry-constrained subspace, yielding improved convergence on both qubit and qudit instances.

Ch.4 presents results from collaborations in which the author contributed in a supporting role. The first part, based on [26], provides an experimental demonstration on IBM superconducting processors that noise-induced barren plateaus are absent under amplitude-damping (T_1) noise, confirming theoretical predictions that non-unital noise preserves a finite gradient signal even at large system sizes. The second part, based on [27], investigates the internal structure of QAOA runs on the Sherrington–Kirkpatrick spin-glass model, finding a universal transient build-up of nonstabilizerness — a *magic barrier* — that scales with circuit depth and correlates with the algorithm's success fidelity.

Ch.5 reports preliminary results on a D_3 lattice gauge theory in 1+1 dimensions, analyzing the ground-state phase structure and the interplay between the electric-field energy, the hopping term, and gauge invariance. The D_3 model constitutes the minimal non-Abelian LGT in which string breaking can be observed without truncating the gauge group, making it an ideal testbed for non-Abelian confinement dynamics.

Ch.6 presents results from two further collaborations. The first, based on [28], studies confinement in pure D_N lattice gauge theories, showing that the structure of the group center determines whether static charges are confined or screened: for even N, the non-trivial $\mathbb{Z}_2$ center produces stable flux strings and confinement; for odd N, the trivial center allows string breaking even without dynamical matter. The second, based on [29], probes the quantum

complexity of ground states by comparing nonstabilizerness, multipartite entanglement, and fermionic non-Gaussianity across $\mathbb{Z}_N$, D_3, and SU(2) gauge theories on a plaquette ladder, finding that the simulation complexity depends non-trivially on the interplay between group structure, superselection sector, and encoding strategy.

Except Ch 5, all the results are based on the following published works

- Bottarelli, Alberto, Sebastian Schmitt, and Philipp Hauke. "Inequality constraints in variational quantum circuits with qudits." Physical Review Research 7.3 (2025): 033202.https://journals.aps.org/prresearch/abstract/10.1103/3l96-41xf
- Santra, Gopal Chandra, et al. "Quantum Resources in Non-Abelian Lattice Gauge Theories: Nonstabilizerness, Multipartite Entanglement, and Fermionic Non-Gaussianity." arXiv preprint arXiv:2510.07385 (2025). https://arxiv.org/abs/2510.07385
- Bottarelli, Alberto, et al. "Symmetry-enhanced counterdiabatic quantum algorithm for qudits." Physical Review Research 7.4 (2025): 043030. https://journals.aps.org/prresearch/abstract/10.1103/6ldg-3w1f
- Capecci, Chiara, et al. "Role of Nonstabilizerness in Quantum Optimization." arXiv preprint arXiv:2505.17185 (2025). https://arxiv.org/abs/2505.17185
- Popov, Pavel P., et al. "When the center matters: color screening and gluelumps in dihedral lattice gauge theories." arXiv preprint arXiv:2512.00152 (2025). https://arxiv.org/abs/2512.00152
- Schmitt, Sebastian, et al. "Experimental demonstration of the absence of noise-induced barren plateaus using information content landscape analysis." arXiv preprint arXiv:2602.22851 (2026). https://arxiv.org/abs/2602.22851

Chapter 1

Motivation and background

1.1 Quantum information with qubits and qudits

In this section, we briefly review the qubit and qudit formalism that underlies the rest of this thesis. The treatment is intentionally self-contained and introductory, with emphasis on the concepts most directly relevant for variational quantum algorithms and their extension to constrained optimisation and lattice gauge theories.

1.1.1 Qubit States and the Bloch Sphere

The fundamental unit of quantum information [30] is the *qubit*: a two-level quantum system whose state is a normalised vector in a two-dimensional complex Hilbert space $\mathcal{H} \simeq \mathbb{C}^2$. A convenient computational basis is provided by the orthonormal pair

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \tag{1.1}$$

An arbitrary pure qubit state is a linear superposition

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle\,, \qquad \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1. \tag{1.2}$$

Because a global phase carries no physical content, the state space is effectively two-dimensional over the reals. A standard choice of coordinates is

$$|\psi(\theta,\phi)\rangle = \cos\!\left(\tfrac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\!\left(\tfrac{\theta}{2}\right)|1\rangle\,, \qquad \theta \in [0,\pi], \quad \phi \in [0, 2\pi), \tag{1.3}$$

which admits a clean geometric interpretation. Each pure state of Eq. (1.3) corresponds uniquely to a point on the surface of the unit sphere in $\mathbb{R}^3$ via the Cartesian coordinates

$$\vec{r} = \big(\sin\theta\cos\phi,\ \sin\theta\sin\phi,\ \cos\theta\big). \tag{1.4}$$

This object is the *Bloch sphere*, shown in Figure 1.1. The north and south poles correspond to $|0\rangle$ and $|1\rangle$, respectively, while the equator contains the equal-weight superpositions. Mixed states of a single qubit are described by density matrices ρ (introduced below) and correspond to points *inside* the sphere; the completely mixed state $\rho = \mathbb{I}/2$ lies at the origin.

When a physical observable is measured on a qubit in state $|\psi\rangle$, the outcome is governed by the *Born rule*: a measurement in the computational basis yields outcome 0 with probability $p_0 = |\alpha|^2 = |\,\langle 0|\psi\rangle\,|^2$ and outcome 1 with probability $p_1 = |\beta|^2 = |\,\langle 1|\psi\rangle\,|^2$. More

generally, a quantum observable is represented by a Hermitian operator $\hat{O} = \hat{O}^\dagger$ acting on $\mathcal{H}$. Its *expectation value* in the state $|\psi\rangle$ is

$$\langle \hat{O} \rangle = \langle \psi | \hat{O} | \psi \rangle = \sum_k o_k \, | \langle o_k | \psi \rangle |^2 \,, \tag{1.5}$$

where o_k and $|o_k\rangle$ are the eigenvalues and eigenstates of $\hat{O}$, respectively, and the second equality follows directly from the Born rule applied to each projector $|o_k\rangle\langle o_k|$. The expectation value represents the average outcome of many identical measurements, and is the primary quantity accessed in quantum algorithms and quantum simulations.

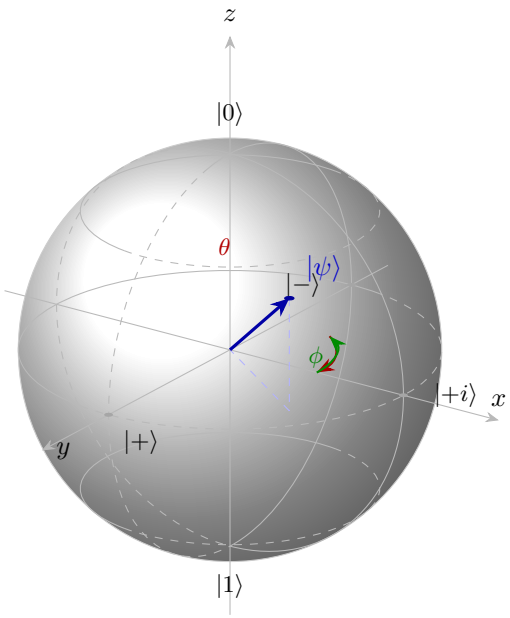

Figure 1.1: The Bloch sphere. Every pure single-qubit state $|\psi(\theta, \phi)\rangle$ of Eq. (1.3) corresponds to a point on the unit sphere parameterized by polar angle $\theta \in [0, \pi]$ (measured from $|0\rangle$ at the north pole) and azimuthal angle $\phi \in [0, 2\pi)$. The computational basis states $|0\rangle$ and $|1\rangle$ sit at the poles; the X-eigenstates $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$ and the Y-eigenstate $|+i\rangle = (|0\rangle + i\,|1\rangle)/\sqrt{2}$ lie on the equator. Mixed states correspond to points *inside* the sphere.

1.1.2 Density Matrices and Mixed States

The state-vector description is sufficient for isolated pure states, but a more general framework is needed when dealing with statistical mixtures or subsystems of a larger entangled system. This is provided by the *density matrix* formalism.

A density matrix is a positive-semidefinite, unit-trace, Hermitian operator of the form

$$\rho = \sum_i p_i \, |\psi_i\rangle \, \langle \psi_i| \,, \qquad p_i \geq 0, \quad \sum_i p_i = 1, \tag{1.6}$$

where $\{p_i\}$ are classical probabilities and $\{|\psi_i\rangle\}$ are pure states. When the ensemble contains only a single term, $\rho = |\psi\rangle \langle \psi|$, the state is *pure*; otherwise it is *mixed*.

For any operator ρ to represent a valid physical state it must satisfy three conditions:

- **Hermiticity:** $\rho^\dagger = \rho$,
- **Positivity:** $\langle\phi|\,\rho\,|\phi\rangle \geq 0$ for all $|\phi\rangle \in \mathcal{H}$,
- **Unit trace:** $\mathrm{Tr}(\rho) = 1$.

The expectation value of any observable O is then compactly written as

$$\langle O\rangle_\rho = \mathrm{Tr}[\rho\, O], \tag{1.7}$$

which reproduces the standard Born-rule result for pure states and defines the measurement postulate for general mixed states.

1.1.3 Composite Systems and the Partial Trace

When two subsystems A and B are combined, their joint Hilbert space is the tensor product

$$\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B. \tag{1.8}$$

The joint state is described by a density matrix ρ_{AB} acting on $\mathcal{H}_{AB}$. The *reduced state* of subsystem A, containing all information accessible to measurements on A alone, is obtained by *tracing out* subsystem B:

$$\rho_A = \mathrm{Tr}_B(\rho_{AB}). \tag{1.9}$$

This operation is uniquely characterised by the requirement that

$$\mathrm{Tr}(\rho_A\, O) = \mathrm{Tr}(\rho_{AB}\,(O \otimes \mathbb{I}_B)) \quad \forall\, O \text{ acting on } \mathcal{H}_A. \tag{1.10}$$

Geometrically, taking the partial trace corresponds to marginalising over the degrees of freedom of B. If ρ_{AB} describes an entangled state, ρ_A will in general be a *mixed* state even when the global state is pure—a hallmark of quantum correlations with no classical analogue.

1.1.4 Operators on Qubits

Physical operations on qubits are represented by linear operators on $\mathcal{H}$. Of particular importance are *unitary* operators, which describe reversible quantum evolution. A central role is played by the three Pauli matrices

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{1.11}$$

which generate rotations of the Bloch vector about the corresponding Cartesian axes via

$$R_\alpha(\theta) = e^{-i\frac{\theta}{2}\sigma_\alpha} = \cos\tfrac{\theta}{2}\,\mathbb{I} - i\sin\tfrac{\theta}{2}\,\sigma_\alpha, \qquad \alpha \in \{x, y, z\}. \tag{1.12}$$

Any single-qubit unitary $U \in \mathrm{SU}(2)$ can be decomposed as [30]

$$U = e^{i\alpha}\, R_z(\beta)\, R_y(\theta)\, R_z(\delta), \tag{1.13}$$

for real parameters $\alpha, \beta, \theta, \delta$.

For an N-qubit system, the *N-qubit Pauli group* is generated by all tensor products of single-qubit Paulis,

$$\mathcal{P}_N = \Big\{ \bigotimes_{i=0}^{N-1} P_i \;\Big|\; P_i \in \{I, X, Y, Z\} \Big\}. \tag{1.14}$$

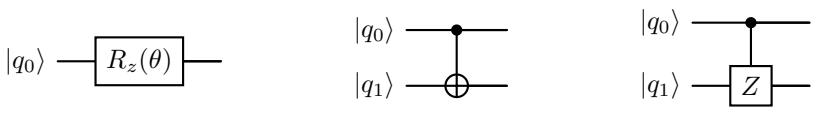

Figure 1.2: Representative quantum gates in circuit notation. *Left:* single-qubit Z-rotation $R_z(\theta)$. *Centre:* controlled-NOT (CNOT) gate. *Right:* controlled-Z (CZ) gate.

Pauli strings from $\mathcal{P}_N$ form a complete operator basis for N-qubit Hamiltonians and appear directly as the elementary observables measured in variational quantum algorithms. In particular, cost functions for the constrained optimisation problems studied in this thesis are expressed as linear combinations of Pauli-string expectation values. Hamiltonians of lattice gauge theories, once mapped to finite-dimensional local Hilbert spaces, likewise decompose naturally into Pauli strings or their higher-dimensional (qudit) analogues.

For multi-qubit systems, local operators are constructed by tensoring single-qubit operators with identities on the remaining qubits. Entangling operations, which act non-trivially on more than one qubit, are essential for generating the non-classical correlations required for quantum advantage.

1.1.5 Quantum Circuit Representation

Quantum algorithms are most naturally described in the *quantum circuit model*. A computation begins in the fiducial state $|0\rangle^{\otimes N}$ and proceeds as a sequence of unitary gates whose ordered product implements the desired transformation.

Graphically, circuits are drawn as a set of horizontal wires representing qubits, with time flowing left to right. Single-qubit gates appear as labelled boxes on individual wires; multi-qubit gates, such as controlled operations, are connected across multiple wires. Three standard examples are shown in Figure 1.2, with matrix representations

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \qquad CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \qquad i\text{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{1.15}$$

The circuit model provides a natural language for quantifying resource requirements such as circuit depth and gate count—both critical figures of merit in the noisy intermediate-scale quantum (NISQ) regime. These considerations are directly relevant for the variational algorithms and lattice gauge theory simulations developed in the following chapters, where the circuit architecture affects both physical fidelity and computational cost.

1.1.6 Qudits: Definition and Gatesets

While qubits are the most common platform for quantum information processing, many physical systems naturally provide access to local Hilbert spaces of dimension $d > 2$. Such systems are called *qudits*, and they play a central role in this thesis: both for formulating lattice gauge theories with finite-dimensional local degrees of freedom [**qudits**, 23], and for designing variational quantum algorithms with enhanced expressive power [31, 32, 33].

A qudit is a quantum system with a d-dimensional Hilbert space $\mathcal{H}_d \simeq \mathbb{C}^d$, spanned by the orthonormal computational basis $\{|0\rangle, |1\rangle, \ldots, |d-1\rangle\}$. An arbitrary pure state takes

the form

$$|\psi\rangle = \sum_{k=0}^{d-1} \alpha_k |k\rangle , \qquad \alpha_k \in \mathbb{C}, \quad \sum_{k=0}^{d-1} |\alpha_k|^2 = 1. \tag{1.16}$$

Relative to qubits, the local state space grows rapidly with d, enabling more compact encodings of information and symmetries. From an algorithmic perspective, enlarging the local Hilbert space can reduce the number of physical subsystems required to represent a given problem, potentially lowering circuit depth and relaxing hardware constraints. In lattice gauge theories, qudits arise naturally when truncating infinite-dimensional gauge degrees of freedom to finite-dimensional representations while preserving the essential algebraic structure.

For N qudits, the joint Hilbert space is $\mathcal{H}_d^{\otimes N}$, with dimension d^N. As in the qubit case, entanglement between subsystems is a key computational resource, generated through multi-qudit interactions.

In the recent years, several research groups have put their effort into realizing qudit-based universal quantum computing with various technologies [34, 34, 35, 11]. Sec. 1.1.7 briefly describes trapped ion qudits, but other realizations include for example optical qudits [34], and superconducting qudits [35, 11].

Generalized Pauli operators. A natural generalization of the qubit Pauli matrices to dimension d is provided by the *Weyl–Heisenberg* operators [36]. Setting $\omega = e^{2\pi i/d}$, the shift and clock operators act on the computational basis as

$$X |k\rangle = |k+1 \bmod d\rangle , \qquad Z |k\rangle = \omega^k |k\rangle , \qquad k = 0, \ldots, d-1, \tag{1.17}$$

and satisfy the Weyl commutation relation $ZX = \omega\, XZ$. The d^2 unitaries $\{X^a Z^b \mid a, b = 0, \ldots, d-1\}$ form an orthogonal operator basis with respect to the Hilbert–Schmidt inner product, and play the same structural role for qudits that Pauli strings play for qubits: they constitute the building blocks of qudit Hamiltonians, observables, and error models, and are particularly convenient for expressing truncated gauge-field operators.

Angular momentum operators. A physically motivated alternative arises when the qudit Hilbert space is identified with a spin-j representation of SU(2), so that $d = 2j + 1$. The local operators are then the angular momentum generators J_x, J_y, J_z, satisfying

$$[J_\alpha, J_\beta] = i\, \epsilon_{\alpha\beta\gamma}\, J_\gamma, \qquad \alpha, \beta, \gamma \in \{x, y, z\}, \tag{1.18}$$

where $\epsilon_{\alpha\beta\gamma}$ is the completely antisimmetryc Levi-Civita symbol. The eigenstates of J_z then provide the computational basis. Polynomial combinations of these operators span the full operator algebra on $\mathcal{H}_d$. This basis is especially relevant for lattice gauge theories, where local degrees of freedom transform under finite-dimensional representations of continuous groups and where gauge generators and Casimir operators admit a direct angular-momentum interpretation.

Unitary Bases and Gatesets for Qudits

Arbitrary single-qudit unitaries belong to SU(d), whose Lie algebra $\mathfrak{su}(d)$ is spanned by the $d^2 - 1$ traceless Hermitian *generalized Gell-Mann matrices* $\{\Lambda_a\}$. Any single-qudit unitary can be written as

$$U = \exp\left(-i \sum_{a=1}^{d^2-1} \theta_a\, \Lambda_a \right) , \qquad \theta_a \in \mathbb{R}. \tag{1.19}$$

In practice, quantum circuits are built from a finite gateset that approximates such unitaries, typically using parametrized rotations generated by a subset of the Λ_a. Universal quantum computation on N qudits is then achieved by combining arbitrary single-qudit unitaries with at least one entangling two-qudit gate, such as a generalized controlled-shift or an interaction generated by a product of Weyl–Heisenberg operators.

The choice of gateset is strongly platform-dependent and directly influences circuit depth, expressivity, and the efficient implementation of symmetries. Experimental progress on qudit platforms has been reported in several recent works [6, 11, 9]; in the following chapters we focus specifically on trapped-ion qudit architectures, which offer high-fidelity control over $d > 2$ levels and natural access to the angular-momentum operator basis discussed above.

The Clifford Group for Qubits and Qudits

As will be shown later, the Clifford group plays a central role in both entanglement and nonstabilizerness resource theories [37, 38, 39]. In order to define it, we start form the Pauli group. These definitions are given for general qudits of dimension d, and the corresponding qubit version can be obtained by setting $d = 2$. The generalized Pauli group for a single qudit is defined as the set

$$\mathcal{P} = \{X^a Z^b | a, b = 0, ..., d-1\}. \tag{1.20}$$

For a system of N qudits, the Pauli group $\mathcal{P}_n$ is defined by the tensor products of all the elements of the single qudit Pauli operators. The Clifford group $\mathcal{C}_N$ is defined as the normalizer [1] of the Pauli group in the unitary group, i.e.

$$\mathcal{C}_N = \{U | UPU^\dagger \in \mathcal{P}_N \quad \text{for all } P \in \mathcal{P}_N\}. \tag{1.21}$$

For qubits, canonical examples of Clifford gates include the Hadamard, phase, and controlled-NOT gates, while an example of non-Clifford gate is the $T-$gate.

The Clifford group maps stabilizer states to stabilizer states and preserves the structure of Pauli operators under conjugation. Starting from this, Clifford circuits can be efficiently simulated classically using the Gottesman–Knill theorem [40]. Moreover, non Clifford gates (such as the T-gate) are required in order to achieve universal quantum computation.

1.1.7 Trapped-ion qudits

Ringbauer et al. [6] demonstrated a universal quantum processor using trapped $^{40}Ca^+$ ions operating as qudits with local Hilbert-space dimensions up to seven. Unlike conventional qubit-based quantum computers that restrict physical systems to two levels, this approach exploits the native multilevel structure of trapped ions to enable more efficient quantum computation for certain applications.

The quantum information is encoded in the $S_{1/2}$ electronic ground state and the metastable $D_{5/2}$ excited state of $^{40}Ca^+$ ions (lifetime $\tau \approx 1.1$ s). A magnetic field of approximately 4.2 G splits these states into Zeeman sublevels: two for the ground state ($m = \pm 1/2$) and six for the excited state manifold ($m = \pm 5/2, \pm 3/2, \pm 1/2$). A graphical representation is shown in Fig 1.3. This creates ten allowed transitions governed by selection rules ($\Delta m = 0, \pm 1, \pm 2$), forming a highly connected eight-level system per ion.

[1] In general, given H a subgroup of G, the normalizer of H with respect to G is defined as $N_G(H) = \{g \in G | g^{-1}hg \in H \ \ \forall \ \ h \in H\}$. In contrast, the centralizer requires a more specific commutation relation: $C_G(H) = \{g \in G | g^{-1}hg = h \ \ \forall \ \ h \in H\}$

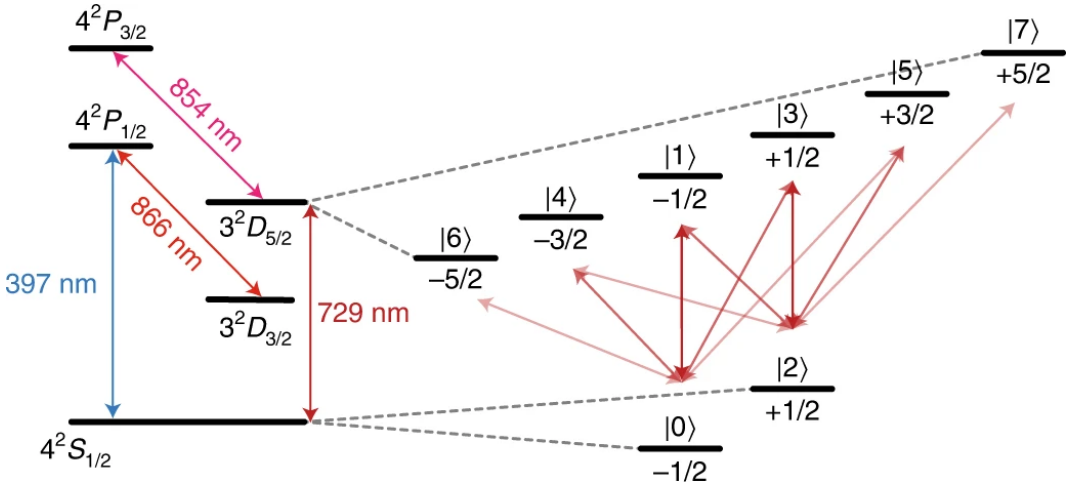

Figure 1.3: Quantum information is encoded in the $S_{\frac{1}{2}}$ and $D_{\frac{5}{2}}$ states, where each transition between S and D is accessible using a single narrowband laser at 729nm. Picture credits to Ringbauer et. al. [6]

The system implements single-qudit gates through equatorial rotations on any S–D transition:

$$R_{i,j}(\theta,\varphi) = \exp\!\left(-i\theta\sigma^{\varphi}_{i,j}/2\right) \tag{1.22}$$

where θ is the rotation angle and i, j label the addressed transition. These rotations enable arbitrary SU(d) operations through decomposition into at most $O(d^2)$ two-level rotations.

Two-qudit entangling gates are realized via the Mølmer–Sørensen interaction:

$$\mathrm{MS}_{i,j}(\theta,\varphi) = \exp\left(-i\frac{\theta}{4}\left(\sigma^{\varphi}_{i,j}\otimes\mathbb{I}+\mathbb{I}\otimes\sigma^{\varphi}_{i,j}\right)^2\right) \tag{1.23}$$

This couples the internal states to the common motional mode using bichromatic laser fields. The implemented universal two qudit gates controlled-exchange (CEX) and controlled-increment (CINC), which act as the following:

$$\mathrm{CEX}_{c,t_1,t_2} : \begin{cases} |c,t_1\rangle \leftrightarrow |c,t_2\rangle \\ |j,k\rangle \to |j,k\rangle \text{for } j\neq c,\ k\neq t_1,t_2. \end{cases} \tag{1.24}$$

$$\mathrm{CINC} : \begin{cases} |j,k\rangle \to |j,k\rangle, & \text{if } j < d-1, \\ |j,k\rangle \to |j,k\oplus 1\rangle, & \text{if } j = d-1, \end{cases} \tag{1.25}$$

Full qudit state discrimination is achieved through sequential electron shelving. The system alternates between fluorescence detection (distinguishing S and D manifolds) and reordering pulses that map different qudit states into the S manifold for measurement. Polarization-gradient cooling between detection steps mitigates scattering-induced heating. For qutrits, this yields a readout error of approximately 3×10^{-3}, compared to 5×10^{-4} for qubits [6], with readout time scaling linearly with qudit dimension.

This qudit approach offers significant benefits for applications naturally formulated in high-dimensional Hilbert spaces, including quantum simulation of lattice gauge theories, quantum chemistry calculations, and higher-spin models. While individual gate fidelities

are slightly lower than optimized qubit systems, the reduction in required quantum operations for these applications provides substantial overall efficiency gains. The system maintains coherence times exceeding 100 ms—three orders of magnitude longer than typical gate times—across all transitions. This platform has found use already in simulating lattice gauge theories [41], study toplogical phases of matter [7] and non ergodic processes [42].

1.1.8 Resource Theories in Quantum Mechanics

The notion of a *resource theory* provides a unifying framework for identifying, quantifying, and manipulating features of quantum systems that are useful for information-processing tasks [43]. The central observation is that not all quantum states and operations are equally powerful: some states carry properties that enable tasks impossible with others, and some operations are harder to implement than others. Resource theories formalize this asymmetry by partitioning states and operations into those that are *free*—easy to prepare or implement within given constraints—and those that are *resourceful*, i.e. not achievable for free.

A well-defined resource theory is specified by three ingredients:

1. a convex set $\mathcal{F}$ of *free states*, which are states preparable without consuming any resource;

2. a class $\mathcal{O}$ of *free operations*, which map $\mathcal{F}$ to itself and define what manipulations are allowed at no cost;

3. one or more *resource monotones* $\mathcal{R}$: real-valued functions on the state space that are non-increasing under $\mathcal{O}$, and vanish on $\mathcal{F}$.

The monotonicity condition ensures that free operations cannot create a resource from nothing, and that $\mathcal{R}$ provides a meaningful ordering of states by their resourcefulness. When $\mathcal{R}$ admits an operational interpretation—such as the maximum probability of success in some task—it also establishes quantitative bounds on what can be achieved with a given state.

Well-known examples of resource theories include entanglement (free states: separable states; free operations: local operations and classical cummunications-LOCC), coherence (free states: diagonal states in a fixed basis; free operations: incoherent operations [44]), and nonstabilizerness (free states: stabilizer states; free operations: Clifford unitaries [45]). Each captures a distinct physical constraint and gives rise to a different notion of what makes a quantum state useful.

Quantifying resources serves both theoretical and practical purposes. Theoretically, resource monotones enable rigorous comparison of quantum states and protocols, and establish limits on interconversion: a state with lower resource content cannot be converted into one with higher content by free operations alone. They also connect to classical simulation hardness, since states with low resource content are often efficiently simulable classically.

Practically, resource measures guide the design of algorithms and experiments by identifying which features must be actively generated and carefully preserved in the presence of noise and hardware constraints. In the context of variational quantum algorithms, resource quantification helps balance expressivity against implementability: an ansatz must be sufficiently resourceful to represent the target state, yet not so deep or entangled as to be overwhelmed by noise in the NISQ regime. Monitoring resource measures throughout a computation therefore provides diagnostics for both performance and classical simulability.

In the remainder of this section we discuss the two resources most directly relevant to this thesis: *entanglement* and *nonstabilizerness*.

Entanglement

Entanglement is one of the most fundamental resources in quantum mechanics, playing a central role in quantum communication, computation, and simulation [46]. In a bipartite system AB, a state ρ_{AB} is called *separable* if it can be written as a convex mixture of product states,

$$\rho_{AB} = \sum_k p_k \, \rho_A^{(k)} \otimes \rho_B^{(k)}, \qquad p_k \geq 0, \quad \sum_k p_k = 1; \tag{1.26}$$

states not of this form are *entangled*. The resource theory of entanglement takes separable states as free and LOCC as free operations.

For pure bipartite states, the structure of entanglement is fully characterised by the *Schmidt decomposition*: any $|\psi_{AB}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ can be written as

$$|\psi_{AB}\rangle = \sum_{k=1}^{r} \lambda_k \, |a_k\rangle \otimes |b_k\rangle \,, \qquad \lambda_k > 0, \quad \sum_k \lambda_k^2 = 1, \tag{1.27}$$

where $\{|a_k\rangle\}$ and $\{|b_k\rangle\}$ are orthonormal bases of $\mathcal{H}_A$ and $\mathcal{H}_B$ respectively, and $r \leq \min(\dim \mathcal{H}_A, \dim \mathcal{H}_B)$ is the *Schmidt rank*. The state is a product state if and only if $r = 1$; otherwise it is entangled. The Schmidt coefficients $\{\lambda_k^2\}$ coincide with the eigenvalues of the reduced density matrix $\rho_A = \mathrm{Tr}_B[|\psi_{AB}\rangle \langle\psi_{AB}|]$, and entanglement is uniquely quantified (up to monotonic transformations) by the *entanglement entropy*

$$E(|\psi_{AB}\rangle) = S(\rho_A) = -\sum_k \lambda_k^2 \log_2 \lambda_k^2 = -\mathrm{Tr}(\rho_A \log_2 \rho_A) \,, \tag{1.28}$$

which vanishes for product states and reaches its maximum of $\log_2(r)$for maximally entangled states.

In variational quantum algorithms and lattice gauge theory simulations, entanglement is both a resource and a challenge: sufficient entanglement is necessary to represent complex many-body correlations, yet excessive or poorly structured entanglement increases circuit depth and sensitivity to noise [47]. Monitoring entanglement entropy across bipartitions therefore serves as a diagnostic tool for assessing both the expressivity of variational ansätze and the complexity of the simulated states.

Nonstabilizerness (Magic)

Entanglement alone does not suffice for universal quantum computation. Operations confined to the *stabilizer formalism*—stabilizer states, Clifford unitaries, and Pauli measurements—can be efficiently simulated classically due to the Gottesman–Knill theorem [48]. As mentioned previously, Clifford unitaries are defined as those that map the Pauli group to itself under conjugation,

$$U \in \mathcal{C} \quad \Longleftrightarrow \quad UPU^\dagger \in \mathcal{P}_N \quad \forall \, P \in \mathcal{P}_N. \tag{1.29}$$

Prominent single-qubit examples are the Hadamard gate and the phase gate,

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \qquad S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \tag{1.30}$$

together with the two-qubit CNOT gate in Eq. (1.15). While Clifford circuits can generate highly entangled states, they are not universal for quantum computation: any circuit composed solely of Clifford gates can be efficiently simulated on a classical computer, regardless of the number of qubits.

The missing ingredient is provided by *non-Clifford* gates, of which the single-qubit T gate is the canonical example,

$$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}. \tag{1.31}$$

Adding T to the Clifford gate set achieves universality. The resource that quantifies the departure from the efficiently simulable stabilizer regime is called *nonstabilizerness*, or *magic*. The corresponding resource theory takes as free states those reachable from $|0\rangle^{\otimes N}$ by Clifford circuits alone—the *stabilizer states*, examples of which include the Bell state

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix}, \tag{1.32}$$

and takes Clifford operations as free operations. Non-stabilizer states carry magic and enable universal quantum computation when supplemented with Clifford gates. A paradigmatic example is the *magic state*

$$|T\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ e^{i\pi/4} \end{pmatrix}, \tag{1.33}$$

which is the output of the T gate acting on $|+\rangle$ and serves as the key resource in magic-state distillation [37].

Wigner negativity. A computable witness of nonstabilizerness that can be easily translated to the qudit formalism is provided by the *discrete Wigner function* [45]. Its construction proceeds in two steps.

Given a point $\vec{v} = v_1 + iv_2$ in the discretised local phase space $\mathbb{Z}_d^{\otimes 2}$, the local *Heisenberg–Weyl* operators are defined as

$$T_{\vec{v}} = \tau^{-v_1 v_2} X^{v_1} Z^{v_2}, \qquad \tau = e^{i(D+1)\pi/D}, \tag{1.34}$$

and extended to N-qudit systems by

$$T_{\vec{V}} = T_{\vec{v}^{(1)}} \otimes \cdots \otimes T_{\vec{v}^{(N)}}, \qquad \vec{V} \equiv \vec{v}^{(1)} \oplus \cdots \oplus \vec{v}^{(N)} \in \mathbb{Z}_D^{\otimes 2N}. \tag{1.35}$$

From these, one constructs the *phase-space point operators*

$$A_{\vec{0}} = \frac{1}{D^N} \sum_{\vec{V}} T_{\vec{V}}, \qquad A_{\vec{V}} = T_{\vec{V}}\, A_{\vec{0}}\, T_{\vec{V}}^{\dagger}, \tag{1.36}$$

which define the *discrete Wigner function* of a state ρ as

$$W_\rho(\vec{V}) \coloneqq \frac{1}{D^N} \operatorname{Tr}\!\left[A_{\vec{V}}\, \rho\right]. \tag{1.37}$$

For stabilizer states, W_ρ is non-negative at every phase-space point [49]; negativity of W_ρ therefore certifies the presence of magic. This is quantified by the *mana* [45],

$$\mathcal{M}(\rho) = \log_2 \left(\sum_{\vec{V}} |W_\rho(\vec{V})| \right), \tag{1.38}$$

which is a monotone under free (Clifford) operations and vanishes exactly on stabilizer states.

Stabilizer Rényi entropies. A versatile family of nonstabilizerness measures is provided by the *Stabilizer Rényi Entropies* (SRE) [50, 51]. For a pure N-qudit state $|\psi\rangle$ and order $n \neq 1$, the SRE is defined as

$$M_n(|\psi\rangle) = \frac{1}{1-n} \log_2 \left[\sum_{P \in \mathcal{P}_N} \frac{\langle \psi | P | \psi \rangle^{2n}}{d^N} \right], \tag{1.39}$$

where $\mathcal{P}_N$ is the N-qudit Pauli group defined in Eq. (1.14) and d^N is the dimension of the Hilbert space (we assume all the qudits have the same local dimension). The SRE is non-increasing under Clifford operations, always positive and vanishes for stabilizer states, i.e. it satisfies all the criteria [51, 52] to be a valid measure. It is considered to be the most relevant measure for nonstabilizerness due to the fact that is can be experimentally computed [53].

A direct evaluation of Eq. (1.39) requires computing d^{2N} Pauli expectation values, which is prohibitive for all but small systems. Efficient approximation schemes based on Monte Carlo sampling of the Pauli spectrum [54] and fast Hadamard transform [55] bring the SRE within reach of medium-sized systems and are employed in the numerical studies of this thesis.

1.2 Quantum optimization

A main focus of application of qudits in this thesis will be to use their increased dimensionality to enhance quantum optimization algorithms. This section covers a general background on classical and quantum optimization, with the description of the typical problems handled and the quantum algorithms typically used to tackle them.

1.2.1 Classical and Quantum Perspectives on Combinatorial Optimization

A large class of optimization problems of practical and theoretical interest involves decision variables that take discrete, often binary or integer, values. Such problems arise naturally in scheduling, routing, resource allocation, and graph-theoretic applications, and are commonly formulated as *integer optimization* problems. In their most general form, these problems can be written as

$$\min_{\mathbf{x}\in\mathbb{Z}^n} C(\mathbf{x}) \quad \text{subject to} \quad \mathbf{x} \in \mathcal{S}, \tag{1.40}$$

where $\mathbf{x}$ denotes a vector of n integer or binary decision variables, $\mathcal{S}$ represents a set of feasible solutions, and $C(\mathbf{x})$ is a real-valued objective function [56]. The number of possible solutions grows exponentially with n, and many such problems are NP-hard.

Constraints restrict the admissible configurations of variables and may encode physical limitations, logical conditions, or combinatorial rules. They are typically expressed as equalities or inequalities which have to be fulfilled for a solution to be feasible

$$g_k(\mathbf{x}) \leq 0, \qquad h_\ell(\mathbf{x}) = 0. \tag{1.41}$$

In classical optimization, constraints are either enforced explicitly during the search or incorporated into the objective function through penalty terms, yielding an effective unconstrained problem.

A widely used unconstrained formulation is *quadratic unconstrained binary optimization* (QUBO), in which all variables are binary, $x_i \in \{0, 1\}$, and the cost function is at most quadratic:

$$C(\mathbf{x}) = \sum_i h_i x_i + \sum_{i<j} J_{ij} x_i x_j. \tag{1.42}$$

In order to ensure that the solutiuon satisfies the constraints, the cost function is enriched with penalty terms for the unfeasible solutions. These terms are chosen such that constraint-violating configurations are energetically suppressed. Many paradigmatic combinatorial optimization problems, including Max-Cut, graph coloring, and the knapsack problem, admit natural QUBO representations [57].

The QUBO structure provides a direct bridge to quantum optimization. Binary variables can be mapped to qubits via

$$x_i \to \frac{1 - \sigma_z^i}{2}, \tag{1.43}$$

where σ_z^i is the Pauli-Z operator acting on qubit i. Under this mapping, the classical cost function becomes a diagonal quantum Hamiltonian

$$H_C = \sum_i h_i \frac{1 - \sigma_z^i}{2} + \sum_{i<j} J_{ij} \frac{(1 - \sigma_z^i)(1 - \sigma_z^j)}{4}, \tag{1.44}$$

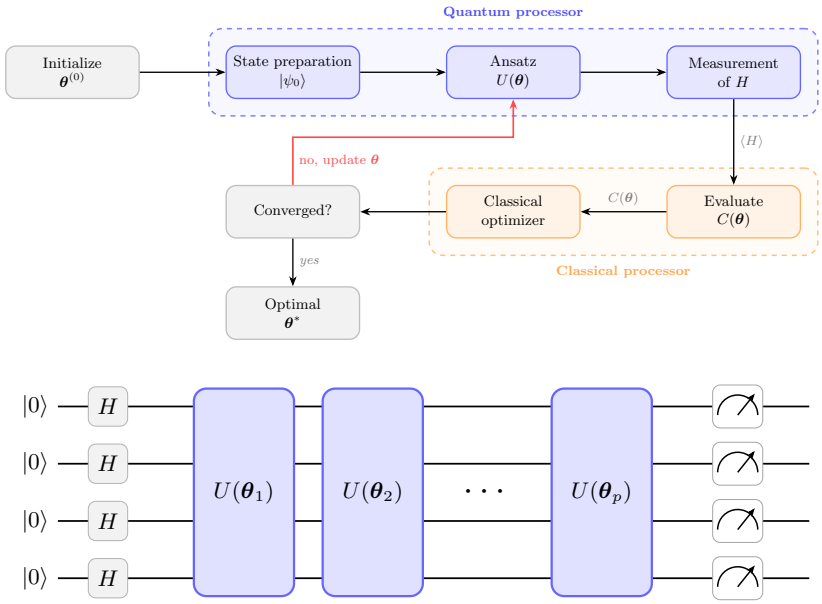

Figure 1.4: **Top:** Algorithm flow for a generic VQE: after preparing an initial state set θ^0 the quantum computer applies a set of parametrized unitaries $U(\theta)$. The state is then measured and from the results the cost function $C(\theta)$ is computed and processed through a classical optimizer. The procedure is then iterated until convergence.
Bottom: Generic ansatz (structured in layers) for a VQA on $n = 4$ qubits. Each layer consists of single-qubit R_y rotations with trainable parameters θ_i, followed by a ladder of CNOT gates to entangle neighbouring qubits. The circuit prepares the parametrized state $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})\,|\psi_0\rangle$, where $|\psi_0\rangle = |+\rangle^{\otimes n}$ after the initial Hadamard layer.

whose ground state encodes the optimal solution of the original problem. This formulation allows combinatorial optimization to be interpreted as a ground-state preparation problem and forms the basis of quantum optimization algorithms.

1.2.2 Variational Quantum Algorithms

Variational Quantum Algorithms (VQAs) are hybrid quantum–classical algorithms designed to address computational problems using near-term quantum hardware [13, 58, 59]. In this framework, a target problem is formulated as the optimization of a cost function defined over a family of parametrized quantum states, while parameter updates are performed on a classical processor.

At the core of a VQA lies a parametrized quantum circuit, or *ansatz*, which prepares states of the form

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})\,|\psi_0\rangle\,, \tag{1.45}$$

where $|\psi_0\rangle$ is a fixed reference state and $U(\boldsymbol{\theta})$ is a unitary operator depending on real parameters $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_p)$.

A problem-specific cost function is defined as the expectation value of a Hermitian op-

erator H,

$$C(\boldsymbol{\theta}) = \langle \psi(\boldsymbol{\theta}) | \, H \, | \psi(\boldsymbol{\theta}) \rangle \,, \tag{1.46}$$

which may represent the energy of a physical system or a classical objective function mapped to a quantum Hamiltonian. The optimization proceeds by iteratively evaluating $C(\boldsymbol{\theta})$ on a quantum device and updating the parameters using a classical optimization routine.

The hybrid structure of VQAs makes them well suited for the noisy intermediate-scale quantum (NISQ) era, as circuit depth can be traded for repeated state preparation and measurement. When the cost Hamiltonian is expressed as a sum of local operators,

$$H = \sum_j c_j H_j, \tag{1.47}$$

the expectation value is reconstructed from measurements of the individual terms. The finite number of measurements required to estimate these expectation values introduces statistical noise, which must be taken into account by the classical optimizer.

Variational quantum algorithms provide a flexible framework encompassing applications ranging from quantum chemistry to combinatorial optimization. In optimization-oriented VQAs, constraints and problem structure can be incorporated either directly into the cost Hamiltonian or through the design of the variational circuit. This perspective sets the stage for the Quantum Approximate Optimization Algorithm, which is tailored specifically to combinatorial optimization problems.

1.2.3 Focus: Quantum Approximate Optimization Algorithm

The Quantum Approximate Optimization Algorithm (QAOA) [60, 12, 61] is a hybrid variational quantum algorithm designed to solve combinatorial optimization problems. Its design is inspired by quantum annealing [62, 63, 64] where one considers a time-dependent Hamiltonian

$$H_a(t) = (1 - \lambda(t)) \, H_M + \lambda(t) H_C, \quad t \in [0, T], \quad \lambda(0) = 0 \lambda(T) = 1 \tag{1.48}$$

which interpolates between a simple mixing Hamiltonian H_M, whose ground state $|\psi_0\rangle$ is easy to prepare, and the cost Hamiltonian H_C encoding the optimization problem. If the system is initialized in $|\psi_0\rangle$ and evolved sufficiently slowly, the adiabatic theorem guarantees that it remains in the instantaneous ground state of $H(t)$, reaching the ground state of H_C at $t = T$. However, the required evolution time scales as

$$T = \mathcal{O}\left(\frac{1}{\Delta_{\min}^2} \right), \tag{1.49}$$

where $\Delta_{\min} = \min_{t \in [0,T]} \Delta(t)$ is the minimum spectral gap along the path, which can be exponentially small or even vanishing for hard problem instances, making the approach generally intractable.

QAOA [12] can be understood as a Trotterized discretization of this adiabatic evolution. Applying a first-order Trotter decomposition to the time-ordered evolution operator $\mathcal{T} \exp\left(-i \int_0^T H(t) \, dt \right)$ and dividing the total time T into p steps yields

$$\mathcal{T} e^{-i \int_0^T H(t) \, dt} \approx \prod_{l=1}^{p} e^{-i \alpha_l H_C} e^{-i \beta_l H_M}, \tag{1.50}$$

where α_l and β_l are determined by the annealing schedule. Rather than fixing these angles to a particular schedule, QAOA treats them as free variational parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{R}^{2p}$ to

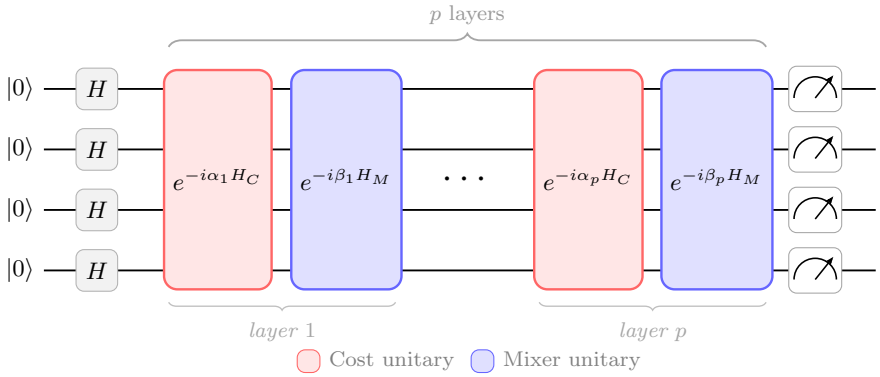

Figure 1.5: High-level circuit representation of the QAOA ansatz at depth p. Each of the p layers consists of a cost unitary $e^{-i\alpha_\ell H_C}$ (red) and a mixer unitary $e^{-i\beta_\ell H_M}$ (blue), applied alternately to the initial state $|+\rangle^{\otimes n}$. The circuit prepares the parametrized state $|\psi(\boldsymbol{\alpha}, \boldsymbol{\beta})\rangle$, whose expectation value $\langle H_C \rangle$ is measured and returned to the classical optimizer for parameter updates.

be optimized classically, trading the adiabatic guarantee for a flexible, hardware-amenable ansatz. Its schematics is represented in Fig. 1.5.

Given a cost function $C(\boldsymbol{x})$ of N classical variables $\boldsymbol{x} = (x_1, \ldots, x_N)$, QAOA aims at finding the solution to

$$\min_{\boldsymbol{x} \in \{0,1\}^N} C(\boldsymbol{x}), \tag{1.51}$$

yielding the optimal configuration $\boldsymbol{x}_{\min}$.

An important class of problems addressed by QAOA consists of quadratic unconstrained binary optimization (QUBO) problems [57], whose cost function can be written as

$$C(\boldsymbol{x}) = 4\boldsymbol{x}^T J \boldsymbol{x} = 4 \sum_{i,j=1}^{N} J_{ij} x_i x_j, \tag{1.52}$$

where J is a real symmetric cost matrix. Many paradigmatic optimization problems can be mapped to this form [65, 66], and QUBOs are equivalent to Ising Hamiltonians.

To implement the optimization on a quantum device, the classical variables are promoted to operators acting on the Hilbert space $\mathcal{H}_C = \mathbb{C}^{2^N}$ according to

$$x_i \to \frac{1 - \sigma_z^i}{2}, \tag{1.53}$$

where σ_z^i is the Pauli-Z operator acting on qubit i. The cost function is thus mapped to a diagonal Hamiltonian H_C whose eigenvalues correspond to classical costs,

$$H_C |\boldsymbol{x}\rangle = C(\boldsymbol{x}) |\boldsymbol{x}\rangle . \tag{1.54}$$

The solution of the optimization problem is encoded in the ground state of the Ising Hamiltonian

$$H_C = \sum_{i,j} J_{ij} \sigma_i^z \sigma_j^z + \sum_i h_i \sigma_i^z, \tag{1.55}$$

with $h_i = 2\sum_{j=1}^{N} J_{ij}$.

The QAOA ansatz prepares a parametrized trial state

$$|\psi(\boldsymbol{\alpha},\boldsymbol{\beta})\rangle = U(\boldsymbol{\alpha},\boldsymbol{\beta})\,|\psi_0\rangle\,, \tag{1.56}$$

where the unitary is constructed from alternating applications of the cost Hamiltonian and a mixing Hamiltonian H_M,

$$U(\boldsymbol{\alpha},\boldsymbol{\beta}) = \prod_{l=1}^{p} e^{i\beta_l H_M} e^{i\alpha_l H_C}. \tag{1.57}$$

The mixer induces transitions between computational basis states and must not commute with H_C. A common choice is

$$H_M = \sum_i \sigma_i^x. \tag{1.58}$$

The expectation value of the cost Hamiltonian,

$$E(\boldsymbol{\alpha},\boldsymbol{\beta}) = \langle\psi(\boldsymbol{\alpha},\boldsymbol{\beta})|\,H_C\,|\psi(\boldsymbol{\alpha},\boldsymbol{\beta})\rangle\,, \tag{1.59}$$

is minimized over the variational parameters $(\boldsymbol{\alpha},\boldsymbol{\beta})$ using a classical optimizer. By the variational principle, this provides an upper bound to the ground-state energy of H_C.

QAOA can be generalized to problems involving d-ary variables $x_i \in \{0,\ldots,d-1\}$ by representing them with qudits of local Hilbert space dimension $\mathbb{C}^d$. In this case, Pauli operators are replaced by angular momentum operators $\{L_x, L_y, L_z\}$ with representation index $\ell = (d-1)/2$, and classical variables are mapped as

$$x_i \to L_z^i + \frac{d-1}{2}. \tag{1.60}$$

To generate universal dynamics in the qudit space, the mixer is augmented to include quadratic terms,

$$H_M = \sum_i \left(\beta L_x^i + \gamma (L_z^i)^2\right), \tag{1.61}$$

which allows the generation of arbitrary $SU(d)$ unitaries [67, 68]. Once the cost Hamiltonian and mixer are defined, the algorithm proceeds analogously to the qubit case.

1.3 Lattice gauge theories

A second field in which qudits are expected to enhance the performance of quantum algorithms is the quantum simulation of lattice gauge theories (LGTs), a formulation of quantum field theories on a discrete spacetime lattice.

Gauge-invariant quantum field theories [69, 70] form the foundation of modern theoretical physics, providing the framework for all known fundamental interactions except gravity. The Standard Model of particle physics[71] is entirely formulated in terms of gauge theories, with quantum electrodynamics, the electroweak theory, and quantum chromodynamics (QCD) describing electromagnetic, weak, and strong interactions, respectively. Despite their remarkable success, gauge theories formulated in the continuum generically suffer from ultraviolet divergences arising from fluctuations at arbitrarily short length scales. The consistent definition of physical observables therefore requires the introduction of regularization and renormalization procedures.

Lattice gauge theories [72] offer a conceptually transparent and fully non-perturbative regularization of gauge theories by discretizing space (or spacetime) onto a lattice while preserving exact local gauge invariance. The finite lattice spacing provides a natural ultraviolet cutoff, rendering all quantities finite at intermediate stages of the calculation. In the context of QCD, the lattice formulation has played a crucial role in establishing genuinely non-perturbative phenomena such as confinement and the emergence of a mass gap [73, 74], which are inaccessible to standard perturbative techniques. The continuum theory is recovered by taking the lattice spacing to zero while appropriately renormalizing the parameters of the theory, making LGTs both a mathematically well-defined regulator and a powerful computational framework for exploring strongly coupled gauge dynamics.

Originally introduced within high-energy physics, lattice gauge theories have since become a central tool for the study of strongly interacting quantum many-body systems and, more recently, a prominent target for quantum simulation and variational quantum algorithms [75, 76, 21, 77]. Their local structure, exact gauge constraints, and finite-dimensional truncations make them particularly well suited for implementation on near-term quantum devices.

In this section, we introduce the basic concepts of lattice gauge theories and present the Hamiltonian formulation due to Kogut and Susskind. The discussion is kept general and group-agnostic, as specific realizations—including SU(2), $\mathbb{Z}_N$, and D_N gauge theories—will be addressed in detail in subsequent sections.

1.3.1 What Is a Lattice Gauge Theory

A gauge theory is defined by the invariance of its dynamics under local transformations belonging to a symmetry group G, referred to as the *gauge group*. In the continuum, gauge fields are defined at every point in space or spacetime, and gauge invariance leads to the introduction of covariant derivatives and gauge connections.

A central role in lattice gauge theories is played by the representation theory of the gauge group G. A (unitary) representation [78] of G is a homomorphism that assigns to each group element $g \in G$ a linear operator $D^J(g)$ acting on a vector space $\mathcal{V}$, such that group multiplication is preserved,

$$D^J(g_1 g_2) = D^J(g_1) D^J(g_2)\,. \tag{1.62}$$

A representation is said to be *irreducible* if the vector space $\mathcal{V}$ admits no nontrivial invariant subspaces under the action of all $D^J(g)$. For compact groups, such as Lie groups and discrete

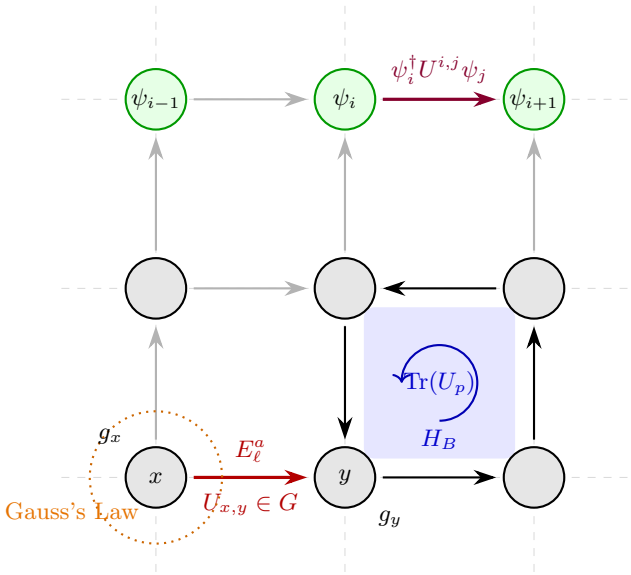

Figure 1.6: Description of LGT

groups, any unitary representation can be decomposed into a direct sum of irreducible representations (irreps), which therefore constitute the fundamental building blocks of the theory.

In a lattice gauge theory, space is discretized into a lattice composed of sites (vertices) and links (edges)[19]. Each vertex $\mathbf{x}$ and link ℓ is endowed respectively with local fields $\psi_{\mathbf{x}}$ and U_ℓ, describing the (usually fermionic) matter and the (bosonic) gauge degrees of freedom. In an LGT framework, the field U_ℓ takes values in the group G and represents the parallel transporter between the two sites connected by the link.

The Hilbert space of the theory is defined through the tensor product of all the local Hilbert spaces defined singularly on the links and on the vertices of the lattice. The Hilbert space of the links is defined by the span of two possible basis: the group element basis and the representation basis.

The group element basis is defined as the set of orthonormal states $\{|g\rangle, g \in G\}$, which we take to be the eigenstates of a link operator U_ℓ^J with eigenvalues the matrix elements of $D^J(g)$ in the irrep J:

$$U^J_{mn,\ell} |g\rangle = D^J_{mn}(g) |g\rangle . \tag{1.63}$$

For the rest of the thesis we will drop the internal indices m, n if not specifically needed and write U_ℓ^J.

The group representation basis is taken to be conjugate to the group element basis through the group's quantum Fourier transform QFT_G:

$$\langle g|Jmn\rangle = \sqrt{\frac{dim(J)}{|G|}} D^J_{nm}(g). \tag{1.64}$$

States are then labelled with a representation index J and sub-indices n, m that distinguishes the projections on J of the left and right components of the gauge field.

Local gauge transformations act independently at each lattice site. Given a group element $g \in G$, the corresponding gauge transformations acting on matter and gauge fields are defined through unitary operators $\Theta_g^\psi, \Theta_g^L, \Theta_g^R$. For the gauge field transformations we have in general that $\Theta_g^L \neq \Theta_g^R$, since for non-abelian groups the left and right multiplication results in a different element of the group: $gh \neq hg, \;\; h, g \in G$. This is reflected on the action of the left and right transformations on group element states:

$$\Theta_g^L |h\rangle = |gh\rangle \qquad\qquad \Theta_g^R |h\rangle = |hg^{-1}\rangle \tag{1.65}$$

The action of these transformation on the link operators is the following [19]

$$\Theta_g^L U_{mn}^J \Theta_g^{L\dagger} = D_{mk}(g^{-1}) U_{kn}^J \qquad\qquad \Theta_g^R U_{mn}^J \Theta_g^{R\dagger} = U_{mk}^J D_{kn}(g) \tag{1.66}$$

wher the implicit summation rund over the sub indices of the irrep J.

Given a directed 2-d rectangular lattice, ingoing links are defined as the one that follow the direction of the lattice, while the opposite ones are referred to as the outgoing ones.

When defining a local gauge transformation on a vertex $\mathbf{x}$ according to a group element $g \in G$, we associate left transformations Θ_g^L with ingoing links and right transformations Θ_g^R with outgoing links.

The overall transformation acting on a vertex $\mathbf{x}$ is then

$$\Theta_g^{\mathbf{x}} = \Theta_g^{\psi,\mathbf{x}} \prod_{i(\mathbf{x})} \Theta_g^{L,i(\mathbf{x})} \prod_{o(\mathbf{x})} \Theta_g^{R,o(\mathbf{x})}, \tag{1.67}$$

where $i(\mathbf{x})$ and $o(\mathbf{x})$ label all the ingoing and outgoing links of $\mathbf{x}$ and must leave all physical quantities invariant, starting from the Hamiltonian of the system. In particular, physical states are the one that are gauge invariant under the action of (1.67), i.e.

$$\Theta_g^{\mathbf{x}} |\psi\rangle = |\psi\rangle \quad \forall \mathbf{x} \tag{1.68}$$

This is the most general form of the gauss law for LGTs.

These features render LGTs particularly attractive for numerical simulations and for implementation on quantum devices with finite-dimensional local Hilbert spaces [6].

While Abelian lattice gauge theories, such as those based on $\mathbb{Z}_N$ or U(1) groups, already display rich physical behavior, non-Abelian LGTs play a fundamental role in the description of nature [71]. In particular, non-Abelian gauge theories underlie the Standard Model of particle physics, with quantum chromodynamics (QCD) being formulated as an SU(3) gauge theory [71].

Non-Abelian LGTs are characterized by non-abelian gauge groups and can describe phenomena such as confinement, asymptotic freedom, and the emergence of complex phase structures [79, 80, 74]. These features make classical simulation extremely challenging, especially in regimes of real-time dynamics or finite density [81, 82, 83].

From the perspective of quantum simulation and variational quantum algorithms, non-Abelian LGTs represent a particularly compelling application. Their local gauge constraints impose strong structural requirements on admissible quantum states, motivating the development of symmetry-preserving encodings and constrained variational Ansätze.

1.3.2 Kogut-Susskind Hamiltonian for LGTs

A widely used formulation of lattice gauge theories in the context of quantum simulation is the Hamiltonian approach introduced by Kogut and Susskind [84]. In this framework, time

remains continuous while space is discretized, and the dynamics is governed by a gauge-invariant Hamiltonian acting on the Hilbert space of link and matter degrees of freedom.

For a general gauge group G, the Kogut–Susskind Hamiltonian can be written as

$$H = H_E + H_B + H_M + H_{hop}, \tag{1.69}$$

where H_E is the electric term, H_B is the magnetic (plaquette) term, H_M the matter contributions and H_{hop} is the kinetic term (hopping).

The electric term is a sum over lattice links and depends on the generators of the gauge group,

$$H_E = \frac{g_0^2}{2} \sum_{\ell} \left(E_\ell\right)^2, \tag{1.70}$$

where g_0^2 is the coupling constant and E_ℓ are the electric field operators associated with link ℓ. In particular, the electric field operator is diagonal on the representation basis and can be written as [19]

$$E^2 = \sum_J \alpha(J) \Pi_J \tag{1.71}$$

where Π_J is a projection over the subspace labelled by the representation J. The group structure can determine the coefficients α_j, as will be shown later.

The magnetic term encodes the dynamics of the gauge field through elementary plaquettes p of the lattice,

$$H_B = -\frac{1}{2g_0^2} \sum_p \mathrm{Tr}\left(U_p\right) + H.C., \qquad U_p = U_1 U_2 U_3^\dagger U_4^\dagger \tag{1.72}$$

Where U_1, U_2, U_3, U_4 are respectively the gauge operators acting on the links of the plaquette, counted clockwise starting from the lower left vertex. This term favors configurations with minimal flux and reduces to the familiar field-strength contribution in the continuum limit.

In order to define the mass and hopping terms of the Hamiltonian, we need to specify over which representation the mass fields are defined. We say that matter fields ψ^J lives in a specific representation J of the gauge group, if

$$\Theta_g^\psi \psi^J \Theta^\psi : M_g^\dagger = D^J(g^{-1})_m n \psi_n, \tag{1.73}$$

i.e. the field transforms according to J with a left group transformation (label J is dropped for the rest of the thesis unless required). This implies that the number of colors of the matter field is going to be equal to $dim(\psi^J) = dim(J)$. It is worth mentionign that it is possible to choose the matter to transorm with right transformations and redefine everything to be consistent. Since all irreducible representations of abelian groups are one-dimensional[2] it is necessary to consider a non abelian gauge group in order to introduce nontrivial color dynamics in the theory. The mass term is defined through staggered fermions [84]

$$H_M = \sum_{\mathbf{x}} (-1)^{x+y} M \psi_{\mathbf{x}}^\dagger \psi_{\mathbf{x}}. \tag{1.74}$$

The kinetic term is defined by combining the matter fields and the parallel transport in the following way

$$H_{hop} = J \sum_{\mathbf{x},\hat{\mu}} \psi_{\mathbf{x}}^\dagger U_{\mathbf{x},\mathbf{x}+\hat{\mu}} \psi_{\mathbf{x}+\hat{\mu}} + H.C. \tag{1.75}$$

[2]Shur's lemma

where $\hat{\mu} \in \{\hat{x}, \hat{y}\}$ are the unitary vectors of the lattice and J is the coupling between matter and fields (the label J label on the fields ψ, U has been dropped).

The Hamiltonian written in this way, considering the transformations of matter (1.73) and field operators (1.66) is explicitly gauge invariant:

$$\Theta_{g,\mathbf{x}} H \Theta^{\dagger}_{g,\mathbf{x}} = H \Rightarrow [H, \Theta_{g,x}] = 0 \quad \forall g, \mathbf{x}. \tag{1.76}$$

The general structure outlined above applies to both Abelian and non-Abelian gauge groups. In the following sections, we will specialize this framework to specific choices of G, highlighting how different group structures lead to distinct local Hilbert spaces and operator algebras [6].

The following sections contain explicit realizations of the Kogut-Susskind Hamiltonian and gauge transformations for specific groups. The gauge group Z_N and $SU(2)$ are introduced in order to present a parallel between abelian and nonabelian groups. The group D_3 is analyzed in detail, as it will be the center of all the results of Chapter 5.

1.3.3 General Models: $\mathbb{Z}_N$ and SU(2)

In this subsection we introduce two broad classes of gauge groups that will serve as reference models throughout this thesis: discrete Abelian groups of the form $\mathbb{Z}_N$ and the continuous non-Abelian group SU(2). These groups capture complementary aspects of lattice gauge theories, ranging from minimal realizations of gauge invariance to genuinely non-Abelian dynamics, and provide a natural bridge between theoretical models and quantum simulation with qubits and qudits [6].

The choice of $\mathbb{Z}_N$ and SU(2) gauge groups is motivated by both theoretical significance and practical considerations related to quantum simulation. Discrete Abelian gauge theories of the form $\mathbb{Z}_N$ represent some of the simplest nontrivial realizations of local gauge invariance on a lattice. Despite their apparent simplicity, these models already capture essential features of gauge theories, such as Gauss's law constraints, confinement–deconfinement transitions, and the emergence of topological sectors [18, 85]. As such, $\mathbb{Z}_N$ gauge theories serve as minimal testbeds for studying fundamental mechanisms of lattice gauge dynamics while remaining analytically and numerically tractable.

From the perspective of quantum simulation, $\mathbb{Z}_N$ models are particularly attractive because their finite-dimensional local Hilbert spaces can be naturally encoded using qubits or low-dimensional qudits. This makes them well suited for near-term quantum hardware, where limitations on coherence times and gate fidelities favor compact local degrees of freedom. $\mathbb{Z}_N$ gauge theories can often be viewed as controlled approximations to continuous Abelian theories ($U(1)$ for example), providing a systematic route toward more realistic models while maintaining experimental feasibility. This is contrapposed to other methods based for example on truncations of the local basis [86, 87].

In contrast, SU(2) represents the simplest non-Abelian Lie group and plays a central role in high-energy physics, most notably as the gauge group of the weak interaction and as a building block for more general SU(N) Yang–Mills theories. Lattice gauge theories based on SU(2) exhibit genuinely non-Abelian phenomena, including self-interacting gauge fields, nontrivial representation theory. Studying SU(2) therefore allows one to address qualitative features that are entirely absent in Abelian or discrete models, making it an essential stepping stone toward the simulation of realistic non-Abelian gauge theories.

At the same time, SU(2) is the perfect example of how even simple non-abelian groups can provide many layers of complexity to a gauge theory. This group has been studied with explicit constructions of lattice Hamiltonians, gauge constraints, and truncated Hilbert

spaces that are amenable to numerical methods and quantum simulation schemes [**silvi**, 23, 88, 89]. In this sense, SU(2) serves as a benchmark non-Abelian model for exploring how concepts such as gauge invariance, symmetry protection, and continuum limits can be implemented and tested in quantum simulators [90, 81, 23].

Taken together, $\mathbb{Z}_N$ and SU(2) gauge theories span a broad spectrum of lattice gauge dynamics—from discrete to continuous and from Abelian to non-Abelian. This complementary character makes them ideal reference models for systematically investigating both foundational aspects of gauge theories and their concrete realization in quantum simulation platforms.

Properties and definitions for $\mathbb{Z}_N$

The cyclic group $\mathbb{Z}_N$ is the simplest example of a finite Abelian gauge group. It consists of N elements,

$$\mathbb{Z}_N = \{0, 1, \ldots, N-1\}, \tag{1.77}$$

with group operation given by addition modulo N. Despite its simplicity, $\mathbb{Z}_N$ gauge theory retains the essential structural features of local gauge invariance and serves as an important testing ground for both classical and quantum simulation methods.

All irreducible representations of $\mathbb{Z}_N$ are one-dimensional and labeled by an integer $k = 0, \ldots, N-1$ [6]:

$$\rho_k(\bullet) : \mathbb{Z}_N \to \mathbb{C} \qquad \rho_k(g) = e^{\frac{i2\pi kg}{N}} \tag{1.78}$$

In a lattice formulation, link degrees of freedom are naturally represented by an N-dimensional local Hilbert space spanned by group element states $|m\rangle$, with $m \in \mathbb{Z}_N$. We can use the generalized Pauli operators introduced earlier in order to define

$$U^J |m\rangle = Z^J |m\rangle = \omega^{Jm} |m\rangle , \quad \omega = e^{2\pi i/N}. \tag{1.79}$$

By defining the representation basis states $|k\rangle$ through Fourier transform (we have no internal labels since all representations are one dimensional)

$$|k\rangle = \sum_{g\in G} \frac{1}{\sqrt{N}} e^{-i\frac{2\pi kg}{N}} |g\rangle \tag{1.80}$$

it's also possible to check that

$$X |k\rangle = e^{-i\frac{2\pi k}{N}} |k\rangle , \tag{1.81}$$

which implies that the generalized X operator for qudits is a proper candidate for the electric field of the $\mathbb{Z}_N$ pure gauge LGT. We can then write

$$H_E = \frac{g^2}{2} \sum_{\ell} \left(1 - \frac{1}{2}(X_\ell + X_\ell^\dagger)\right), \tag{1.82}$$

while the magnetic term couples links around elementary plaquettes,

$$H_B = -\frac{1}{2g^2} \sum_{p} \left(Z_p + Z_p^\dagger\right), \tag{1.83}$$

where Z_p denotes the ordered product of Z operators around plaquette p as in (1.72).

Notice how the trace operation introduced in (1.72) is not present, since the group is abelian.

The special case $N = 2$ corresponds to a $\mathbb{Z}_2$ lattice gauge theory, where the local Hilbert space on each link is two-dimensional and can be represented by a qubit [6]. In this case, the generalized Pauli operators reduce to the standard Pauli matrices,

$$X \to \sigma_x, \qquad Z \to \sigma_z, \tag{1.84}$$

and the Hamiltonian becomes equivalent to an Ising-type gauge theory. Owing to its minimal local dimension and direct qubit encoding, $\mathbb{Z}_2$ gauge theory is widely used as a benchmark model for quantum simulation and variational algorithms [91, 92].

Properties and definitions for SU(2)

The group SU(2) is the simplest non-Abelian Lie group and plays a central role in both high-energy and condensed-matter physics [6]. It consists of all 2×2 unitary matrices with unit determinant:

$$\mathrm{SU}(2) = \{U \in GL_2(\mathbb{C}) s.t. UU^\dagger = 1, \det(U) = 1\}. \tag{1.85}$$

Being a continuous Lie group, its elements can be written using the generators of the $\mathfrak{su}(2)$ Lie algebra as

$$g \in \mathrm{SU}(2), \qquad g = e^{i\alpha_a T^a}, \qquad [T^a, T^b] = i f^{abc} T^c \tag{1.86}$$

where f^{abc} is the completely antisymmetric group structure constant.

Irreducible representations of SU(2) are labeled by a half-integers $j \in \{0, \frac{1}{2}, 1, \frac{3}{2}, \dots\}$ and have dimension $2j + 1$ [6]. The corresponding representation spaces can be interpreted as spin-j systems, with basis states $|j, mn\rangle$, where $m, n = -j, \dots, j$ are eigenvalues of the J_z generator. Since the transformations are elements of the group we can write

$$\Theta_g^{\psi} = e^{ia_i T^i} \qquad \Theta_g^L = e^{ia_{L,i} T_L^i} \qquad \Theta_g^R = e^{ia_{R,i} T_R^i} \tag{1.87}$$

where T_R^i, T_L^i are the sets of left and right generators of the gauge transformations that satisfy (1.86) and

$$[T_R^i, T_L^i] = 0 \forall i = 1, 2, 3 \tag{1.88}$$

From this, it's possible to write the infinitesimal generator of the Gauss law:

$$\mathbf{G}_n = \sum_i \mathbf{T}_R^{i,n} + \sum_o \mathbf{T}_L^{i,n} + \mathbf{T}^n = 0 \qquad \mathbf{G}_n |\psi\rangle = 0 \qquad \forall n \tag{1.89}$$

On a lattice, SU(2) gauge fields are associated with link variables taking values in the group, while the electric field operators correspond to left- and right-invariant generators acting on each link.[6]. In particular, we have that the local electric field can be written in terms of the Casimir operators $E^2 = \mathbf{T}_R^2 = \mathbf{T}_L^2$ with eigenvalue $j(j + 1)$

$$E^2 |j, nm\rangle = j(j + 1) |j, nm\rangle . \tag{1.90}$$

A particularly relevant representation for $SU(2)$ is the fundamental representation, i.e. the representation with $j = \frac{1}{2}$ and dimension 2 that maps every element $g \in SU(2)$ to itself:

$$\rho^{\frac{1}{2}}(\bullet) : SU(2) \to SU(2) \;\; \rho^{\frac{1}{2}}(g) = g \;\; \forall g \in SU(2) \tag{1.91}$$

Matter is usually taken to be in the fundamental representation, i.e. the parallel transport $U^J = U^{\frac{1}{2}}$ carries internal sub-indices $s, s' \in (r, g)$, such that

$$U^{\frac{1}{2}}_{s,s'} |g\rangle = g_{s,s'} |g\rangle . \tag{1.92}$$

In the quantum simulation context, the infinite-dimensional local Hilbert space of SU(2) gauge fields is typically truncated to a finite set of representations, leading naturally to qudit realizations consistent with the framework introduced earlier [6]. In particular, the Hilbert space is often truncated in the representation basis to the values $j = 0, \frac{1}{2}$, which leads to only 5 possible states for each link. By labeling r, g (red,green) the sub-indices of the $j = \frac{1}{2}$ representation, the states are the following

$$|0\rangle, |rr\rangle, |rg\rangle, |gr\rangle, |gg\rangle \tag{1.93}$$

where we omitted the $j = \frac{1}{2}$ label for the colored states.

The non-Abelian nature of SU(2) gives rise to intrinsically richer dynamics compared to Abelian models, including nontrivial commutation relations between electric field components and genuinely interacting plaquette terms.

1.3.4 Dihedral LGTs

Dihedral lattice gauge theories are based on the *dihedral groups* D_N, which are discrete non-Abelian groups describing the symmetries ofregular N-sided polygons, including rotations and reflections [67]. Dihedral LGTs provide a natural intermediate step between Abelian $\mathbb{Z}_N$ gauge theories and continuous non-Abelian theories such as SU(2), making them particularly useful for testing constrained variational quantum algorithms and qudit encodings.

The dihedral group D_N consists of $2N$ elements:

$$D_N = \{r^k, r^k s \,|\, k = 0, \ldots, N-1\}, \tag{1.94}$$

where r represents a rotation by an angle $2\pi/N$ and s represents a reflection. The group multiplication rules are

$$r^N = e, \quad s^2 = e, \quad r^i r^j = r^{i+j}, \tag{1.95}$$

with e being the identity element [67]. These relations encode the non-Abelian nature of D_N for $N > 2$, since the rotation and reflection operations do not commute.

In the following we specifically talk about the case $N = 3$, which is the most studied for the rest of the chapter, and will mention the case $N = 4$ briefly in the last section.

Following what stated previously, we have $D_3 = \{e, r, r^2, s, sr, sr^2\}$ i.e. the group is composed of three rotations and three reflections. These elements are graphically represented in 1.7. In particular, considering that $r^3 = e$, each element can be constructed from (r, s) only.

The group is divided in three conjugacy classes: [3]

$$D_3 = \{e\} \cup \{r, r^2\} \cup \{s, sr, sr^2\}. \tag{1.96}$$

From representation theory of finite groups, we know that given a finite group G we have

$$|G| = \sum_j \dim(j)^2. \tag{1.97}$$

The only solution of this equation for $|G| = |D_3| = 6$ is to have two trivial $(1-d)$ representations a single two-dimensional irrep.

These are:

- Trivial representation $I(g) = D^I(g) = 1 \ \ \forall g$, $\dim(I) = 1$

[3] given a group G, a conjugacy class is given by $C = \{a | g^{-1} a g \in C \ \ \forall \ \ g \in G\}$

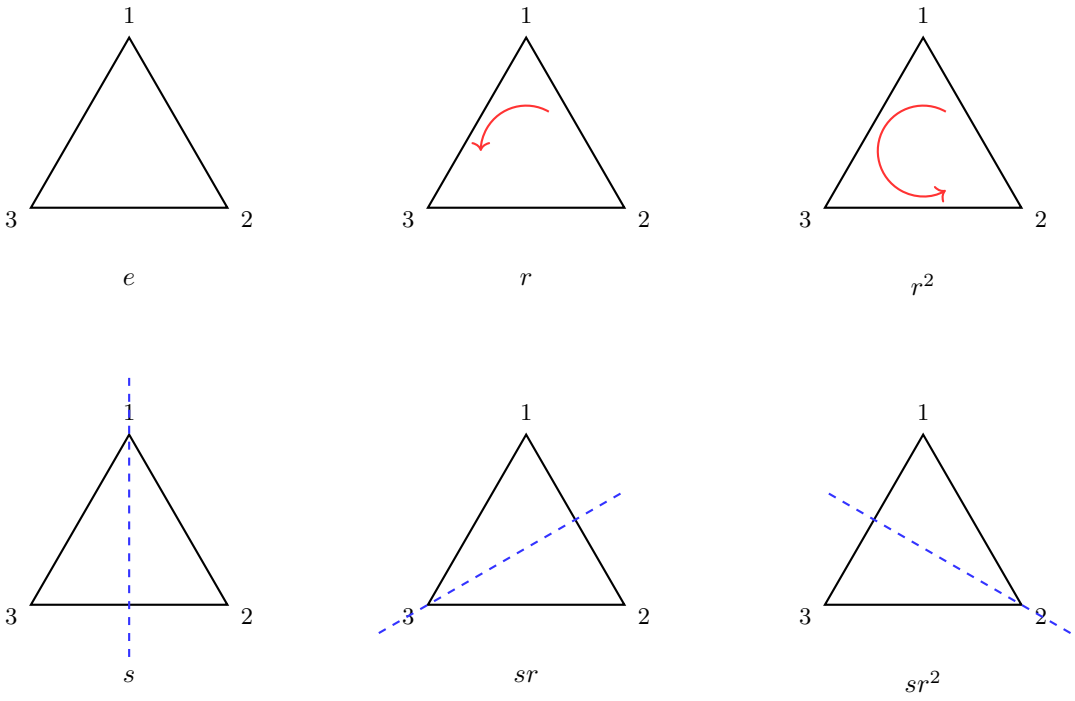

Figure 1.7: Graphical representation of the elements of the symmetry group D_3. On top, three rotations of respectively 0, $\frac{2\pi}{3}$, $\frac{4\pi}{3}$. On the bottom, three reflections across three axes.

- Parity representation $p(g) = D^p(g) = (-1)^s$, $\dim(p) = 1$
- Faithful representation $\tau(g) = D^\tau(g)$, $\dim(\tau) = 2$

where $D^\tau(r) = \begin{pmatrix} \zeta^+ & 0 \\ 0 & \zeta^- \end{pmatrix}$ and $D^\tau(s) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $\zeta^+ = e^{i\frac{2\pi}{3}}$ $\zeta^- = (\zeta^+)^*$and the other matrices can be constructed through group composition.

As mentioned before, states on the links represent the degrees of freedom of the gauge fields. In general, they can be expressed either in the group element basis $\{|g\rangle\}_{g=1}^{|G|}$ or in the representation basis $\{|Jlm\rangle\}$ [19].

For this group, we can explicitly write the matrix form of the Fourier transform:

$$Q \equiv QFT_{D_3} = \frac{1}{\sqrt{3}} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ 1 & \zeta^+ & \zeta^- & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & \zeta^+ & \zeta^- \\ 0 & 0 & 0 & 1 & \zeta^- & \zeta^+ \\ 1 & \zeta^- & \zeta^+ & 0 & 0 & 0 \end{pmatrix} \tag{1.98}$$

We consider fermions living in the faithful representation, i.e.

$$\psi = \begin{pmatrix} \psi_r \\ \psi_g \end{pmatrix} \qquad \Theta_g^\psi \psi \Theta_g^{\psi,\dagger} = D^\tau(g)\psi \qquad U = \sum_{g\in G} D^\tau(g)\,|g\rangle\langle g|\,. \tag{1.99}$$

Now that we have explicitly chosen a representation for the matter fields, we can explain better why we are interested specifically in this group.

We have two colors of Fermions and considering that each fermionic field $\psi_{r(g)}$ has dimension two, we have that the matter field is represented by a four dimensional vector. The matter fields can thus be represented with a qudit of dimension four. The local dimension of the link sites is equal to the cardinality of the group, thus the link variables(fields) are represented by a six element vector which corresponds a to a qudit of dimension six. For a D_3 LGT with fermionic matter in the faithful representation, the composite Hilbert space is then

$$\mathcal{H} = (\mathbb{C}^4)^{\otimes N_{\mathrm{v}}} \otimes (\mathbb{C}^6)^{\otimes N_{\mathrm{l}}} \qquad \dim(\mathcal{H}) = 4^{N_{\mathrm{v}}} 6^{N_{\mathrm{l}}} \tag{1.100}$$

where v and l are respectively the number of vertices and links in the lattice. This system can be efficiently represented in the degrees of freedom of the Ca_{40}^{+} trapped ion machine described in Sec. 1.1.7 , without requiring costly encodings of the theory's degrees of freedom into many qubits.

Relevant operators and states for D_3 theories

We have mentioned before that we take the matter to be in the faithful representation. This means that the mass field describes two colorful fermions, each with dimension two (the internal degrees of freedom of each color can be thought of as labeling if a fermion is present or not in a given state). The density operator for a an even (matter) site is then defined as

$$n_{\mathrm{even}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}. \tag{1.101}$$

Due to the staggering of the mass , the density operator on an odd (antimatter) site is defined as $\bar{n} = 2 - n$.

Given matter in the faithful representation, we define four orthogonal mass states based on occupation numbers such that $\langle i|n|i\rangle = i$:

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} \qquad |r\rangle = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \tag{1.102}$$

$$|g\rangle = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \qquad |rg\rangle = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} \tag{1.103}$$

These states transform in the following way under the generators of the group (r, s):

$$\theta(r)\,|0\rangle = |0\rangle \qquad \theta(s)\,|0\rangle = |0\rangle \tag{1.104}$$

$$\theta(r)\,|r\rangle = \zeta^{+}\,|r\rangle \qquad \theta(s)\,|r\rangle = |g\rangle \tag{1.105}$$

$$\theta(r)\,|g\rangle = \zeta^{-}\,|g\rangle \qquad \theta(s)\,|g\rangle = |r\rangle \tag{1.106}$$

$$\theta(r)\,|rg\rangle = |gr\rangle \qquad \theta(s)\,|rg\rangle = -\,|rg\rangle \tag{1.107}$$

Given the staggered nature of the mass, we identify the void and double occupied states in the following way:

$$|\emptyset\rangle_{\mathrm{M}} = |0\rangle \qquad |\emptyset\rangle_{\mathrm{AM}} = |rg\rangle \tag{1.108}$$

$$|2\rangle_{\mathrm{M}} = |rg\rangle \qquad |2\rangle_{\mathrm{AM}} = |0\rangle \tag{1.109}$$

$$\tag{1.110}$$

We can work similarly with the field states. We explicitly define the group element basis as the following set of states:

$$|e\rangle = \begin{pmatrix}1\\0\\0\\0\\0\\0\end{pmatrix} \qquad |r\rangle = \begin{pmatrix}0\\1\\0\\0\\0\\0\end{pmatrix} \tag{1.111}$$

$$|r^2\rangle = \begin{pmatrix}0\\0\\1\\0\\0\\0\end{pmatrix} \qquad |s\rangle = \begin{pmatrix}0\\0\\0\\1\\0\\0\end{pmatrix} \tag{1.112}$$

$$|sr\rangle = \begin{pmatrix}0\\0\\0\\0\\1\\0\end{pmatrix} \qquad |sr^2\rangle = \begin{pmatrix}0\\0\\0\\0\\0\\1\end{pmatrix} \tag{1.113}$$

With this definition, the representation basis corresponds to the rows of the Quantum Fourier transofrm Q:

$$|e\rangle = \frac{1}{\sqrt{6}}\begin{pmatrix}1\\1\\1\\1\\1\\1\end{pmatrix} \qquad |p\rangle = \frac{1}{\sqrt{6}}\begin{pmatrix}1\\1\\1\\-1\\-1\\-1\end{pmatrix}$$

$$|\tau, rr\rangle = \frac{1}{\sqrt{3}}\begin{pmatrix}1\\\xi^+\\\xi^-\\0\\0\\0\end{pmatrix} \qquad |\tau, rg\rangle = \frac{1}{\sqrt{3}}\begin{pmatrix}1\\\xi^-\\\xi^+\\0\\0\\0\end{pmatrix}$$

$$|\tau, gr\rangle = \frac{1}{\sqrt{3}}\begin{pmatrix}0\\0\\0\\1\\\xi^-\\\xi^+\end{pmatrix} \qquad |\tau, gg\rangle = \frac{1}{\sqrt{3}}\begin{pmatrix}0\\0\\0\\1\\\xi^+\\\xi^-\end{pmatrix}$$

These states transform (under left transformations) as follows:

$$\theta^L(r)\,|e\rangle = |e\rangle \qquad \theta^L(s)\,|e\rangle = |e\rangle \tag{1.114}$$
$$\theta^L(r)\,|p\rangle = |p\rangle \qquad \theta^L(s)\,|p\rangle = -\,|p\rangle \tag{1.115}$$
$$\theta^L(r)\,|\tau, rr\rangle = \zeta^+\,|\tau, rr\rangle \qquad \theta^L(s)\,|\tau, rr\rangle = |\tau, rr\rangle \tag{1.116}$$
$$\theta^L(r)\,|\tau, rg\rangle = \zeta^-\,|\tau, gr\rangle \qquad \theta^L(s)\,|\tau, rg\rangle = -\,|\tau, gr\rangle \tag{1.117}$$
$$\theta^L(r)\,|\tau, gr\rangle = \zeta^+\,|\tau, rg\rangle \qquad \theta^L(s)\,|\tau, gr\rangle = |\tau, rg\rangle \tag{1.118}$$
$$\theta^L(r)\,|\tau, gg\rangle = \zeta^-\,|\tau, gg\rangle \qquad \theta^L(s)\,|\tau, gg\rangle = -\,|\tau, gg\rangle \tag{1.119}$$

Notice how the representation states all transform according to the left internal sub-index on the representation. Similarly, in order to understand how the right transformation act on these basis states it's enough to have the state transform as the right internal label of the representation (on trivial and parity representation, the action is the same).

Having defined this basis, it's possible to give some explicit examples of gauge invariant states for a single link in position n on a 1-D lattice, along with their graphical representation. A complete list s given in App. D

$|e0e\rangle =$ 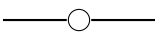
$|e2p\rangle =$ 
$|p0p\rangle =$ 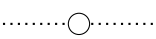
$|p2e\rangle =$ 
$|e\rangle \otimes (|r, rg\rangle + |g, gr\rangle) =$ 
$|p\rangle \otimes (|r, rg\rangle - |g, gr\rangle) =$ 
$|rr, 2, gg\rangle - |gg, 2, rr\rangle =$ 

where we omitted the τ label for the link representation states and only label them with sub indices. In general, the label $|\tau 1 \tau\rangle$ will be used to indicate a gauge invariant superposition of states with links in the fundamental representation and single occupation on the vertices of the lattice.

Kogut–Susskind Hamiltonian for a D_3 LGT

We now specialize the Kogut–Susskind Hamiltonian to the case of a lattice gauge theory with gauge group D_3 and matter transforming in the fundamental representation. Since the general form of the Hamiltonian has already been introduced in the previous section, here we explicitly write the magnetic and electric contributions for this specific discrete group.

For a plaquette composed of four links with group elements $g_1, g_2, g_3, g_4 \in G$, the magnetic contribution depends on the group flux through the plaquette, given by the ordered product

$$g_p = g_1 g_2 g_3^{-1} g_4^{-1}. \tag{1.120}$$

The magnetic term of the Hamiltonian can therefore be written as

$$2g^2 H_B = \sum_{g_1,g_2,g_3,g_4 \in G} \mathrm{Tr}\left[D(g_1 g_2 g_3^{-1} g_4^{-1})\right] \left|g_1 g_2 g_3^{-1} g_4^{-1}\right\rangle\!\!\left\langle g_1 g_2 g_3^{-1} g_4^{-1}\right| + H.C. \tag{1.121}$$
$$= \sum_{g_1,g_2,g_3,g_4 \in G} \chi(g_1 g_2 g_3^{-1} g_4^{-1}) \left|g_1 g_2 g_3^{-1} g_4^{-1}\right\rangle\!\!\left\langle g_1 g_2 g_3^{-1} g_4^{-1}\right| + H.C., \tag{1.122}$$

where $D(g)$ denotes the fundamental representation of the group element g and $\chi(g) = \mathrm{Tr}[D(g)]$ is the corresponding character.

This expression shows that, in the group element basis, the plaquette operator is diagonal and its eigenvalues are given by the character of the resulting plaquette flux in the fundamental representation.

For the group D_3, the characters evaluated on the three conjugacy classes in the fundamental representation are

$$\chi(e) = 2 \qquad \chi(\{r, r^2\}) = 2\cos\left(\frac{2\pi}{3}\right) \qquad \chi(\{s, sr, sr^2\}) = 0 \tag{1.123}$$

For discrete groups, it is convenient to express the electric contribution in terms of projectors onto irreducible representations. The projector onto the representation J can be written as [93]

$$\Pi_J = \sum_{g\in G} \frac{\dim(J)}{|G|} \chi^J(g)\, \Theta_g^L. \tag{1.124}$$

The electric energy operator can then be written as a weighted sum of these projectors,

$$E^2 = \sum_J \alpha(J) \sum_{g\in G} \frac{\dim(J)}{|G|} \chi^J(g)\, \Theta_g^L. \tag{1.125}$$

Since D_3 (and more generally D_N) is a discrete group, there is no underlying Lie algebra structure fixing the spectrum of the electric field. Consequently, the coefficients $\alpha(J)$ appearing in Eq. (1.71) are not determined by the group structure and can be chosen freely.

This freedom will play an important role in Chapter 5.3, where different choices of $\alpha(J)$ will be exploited to explore the phase diagram of the model.

The mass term is given by the staggered local densities

$$H_M = \sum_{\mathrm{x}} (-1)^{\mathrm{x}} n_{\mathrm{x}} \tag{1.126}$$

which counts he number of fermions or antifermions for each site.

Chapter 2

Inequality constraints in QAOA with qudits

As mentioned in the introduction, many optimization problems are subject to constraints. Some of these are particularly difficult to include in optimisation algorithms, and qudits can help with the required gates decompositions.

This chapter addresses the problem of utilizing qudits to enforce inequality constraints in quantum optimisation algorithms, specifically QAOA. The material presented here is largely based on [24]. Two main strategies are explored: first, a generalization of the concept of slack variables [56]; second, a direct energy penalty applied to unfeasible states, which requires no ad-hoc auxiliary variables but only a single ancilla qudit.

2.1 Theoretical Description

The standard approach for introducing inequality constraints into quantum optimisation is via slack variables [94, 66].

Consider a problem defined through a set of R inequality constraints of the form

$$P_r(\boldsymbol{x}) \leq 0 \qquad \forall \quad r = 1, \ldots, R. \tag{2.1}$$

The function $P_r(\boldsymbol{x})$ is assumed to take N_r possible values, i.e., $P_r(\boldsymbol{x}) \in \{p_1, \ldots, p_{N_r}\}$ for all $\boldsymbol{x}$. Of these values, the first N_r^{feas} are non-positive (feasible solutions) and the remaining $N_r - N_r^{\text{feas}}$ are strictly positive (unfeasible solutions).

Equality constraints can be incorporated directly as quadratic penalty terms in the cost function [95], preserving the QUBO form when the constraints are linear. Inequality constraints, however, cannot be treated in this way directly. The standard remedy is to transform each inequality constraint into an equality constraint through the introduction of a slack variable s_r,

$$P_r(\boldsymbol{x}) \leq 0 \;\Leftrightarrow\; P_r(\boldsymbol{x}) + s_r = 0\,. \tag{2.2}$$

The slack variables are restricted to non-negative values ($s_r \geq 0$), and their range is determined by the values attained by the constraint function for feasible configurations: a valid s_r satisfying Eq. (2.2) exists if and only if the constraint is fulfilled. Technically, $s_r \in \{-P_r(\boldsymbol{x}) : \boldsymbol{x} \text{ feasible}\}$, so each s_r can assume N_r^{feas} distinct values.

The transformed equality constraints of Eq. (2.2) are then added as quadratic penalty terms to the cost function, yielding the penalised cost function

$$C_{\text{slack}}(\boldsymbol{x}) = C(\boldsymbol{x}) + \sum_{r=1}^{R} \lambda_r \big(P_r(\boldsymbol{x}) + s_r\big)^2 , \tag{2.3}$$

where $\lambda_r > 0$ is the penalty factor for the r-th constraint. In the standard QUBO formulation each slack variable must be encoded using multiple binary variables, requiring at least $\sum_r \log_2 N_r^{\text{feas}}$ auxiliary binary variables [94, 56]. Since N_r^{feas} is typically larger than two, the associated spatial resource overhead is considerable, particularly on current NISQ devices.

This overhead can be reduced by representing the slack variables with qudits. When the qudit dimension d matches N_r^{feas}, the computational basis is extended by one qudit state per constraint (i.e. per slack variable),

$$|\boldsymbol{x}\rangle \;\Rightarrow\; |\boldsymbol{x}, \{s_1, \ldots, s_R\}\rangle \equiv |\boldsymbol{x}, \boldsymbol{s}\rangle \; . \tag{2.4}$$

In the quantum formulation the constraint function and the slack variable are represented by operators with the appropriate eigenvalues,

$$\big(\hat{P}_r + \hat{S}_r\big)\, |\boldsymbol{x}, \boldsymbol{s}\rangle = \big(P_r(\boldsymbol{x}) + s_r\big)\, |\boldsymbol{x}, \boldsymbol{s}\rangle \; . \tag{2.5}$$

A feasible configuration of the search variables $\boldsymbol{x}$ is identified by locating the basis state in which the qudit slack variable takes the value that yields a zero eigenvalue, i.e. $P_r(\boldsymbol{x}) + s_r = 0$.

It is worth noting that the number of additional dimensions introduced by the auxiliary qudits is proportional to N_r^{feas}: the more feasible configurations exist, the larger the required auxiliary Hilbert space. The dimension of the full Hilbert space including the slack qudits grows as

$$\dim(\mathcal{H}_{\text{slack}}) = 2^N \prod_{r=1}^{R} N_r^{\text{feas}}, \qquad \mathcal{H}_{\text{slack}} = \mathcal{H}_C \otimes \prod_{r=1}^{R} \mathbb{C}^{N_r^{\text{feas}}} . \tag{2.6}$$

For a given feasible configuration $|\boldsymbol{x}\rangle$, however, only one state of the extended Hilbert space, $|\boldsymbol{x}, \boldsymbol{s}\rangle$, is feasible (the one with eigenvalue $s_r = -P_r(\boldsymbol{x})$), while the majority of the added quantum states are unfeasible. Consequently, for lightly constrained problems (large N_r^{feas}), the effective optimisation problem is harder because the fraction of feasible solutions decreases by a factor $1/N_r^{\text{feas}}$ for each slack variable included.

The final constrained Hamiltonian takes the form

$$H_{\text{slack}} = H_C + \sum_{r=1}^{R} \lambda_r \big(\hat{P}_r + \hat{S}_r\big)^2 . \tag{2.7}$$

The quadratic form of the penalty terms guarantees that they vanish for all feasible configurations.

The inclusion of qudit slack variables requires a modification of the QAOA mixer. Explicitly, the mixer adopted here is

$$H_M = \beta \left(\sum_{i=1}^{N} \sigma_x^i + \sum_{r=1}^{R} L_x^{s_r} \right) + \gamma \sum_{r=1}^{R} \big(L_z^s\big)^2 , \tag{2.8}$$

where β and γ are variational parameters associated with the x- and z-components of the mixer, respectively. The second term is necessary because of the higher dimensionality of

the qudit slack variables: as explained in Eq. (1.61), the additional $(L_z^s)^2$ term ensures universality, i.e. the ability to reach all qudit states. Beyond this modification to the mixing operator, the QAOA procedure follows the description of Sec. 1.2.3 with the cost Hamiltonian given by Eq. (2.7).

An alternative formulation that avoids the drawbacks of slack variables is a penalised Hamiltonian of the form

$$H_{\text{penal}} = H_C + \sum_{r=1}^{R} \lambda_r \hat{G}_r \,, \tag{2.9}$$

where $\lambda_r > 0$ are penalty factors and the operators $\hat{G}_r$ have eigenvalues determined by the corresponding constraint function (Eq. (2.1)),

$$\hat{G}_r \,|\boldsymbol{x}\rangle = g\big(P_r(\boldsymbol{x})\big)\,|\boldsymbol{x}\rangle \,. \tag{2.10}$$

An explicit construction of $\hat{G}$ is given below. The function $g(\cdot)$ is a penalty function in the classical-optimisation sense: it increases the energy of unfeasible solutions while leaving the energy of feasible solutions unchanged. Formally, $g(P_r(\boldsymbol{x})) = 0$ when $P_r(\boldsymbol{x}) \leq 0$ (feasible) and $g(P_r(\boldsymbol{x})) > 0$ when $P_r(\boldsymbol{x}) > 0$ (unfeasible). In order to facilitate a gradient toward the feasible subspace, $g(\cdot)$ is moreover required to be non-decreasing for positive arguments. A typical choice satisfying these requirements is [96]

$$g(y) = y^a\,\Theta(y)\,, \tag{2.11}$$

where $\Theta(y)$ is the Heaviside step function and $a \geq 0$ is a freely chosen exponent. Examples of such penalty functions are illustrated in Fig. 2.1.

Including these penalties, the unitary operator generating the trial wave function becomes

$$U(\boldsymbol{\alpha},\boldsymbol{\beta}) = \prod_{l=1}^{p} e^{i\beta_l H_M}\, e^{i\alpha_l H_{\text{penal}}}\,, \tag{2.12}$$

$$e^{i\alpha_l H_{\text{penal}}} = e^{i\alpha_l \left(H_C + \sum_{r=1}^{R} \lambda_r \hat{G}_r\right)}\,. \tag{2.13}$$

The rotation generated by the penalising Hamiltonian imprints phases of order $\alpha_l\lambda_r$ onto the unfeasible states, enabling the QAOA procedure to suppress them.

Since the constraint terms are diagonal in the computational basis, efficient constructions for implementing $e^{i\alpha_l \sum_{r=1}^{R} \lambda_r \hat{G}_r}$ exist via Fourier analysis or ancilla registers [97, 98]. These methods operate purely within qubit systems, and their efficiency depends on the problem type and instance as well as on the target hardware platform.

Here, an efficient implementation of diagonal unitaries of the form of Eqs. (2.12)–(2.13) is proposed for constraint functions $P_r(\boldsymbol{x})$ that depend only on the Hamming weight, i.e. the magnetisation m_r of a subset $\mathcal{I}_r$ of problem qubits,

$$P_r(\boldsymbol{x}) = P_r\big(m_r(\boldsymbol{x})\big)\,, \tag{2.14}$$

$$m_r(\boldsymbol{x}) = \sum_{x_i \in \mathcal{I}_r} x_i\,, \tag{2.15}$$

where $\mathcal{I}_r$ denotes a subset of the N qubits. Generalisation to summations of x_i and $\overline{x}_i$ is straightforward.

In contrast to the slack-variable approach, the ancilla qudit does not enter the cost function but serves solely to imprint appropriate phases according to Eq. (2.13) onto the

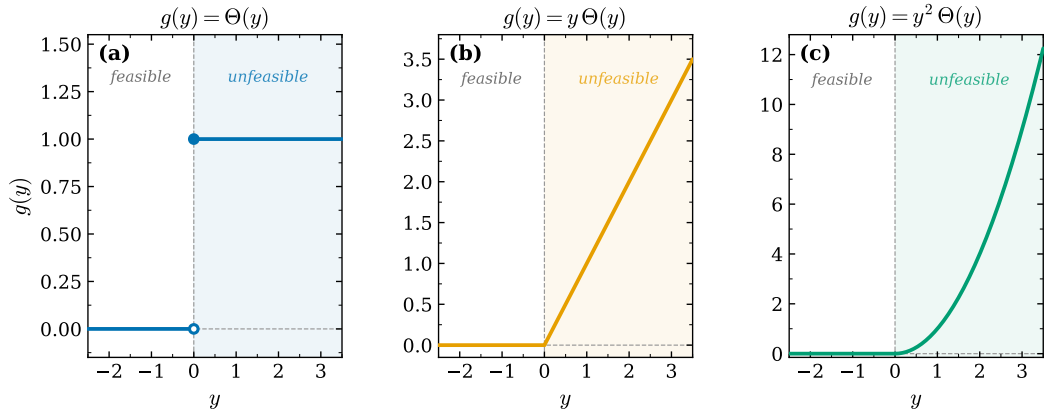

Figure 2.1: Penalty functions $g(y) = y^a\,\Theta(y)$ for the three choices of exponent considered in this work: (a) flat penalty ($a = 0$), (b) linear penalty ($a = 1$), and (c) quadratic penalty ($a = 2$). Here $\Theta(y)$ denotes the Heaviside step function and $y \equiv P_r(\boldsymbol{x})$ is the constraint function evaluated at a candidate solution $\boldsymbol{x}$. The shaded region ($y > 0$) identifies unfeasible configurations, for which $g(y) > 0$ and an energy penalty is applied. In the feasible region ($y \leq 0$) the penalty vanishes identically, leaving the cost function unaffected. All three forms are non-decreasing for positive arguments, which ensures a gradient toward the feasible subspace during the QAOA optimisation. The open and filled circles in panel (a) indicate the jump discontinuity of the flat penalty at $y = 0$.

trial wave function. A key advantage of this method is that the dimension of the Hilbert space in which solutions are searched does not increase. In particular, the fraction of feasible solutions remains constant regardless of the number of constraints — a sharp contrast to the slack-variable approach, where each constraint inflates the Hilbert space by a factor N_r^{feas} and, more critically, increases the number of infeasible solutions by a factor $(N_r^{\text{feas}} - 1)$, since only one configuration of the slack variable per constraint is feasible. This constitutes the main bottleneck of the slack-variable approach.

The bottleneck of the proposed procedure is the dimensionality of the ancilla qudit, which must be at least $N + 1$ for a constraint on the magnetisation of N qubits. Since the same ancilla qudit can be reused for multiple constraints, the method is expected to be particularly advantageous when many constraints are present simultaneously, each involving only a restricted number of problem qubits and a small-to-intermediate set of possible values P_r.

For concreteness, consider a single constraint of the form of Eq. (2.14) depending on the total Hamming weight (magnetisation),

$$P_r(\boldsymbol{x}) \equiv P_r\big(m(\boldsymbol{x})\big), \qquad m(\boldsymbol{x}) = \sum_{i=1}^{N} x_i\,, \tag{2.16}$$

of an N-qubit basis state $|\boldsymbol{x}\rangle = |x_1, x_2, \ldots, x_N\rangle$ with $x_i \in \{0, 1\}$. The constraint label r is dropped for notational convenience; the generalisation to multiple constraints, each acting on a subset of qubits, is straightforward. The penalty function $P(m)$ specifies which values of m are feasible ($P(m) \leq 0$) and which are unfeasible ($P(m) > 0$). The set of unfeasible values is denoted $\mathcal{I} = \{m : P(m) > 0\}$ and its cardinality $N_{\mathcal{I}} = |\mathcal{I}|$.

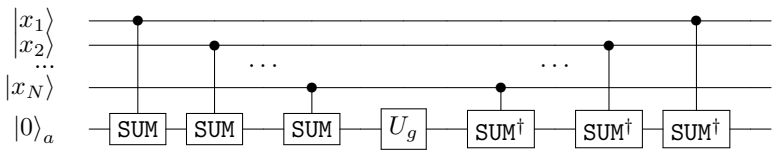

Figure 2.2: Circuit showing how an arbitrary phase $g(m)$ as function of the magnetization $m(\boldsymbol{x}) = \sum_i x_i$ can be applied on the qubit register using one ancilla qudit $|y\rangle_a$.

For an N-qubit state the Hamming weight takes the $N+1$ integer values $m \in \{0, 1, \ldots, N\}$, fixing the ancilla qudit dimension to $N + 1$,

$$|y\rangle_a , \qquad y \in \{0, 1, \ldots, N\} . \tag{2.17}$$

The corresponding circuit is depicted in Fig. 2.2.

The procedure begins from a product state of the problem qubits (in an arbitrary superposition) and the ancilla qudit initialised to $|0\rangle_a$,

$$|\psi\rangle |0\rangle_a = \sum_{\boldsymbol{x}} c_{\boldsymbol{x}} |\boldsymbol{x}\rangle |0\rangle_a . \tag{2.18}$$

A SUM gate is then applied, adding the Boolean value of each qubit sequentially onto the ancilla, so that after all N applications the ancilla holds the value $m = m(\boldsymbol{x}) = \sum_i x_i$,

$$\prod_i \mathtt{SUM}_{i\to a} |\psi\rangle |0\rangle_a = \sum_{\boldsymbol{x}} c_{\boldsymbol{x}} |\boldsymbol{x}\rangle |m\rangle_a . \tag{2.19}$$

A phase-shift unitary U_g is then applied to the ancilla qudit, adding a phase only to unfeasible states according to the penalty function $g(\cdot)$ of Eq. (2.11): $U_g |m\rangle_a = e^{ig(P(m))} |m\rangle_a$. This operation is straightforwardly realised on quantum devices, as it requires only the ability to apply a different phase selectively to each level of a single qudit. The state becomes

$$\sum_{\boldsymbol{x}} c_{\boldsymbol{x}} |\boldsymbol{x}\rangle U_g |m\rangle_a = \sum_{\substack{\boldsymbol{x} \\ m(\boldsymbol{x})\notin\mathcal{I}}} c_{\boldsymbol{x}} |\boldsymbol{x}\rangle |m\rangle_a + \sum_{\substack{\boldsymbol{x} \\ m(\boldsymbol{x})\in\mathcal{I}}} c_{\boldsymbol{x}} e^{ig(P(m))} |\boldsymbol{x}\rangle |m\rangle_a . \tag{2.20}$$

After undoing the SUM gates the final state is

$$\left(\sum_{\substack{\boldsymbol{x} \\ m(\boldsymbol{x})\notin\mathcal{I}}} c_{\boldsymbol{x}} |\boldsymbol{x}\rangle + \sum_{\substack{\boldsymbol{x} \\ m(\boldsymbol{x})\in\mathcal{I}}} c_{\boldsymbol{x}} e^{ig(P(m))} |\boldsymbol{x}\rangle \right) |0\rangle_a , \tag{2.21}$$

which realises the desired unitary: a phase shift parameterised by $g(\cdot)$ is applied exclusively to the infeasible basis states, as indicated by $P(m)$.

As a concrete example, consider the constraint that the magnetisation must not exceed a given value m_0, i.e. $P(m) = m - m_0 \leq 0$. With the penalty function of Eq. (2.11), the conditional phases are

$$g(P(m)) = \Theta(m - m_0)\,(m - m_0)^a , \tag{2.22}$$

and the final state becomes

$$\left(\sum_{\substack{\boldsymbol{x} \\ m(\boldsymbol{x}) \notin \mathcal{I}}} c_{\boldsymbol{x}} \left| \boldsymbol{x} \right\rangle + \sum_{\substack{\boldsymbol{x} \\ m(\boldsymbol{x}) \in \mathcal{I}}} c_{\boldsymbol{x}} e^{i(m(\boldsymbol{x})-m_0)^a} \left| \boldsymbol{x} \right\rangle \right) \left| 0 \right\rangle_a \,, \tag{2.23}$$

adding a phase to all states with $m \in \mathcal{I} = \{m_0 + 1, m_0 + 2, \ldots, N\}$.

The entire procedure for arbitrary $g(P(m))$ requires only $2N$ qudit-controlled `SUM` gates and a single qudit phase-shift unitary U_g. The sole hardware requirement is the availability of qudits with at least $N+1$ levels, where N is the number of qubits involved in the constraint.

The procedure becomes particularly advantageous when many constraints must be implemented simultaneously . Each constraint is handled sequentially by running the circuit above on the relevant subset of qubits with the corresponding constraint function. A single ancilla qudit suffices, as it can be reused across all constraints; employing multiple ancillas allows parallelisation.

The approach can also be generalised to penalties depending on functions of the form

$$m^*(\boldsymbol{x}) = \sum_{i \in N} x_i + \sum_{i \in \bar{N}} \bar{x}_i \,, \tag{2.24}$$

where N and $\bar{N}$ are two distinct subsets of the system qubits and $\bar{x}_i$ is the Boolean complement of x_i. This requires only the insertion of a π-rotation (σ_x) for each qubit in $\bar{N}$ before (and after) the `SUM`-gate (`SUM`†-gate).

2.2 Numerical Tests

In this section, numerical results for QAOA are presented for several scenarios in which inequality constraints play a role. First, the ability of QAOA to obtain feasible solutions for a generic random spin model [99] is analysed. Subsequently, the construction of a constrained initial state for warm-starting the QAOA procedure is examined. Finally, the performance on the industrially relevant problem of electric vehicle (EV) charging [100], which is subject to multiple constraints, is evaluated. In all cases, the constraint-handling methods described in the previous section are compared.

The quantum part of the QAOA procedure is simulated using exact state-vector simulations. The variational parameters defining the trial wave functions are updated using the Powell classical optimiser from the `scipy` library [101] with default options, which runs until the predefined convergence criteria are met. For each problem instance, $N_{\text{runs}} = 50$ executions of QAOA are performed as described in Sec. 1.2.3, each with initial parameters drawn uniformly from $[0, 2\pi]$. A single execution of QAOA for a given setup and initial parameter set is referred to as a *run*. During the optimisation loop, the cost function is evaluated as the expectation value of the cost Hamiltonian over the parametrised state. For each run, classical solutions are extracted by sampling $N_S = 64$ solutions (measurement shots) from the final optimised quantum state $|\psi(\alpha^*, \beta^*)\rangle$.

A finite number of shots was also examined for estimating the cost function expectation value during optimisation. The results are quantitatively equivalent to those obtained with exact expectation values (see, e.g., Fig. 2.3); however, since the associated simulation runtimes are substantially longer, exact expectation values are used throughout the numerical experiments reported here.

2.2.1 Metrics

Several figures of merit are employed to assess the performance of the QAOA protocols under different constraint-handling techniques.

The first metric is the *approximation ratio*,

$$R = \min_{\text{samples } s} \frac{E_s - E_0}{|E_0|}\,, \tag{2.25}$$

where E_s ($s = 1, \ldots, N_S$) is the energy of a state sampled from the final quantum state and E_0 is the energy of the optimal (constrained) solution. This metric is relevant for industrial optimisation, where the goal is the lowest-energy single configuration that can be obtained.

The second metric is the *success rate*,

$$r = \frac{\sum_{i=1}^{N_{\text{runs}}} X_i}{N_{\text{runs}}}\,, \tag{2.26}$$

where $X_i = 1$ if at least one of the (possibly degenerate) optimal states is sampled among the N_S samples of run i, and $X_i = 0$ otherwise.

The third metric is the *total weight of feasible solutions* in the final state,

$$W = \sum_{\boldsymbol{x} \in \text{feasible}} \left| c_{\boldsymbol{x}}(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*) \right|^2 , \tag{2.27}$$

where $c_{\boldsymbol{x}}(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*)$ are the amplitudes in the final QAOA state $|\psi(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*)\rangle = \sum_{\boldsymbol{x}} c_{\boldsymbol{x}}(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*)\,|\boldsymbol{x}\rangle$. This metric indicates the efficiency of the algorithm in concentrating probability weight on feasible solutions.

2.2.2 Random Spin Hamiltonian

The first benchmark consists of finding the ground state of a model of randomly interacting spin-$\frac{1}{2}$ particles in a random longitudinal field, comparing the qudit slack-variable approach with the direct penalty approach. The cost Hamiltonian is

$$H_C = \sum_{i=1}^{N} h_i \sigma_i^z + \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z\,, \tag{2.28}$$

with independent Gaussian-distributed parameters $J_{ij}, h_i \sim \mathcal{N}(0, 1)$.

A single constraint acting on all spins is considered,

$$P(\boldsymbol{\sigma}_z) = \sum_{i=1}^{N} \sigma_i^z - m_0 \equiv S^z_{\text{tot}} - m_0 \leq 0\,, \tag{2.29}$$

which selects all states with total spin $S^z_{\text{tot}} = \sum_{i=1}^{N} \sigma_i^z$ below a target value $m_0 \in \{-\frac{N}{2}, \ldots, \frac{N}{2}\}$.

Three choices of the exponent a are considered for the direct penalty approach, giving

$$H_{a=0} = \lambda\, \Theta(S^z_{\text{tot}} - m_0)\,, \tag{2.30a}$$

$$H_{a=1} = \lambda\, \Theta(S^z_{\text{tot}} - m_0)\,(S^z_{\text{tot}} - m_0)\,, \tag{2.30b}$$

$$H_{a=2} = \lambda\, \Theta(S^z_{\text{tot}} - m_0)\,(S^z_{\text{tot}} - m_0)^2\,. \tag{2.30c}$$

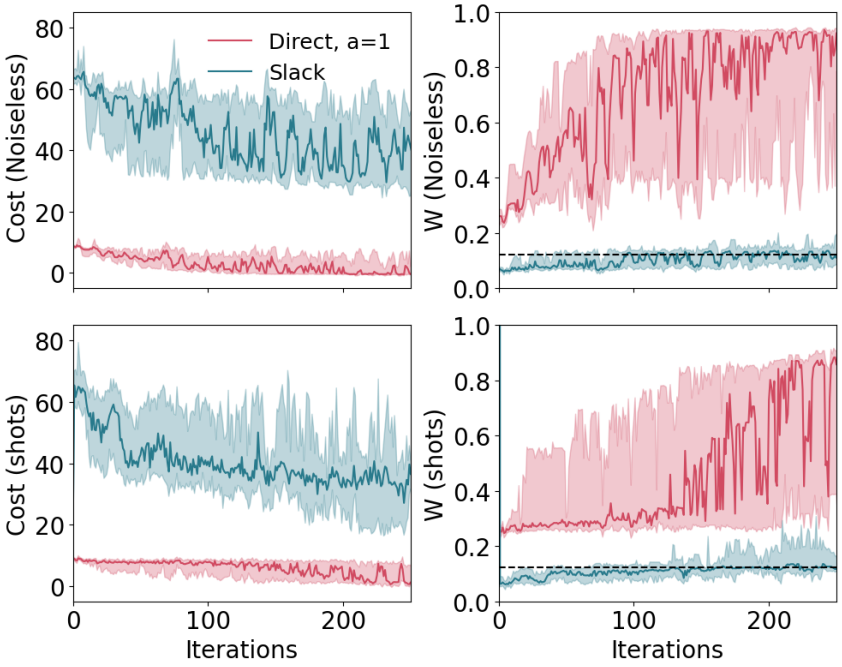

Figure 2.3: Convergence of the optimised cost $\langle H_C \rangle$ (left column) and feasible weight W (right column) over optimisation steps for a three-layer QAOA with $m_0 = -1.5$ and $N = 9$ qubits. Top row: exact state-vector evaluation of the cost function. Bottom row: noisy cost estimation with a finite number of $N_{\text{shots}} = 128$ samples per iteration. Red curves correspond to the direct penalty implementation with $a = 1$; blue curves use slack variables. Results are obtained from $N_{\text{runs}} = 20$ optimisations for one specific cost Hamiltonian. Thick lines indicate the median; shaded regions denote the 25th–75th percentiles. The dashed black line indicates the baseline probability of feasible states, i.e. the fraction of feasible over total states.

For the slack-variable approach, the complete cost Hamiltonian is

$$H_{\text{C,slack}} = H_C + \lambda \, (S^z_{\text{tot}} - m_0 + \hat{S})^2 \, , \tag{2.31}$$

where $\hat{S}$ acts on a qudit of dimension $d = \frac{N}{2} + 1 + m_0$.

The analysis begins with a single instance of the Hamiltonian of Eq. (2.28) for $N = 9$ qubits, examining the convergence of the optimiser for different setups..

Figure 2.3 compares the convergence obtained with exact state-vector evaluation of the cost function (upper row) against estimation from $N_{\text{shots}} = 128$ samples per iteration (lower row). Even with a relatively low shot count, the optimiser converges within the same number of iterations for both evaluation methods. The same behaviour was observed for all problems considered in this work; consequently, the remaining simulations are restricted to exact state-vector evaluations.

Figure 2.3 also juxtaposes the proposed direct penalty method (red) with the standard slack-variable approach (blue). Both approaches reach convergence within the displayed iterations; however, the slack-variable approach yields cost-function values an order of magnitude larger than the proposed approach, indicating a failure to identify good low-energy

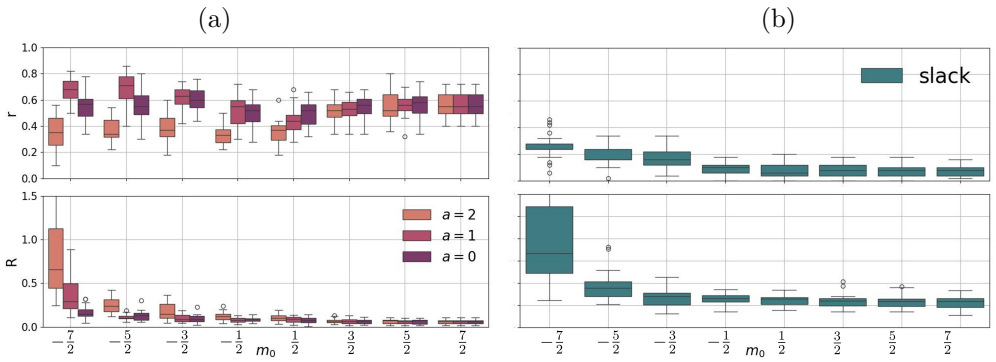

Figure 2.4: Success rate r (top row) and approximation ratio R (bottom row) for constrained QAOA with a single layer ($p = 1$) and $N = 9$ qubits. The box plots show statistics averaged over 20 different instances of the random Hamiltonian of Eq. (2.28) and 50 runs per instance, as a function of the maximum allowed total spin m_0. (a) Results for the proposed approach implementing the penalty functions of Eqs. (2.30a)–(2.30c). The linear form ($a = 1$) consistently achieves the highest success rate r, while the flat penalty yields the best approximation ratio R. The differences between penalty forms diminish as m_0 increases, since the problem becomes progressively less constrained. (b) Results for the slack-variable approach. The direct penalty approaches in (a) clearly outperform the slack-variable approach, achieving markedly higher success rates and lower approximation ratios.

solutions. Correspondingly, the feasible-state weight for the proposed approach converges above 90%, whereas the slack-variable approach converges at only around 10%. These results provide a first, clear indication that the proposed approach produces far more feasible solutions with significantly lower cost values.

The behavior is further analysed for a system of $N = 9$ qubits by studying the metrics defined in Sec. 2.2.1 for a $p = 1$ layer QAOA as a function of m_0. Twenty random realisations of the cost Hamiltonian are generated and $N_{\text{runs}} = 50$ simulations are performed for each. All results shown are obtained with a penalty factor $\lambda = 4$; higher values of λ yield qualitatively similar results.

The metrics of Eqs. (2.25), (2.26), and (2.27) are reported in Figs. 2.4 and 2.5 for increasing feasible subspace dimension. Across all metrics, the proposed approach consistently outperforms the slack-variable implementation. For example, the success rate r for the direct penalty functions is always significantly higher; as m_0 approaches its maximum value, the proposed methods converge to $r \gtrsim 0.5$, whereas the slack-variable approach converges to $r \approx 0.1$. Similarly, the approximation ratio R for the slack-variable approach fluctuates near 0.3 for $m_0 \gtrsim -2.5$, while for direct penalisation R rapidly approaches 0 as the problem becomes less constrained.

Among the different penalty forms, the linear penalty ($a = 1$) delivers the best success rate r. For highly constrained problems ($m_0 \lesssim -1.5$), the flat penalty ($a = 0$) achieves the best approximation ratio R. This latter observation is a direct consequence of the penalty terms themselves: for highly constrained problems, some sampled low-energy states are unfeasible, and linear or quadratic penalties naturally assign these states higher energies, thereby worsening the raw approximation ratio metric.

Figure 2.5 shows the total weight of the feasible states as a function of m_0. Solid lines indicate the baseline weight for a uniformly distributed quantum state, equal to the ratio

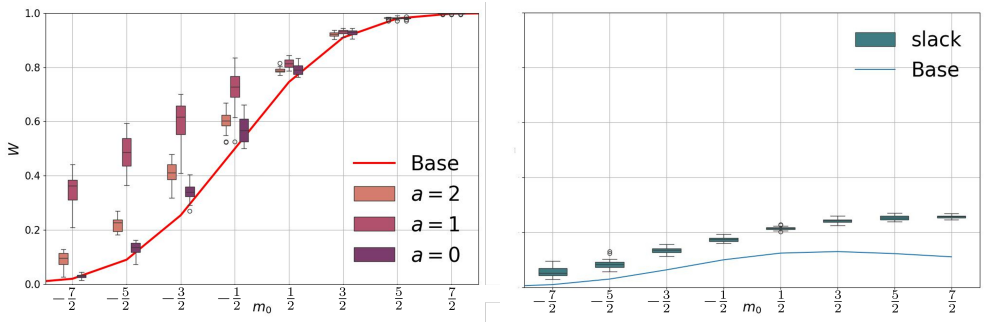

Figure 2.5: Total weight W of all feasible states for single-layer QAOA with the constrained cost functions of Eqs. (2.30a)–(2.30c) (left) and the slack-variable approach of Eq. (2.31) (right), for the same setup as Fig. 2.4. Solid lines show the baseline probability of finding a feasible state from an equal superposition of all basis states. The slack-variable approach consistently yields a lower probability of sampling a feasible configuration from the final state; even when almost all qubit configurations are feasible (large m_0), it saturates around 25%, while the direct penalty method reaches 100%. This trend mirrors Fig. 2.4 and serves as a further indicator of the limitations of the slack-variable approach for inequality-constrained optimisation.

of feasible to total states. For the proposed approach, the baseline increases monotonically from $1/2^N$ at $m_0 = -\frac{N}{2}$ toward 1 for the unconstrained problem at $m_0 = \frac{N}{2}$. For the slack-variable approach, the baseline does not approach unity with increasing m_0; instead, it peaks at an intermediate value $m_0 \approx 1.5$ and slightly decreases for larger m_0, due to the growing dimension of the slack variable with m_0. All methods tested increase the probability of sampling a feasible state above the baseline; however, the proposed direct penalty methods offer substantially larger improvements, particularly for the linear penalty and at larger m_0.

2.2.3 Constrained State Preparation and Sampling of Feasible States

In certain scenarios the goal of the algorithm is to sample from the feasible subspace without preference among the states within that subspace. Relevant examples include topological models and lattice gauge theories (LGTs) [102], spin ice [103], and the sampling of polymer melts [104, 105]. A related scenario arises when a superposition of all (or many) feasible states is required as the initial state for approaches that employ constraint-preserving mixing operators [106, 107, 108].

Motivated by these considerations, a QAOA is examined in which the cost function consists solely of the penalty terms encoding the constraint of Eq. (2.29). The cost Hamiltonian then takes the form of Eqs. (2.30a)–(2.30c) for the proposed approach and of Eq. (2.31) for the slack-variable approach. The QAOA is run for up to $p = 5$ layers with $N_{\text{runs}} = 50$ per setup (penalty type, number of layers, and constraint value m_0).

Figure 2.6 shows the total weight of feasible configurations in the final quantum state for $p = 1$, 3, and 5 layers. The linear penalty term substantially outperforms both the slack-variable approach and the other penalty forms, reaching close to $W = 1$ already after a single layer for nearly all values of m_0. As in the preceding example, all approaches improve over the baseline given by the weight of the feasible states in the initial equal-superposition state (dashed line). The proposed direct penalty approach, however, offers

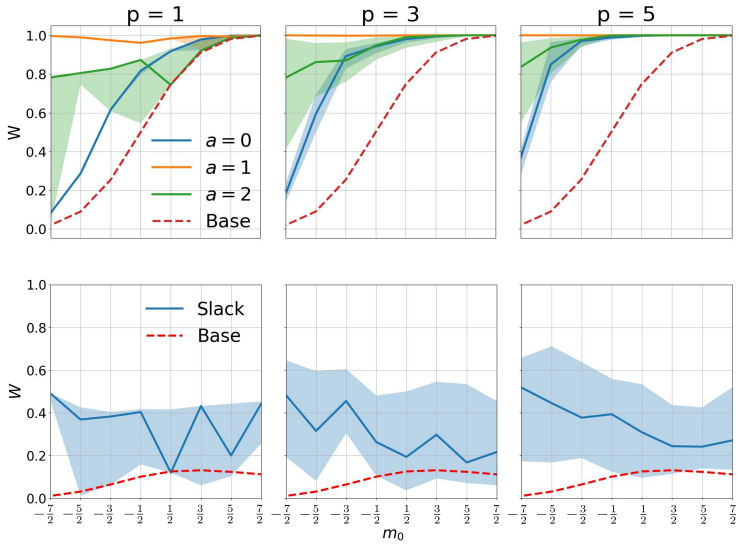

Figure 2.6: Probability W of sampling a feasible state from the final QAOA state with a constraint-only cost function, for the direct penalty approach (top row) and the slack-variable approach (bottom row), for $p = 1$, 3, and 5 QAOA layers (left to right panels). Solid lines show median values; shaded areas indicate the 20th–80th percentiles over 50 runs. Dashed red lines labelled 'Base' denote the baseline probabilities from the equal-superposition state. The direct penalty functions substantially outperform the slack-variable approach. In particular, the penalty with $a = 1$ achieves close to 100% probability for all constraint values already at $p = 1$.

far larger improvements over the baseline than the slack-variable approach. Furthermore, increasing the number of layers p consistently improves the performance of the direct penalty approach, a trend not observed for the slack-variable method.

2.2.4 EV Charging Problem

To connect the proposed approach to QAOA applied to industrially relevant combinatorial problems and to test its behaviour under multiple constraints, its performance is evaluated on an electric vehicle (EV) charging problem (schematically illustrated in Fig. 2.7 [100]. The objective is to find the optimal charging schedule for a fleet of EVs, i.e. the charging power assigned to each vehicle at each time step. The problem is naturally formulated in terms of qudit variables [109] when more than two charging levels are considered. Crucially, it involves many inequality constraints: each EV has a minimum energy requirement, and at each time step the total charging power must not exceed the fuse limit. This makes it a useful benchmark for studying the behaviour of different constraint-handling methods under multiple simultaneous constraints.

The cost function is

$$C(\boldsymbol{x}) = \sum_{t=1}^{T} c_t \sum_{n=1}^{N_{\mathrm{EV}}} x_{n,t}\,, \tag{2.32}$$

where $x_{n,t} \in \{x_{\min}, \ldots, x_{\max}\}$ is the amount of energy charged (or discharged, if $x_{\min} < 0$)

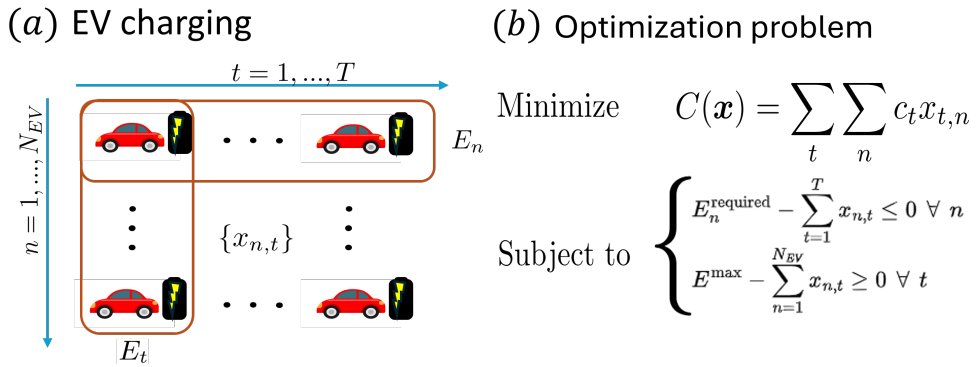

Figure 2.7: Schematic of the electric vehicle (EV) charging optimisation problem. (a) The decision variables $x_{n,t}$ are arranged in an $N_{\text{EV}} \times T$ matrix, where rows correspond to vehicles $n = 1, \ldots, N_{\text{EV}}$ and columns to time steps $t = 1, \ldots, T$. Each entry $x_{n,t}$ encodes the charging power delivered to vehicle n at time step t. The orange boxes highlight the two types of constraints: the per-column fuse limit E_t, which caps the total power drawn at each time step, and the per-row energy requirement E_n, which enforces a minimum total energy delivered to each vehicle. (b) Formal statement of the optimisation problem. The objective is to minimise the total charging cost $C(\boldsymbol{x}) = \sum_t \sum_n c_t x_{n,t}$, subject to the energy-requirement constraint $E_n^{\text{required}} - \sum_{t=1}^{T} x_{n,t} \leq 0$ for all n and the fuse-limit constraint $E^{\max} - \sum_{n=1}^{N_{\text{EV}}} x_{n,t} \geq 0$ for all t.

to EV n at time step t, and c_t are time-dependent energy cost coefficients. The two types of constraints considered are

$$E_n^{\text{required}} - \sum_{t=1}^{T} x_{n,t} \leq 0 \quad \forall\, n\,, \tag{2.33}$$

$$\sum_{n=1}^{N_{\text{EV}}} x_{n,t} - E^{\max} \leq 0 \quad \forall\, t\,, \tag{2.34}$$

representing, respectively, the minimum energy requirement for each vehicle and the fuse limit at each time step. These constitute $N_{\text{EV}} + T$ linear constraints in total, each couplingc a specific subset of variables; collectively they couple all variables.

For demonstration purposes, a simple instance is considered: two binary charging levels $x_{n,t} \in \{0, 1\}$, two vehicles ($N_{\text{EV}} = 2$), and four time steps ($T = 4$), with $E^{\max} = 1$ and $E_0^{\text{required}} = E_1^{\text{required}} = 2$. The feasible configurations are those in which each vehicle is charged during exactly two (not necessarily consecutive) time steps, subject to the restriction that the two vehicles are never charged simultaneously. The problem Hilbert space has dimension $\dim(\mathcal{H}) = 2^{T \cdot N_{\text{EV}}} = 2^8 = 256$, of which only 12 basis states are feasible; by symmetry between the two vehicles, the optimal solution is always doubly degenerate.

Incorporating the constraints via slack variables requires six auxiliary variables in total: two qudits of dimension $d = 3$ for the energy constraints of Eq. (2.33) and four qubits ($d = 2$) for the fuse constraints of Eq. (2.34). The Hilbert space dimension then becomes $\dim(\mathcal{H}_{\text{slack}}) = 2^8 \cdot 3^2 \cdot 2^4 = 36\,864$, which is two orders of magnitude larger than the original Hilbert space.

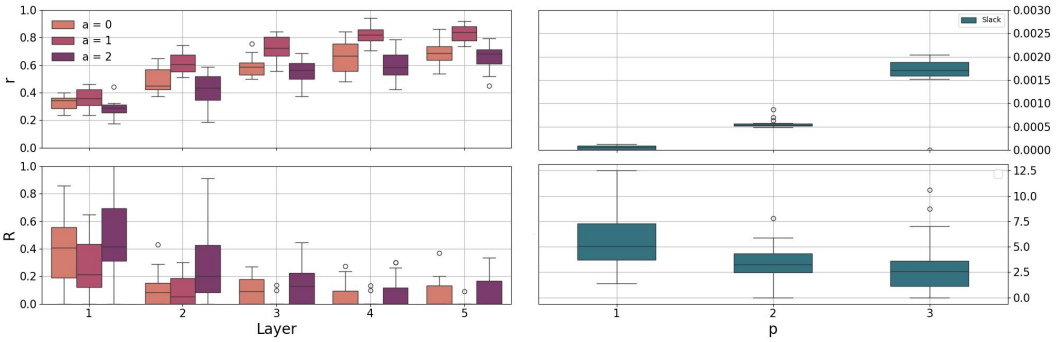

Figure 2.8: Results for the EV charging problem of Eqs. (2.32)–(2.34) with $N_{\mathrm{EV}} = 2$, $T = 4$, $E^{\max} = 1$, and $E_n^{\mathrm{required}} = 2$, for direct penalties (left) and slack variables (right) as a function of the number of QAOA layers p. The direct penalty approach vastly outperforms the slack-variable approach (note the different scales on the y-axes). Among the direct penalty approaches, the linear case $a = 1$ performs best: for $p \geq 3$, both the median and the 20th–80th quantiles of the approximation ratio R vanish, indicating that an optimal solution is found in every run.

Performance is studied as a function of the number of layers for different random realisations of the Hamiltonian of Eq. (2.32), where the prices are drawn from a uniform distribution $c_t \sim \mathcal{U}(0, 1)$. $N_{\mathrm{runs}} = 50$ runs are performed per instance. For the direct penalty approach, $p = 1, \ldots, 5$ layers are considered; for the slack-variable approach, simulations are limited to $p \leq 3$ layers owing to the substantially larger Hilbert-space dimension and the correspondingly longer runtimes.

The results, summarised in Fig. 2.8, corroborate and extend the findings of the previous benchmarks. The direct penalty approach achieves success rates of the order of 30% already at $p = 1$, rising to 60%–90% for $p = 5$. The approximation ratio rapidly approaches zero with increasing p, particularly for the best-performing case $a = 1$.

By contrast, the slack-variable results indicate a fundamental failure of the method in this multi-constraint regime. Although the success rate does increase slightly with the number of layers, it remains at the scale of 10^{-3}. The approximation ratio is an order of magnitude worse than for the direct penalty approach, indicating that the algorithm never approaches the ground-state energy.

The origin of this failure is the dramatic increase in Hilbert-space dimension required to encode all constraints via slack variables. For each slack variable only one configuration represents a feasible state, so the majority of the added states are unfeasible. In the example above, the fraction of feasible solutions falls from $2/2^8 \approx 4 \times 10^{-3}$ (without slack variables) to $2/(2^8 \cdot 3^2 \cdot 2^4) \approx 5 \times 10^{-5}$ (with slack variables), making the corresponding search problem prohibitively harder for QAOA.

2.3 Conclusions

This chapter has presented and benchmarked two strategies for incorporating inequality constraints into QAOA via the use of qudits: the qudit slack-variable approach and the direct penalty approach.

The main finding is that the direct penalty approach consistently and substantially out-

performs the slack-variable approach across all problem instances and metrics considered — success rate, approximation ratio, and feasible-state weight. The root cause of the slack-variable approach's limitations is structural: each slack variable inflates the total Hilbert space by a factor equal to the number of feasible values for the corresponding constraint, while only one of those additional states per constraint is actually feasible. The resulting dilution of the feasible subspace grows exponentially with the number of constraints, rendering the approach impractical for problems with many simultaneous inequality constraints, as illustrated by the EV charging benchmark.

The direct penalty approach avoids this problem entirely. By using a single ancilla qudit only to imprint conditional phases onto unfeasible basis states, the dimension of the solution space is left unchanged and the fraction of feasible solutions remains constant regardless of the number of constraints. The circuit implementation requires only $2N$ qudit-controlled `SUM` gates and a single qudit phase-shift unitary per constraint, and a single ancilla qudit suffices for arbitrarily many constraints applied sequentially. These properties make the method particularly well suited to problems where many constraints act simultaneously on restricted subsets of qubits.

Among the penalty forms examined (flat $a = 0$, linear $a = 1$, quadratic $a = 2$), the linear penalty consistently achieves the best success rate across all benchmarks and QAOA depths. For strongly constrained problems the flat penalty can offer a marginal advantage in the approximation ratio, an artefact of how energy is assigned to unfeasible states rather than a reflection of superior convergence. These observations suggest that $a = 1$ is the default choice of practical interest.

A secondary finding concerns constrained state preparation: even in the absence of a cost Hamiltonian, the direct penalty approach concentrates almost all quantum weight on the feasible subspace within a single QAOA layer, providing an efficient route to initialise constraint-preserving algorithms or to sample feasible configurations for applications in lattice gauge theories, topological models, and combinatorial sampling.

Looking forward, the bottleneck of the direct penalty method — the requirement for a qudit with at least $N+1$ levels for a constraint on N qubits — is relevant but not prohibitive given rapid experimental progress on qudit-based quantum hardware. Furthermore, the ability to parallelise constraint implementation by using multiple ancilla qudits offers a natural path to scaling the approach to larger problem instances. Overall, the results demonstrate that qudit-assisted direct penalisation is a practically viable and resource-efficient strategy for constrained quantum optimisation, and motivate its adoption as an alternative to slack variables in future QAOA implementations on near-term devices.

Chapter 3

Digitalized Counter Diabatic Optimization with qudits

This chapter presents the results obtained by adapting the counterdiabatic approach for adiabatic optimization to variational quantum algorithms for qudits. The material presented here is largely based on [25].

3.1 Digitalised Counterdiabatic Protocols for QAOA

As discussed in the Introduction, QAOA takes inspiration from quantum annealing. A key challenge in adiabatic quantum algorithms for ground-state search is the avoidance of diabatic transitions to higher excited states. One established strategy to mitigate this is counterdiabatic (CD) driving [110, 111], in which the Hamiltonian $H_a(t)$ of Eq. (1.48) is augmented to $H(t)$,

$$H(t) = H_a(t) + \dot{\lambda}(t)\, A_\lambda^{(\ell)}\,, \tag{3.1}$$

where $A_\lambda^{(\ell)} = i\sum_{k=1}^{\ell} \alpha_k(t)\, Q_{2k-1}$ is the auxiliary CD term, with $Q_k = [H_a(t), Q_{k-1}]$, $Q_0 = \partial_\lambda H_a(t)$, and ℓ denoting the order of approximation. The term $\dot{\lambda}(t)A_\lambda^{(\ell)}$ introduces the CD correction that suppresses non-adiabatic transitions.

The coefficients $\alpha_k(t)$ can be determined in several ways, including action minimisation [112, 111], variational optimisation [113], machine learning [114, 115], and Krylov subspace methods [116]. Throughout this work the action minimisation approach is adopted, requiring

$$\min_{\alpha_k} \mathrm{Tr}(G_\ell^2)\,, \tag{3.2}$$

with $G_\ell = Q_0 - i[H_a(t), A_\lambda^{(\ell)}]$. The structure of G_ℓ contains a term $\partial_\lambda H_0$ independent of the CD weights, together with a sum of terms that depend on those weights. This allows G_ℓ to be rewritten as $G_\ell = B_0 + \sum_k \alpha_k B_k$. For first-order CD terms ($\ell = 1$), the operators B_k contain at most two-body terms when H_P is two-local and H_0 is local, enabling simulation of the evolution under the CD Hamiltonian using the single- and two-body native gate set described in Sec. 1.1. This setup is adopted throughout.

Once the optimal CD terms are determined, the evolution is discretised using a first-order

Trotter decomposition,

$$U_{\text{Trott}}(0,T) = \prod_{j=1}^{n} \prod_{m=1}^{M} \exp\Big[-i\,\Delta t\, C_m(j\Delta t)\, H^{(m)}\Big] + \mathcal{O}(\Delta t^2)\,, \tag{3.3}$$

where n is the total number of Trotter steps ($T = n\Delta t$), $H^{(m)}$ are the individual terms of $H(t)$, $C_m(t)$ are their corresponding variational parameters, and M is the total number of terms in $H(t)$. A hybrid quantum-classical variant, known as digitised counterdiabatic QAOA (DCQAOA), is also considered: the coefficients of each term in Eq. (3.1) are treated as trainable parameters optimised classically [117]. An additional unitary $U_{\text{CD}} = e^{-i\sum_k \theta_k P_k}$ corresponding to the CD term is included in each layer, where P_k denotes the k-th term in the linear combination of Pauli operator tensor products corresponding to $A_\lambda^{(\ell)}$. The resulting parameterised evolution takes the form

$$U(\boldsymbol{\theta}) = \prod_p e^{-i\theta_{1,p}H_0}\, e^{-i\theta_{2,p}H_P}\, e^{-i\sum_{k=1}\theta_{2+k,p}P_k}\,. \tag{3.4}$$

In the impulse regime, where $|\alpha_k(t)\dot{\lambda}(t)| \gg |\lambda(t)|$, the adiabatic Hamiltonian $H_a(t)$ can be neglected and only the CD terms need to be implemented [118, 119]. The unitary evolution for this CD-inspired ansatz is then

$$U(\boldsymbol{\theta}) = \prod_p e^{-i\sum_{k=1}\theta_{k,p}\,P_k}\,. \tag{3.5}$$

Both DCQAOA and CD-inspired ansätze share a parameter scaling of $\mathcal{O}(N^2)$. For many optimisers the number of circuit evaluations per iteration depends on the number of parameters. To address this challenge, symmetries of the problem are exploited in the following section to reduce the number of trainable parameters in a problem-specific manner.

3.2 Symmetry-Enhanced Ansatz

This section introduces the use of symmetries to reduce the number of variational parameters in the DCQAOA algorithm.

The symmetries acting on the system — and in particular on the operator G_ℓ — are considered. The focus is restricted to spatial symmetries [120], i.e. permutations π that act on a Hermitian operator O by relabelling the qudit positions such that $\pi O\pi^{-1} = O$: the symmetry changes the labels of the qudits but leaves the operator invariant.

The set of symmetries of H_P is denoted $S_P = \{\pi_q\}$, where q labels the elements. Since H_0 is symmetric under the full permutation group S_n, it is also symmetric under $S_P \subseteq S_n$. It follows that the adiabatic Hamiltonian shares the same symmetries as H_P. The CD Hamiltonian inherits these symmetries as well, since $\partial_\lambda H(t)$ preserves them: $\partial_\lambda H(t) = \pi_q\, \partial_\lambda H(t)\, \pi_q^{-1} = \partial_\lambda(\pi_q H(t)\pi_q^{-1})$.

Denoting the CD terms by $\{A_k\}$, the action of any symmetry $\pi_q \in S_P$ merely relabels the indices ($k \to \pi_q(k) = k_q$),

$$A_\lambda^{(\ell)} = \sum_k \alpha_k A_k = \pi_q\left(\sum_k \alpha_k A_k\right)\pi_q^{-1} = \sum_{k_q} \alpha_{k_q} A_{k_q}\,. \tag{3.6}$$

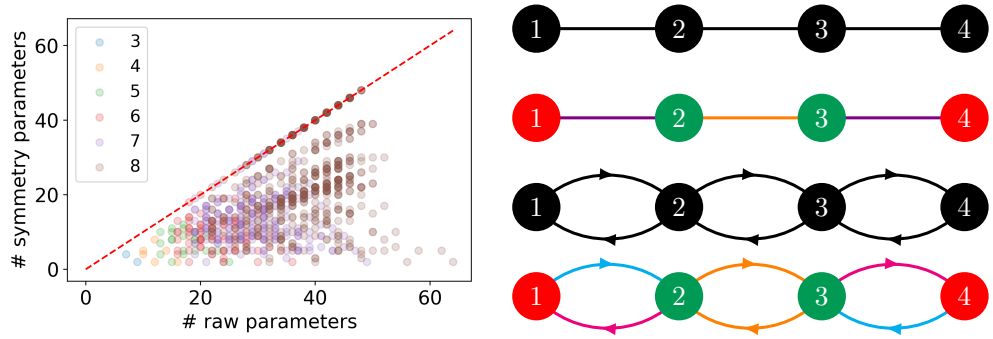

Figure 3.1: Left: Number of unique parameters in the CD ansatz after exploiting spatial permutation symmetries versus the total number of unique parameters without symmetry reduction. All ZZ-Ising Hamiltonians of the form of Eq. (3.9) with connected graph structure up to 8 nodes from the database of Ref. [121] are included. Labels indicate the system size. Right: Example graphs and their symmetries. From top to bottom: $\mathcal{G}$, $\mathcal{G}$ with vertex and edge orbits, $\mathcal{H}$, and $\mathcal{H}$ with vertex and arc orbits. Colours indicate the distinct elements of each vertex/arc orbit, which determine the parameter groupings.

Identifying each term, one sees that the only effect of applying a symmetry is a relabelling of the parameter indices. Consequently, for any π_q,

$$A_\lambda^{(\ell)} = \pi_q A_\lambda^{(\ell)} \pi_q^{-1} = \sum_k \alpha_{k_q} A_k \,. \tag{3.7}$$

This observation leads to the central property that underpins the parameter reduction. Informally, a set of parameters $\{\alpha_k\}_g$ can be *grouped* (i.e. constrained to share the same value) if their labels k form a closed set under the action of all symmetries in S_P. Formally, $\{\alpha_k\}_g$ is a grouped set if $\pi_q(\alpha_k \in \{\alpha_k\}_g) \in \{\alpha_k\}_g$ for all $\pi_q \in S_P$.

In the context of the action minimisation problem, the operators B_k in G_ℓ can be identified with the CD terms A_k. Since $B_k \sim [H_0, A_k]$ and both H_0 and A_k share the same symmetries, the operators B_k inherit them as well. Grouping parameters therefore reduces the number of unknowns in the action minimisation by $|\{\alpha_k\}_g| - 1$ for each group.

To quantify the reduction numerically, the ZZ-Ising Hamiltonian

$$H_P = \sum_i \sigma_i^z + \sum_{i,j \in \mathcal{G}} \sigma_i^z \sigma_j^z \tag{3.8}$$

with local mixer $H_0 = \sum_i \sigma_i^x$ is considered as a test case. The number of parameters after applying the reduction technique is computed for all non-trivial graph instances up to 8 qubits from a reference database [121]. As shown in Fig. 3.1, the average reduction of parameters over the considered setof graphs amounts to $59 \pm 19\%$ of the original number.

After computing the optimized CD parameters for both the standard and symmetry-enhanced methods, the results converge to the optimal CD coefficients, confirming that the reduction does not degrade the quality of the approximation. This reduction can therefore be extended to CD-inspired VQAs as well. Detailed derivations are provided in App. A.1.

To illustrate the parameter reduction procedure concretely, consider an $N = 4$ qudit

Algorithm 1 Algorithm for obtaining the arc orbits of the directed graph $\mathcal{H}$

Obtain the undirected graph $\mathcal{G}$ from $\mathcal{H}$
Compute vertex and edge orbits of $\mathcal{G}$
Define arc orbits as tuples (*arcs*, *tail vertex orbit*, *head vertex orbit*, *edge orbit*)
for each edge orbit **do**
  Identify arcs (h,t) of $\mathcal{H}$ such that $(h,t)=(i,j)$ or $(h,t)=(j,i)$, where (i,j) are edges of $\mathcal{G}$
  for each arc **do**
    Identify the vertex orbit of the tail and head
    Add the arc to the arc orbit with matching tail vertex, head vertex, and edge orbits
  end for
end for

system with the problem Hamiltonian

$$H_P = \sum_{i=1}^{4} L_{zi} + \sum_{i=1}^{3} L_{zi}L_{zi+1}\,, \tag{3.9}$$

corresponding to a linear-path connectivity graph $\mathcal{G}=(V_\mathcal{G}, E_\mathcal{G})$ with $V_\mathcal{G}=[1,2,3,4]$ and $E_\mathcal{G}=[(1,2),(2,3),(3,4)]$. With the mixer $H_0=\sum_{i=1}^{4} L_{xi}$, the first-order CD Hamiltonian is

$$A_\lambda^{(1)} = \sum_{i\in V_\mathcal{G}} \alpha_i\, L_{y_i} + \sum_{\{i,j\}\in E_\mathcal{G}} \Big(\alpha_{i,j}\, L_{y_i}L_{zj} + \alpha_{j,i}\, L_{y_j}L_{zi}\Big)\,. \tag{3.10}$$

Notably, even though the interactions in the problem Hamiltonian are undirected, the CD Hamiltonian has directionality: $L_{y_j}L_{zi}$ and $L_{y_i}L_{zj}$ are distinct terms with different parameters. The CD terms are therefore naturally associated with a directed graph $\mathcal{H}$ in which each undirected edge of $\mathcal{G}$ corresponds to two directed arcs, one in each direction (Fig. 3.1):

$$A_\lambda^{(1)} = \sum_{i\in V_H} \alpha_i\, L_{y_i} + \sum_{\{i,j\}\in A_H} \alpha_{i,j}\, L_{y_i}L_{zj}\,. \tag{3.11}$$

To reduce the number of parameters, the parameter groups must be identified. The grouping defined above coincides with the notion of graph orbits [122, 123]. For directed graphs, the edge-orbit concept is extended to arc orbits following Ref. [124]. The procedure for computing arc orbits is detailed in Alg. 1.

For the linear-path graph considered here, the vertex orbits of $\mathcal{G}$ are $[1,4]$ and $[2,3]$, and the edge orbits are $[(1,2),(3,4)]$ and $[(2,3)]$. Applying Alg. 1 yields the arc orbits $[(1,2),(4,3)]$, $[(2,1),(3,4)]$, and $[(2,3),(3,2)]$. The total number of variational parameters is thereby reduced from 10 to 5, since the symmetries enforce $\alpha_1=\alpha_4$, $\alpha_2=\alpha_3$, $\alpha_{12}=\alpha_{43}$, $\alpha_{21}=\alpha_{34}$, $\alpha_{23}=\alpha_{32}$.

In summary, the procedure for generating a symmetry-enhanced CD-inspired ansatz proceeds as follows. First, the CD terms $A_\lambda^{(\ell)}$ are obtained from H_P and H_0. Next, the arc and vertex orbits of the associated graph are computed. Finally, when parameterising the ansatz, the resulting parameter groups are imposed as equality constraints on the variational parameters. It is worth noting that while symmetry exploitation can sometimes hinder convergence in tensor networks and DMRG [125, 126], in quantum systems it consistently

reduces the number of required measurements — an important practical advantage for near-term devices. The following section illustrates with numerical examples that this reduction translates into concrete improvements in key figures of merit of quantum optimisation.

3.3 Numerical Results

To test the validity of the proposed techniques for practical quantum optimisation, problem encodings that are efficient in terms of qudit count, naturally map onto single- and two-body interaction Hamiltonians, and exhibit non-trivial symmetries are required. These properties are satisfied by several combinatorial problems defined over lattices or graphs, which are adopted as benchmarks in the following.

3.3.1 Max-3-Cut

As a paradigmatic combinatorial optimisation benchmark, the Max-k-Cut problem [127] is considered, focusing on instances for which no trivial or greedy solution exists. The objective is to partition the vertices of a graph into k colour classes such that the number of edges connecting vertices of different colours is maximised,

$$\max_{\mathbf{c}} \sum_{\{i,j\}\in E} \begin{cases} 1 & \text{if } c_i \neq c_j\,, \\ 0 & \text{otherwise}\,, \end{cases} \tag{3.12}$$

where $\mathbf{c} = \{c_i\}^N$ assigns each vertex i to a colour $1 \leq c_i \leq k$ and N denotes the number of vertices. Since each instance is uniquely defined by a graph, the graph's symmetries are used directly to reduce the number of variational parameters. The quantum encoding for the Max-3-Cut problem reads

$$H_P = \sum_{\{i,j\}\in E} L_{zi}L_{zj} - 2(L_{zi}^2 + L_{zj}^2) + 3L_{zi}^2 L_{zj}^2\,. \tag{3.13}$$

Although the numerical tests are limited to Max-3-Cut, the approach is immediately extensible to arbitrary k: the Hamiltonian for Max-k-Cut is obtained by solving a linear system of $\mathcal{O}(k^2)$ equations, as detailed in App. A.2.

The performance of the proposed techniques is assessed using four $p = 1$ layered ansätze: the first-order DCQAOA ansatz of Ref. [117], a symmetry-enhanced (grouped) variant of the same, a CD-inspired ansatz [118], and its symmetry-enhanced counterpart. As a baseline, the standard qudit-QAOA [12, 109] is employed. Using qudits for all algorithms ensures that comparisons are made within the same Hilbert space of dimension d^N.

Six instances of the smallest non-trivial size are considered: $N = 6$ qutrits ($d = 3$), with graphs selected from the database of Ref. [121] (HoG IDs 728, 220, 730, 758, 752, and 748). The selection criterion is that graphs have 6 nodes, are fully connected, and are non-planar [128]. Up to 500 iterations of the COBYLA optimiser [129] are performed for each classical update loop, as this algorithm does not require the computation of costly gradients. Multiple random initialisations are used for each problem instance.

The results for the expected energy on the first graph instance are shown in Fig. 3.2; the full comparison across all instances is provided in App. A.2. Approximation ratios $\mathcal{R} = \langle H_P \rangle / E_0$ (with E_0 the ground-state energy) are reported in Tab. 3.1.

All four CD-based algorithms reach lower energies than qudit-QAOA within far fewer iterations. The symmetry-enhanced ansatz converges rapidly to near-optimal energies, while

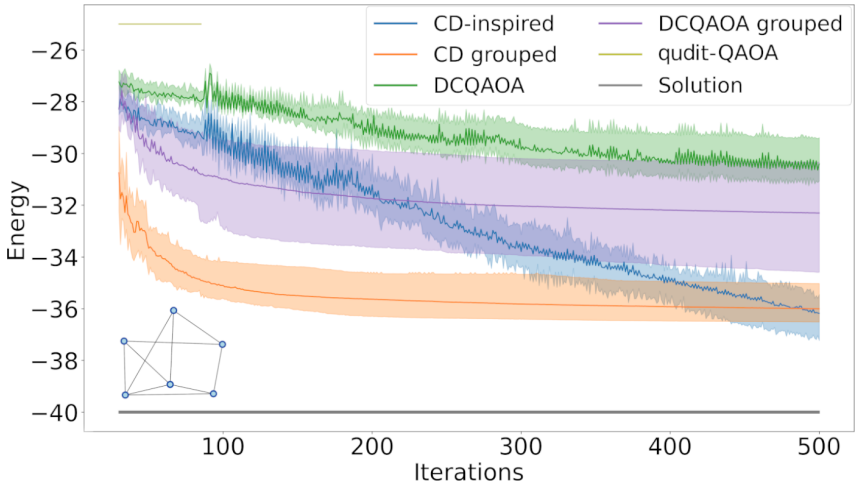

Figure 3.2: Expected energy at each iteration of the variational procedure for a Max-3-Cut instance, shown for all ansätze at a single layer ($p = 1$). The first 30 iterations are omitted to suppress initial fluctuations. The symmetry-enhanced (grouped) CD ansatz exhibits faster convergence than the fully parameterised CD ansatz. After many iterations, the fully parameterised ansatz achieves lower final energies owing to its higher expressivity. The same qualitative behaviour holds for the DCQAOA variants. Qudit-QAOA fails to solve the problem within the given circuit depth due to its limited expressivity.

the fully parameterised ansatz achieves marginally better final results at the cost of a harder optimisation landscape. For both grouped and ungrouped variants, the DCQAOA ansatz requires more iterations to converge — a counter-intuitive finding given that setting the QAOA parameters to zero reduces it to the CD ansatz. This slower convergence is likely attributable to barren plateaus [130] and over-parameterisation effects [131]. It is also noteworthy that, although the grouped ansatz performs slightly below the ungrouped one on average, competitive approximation ratios are achieved with a substantially reduced parameter count. Across all four problem instances examined, the grouping method consistently improves convergence speed, and the primary practical benefit of the approach is the reduction in the number of required circuit evaluations.

3.3.2 W State Preparation

As a second benchmark, the ability of the symmetry-enhanced CD ansatz to prepare the W state of a system of N qutrits is examined. The target state is

$$|W\rangle = \frac{1}{\sqrt{N}} \sum_{i=1}^{N} P_i \left(|2\rangle^{\otimes N-1} |0\rangle \right), \tag{3.14}$$

where P_i indicates a permutation of the local basis labels. The W state is equivalently a qutrit Dicke state D_k^N with $k = 1$, as shown in App. A.3. It constitutes an equal superposition of all possible ground states of the Hamiltonian

$$H_P = \left(\frac{N}{2} - 1 \right) \sum_{i=1}^{N} L_{zi} + \frac{1}{2} \sum_{j<i} L_{zi} L_{zj}, \tag{3.15}$$

Table 3.1: Mean approximation ratios and standard deviations at the final iteration for all algorithms and problem instances. Bold values indicate results exceeding the classical approximation ratio of 0.800217 [132]. Underlined values mark the best mean approximation ratio $\mathcal{R}$ for each problem. Qudit-QAOA converged to the same energy for each graph up to machine precision.

Problem graph Id	728	220	730	758	752	748
CD all	**0.90 ± 0.04**	**0.96 ± 0.01**	**0.85 ± 0.03**	**0.91 ± 0.03**	**0.92 ± 0.02**	**0.83 ± 0.01**
CD grouped	**0.90 ± 0.05**	**0.95 ± 0.03**	**0.92 ± 0.06**	**0.92 ± 0.03**	**0.93 ± 0.03**	**0.85 ± 0.03**
DCQAOA all	0.76 ± 0.03	**0.91 ± 0.04**	0.74 ± 0.04	**0.81 ± 0.01**	**0.87 ± 0.03**	**0.81 ± 0.01**
DCQAOA grouped	**0.81 ± 0.06**	**0.90 ± 0.04**	**0.81 ± 0.04**	**0.87 ± 0.06**	**0.85 ± 0.02**	**0.87 ± 0.06**
qudit-QAOA	0.63	0.76	0.63	0.69	0.75	0.68

which corresponds to a fully connected graph with a local potential on each vertex. For this fully connected structure, the parameter grouping is trivial: the CD Hamiltonian of Eq. (3.10) reduces to the case $\alpha_i = \alpha$ and $\alpha_{i,j} = \tilde{\alpha}$ for all i, j, yielding only two independent parameters. Crucially, this grouping enforces the symmetry of the unitary transformation on the system — a symmetry that would be broken if all parameters were left independent. Since the W state is an equal superposition of the N degenerate ground states, the symmetry enforcement is essential to guide the evolution toward this specific target rather than a generic superposition of ground states.

The efficiency of the state preparation is measured by the fidelity of the final state with the target state of Eq. (3.14). For pure states, the fidelity reduces to

$$\mathcal{F} = \left|\langle W|U(\boldsymbol{\theta})|\psi_0\rangle\right|^2, \tag{3.16}$$

where the initial state is chosen as

$$|\psi_0\rangle = \frac{1}{\sqrt{d}} \bigotimes_{j=1}^{N} \sum_{i_j=0}^{d-1} |i_j\rangle \ . \tag{3.17}$$

The cost function minimised during the algorithm is the expectation value of H_P, which drives the state towards lower energies and, consequently, towards the feasible subspace. Results obtained by directly maximising the fidelity are provided in App. A.3 for comparison.

The algorithm is executed from $p = 1$ to $p = 5$ layers for $N = 3$ qutrits, and the results are displayed in Fig. 3.4.

The symmetry-enhanced CD ansatz outperforms qudit-QAOA for a small number of layers; however, this advantage diminishes as p increases. This behaviour is consistent with the transition from the impulse regime (low p) to the adiabatic regime (large p). At the same time, the fully parameterised ansätze approximate the ground-state energy more rapidly but achieve lower fidelity with the target W state. Already for $p = 2$, the symmetry-enhanced ansatz achieves a fidelity $\mathcal{F} \sim 71\%$ with the W state.

3.3.3 A consideration on the usefulness of Symmetries in CD Ansatze

The preceding two examples demonstrate that the use of symmetries can be beneficial for quantum optimisation, particularly by improving the convergence speed of the classical

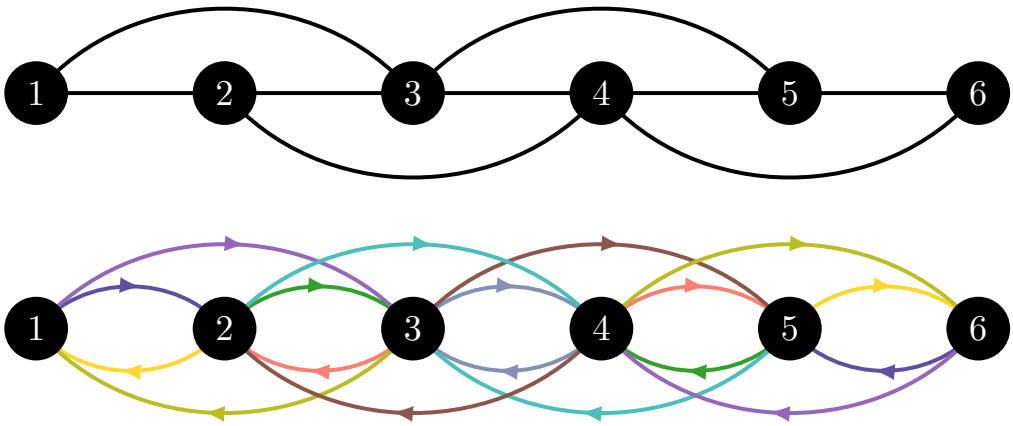

Figure 3.3: Top: Interaction scheme for Hamiltonian (3.18). Bottom: Parameter reduction for the symmetrised DC ansatz. Each line represents a CD term in the Hamiltonian; lines of the same colour share the same variational parameter. This grouping reduces the number of optimised parameters from 18 to 9.

optimisation loop while maintaining competitive performance. This improvement originates from the reduction in the number of variational parameters, which simplifies the optimisation landscape. A remaining question concerns the expressivity of the symmetry-constrained circuits.

In the MAX-3-CUT example, the symmetry-enhanced ansatz converges significantly faster, indicating an easier optimisation landscape. At the same time, the fully parameterised CD ansatz reaches better final solutions, reflecting higher expressivity at the cost of a harder landscape.

An intuitive explanation is the following. When the problem Hamiltonian contains both single- and two-body terms, the leading order of the nested commutator expansion (3.1) generates both single- and two-body contributions. Including these in the approximated ansatz of Eq. (3.4) provides sufficient flexibility to compensate for higher-order diabatic contributions, even with a small number of layers. For problem Hamiltonians containing only two-body terms, however, the leading-order commutator with the driver Hamiltonian generates only two-body contributions. Restricting the ansatz to these terms reduces its ability to compensate for higher-order effects. As demonstrated below, supplementing the ansatz with additional single-body terms significantly improves performance in this regime.

In terms of optimisation hardness, reducing the number of parameters is most beneficial for single- and two-body Hamiltonians, where the landscape simplification outweighs the loss of expressivity. For Hamiltonians containing only two-body terms, both the symmetry-enhanced and the fully parameterised ansätze have a low parameter count to begin with, so the additional restriction imposed by the symmetries can further degrade expressivity without a commensurate improvement in convergence.

As a practical guideline: for Hamiltonians with single- and two-body terms, there is a trade-off between expressivity and optimisation hardness, and the symmetry reduction can be chosen depending on which property is prioritised. For purely two-body Hamiltonians, the marginal gain in convergence speed from symmetry reduction does not justify the associated loss of expressivity.

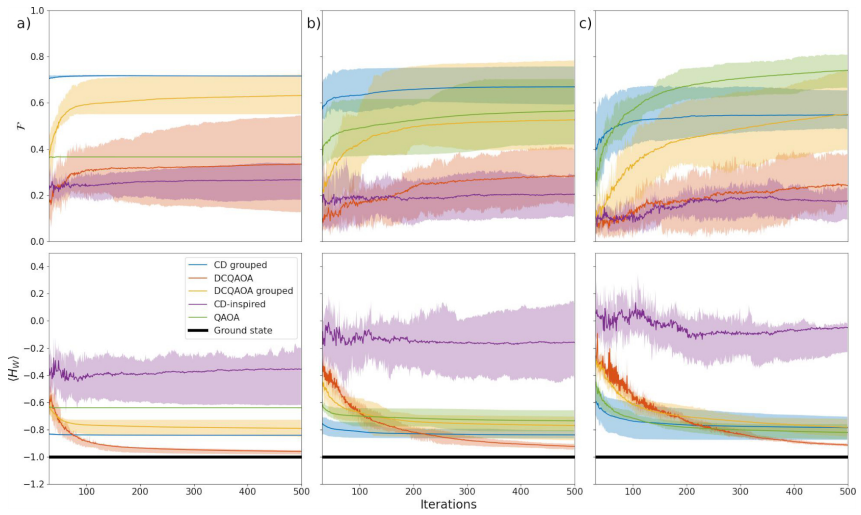

Figure 3.4: Fidelity defined in Eq. (3.16) (top row) and expectation value of the Hamiltonian of Eq. (3.15) (bottom row) as a function of optimizer iteration, for $N = 3$ and (a) $p = 2$, (b) $p = 3$, and (c) $p = 5$ layers (left to right). The first 30 iterations are omitted to suppress initial fluctuations. The symmetry-enhanced grouped CD ansatz outperforms the other ansätze at low layer counts, as expected from the symmetry argument.

To verify this intuition, the following subsection examines an antiferromagnetic Ising Hamiltonian with only two-body interactions, supplemented by artificial single-body terms to test the role of expressivity.

3.3.4 Antiferromagnetic Second-Neighbour Ising Hamiltonian

The trade-off described above is further examined by studying the ability of the symmetry-enhanced CD ansatz to optimise a next-to-nearest-neighbour Ising model of N qutrits with open boundary conditions,

$$H_{AF} = \sum_{\langle\langle ij \rangle\rangle} L_{z_i} L_{z_j} + \sum_{\langle ij \rangle} L_{z_i} L_{z_j} \,. \tag{3.18}$$

The Hamiltonian is symmetric under the index permutation $i \to N-i$, and the corresponding parameter grouping is illustrated in Fig. 3.3.

The cost function minimised during the optimisation is the expectation value of H_{AF} over the parametrised state,

$$\mathcal{C} = \langle \psi_0 | \, U^\dagger(\boldsymbol{\theta}) \, H_{AF} \, U(\boldsymbol{\theta}) \, | \psi_0 \rangle \,. \tag{3.19}$$

The results in Fig. 3.5 confirm the expected behaviour. When single-body terms are not included, a substantial performance gap exists between the grouped and fully parameterised ansätze, particularly at lower layer counts ($p = 2, 3$). The introduction of an additional layer of single-body rotations $U = \prod_{i=0}^{N-1} e^{i\alpha_i L_{y_i}}$ substantially narrows this gap for the grouped ansätze, while the performance at $p = 4$ becomes comparable across all variants. These observations are consistent with the intuition developed in the previous section and confirm that single-body terms play an important role in enhancing expressivity for purely two-body Hamiltonians.

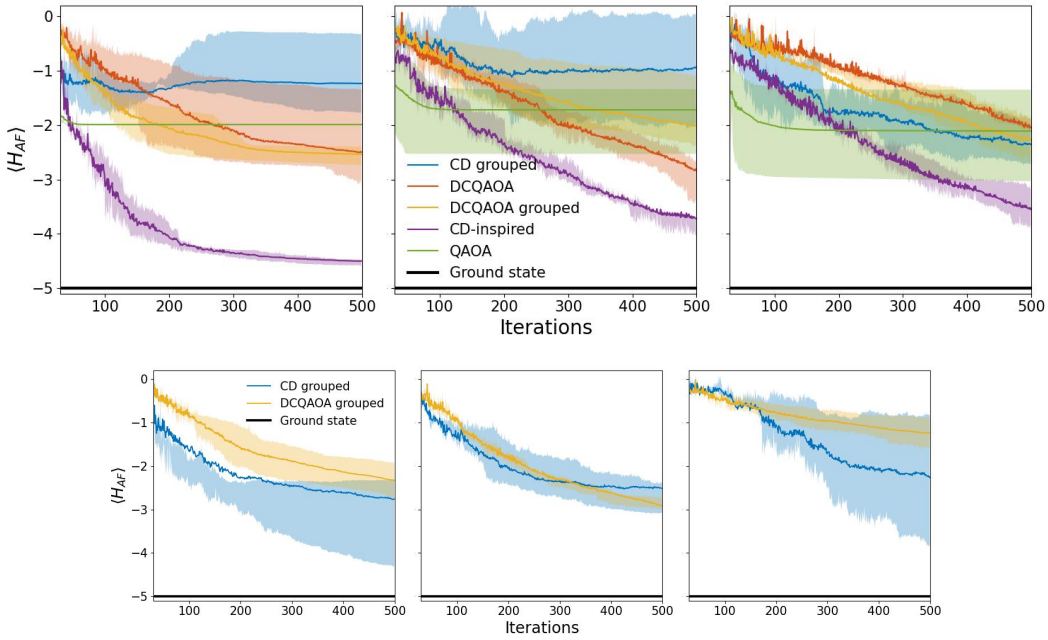

Figure 3.5: Top: Expectation value of Hamiltonian (3.18) over optimiser iterations for all ansätze without artificial single-body terms, for $p = 2, 3, 4$ layers (left to right). Bottom: Same figure of merit for the grouped ansätze supplemented with an additional single-body layer $U = \prod_{i=0}^{N-1} e^{i\alpha_i L_{y_i}}$. The fully parameterised CD-inspired ansatz converges to a value closer to the ground state in fewer iterations compared to the grouped CD and grouped DCQAOA ansätze. This difference is reduced when the grouped ansätze are enriched with single-body terms, although the fully parameterised ansatz still converges faster.

3.4 Conclusions

This chapter has introduced and benchmarked a symmetry-enhanced framework for CD-inspired variational quantum algorithms operating on qudit systems. The central contribution is the identification of a systematic procedure that exploits the spatial symmetries of the problem Hamiltonian to reduce the number of trainable parameters in the DCQAOA and CD-inspired ansätze, without requiring modifications to the circuit structure itself.

The parameter reduction is achieved by identifying groups of CD terms whose indices form orbits under the symmetry group of the problem Hamiltonian. Terms belonging to the same orbit are constrained to share a single variational parameter. For the ZZ-Ising Hamiltonian benchmarked across a set of non-trivial graphs with up to 8 qubits, this procedure yields a reduction to $59 \pm 19\%$ of the original parameter count, with no loss in the quality of the CD approximation. The reduction is directly reflected in the number of circuit evaluations required per optimisation iteration, which is the dominant cost on near-term quantum hardware.

Three numerical benchmarks are used to characterise the benefits and limitations of the approach. For the Max-3-Cut problem on qutrit systems, the symmetry-enhanced CD ansatz consistently outperforms qudit-QAOA and converges significantly faster than the fully parameterised CD ansatz, achieving approximation ratios that exceed the best known

classical bound of 0.800217 across all tested instances. For the W state preparation task, the symmetry enforcement is not merely beneficial but necessary: without it, the ansatz does not preserve the permutation symmetry of the target state, and the optimisation drives the system toward a generic superposition of ground states rather than the equal-weight W state. Already at $p = 2$ layers, a fidelity of approximately 71% is achieved. For the antiferromagnetic next-to-nearest-neighbour Ising model — a purely two-body Hamiltonian — the results reveal the limitations of the approach: the reduction in parameter count imposes an expressivity cost that is not offset by the improvement in optimisation landscape when no single-body terms are present in the CD expansion. Supplementing the grouped ansatz with additional single-body rotation layers substantially narrows the performance gap, confirming that single-body terms are a key ingredient for expressivity in this regime.

A practical guideline emerges from these findings. For Hamiltonians containing both single- and two-body terms, symmetry-enhanced ansätze offer an attractive balance between convergence speed and approximation quality, and are recommended as the default choice. For purely two-body Hamiltonians, the symmetry reduction should be complemented by the inclusion of additional single-body layers to retain sufficient expressivity. In both cases, the reduction in the number of circuit evaluations constitutes a concrete and hardware-relevant advantage for near-term quantum devices.

Looking ahead, the framework presented here is directly extensible to higher-order CD terms, larger system sizes, and other combinatorial problems defined on graphs. The combination of qudit encoding — which provides a compact Hilbert space representation — with symmetry-informed parameter reduction and CD-inspired circuit design represents a promising direction for scaling quantum optimisation beyond the capabilities of qubit-based approaches.

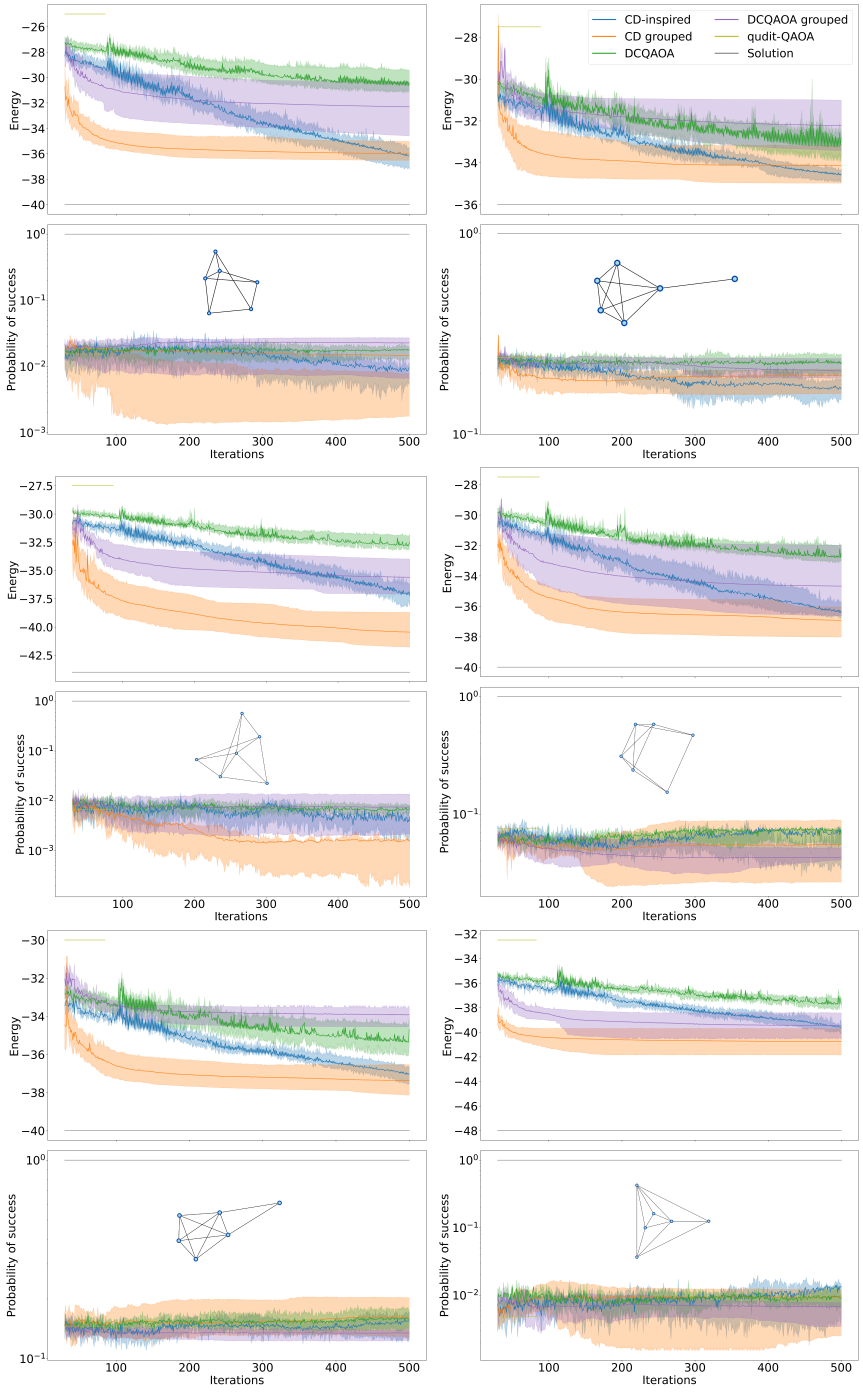

Figure 3.6: Results for all algorithms and instances of the Max-3-Cut problem. Solid lines correspond to mean values; coloured areas indicate the interquartile range. Upper panels: energy at each iteration. Middle panels: probability of measuring any state corresponding to the correct solution. Inset graphs depict the connectivity graph for each instance. The numbers of runs are 20, 12, 12, 10, 11, and 9, respectively, each using uniformly random initial parameters.

Chapter 4

Barren plateaus and resources in quantum optimization

This chapter is focused on describing results obtained in collaborations in which the author of this thesis was not the principal investigator. Both of the works presented here focus on the feasibility of executing quantum optimization algorithms on real quantum hardware.

The first work focuses on noise induced barren plateaus, or rather their absence, in variational quantum algorithms run on real hardware. These results help understanding the impact of different types of noise when executing variational quantum algorithms.

The second work focuses on quantifying the nonstabilizerness generated in quantum optimization protocols, in order to understand if those are easily classically simulatable or if the present genuine quantum complexity.

For both of the sections of this chapter, the author helped develop the the theory and provided numerical support. In the specific, for section 4.1 the author helped frame the theory of non unital noise in the context of the hardware computation and performed the density matrix simulations mentioned in the main text and in App B.4. For the analysis of resources in quantum optimization, the author proveded numerical simulations of QAOA, along with explicit computations of analytical values of nonstabilizerness for specific states described in the main text and in App. C.8

4.1 Absence of Noise Induced Barren Plateaus In Hardware

This section presents an experimental study of Noise-Induced Barren Plateaus (NIBP) on IBM quantum hardware, demonstrating their absence under amplitude-damping noise characterised by the qubit T_1 coherence time, based on [26].

Variational Quantum Algorithms (VQAs) [59, 13, 14] are among the most promising near-term quantum computing approaches, yet are fundamentally limited by Barren Plateaus (BPs) — the exponential vanishing of cost-function gradients with system size [15, 16]. Of particular practical relevance are Noise-Induced Barren Plateaus (NIBPs) [133], which arise independently of circuit structure or observable, depending solely on gate error accumulation. The original NIBP theory assumes depolarising noise, under which deep circuits converge to the maximally mixed state, making all gradient estimates exponentially small.

However, recent work has shown that non-unital noise — specifically amplitude damping via T_1 relaxation — does *not* produce NIBPs [134, 135, 136]: the output state instead converges to a noise-induced limit set that retains parameter dependence, preserving a finite gradient signal. Since amplitude damping is unavoidable in physical hardware, NIBPs may never be observable in practice.

This chapter provides the first experimental confirmation of this prediction. Running circuits with up to $N = 102$ qubits on IBM superconducting processors and estimating gradient norms via Information Content Landscape Analysis (ICLA) [137], we find that the gradient signal saturates at a constant value rather than decaying exponentially. An extracted effective coherence time T_1^{eff} falls well below the mean device T_1, suggesting that the worst-performing qubits dominate the onset of amplitude-damping effects.

4.1.1 Information Content Landscape Analysis

As described in Section 1.2.3, the cost function of a VQA is the expectation value of an observable $\hat{O}$ over a parametrised state $\rho(\vec{\theta}) = U(\vec{\theta})\rho_0 U(\vec{\theta})^\dagger$:

$$C(\vec{\theta}) = \mathrm{Tr}[\hat{O}\,\rho(\vec{\theta})]\,. \tag{4.1}$$

BP are a property of the gradient ∇C, defined as the exponential decay of its variance with system size N [15]:

$$\mathrm{Var}_{\vec{\theta}}[\partial C(\vec{\theta})] \in \mathcal{O}\left(\frac{1}{b^N}\right), \quad b > 1\,. \tag{4.2}$$

Estimating the gradient norm directly on hardware requires many circuit evaluations. Information Content Landscape Analysis (ICLA) [137, 138] provides an efficient alternative: it is analytically proven to estimate the average normalised gradient norm $\|\nabla C\|/C_0$ using only $M = \mathcal{O}(m)$ parameter samples, where m is the number of variational parameters.

The procedure is summarised in Algorithm 2. A discrete cost-function landscape $\Omega = \{C_R(\vec{\theta}_1), \ldots, C_R(\vec{\theta}_M)\}$ is obtained by measuring the cost function for each sampled parameter vector. A random walk over Ω is used to compute finite-difference gradient approximations ΔC_i, which are then symbolised into a ternary sequence $\phi(\epsilon) \in \{-, \odot, +\}$ according to a threshold ϵ. The empirical information content (IC) is computed from the transition probabilities between consecutive symbols:

$$H(\epsilon) = \sum_{a \neq b} -p_{ab} \log_6 p_{ab}\,. \tag{4.3}$$

The threshold $\epsilon_{\mathrm{M}} = \mathrm{argmax}_\epsilon H(\epsilon)$ signals maximal landscape information content, and the gradient norm estimate is:

$$\|\nabla C\|/C_0 \simeq \epsilon_{\mathrm{M}}\sqrt{m}\,. \tag{4.4}$$

Algorithm 2 Information Content Landscape Analysis (ICLA)

Input: number of parameters m, oversampling factor F, cost function C
Sample $M = F \cdot m$ points $\Theta = \{\vec{\theta}_1, \ldots, \vec{\theta}_M\} \in [0, 2\pi)^m$
Measure $C(\vec{\theta}_i)$ on a quantum computer for each $\vec{\theta}_i \in \Theta$
Store the landscape $\Omega = \{C(\vec{\theta}_1), \ldots, C(\vec{\theta}_M)\}$
Generate a random walk W of $S + 1 < M$ steps over Θ
Compute the finite-difference gradient at each step i: $\Delta C_i = \dfrac{C(\vec{\theta}_{i+1}) - C(\vec{\theta}_i)}{\|\vec{\theta}_{i+1} - \vec{\theta}_i\|}$
for each value of ϵ **do**
    Map each ΔC_i to a symbol in $\{-, \odot, +\}$:

$$\phi_i(\epsilon) = \begin{cases} - & \text{if } \Delta C_i < -\epsilon \\ \odot & \text{if } |\Delta C_i| \leq \epsilon \\ + & \text{if } \Delta C_i > \epsilon \end{cases}$$

    Compute transition probabilities p_{ab} between consecutive symbols
    Compute the information content $H(\epsilon) = \sum_{a \neq b} -p_{ab} \log_6 p_{ab}$
end for
Find $\epsilon_{\mathrm{M}} = \mathrm{argmax}_\epsilon \; H(\epsilon)$
Output: $\|\nabla C\|/C_0 = \epsilon_{\mathrm{M}} \sqrt{m}$

4.1.2 Noise-Induced Barren Plateaus and Their Absence

As introduced in Section 4.1, the T_1 coherence time parametrises the decay of a qubit from the excited state $|1\rangle$ to the ground state $|0\rangle$:

$$p_A(t) = 1 - e^{-t/T_1} \,. \tag{4.5}$$

This amplitude-damping channel is non-unital: its fixed point is the pure state $|0\rangle\langle 0|$, not the maximally mixed state.

The original NIBP analysis of Ref. [133] assumes a depolarising error model. Under depolarising noise, for a circuit of depth L with single-qubit error probability p, the output state approaches the maximally mixed state exponentially:

$$\mathcal{E}_{\mathrm{dep}}(\rho) = (1 - p)\rho + \frac{p}{3} \sum_{P \in \{X,Y,Z\}} P\rho P \,. \tag{4.6}$$

Once the maximally mixed state is reached, all cost-function values concentrate around a single constant, and any gradient estimator returns a value exponentially close to zero — the NIBP.

Recent theoretical results [134, 135, 136] show that non-unital channels such as amplitude damping prevent this concentration. Instead, the dynamics converge to a *noise-induced limit set* — a structured sub-manifold of state space whose elements retain dependence on the variational parameters, thereby preserving a finite gradient. Specifically, Ref. [134] proves that NIBP are absent once the amplitude-damping probability per qubit reaches $p_{\mathrm{non-unital}} = 3/4$.

4.1.3 Experimental Setup

To probe the impact of amplitude damping on NIBP experimentally, the nearest-neighbour Ising chain is used as the model Hamiltonian:

$$H_C = \sum_{i=0}^{N-2} \frac{J_{i,i+1}}{2} \sigma_z^i \sigma_z^{i+1} + \sum_{i=0}^{N-1} \frac{h_i}{2} \sigma_z^i \,, \tag{4.7}$$

with parameters drawn randomly from fixed sets $J_{i,i+1} \sim \mathcal{U}(\{\pm 2, \pm 1.2, \pm 0.8, \pm 0.4\})$ and $h_i \sim \mathcal{U}(\{0.8, \pm 0.4, \pm 0.24, \pm 0.16, -0.08\})$. This Hamiltonian is chosen because it is 2-local (avoiding cost-function BP) and its coupling topology is compatible with the IBM heavy-hex hardware layout (no additional two-qubit gate overhead).

A QAOA ansatz [12, 60] is adopted, as described in Section 1.2.3, with a number L of alternating layers:

$$U(\vec{\theta}) = \prod_{l=1}^{L} e^{-i\theta_l^{(1)} H_M} \, e^{-i\theta_l^{(2)} H_C} \,, \quad H_M = \sum_{i=0}^{N-1} X_i \,, \tag{4.8}$$

giving $m = 2L$ variational parameters. The initial state is the equal superposition $\rho_0 = |+\cdots+\rangle\langle+\cdots+|$.

For a given parameter set $\vec{\theta}$, all qubits are measured in the computational basis R times. Since H_C contains only commuting σ_z^i and $\sigma_z^i \sigma_z^{i+1}$ terms, a single measurement setting suffices. The energy of a single shot $z^{(r)}$ is

$$E(z^{(r)}) = \sum_{i=0}^{N-2} \frac{J_{i,i+1}}{2} s_i^{(r)} s_{i+1}^{(r)} + \sum_{i=0}^{N-1} \frac{h_i}{2} s_i^{(r)} \,, \tag{4.9}$$

where $s_i^{(r)} = 2z_i^{(r)} - 1 \in \{-1, 1\}$.

ICLA is applied to estimate $\|\nabla C\|/C_0$ for $N \in \{8, 20, 45, 65, 102\}$ qubits and up to $L = 120$ layers. The landscape size is capped at $M = \min(10m, 200)$; robustness of this truncation is verified in App. B.2. $R = 16384$ measurement shots are used per parameter set, with $R = 32768$ or $R = 65536$ for cases requiring lower shot-noise floors (see App. B.3).

Standard error mitigation techniques (Pauli twirling, dynamical decoupling, zero-noise extrapolation) are explicitly confirmed to be ineffective against non-unital amplitude-damping errors and are not applied.

Experiments are run on IBM 127-qubit Falcon processors `ibm_brisbane`, `ibm_kyiv`, and `ibm_sherbrooke`, and the 156-qubit Heron processor `ibm_fez`. Calibration data — T_1, T_2, ECR-gate error, and readout error — are recorded for each experimental run. Details on circuit structure, gate counts, depth, and runtime scaling, as well as full calibration tables, are provided in App. B.1.

4.1.4 Results

Figure 4.1 shows the normalized gradient norm $\|\nabla C\|/C_0$ as a function of circuit runtime t_{cir} for all experimental configurations and for classical density-matrix simulations at $N = 8$. The normalization factor $C_0 = \sqrt{\sum_i (J_{i,i+1}^2 + h_i^2)} \sim \sqrt{N}$ accounts for the scaling of H_C with qubit number.

Classical simulations ($N = 8$). The noiseless simulation (dark purple) shows $\|\nabla C\|/C_0$ levelling off around 0.05 for runtimes above $\sim 150\,\mu s$, consistent with the use of a local observable. Under depolarising noise ($p = 0.025$, light green), $\|\nabla C\|/C_0$ decays exponentially

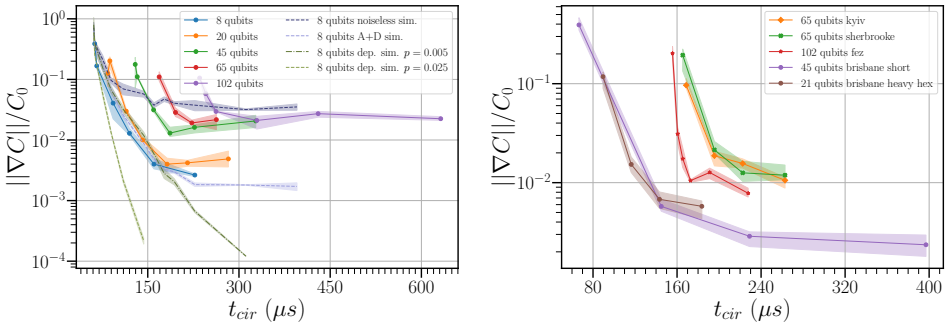

Figure 4.1: **Left:**Average normalized gradient norm $\|\nabla C\|/C_0$ versus circuit runtime t_{cir}. Solid lines with markers: experimental data from IBM quantum hardware for various qubit numbers N. Dashed lines: 8-qubit density-matrix simulations under noiseless, depolarising (label 'dep.'), and amplitude-damping-plus-dephasing (label 'A+D') noise models. Depolarising simulations show clear NIBP (exponential decay); all hardware experiments and amplitude-damping simulations instead show the gradient saturating at a constant value for large t_{cir}.
Right: Normalised gradient norm $\|\nabla C\|/C_0$ as a function of circuit runtime t_{cir} for various hardware setups. 'short' refers to a parallel compilation of the same logical circuit; 'heavy hex' denotes a Hamiltonian with the native hardware coupling topology. The flattening behaviour is consistent across all platforms and setups.

from 1 to 10^{-4} within $t_{cir} \lesssim 150\,\mu s$, confirming the NIBP prediction of Ref. [133]. Under the realistic T_1/T_2 noise model

$$\rho_A(t) = \begin{pmatrix} \rho_{00} + (1 - e^{-\frac{t}{T_1}})\rho_{11} & e^{-\frac{t}{2T_1}}\rho_{10} \\ e^{-\frac{t}{2T_1}}\rho_{01} & e^{-\frac{t}{T_1}}\rho_{11} \end{pmatrix}. \tag{4.10}$$

with$T_1 \sim \mathcal{N}(244\,\mu s, 74\,\mu s)$ and $T_2 \sim \mathcal{N}(159\,\mu s, 93\,\mu s)$ (light blue), the gradient initially decreases similarly to the depolarising case, then changes behaviour around $t_{cir} \sim 170\,\mu s$ and fully flattens at $\|\nabla C\|/C_0 \approx 0.002$ for $t_{cir} \gtrsim 230\,\mu s$. This confirms that amplitude damping prevents NIBP on the timescale of $t_{cir} \sim T_1$.

A spectral analysis of the simulated density matrices (App. B.4) corroborates this picture: depolarising noise drives the spectrum toward a uniform distribution (maximally mixed state), while amplitude damping produces a structured, non-uniform spectrum consistent with convergence to a noise-induced limit set.

Hardware experiments. All experimental curves show the same qualitative behaviour: $\|\nabla C\|/C_0$ decreases initially, then levels off at a constant value. The $N = 8$ case (blue) shows remarkable quantitative agreement with the realistic simulation, both flattening around $t_{cir} \approx 230\,\mu s$ at $\|\nabla C\|/C_0 \approx 0.002$. For $N = 20$ and $N = 45$, the gradient levels off around $t_{cir} \approx 180$–$185\,\mu s$, reaching $\|\nabla C\|/C_0 \approx 0.004$–$0.01$. For $N = 65$ and $N = 102$, flattening occurs around $t_{cir} \approx 220$–$270\,\mu s$, with $\|\nabla C\|/C_0 \approx 0.025$. The absence of NIBP is robustly confirmed up to $t_{cir} = 630\,\mu s$ for the $N = 102$ case.

Small upward trends observed for $N = 20$ and $N = 45$ at long runtimes are attributed to different physical qubits being used for different circuit lengths, leading to slightly different mean T_1 values that introduce additional landscape structure.

The same experiments are repeated on `ibm_kyiv`, `ibm_sherbrooke`, and `ibm_fez`, and

Table 4.1: Mean T_1 and effective T_1^{eff} for all experimental configurations on `ibm_brisbane`, together with the fraction of qubits below each threshold and the extracted flattening times t_{cir}^{flat}.

N	$\langle T_1 \rangle$ $[\mu s]$	% qubits	t_{cir}^{flat} $[\mu s]$	T_1^{eff} $[\mu s]$	% qubits
8	216 ± 39	32	230 ± 20	166 ± 33	13
20	256 ± 65	52	180 ± 20	130 ± 36	2.8
45	240 ± 76	50	185 ± 15	134 ± 44	7.4
65	247 ± 77	56	220 ± 15	166 ± 51	11
102	236 ± 68	51	260 ± 30	188 ± 58	24

with a more efficient parallel circuit compilation (labelled 'short') and a heavy-hex-native Hamiltonian coupling topology. As shown in Fig. 4.1, all configurations exhibit the same qualitative flattening behaviour, confirming the robustness of the result across hardware platforms, circuit types, and Hamiltonian topologies.

4.1.5 Effective T_1 from the Absence of NIBP

The theoretical result of Ref. [134] predicts that NIBP are absent once the amplitude-damping probability per qubit reaches $p_{\text{non-unital}} = 3/4$. Assuming that the flattening of the gradient signal marks the onset of this regime, the flattening time t_{cir}^{flat} can be related to an effective coherence time via Eq. (4.5):

$$T_1^{\text{eff}} = \frac{-t_{cir}^{\text{flat}}}{\ln[1 - p_{\text{non-unital}}]} \approx 0.72\, t_{cir}^{\text{flat}} \,. \tag{4.11}$$

Table 4.1 reports the extracted values alongside the mean device T_1. The effective T_1^{eff} is substantially smaller than the mean T_1 in all cases. As shown in Fig. 4.2, the T_1^{eff} corresponds to only 3%–25% of the qubits having a coherence time at or below this value, indicating that it is the $\sim 20\%$ of qubits with the shortest T_1 that determine the onset of the gradient flattening.

Figure 4.2 confirms this finding across all platforms and setups: all extracted T_1^{eff} lie in the range $0.4 \lesssim T_1^{\text{eff}}/\langle T_1 \rangle \lesssim 0.9$, well below the mean coherence time.

4.1.6 Conclusions

This chapter has presented the first experimental demonstration of the absence of Noise-Induced Barren Plateaus (NIBP) on real quantum hardware, corroborating recent theoretical predictions [134, 135].

The central finding is that the normalised gradient norm $\|\nabla C\|/C_0$ does not decay exponentially with circuit depth, as NIBP theory would predict under depolarising noise. Instead, $\|\nabla C\|/C_0$ saturates at a constant positive value for circuit runtimes beyond a characteristic timescale determined by the T_1 coherence time. This behaviour is observed robustly across circuit sizes from $N = 8$ to $N = 102$ qubits, multiple IBM hardware platforms, different circuit compilations, and different Hamiltonian topologies.

The ICLA method is used throughout as an efficient gradient estimator, requiring only $M = \mathcal{O}(m)$ circuit evaluations — a substantial reduction compared to direct gradient estimation — and is shown to be robust even when capped at $M = 200$ samples.

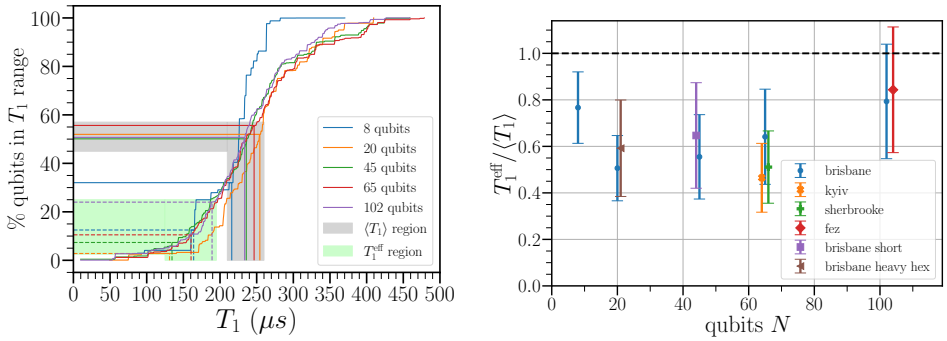

Figure 4.2: (Left)Cumulative distribution function (CDF) of T_1 values from `ibm_brisbane` against the percentage of qubits below each value. Solid (dashed) vertical and horizontal lines indicate the mean (effective) T_1 for each circuit size. Gray (green) shaded areas mark the mean (T_1^{eff}) regions. The effective coherence time corresponds to only 3%–25% of qubits, indicating that the worst-performing qubits dominate the onset of amplitude-damping effects. (Right) T_1^{eff} relative to $\langle T_1 \rangle$ as a function of qubit number N for all experimental setups. Colors denote different IBM devices; 'short' and 'heavy hex' indicate different circuit compilation and Hamiltonian topology, respectively. All T_1^{eff} are significantly below $\langle T_1 \rangle$.

An effective coherence time $T_1^{\text{eff}} \approx 0.72\, t_{cir}^{\text{flat}}$ is extracted from the flattening point of the gradient signal. In all experiments, T_1^{eff} is found to be substantially smaller than the mean device T_1, with only $\sim$ 3%–25% of qubits having coherence times at or below T_1^{eff}. This implies that the worst-performing qubits in the circuit dominate the onset of amplitude-damping effects, and that standard calibration metrics based on mean T_1 values systematically overestimate the useful circuit runtime.

These findings have important practical implications. First, NIBP are not a realistic concern for quantum hardware with any non-trivial amount of T_1 relaxation — which encompasses all existing physical platforms. Second, the emergence of a noise-induced limit set under amplitude damping imposes a finite effective circuit depth p_{eff}: circuits deeper than this depth are effectively equivalent to shallower ones, limiting expressiveness. This calls for re-evaluating warm-starting and parameter-transfer schemes in QAOA and related algorithms. Third, the methodology introduced here — combining ICLA with the analytical lower bound from Ref. [134] — provides a practical diagnostic for estimating the effective T_1 of a device under realistic computational conditions, complementing standard calibration protocols.

4.2 Resources in quantum optimization

This section investigates the resource requirements of QAOA through the lens of the resource theory of nonstabilizerness, also known as magic. The material presented here is based on [27].

As discussed in Sec. 1.1, entanglement has been extensively studied as a quantum resource in the context of QAOA and quantum annealing [139, 140, 141, 142, 143, 144, 145, 146, 147, 47]. However, entanglement alone does not ensure a quantum advantage. Stabilizer states can exhibit large entanglement yet remain efficiently simulable classically via Clifford circuits, as guaranteed by the Gottesman–Knill theorem [40, 48, 148].

As introduced in Section 1.1, nonstabilizerness — the degree to which a quantum state deviates from the set of stabilizer states — is a crucial ingredient for computational quantum advantage [149, 150, 151, 152, 153]. Despite this, its role in variational quantum optimisation remains largely unexplored.

This chapter addresses this gap by analysing nonstabilizerness in QAOA applied to the paradigmatic Sherrington–Kirkpatrick (SK) spin-glass model [154]. The main findings are:

- the existence of a *magic barrier* — a transient build-up of nonstabilizerness during the QAOA run, akin to previously observed entanglement barriers [140, 155, 47];
- a universal scaling collapse of the magic curves across different circuit depths;
- a characteristic and analytically explicable relationship between the final nonstabilizerness and the success fidelity;
- the presence of an analogous magic barrier in adiabatic quantum annealing.

Both qubit and qutrit versions of the algorithm are studied, using both the Stabilizer Rényi Entropy (SRE) and Mana as measures of nonstabilizerness.

The results have practical relevance: measurements of nonstabilizerness can serve as a figure of merit for estimating the performance of variational algorithms on concrete quantum hardware, including superconducting qubits [156], trapped ions [53], and Rydberg atoms [8].

4.2.1 QAOA on the Sherrington–Kirkpatrick Model

As described in Section 1.2.3, QAOA prepares a trial state $|\psi(\boldsymbol{\beta},\boldsymbol{\gamma})\rangle$ by applying p layers of alternating unitaries $U = \exp\left(-i\beta_p \hat{H}_M\right)\exp\left(-i\gamma_p \hat{H}_C\right)$, starting from the ground state of the mixer Hamiltonian $\hat{H}_M$. For $\hat{H}_M = -\sum_i \hat{X}_i$, the initial state is $|+\rangle^{\otimes N}$. A classical optimiser then minimises $\langle\psi|\hat{H}_C|\psi\rangle$ by tuning the variational parameters. QAOA can be generalised to qudit systems with any local dimension d [109], as also described in Section 1.2.3.

Since the initial state $|+\rangle^{\otimes N}$ is a stabilizer state ($\mathcal{M}(|+\rangle) = 0$) and the solution state of a classical problem with a unique optimum is also non-magical, QAOA offers a natural setting to probe the build-up and role of nonstabilizerness during quantum optimisation.

The model of interest is the SK model, a variant of the Ising model with all-to-all interactions drawn from independent Gaussian distributions, introduced as a solvable mean-field model for the spin-glass phase [157, 154] and serving as a prototype for real-world optimisation problems [158, 159]. The SK Hamiltonian generalised to qudits reads

$$\hat{H}_{\mathrm{SK}} = \sum_{i\neq j} J_{ij}(Z_i Z_j^\dagger + Z_i^\dagger Z_j) + \hat{H}_{\mathrm{bias}}\,, \tag{4.12}$$

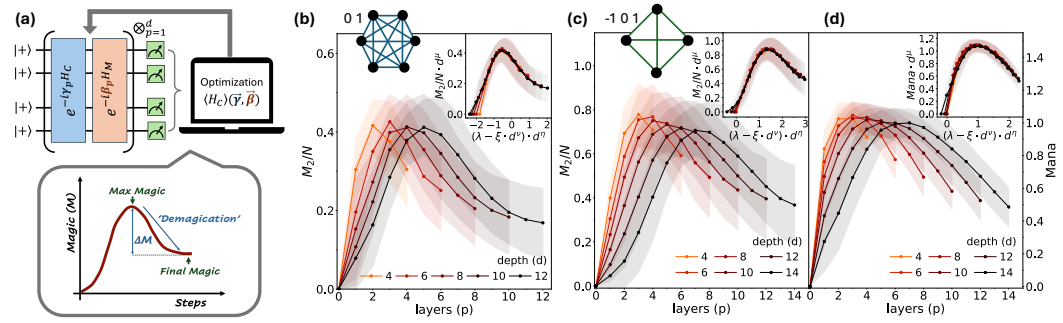

Figure 4.3: (a) QAOA scheme and a pictorial representation of the magic barrier: initially, nonstabilizerness rises to a maximum, after which it decreases — an effect referred to as "demagication." (b) SRE density as a function of the layer index during a QAOA protocol, for a system of 6 qubits and different depths (i.e. total number of layers). Inset: using a simple scaling function, the magic barrier collapses onto a universal curve. (c,d) Same as panel (b) for a system of 4 qutrits, showing SRE density (c) and Mana (d) as a function of the layer index. The universal collapses in the insets of (b–d) use the scaling function $\mathcal{M} = d^{-\mu} f[(\lambda - \xi \cdot d^{\nu}) \cdot d^{\eta}]$, with (μ, ξ, ν, η) equal to $(5.11 \cdot 10^{-3}, 1.28, -0.34, 0.60)$, $(-8.20 \cdot 10^{-2}, 1.19 \cdot 10^{-5}, 3.49, 0.46)$, $(-2.86 \cdot 10^{-2}, 3.69 \cdot 10^{-8}, 5.50, 0.34)$, respectively.

where the coefficients J_{ij} are drawn from a normal distribution with unit variance, Z_i is the generalised phase operator, and the Hermitian conjugate term ensures Hermiticity in the qutrit case. Bias magnetic-field terms $\hat{H}_{\text{bias}}$ are included to remove ground-state degeneracies; their explicit construction is detailed in App. C.2.

For each random realisation of the SK model, QAOA is simulated using exact numerics. The best result from 20 independent runs initialised using Trotterised quantum annealing [160] is selected, and the results are averaged over 50 realisations.

4.2.2 The Magic Barrier

When targeting the non-degenerate ground state of the SK model, QAOA starts — and ideally ends — in stabilizer states with zero magic. The relevant performance metrics — relative energy $(E_{\text{QAOA}}(p) - E_{\text{exact}})/E_{\text{exact}}$ and fidelity with the ground state $\mathcal{F}(p) = |\langle \psi_{\text{exact}} | \psi_{\text{QAOA}}(p) \rangle|^2$ — improve monotonically with the number of QAOA layers p (see App. C.5).

In contrast, the evolution of nonstabilizerness follows a non-monotonic pattern: it initially rises rapidly as the first layers are applied, reaches a maximum, and then decreases as the algorithm converges. This structure is defined as the *magic barrier*.

A pictorial representation of the magic barrier is shown in Fig. 4.3(a). This phenomenon is observed consistently across different systems and for both SRE and Mana. Figure 4.3(b) depicts the evolution of the SRE density M_2/N for a 6-qubit system at varying QAOA depths. A distinct peak emerges at approximately half the total depth, indicating the point of maximal nonstabilizerness. A comparable scenario arises in the 4-qutrit system, for both SRE and Mana [Fig. 4.3(c,d)].

Two notable properties of the magic barrier deserve attention. First, for a fixed system size, the maximum value of nonstabilizerness remains approximately constant across different QAOA circuit depths, suggesting that the magic peak is independent of circuit depth.

Second, even at its peak, the generated magic remains below that of a Haar-random state in the corresponding Hilbert space — $M_2^{\text{Haar}} = -\log_2\left(\frac{4}{2^N+3}\right)$ for qubits and $-\log_2\left(\frac{3}{3^N+2}\right)$ for qutrits [161, 162] — suggesting that a device running QAOA needs to build up only a limited amount of magic.

4.2.3 Universal Scaling of the Magic Barrier

The systematic and consistent presence of the magic barrier motivates the search for a scaling law capturing the universal features of magic evolution in QAOA. As shown in the insets of Fig. 4.3(b–d), the data for different depths collapse onto a single curve under the rescaling

$$\frac{\mathcal{M}\big(|\psi_{\text{QAOA}_d}(p)\rangle\big)}{N} = d^{-\mu}\, f[(\lambda - \xi \cdot d^{\nu}) \cdot d^{\eta}]\,, \quad \mu, \xi, \nu, \eta \in \mathbb{R}, \quad \lambda = \frac{p}{d}\,. \tag{4.13}$$

For qubit systems, the scaling exponent μ is close to zero, suggesting that the peak SRE is independent of both circuit depth and system size. For qutrit systems, μ is slightly larger, potentially indicating a mild dependence on circuit depth. The coefficients ξ and ν appear to be inversely correlated. For the qubit case, the scaling simplifies by replacing $\xi \cdot d^{\nu}$ with a constant critical point λ_c, yielding $d^{-\mu} f[(\lambda - \lambda_c) \cdot d^{\eta}]$. The critical point $\lambda_c \approx 0.2$ is robust across different system sizes (see App. C.3). The exponent η, governing the width rescaling, lies within a narrow range and increases with system size, suggesting a sharper magic barrier in larger systems. Detailed scaling tables for $N = 4, 6, 8$ qubits are provided in App. C.3.

4.2.4 Final Nonstabilizerness and Fidelity

The ultimate measure of QAOA success is the fidelity of the final state with the target solution. As the algorithm approaches the final layer, successful instances reach high overlap with a classical state, implying — in the absence of degenerate solutions — low magic. This behavior is indeed observed in randomly sampled instances for different system sizes and depths, shown in Fig. 4.4. The converse, however, does not hold: a QAOA sweep can end in a stabilizer state distinct from the solution and thus reach low fidelity with low magic.

Focusing on instances with $\mathcal{F} \geq 0.5$ [Fig. 4.4(b)], a void region at the lower axis is apparent. At the end of a rather successful QAOA run, the state is expected to consist of the solution $|\phi_0\rangle$ plus a small contamination, typically from the first excited state $|\phi_1\rangle$. The following two-state *ansatz* is therefore considered:

$$|\psi_{\text{II}}\rangle = \sqrt{\mathcal{F}}\,|\phi_0\rangle + e^{i\theta}\sqrt{1-\mathcal{F}}\,|\phi_1\rangle\,, \quad \theta \in [0, 2\pi]\,, \tag{4.14}$$

where $\sqrt{\mathcal{F}}$ is the probability amplitude to reach the desired solution and $e^{i\theta}$ is a relative phase. Using the additivity of M_2 and permutational invariance, it can be shown that for any number of qubits (see App. C.4),

$$M_2(|\psi_{\text{II}}\rangle) = -\log_2\big[1 - 4\mathcal{F}(1-\mathcal{F}) + 2\mathcal{F}^2(1-\mathcal{F})^2(7+\cos 4\theta)\big]\,. \tag{4.15}$$

This result is symmetric around $\mathcal{F} = 0.5$. The minima ($\theta = 0$) and maxima ($\theta = \pi/4$) define a bounded region in the fidelity–magic plane that is independent of qubit number and circuit depth, within which all instances matching this *ansatz* must lie.

A refined *ansatz* with three computational basis states,

$$\begin{aligned} |\psi_{\text{III}}\rangle = \sqrt{\mathcal{F}}\,|\phi_0\rangle + e^{i\theta_1}\sqrt{p}\,|\phi_1\rangle + e^{i\theta_2}\sqrt{1-\mathcal{F}-p}\,|\phi_2\rangle\,, \\ p \in [0, 1-\mathcal{F}],\ \theta_{1/2} \in [0, 2\pi]\,, \end{aligned} \tag{4.16}$$

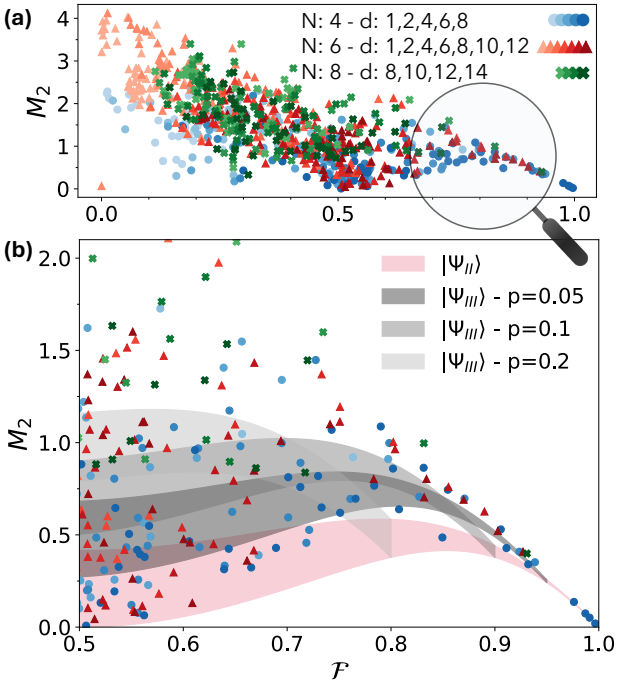

Figure 4.4: Nonstabilizerness of the final QAOA state as a function of final fidelity. (a) Randomly sampled data for varying numbers of qubits and QAOA depths. (b) Zoom into the region $\mathcal{F} \geq 0.5$. The red shaded region delimits the SRE of any n-qubit system described by a superposition of two computational basis states. The grey regions collect the possible SRE values for a wavefunction spanning three basis states. The empty region at the bottom is analytically forbidden, as explained in the text.

for which M_2 is also computed analytically (App. C.4), further fills the accessible region. Since *ansätze* with more basis states generate higher magic, these considerations explain the numerically observed empty region at the bottom of Fig. 4.4(b), and further suggest that obtaining a medium-to-high fidelity state $\mathcal{F} \sim 0.6$–0.9 with low magic is unlikely, reinforcing the conclusion that QAOA must traverse high-magic states to reach good solutions.

4.2.5 Demagication and Success Probability

A natural question is whether the *demagication* — the drop in SRE after the barrier, $\Delta M = M_2^{\max} - M_2^{\text{final}}$ — is correlated with the success fidelity. To address this, the conditional probability of obtaining fidelity above a threshold f_{th}, given a minimum demagication $\Delta M > \epsilon$, is computed:

$$\mathcal{P}_{\text{cond}}\left[F > f_{\text{th}} \,\middle|\, \Delta M > \epsilon\right] = \frac{\mathcal{P}[(F > f_{\text{th}}) \cap (\Delta M > \epsilon)]}{\mathcal{P}[\Delta M > \epsilon]} \,. \tag{4.17}$$

As shown in Fig. 4.6, a larger demagication (larger ϵ) results in a higher probability of solving the QAOA with high fidelity. The probability of reaching or surpassing a given fidelity also increases with depth, as expected. Analogous results for the qutrit case are reported in App. C.8.

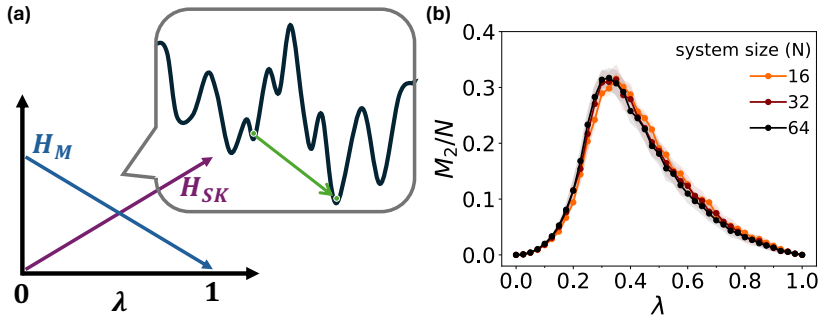

Figure 4.5: (a) Scheme of the quantum annealing protocol: as λ increases, the cost Hamiltonian H_{SK} gains strength while the mixer H_M weakens. Pictorially, quantum annealing reaches the solution by tunnelling between local minima. (b) SRE density computed on the ground state of the instantaneous Hamiltonian — corresponding to ideal adiabatic quantum annealing — as a function of the normalised annealing time $\lambda \in [0, 1]$. A magic barrier is evident around $\lambda \sim 0.35$. The near-overlapping curves for $N = 16, 32, 64$ indicate that the SRE density traces a barrier largely independent of system size, consistent with the QAOA observations.

4.2.6 Magic Barrier in Quantum Annealing

To assess the generality of the magic barrier beyond QAOA, its emergence is examined in continuous-time quantum annealing protocols [163, 164]. Quantum annealing initializes the system in the ground state of the transverse-field mixer H_M and gradually transforms the Hamiltonian into the problem-specific cost Hamiltonian H_C, whose ground state encodes the solution [163, 165]:

$$H(\lambda) = (1 - \lambda)H_M + \lambda H_C \,, \tag{4.18}$$

where $\lambda \in [0, 1]$ parameterizes progress along the annealing path. In the adiabatic limit, the system remains in the instantaneous ground state of $H(\lambda)$ throughout.

The stationary regime is simulated using matrix product state (MPS) methods [166, 167, 168, 169], which allow numerically exact computation of the ground state $|\psi(\lambda)\rangle$ of the instantaneous Hamiltonian at arbitrary points along the interpolation (bond dimension $\chi = 60$). SRE is computed efficiently using the Pauli Matrix Product State (Pauli MPS) formalism [170], with bond dimension up to $\chi_P = 1024$. The fully connected spin-glass interactions are truncated to fifth-neighbour couplings; this truncation does not qualitatively affect the key features of the observed magic dynamics.

The resulting SRE profile, shown in Fig. 4.5, reveals a clear magic barrier: nonstabilizerness rises during the early sweep, peaks at intermediate λ, and decreases as the system approaches the classical solution. This behaviour closely mirrors the magic evolution in QAOA. Moreover, based on earlier studies of entanglement growth in non-adiabatic protocols [140], deviation from perfect adiabaticity would generally enhance the amount of magic generated during the sweep.

4.2.7 Conclusions

This chapter has analysed the role of nonstabilizerness in quantum optimisation, focusing on QAOA applied to the SK spin-glass model and extended to continuous-time quantum

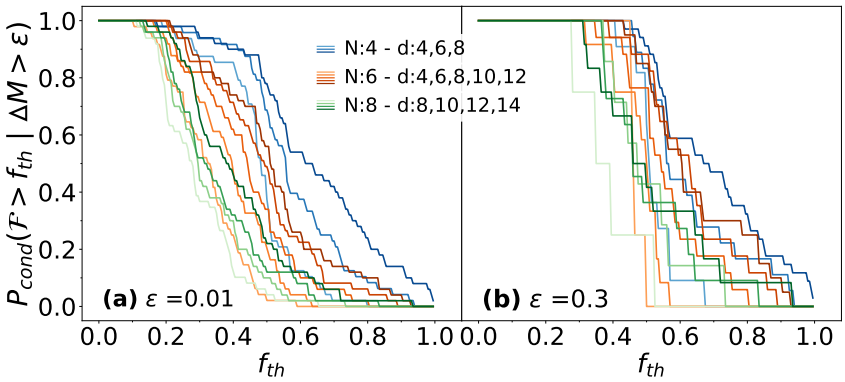

Figure 4.6: Conditional probability of reaching a high final fidelity $\geq f_{\text{th}}$ given a demagication ΔM larger than a fixed value ϵ, for various system sizes and QAOA depths. (a) $\epsilon = 0.01$; (b) $\epsilon = 0.3$. The rightward shift from panel (a) to (b) indicates that stronger demagication leads to a higher probability of achieving high fidelity. Darker shades denote higher depths.

annealing.

The central finding is the existence of a *magic barrier*: irrespective of the circuit depth, nonstabilizerness rises to a peak near the middle of the QAOA protocol and, in well-performing runs, falls towards the end. This barrier is consistently observed for both SRE and Mana, in qubit and qutrit systems alike. Upon rescaling, the average magic collapses onto a universal curve across different depths, described by the scaling form of Eq. (4.13). For qubit systems, a robust critical point $\lambda_c \approx 0.2$ is found, consistent across system sizes.

The relationship between final nonstabilizerness and fidelity is non-trivial: while high fidelity implies low final magic, low magic does not guarantee high fidelity. An analytically tractable *ansatz* based on superpositions of two or three computational basis states explains the characteristic structure of the fidelity–magic plane and the existence of a forbidden region at low magic and high fidelity. The conditional probability analysis further reveals that larger demagication is associated with higher probability of achieving high fidelity, providing a practical figure of merit for assessing algorithmic performance.

The magic barrier is not specific to QAOA: an analogous barrier is observed in ideal adiabatic quantum annealing, with a peak around $\lambda \sim 0.35$ that is largely independent of system size.

These findings have direct practical implications. In early fault-tolerant devices, where non-Clifford gates represent a major resource bottleneck [149, 171, 39, 172], the magic barrier constitutes an unavoidable extra resource cost. This motivates the question of whether QAOA can be parameterised to minimise magic generation while still reaching the solution. Measurements of nonstabilizerness on hardware — accessible, e.g., on trapped ions [53] — can provide a concrete diagnostic for algorithmic performance, complementing energy-based figures of merit.

Chapter 5

First order phase transition in a discrete nonabelian LGT

In this chapter, we report the findings obtained by analyzing the ground state properties of a LGT with symmetry group D_3. After introducing the model, we identify the possible ground states across different parameter regimes perform a numerical analysis of the ground state, from which we extract the relevant quantities needed to characterize the phase transition. All the results presented in this chapter are preliminary. Discussion on how to complement and expand them is presented in the conclusions of the chapter.

The D_3 gauge theory is motivated by two key physical reasons. First, it constitutes the minimal lattice gauge theory model in which non-Abelian string breaking can be observed without truncating the gauge group, making it an ideal testbed for studying confinement phenomena beyond the Abelian setting. Second, the tuning of the representation content of the model allows one to investigate false vacuum decay, confinement and string breaking dynamics in a genuinely non-Abelian context, generalizing the framework already developed for abelian groups [173] to the non-Abelian case. Together, these features make the D_3 model a natural step towards a broader understanding of non-Abelian confinement dynamics in low-dimensional quantum field theories.

5.1 Model

In this section, we focus on a 1+1D LGT with gauge group D_3. We take matter to be in the faithful representation, which we label as τ, with $\dim(\tau) = 2$. The system is governed by a Kogut-Susskind Hamiltonian given in Eq. (1.69), with the exception that the plaquette term is not present given the 1D spatial nature of the problem

$$H = M \sum_i (-1)^i \psi_i^\dagger \psi_i + J \sum_i \psi_i U_{i,i+1} \psi_{i+1}^\dagger + \sum_i E_i^2. \tag{5.1}$$

Let us focus on the electric field term. We have seen that it is uniquely defined by the coefficients $\alpha(j)$ present in (**??**). In particular, these values represent the eigenvalues of the states in the different representations and can be tuned in order to set a specific representation as the one with the lowest energy. For the case $J = 0$, we can find the exact ground state of the theory based on the ordering of the different representations.

We consider two specific orderings:

$$\alpha(e) < \alpha(\tau) \ll \alpha(p) \tag{5.2}$$
$$\alpha(\tau) < \alpha(e) \ll \alpha(p) \tag{5.3}$$

In the first case, the field favors the trivial representation and thus the state that minimizes the energy is

$$|\psi_e\rangle = |0+2\rangle^{\otimes \frac{N}{2}}, \qquad H|\psi_e\rangle = ((N-1)\alpha(e) - NM)|\psi_e\rangle \equiv E_e|\psi_e\rangle. \tag{5.4}$$

This situation is rather trivial, as the matter states are empty and the field always transforms trivially. A more interesting case is the one corresponding to the second ordering. In this case, the field favors the fundamental representation, but due to gauge invariance we have two candidates for the ground state. Depending on the specific values of $\alpha(e)$ and $\alpha(\tau)$, the possible choice for the ground state is between the two following states:

$$|\psi_e\rangle = |0+2\rangle^{\otimes \frac{N}{2}} \tag{5.5}$$
$$|\psi_\tau\rangle = |1\tau 2\rangle \otimes |\tau 0\tau 2\rangle^{\otimes \frac{N-4}{2}} |\tau 0\tau 1\rangle \tag{5.6}$$

where the second state has energy $E_\tau = (N-1)\alpha(e) - (N-4)M$. This state has all the fields polarized along the faithful representation and is empty on all the mass sites of the bulk and due to gauge invariance is single occupied on the edges. For $|\psi_\tau\rangle$ to be the ground state of the theory, we require $E_\tau < E_e$ or

$$\alpha(e) - \alpha(\tau) > \frac{4M}{N-1}. \tag{5.7}$$

In the following, we parametrize $\alpha(e)$ and $\alpha(\tau)$ and study the transition of the ground state for $J \neq 0$. We consider $\alpha(e) = -\alpha(\tau) = -\alpha, \alpha_p = 5$, which transforms (5.7) into

$$\alpha < -\frac{2M}{N-1} \equiv \alpha^*. \tag{5.8}$$

5.2 Numerical analysis

Hamiltonian (6.11) was implemented using the ITensors package [174] in the Julia programming language, which provides a high-level interface for tensor network construction and manipulation. Ground state simulations were then carried out using the density matrix renormalization group (DMRG) algorithm, as implemented in the built-in `dmrg` function of the ITensors library. The interplay between the different ground states was studied by varying the parameter α over a chosen set, once a fixed value of J was fixed. In systems undergoing a first-order phase transition, the free energy landscape generically develops two distinct local minima separated by a potential barrier of higher energy. The coexistence of such metastable configurations implies that the ground state reached by the algorithm may depend on the history of the sweep rather than being uniquely determined by the instantaneous value of α. This sensitivity to the adiabatic path is a hallmark of first-order transitions and necessitates probing the parameter space in both directions in order to detect metastable branches and to delineate the hysteresis region within which the two sweep directions yield inequivalent ground states. Two complementary sweeping sequences over the parameter set $\{\alpha_{\min}, \ldots, \alpha_{\max}\}$ were therefore performed. In the first sequence, α was initialized at its

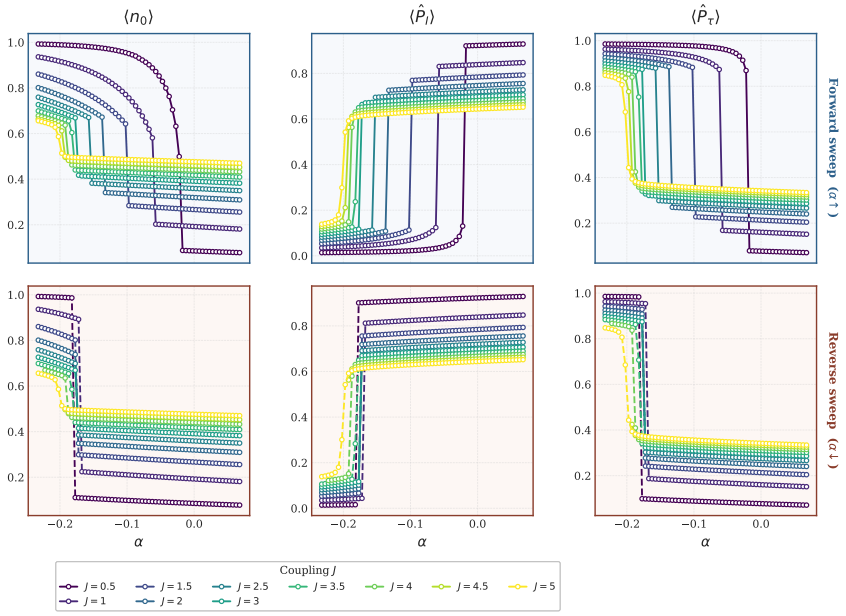

Figure 5.1: Hysteresis of the edge occupation and symmetry-sector projections for a chain of $n = 32$ sites with fixed parity, as a function of the parameter α. The top row shows results obtained by sweeping α upward (forward sweep), while the bottom row corresponds to the reverse sweep (α decreasing). Each panel displays a different observable: the average edge-site occupation $\langle n_0 \rangle$ (left), the projection onto the identity sector $\langle \hat{P}_I \rangle$ (center), and the projection onto the τ-parity sector $\langle \hat{P}_\tau \rangle$ (right). Colors encode the coupling strength J as indicated in the legend. The persistence of distinct values between the two sweeps signals hysteretic behavior and the coexistence of competing phases in the region of parameter space explored.

maximum value $\alpha_{\max}$ and decreased monotonically to its minimum $\alpha_{\min}$; in the second, the procedure was reversed, starting from $\alpha_{\min}$ and increasing it to $\alpha_{\max}$. In both cases, the converged ground state wavefunction obtained at each value of α was used as the initial state — that is, as the initial matrix product state (MPS) ansatz — for the subsequent DMRG run at the next value of α. This adiabatic-continuation-like strategy improves convergence and helps track the ground state continuously across the parameter space, reducing the risk of the algorithm converging to excited states or undesired metastable configurations.

The DMRG calculations were performed with a maximum bond dimension of $\chi = 200$. Convergence was assessed by monitoring the truncation error and the energy variance across a minimum of 20 sweeps; convergence threshold for the energy was set to 10^{-6}. Simulations were performed with maximum DMRG noise equal to 10^{-3}.

The procedure was repeated for three different sizes: $N = 16, 32, 64$ and for various values of J.

Once DMRG finds the ground state, we compute three observables: the local density on the first and last site (which we expect to change at the transition, differently than the one in the bulk) and the mean projection of the link states over the different representations. In terms of on-site operators, they are defined as

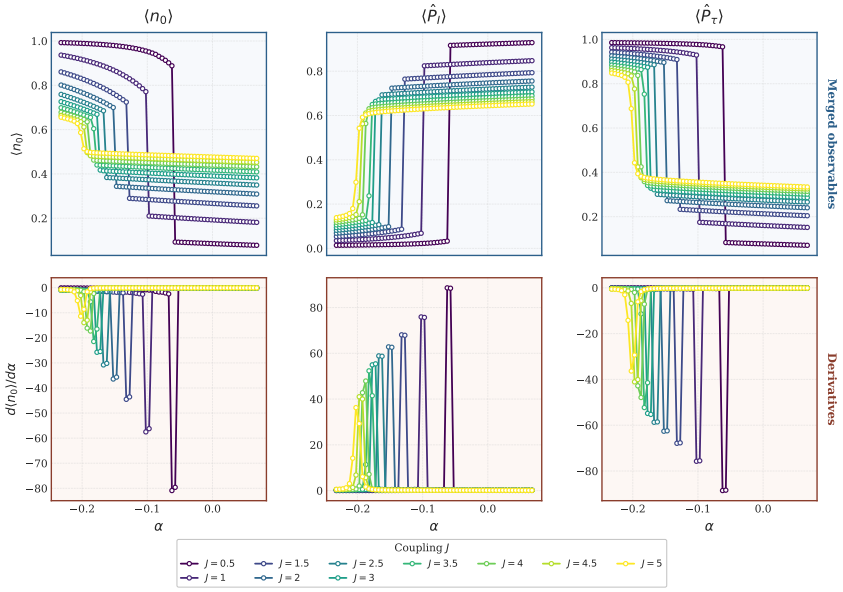

Figure 5.2: **Top:** Key observables as a function of α for several values of J. From left to right: occupation number on the leftmost site; mean electric field projected onto the trivial representation; mean electric field projected onto the faithful representation.
Bottom: Numerical derivatives of the corresponding curves above, in the same order. The plots refer to system size $N = 32$.
The figure shows how the transition point is decreased as the strength of the hopping J is increased. Moreover, the increase in J also reduces the jump in the observables, indicating a suppression of the transition towards $J \to +\infty$.

$$n = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix} \qquad P_e = \mathrm{Q} \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \mathrm{Q}^\dagger \tag{5.9}$$

$$P_\tau = \mathrm{Q} \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \mathrm{Q}^\dagger \qquad P_p = \mathrm{Q} \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \mathrm{Q}^\dagger \tag{5.10}$$

Fig. 5.1 reports the ground state expectation values of the observables as a function of the sweeping parameter α, for both directions of the sweep and $N = 32$. We note that the two sweeps (forward and reverse) over the same parameter set provide different curves for the expectation values of the observables mentioned in the previous section. In particular we see that even if the state is supposed to have crossed the transition point, it remains

stuck in a local minima corresponding to the ground state of the previous phase. This can be easily explained by considering that the there is a region of the α, J parameter space where the candidates ground states correspond to local minima also after the transition. Once this regime is exited, the two curves collapse.

It is possible to merge the two curves by selecting the state of lowest free energy for each combination of α and J, thereby identifying the true thermodynamic transition point. The resulting phase boundary is shown in Fig. 5.2, which illustrates how the transition point shifts from α^* given by (5.8) as a function of the coupling strength J.

The extent of the metastable region depends on both the system size N and the coupling parameter J, as shown in Fig. 5.3. For small values of N, the metastable region is narrow and the two branches of the transition curve converge rapidly, indicating that the free-energy barriers separating the competing minima remain shallow. As N increases, the metastable region broadens progressively, reflecting the deepening of the respective free-energy wells and the consequent increase in the value of J at which the two minima become degenerate. Furthermore, the value of α^* at which the merged curve asymptotically stabilizes is observed to scale as $\mathcal{O}(1/N)$, consistent with the typical scaling given by finite-size corrections. An exception is noted for $N = 16$, where deviations from this scaling are more pronounced, likely attributable to boundary effects that are disproportionately significant at small system sizes.

In order to identify the phase transition and study its scaling with the system size, we compare the first and second derivatives of the energy as a function of the sweeping parameter α for the three system sizes considered. Results are reported in Fig. 5.4 and explained below.

$N = 16$

For the smallest system size considered, $N = 16$, the first derivative of the ground-state energy per site $dH/d\alpha \,/\, N$ already reveals a clear signature of the transition. Across all coupling strengths J, the derivative exhibits a pronounced step-like drop as α increases toward α^*, transitioning from a plateau near unity to a lower plateau whose value depends on J. Weaker couplings (small J) show a smoother and more gradual descent occurring at more negative values of α, while stronger couplings drive the transition closer to $\alpha = \alpha^*$ and with greater steepness. The second derivative $d^2H/d\alpha^2 \,/\, N$ corroborates this picture, displaying sharp negative peaks localised at the transition point. At this system size the peaks, though pronounced, remain finite and somewhat broadened, reflecting the rounding of the singularity due to finite-size effects.

$N = 32$

Increasing the system size to $N = 32$ sharpens the signatures observed at $N = 16$. The step in the first derivative becomes more abrupt and the transition point shifts noticeably closer to $\alpha = \alpha^*$ for all values of J, consistent with the expected finite-size drift toward the thermodynamic critical point. The plateaux on either side of the transition are better resolved, and the dependence on J remains systematic: larger couplings push the critical α toward α^* and deepen the post-transition plateau. In the second derivative, the negative peaks grow significantly in magnitude compared to $N = 16$ and become more localised, indicating that the underlying singularity is becoming better resolved as the system size increases.

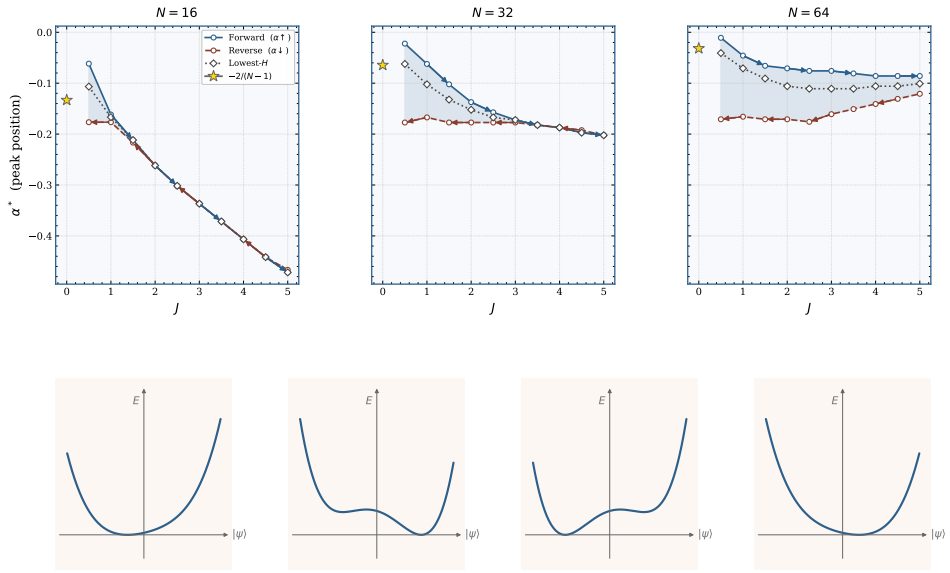

Figure 5.3: **Top:** Phase transition boundary for system sizes $N = 16$ (left), $N = 32$ (center), and $N = 64$ (right), plotted as a function of the parameter set $\{\alpha\}$. The red (blue) curve represents the transition point obtained via a forward (reverse) sweep in parameter space. The two curves differ due to the presence of metastable local minima distinct from the global minimum, causing the system to become trapped in a local potential well during the sweep. The gray line is instead obtained by considering both of the curves and selecting only the points for which the energy of the system was lower (i.e. always considering the true ground state). The star in each plot indicates the value of α^* for a the given size.
Bottom: Schematic illustration of the putative effective potential landscape of the theory. At large values of J (leftmost and rightmost panels), the secondary minimum is suppressed and the system exhibits no hysteresis. At intermediate values of J, the potential develops two competing minima corresponding to the states $|\psi_e\rangle$ and $|\psi_\tau\rangle$. In this regime, the DMRG algorithm becomes trapped in one of the two minima, and the transition is only triggered once the potential barrier becomes sufficiently large that the metastable minimum is destabilized.

$N = 64$

At $N = 64$ the finite-size rounding is further suppressed and the derivatives approach the behaviour expected in the thermodynamic limit. The first derivative displays an increasingly discontinuous drop, with the transition concentrated in a very narrow window of α near α^*. The second derivative develops sharp, deeply negative peaks whose magnitude is substantially larger than at smaller system sizes, consistent with a divergence in the thermodynamic limit that is characteristic of a quantum phase transition. The systematic trend across $N = 16$, 32, and 64 — shrinking transition width, growing peak magnitude, and drift of the critical point toward $\alpha = \alpha^*$ — provides strong numerical evidence that the observed features are genuine critical signatures rather than finite-size artefacts.

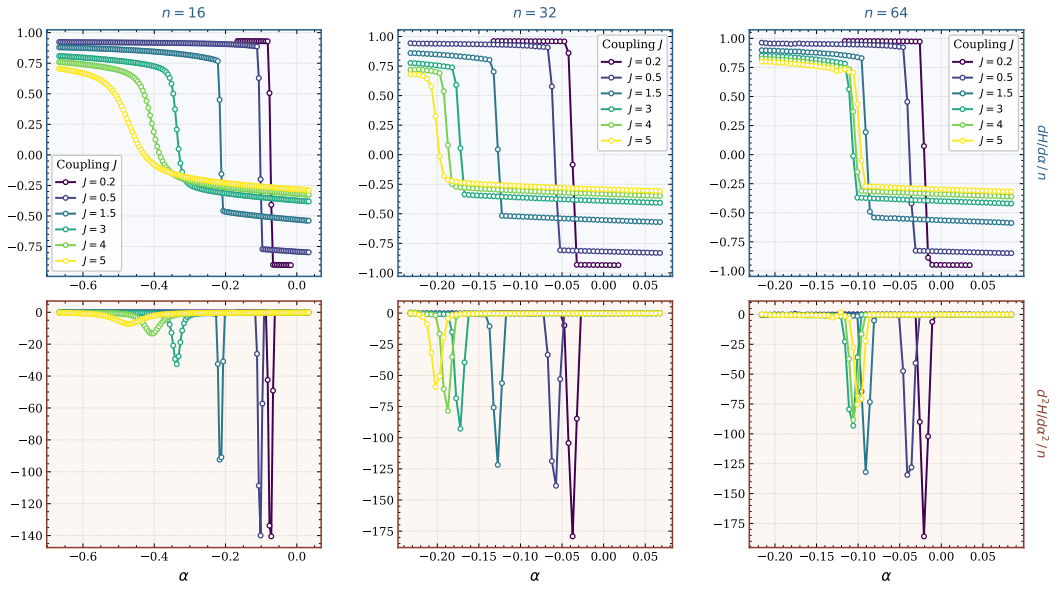

Figure 5.4: First and second derivatives of the ground-state energy per site with respect to the detuning parameter α, computed from the lowest-energy state obtained by merging forward and reverse hysteresis sweeps. Results are shown for system sizes $n = 16$, 32, and 64 (columns) and coupling strengths $J = 0.2, 0.5, 1.5, 3, 4, 5$ (colours, viridis palette).
Top row: $dH/d\alpha \,/\, N$; The first derivative exhibits a sharp step-like drop whose position shifts toward $\alpha = \alpha^*$ with increasing n, indicating a transition that sharpens with system size.
Bottom row: $d^2H/d\alpha^2 \,/\, N$. The second derivative develops pronounced negative peaks at the same locations, whose magnitude grows with both n and J, consistent with the divergence expected at a phase transition in the thermodynamic limit.

5.3 Conclusions and Outlook

In this chapter we have studied the ground state properties of a $1+1$D lattice gauge theory with non-Abelian gauge group D_3, with matter in the faithful two-dimensional representation τ. Starting from the exact solution at $J = 0$, we identified two competing ground states, $|\psi_e\rangle$ and $|\psi_\tau\rangle$, whose relative stability is controlled by the electric field parameter α through the condition (5.7). The parametrisation $\alpha(e) = -\alpha(\tau) = -\alpha$ reduces this to a simple criterion involving the system size N, providing a clean handle on the transition.

For $J \neq 0$, the transition was studied numerically using DMRG across three system sizes $N = 16, 32, 64$ and a range of coupling strengths J. The use of bidirectional sweeps in α revealed a pronounced hysteretic response, whereby the ground state found by the algorithm depends on the history of the sweep rather than the instantaneous value of α alone. This behaviour is consistent with a first-order quantum phase transition, in which the free energy landscape develops two competing local minima separated by a barrier. The metastable region widens with increasing N and varies non-monotonically with J, reflecting the interplay between hopping-induced fluctuations and the electric field energy scale.

By merging the two sweep branches and retaining the lowest-energy state at each (α, J) point, we reconstructed the true ground-state energy and computed its first and second

derivatives with respect to α. The first derivative exhibits a sharp step whose position drifts toward the analytically predicted α^* as N increases, while the second derivative develops increasingly large and localised negative peaks at the same location. The systematic sharpening of these signatures across $N = 16$, 32, and 64 is consistent with the divergence expected at a first-order transition in the thermodynamic limit, and provides strong numerical evidence that the observed features are genuine critical phenomena rather than finite-size artefacts.

Several directions remain open. A finite-size scaling analysis of the peak positions and magnitudes in $d^2H/d\alpha^2$ would allow a quantitative extrapolation of the critical point $\alpha_c(J)$ to the thermodynamic limit and a more precise characterisation of the transition. The dependence of the metastable region on J deserves further investigation, in particular to determine whether the transition is suppressed entirely at large J or whether a residual discontinuity persists. On the operator side, computing string order parameters and entanglement entropy across the transition would help to characterise the nature of the two phases in terms of confinement and string breaking, connecting the present results to the broader context of non-Abelian confinement dynamics that motivates the D_3 model. Finally, extending the analysis to larger system sizes and higher bond dimensions would strengthen the finite-size scaling conclusions and test the robustness of the DMRG results against truncation errors.

Chapter 6

Confinement and resources in nonabelian LGTs

This chapter presents an analysis of static confinement and the quantification of quantum resources in the preparation of ground states of Lattice Gauge Theories (LGTs), stemming from collaborations in which the author contributed in a supporting role. The results are largely based on [28] and [29].

On the topic of confinement, the author's contributions include the development of the theoretical framework for dihedral LGTs, encompassing field formulations in both the representation and group bases, as well as their decomposition in terms of rishon operators.

Concerning quantum resources, the author contributed the theoretical background on the symmetry groups under consideration, the dual models, and the mappings relevant to the $SU(2)$ case. Additional contributions include numerical computations of $\mathbb{Z}_N$ ground states and the corresponding resource quantifiers.

6.1 Confinement

This chapter studies the confinement properties of pure D_N lattice gauge theories in the strong-coupling regime, focusing on the role of the group center and fusion rules in determining whether static charges are confined or screened. The material is based on [28].

Among the central goals of high-energy theoretical physics is the understanding of the Standard Model of particle physics from first principles. In this context, the solution of non-Abelian gauge theories with SU(3) symmetry — the gauge group of quantum chromodynamics (QCD) — represents a long-standing challenge. In particular, the derivation of color confinement, namely the absence of isolated color-charged particles such as quarks, remains one of the major open problems in the field [73].

The non-perturbative nature of confinement has motivated the development of powerful classical numerical approaches, most notably Monte Carlo simulations of lattice gauge theories (LGTs) [175, 17, 176, 177, 178]. While these methods have provided invaluable insights into the equilibrium properties of gauge theories, they are subject to intrinsic limitations. In particular, they face severe difficulties in addressing real-time dynamics and regimes of large fermionic density, where the sign problem renders simulations impractical. Alternative non-perturbative approaches based on Hamiltonian formulations of LGTs have accordingly gained increasing attention, including quantum simulation platforms and tensor-network

methods, which are free from the sign problem and well suited to the study of real-time dynamics [19, 20, 179, 180, 181, 182, 21, 22, 89].

As introduced in Section 1.3, Hamiltonian approaches require suitable approximations that encode the continuous degrees of freedom of the gauge group into a finite-dimensional Hilbert space. Common strategies include truncations in representation space [183, 184, 185, 186, 100] or formulations in dual variables [187, 188, 189, 86, 87]. An alternative approach consists in replacing continuous gauge symmetries with discrete ones, giving rise to gauge theories based on finite groups that approximate their continuous counterparts [190, 191, 192, 193, 194, 195].

Discrete non-Abelian gauge theories share close connections with topologically ordered systems in two dimensions [18, 196, 197, 198, 199] and have recently enabled the first quantum simulations of dynamical string breaking in two spatial dimensions [200, 201, 202]. Among these models, gauge theories based on dihedral groups D_N provide one of the simplest non-Abelian settings, making them particularly attractive for both numerical and experimental investigations [19, 203, 204, 205, 206].

This chapter addresses the question of how deeply non-perturbative phenomena — such as confinement and string breaking — manifest in D_N gauge theories, and how these phenomena relate to the phenomenology of more complex theories such as QCD. The confinement properties of pure D_N LGTs in the strong-coupling regime are examined and it is shown that confinement is intimately connected to the structure of the group center. For even N, the dihedral group D_N possesses a non-trivial $\mathbb{Z}_2$ center, leading to the formation of stable non-Abelian electric flux strings that confine static charges over arbitrary distances. For odd N, the center is trivial and the electric string breaks even in the absence of dynamical matter: static charges are screened by gluonic excitations, forming bound states known as *gluelumps* in the context of QCD [73].

The minimal geometry required to observe this qualitative distinction is provided by a ladder lattice [93, 207, 208, 209] , which retains both electric and magnetic contributions while remaining amenable to quasi-one-dimensional numerical simulations. Tensor-network results are presented for D_3 and D_4 in the strong-coupling limit, supplemented by a scaling analysis of D_3 gluelumps that characterises their spatial extent and continuum-limit behaviour.

6.1.1 Fusion Rules and Center Symmetry

The mathematical foundation of the results presented in this chapter lies in the structure of the fusion rules of the dihedral groups D_N, which differ qualitatively depending on the parity of N. For finite groups, Clebsch–Gordan (CG) coefficients — and therefore fusion rules — are computed exactly from the characters of the group representations (see App. E.1).

The central object is the *screening fusion rule*, which holds exclusively for odd values of N:

$$\exists \text{ irrep } j \mid \quad j \otimes \tau \ni j \,, \tag{6.1}$$

where j is a non-Abelian irreducible representation ($\dim(j) > 1$) and τ denotes the fundamental representation. This rule implies that the gauge-field irrep j is left invariant when fused with a fundamental charge τ, so the corresponding electric flux can pass through the charge without being altered. For dihedral groups, Eq. (6.1) is satisfied by the faithful antifundamental irrep $j = \tilde{\tau}$ (see App. E.1). The associated CG coefficients are determined by the structure of the group center.[1] For odd N, the center of D_N is trivial, $Z = \{\mathbb{I}\}$, and the

[1]The center Z of a group G is the Abelian subgroup of elements commuting with all elements of G:

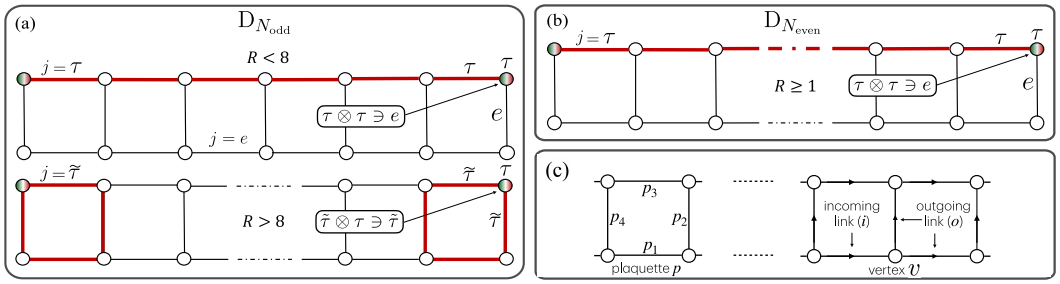

Figure 6.1: **(a)** GS at strong coupling of $D_{N_{\text{odd}}}$ LGTs on a ladder with static color charges on the corners. The string breaks when the charge separation R, in lattice units, is greater than 8, making the anti-fundamental $\tilde{\tau}-$glueballs energetically favorable. This si allowed thnks to the screening fusion rule of Eq. (6.1). **(b)** GS at strong coupling for a $D_{N_{\text{even}}}$ LGT. No strng is formed due to the absence of the screening fusion rule. **(c)** Ordering of links in a plaquette operator

relevant CG coefficient equals one. For even N, the center is $\mathbb{Z}_2$ and the same coefficient vanishes, rendering the screening fusion rule invalid.

This distinction has direct and far-reaching consequences for the physics of dihedral LGTs. For static probe charges in the fundamental representation τ, the screening fusion rule in $D_{N_{\text{odd}}}$ allows such charges to act simultaneously as sources and sinks of electric flux lines in the anti-fundamental representation $\tilde{\tau}$. Consequently, at strong coupling and large separations, static charges are individually screened by localised gluonic excitations — forming composite, gauge-invariant objects interpretable as glueball-dressed charges. For $D_{N_{\text{even}}}$ groups, the screening fusion rule does not hold: gauge invariance requires any pair of static fundamental charges to be connected by an electric flux string in representation τ, and confinement persists at arbitrarily large distances. These two qualitatively different configurations are illustrated in Fig. 6.1(a–b).

6.1.2 Models and Numerical Method

To investigate these phenomena quantitatively, the ground-state properties of D_3 and D_4 lattice gauge theories are studied using density matrix renormalization group (DMRG) simulations. Both groups possess a single two-dimensional fundamental representation, implying $\tilde{\tau} = \tau$ in both cases. Two sectors are considered: the neutral sector without static charges ($q = 0$), and a charged sector ($q = 1$) where two static charges occupy the upper corners of the ladder geometry. All simulations employ the ITensor Julia library [174, 210], with a maximum MPS bond dimension $\eta = 200$ and a truncation cutoff $\epsilon = 10^{-11}$.

The Hamiltonian formulation follows the Kogut–Susskind construction introduced in Section 1.3, with the electric and magnetic contributions described in App. E.3. For the electric Casimir energies, the values $\alpha_0 = 0$, $\alpha_\tau = 1$, $\alpha_p = 10$ are adopted throughout.

6.1.3 Phase Diagram and Crossover

In fully two-dimensional geometries, discrete LGTs are expected to undergo a phase transition between a confined phase at strong coupling and a topologically ordered deconfined

$Z = \{h \in G : \forall g \in G,\ hg = gh\}$. It is also a normal subgroup, satisfying $ghg^{-1} \in Z$ for all $g \in G$ and $h \in Z$.

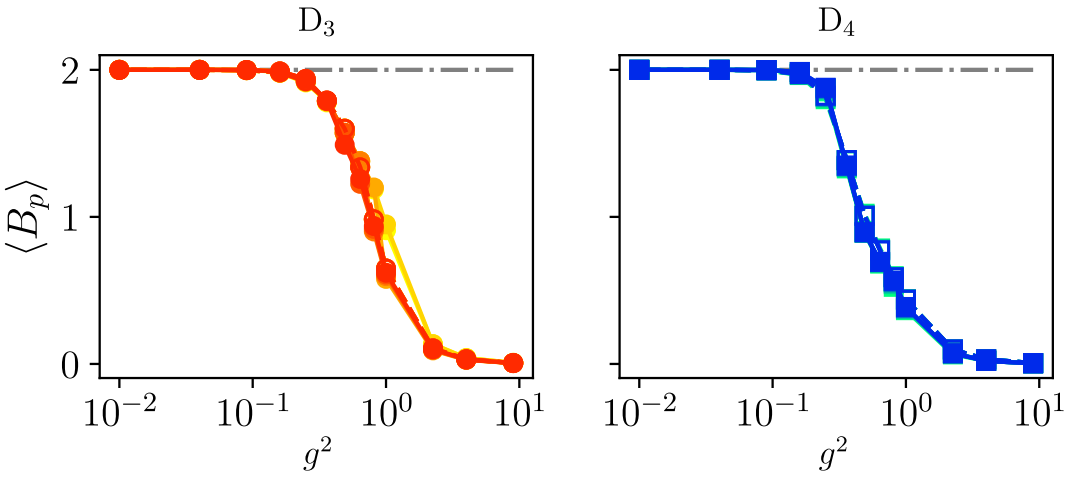

Figure 6.2: Magnetic energy density $\langle B_p \rangle$ for D_3 and D_4 lattice gauge theories on a ladder. Continuous lines with filled markers: sector with two static charges at the upper corners ($q = 1$). Dashed lines with empty markers: neutral sector ($q = 0$). Different shades correspond to different numbers of rungs $N_r \in [4, 60]$, with darker colours indicating larger systems. The value $\langle B_p \rangle \sim 2$ in the limit $g^2 \to 0$ corresponds to a flux-free configuration.

phase at weak coupling [211, 18, 212, 199, 213]. On a ladder geometry, however, boundaries induce an effective electric field that suppresses a true phase transition at finite coupling [207, 209]. Instead, the system exhibits a smooth crossover between regimes dominated by electric or magnetic energy contributions.

This crossover is characterised by the plaquette expectation value $B_p = \Re[\mathrm{Tr}(U^j_{p_1} U^j_{p_2} U^{j\dagger}_{p_3} U^{j\dagger}_{p_4})]$, which measures the magnetic energy density. As shown in Fig. 6.2, both D_3 and D_4 models display the same qualitative behaviour: $\langle B_p \rangle$ vanishes at strong coupling and approaches the character of the identity in the fundamental representation, $\chi^\tau(\mathbb{I}) = 2$, as $g^2 \to 0$. This limiting value signals the deconfined, flux-free regime — restricted to $g^2 = 0$ on the ladder geometry. The presence of static charges only weakly perturbs the magnetic sector.

6.1.4 String Tension and Gluelump Formation

The qualitative distinction between the two groups becomes clear through the string tension,

$$\sigma(g, R) = \frac{E_{q=1}(g, R) - E_{q=0}(g, R)}{R}\,, \tag{6.2}$$

where $E_{q=1}$ ($E_{q=0}$) is the ground-state energy in the presence (absence) of static charges separated by a distance R.

At strong coupling, the dominant contribution to the energy difference is linear in R, giving $\sigma \propto g^2(\alpha^\tau - \alpha^e)$. This is clearly observed in the D_4 case, where the string tension shows no appreciable R-dependence (Fig. 6.3). For D_3, by contrast, the string tension decreases linearly with increasing separation at strong coupling, rendering the total energy difference $E_{q=1} - E_{q=0} = \sigma R$ independent of R (inset of Fig. 6.3). The asymptotic value

$$E_{q=1}(g, R \to \infty) - E_{q=0}(g, R \to \infty) = 2M(g) \tag{6.3}$$

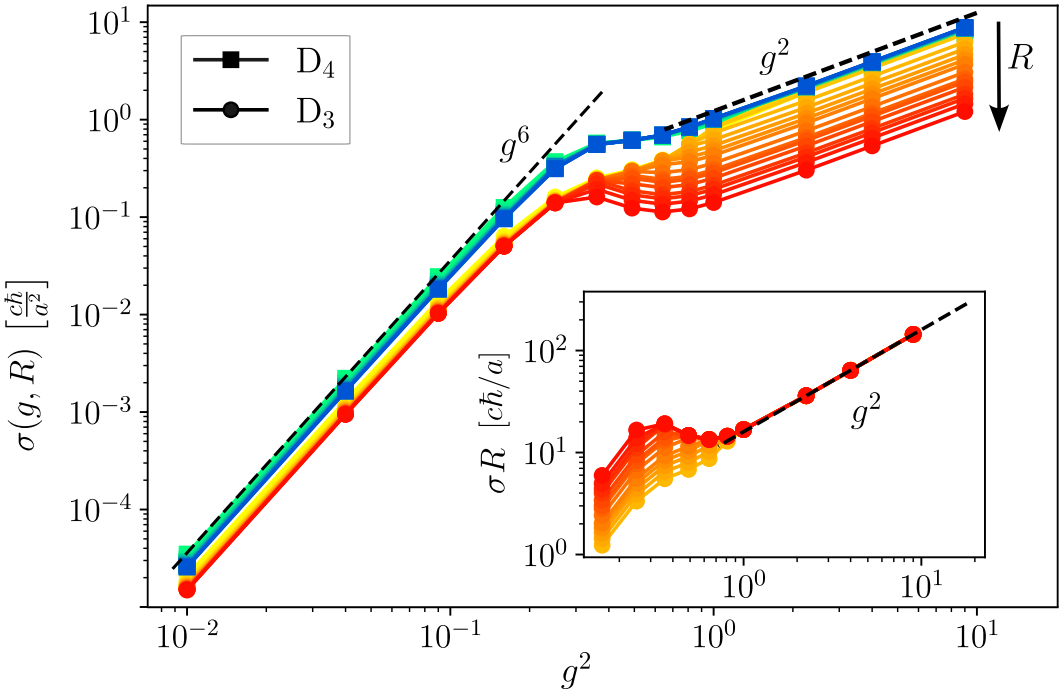

Figure 6.3: String tension $\sigma(g, R)$ as a function of the coupling g^2 for different charge separations R (colour shading) in D_3 (warm colours, circles) and D_4 (cold colours, squares). At weak coupling ($g^2 \lesssim 0.2$), both models exhibit weak confinement with $\sigma \propto g^6$. For D_4, the string tension is R-independent, consistent with a stable flux string. For D_3, the string tension decreases with increasing R at strong coupling, signalling string breaking. Inset: total energy difference $\sigma(g, R)R$, highlighting its saturation to $2M_0(g)$ in the strong-coupling regime.

is interpreted as twice the renormalized mass of a single gluelump.[2] Since $\alpha^\tau = 1$ and $\alpha^e = 0$, this gives $M(g) \simeq M_0(g) = 4g^2$ for $g^2 > 1$.

At weak coupling and short distances, both groups exhibit weak confinement due to the background electric field, similarly to $\mathbb{Z}_N$ gauge theories [207] (see App. E.4). Second-order perturbation theory predicts $\sigma(g, R) \propto g^6$, in agreement with the numerical results for both D_3 and D_4. The scaling behaviours in both coupling regimes are summarised in Table 6.1.

6.1.5 Gluelump Scaling and Continuum Limit

The R-dependent crossover between the g^2 and g^6 regimes in D_3 is attributed to the growing spatial extent of the dressed quasiparticles as the coupling is reduced. Enhanced electric-field fluctuations at smaller g^2 spread the gluelump wavefunctions along the ladder, affecting their interaction range. This is quantified through the interaction potential between two gluelumps:

$$V(g, R) = E_{q=1}(g, R) - E_{q=0}(g, R) - 2M(g)\,. \tag{6.4}$$

[2] In practice, the rapid convergence of the gluelump mass allows $M(g)$ to be extracted from simulations at a fixed large separation accessible to DMRG (here $R = 59$).

	Strong coupling	Weak coupling
D_3	$\min(8, R)\, g^2$	$R\, g^6$
D_4	$R\, g^2$	$R\, g^6$

Table 6.1: Scaling of the string energy $R\sigma(g, R)$ in the two coupling regimes for D_3 and D_4. Distances are measured in lattice units a.

DMRG data at different coupling strengths collapse onto a universal curve for the rescaled potential $V(g, Rg^\kappa)/M(g)$, as shown in Fig. 6.4. The best collapse is obtained for $\kappa = 2.5(1) \simeq 5/2$, indicating a finite screening length $R_s \sim g^{-5/2}$ at any nonzero coupling. The gluelump mass follows the scaling $M(g) \simeq \gamma g^{-\upsilon} - M_0(g)$, with $\upsilon = 5.99(5) \simeq 6$ (inset of Fig. 6.4).

These results provide insight into the scaling behaviour of the theory towards the continuum limit. Requiring that the physical interaction range r_{phys} remains finite as the lattice spacing $a \to 0$ yields

$$g \simeq \left(\zeta \frac{a}{r_{\text{phys}}}\right)^{1/\kappa}, \qquad a \ll r_{\text{phys}}\,, \tag{6.5}$$

where $\zeta = 21.2(4)$ is the crossover scale extracted from Fig. 6.4. Since $\kappa > 0$, the continuum limit is reached at small g, consistent with the growing gluelump extent as magnetic fluctuations become dominant.

A second condition requires the physical gluelump excitation energy $\mathcal{E}_{\text{phys}}$ to remain finite in the continuum limit. Neglecting the subleading strong-coupling term $M_0(g)$, this fixes the scaling of lattice energies as $M(g) \propto a^{-\upsilon/\kappa}$. Writing the dimensionless lattice mass as

$$M(g(a)) \simeq \xi \, \frac{\mathcal{E}_{\text{phys}}\, r_{\text{phys}}}{c\hbar} \left(\frac{r_{\text{phys}}}{a}\right)^{\upsilon/\kappa}, \tag{6.6}$$

and using $M(g) \simeq \gamma g^{-\upsilon}$ together with Eq. (6.5), one finds

$$\xi = \left(\frac{c\hbar}{\mathcal{E}_{\text{phys}}\, r_{\text{phys}}}\right) \frac{\gamma}{\zeta^{\upsilon/\kappa}}\,, \tag{6.7}$$

providing a consistent conversion between lattice and physical units in the continuum limit, in close analogy with running-coupling analyses [214, 215].

Remarkably, the amplitude of the interaction potential V exhibits the same emergent scaling with g — and therefore with a — as the gluelump mass $M(g)$. This implies that the interaction between the quasiparticles remains finite in the continuum limit, leading to well-defined elastic scattering amplitudes in the effective field theory.

6.1.6 Comparison with SU(N) Gauge Theories

A comparison between dihedral gauge theories and continuous non-Abelian gauge theories such as SU(N) is instructive. In SU(N), the non-trivial center forbids the screening fusion rule for the fundamental representation, but allows it for representations with zero N-ality.[3] A prominent example is the adjoint representation of SU(3), which permits the formation of color-neutral gluelumps [73, 216, 217]. While these excitations are theoretically well defined, their masses are difficult to determine unambiguously since only energy splittings are scheme independent [216, 218].

[3] N-ality characterises the representation of the center subgroup $\mathbb{Z}_N$ in a given SU(N) representation. Zero N-ality corresponds to a trivial action of the center.

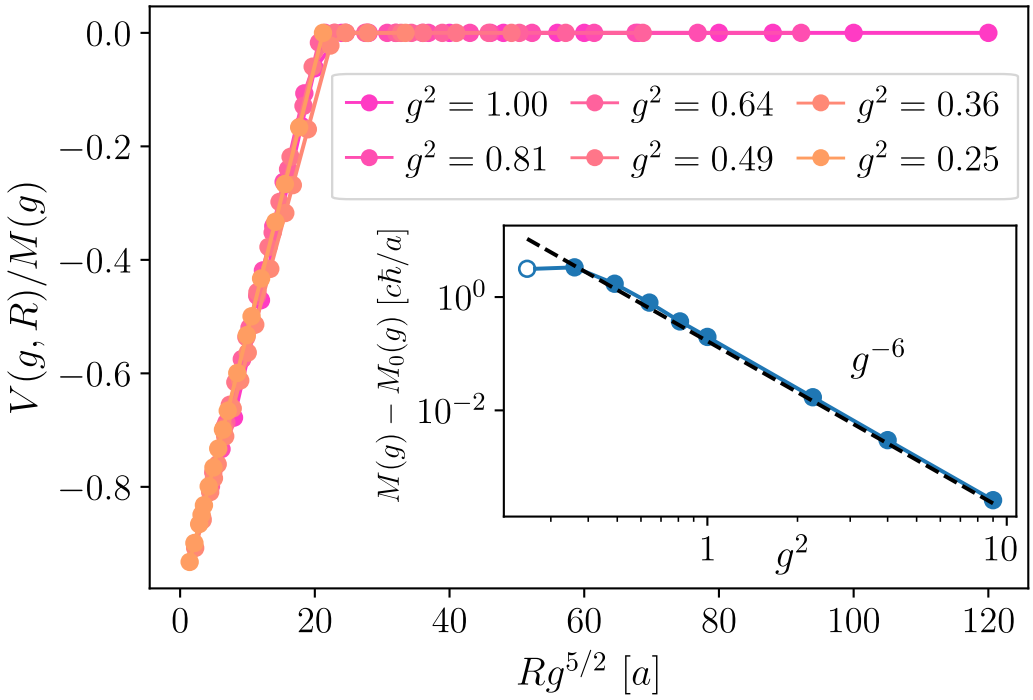

Figure 6.4: Rescaled interaction potential $V(g, Rg^{\kappa})$ with $\kappa = 5/2$. The collapse of data from different coupling strengths onto a single curve indicates a screening length scaling as $R_s \sim g^{-5/2}$. Inset: deviation of the renormalised gluelump mass $M(g)$ from the strong-coupling behaviour $M_0(g) = 4g^2$, fitted with $M(g) \simeq \gamma g^{-\upsilon}$, $\upsilon \simeq 6$.

A similar situation arises in $D_{N_{\text{even}}}$ gauge theories with $N_{\text{even}} \geq 6$, where non-faithful two-dimensional representations with zero N-ality exist. However, important differences emerge once the gauge field is coupled to dynamical matter. In QCD, adjoint charges are composite objects, typically realised as quark–antiquark pairs, and gluelumps can be created and annihilated via vacuum fluctuations. In the D_3 gauge theory with staggered fermions in the fundamental representation τ, by contrast, the charge component of a gluelump carries a conserved global U(1) quantum number, preventing the absorption of gluelumps into the vacuum and rendering them stable quasiparticles.

6.1.7 Conclusions

This chapter has investigated confinement and color screening in pure D_N lattice gauge theories on a ladder geometry, focusing on the qualitative distinction between D_3 (odd N, trivial center) and D_4 (even N, $\mathbb{Z}_2$ center).

The central finding is that the group center determines whether static fundamental charges are confined or screened. For D_4, the screening fusion rule does not hold and the string tension is R-independent at strong coupling, indicating permanent confinement via a stable electric flux string. For D_3, the trivial center permits the screening fusion rule, and the string breaks at large separations even without dynamical matter: static charges are screened by localised gluonic excitations forming gluelumps.

At weak coupling, both groups exhibit geometric confinement with string tension $\sigma \propto g^6$, a perturbative effect arising from the ladder geometry (see App. E.4). The scaling of the D_3 gluelump mass and interaction range with the coupling admits a consistent continuum limit, with the screening length $R_s \sim g^{-5/2}$ remaining finite for any nonzero coupling. The same power-law scaling governs both the gluelump mass and the interaction amplitude, ensuring well-defined physical scattering amplitudes in the continuum.

These results demonstrate that discrete non-Abelian gauge theories — even for groups of small order — constitute a controlled and tractable framework for investigating the interplay between confinement, center symmetry, and fusion rules. The close analogy with QCD gluelumps, and the qualitative differences that emerge when dynamical matter is introduced, suggest that D_N LGTs offer a valuable laboratory for understanding the non-perturbative physics of more complex gauge theories.

6.2 Quantum resources in nonabelian lattice gauge theories

This section investigates the connection between the symmetry group of a lattice gauge theory (LGT) and the quantum resources encoded in its ground state. Three complementary resource measures are studied: nonstabilizerness (stabilizer Rényi entropy), multipartite entanglement (generalised geometric measure), and fermionic non-Gaussianity (fermionic anti-flatness). The material presented here is based on [27]. The intricate structure of non-Abelian symmetries suggests that computational complexity in these systems may unfold differently from their Abelian counterparts. Yet it remains unknown whether this translates into systematically larger quantum resources. As discussed in Section 6.2.2, nonstabilizerness is a crucial ingredient for computational quantum advantage beyond what Clifford circuits can achieve [48, 219]. In classical stabilizer-tableau simulations and fault-tolerant quantum computation [37, 220], entanglement is no longer the scarce resource — magic states are. Complementing this, fermionic non-Gaussianity characterises the resource for achieving exponential hardness in matchgate-circuit simulations [221, 222].

These quantum resources have been employed to study many-body physics [223, 224, 225, 226, 227], Abelian scattering processes [228, 229], and quantum computation [47, 162, 27], providing a natural framework for probing ground-state complexity in LGTs. Reference [226] illustrates the subtle connection between nonstabilizerness and critical phenomena in an Abelian U(1) LGT, but leaves open how the group structure relates to simulation complexity.

This section addresses this question by comparing three paradigmatic quantum resources across pure-gauge theories on a ladder geometry, with Abelian $\mathbb{Z}_N$ as well as non-Abelian D_3 and SU(2) gauge symmetries. The main findings are the following:

- non-Abelian symmetries do not necessarily generate more ground-state complexity than Abelian ones;
- the required quantum resources depend non-trivially on the interplay between group structure, superselection sector, and encoding strategy;
- both Abelian and non-Abelian LGTs display resource-demanding regimes relevant to NISQ and fault-tolerant scenarios.

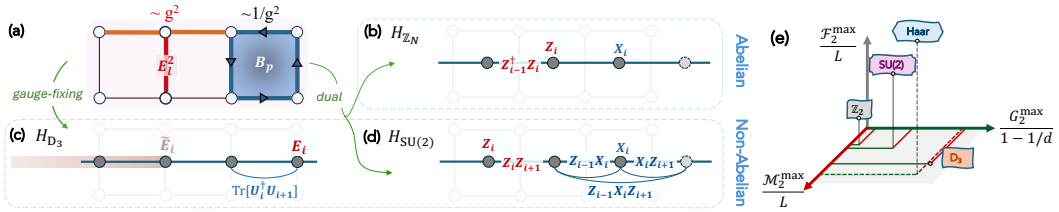

Figure 6.5: (a) Extract of the 2D LGT flux ladder and the three different mappings to one-dimensional chains, depending on the underlying gauge symmetry: (b) $\mathbb{Z}_N$, (c) D_3, (d) SU(2). For D_3, $\tilde{E}$ represents the non-local electric string after gauge-fixing, see Eq. (6.11). (e) Sketch of the maximum values of the resources across the phase diagrams: multipartite entanglement G_2, stabilizer Rényi entropy $\mathcal{M}_2$, and fermionic antiflatness $\mathcal{F}_2$ densities. All values except the maximal GGM for D_3 remain below those of Haar random states.

6.2.1 Models

To probe the connection between quantum resources and different gauge symmetry groups, pure-gauge theories in the quasi-2D geometry of a plaquette (flux) ladder are studied, as depicted in Fig. 6.5(a). This is the simplest scenario to observe the competition between electric and magnetic fields — a key ingredient for quantum simulations of LGTs [230, 231, 232, 233, 234, 86].

As introduced in Section 6.2.1, the Kogut–Susskind Hamiltonian [84] for a generic gauge group $\mathcal{G}$ reads

$$H = \frac{g_0^2}{2} \sum_{\text{links}} \mathbf{E}_l^2 - \frac{2}{g_0^2} \sum_{\text{plaquettes}} B_p \,, \tag{6.8}$$

where g_0 is the dimensionless coupling constant, $\mathbf{E}_l^2 = \sum_J \alpha^J \hat{P}_l^J$ is the weighted sum of projectors $\hat{P}_J$ onto the irreducible representations (irreps) J of the group, and α^J is the associated energy (e.g. $\alpha^J = J(J+1)$ for SU(2)). The plaquette operator is $B_p = \mathrm{Tr}(U_{p_1} U_{p_2} U_{p_3}^\dagger U_{p_4}^\dagger + \text{H.c.})$, where U_{p_i} is the parallel transporter in the faithful irrep acting on the p_i-th link of plaquette p. At each vertex $\mathbf{v}$, the local gauge transformation associated with group element $h \in \mathcal{G}$ is $\Theta_{\mathbf{v}}(h) = \prod_i \theta_i^R(h) \prod_o \theta_o^L(h)$, where i (o) labels ingoing (outgoing) links. Physical states satisfy $\Theta_{\mathbf{v}}(h) \left|\psi\right\rangle_{\text{phys}} = \left|\psi\right\rangle_{\text{phys}}$ for all $h \in \mathcal{G}$ and all vertices $\mathbf{v}$ (Gauss's law, see Section 6.2.1).

At strong coupling $g^2 \gg 1$, the system is in a confined phase and the ground state is well approximated by a product state with the identity irrep on every link. At weak coupling $g^2 \ll 1$, the ground state minimises magnetic vortices, generating complex entangled structures whose specific character depends on the symmetry group. In this regime, phenomena such as deconfinement, topological order, spin-liquid phases, and anyonic excitations can appear [235, 85, 196, 197, 199].

Three gauge groups are compared: SU(2), D_3, and $\mathbb{Z}_N$.

SU(2) with its smallest nontrivial truncation to irreps $J = 0$ and $J = 1/2$ — the hardcore-gluon approximation [81] — is adopted, yielding a spin-$\frac{1}{2}$ Hamiltonian (see App. F.1).

D_3 is the smallest discrete non-Abelian group (symmetries of an equilateral triangle), offering the prospect of exact implementation on qudit platforms [193, 192, 236]. It is associated with non-Abelian topological order and universal anyons [18, 199].

$\mathbb{Z}_N$ models are well-studied Abelian workhorses that interpolate between $\mathbb{Z}_2$ (relevant for the toric code [237]) and U(1) (the gauge symmetry of QED) [190, 238, 195].

Using local gauge symmetries, each model is mapped onto an effective one-dimensional Hamiltonian. For SU(2) and $\mathbb{Z}_N$, this follows the dual-lattice mapping of Refs. [239, 240, 209]. For D_3, gauge-fixing on the ladder legs eliminates redundant degrees of freedom, yielding a 1D chain with dynamical variables on the rungs only [93]. The mappings are sketched in Fig. 6.5(a–c) and detailed in App. F.1. The effective Hamiltonians are:

$$H_{\mathrm{SU}(2)} = g^2 \sum_i \left[-\tfrac{3}{2} Z_i Z_{i+1} + 3 Z_i\right] - \tfrac{1}{g^2} \sum_i (1-3Z_{i-1}) X_i (1-3Z_{i+1}) \,, \tag{6.9}$$

$$H_{\mathbb{Z}_N} = -\tfrac{g^2}{2} \sum_i \left[Z_{i-1}^\dagger Z_i + (1+\omega^k) Z_i + \mathrm{H.c.} \right] + \tfrac{1}{2g^2} \sum_i \left(X_i + X_i^\dagger \right) \,, \tag{6.10}$$

$$H_{\mathrm{D}_3} = -g^2 \sum_i \sum_J \alpha_J \left[\hat{P}_i^J + 2 \prod_{i'<i} \hat{P}_{i'}^J \right] - \tfrac{1}{2g^2} \sum_i \left(\mathrm{Tr}[U_i^\dagger U_{i+1}] + \mathrm{H.c.} \right) . \tag{6.11}$$

For SU(2) and $\mathbb{Z}_2$, qubit Pauli matrices are used. For $\mathbb{Z}_N$, qudits of local dimension $d = N$ with clock operators satisfying $ZX = \omega XZ$ ($\omega = e^{2\pi i/N}$) are employed, and k labels the superselection sector corresponding to background field ω^k. For D_3, qudits with $d = 6$ are required.

6.2.2 Quantum Resource Measures

Three complementary measures of quantum resources are employed.

Generalised Geometric Measure of Entanglement

The Generalised Geometric Measure (GGM) extends the geometric measure of entanglement to quantify genuine multipartite entanglement [241, 242, 243, 244, 245]. It measures the minimal distance between a quantum state $|\psi\rangle$ and the set $\mathcal{S}_2$ of 2-separable product states, and is computable for pure states as

$$G_2(|\psi\rangle) = 1 - \max_{|\pi\rangle \in \mathcal{S}_2} |\langle \pi | \psi \rangle|^2 = 1 - \max_{\mathcal{A}:\mathcal{B}} (\lambda_{\mathcal{A}:\mathcal{B}}^{\max})^2 \,, \tag{6.12}$$

where $\lambda_{\mathcal{A}:\mathcal{B}}^{\max}$ is the largest Schmidt coefficient across the bipartition $\mathcal{A}:\mathcal{B}$, and the maximisation runs over all nontrivial bipartitions. The maximum possible value is $G_2^{\max} = 1 - 1/d$, where d is the local Hilbert space dimension.[4]

Stabilizer Rényi Entropy

The SRE measures nonstabilizerness as described in Section 6.2.2. For the present setting of N qudits of local dimension d, the SRE is

$$\mathcal{M}_k(|\psi\rangle) = \frac{1}{1-k} \log \left[\sum_{P \in \mathcal{P}_N} \frac{|\langle \psi | P | \psi \rangle|^{2k}}{d^N} \right] , \tag{6.13}$$

where $\mathcal{P}_N$ is the N-qudit Pauli group defined in Section 1.1. $\mathcal{M}_k$ is non-negative, vanishes if and only if $|\psi\rangle$ is a stabilizer state [51, 52], and is experimentally accessible [53]. Throughout this chapter, $\mathcal{M}_2$ is used.

[4] Stabilizer states are restricted to $G_2(|\psi_{\mathrm{STAB}}\rangle) \in \{0, 1 - 1/d\}$ [246].

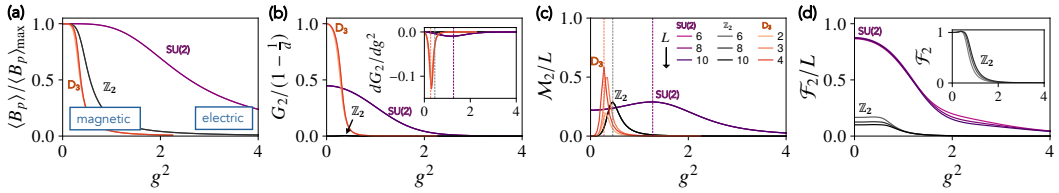

Figure 6.6: (a) Expectation value of the plaquette operator, (b) GGM, (c) SRE density, and (d) FAF density, for the ground states of SU(2), $\mathbb{Z}_2$, and D_3 pure LGTs. Results converge rapidly with increasing system size. The derivative of G_2 (inset of panel b) peaks at the same position as $\mathcal{M}_2$. SU(2) displays sizable values of all resources at small coupling. Discrete groups display mixed regimes — easy in one resource, hard in another. FAF is computed only for SU(2) and $\mathbb{Z}_2$; the inset of panel (d) shows the FAF itself (not its density) for $\mathbb{Z}_2$.

Fermionic Anti-Flatness

Matchgate circuits are classically simulable [222] but become classically intractable when combined with fermionic non-Gaussian states [221]. The fermionic anti-flatness (FAF) [227] quantifies this resource:

$$\mathcal{F}_k(|\psi\rangle) = N - \tfrac{1}{2}\mathrm{tr}[(M^T M)^k]\,, \qquad M_{mn} = -\tfrac{i}{2}\langle\psi|[\gamma_m, \gamma_n]|\psi\rangle\ , \tag{6.14}$$

where M is the covariance matrix of the $2N$ Majorana operators $\{\gamma_n\}$. The FAF vanishes if and only if $|\psi\rangle$ is a fermionic Gaussian state, is invariant under Gaussian unitaries, and is efficiently computable [225, 227] and experimentally accessible [227]. Haar random states achieve nearly maximal values of all three measures [247, 161, 227].

6.2.3 Resources and Group Structure

All models display a crossover between strong-coupling (electric ordering) and weak-coupling (magnetic ordering) regimes, identifiable via the plaquette expectation value $\langle B_p\rangle$, shown in Fig. 6.6(a). In the ladder geometry, no appreciable finite-size dependence is expected in either the order parameter or the minimum energy gap (see App. F.2), since the boundaries induce an effective longitudinal field connecting the two regimes adiabatically. The resulting fast convergence makes the relevant physics accessible at numerically tractable system sizes.[5]

The behavior of all three resource measures across the crossover is shown in Fig. 6.6.

Multipartite entanglement (GGM) vanishes in the electric limit for all three models — both SU(2) and $\mathbb{Z}_2$ approach an Ising model with a longitudinal field whose ground state is a product state, and the D_3 ground state at strong coupling is likewise a product of trivial-irrep states. In the magnetic regime, entanglement vanishes only for $\mathbb{Z}_2$, where all links become polarised by the local X field. In both non-Abelian theories, the plaquette terms correlate different sites, leading to maximal entanglement at weak coupling.

Nonstabilizerness (SRE), shown in Fig. 6.6(c), tells a different story. For the discrete gauge groups, nonvanishing SRE appears only in the crossover regime: deep in the magnetic region, the $\mathbb{Z}_N$ LGT maps to a $\mathbb{Z}_N$ paramagnet and the D_3 ground state takes the form of a generalised GHZ state — both stabilizer states. Deep in the electric region, the separable product state also has vanishing SRE. The SRE peak in the crossover regime qualitatively

[5] In extended 2D systems, the crossover becomes a continuous phase transition, at least for $\mathbb{Z}_N$ LGTs mappable to Potts models.

resembles the derivative of GGM [inset of Fig. 6.6(b)]. For SU(2), by contrast, SRE remains finite throughout the magnetic regime: at $g \to 0$, the SU(2) dual Hamiltonian is close to a cluster Ising model [248] — which has a stabilizer ground state — but further terms destroy this structure. This hints that continuous Lie groups may be harder to simulate than discrete ones.

Fermionic non-Gaussianity (FAF), shown in Fig. 6.6(d), is computed for SU(2) and $\mathbb{Z}_2$, which both admit a natural Jordan–Wigner mapping to qubit registers. In both cases, the FAF shows qualitative similarity to $\langle B_p \rangle$, plateauing in the magnetic regime and decreasing into the crossover. A crucial difference emerges: in SU(2), the FAF density converges rapidly with system size; in $\mathbb{Z}_2$, the FAF density plateau decreases as $1/N$ (i.e., the FAF itself is system-size independent). In the magnetic regime, the $\mathbb{Z}_2$ Hamiltonian maps to nearly free fermions, while the SU(2) plaquette term is not quadratic in fermionic operators, generating substantially larger non-Gaussianity.

6.2.4 Group Order and Superselection Sectors

To probe how sensitive the above findings are to details beyond the gauge symmetry group, $\mathbb{Z}_N$ LGTs with varying N and superselection sector k are examined. The parameter k labels the background charge on the ladder boundaries, inducing a uniform $\mathbb{Z}_N$ electric field $\omega^k = e^{i2\pi k/N}$.

Without background charges ($k = 0$), the ground state is fully polarised in both limits, leading to vanishing SRE and entanglement independently of N. Around the crossover, both quantities display a peak whose magnitude increases monotonically with N, as shown in Fig. 6.7(a_0, b_0).

With background charges ($k > 0$), the physics changes qualitatively. At weak coupling, the electric fluctuations screen the background field and the ground-state resources are independent of k. At strong coupling, specific combinations of N and k induce degeneracies associated with regions of extensive SRE and GGM.

For $k = N/2$, the longitudinal field vanishes [209], restoring a global $\mathbb{Z}_2$ symmetry. The strong-coupling ground state is a $\mathbb{Z}_N$ GHZ state with $G_2 = 1 - 1/N$ but vanishing SRE. This is illustrated for $\mathbb{Z}_2$ ($k = 1$) and $\mathbb{Z}_4$ ($k = 2$) in Fig. 6.7(a1–b2).

Beyond this special case, the electric Hamiltonian has degenerate pairs of eigenvalues at angles $2\pi a/N$ and $2\pi b/N$ with $a + b = N - k$. When such a pair minimises the electric energy, the strong-coupling ground state takes the form $|\psi\rangle = (|a\rangle^{\otimes N} + |b\rangle^{\otimes N})/\sqrt{2}$, which (except for $N = 2$) is not a GHZ state and carries both finite entanglement ($G_2 = 1/2$) and finite nonstabilizerness ($\mathcal{M}_2 = 0.32$, computed analytically in App. F.3). Figure 6.7(a1–b2) clearly illustrates this: at $k = 1$, all $\mathbb{Z}_N$ models except $\mathbb{Z}_2$ display finite nonstabilizerness and entanglement at strong coupling, consistent with the analytical prediction.

6.2.5 Remarks on Extended 2D LGTs

The analysis focused on two-leg ladders, as extended 2D systems pose a significant numerical challenge. Several statements can nonetheless be made in limiting cases.

At $g \to \infty$, independently of $\mathcal{G}$ and without background charges, the ground state is a product of polarised links with vanishing resources.

At $g \to 0$, the ground state for discrete groups is related to the associated surface code [237]. For Abelian $\mathbb{Z}_N$, the set of stabilizers (plaquette and vertex operators) uniquely defines each eigenstate. Analogously to the $\mathbb{Z}_2$ case, these states have large entanglement

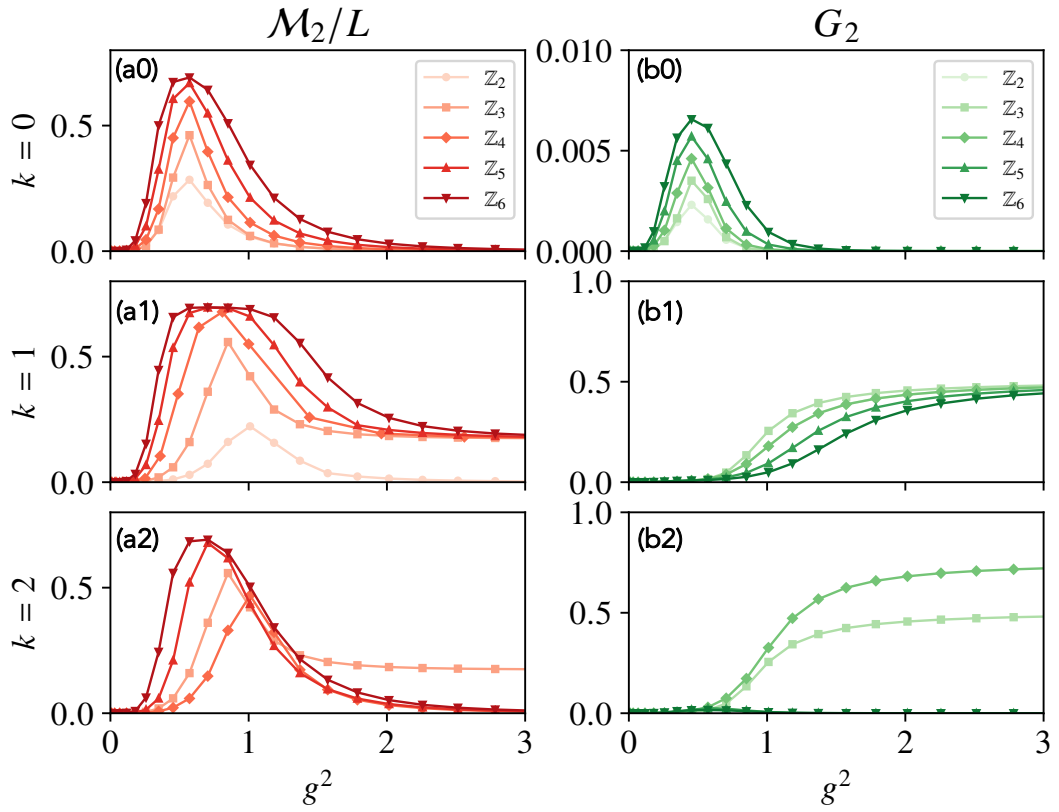

Figure 6.7: Impact of superselection sector k and group order N on quantum resources, for a $\mathbb{Z}_N$ LGT with $N = 2, 3, 4, 5, 6$ and $k = 0, 1, 2$, at fixed $N = 4$. (a) SRE peaks whenever electric and magnetic terms compete. (b) GGM has a significant crossover peak only at $k = 0$. Both quantities may reach a strong-coupling plateau whose height depends critically on the combination of k and N.

in the full Hilbert space but vanishing SRE, since nonlocal transformations between representations map Pauli strings to Pauli strings. Non-Abelian topological order breaks this picture [199], potentially leading to larger quantum resources independently of the chosen representation. On the ladder, this is already hinted at by SU(2), which maintains large $\mathcal{M}_2$ at weak coupling even after the dual mapping. Its complexity is expected to increase in a full 2D geometry.

6.2.6 Conclusions

This chapter has investigated the connection between gauge symmetry group and ground-state complexity through three complementary quantum resource measures — nonstabilizerness, multipartite entanglement, and fermionic non-Gaussianity — in pure-gauge LGTs on a ladder geometry.

An intricate picture emerges. Quantum resources depend simultaneously on the gauge group, the background charge (superselection sector), and the strategy used to project the Hamiltonian onto independent degrees of freedom.

Perhaps counterintuitively, non-Abelian symmetries are not necessarily associated with larger nonstabilizerness. In the discrete groups $\mathbb{Z}_N$ and D_3, nonstabilizerness is confined to the crossover regime and vanishes in both limiting phases. For SU(2), by contrast, nonstabilizerness remains finite throughout the magnetic phase, suggesting that continuous Lie groups may be inherently harder to simulate than discrete ones in the fault-tolerant sense. Non-Abelian groups do, however, consistently induce larger multipartite entanglement than Abelian ones in the magnetic phase.

In $\mathbb{Z}_N$ LGTs, the combination of background charge k and group order N plays a decisive role in determining ground-state complexity. Specific combinations of k and N can restore degeneracies at strong coupling, generating a finite plateau of both SRE and GGM that would not be present in the zero-charge sector. The analytical expression derived for the SRE of the resulting two-component superposition state confirms this picture quantitatively.

From a practical perspective, the results indicate that quantum simulations of LGTs will face resource-demanding regimes for both NISQ and fault-tolerant implementations, associated with large entanglement (CNOT gate count) and nonstabilizerness (T-gate count), respectively. Future experiments and simulation protocols will need to carefully account for the chosen superselection sector, group order, and coupling regime.

Several open questions follow naturally from this work, including the role of dynamical matter, the truncation of continuous groups, and the effect of spatial dimensionality. The comparison between qubit and qudit encodings also deserves further attention: the results show that both SRE and GGM increase with group order in $\mathbb{Z}_N$ LGTs, suggesting that qudit encodings may be advantageous — or disadvantageous — depending on the specific resource bottleneck of interest.

Conclusions

This thesis has explored the role of qudits in two complementary research programmes: quantum optimization and the quantum simulation of lattice gauge theories.

In the domain of quantum optimization, the central contributions concern the use of qudit ancillae to handle constrained problems and the improvement of variational algorithms through counterdiabatic driving. The direct penalty approach developed in Chapter 2 demonstrates that a single ancilla qudit of dimension $N + 1$ suffices to enforce Hamming-weight constraints in QAOA without expanding the solution space — a concrete advantage over binary slack-variable encodings, which inflate the infeasible sector with each additional constraint. Chapter 3 showed that symmetry-enhanced counterdiabatic driving reduces the number of variational parameters to a symmetry-constrained subspace, improving both convergence speed and solution quality on qudit instances. These results suggest that incorporating the algebraic structure of both the problem and the hardware into the ansatz design is a productive direction for NISQ-era algorithms.

Chapter 4 addressed two further aspects of variational algorithm performance. The experimental confirmation that amplitude-damping noise does not induce barren plateaus is an encouraging result for near-term hardware, as it implies that the dominant noise model of superconducting processors does not suppress trainability. The analysis of the Sherrington–Kirkpatrick spin-glass model revealed a universal transient structure in QAOA runs — a build-up and partial relaxation of nonstabilizerness — that tracks the algorithm's progress toward the solution and is present in both qubit and qutrit versions of the protocol. While nonstabilizerness is not the primary object of study in this work, its behaviour provides a useful diagnostic for understanding what happens inside a variational circuit at a structural level.

On the lattice gauge theory side, the analysis of D_N models illustrates how discrete non-Abelian symmetries give rise to qualitatively different physical phenomena depending on the group structure. The connection between the group center and confinement established in Chapter 6 — confinement for even N, gluon-induced screening for odd N — provides a clear organizing principle for the phenomenology of dihedral LGTs and contributes to the broader effort of understanding non-Abelian confinement in tractable low-dimensional models. The D_3 ground-state analysis in Chapter 5, while preliminary, identifies parameter regimes with distinct phase structure and lays the groundwork for studying string-breaking dynamics in a genuinely non-Abelian setting that is accessible to both current tensor-network methods and near-term quantum hardware.

The comparison of ground-state complexity across $\mathbb{Z}_N$, D_3, and SU(2) gauge theories in Chapter 6 yields a counterintuitive finding: non-Abelian symmetry does not automatically make ground states harder to simulate. The simulation cost depends instead on the interplay between group structure, superselection sector, and the choice of encoding — a result with practical implications for the design of quantum simulation experiments and the assessment

of their classical hardness.

Taken together, these results point to qudits as a useful extension of the qubit paradigm for specific tasks — constraint handling, group encoding, and symmetry exploitation — rather than a universal replacement. Several directions remain open. The D_3 results of Chapter 5 call for a systematic finite-size scaling analysis and an extension to dynamical observables such as real-time string-breaking rates, which would connect directly to ongoing trapped-ion and Rydberg experiments. The universality of the magic barrier in QAOA invites analytical treatment, in particular whether a closed-form expression for its peak and scaling can be derived in the large-N limit. More broadly, the question of how group structure shapes the classical simulability of LGT ground states — beyond the specific resource measures studied here — remains an open and physically important problem. These questions place qudit-based quantum information at the intersection of quantum algorithms, quantum simulation, and many-body physics, a territory that current and near-future quantum hardware is well positioned to explore.

From a broader perspective, the results obtained in this thesis can be situated within the current state of the art in both quantum optimization and quantum simulation. In quantum optimization, significant effort has recently focused on improving the trainability and expressivity of variational algorithms under realistic hardware constraints. Within this context, the use of qudit ancillae for constraint enforcement provides an alternative to standard qubit-based encodings that aligns with ongoing attempts to reduce circuit overhead and mitigate barren plateau phenomena. Similarly, the symmetry-informed reduction of variational parameters connects with a growing body of work emphasizing problem-inspired ansätze and hardware-efficient implementations. On the simulation side, the study of lattice gauge theories with finite groups contributes to a well-established programme aimed at identifying minimal models that capture essential features of non-Abelian physics while remaining accessible to both tensor-network and quantum hardware approaches. In particular, the identification of the role of the group center in determining confinement properties complements existing analytical and numerical studies, offering a unifying viewpoint that bridges discrete and continuous gauge theories.

Looking ahead, several promising research directions emerge from this work. First, the integration of qudit-based encodings with error-mitigation and error-correction strategies represents a natural next step, especially in light of recent progress in noise-resilient quantum algorithms. Understanding whether the structural advantages of qudits persist in the presence of realistic error models will be crucial for their practical adoption. Second, the extension of symmetry-enhanced variational principles to dynamical settings — for instance, in variational quantum simulation of real-time evolution — could provide new tools for probing out-of-equilibrium phenomena in many-body systems. On the lattice gauge theory side, a systematic exploration of larger non-Abelian discrete groups and higher-dimensional lattices would clarify the extent to which the observed interplay between group structure and genuine non-perturbative phenomena generalizes. More ambitiously, combining qudit-based hardware platforms with analog-digital hybrid approaches may enable the study of regimes that are currently beyond typical quantum strategies (e.g. truncations of Hilbert spaces), thereby strengthening the role of qudit quantum simulators in the context of non abelian LGTs.

Appendix A

Appendix CD

This appendix collects supplementary material for Chapter 3, which develops the symmetry-enhanced counterdiabatic approach for qudit QAOA. Section A.1 provides the formal proof that the counterdiabatic terms inherit the symmetries of the problem Hamiltonian, together with an explicit example of how those symmetries reduce the number of independent variational parameters. Section A.2 gives the explicit qudit encodings of the MAX-k-CUT Hamiltonian for $k = 2, 3, 4$, along with the corresponding first-order counterdiabatic operator pools and numerical results for six-vertex MAX-3-CUT instances. Section A.3 presents a complementary benchmark in which the variational ansätze are trained by directly maximising the fidelity with the target W state, rather than minimising the energy, confirming the superiority of counterdiabatic circuits over standard qudit-QAOA at low circuit depth

A.1 Theory of Symmetry-Grouped Counterdiabatic Terms

Consider a problem Hamiltonian H_P that is invariant under a set of symmetries $S = \{\pi_k\}$ drawn from the permutation group:

$$H_P = \pi_k H_P \pi_k^{-1} \quad \forall\, k\,. \tag{A.1}$$

It follows directly that the adiabatic Hamiltonian $H_a(t)$ inherits these symmetries,

$$H_a(t) = \pi_k H_a(t) \pi_k^{-1}\,. \tag{A.2}$$

The CD term at first order is $Q_1(t) = [H_a(t), \partial_t H_a(t)]$. The symmetries of H_P are inherited by the CD Hamiltonian exactly, as can be seen from

$$\begin{aligned} \pi_k Q_1(t) \pi_k^{-1} &= \pi_k [H_a(t), \partial_t H_a(t)] \pi_k^{-1} \\ &= [\pi_k H_a(t) \pi_k^{-1},\, \pi_k \partial_t H_a(t) \pi_k^{-1}] \\ &= [\pi_k H_a(t) \pi_k^{-1},\, \partial_t (\pi_k H_a(t) \pi_k^{-1})] \\ &= [H_a(t), \partial_t H_a(t)] = Q_1(t)\,. \end{aligned} \tag{A.3}$$

The only effect of applying a permutation symmetry is therefore a relabelling of the CD terms. Parameters whose labels form a closed set under the action of all symmetries in S can be grouped, i.e. constrained to share a single value, reducing the total number of independent variational parameters.

Example. Consider a CD Hamiltonian with seven distinct terms and a single symmetry π. The first-order CD operator reads

$$Q_1 = \alpha_1 B_1 + \alpha_2 B_2 + \alpha_3 B_3 + \alpha_4 B_4 + \alpha_5 B_5 + \alpha_6 B_6 + \alpha_7 B_7 \,. \tag{A.4}$$

Applying the symmetry and using $Q_1 = \pi Q_1 \pi^{-1}$ gives

$$\pi Q_1 \pi^{-1} = \alpha_3 B_1 + \alpha_2 B_2 + \alpha_1 B_3 + \alpha_7 B_4 + \alpha_6 B_5 + \alpha_5 B_6 + \alpha_4 B_7 \,. \tag{A.5}$$

Matching terms between the two expressions yields the identifications $\alpha_1 = \alpha_3$, $\alpha_4 = \alpha_7$, and $\alpha_5 = \alpha_6$, reducing the number of independent parameters from seven to four.

A.2 Formulation of Max-k-Cut on Qudits and Results for $k = 3$

The Max-k-Cut problem seeks to partition the vertices of a graph into k colour classes so as to maximise the number of edges connecting vertices of different colours. The quantum encoding is obtained by expressing the per-edge contribution to the cost Hamiltonian in terms of angular-momentum operators for qudits of local dimension $d = k$.

For Max-2-Cut ($d = 2$ qudits), the per-edge term is

$$H_{i,j} = 4 L_{z,i} L_{z,j} \,, \tag{A.6}$$

and the first-order CD pool is $Q_1 = \{L_{y,i} L_{z,j}\}$.

For Max-3-Cut ($d = 3$ qutrits):

$$H_{i,j} = L_{z,i} L_{z,j} - 2(L_{z,i}^2 + L_{z,j}^2) + 3 L_{z,i}^2 L_{z,j}^2 \,, \tag{A.7}$$

$$Q_1 = \left\{ L_{y,i} L_{z,j},\ L_{z,i} L_{y,j} + L_{y,i} L_{z,j},\ L_{z,i} L_{y,j} + L_{y,j} L_{z,i},\ L_z^2 \right\} . \tag{A.8}$$

For Max-4-Cut ($d = 4$ ququarts):

$$H_{i,j} = \frac{365}{72} L_{z,i} L_{z,j} - \frac{5}{8}(L_{z,i}^2 + L_{z,j}^2) + \frac{1}{2} L_{z,i}^2 L_{z,j}^2 - \frac{41}{18}(L_{z,i}^3 L_{z,j} + L_{z,i} L_{z,j}^3) + \frac{10}{9} L_{z,i}^3 L_{z,j}^3 \,. \tag{A.9}$$

All terms in these problem Hamiltonians contain an even number of L_z operators per qudit, with the maximum power of L_z on a single site equal to $k - 1$. The Hamiltonian coefficients for general k are obtained by solving a linear system of $\sum_{d=1}^{k} \lceil d/2 \rceil - 1$ equations, which scales as $\mathcal{O}(k^2)$.

Figure 3.6 presents the numerical results for the Max-3-Cut problem on six instances defined by non-planar graphs with six vertices. Results are shown in terms of the cost function value (expected energy) at each optimisation iteration, which is the quantity minimised in the classical loop, as well as the probability of sampling an optimal solution. In all instances, the success probability is at least one order of magnitude above random guessing. This behaviour is consistent with the use of the energy, rather than the fidelity, as the cost function: energy minimisation does not directly penalise weight on non-optimal states of equal or near-equal energy, yet the CD-enhanced ansätze still achieve a substantially elevated probability of finding the correct solution.

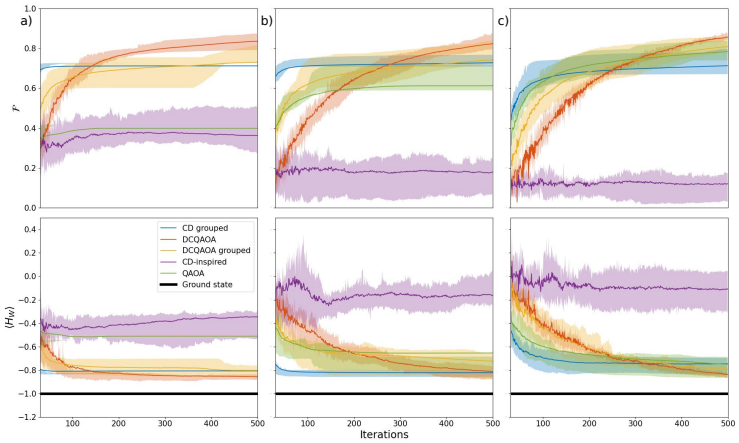

Figure A.1: Fidelity with the W state as a function of optimiser iteration, obtained by directly maximising $\mathcal{F}$ for $N = 3$ qutrits. CD-based ansätze outperform standard qudit-QAOA for low layer counts ($p \leq 4$), achieving a best-case fidelity of $\mathcal{F} \sim 85\%$.

A.3 Dicke States and Fidelity Maximisation for W State Preparation

Dicke states are important metrological states with connections to quantum communication and multipartite entanglement [249]. For a system of N qudits of local dimension d, the Dicke state with excitation number κ is defined as

$$D_\kappa^N = \binom{N}{\kappa}^{-1/2} \sum_i P_i \Big(|d-1\rangle^{\otimes N-\kappa} |0\rangle^{\otimes \kappa} \Big) , \tag{A.10}$$

where P_i runs over all distinct permutations of the local basis labels. This state is the ground state of the Hamiltonian

$$H_\kappa^N = \left(\frac{N}{2} - \kappa \right) \sum_{i=1}^{N} L_{zi} + \frac{1}{2} \sum_{j<i} L_{zi} L_{zj} . \tag{A.11}$$

For $\kappa = 1$ and qutrits ($d = 3$), this reduces to the Hamiltonian of Eq. (3.15), and the corresponding Dicke state is

$$D_1^N = \frac{1}{\sqrt{N}} \sum_i P_i \Big(|2\rangle^{\otimes N-1} |0\rangle \Big) , \tag{A.12}$$

which is precisely the W state of Eq. (3.14).

In Section 3.3, the variational algorithms are trained by minimising the expectation value of the energy of Eq. (3.15). Here, a complementary approach is tested by instead directly maximising the fidelity of Eq. (3.16) with the target W state, using the same variational ansätze. The cost function takes the form $\mathcal{F} = \mathrm{Tr}[U(\boldsymbol{\theta}) |\psi_0\rangle\langle\psi_0| U(\boldsymbol{\theta})^\dagger |W\rangle\langle W|]$. Since the W state is a ground state of Eq. (3.15), maximising the fidelity is consistent with minimising the energy, and both objectives guide the optimisation toward the same target subspace.

As shown in Fig. A.1, direct fidelity maximisation yields better approximation to the W state than energy minimisation in the main text, reaching $\mathcal{F} \sim 85\%$ in the best case. The results further confirm that CD-based ansätze consistently outperform standard qudit-QAOA for a small number of layers ($p \leq 4$), in line with the findings reported in Section 3.3.

Appendix B

Experimental Demonstration of the Absence of NIBP

This appendix provides supplementary material for Sec. 4.1, which experimentally demonstrates the absence of noise-induced barren plateaus under amplitude-damping noise on IBM superconducting processors. Sec. B.1 describes the two circuit compilation strategies used (ladder and parallel), details the resulting two-qubit gate counts, circuit depths, and runtimes as a function of system size and number of layers, and discusses the effective coherence time extracted from the hardware data. Further details on the Information Content Landscape Analysis (ICLA) procedure and the statistical analysis of the gradient norm estimates are also reported.

B.1 Details on Hardware Experiments

B.1.1 Circuit Structure

Figure B.1 shows examples of the two circuit types used on IBM hardware, differing only in gate scheduling. In the ladder implementation (left), two-qubit terms are executed sequentially, resulting in larger circuit depth. In the parallel ('short') implementation (right), commuting two-qubit terms are executed simultaneously.

After transpilation, both circuit types have exactly the same number of two-qubit gates, $2(N-1)L$, where the factor of 2 accounts for the realisation of each R_{zz} gate with two `ECR` gates on the Falcon platform. The number of single-qubit gates scales approximately as $13.5\,NL$ for both implementations.

B.1.2 Circuit Depth and Runtimes

The circuit depth follows the linear relation

$$d(N, L) = aN + bL + c\,, \tag{B.1}$$

with $(a, b, c) = (11.1,\ 24.0,\ -29.6)$ for the ladder architecture, $(0,\ 24.0,\ -6.0)$ for the short circuits on Falcon devices, and $(7.3,\ 18.5,\ 1.0)$ for `ibm_fez`.

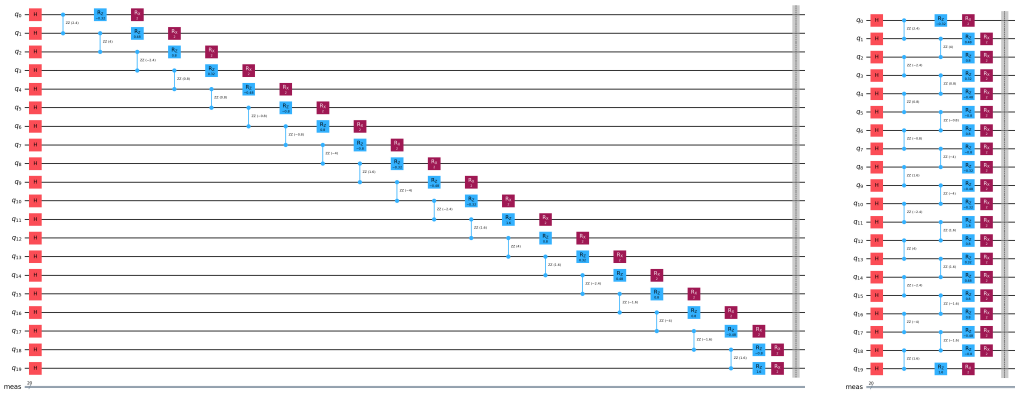

Figure B.1: Examples of the same logical $L = 1$ single-layer circuit with $N = 20$ qubits transpiled in two different ways. Left: ladder structure (default compilation used throughout the experiments). Right: parallel execution ('short' compilation), where commuting two-qubit terms are grouped and executed simultaneously to minimise circuit depth.

The net circuit runtime (excluding the $250\,\mu s$ reset time between shots) as a function of N and L is:

$$t_{cir}^{\texttt{Falcon}}(N, L) = 1.8\,\mu s\,N + 3.3\,\mu s\,L + 45\,\mu s\,, \tag{B.2}$$

$$t_{cir}^{\texttt{Falcon,short}}(N, L) = 2.8\,\mu s\,L + 61\,\mu s\,, \tag{B.3}$$

$$t_{cir}^{\texttt{fez}}(N, L) = 0.24\,\mu s\,N + 0.61\,\mu s\,L + 130\,\mu s\,. \tag{B.4}$$

The parallel circuit runtime is independent of N, as expected. These relations are used to compute all circuit runtimes reported in the main chapter.

B.2 Robustness of ICLA

To reduce the number of required hardware evaluations, the cost-function landscape is capped at $M = \min(10m, 200)$ sample points rather than the full $M = \mathcal{O}(m)$ prescription of Ref. [137]. Figure B.4 shows ICLA results for various values of M for $N = 8$ and $N = 102$ qubit circuits. All results are within their respective error bars for $M \geq 100$, confirming that the truncation does not affect the conclusions. The only exception is $M = 50$ at $N = 102$, where larger variations appear. For deep circuits (large L), this truncation reduces the number of required circuit evaluations from $M = 20L = 1200$ to 200 for $L = 60$, representing a substantial saving.

Device	$\langle T_1 \rangle\,[\mu s]$	$\langle T_2 \rangle\,[\mu s]$	$\langle$`readout_error`$\rangle\,[\%]$	$\langle$`ECR_error`$\rangle\,[\%]$
`ibm_brisbane`	236 ± 67	156 ± 93	3.0 ± 3.0	1.2 ± 1.4
`ibm_kyiv`	302 ± 91	155 ± 113	2.2 ± 2.7	1.4 ± 1.4
`ibm_sherbrooke`	310 ± 90	181 ± 136	3.3 ± 3.4	1.2 ± 1.4
`ibm_fez`	146 ± 46	88 ± 54	2.0 ± 3.2	0.7 ± 0.9

Table B.1: Summary of calibration statistics across all hardware platforms used. Errors are standard deviations. Note that the distributions are highly non-Gaussian; the stated values indicate the order of magnitude of variation only.

B.3 Experimental Details on Cost Functions and Noise

B.3.1 Estimating Cost Function Values

Figure B.5 shows the sorted cost values obtained from Eq. (**??**) for $N = 20$ and $N = 102$ qubits at several circuit depths. The distributions display a characteristic S-shape, with most cost values near zero and some outliers reaching the scale of the Hamiltonian energies.

The typical width of the distributions is set by $C_0 = \sqrt{\langle H_C^2 \rangle} = \frac{\sqrt{N}}{2}\sqrt{\frac{N-1}{N}\overline{J^2} + \overline{h^2}} \sim \sqrt{N}$, where $\overline{J^2}$ and $\overline{h^2}$ are the mean squared Hamiltonian parameters. Despite this width, the circuit parameters have a pronounced effect by shifting the centre of the distribution — confirming that the gradient signal reflects genuine variational structure rather than random noise.

For $N = 20$, the absolute cost values diminish with increasing L, consistent with NIBP onset. For $N = 102$, the cost values retain their S-shape structure across all depths, consistent with the observed gradient saturation.

B.3.2 Shot-Noise Floor

The shot-noise floor is estimated by replacing the actual cost-function landscape with random samples from the Hamiltonian spectrum and computing the ICLA from this surrogate landscape. As shown in Fig. B.6, the shot-noise floor is substantially below the observed gradient signal in all cases, confirming that the finite gradient is not an artefact of measurement noise.

B.3.3 Reproducibility Across Calibrations

To assess the influence of calibration drift, the same 600-parameter landscape is measured twice on `ibm_brisbane` for $N = 102$, $L = 30$, with four days between the two runs. The same physical qubits and logical-to-physical mapping are used throughout; calibration data for both runs are reported in Table B.2.

As shown in Fig. B.7, the cost function values from the two experiments are highly correlated, confirming that the observed gradient signal is reproducible and not an artefact of random calibration fluctuations. The inter-experiment variation exceeds the shot-noise error bars but is clearly structured, consistent with a deterministic influence of the variational parameters.

	$\langle T_1 \rangle$	$\langle T_2 \rangle$	$\langle$`readout_error`$\rangle$	$\langle$`ECR_error`$\rangle$	Num. circs
exp_1, calib_1	$237\,\mu s$	$157\,\mu s$	2.88%	3.18%	502
exp_1, calib_2	$237\,\mu s$	$157\,\mu s$	3.02%	3.26%	81
exp_1, calib_3	$237\,\mu s$	$157\,\mu s$	3.02%	3.18%	17
exp_2, calib_1	$236\,\mu s$	$157\,\mu s$	3.08%	3.06%	600

Table B.2: Calibration statistics for the two reproducibility experiments on `ibm_brisbane` ($N = 102$, $L = 30$). All values are computed from the actual physical qubits used, which are the same for all circuits.

B.4 Density Matrix Analysis

Classical density-matrix simulations are performed for $N = 8$ qubits using Qiskit, considering (i) amplitude-damping noise emulating T_1 relaxation and (ii) single-qubit depolarising noise, applied after transpilation to the native IBM gate set.

Under depolarising noise, the eigenvalue spectrum of the output density matrix converges to a uniform distribution concentrated at $1/2^N$, consistent with convergence to the maximally mixed state and the emergence of NIBP.

Under amplitude damping, the spectrum remains structured with a small number of dominant eigenvalues, even at large circuit depth. This is consistent with the noise-induced limit set picture of Ref. [134]: the output state contracts to a specific sub-manifold determined by the noise channel, which retains dependence on the variational parameters and thereby preserves finite gradient information.

The essential features of this spectral structure can be understood through a simple analytical model. For an initial state $\rho_0 = |1\rangle\langle 1|^{\otimes N}$ decayed under amplitude damping with per-qubit probability p, the output eigenvalues and their multiplicities are:

$$\rho_n = p^n (1-p)^{N-n}\,, \tag{B.5}$$

$$S(\rho_n) = \binom{N}{n}, \quad n = 0, 1, \ldots, N\,. \tag{B.6}$$

This reproduces the qualitative form of the amplitude-damping spectrum in Fig. B.8 (top right panel), providing an intuitive explanation for why non-unital noise preserves spectral structure even at large L.

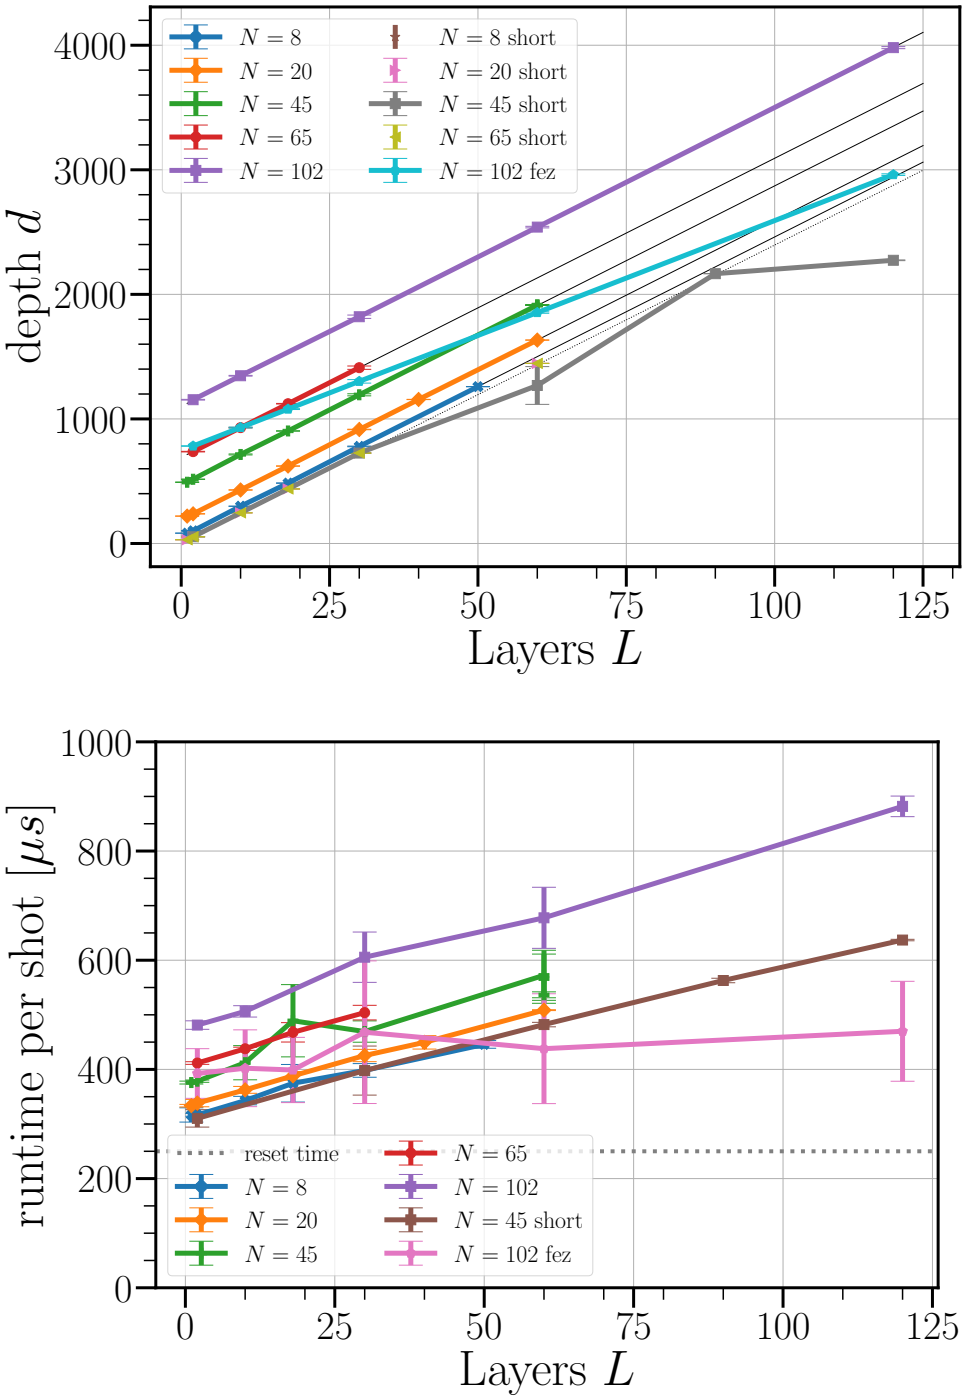

Figure B.2: Circuit depth (top) and circuit runtimes (bottom panels) as a function of layers L for different qubit numbers N, from `ibm_brisbane` unless otherwise indicated in the legend.

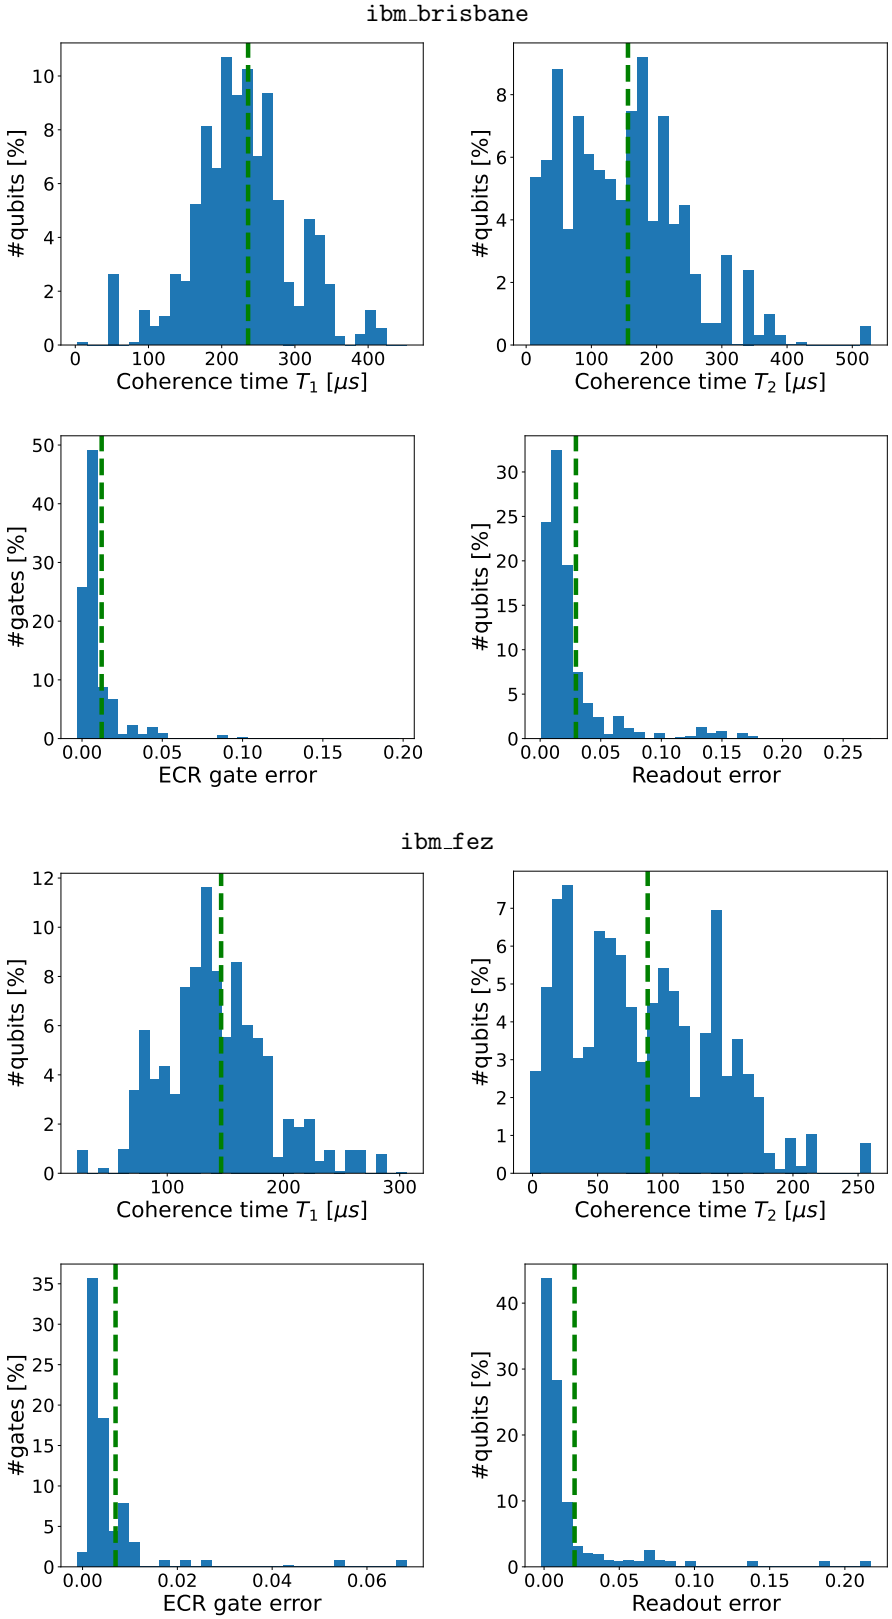

Figure B.3: Distributions of hardware properties for `ibm_brisbane` (top) and `ibm_fez` (bottom) for the $N = 102$ qubit circuits. From left to right, top to bottom: T_1, T_2, ECR-gate error, and readout error. Green vertical lines indicate mean values (reported in Table B.1). All distributions are highly asymmetric and far from Gaussian.

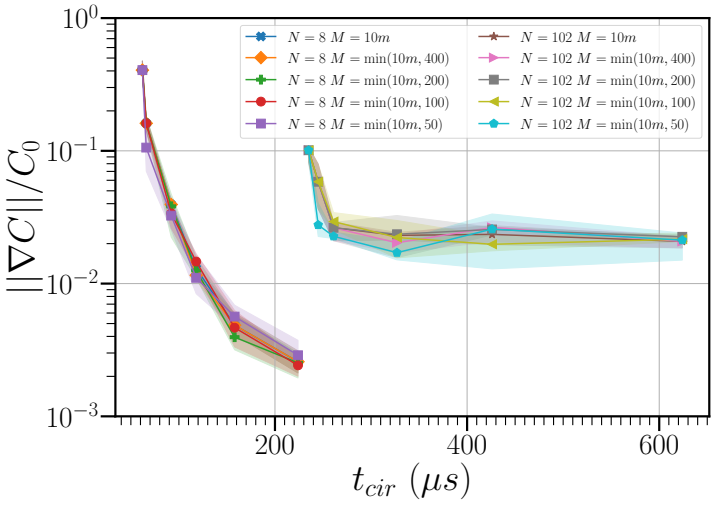

Figure B.4: Dependence of ICLA results on the number of sample points M used for the cost-function landscape, for $N = 8$ and $N = 102$ qubit circuits on `ibm_brisbane`. Results are consistent across all values of M tested, with larger variations only for $M = 50$ at $N = 102$.

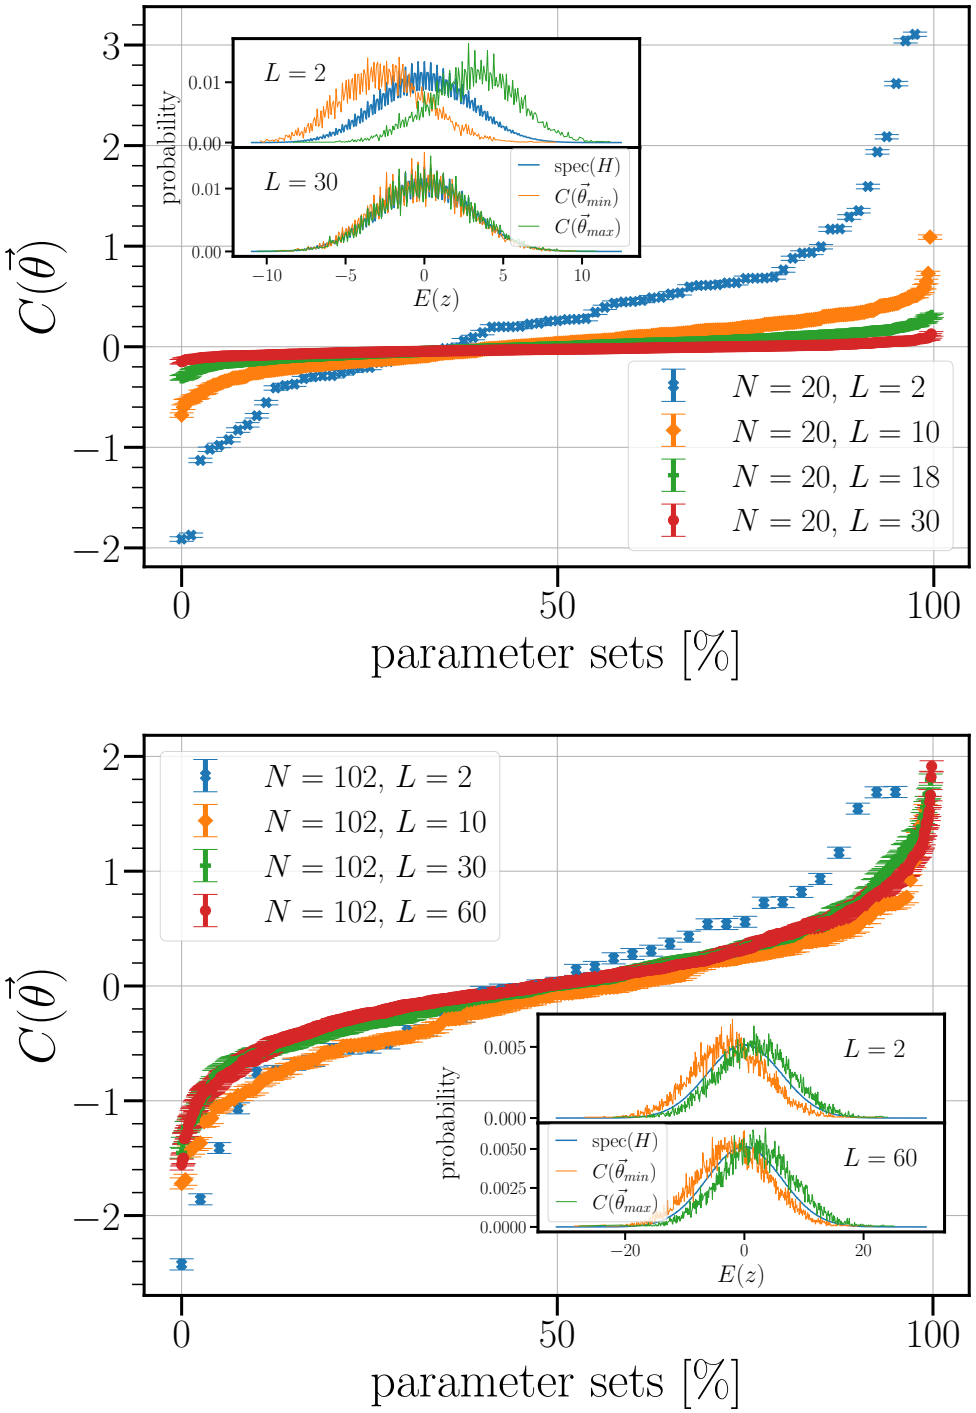

Figure B.5: Sorted cost values for $N = 20$ (top) and $N = 102$ (bottom) from `ibm_brisbane` experiments. Error bars indicate shot-noise variation for $R = 16384$ shots. Insets: cost spectra for parameters producing minimal (orange) and maximal (green) cost values at small and large layer numbers. The blue curve shows the Hamiltonian eigenvalue spectrum.

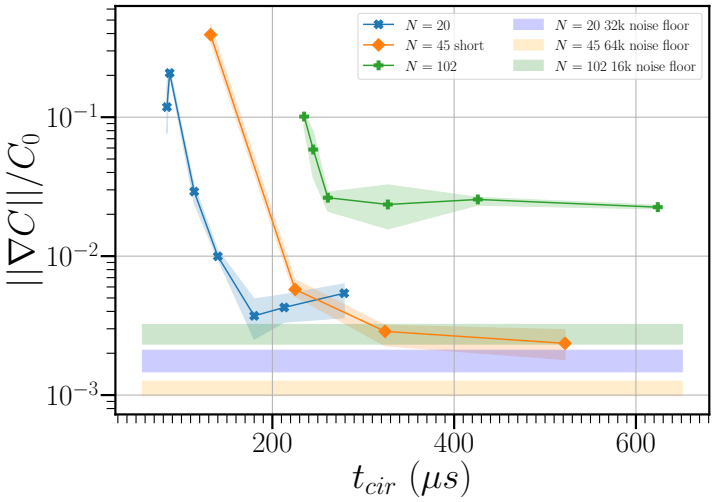

Figure B.6: Gradient signal and corresponding shot-noise floor for several experimental configurations. The shot-noise floor (dashed lines) is computed by replacing the actual cost values with random samples from the Hamiltonian spectrum. The observed gradient signal is always substantially above the shot-noise floor.

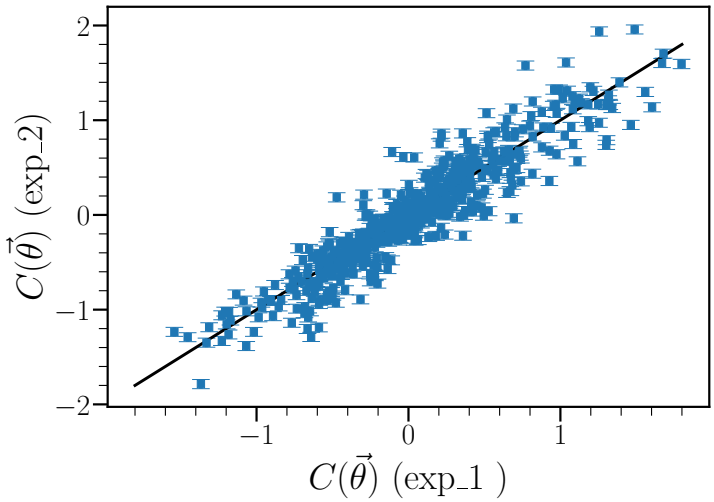

Figure B.7: Mean cost function values for $N = 102$ and $L = 30$ measured twice on `ibm_brisbane` with four days between experiments (experiment 1 and experiment 2, each with 600 parameter vectors and $R = 16384$ shots). The strong correlation between the two runs confirms that the observed signal is reproducible and reflects the deterministic influence of the variational parameters.

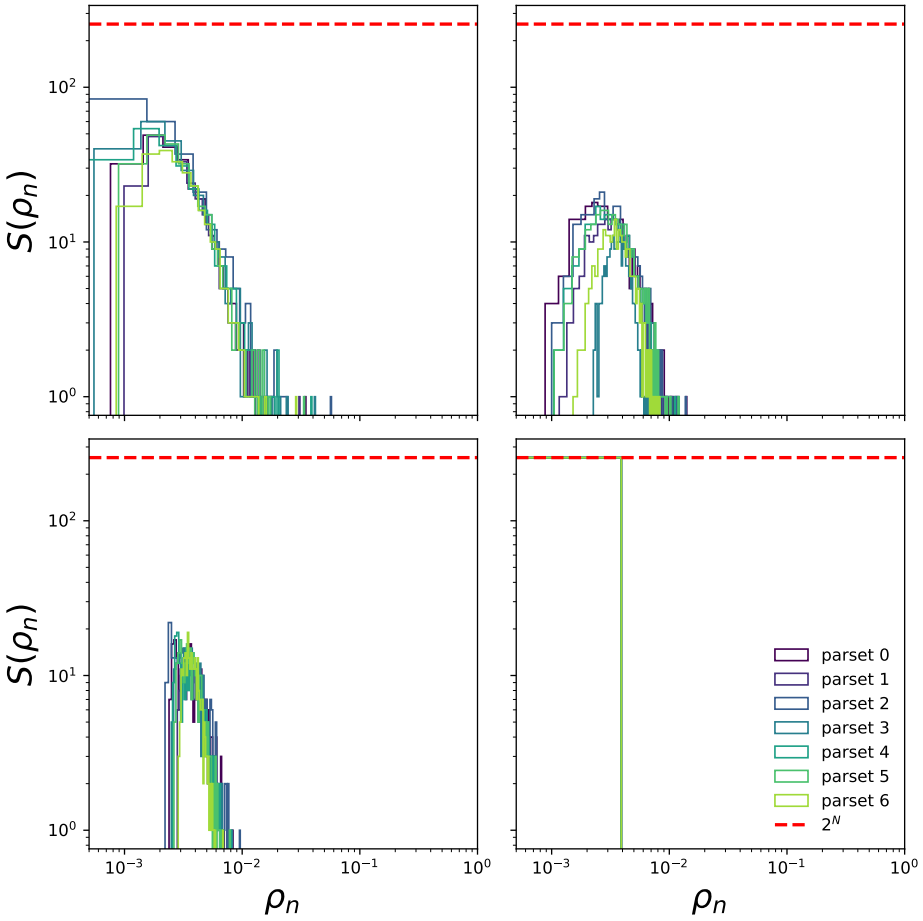

Figure B.8: Eigenvalue distributions of simulated density matrices for $N = 8$ qubits at $L = 10$ (left) and $L = 50$ (right) layers. Top row: amplitude-damping noise ($p_1 = 0.05$). The persistence of dominant eigenvalues indicates convergence to a noise-induced limit set. Bottom row: depolarising noise ($p_2 = 0.05$). The spectrum concentrates near $1/2^N$, consistent with convergence to the maximally mixed state and the onset of NIBP.

Appendix C

Extra information on Resources in Quantum optimization

This appendix collects additional details for Section 4.2, which investigates the build-up of nonstabilizerness in QAOA applied to the Sherrington–Kirkpatrick spin-glass model. Section C.1 specifies the classical optimiser settings and parameter initialisation strategy used in all simulations, and discusses the computational cost of evaluating the Stabilizer Rényi Entropy exactly. The remaining sections provide supplementary numerical results, including the degeneracy-breaking procedure for the SK Hamiltonian and additional data for the qutrit version of the protocol, supporting the scaling and universality claims made in the main text.

C.1 QAOA Optimisation Details

All QAOA simulations are executed using a custom implementation. The optimal parameters $\{\boldsymbol{\beta}, \boldsymbol{\gamma}\}$ are found using the Constrained Optimisation BY Linear Approximation (COBYLA) algorithm [250], a gradient-free method that solves the problem by linearly approximating the cost function. Parameters are initialised using the Trotterised Quantum Annealing (TQA) strategy [160]. The classical optimiser stops when the trust-region radius decreases below the tolerance 10^{-4}. The SRE is computed exactly, with computational cost scaling as D^{2N} due to the exponential growth in the number of Pauli strings, where D is the qudit dimension and N is the number of qudits.

C.2 Degeneracy Breaking in the SK Model

The SK Hamiltonian of Eq. (4.12) possesses symmetries that lead to degenerate ground states, which could introduce magic through superpositions of equally valid solutions rather than through the QAOA dynamics itself. The base Hamiltonian (without bias) reads

$$\hat{H} = \sum_{i \neq j} J_{ij} (Z_i Z_j^\dagger + Z_i^\dagger Z_j) \,, \tag{C.1}$$

which is symmetric under reflection $Z_i \leftrightarrow -Z_i$, creating two degenerate states. This degeneracy is lifted by a longitudinal field $\hat{H}_{\text{bias},2} = \sum_i h_i (Z_i + Z_i^\dagger)$, which reduces to $\sum_i 2h_i Z_i$ for qubits.

For qutrits, the Hamiltonian is additionally symmetric under $Z_i \to cZ_i$ with $c \in \{\pm 1, \pm\omega, \pm\omega^2\}$, introducing four further symmetries and resulting in six degenerate eigenstates (corresponding to the $3! = 6$ permutations of the local qutrit basis). The field $\hat{H}_{\text{bias},2}$ removes the degeneracy between $|0\rangle$ and $\{|1\rangle, |2\rangle\}$ — since $(Z + Z^\dagger) = \text{diag}(2, -1, -1)$ — but retains the degeneracy between $|1\rangle$ and $|2\rangle$. To lift this residual degeneracy, the term $\sum_i i h_i'(Z_i - Z_i^\dagger)$ is added, since $i(Z - Z^\dagger) = \text{diag}(0, -\sqrt{3}, +\sqrt{3})$. In all numerical experiments, h_i and h_i' are drawn from a normal distribution with variance 0.3.

C.3 Scaling of Nonstabilizerness

Although the states generated by QAOA strongly depend on circuit depth, optimisation parameters, and system size, a qualitatively uniform behaviour of the magic barrier is observed. For fixed N, the barrier height remains nearly constant while the curve stretches horizontally with increasing depth, motivating a scaling collapse.

Table C.1 reports the fitted parameters μ, ξ, ν, and η in Eq. (4.13) for $N = 4, 6, 8$ qubit systems.

N	μ	ξ	ν	η
4	$-6.47 \cdot 10^{-2}$	$1.64 \cdot 10^{-2}$	0.76	0.23
6	$5.11 \cdot 10^{-3}$	1.28	-0.34	0.60
8	$-6.83 \cdot 10^{-2}$	$1.38 \cdot 10^{-2}$	0.95	0.86

Table C.1: Scaling parameters μ, ξ, ν, and η for the data collapse at fixed system size N, fitted from QAOA magic curves for $N = 4, 6, 8$ qubits.

For systems exhibiting universality, the critical scaling point becomes a constant, motivating the simplified form

$$\frac{\mathcal{M}\big(|\psi_{\text{QAOA}_d}(p)\rangle\big)}{N} = d^{-\mu}\, f[(\lambda - \lambda_c) \cdot d^{\eta}], \quad \mu, \lambda_c, \eta \in \mathbb{R}, \quad \lambda = \frac{p}{d}. \tag{C.2}$$

Table C.2 reports the corresponding fitted parameters. Remarkably, the extracted λ_c values are consistent across system sizes for qubits, suggesting a universal critical point. This universality is not observed in the qutrit case.

N	μ	λ_c	η
4	$-6.51 \cdot 10^{-2}$	0.27	0.23
6	$6.50 \cdot 10^{-3}$	0.27	0.59
8	$-9.99 \cdot 10^{-2}$	0.29	0.92

Table C.2: Fitted scaling parameters μ, λ_c, and η using the critical-point form of Eq. (C.2), for $N = 4, 6, 8$ qubits.

C.4 Analytical Relation Between Nonstabilizerness and Fidelity

This section provides the full derivation of the analytical expressions for the SRE of the *ansatz* states discussed in Section 4.2.4.

Qubit case

The derivation of Eq. (4.15) begins with a single qubit and the state of Eq. (4.14). The four Pauli expectation values are:

$$\langle\psi|I|\psi\rangle = 1, \tag{C.3}$$
$$\langle\psi|Z|\psi\rangle = 2\mathcal{F} - 1, \tag{C.4}$$
$$\langle\psi|X|\psi\rangle = 2\sqrt{\mathcal{F}(1-\mathcal{F})}\cos\theta, \tag{C.5}$$
$$\langle\psi|Y|\psi\rangle = 2i\sqrt{\mathcal{F}(1-\mathcal{F})}\sin\theta\,. \tag{C.6}$$

Summing and simplifying yields

$$\sum_P |\,\langle\psi|P|\psi\rangle\,|^4 = 2\big(1 - 4\mathcal{F}(1-\mathcal{F}) + 14\mathcal{F}^2(1-\mathcal{F})^2 + 2\mathcal{F}^2(1-\mathcal{F})^2\cos 4\theta\big)\,. \tag{C.7}$$

The result generalises to any number of qubits n as follows. Since $|\phi_0\rangle$ and $|\phi_1\rangle$ are in the computational basis with Hamming distance $h < n$, the state factorises as $|\psi_{\mathrm{II}}\rangle = |\psi'_{\mathrm{II}}\rangle_h \otimes |d\rangle_{n-h}$, where $|d\rangle$ is a computational basis state. By additivity of SRE, the non-trivial contribution comes solely from the h-qubit factor. Using the invariance of the Pauli group under X-conjugation, one can reduce the calculation to the case $\sqrt{\mathcal{F}}\,|0\rangle^{\otimes n} + e^{i\theta}\sqrt{1-\mathcal{F}}\,|1\rangle^{\otimes n}$ without loss of generality. The contributions from $\{I, Z\}$-type Paulis and from $\{X, Y\}$-type Paulis are then:

$$\sum_{k=0}^{n}\binom{n}{k}\big|\mathcal{F} + (-1)^k(1-\mathcal{F})\big|^4 = 2^{n-1}\big[1 + (2\mathcal{F}-1)^4\big]\,, \tag{C.8}$$

$$\sum_{k=0}^{n}\binom{n}{k}\mathcal{F}^2(1-\mathcal{F})^2|e^{i\theta}(-i)^k + e^{-i\theta}(+i)^k|^4 = 2^{n-1}\mathcal{F}^2(1-\mathcal{F})^2\cdot 16[\cos^4\theta + \sin^4\theta]\,. \tag{C.9}$$

Combining gives the n-independent result:

$$M_2(|\psi_{\mathrm{II}}\rangle) = -\log_2\big[1 - 4\mathcal{F}(1-\mathcal{F}) + 2\mathcal{F}^2(1-\mathcal{F})^2(7 + \cos 4\theta)\big]\,. \tag{C.10}$$

The n-independence holds because the magic depends on the number of states in the superposition, not on the system dimension or the specific basis states.

For the three-state *ansatz* of Eq. (4.16), an analogous calculation yields the N-independent expression:

$$\begin{aligned} M_2 = -\log_2\Big[&1 + 14f^4 + 28f^3(p-1) + 2(p-1)p\big(2 + 7p(p-1)\big) + 6f^2\big(3 + p(7p-6)\big) \\ &+ 4f(p-1)\big(1 + p(7p-6)\big) + 2f^2(f+p-1)^2\cos 4\phi \\ &+ 4p^2f^2\cos 2\phi\cos(2\phi - 4\theta) + 2p^2(p-1)(2f+p-1)\cos(4\phi - 4\theta)\Big]\,. \end{aligned} \tag{C.11}$$

The validity of both expressions is verified numerically in Fig. C.1.

Qutrit case

For qutrits, the Pauli group is $\{X^rZ^j \mid r, j \in \{0,1,2\}\}$. Repeating the calculation for the state $|\psi\rangle = \sqrt{\mathcal{F}}\,|a\rangle + e^{i\theta}\sqrt{1-\mathcal{F}}\,|b\rangle$ gives contributions from Z-type and XZ-type Paulis:

$$\begin{aligned} e^{-M_2} = 1 &+ |\omega^a\mathcal{F} + (1-\mathcal{F})\omega^b|^4 + |\omega^{2a}\mathcal{F} + (1-\mathcal{F})\omega^{2b}|^4 \\ &+ \mathcal{F}^2(1-\mathcal{F})^2\sum_{r=1,2;\,j=0,1,2}|e^{i\theta}\omega^{jb}\delta_{a,b\oplus r} + e^{-i\theta}\omega^{ja}\delta_{b,a\oplus r}|^4\,. \end{aligned} \tag{C.12}$$

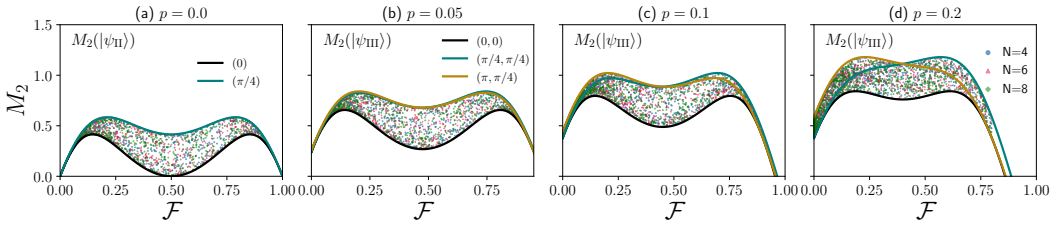

Figure C.1: SRE vs. fidelity for superpositions of two (a) and three (b,c,d) states at different values of p. Continuous lines represent the analytic bounds of Eqs. (C.10) and (C.11). Dots are numerically computed SRE values for random states of the form $|\psi_{\mathrm{II}}\rangle$ (a) and $|\psi_{\mathrm{III}}\rangle$ (b,c,d). All points lie within the analytic bounds, confirming the validity of the expressions.

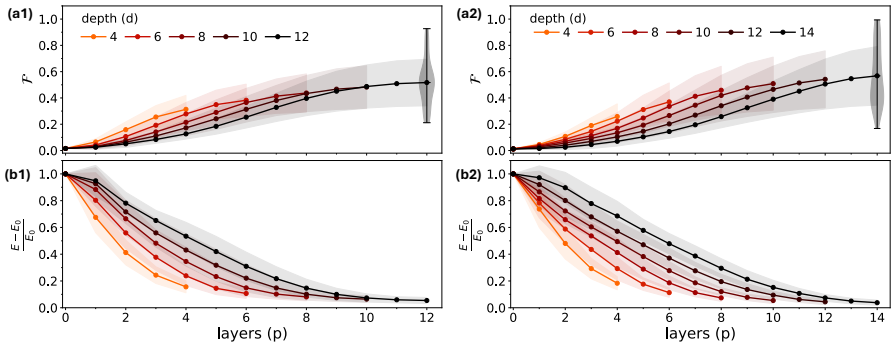

Figure C.2: (a1) Mean fidelity over 50 different realisations, with relative standard deviation, as a function of layer index for different depths of the 6-qubit system. The violin plot for depth 12 highlights the high fidelity achieved by some realisations. (b1) Mean relative energy under the same conditions. (a2) and (b2) show the corresponding results for a 4-qutrit system.

For qutrits, $a = b \oplus r$ and $b = a \oplus r$ cannot be satisfied simultaneously, removing any dependence on the relative phase θ. For $a = 0$, $b = 1$:

$$e^{-M_2} = 1 + |\mathcal{F} + (1-\mathcal{F})\omega|^4 + |\mathcal{F} + (1-\mathcal{F})\omega^2|^4 + 3\mathcal{F}^2(1-\mathcal{F})^2\,. \tag{C.13}$$

C.5 Energy and Fidelity as a Function of Layer Index

For completeness, the evolution of the approximation ratio $(E - E_0)/E_0$ and fidelity $\mathcal{F}$ of the optimised QAOA state with the problem solution are shown in Fig. C.2. Both quantities improve monotonically with the layer index, despite the non-monotonic behaviour of nonstabilizerness. This confirms that the magic barrier is a genuine feature of the QAOA dynamics and not an artefact of suboptimal convergence. It also implies that the quantum computational resources required during the intermediate steps of QAOA exceed those needed for direct simulation of the final target state — an important consideration for fault-tolerant device benchmarking.

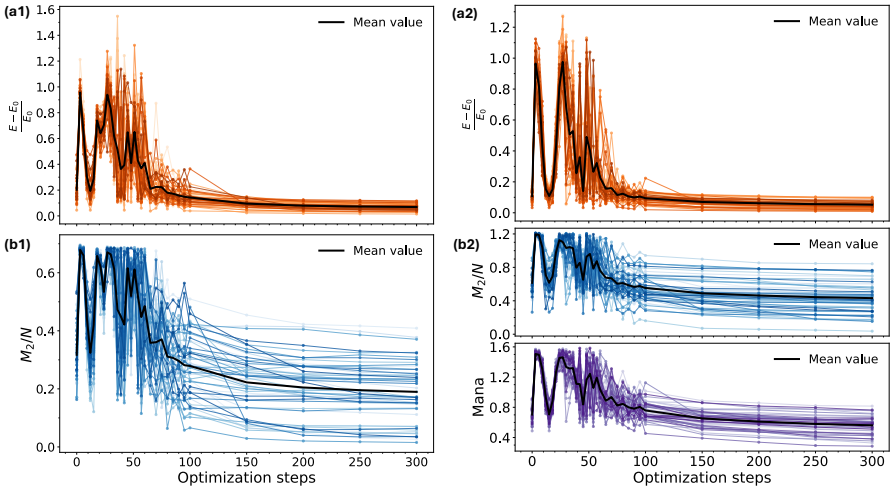

Figure C.3: Energy and nonstabilizerness during the optimisation loop. Black curves show mean values over 50 realisations. (a1), (a2) Energy ratio as a function of optimisation steps for 6-qubit and 4-qutrit systems, respectively. (b1) SRE as a function of optimisation steps for the 6-qubit system. (b2) SRE (upper) and Mana (lower) for the 4-qutrit system.

C.6 Nonstabilizerness During the Variational Optimisation Loop

To characterise the nonstabilizerness cost during the classical optimisation loop, its evolution over the course of parameter updates is analysed. 50 problem instances are considered, with 50 QAOA runs per instance initialised with annealing-inspired parameters (decreasing γ, increasing β) [160]. For each instance, the run achieving the lowest energy is selected.

Figure C.3 shows energy and nonstabilizerness as a function of optimisation steps for the best run of each of the 50 realisations, for a 6-qubit system at depth 12 and a 4-qutrit system at depth 14. During optimisation, the energy initially increases, crosses a barrier, and converges close to E_0. This energy barrier is accompanied by a peak in nonstabilizerness, suggesting that crossing a high-magic region is required for convergence. A similar pattern is observed for Mana in the 4-qutrit case.

The final magic values remain widely spread across instances, even as the energy approximation ratio consistently approaches zero. This highlights a key distinction: energy convergence does not imply convergence to the exact solution. Even near the ground-state energy, QAOA may converge to a superposition of low-lying eigenstates when energy level splittings are small — an effect explored further in the fidelity analysis of Section 4.2.4.

C.7 Final Results as a Function of QAOA Depth

Figure C.4 presents final energy, fidelity, and nonstabilizerness as a function of circuit depth for 6-qubit and 4-qutrit systems. The energy converges to near-ground-state values for sufficiently high depth. Fidelity is low for shallow circuits and increases with depth, though never reaching unity. SRE (and Mana) rises rapidly at small depths, reaches a maximum, and then decreases without reaching zero — a consequence of the algorithm converging to

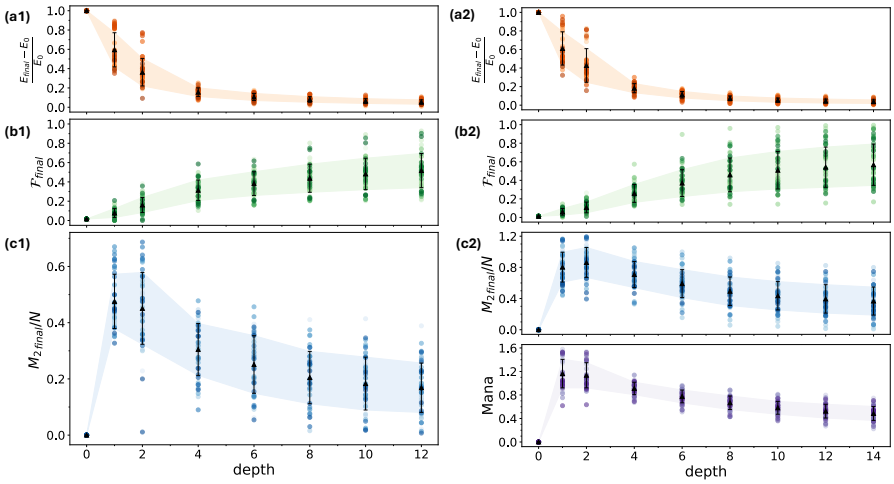

Figure C.4: Final results from the optimised QAOA state as a function of circuit depth, for a 6-qubit system and a 4-qutrit system across 50 random SK realisations. (a1), (b1), (c1): energy ratio (orange), fidelity (green), and SRE density (blue) for the qubit system. (a2), (b2), (c2): energy ratio (orange), fidelity (green), SRE density (blue), and Mana (violet) for the qutrit system. Coloured dots: individual realisations; black triangles: averages; shaded regions: standard deviation.

superpositions of low-lying eigenstates rather than the exact ground state.

Figure C.5 further illustrates that many optimised states reach near-ground-state energy while retaining nonzero nonstabilizerness, confirming that energy convergence does not guarantee collapse onto the true solution state.

C.8 Additional Results for Qutrit Systems

This section presents complementary results for the 4-qutrit system, paralleling the qubit results in the main chapter.

Final Nonstabilizerness and Fidelity

Figure C.6 shows the nonstabilizerness versus fidelity for the 4-qutrit system. The qualitative structure is analogous to the qubit case of Fig. 4.4: high fidelity is associated with low final nonstabilizerness, while low nonstabilizerness does not guarantee high fidelity.

Demagication and Success Probability

Figure C.7 shows the conditional probability analysis of Eq. (4.17) for the 4-qutrit system. The results are consistent with the qubit case: larger demagication is associated with higher probability of achieving high fidelity, confirming that the connection between demagication and algorithmic success is independent of the local Hilbert space dimension.

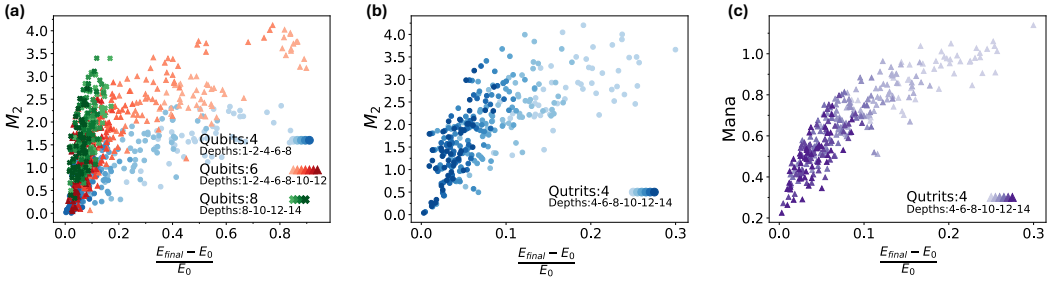

Figure C.5: Final SRE (and Mana) as a function of the final approximation ratio, for all data collected across different system sizes and circuit depths. (a) Qubit systems. (b,c) 4-qutrit system. Many solutions achieve near-ground-state energy while retaining nonzero nonstabilizerness.

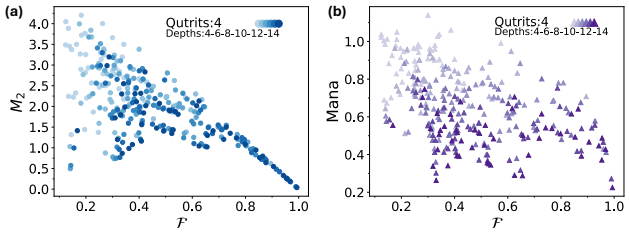

Figure C.6: SRE (a) and Mana (b) as a function of fidelity for a system of four qutrits. Different shades represent different QAOA depths. The qualitative structure mirrors the qubit case shown in Fig. 4.4.

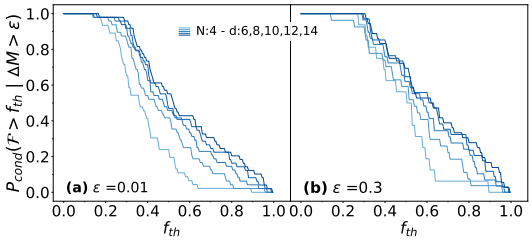

Figure C.7: Conditional probability of achieving fidelity above a given threshold, given a minimum demagication, for a 4-qutrit system at QAOA depths 6–14. (a) $\epsilon = 0.01$; (b) $\epsilon = 0.03$. As in the qubit case (Fig. 4.6), higher demagication thresholds correspond to higher success probabilities.

Appendix D

Further details on D_3 lattice gauge theory

This appendix provides technical details complementing Chapter 5, which analyses the ground-state structure and phase transitions of the D_3 lattice gauge theory in 1+1 dimensions. Section D.1 lists the explicit 6×6 matrix form of the left and right gauge transformation operators for the generators r and s of D_3 in the group-element basis. Section D.2 gives the corresponding transformation matrices for the fermionic matter fields in the faithful representation τ. Section D.3 reports the full parallel-transporter matrix U^τ in the representation basis, from which the hopping term of the Kogut–Susskind Hamiltonian is constructed.

D.1 Explicit form of the gauge tranformaions

Here we list the explicit form of the left and right transformation operators in the group element basis for the electric fields. The matrices shows are only for the generators r and s elements of the group, and the others can be obtained by usual matrix multiplications via $\Theta(gh) = \Theta(g)\Theta(h)$

$$\Theta^L(r) = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \qquad \Theta^L(s) = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\Theta^R(r) = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \qquad \Theta^R(s) = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

D.2 Fermion transformations

$$\Theta^{\psi}(r^j) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{i\frac{2j\pi}{3}} & 0 & 0 \\ 0 & 0 & e^{-i\frac{2j\pi}{3}} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \qquad \Theta^{\psi}(sr^j) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & e^{-i\frac{2j\pi}{3}} & 0 \\ 0 & e^{i\frac{2j\pi}{3}} & 0 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \tag{D.1}$$

D.3 Explicit form of the parallel transporter

Here we list all the components of U^{τ} in the representation basis (the order of the columns is e, p, rr, rg, gr, gg)

$$U_{ab} = \left(\begin{array}{cc|cccc} 0 & 0 & \frac{1}{\sqrt{2}}d_{rr} & \frac{1}{\sqrt{2}}d_{rg} & \frac{1}{\sqrt{2}}d_{gr} & \frac{1}{\sqrt{2}}d_{gg} \\ \hline 0 & 0 & \frac{1}{\sqrt{2}}d_{rr} & \frac{-1}{\sqrt{2}}d_{rg} & \frac{-1}{\sqrt{2}}d_{gr} & \frac{1}{\sqrt{2}}d_{gg} \\ \hline \frac{1}{\sqrt{2}}d_{gg} & \frac{1}{\sqrt{2}}d_{gg} & 0 & 0 & 0 & d_{rr} \\ \frac{1}{\sqrt{2}}d_{gr} & \frac{-1}{\sqrt{2}}d_{gr} & 0 & 0 & d_{rg} & 0 \\ \frac{1}{\sqrt{2}}d_{rg} & \frac{-1}{\sqrt{2}}d_{rg} & 0 & d_{gr} & 0 & 0 \\ \frac{1}{\sqrt{2}}d_{rr} & \frac{1}{\sqrt{2}}d_{rr} & d_{gg} & 0 & 0 & 0 \end{array}\right), \tag{D.2}$$

where $d_{\alpha\beta} = \delta_{a\alpha}\delta_{b\beta}$.

Appendix E

Confinement in Dihedral Lattice Gauge Theories

This appendix contains the analytical derivations supporting Section 6.1, which studies confinement and color screening in pure D_N lattice gauge theories and relates these phenomena to the structure of the group center. Section E.1 provides a rigorous proof, based on the character theory of dihedral groups, that the screening fusion rule holds for odd N — allowing gluonic excitations to break the electric string — while it fails for even N, where the non-trivial $\mathbb{Z}_2$ center enforces confinement over arbitrary distances. The proof covers the general case of l-fold tensor products of the highest-dimensional irreducible representation τ_h, unifying the even- and odd-N phenomenology within a single character-theoretic framework.

E.1 Fusion Rules of Dihedral Groups

This section provides a proof that the screening fusion rule of Eq. (6.1) holds for odd N but not for even N, as claimed in Section 6.1.1. The argument is based on the character theory of the dihedral groups D_N.

By the great orthogonality theorem for finite groups, any representation J of G decomposes as a direct sum over irreducible representations:

$$J = \bigoplus_{j \text{ irrep}} j^{\otimes a_j} \,, \tag{E.1}$$

where a_j is the multiplicity of the j-th irrep. The multiplicities are computed via the inner product of characters,

$$a_j = (\chi_j, \chi_J) := \frac{1}{|G|} \sum_{g \in G} \overline{\chi_j(g)}\, \chi_J(g) = \frac{1}{|G|} \sum_{C} |C|\, \overline{\chi_j(C)}\, \chi_J(C) \,, \tag{E.2}$$

where $\chi_J(g)$ is the character of element g in representation J, the bar denotes complex conjugation, and the second equality exploits the fact that characters are constant on conjugacy classes C. When $a_j = 0$, the irrep j does not appear in the decomposition of J.

The aim is to show that the decomposition of $J = \tau_h \otimes \tau_h$ contains the fundamental irrep τ — i.e. $a_\tau = 1$ — for odd N, while $a_\tau = 0$ for even N. Since the D_N characters are real,

the fusion rule can be rewritten as $\tau_h \otimes \tau_h \ni \tau$. The more general case $J = \tau_h^{\otimes l}$ with l even is considered, requiring the evaluation of

$$(\chi_\tau, \chi_{\tau_h^{\otimes l}}) = \frac{1}{|D_N|} \sum_C |C| \, \overline{\chi_{\tau_h}(C)}^{\,l} \, \chi_\tau(C) \,, \tag{E.3}$$

where the identity $\chi_{\tau_h^{\otimes l}} = \chi_{\tau_h}^l$ has been used.

For N even, the sum over conjugacy classes gives

$$\begin{aligned} |D_N|(\chi_\tau, \chi_{\tau_h^{\otimes l}}) = {} & \chi_\tau(\mathbb{1})\chi_{\tau_h}^l(\mathbb{1}) + \chi_\tau(r^{N/2})\chi_{\tau_h}^l(r^{N/2}) \\ & + 2 \sum_{k=1}^{N/2-1} \chi_\tau(r^k)\chi_{\tau_h}^l(r^k) + \frac{N}{2}\chi_\tau(s)\chi_{\tau_h}^l(s) + \frac{N}{2}\chi_\tau(sr)\chi_{\tau_h}^l(sr) \,. \end{aligned} \tag{E.4}$$

Substituting the D_N character values (Eq. (**??**)):

$$|D_N|(\chi_\tau, \chi_{\tau_h^{\otimes l}}) = 2^{l+1} + 2^{l+1}(-1)^{hl+1} + 2^{l+2} \sum_{k=1}^{N/2-1} \cos^l\!\left(\frac{2\pi hk}{N}\right) \cos\!\left(\frac{2\pi k}{N}\right) . \tag{E.5}$$

The second term carries an overall negative sign (since l is even) and cancels exactly the first term. Both contributions arise from the conjugacy classes of the $\mathbb{Z}_2$ center $Z = \{\mathbb{1}, r^{N/2}\}$, and their cancellation is a direct consequence of the non-trivial center.

The remaining sum is identically zero for all allowed h, by the following case analysis:

- If $N/2 - 1$ is even, the summands cancel pairwise.
- If $N/2 - 1$ is odd, all but the $k = N/4$ summand cancel pairwise, and the surviving term vanishes since $\cos(\pi/2) = 0$.

Therefore, $a_\tau = 0$ for all even N.

For N odd, the element $r^{N/2}$ is not in the group, so the $\mathbb{Z}_2$ center is absent and the cancellation no longer occurs. The inner product reduces to

$$|D_N|(\chi_\tau, \chi_{\tau_h^{\otimes l}}) = \chi_\tau(\mathbb{1})\chi_{\tau_h}^l(\mathbb{1}) + 2 \sum_{k=1}^{(N-1)/2} \chi_\tau(r^k)\chi_{\tau_h}^l(r^k) + N\chi_\tau(s)\chi_{\tau_h}^l(s) \,, \tag{E.6}$$

which, after substituting character values, gives

$$|D_N|(\chi_\tau, \chi_{\tau_h^{\otimes l}}) = 2^{l+1} + 2^{l+2} \sum_{k=1}^{(N-1)/2} \cos^l\!\left(\frac{2\pi hk}{N}\right) \cos\!\left(\frac{2\pi k}{N}\right) . \tag{E.7}$$

The uncancelled 2^{l+1} term shows that the inner product can be non-zero.

Case $N = 3$. For $N = 3$, $l = 2$, and $h = 1$ (the unique two-dimensional irrep $\tau_{h=1} \equiv \tau$):

$$|D_3|(\chi_\tau, \chi_{\tau^{\otimes 2}}) = 2^3 + 2^4 \left(-\frac{1}{8}\right) = 6 \,, \tag{E.8}$$

giving $a_\tau = 1$.

Figure E.1: Comparison of the fusion rules of D_3 (left) and D_4 (right). Arrows indicate non-vanishing Clebsch–Gordan coefficients. For D_3, the fundamental representation τ has a non-vanishing CG coefficient with itself, permitting the screening fusion rule. For D_4, only terms connecting τ with the one-dimensional representations are non-zero, and the rule is forbidden by the $\mathbb{Z}_2$ center.

General odd $N > 3$. For $N > 3$, the fundamental representation $\tau \equiv \tau_{h=1}$ does not satisfy the screening fusion rule — $(\chi_\tau, \chi_{\tau^{\otimes 2}}) = 0$ — but the faithful representation with the largest index $h = (N-1)/2$, the *anti-fundamental* representation $\tilde{\tau}$, does. Setting $h = (N-1)/2$ and $l = 2$ in Eq. (E.7) yields

$$(\chi_\tau, \chi_{\tilde{\tau}^{\otimes 2}}) = 1 \tag{E.9}$$

for all odd N. This means that a vertex state with two semilinks in the $\tilde{\tau}$ representation is gauge invariant, satisfying the screening fusion rule of Eq. (6.1).

For N even, $\tau_h \otimes \tau_h \not\ni \tau$, and the screening fusion rule does not hold. For N odd, $\tau_h \otimes \tau_h = \cdots \oplus \tau$ when $h = (N-1)/2$, and the rule is satisfied by the anti-fundamental representation $\tilde{\tau}$. The fusion rules of D_3 and D_4 are compared graphically in Fig. E.1.

The $\tilde{\tau}$-gluelumps are dynamically connected to the τ-flux string connecting the two static charges. This is established by evaluating

$$|D_N|(\chi_{\tau^{\otimes l+1}}, \chi_{\tilde{\tau}}) = 2^{l+2} + 2^{l+3} \sum_{k=1}^{(N-1)/2} \cos\left(\frac{\pi k(N-1)}{N}\right) \cos^{l+1}\left(\frac{2\pi k}{N}\right). \tag{E.10}$$

The smallest integer l for which the inner product $(\chi_\tau, \chi_{\tilde{\tau}^{\otimes 2}}) = 1$ is

$$l = \begin{cases} 0 & N = 3\,, \\ \dfrac{N-3}{2} & N > 3\,. \end{cases} \tag{E.11}$$

This has a direct physical interpretation: the transition from the fundamental flux string to the $\tilde{\tau}$ loops requires the application of the group connection U^τ (as part of plaquette operators) exactly $l(N)$ times. String breaking therefore becomes a higher-order process as N increases, rendering it non-perturbative in the limit of odd $N \to \infty$.

E.2 Gauge Transformations, Static Charges, and Fusion Rules

The decomposition $J = \bigoplus_{j\,\mathrm{irrep}} j^{\otimes a_j}$ of general representations into irreducible representations has direct implications for the description of the ladder in the presence of static charges at its corners.

Consider the top-left corner of Fig. 1(a) of the main text (the static charge at the opposite corner has analogous behaviour). In the presence of a static charge τ, there is a multiplet of physical states $|\Psi_{\mathrm{phys},m}\rangle$ ($m = 1, \ldots, \dim(\tau)$) transforming as the τ irreducible representation under the local gauge transformation Θ_{h,v_τ}, where v_τ labels the top-left corner:

$$\theta^R_{h,r}\theta^L_{h,l}\,|\Psi_{\mathrm{phys},m}\rangle = \sum_n D^\tau_{n,m}(h)\,|\Psi_{\mathrm{phys},n}\rangle\ , \tag{E.12}$$

where r labels the first rung, l is the first link in the upper leg, and $D^\tau_{n,m}(h)$ is the 2×2 matrix representing the group element h in the fundamental representation τ. Expanding $|\Psi_{\mathrm{phys},m}\rangle$ in terms of the link states $|j_i, m_i, n_i\rangle$ ($i = r, l$) and the static-charge state $|\tau, m\rangle$:

$$|\Psi_{\mathrm{phys},m}\rangle = \sum_{j_r,j_l,n_r,m_l} A^{j_r,j_l}_{n_r,n_l}\,|j_r, m_r, n_r\rangle\,|j_l, m_l, n_l\rangle\,|\tau, m\rangle\ . \tag{E.13}$$

Applying the gauge transformation to the left-hand side of Eq. (E.12) and using the Clebsch–Gordan series relation [251],

$$D^{j_r}_{n'_r,n_r}(h^{-1})D^\tau_{m',m}(h^{-1}) = \sum_{I,n,n'} \langle j_r n'_r\,\tau m'|In'\rangle\langle In|j_r n_r\,\tau m\rangle\,D^I_{n'n}(h^{-1})\,, \tag{E.14}$$

the gauge-invariance condition is satisfied if and only if the coefficients A take the form [185]

$$A^{j_r,j_l}_{n_r,n_l} = \alpha(j_r, j_l)\,\langle j_r n_r\,\tau m \mid j_l m_l\rangle\ , \tag{E.15}$$

where α depends only on the irreps j_l, j_r and the angle bracket denotes the conjugate of a CG coefficient. This establishes the fusion rule $\tau \otimes j_r \ni j_l$: the ingoing representations j_r and τ must fuse into j_l.

To realise the gluelump state, both links must be in the same irreducible representation and combine with the charge into a group singlet, which is verified for $j_r = j_l = \tilde{\tau}$ and gives rise to Eq. (4) of the companion paper. An analogous construction applies at vertices with three links in the bulk of the ladder, where τ is replaced by the irrep of the additional link (see Ref. [185] for the 2D case).

E.3 Rishons and Gauge-Invariant Basis

This appendix describes in detail the local gauge-invariant basis used for the numerical calculations and its relation to gauge transformations and the Hamiltonian.

E.3.1 Local Dressed Basis in the Rishon Formalism

The D_N LGT is reformulated in a local dressed-site basis, enabling efficient tensor-network implementation [185, 186, 252, 89, 23]. The non-Abelian Gauss's law on each vertex is solved by splitting each link into two halves (semilinks), each attached to its adjacent vertex. The

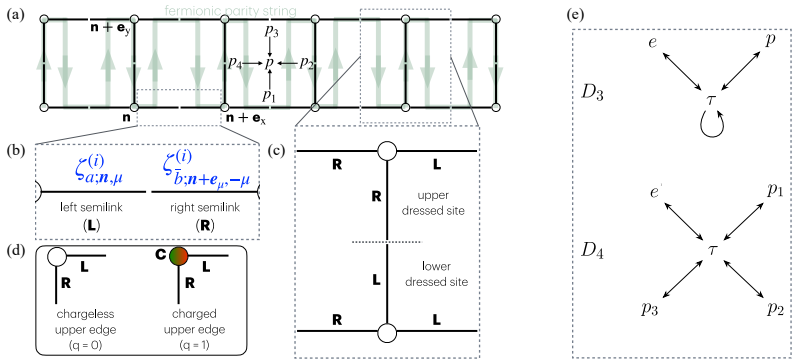

Figure E.2: (a) Sketch of the ladder geometry, highlighting the plaquette operator H_B. The shaded path indicates the ordering of semilinks used to define the fermionic rishon parity in Sec. E.3.3. (b) Splitting of a gauge field on a link into two rishon operators. (c) Dressed sites of the ladder, each consisting of three semilinks. Sites on the upper legs have two ingoing R- and one outgoing L-semilink; sites on the lower legs have the opposite. (d) The upper corners where static charges are inserted to probe confinement. (e) Comparison of the fusion rules of D_3 (left) and D_4 (right). Arrows indicate non-vanishing CG coefficients. For D_3, the fundamental representation τ has a non-vanishing CG coefficient with itself; for D_4, only terms connecting τ with the one-dimensional representations are non-zero.

split is reversed once the selection rules are imposed on semilinks corresponding to the same link.

In the electric basis, the parallel transporter U_{ab} in the fundamental representation τ can be expressed as [**Zohar'PRD2015**]

$$\begin{aligned}\langle j_1; m_1, n_1 |\, U_{ab} \,| j_2; m_2, n_2 \rangle &= \sqrt{\frac{\dim(j_2)}{\dim(j_1)}}\, \langle j_2, m_2; \tau, a \rangle\, j_1, {m_1}^* \, \langle j_2, n_2; \tau, b \rangle\, j_1, n_1 \\ &= \langle \tau, a; \bar{j}_1, \bar{m}_1 \rangle\, \bar{j}_2, {\bar{m}_2}^* \, \langle j_2, n_2; \tau, b \rangle\, j_1, n_1 \,,\end{aligned} \tag{E.16}$$

expressed via CG coefficients (or $3J$-symbols). The fusion rules factorize over the two indices of multi-dimensional representations. Each link is therefore split into left and right semilinks, $|j, m_L, m_R\rangle = |j_1, m_L\rangle \otimes |j_2, m_R\rangle\, \delta_{j_1,j_2} \delta_{j,j_1}$, and the parallel transporter factorises into rishon operators acting separately on each semilink Hilbert space [Fig. E.2(b)]:

$$U_{ab;\mathbf{n},\mathbf{n}+\mathbf{e}_\mu} = \sum_i \zeta^{(i)}_{a;\mathbf{n},\mu}\, \zeta^{(i)\dagger}_{\bar{b};\mathbf{n}+\mathbf{e}_\mu,-\mu} \,, \tag{E.17}$$

where i labels the rishon modes and $\pm\mu$ indicates the position of the semilink relative to site $\mathbf{n}$.

The left and right multiplication operators in the group-element basis are

$$\theta^L_h\, |g\rangle = |hg\rangle \;, \qquad \theta^R_h\, |g\rangle = |gh^{-1}\rangle \;, \tag{E.18}$$

and on the rishon spaces they act as

$$\theta^L_g = \Lambda^L_g \otimes \mathbb{1}_d \,, \qquad \theta^R_g = \mathbb{1}_d \otimes \Lambda^R_g \,. \tag{E.19}$$

Upper dressed sites are constructed from two right (R) and one left (L) semilink, and lower dressed sites from one right and two left semilinks [Fig. E.2(c)].

D_3 representation basis and rishon operators. The representation basis for D_3 is labelled $\{|0\rangle, |p\rangle, |rr\rangle, |rg\rangle, |gr\rangle, |gg\rangle\}$, where $|0\rangle$ is the trivial irrep, $|p\rangle$ the parity irrep, and $|\tau ab\rangle$ the fundamental irrep with colour indices 'red' and 'green'. The parallel transporter is

$$U_{ab} = \left(\begin{array}{c|c|cccc} 0 & & \frac{1}{\sqrt{2}}d_{rr} & \frac{1}{\sqrt{2}}d_{rg} & \frac{1}{\sqrt{2}}d_{gr} & \frac{1}{\sqrt{2}}d_{gg} \\ \hline & 0 & \frac{1}{\sqrt{2}}d_{rr} & \frac{-1}{\sqrt{2}}d_{rg} & \frac{-1}{\sqrt{2}}d_{gr} & \frac{1}{\sqrt{2}}d_{gg} \\ \hline \frac{1}{\sqrt{2}}d_{gg} & \frac{1}{\sqrt{2}}d_{gg} & & & & d_{rr} \\ \frac{1}{\sqrt{2}}d_{gr} & \frac{-1}{\sqrt{2}}d_{gr} & & & d_{rg} & \\ \frac{1}{\sqrt{2}}d_{rg} & \frac{-1}{\sqrt{2}}d_{rg} & & d_{gr} & & \\ \frac{1}{\sqrt{2}}d_{rr} & \frac{1}{\sqrt{2}}d_{rr} & d_{gg} & & & \end{array}\right), \tag{E.20}$$

where $d_{\alpha\beta} = \delta_{a\alpha}\delta_{b\beta}$. Notably, U_{ab} is not block-off-diagonal: acting with the parallel transporter on a link in the fundamental representation τ can still result in the same irrep, reflecting the screening fusion rule. Since D_3 has three distinct flux channels, three rishon modes are required:

$$\zeta_a^{(1)} = 2^{-1/4}\left(\begin{array}{c|c|cc} 0 & & \delta_{ar} & \delta_{ag} \\ \hline & 0 & 0 & 0 \\ \hline \delta_{ag} & 0 & 0 & \\ \delta_{ar} & 0 & & 0 \end{array}\right), \quad \zeta_a^{(2)} = 2^{-1/4}\left(\begin{array}{c|c|cc} 0 & & 0 & 0 \\ \hline & 0 & \delta_{ar} & -\delta_{ag} \\ \hline 0 & \delta_{ag} & 0 & \\ 0 & -\delta_{ar} & & 0 \end{array}\right),$$

$$\zeta_a^{(3)} = \left(\begin{array}{c|c|cc} 0 & & 0 & 0 \\ \hline & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & \delta_{ar} \\ 0 & 0 & \delta_{ag} & 0 \end{array}\right). \tag{E.21}$$

Gauge transformations in the rishon basis. All D_N transformations are generated by powers of r and s. In the full 6-dimensional basis, the transformations $\theta_r^{L(R)}$ and $\theta_s^{L(R)}$ are

$$\begin{aligned} &\theta_r^L = \mathrm{diag}(1,1,\omega,\omega,\omega^{-1},\omega^{-1}), \quad \omega = e^{2\pi i/3}, \\ &\theta_s^L: \quad |rr\rangle \leftrightarrow |gr\rangle\,,\ |rg\rangle \leftrightarrow |gg\rangle\,,\ |p\rangle \to -|p\rangle\,, \end{aligned} \tag{E.22}$$

with θ_r^R and θ_s^R defined analogously (see the companion paper for the full matrices). In the 4-dimensional rishon basis $\{|0\rangle, |p\rangle, |r\rangle, |g\rangle\}$, the reduced transformation operators are:

$$\Lambda_r^L = \mathrm{diag}(1,1,\omega,\omega^{-1}), \quad \Lambda_r^R = \mathrm{diag}(1,1,\omega^{-1},\omega), \quad \Lambda_s^L = \Lambda_s^R = \begin{pmatrix} 1 & & & \\ & -1 & & \\ & & 0 & 1 \\ & & 1 & 0 \end{pmatrix}. \tag{E.23}$$

The dressed-site transformations are

$$\begin{aligned} \Lambda_{r/s}^{\mathrm{up}} &= \Lambda_{r/s}^R \otimes \Lambda_{r/s}^R \otimes \Lambda_{r/s}^L\,, \\ \Lambda_{r/s}^{\mathrm{down}} &= \Lambda_{r/s}^R \otimes \Lambda_{r/s}^L \otimes \Lambda_{r/s}^L\,, \end{aligned} \tag{E.24}$$

and for the four corners:

$$\Lambda^{\text{up,left}}_{r/s} = \Lambda^{R}_{r/s} \otimes \Lambda^{L}_{r/s}, \qquad \Lambda^{\text{up,right}}_{r/s} = \Lambda^{R}_{r/s} \otimes \Lambda^{R}_{r/s},$$
$$\Lambda^{\text{down,left}}_{r/s} = \Lambda^{L}_{r/s} \otimes \Lambda^{L}_{r/s}, \qquad \Lambda^{\text{down,right}}_{r/s} = \Lambda^{R}_{r/s} \otimes \Lambda^{L}_{r/s}. \tag{E.25}$$

Gauge-invariant dressed basis. The physical states in the neutral sector are projected onto the eigenspace of Eqs. (E.24)–(E.25) with eigenvalue 1:

$$\Lambda_r \Lambda_s \left|\psi_{\text{phys}}\right\rangle = \left|\psi_{\text{phys}}\right\rangle . \tag{E.26}$$

In D_3, the resulting gauge-invariant dressed basis for any bulk vertex on the upper leg (semilinks ordered as left, below, right) is 11-dimensional:

$$\begin{aligned}
&|1\rangle = |0,0,0\rangle\,, \quad |2\rangle = \tfrac{1}{\sqrt{2}}(|r,0,r\rangle + |g,0,g\rangle),\\
&|3\rangle = \tfrac{1}{\sqrt{2}}(|0,r,r\rangle + |0,g,g\rangle), \quad |4\rangle = \tfrac{1}{\sqrt{2}}(|r,g,0\rangle + |g,r,0\rangle),\\
&|5\rangle = \tfrac{1}{\sqrt{2}}(|r,r,g\rangle + |g,g,r\rangle), \quad |6\rangle = \tfrac{1}{\sqrt{2}}(|r,p,r\rangle - |g,p,g\rangle),\\
&|7\rangle = \tfrac{1}{\sqrt{2}}(|p,r,r\rangle - |p,g,g\rangle), \quad |8\rangle = \tfrac{1}{\sqrt{2}}(|r,g,p\rangle - |g,r,p\rangle),\\
&|9\rangle = |p,p,0\rangle\,, \quad |10\rangle = |p,0,p\rangle\,, \quad |11\rangle = |0,p,p\rangle\ .
\end{aligned} \tag{E.27}$$

For dressed sites on the lower leg, the basis is obtained by conjugating the colour indices $r \leftrightarrow g$ in Eq. (E.27).

In the presence of static charges (fundamental representation), the physical condition is modified to

$$\Lambda_r \theta^{\psi}_r \Lambda_s \theta^{\psi}_s \left|\psi_{\text{phys}}\right\rangle = \left|\psi_{\text{phys}}\right\rangle , \tag{E.28}$$

where the matter transformation laws coincide with those of the right rishons. The gauge-invariant local dimension is 5 (with static charges) or 3 (without) for the D_3 corners, and 8 or 5 for D_4. The full D_4 dressed basis follows the same construction and yields a 11-dimensional bulk dressed-site space, increasing to 28 dimensions per site.

E.3.2 Hamiltonian in the Rishon Basis

The Hamiltonian is projected onto the gauge-invariant dressed basis constructed above. Each semilink contributes half the electric-field energy. The electric Hamiltonian $H_E = \frac{c\hbar}{a} g^2 \sum_{l\in\text{links}} \sum_j \alpha_j \Pi_l^j$ therefore becomes

$$H_E = \frac{c\hbar}{a} g^2 \sum_{sl\in\text{semilinks}} \sum_j \frac{\alpha_j}{2} \Pi^j_{sl} = \sum_{\mathbf{n}} H^{(1)}_{E,\mathbf{n}}\,, \tag{E.29}$$

where $H_E^{(1)}$ is the single-site electric Hamiltonian. In the 11-dimensional D_3 basis of Eq. (E.27):

$$\begin{aligned}
H_E^{(1)} = \frac{c\hbar}{a}\frac{g^2}{2} \operatorname{diag}(&3\alpha_0,\ 2\alpha_\tau+\alpha_0,\ 2\alpha_\tau+\alpha_0,\ 2\alpha_\tau+\alpha_0,\ 3\alpha_\tau,\\
&2\alpha_\tau+\alpha_p,\ 2\alpha_\tau+\alpha_p,\ 2\alpha_\tau+\alpha_p,\ 2\alpha_p+\alpha_0,\\
&2\alpha_p+\alpha_0,\ 2\alpha_p+\alpha_0)\,.
\end{aligned} \tag{E.30}$$

The numerical values $\alpha_0 = 0$, $\alpha_\tau = 1$, $\alpha_p = 10$ are adopted throughout, consistent with the group Laplacian construction of Ref. [194].

The plaquette operator is expressed in terms of gauge-invariant corner operators,

$$C^{ij}_{\mathbf{n},\mu_1,\mu_2} \equiv \sum_a \zeta^{(i)}_{a;\mathbf{n},\mu_1} \zeta^{(j)\dagger}_{a;\mathbf{n},\mu_2} \quad (a \to \bar{a} \text{ when } \mu = -x,-y)\,, \tag{E.31}$$

giving the magnetic Hamiltonian in the dressed-site basis:

$$H_B = -\frac{1}{ag^2} \sum_{\mathbf{n};i,j,k,l} C^{li}_{\mathbf{n},x,y}\, C^{ij}_{\mathbf{n}+\mathbf{e}_x,y,-x}\, C^{jk}_{\mathbf{n}+\mathbf{e}_x+\mathbf{e}_y,-x,-y}\, C^{kl}_{\mathbf{n}+\mathbf{e}_y,-y,x}\,. \tag{E.32}$$

The non-Abelian Gauss's law is automatically satisfied within each dressed site. The additional Abelian selection rules — enforcing that the left and right semilinks of each physical link are in the same irrep — are imposed as penalty terms in the effective Hamiltonian:

$$\begin{aligned} (\hat{D}^R_{\mathbf{n}} \hat{D}^L_{\mathbf{n}+\mathbf{e}_x} - 1)\,|\psi_{\text{phys}}\rangle = 0 \quad \forall\, \mathbf{n}\,, \\ (\hat{D}^C_{\mathbf{n}} \hat{D}^{C'}_{\mathbf{n}+\mathbf{e}_y} - 1)\,|\psi_{\text{phys}}\rangle = 0 \quad \forall\, \mathbf{n}\,. \end{aligned} \tag{E.33}$$

The diagonal operators $\hat{D}^{L,R,C,C'}$ are defined in the dressed basis such that Eqs. (E.33) are satisfied only when the two matched semilinks are in the same irrep:

$$\begin{aligned} \hat{D}^L &= \text{diag}(1,2,2,1,2,2,2,-1,1,-1,-1)\,, \\ \hat{D}^R &= \text{diag}(1,\tfrac{1}{2},1,\tfrac{1}{2},\tfrac{1}{2},\tfrac{1}{2},-1,\tfrac{1}{2},-1,-1,1)\,, \\ \hat{D}^C &= \text{diag}(1,1,2,2,2,-1,2,2,-1,1,-1)\,, \\ \hat{D}^{C'} &= \text{diag}(1,1,\tfrac{1}{2},\tfrac{1}{2},\tfrac{1}{2},-1,\tfrac{1}{2},\tfrac{1}{2},-1,1,-1)\,, \end{aligned} \tag{E.34}$$

where the values $(1,-1,2)$ are assigned to the trivial, parity, and fundamental representations in $\hat{D}^L$ and $\hat{D}^C$, respectively.

E.3.3 Fermionic Parity of the Rishons

An important consequence of the absence of a group center in $D_{N_{\text{odd}}}$ is that the gauge group does not preserve local $\mathbb{Z}_2$ parity. This implies that, when coupled to fermionic matter, $D_{N_{\text{odd}}}$ LGTs cannot be defermionised: no local Hamiltonian without fermionic degrees of freedom can be obtained, independently of the dimensionality. In $D_{N_{\text{even}}}$, by contrast, the rishons are fermions and can be combined with fermionic matter to form singlets behaving as effective bosons.

In the $D_{N_{\text{even}}}$ case, the parity operator on a semilink is

$$P = \text{diag}(1,1,\ldots,-\mathbb{1}_2)\,, \tag{E.35}$$

where all one-dimensional representations carry even parity and all two-dimensional representations carry odd parity.

For D_4, four rishon modes are required (see Fig. E.2(e)), with operators:

$$\zeta_a^{(1)} = 2^{-1/4}\begin{pmatrix} 0 & & & & \delta_{ar} & \delta_{ag} \\ & 0 & & & & \\ & & 0 & & & \\ & & & 0 & & \\ \delta_{ag} & & & & 0 & \\ \delta_{ar} & & & & & 0 \end{pmatrix}, \quad \zeta_a^{(2)} = 2^{-1/4}\begin{pmatrix} 0 & & & & \delta_{ar} & -\delta_{ag} \\ & 0 & & & & \\ & & 0 & & & \\ & & & 0 & & \\ \delta_{ag} & & & & 0 & \\ -\delta_{ar} & & & & & 0 \end{pmatrix},$$

$$\zeta_a^{(3)} = 2^{-1/4}\begin{pmatrix} 0 & & & & & \\ & 0 & & & & \\ & & 0 & & \delta_{ag} & -\delta_{ar} \\ & & & 0 & & \\ & & \delta_{ar} & & 0 & \\ & & -\delta_{ag} & & & 0 \end{pmatrix}, \quad \zeta_a^{(4)} = 2^{-1/4}\begin{pmatrix} 0 & & & & & \\ & 0 & & & & \\ & & 0 & & & \\ & & & 0 & \delta_{ag} & \delta_{ar} \\ & & & \delta_{ar} & 0 & \\ & & & \delta_{ag} & & 0 \end{pmatrix}. \tag{E.36}$$

These rishons anti-commute with P on the same semilink and commute on different semilinks. Fermionic rishon operators are defined by attaching a Jordan–Wigner-like parity string:

$$\zeta^{(i)}_{F,a;\mathbf{n},\mu} = \left(\bigotimes_{(\mathbf{n}',\mu')\in\text{path}} P_{\mathbf{n}',\mu'}\right) \otimes \tilde{\zeta}^{(i)}_{a;\mathbf{n},\mu} \otimes \mathbb{1}, \tag{E.37}$$

where

$$\tilde{\zeta}^{(i)}_{a;\mathbf{n},\mu} = \begin{cases} \zeta^{(i)}_{a;\mathbf{n},\mu} P_{\mathbf{n},\mu} & \text{on left semilinks,} \\ \zeta^{(i)}_{a;\mathbf{n},\mu} & \text{on right semilinks,} \end{cases} \tag{E.38}$$

and the path follows the ordering illustrated in Fig. E.2(a). The fermionic rishons satisfy $\{\tilde{\zeta}^{(i)}_{F,a;\mathbf{n},\mu}, \tilde{\zeta}^{(i)}_{F,a;\mathbf{n}',\mu'}\} = 0$ for $(\mathbf{n},\mu) \neq (\mathbf{n}',\mu')$, confirming their fermionic nature. They can directly replace the local rishons in Eq. (E.17):

$$U_{ab;\mathbf{n},\mathbf{n}+\mathbf{e}_\mu} = \sum_i \zeta^{(i)}_{F,a;\mathbf{n},\mu}\, \zeta^{(i)\dagger}_{F,\bar{b};\mathbf{n}+\mathbf{e}_\mu,-\mu}\,. \tag{E.39}$$

With dynamical fermionic matter, the hopping term becomes

$$\psi^\dagger_{a,\mathbf{n}} U_{ab;\mathbf{n},\mathbf{n}+\mathbf{e}_\mu} \psi_{b,\mathbf{n}+\mathbf{e}_\mu} + \text{h.c.} = \sum_i \psi^\dagger_{a,\mathbf{n}} \zeta^{(i)}_{F,a;\mathbf{n},\mu} \zeta^{(i)\dagger}_{F,\bar{b};\mathbf{n}+\mathbf{e}_\mu,-\mu} \psi_{b,\mathbf{n}+\mathbf{e}_\mu} + \text{H.c.} \tag{E.40}$$

The right-hand side commutes with the fermionic parity on each dressed site and can therefore be defermionised.

In $D_{N_{\text{odd}}}$, the screening fusion rule prevents assigning odd parity to the fundamental representation. Regardless of the parity assignment to $\tilde{\tau}$, the product on the left-hand side of $j \otimes \tau \ni j$ is of even parity, forcing the fundamental representation to carry even parity and rendering some rishons bosonic. Full defermionisation as in the $D_{N_{\text{even}}}$ case is therefore not possible.

E.4 Weak Confinement in the $\mathbb{Z}_2$ Ladder

To demonstrate that the quasi-1D ladder geometry induces weak confinement even in the weak-coupling regime, the case of a $\mathbb{Z}_2$ LGT ladder is analysed [**Nyhegn2021**, **Pradhan˙PRB2024**].

The full Hamiltonian is

$$H_{\mathbb{Z}_2} = -g^2 \sum_l \sum_{s=\uparrow,\downarrow,r} Z_{l,s} - \frac{1}{g^2} \sum_l X_{l,\downarrow} X_{l+1,r} X_{l,\uparrow} X_{l,r} \,, \tag{E.41}$$

where l labels plaquettes, $\uparrow$ / $\downarrow$ the upper/lower legs, and r the left rung. On the dual lattice, gauge symmetry with background field $Z_{\text{bg}} = +1$ imposes $Z_{l,r} = Z_{l-1} Z_l$ (with $Z_l = Z_{l,\uparrow}$) and $Z_{l,\downarrow} = \pm Z_l$ depending on whether static charges are present ($q = 1$) or absent ($q = 0$). The plaquette operator acts as a Pauli X_l on the dual lattice, giving

$$\tilde{H}_{\mathbb{Z}_2} = -g^2 \sum_{l=1}^{R-1} Z_l Z_{l+1} - 2g^2 \cdot 2(1-q) \sum_{l=1}^{R-1} Z_l + g^2 (1-q)(Z_1 + Z_R) - \frac{1}{g^2} \sum_l X_l \,, \tag{E.42}$$

where R is the number of plaquettes.

In the weak-coupling regime, the ground-state energy is computed in perturbation theory over the polarised state $\bigotimes_l |+\rangle_l$:

$$E_{q=0} = -\frac{R}{g^2} - g^6 \frac{9R+19}{4} + o(g^6) \,, \tag{E.43}$$

$$E_{q=1} = -\frac{R}{g^2} - g^6 \frac{R+3}{4} + o(g^6) \,. \tag{E.44}$$

The energy difference is therefore $E_{q=1} - E_{q=0} = 2g^6(R+2)$, proportional to the charge separation R, indicating weak confinement in the magnetic regime. This contrasts the fully 2D scenario, where $\mathbb{Z}_N$ LGTs undergo a deconfinement phase transition at finite coupling g_c, with the string tension vanishing exponentially for $g < g_c$ [253]. The weak confinement on the ladder is therefore of purely geometric origin, arising from the effective longitudinal field induced by the boundaries. The fact that the dihedral groups D_3 and D_4 display exactly the same g^6 scaling of the string tension at weak coupling (see Section 6.1.4) confirms that this behaviour is a perturbative, geometry-induced effect common to all models studied.

Appendix F

Quantum Resources in Non-Abelian Lattice Gauge Theories

This appendix supplements Section 6.2, which compares nonstabilizerness, multipartite entanglement, and fermionic non-Gaussianity across $\mathbb{Z}_N$, D_3, and SU(2) pure-gauge theories on a plaquette ladder. Section F.1 gives a detailed account of the Kogut–Susskind Hamiltonian in the electric and magnetic limits, explaining how the strong- and weak-coupling ground states are constructed for each gauge group and how the plaquette operator generates entanglement through the group's Clebsch–Gordan coefficients. The remaining sections report the definitions and computational details of the three resource measures used, additional numerical data for the SU(2) and $\mathbb{Z}_N$ cases, and further discussion of the encoding-dependence of the results.

F.1 Details on the Pure-Gauge Flux Ladders

The Kogut–Susskind Hamiltonian of Eq. (6.8) contains two competing terms that favour states with opposite properties. The electric contribution $E_l^2 = \sum_J \alpha^J P_l^J$ is a single-body term diagonal in the irrep basis. The energies α^J are typically derived from the group Laplacian — a natural construction for Lie groups but less immediate for discrete ones [194]. The lowest-energy irrep is always the identity $J = e$, which completely fixes the strong-coupling ground state:

$$|GS(g \to \infty)\rangle = \bigotimes_{\text{links}} |e\rangle_l \ . \tag{F.1}$$

In the group-element basis, this reads

$$|GS(g \to \infty)\rangle = \bigotimes_{\text{links}} \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle_l \ , \tag{F.2}$$

a product state with vanishing entanglement and vanishing SRE (it is an eigenstate of a generalized clock operator, hence a stabilizer state).

The plaquette operator $B_p = \mathrm{Tr}(U_1 U_2 U_3^\dagger U_4^\dagger) + \text{H.c.}$ is diagonal in the group-element basis and acts on the irrep basis according to the Clebsch–Gordan coefficients of the group. It

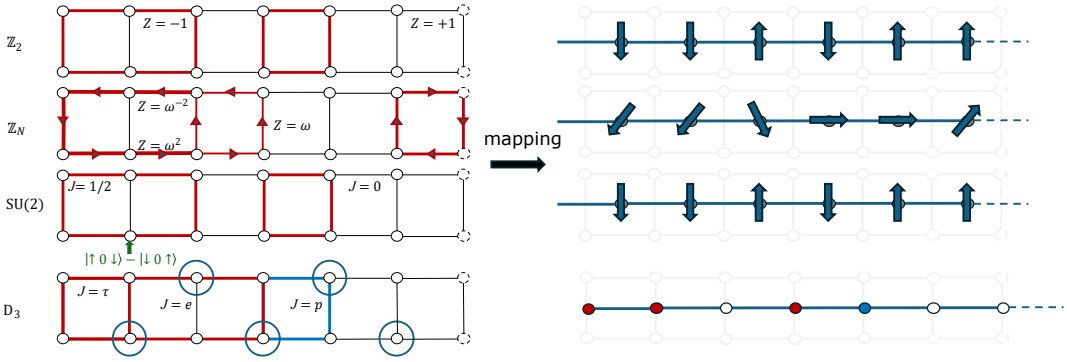

Figure F.1: Examples of gauge-invariant configurations in the representation (electric) basis for the different models. The local basis for $\mathbb{Z}_2$ and SU(2) is $\{|\uparrow\rangle, |\downarrow\rangle\}$ (eigenstates of Pauli-Z). For $\mathbb{Z}_N$, the eigenstates of the generalized phase operator Z are shown as Nth roots of unity on the unit circle. For D_3, circles identify the combinations of irreps — identity e, parity p, fundamental τ — that form group singlets; the internal structure of 2D representations is not depicted.

is responsible for four-body interactions and generates complex entanglement in the ground state. Plaquette excitations correspond to magnetic vortices; the magnetic ground state is associated with vanishing flux. The resulting electric excitations must form group singlets at every vertex, constrained by the gauge symmetry — this generates the large entanglement at weak coupling. The celebrated example is $\mathbb{Z}_2$, whose $g \to 0$ ground state on a square lattice is the toric code, a uniform superposition of all Wilson loops:

$$|GS(g \to 0)\rangle = \frac{1}{\sqrt{\mathcal{N}}} \sum_{\Gamma} \mathcal{W}_\Gamma \, |GS(g \to \infty)\rangle \,, \tag{F.3}$$

where $\mathcal{N}$ is the number of closed paths Γ and $\mathcal{W}_\Gamma = \bigotimes_{l\in\Gamma} X_l$ is the corresponding Wilson loop. Analogous constructions hold for other LGTs; the structure becomes richer as it reflects the more complex fusion rules of the group. Figure F.1 illustrates examples of gauge-invariant states for all models studied.

F.1.1 SU(2)

The pure-gauge Kogut–Susskind Hamiltonian for SU(2), with colour index a summed implicitly, reads [84]

$$H = \frac{g^2}{2} \sum_{\text{links}} (E_l^a)^2 - \frac{2}{g^2} \sum_{\text{plaquettes}} B_p \,, \tag{F.4}$$

where $B_p = \mathrm{Tr}(U_1 U_2 U_3^\dagger U_4^\dagger + \text{H.c.})$ and U_l is the parallel transporter in the $J = 1/2$ representation. The gauge-field operators fulfil

$$\begin{aligned} [L_l^a, U_{l'}^{\alpha\beta}] &= -\delta_{l,l'} \frac{\sigma^a_{\alpha\beta}}{2} U_l^{\beta\gamma} \,, \\ [R_l^a, U_{l'}^{\alpha\beta}] &= \delta_{l,l'} U_l^{\alpha\beta} \frac{\sigma^a_{\beta\gamma}}{2} \,, \end{aligned} \tag{F.5}$$

k, N	2	3	4	5	6
0		$\{\|1\rangle, \|2\rangle\}$	$\{\|1\rangle, \|3\rangle\}$	$\{\|1\rangle, \|4\rangle\}, \{\|2\rangle, \|3\rangle\}$	$\{\|1\rangle, \|5\rangle\}, \{\|2\rangle, \|4\rangle\}$
1	$\{\|0\rangle, \|1\rangle\}$	$\{\|0\rangle, \|2\rangle\}$	$\{\|1\rangle, \|2\rangle\}, \{\|0\rangle, \|3\rangle\}$	$\{\|0\rangle, \|4\rangle\}, \{\|1\rangle, \|3\rangle\}$	$\{\|1\rangle, \|4\rangle\}, \{\|2\rangle, \|3\rangle\}$
2		$\{\|0\rangle, \|1\rangle\}$	$\{\|0\rangle, \|1\rangle, \|2\rangle, \|3\rangle\}$	$\{\|1\rangle, \|2\rangle\}, \{\|0\rangle, \|3\rangle\}$	$\{\|0\rangle, \|4\rangle\}, \{\|1\rangle, \|3\rangle\}$

Table F.1: Degenerate states for each pair (N, k). Red entries mark pairs that dominantly contribute to the ground state at strong coupling.

with $\mathbf{L}_l$ and $\mathbf{R}_l$ the left and right generators of SU(2) gauge transformations, entering the electric field as $|\mathbf{E}_l|^2 = |\mathbf{R}_l|^2 = |\mathbf{L}_l|^2$. Gauss's law reads $[H, G_{\mathbf{v}}] = 0$ with $G_{\mathbf{v}} = \sum_i \mathbf{L}_i + \sum_o \mathbf{R}_o$.

Truncating the irreps at $J = 0$ and $J = 1/2$ and following the mapping of Refs. [239, 240], the Hamiltonian reduces to the interacting spin-$\frac{1}{2}$ chain:

$$H_{\mathrm{SU(2)}} = \sum_{i=0}^{N-1} \left[h_{zz} Z_i Z_{i+1} + h_z Z_i + h_x (1-3Z_{i-1}) X_i (1-3Z_{i+1}) \right], \tag{F.6}$$

where Z_i counts the electric energy in the hardcore-gluon approximation and X_i is the plaquette operator switching between $J = 0$ and $J = 1/2$. The prefactors $(1-3Z_{i\pm1})$ recover the correct SU(2) Clebsch–Gordan coefficients [19, 240]. The couplings are $h_{zz} = -3g^2/16$, $h_z = -2h_{zz}$, and $h_x = -1/(8g^2)$.

F.1.2 $\mathbb{Z}_N$

$\mathbb{Z}_N$ LGTs on the plaquette ladder with periodic boundary conditions map onto one-dimensional clock models [209]. Labelling the upper, lower, and left-rung links of the i-th plaquette as 1, 2, and 0, the plaquette operator reads $B_i = U_{i,2} U_{i+1,0} U_{i,1}^\dagger U_{i,0}^\dagger$, with $U_l |e_{k,l}\rangle = |e_{k+1,l}\rangle$ and $E_l |e_{k,l}\rangle = \omega^k |e_{k,l}\rangle$. The commutation relations

$$\begin{aligned} B_i E_{i,2} &= \omega^{-1} E_{i,2} B_i, \quad & B_i E_{i,1} &= \omega E_{i,1} B_i, \\ E_{i,0} B_i &= \omega^{-1} B_i E_{i,0}, \quad & E_{i,0} B_{i-1} &= \omega B_{i-1} E_{i,0} \end{aligned} \tag{F.7}$$

are preserved by the mapping $B_i \to X_i$, $E_{i,0} \to Z_{i-1} Z_i^\dagger$. Enforcing the Gauss constraints then yields $E_{i,1} \to Z_i^\dagger$, $E_{i,2} \to \omega^k Z_i$, giving the 1D Hamiltonian of Eq. (6.10).

In the magnetic limit ($g \to 0$), the plaquette operators reduce to a local transverse field, giving a $\mathbb{Z}_N$ paramagnet polarised in the X direction — a stabilizer state. In the electric limit ($g \to \infty$), degeneracies can appear depending on N and k. The longitudinal term has eigenvalue

$$E_a = -4 \cos\left(\tfrac{\pi k}{N}\right) \cos\left(\tfrac{2\pi}{N}\left(\tfrac{k}{2} + a\right)\right) . \tag{F.8}$$

Degenerate pairs $\{a, b\}$ satisfy $a + b = N - k$ and may constitute the ground state. Table F.1 lists all degeneracies for $N = 2, \ldots, 6$ and $k = 0, 1, 2$.

F.1.3 D_3

D_3 is the smallest non-Abelian group (the six symmetries of an equilateral triangle). Following Ref. [93], the plaquette ladder is mapped to a 1D chain where each site represents one rung. A gauge transformation at each vertex fixes the left-leg links to the identity in the group-element basis. After gauge-fixing, the plaquette term is simplified from a four-body

to a two-body interaction. The electric energy splits into a rung term and a non-local leg term arising from the gauge fixing:

$$H_E^{\text{rungs}} = -g^2 \sum_i \sum_J \alpha^J \hat{P}_i^J \,, \tag{F.9}$$

$$\hat{P}_i^J = \frac{\dim(J)}{|\mathcal{G}|} \sum_h \chi^J(h)\, \theta_i^L(h) \,, \tag{F.10}$$

where $\chi^J(h)$ is the character of h in irrep J. Using the local symmetries, the leg contributions can be rewritten in terms of rung operators, giving:

$$\begin{aligned} H_{\mathrm{D}_3} = &-g^2 \sum_i \sum_J \alpha^J \hat{P}_i^J - 2g^2 \sum_i \sum_J \alpha^J \sum_h \frac{\dim(J)}{|\mathcal{G}|} \chi^J(h) \prod_{i'<i} \theta_{i'}^L(h) \\ &- \frac{1}{2g^2} \sum_i \Big(\mathrm{Tr}[U_i^\dagger U_{i+1}] + \mathrm{H.c.} \Big) \,. \end{aligned} \tag{F.11}$$

Using $\prod_{i'<i} \frac{\dim(J)}{|\mathcal{G}|} \sum_h \chi^J(h) \theta_{i'}^L(h) = \prod_{i'<i} \hat{P}_{i'}^J$, this reduces to Eq. (6.11) of the main chapter. The results shown use $\alpha_J = (1, 0, 0)$ for the trivial, parity, and fundamental representations, respectively.

F.2 Energy Gaps for Different LGTs

Figure F.2(a) displays the energy gap for the $\mathbb{Z}_2$ and SU(2) models as a function of coupling strength. No significant scaling with system size L is observed in either case, consistent with the absence of a phase transition in the ladder geometry.

Figure F.2(b) shows the gap for the D_3 theory. The system retains a global D_3 gauge symmetry, inducing a $|\mathrm{D}_3| = 6$-fold degeneracy in the ground state at $g^2 = 0$. At finite coupling, the ground state transforms trivially under the global symmetry, and the relevant gap is $E_6 - E_0$ (between the ground state and the first excited state in the same symmetry sector). This gap decreases with system size in the crossover regime but converges to a finite minimum, which occurs around a level crossing between different symmetry sectors. This behavior confirms the smooth crossover between electric and magnetic regimes, without a phase transition in the thermodynamic limit.

F.3 Analytical Computation of SRE for $\mathbb{Z}_N$ LGTs at Strong Coupling

As discussed in Section 6.2.4, specific combinations of N and k induce a degenerate two-component ground state at strong coupling (see Table F.1). In this case, the strong-coupling ground state for L qudits is

$$|\psi(g \to \infty)\rangle = \frac{|a\rangle^{\otimes L} + |b\rangle^{\otimes L}}{\sqrt{2}} \,. \tag{F.12}$$

The SRE of this state is computed analytically by evaluating $\sum_{P_i} |\langle\psi|P_i|\psi\rangle|^4$. For a

single Pauli string $P_i = \bigotimes_{j=1}^{L} X_j^{r_j} Z_j^{s_j}$, the expectation value is

$$\langle\psi|P_i|\psi\rangle = \frac{1}{2}\Bigg[\prod_{j=1}^{L}\delta_{0,r_j}\left(\omega^{a\sum s_j}+\omega^{b\sum s_j}\right) + \prod_{j=1}^{L}\delta_{b,r_j\oplus a}\,\omega^{a\sum s_j}+\prod_{j=1}^{L}\delta_{a,r_j\oplus b}\,\omega^{b\sum s_j}\Bigg], \tag{F.13}$$

where $Z^s|a\rangle = \omega^{as}|a\rangle$, $\langle b|X^r|a\rangle = \delta_{b,r\oplus a}$, and $\oplus$ denotes addition modulo N. Since $a = b\oplus r$ and $b = a\oplus r$ cannot both hold unless $2r \equiv 0 \pmod N$ (which would require the degenerate pair to differ by $N/2$, a configuration that does not correspond to the ground state), the two off-diagonal terms contribute independently.

Introducing the auxiliary variable $x = \sum_{j=1}^{L} s_j$ and its multiplicity $c_x = \sum_{k=0}^{x/d}(-1)^k\binom{L}{k}\binom{x-kd+L-1}{L-1}$ with $d = b-a$, the sum over Pauli strings reduces to:

$$\sum_{P_i}\left|\langle\psi|P_i|\psi\rangle\right|^4 = \sum_{x=0}^{L(N-1)} c_x\left[\cos^4\left(\tfrac{\pi dx}{N}\right)+1\right]. \tag{F.14}$$

Inserting Eq. (F.14) into the definition of $\mathcal{M}_2$ gives, for all pairs (N,k) listed in red in Table F.1 (i.e., $N = 2,\ldots,6$ with $k = 1$):

$$\mathcal{M}_2 = 0.32\,. \tag{F.15}$$

This value is independent of N and L for the dominant degenerate pairs, confirming the analytical prediction discussed in Section 6.2.4.

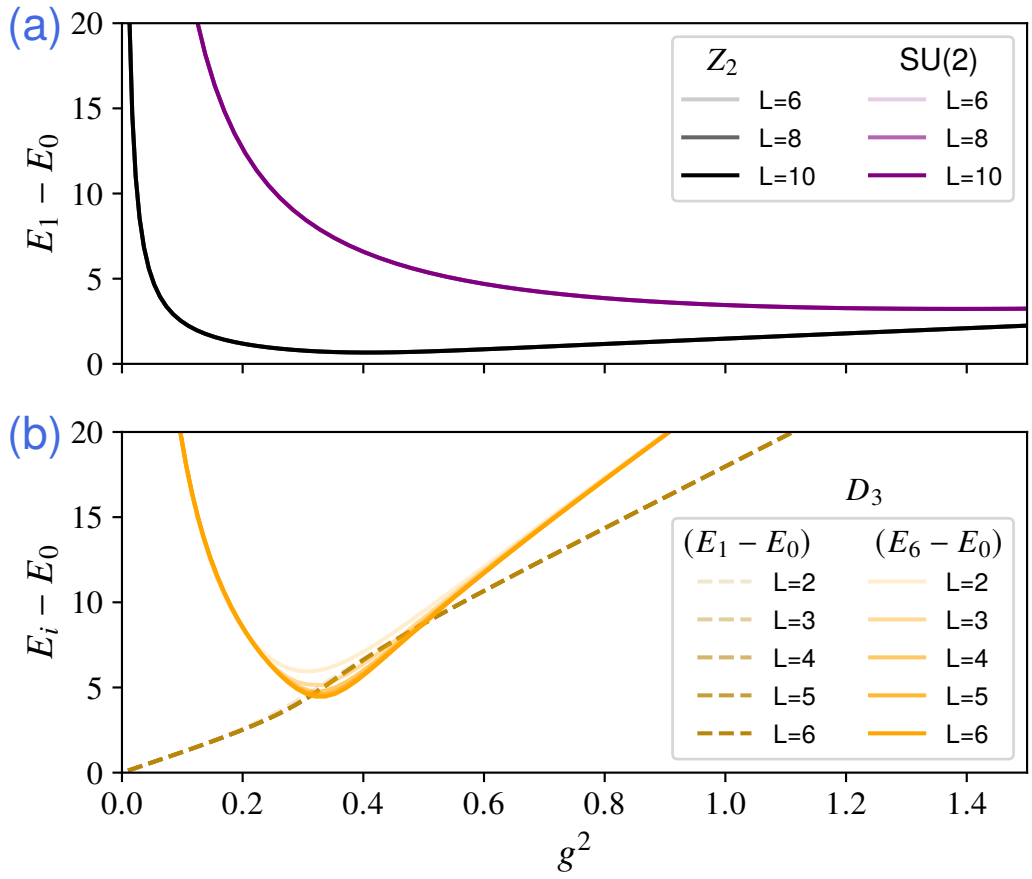

Figure F.2: Energy gaps vs. coupling g^2 for $\mathbb{Z}_2$ and SU(2) (a), and D_3 (b). Neither displays significant finite-size scaling, confirming the absence of a phase transition in the ladder geometry. Panel (b) highlights the level crossing between different symmetry sectors of the global D_3 gauge transformation. Solid lines: gap to the first excited state in the same symmetry sector as the ground state. Dashed lines: gap to the manifold that transforms nontrivially under the global D_3 symmetry and becomes degenerate at $g^2 = 0$.

Bibliography

[1] Richard P Feynman. "Simulating physics with computers". In: *Feynman and computation*. cRc Press, 2018, pp. 133–153.

[2] David Deutsch. "Quantum theory, the Church–Turing principle and the universal quantum computer". In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 400.1818 (1985), pp. 97–117.

[3] Peter W Shor. "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer". In: *SIAM J. Comput.* 41.2 (1999), pp. 303–332. DOI: 10.1137/S0097539795293172. URL: https://doi.org/10.1137/S0097539795293172.

[4] Lov K Grover. "A fast quantum mechanical algorithm for database search". In: *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing* (1996), pp. 212–219. URL: https://dl.acm.org/doi/pdf/10.1145/237814.237866.

[5] Frank Arute et al. "Quantum supremacy using a programmable superconducting processor". In: *Nature* 574.7779 (2019), pp. 505–510. DOI: 10.1038/s41586-019-1666-5.

[6] Martin Ringbauer et al. "A universal qudit quantum processor with trapped ions". In: *Nature Physics* 18.9 (2022), pp. 1053–1057. DOI: https://doi.org/10.1038/s41567-022-01658-0. URL: https://www.nature.com/articles/s41567-022-01658-0#article-info.

[7] C.L. Edmunds et al. "Symmetry-Protected Topological Haldane Phase on a Qudit Quantum Processor". In: *PRX Quantum* 6 (2 June 2025), p. 020349. DOI: 10.1103/PRXQuantum.6.020349. URL: https://link.aps.org/doi/10.1103/PRXQuantum.6.020349.

[8] Dolev Bluvstein et al. "Logical quantum processor based on reconfigurable atom arrays". In: *Nature* 626.7997 (2024), pp. 58–65. URL: https://www.nature.com/articles/s41586-023-06927-3.

[9] Yulin Chi et al. "A programmable qudit-based quantum processor". In: *Nature Communications* 13.1 (Mar. 2022), p. 1166. ISSN: 2041-1723. DOI: 10.1038/s41467-022-28767-x. URL: https://doi.org/10.1038/s41467-022-28767-x.

[10] John Preskill. "Quantum computing in the NISQ era and beyond". In: *Quantum* 2 (2018), p. 79. URL: https://quantum-journal.org/papers/q-2018-08-06-79/.

[11] Laurin E Fischer et al. "Universal qudit gate synthesis for transmons". In: *PRX Quantum* 4.3 (2023), p. 030327. DOI: 10.1103/PRXQuantum.4.030327. URL: https://journals.aps.org/prxquantum/abstract/10.1103/PRXQuantum.4.030327.

[12] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. "A Quantum Approximate Optimization Algorithm". In: (2014). DOI: 10.48550/arXiv.1411.4028.

[13] M. Cerezo et al. "Variational Quantum Algorithms". In: *Nat. Rev. Phys.* 3.9 (Sept. 2021), pp. 625–644. ISSN: 2522-5820. DOI: 10.1038/s42254-021-00348-9. URL: https://www.nature.com/articles/s42254-021-00348-9.

[14] Kishor Bharti et al. "Noisy Intermediate-Scale Quantum Algorithms". In: *Rev. Mod. Phys.* 94.1 (Feb. 2022), p. 015004. DOI: 10.1103/RevModPhys.94.015004. URL: https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.94.015004.

[15] Jarrod R. McClean et al. "Barren Plateaus in Quantum Neural Network Training Landscapes". In: *Nature Communications* 9.1 (Nov. 2018), p. 4812. ISSN: 2041-1723. DOI: 10.1038/s41467-018-07090-4. URL: https://www.nature.com/articles/s41467-018-07090-4.

[16] Martín Larocca et al. "Barren plateaus in variational quantum computing". In: *Nature Reviews Physics* 7.4 (Apr. 2025), pp. 174–189. ISSN: 2522-5820. DOI: 10.1038/s42254-025-00813-9. URL: https://doi.org/10.1038/s42254-025-00813-9.

[17] John B. Kogut. "The lattice gauge theory approach to quantum chromodynamics". In: *Rev. Mod. Phys.* 55 (3 July 1983), pp. 775–836. DOI: 10.1103/RevModPhys.55.775. URL: https://link.aps.org/doi/10.1103/RevModPhys.55.775.

[18] Alexei Kitaev. "Topological quantum codes and anyons". In: *Proceedings of Symposia in Applied Mathematics*. Vol. 58. 2002, pp. 267–272. URL: https://userpages.cs.umbc.edu/lomonaco/ams/lecturenotes/Kitaev.pdf.

[19] Erez Zohar and Michele Burrello. "Formulation of lattice gauge theories for quantum simulations". In: *Phys. Rev. D* 91 (5 Mar. 2015), p. 054506. DOI: 10.1103/PhysRevD.91.054506. URL: https://link.aps.org/doi/10.1103/PhysRevD.91.054506.

[20] Marcello Dalmonte and Simone Montangero. "Lattice gauge theory simulations in the quantum information era". In: *Contemporary Physics* 57.3 (2016), pp. 388–412. URL: https://doi.org/10.1080/00107514.2016.1151199.

[21] Christian W. Bauer et al. "Quantum Simulation for High-Energy Physics". In: *PRX Quantum* 4 (2 May 2023), p. 027001. DOI: 10.1103/PRXQuantum.4.027001. URL: https://link.aps.org/doi/10.1103/PRXQuantum.4.027001.

[22] Alberto Di Meglio et al. "Quantum Computing for High-Energy Physics: State of the Art and Challenges". In: *PRX Quantum* 5 (3 Aug. 2024), p. 037001. DOI: 10.1103/PRXQuantum.5.037001. URL: https://link.aps.org/doi/10.1103/PRXQuantum.5.037001.

[23] Giuseppe Calajó et al. "Digital Quantum Simulation of a (1+1)D SU(2) Lattice Gauge Theory with Ion Qudits". In: *PRX Quantum* 5 (4 Oct. 2024), p. 040309. DOI: 10.1103/PRXQuantum.5.040309. URL: https://link.aps.org/doi/10.1103/PRXQuantum.5.040309.

[24] Alberto Bottarelli, Sebastian Schmitt, and Philipp Hauke. "Inequality constraints in variational quantum circuits with qudits". In: *Phys. Rev. Res.* 7 (3 Aug. 2025), p. 033202. DOI: 10.1103/3l96-41xf. URL: https://link.aps.org/doi/10.1103/3l96-41xf.

[25] Alberto Bottarelli et al. "Symmetry-enhanced counterdiabatic quantum algorithm for qudits". In: *Phys. Rev. Res.* 7 (4 Oct. 2025), p. 043030. DOI: 10.1103/6ldg-3w1f. URL: https://link.aps.org/doi/10.1103/6ldg-3w1f.

[26] Sebastian Schmitt et al. *Experimental demonstration of the absence of noise-induced barren plateaus using information content landscape analysis*. 2026. arXiv: 2602.22851 [quant-ph]. URL: https://arxiv.org/abs/2602.22851.

[27] Chiara Capecci et al. "Role of Nonstabilizerness in Quantum Optimization". In: *arXiv:2505.17185* (2025). URL: https://arxiv.org/abs/2505.17185.

[28] Pavel P. Popov et al. *When the center matters: color screening and gluelumps in dihedral lattice gauge theories*. 2025. arXiv: 2512.00152 [hep-lat]. URL: https://arxiv.org/abs/2512.00152.

[29] Gopal Chandra Santra et al. "Quantum Resources in Non-Abelian Lattice Gauge Theories: Nonstabilizerness, Multipartite Entanglement, and Fermionic Non-Gaussianity". In: *arXiv preprint arXiv:2510.07385* (2025).

[30] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010.

[31] Vladimir Vargas-Calderón et al. "Many-Qudit Representation for the Travelling Salesman Problem Optimisation". In: *Journal of the Physical Society of Japan* 90.11 (2021), p. 114002. DOI: 10.7566/JPSJ.90.114002. URL: https://doi.org/10.7566/JPSJ.90.114002.

[32] Mohammad HS Amin, Neil G Dickson, and Peter Smith. "Adiabatic quantum optimization with qudits". In: *Quantum Inf. Proccess.* 12 (2013), pp. 1819–1829. DOI: 10.1007/s11128-012-0480-x.

[33] Sebastián Roca-Jerat, Juan Román-Roche, and David Zueco. "Qudit machine learning". In: *Machine Learning: Science and Technology* 5.1 (2024), p. 015057. URL: https://iopscience.iop.org/article/10.1088/2632-2153/ad360d/meta.

[34] Zhe Meng et al. "Experimental realization of high-dimensional quantum gates with ultrahigh fidelity and efficiency". In: *Phys. Rev. A* 109 (2 Feb. 2024), p. 022612. DOI: 10.1103/PhysRevA.109.022612. URL: https://link.aps.org/doi/10.1103/PhysRevA.109.022612.

[35] Noah Goss et al. "High-fidelity qutrit entangling gates for superconducting circuits". In: *Nature communications* 13.1 (2022), p. 7481.

[36] Yuchen Wang et al. "Qudits and high-dimensional quantum computing". In: *Front. Phys.* 8 (2020), p. 589504. DOI: 10.3389/fphy.2020.589504.

[37] Sergey Bravyi and Alexei Kitaev. "Universal quantum computation with ideal Clifford gates and noisy ancillas". In: *Phys. Rev. A* 71 (2 Feb. 2005), p. 022316. DOI: 10.1103/PhysRevA.71.022316. URL: https://link.aps.org/doi/10.1103/PhysRevA.71.022316.

[38] Naga Dileep Varikuti, Soumik Bandyopadhyay, and Philipp Hauke. "Impact of Clifford operations on non-stabilizing power and quantum chaos". In: *arXiv:2505.14793* (2025). URL: https://arxiv.org/abs/2505.14793.

[39] Michael Beverland et al. "Lower bounds on the non-Clifford resources for quantum computations". In: *Quantum Science and Technology* 5.3 (2020), p. 035009. URL: https://iopscience.iop.org/article/10.1088/2058-9565/ab8963/meta?casa_token=J4VJf4bJu-8AAAAA:W8SlvH7BTutwCfLu8Iris9AonrTvDD8-C69UAB_xBSq2bli_GB85FOR2_N5z4y5dLIW7GNaLK3djVCOEk_P2ikLBFOqnMg&casa_token=k-P2s554PtkAAAAA:b3gFLU3bXk3O-0B91RlHZnG0D6rxLfXalY-YR59K_TYCdxUmyFVs9-g0wbfoIxLYZmXMGoxsLg4aQHtewOBM4teYqtdNKA.

[40] Daniel Gottesman. "Stabilizer codes and quantum error correction". In: *arXiv: quant-ph/9705052* (1997). URL: https://arxiv.org/abs/quant-ph/9705052.

[41] Michael Meth et al. "Simulating two-dimensional lattice gauge theories on a qudit quantum computer". In: *Nature Physics* 21.4 (2025), pp. 570–576.

[42] Gonzalo Camacho et al. "Observing dynamical localization on a trapped-ion qudit quantum processor". In: *arXiv preprint arXiv:2412.13141* (2024).

[43] Eric Chitambar and Gilad Gour. "Quantum resource theories". In: *Rev. Mod. Phys.* 91 (2 Apr. 2019), p. 025001. DOI: 10.1103/RevModPhys.91.025001. URL: https://link.aps.org/doi/10.1103/RevModPhys.91.025001.

[44] T. Baumgratz, M. Cramer, and M. B. Plenio. "Quantifying Coherence". In: *Phys. Rev. Lett.* 113 (14 Sept. 2014), p. 140401. DOI: 10.1103/PhysRevLett.113.140401. URL: https://link.aps.org/doi/10.1103/PhysRevLett.113.140401.

[45] Victor Veitch et al. "The resource theory of stabilizer quantum computation". In: *New Journal of Physics* 16.1 (Jan. 2014), p. 013009. DOI: 10.1088/1367-2630/16/1/013009. URL: https://dx.doi.org/10.1088/1367-2630/16/1/013009.

[46] Ryszard Horodecki et al. "Quantum entanglement". In: *Rev. Mod. Phys.* 81 (2 June 2009), pp. 865–942. DOI: 10.1103/RevModPhys.81.865. URL: https://link.aps.org/doi/10.1103/RevModPhys.81.865.

[47] Gopal Chandra Santra et al. "Genuine multipartite entanglement in quantum optimization". In: *Phys. Rev. A* 111 (2 Feb. 2025), p. 022434. DOI: 10.1103/PhysRevA.111.022434. URL: https://link.aps.org/doi/10.1103/PhysRevA.111.022434.

[48] Daniel Gottesman. "The Heisenberg representation of quantum computers". In: *arXiv preprint quant-ph/9807006* (1998). URL: https://arxiv.org/abs/quant-ph/9807006.

[49] Francisco Soto and Pierre Claverie. "When is the Wigner function of multidimensional systems nonnegative?" In: *Journal of Mathematical Physics* 24.1 (1983), pp. 97–100. URL: https://pubs.aip.org/aip/jmp/article-abstract/24/1/97/226088/When-is-the-Wigner-function-of-multidimensional?redirectedFrom=fulltext.

[50] Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hamma. "Stabilizer Rényi Entropy". In: *Phys. Rev. Lett.* 128 (5 Feb. 2022), p. 050402. DOI: 10.1103/PhysRevLett.128.050402. URL: https://link.aps.org/doi/10.1103/PhysRevLett.128.050402.

[51] Tobias Haug and Lorenzo Piroli. "Stabilizer entropies and nonstabilizerness monotones". In: *Quantum* 7 (2023), p. 1092. DOI: 10.22331/q-2023-08-28-1092. URL: https://doi.org/10.22331/q-2023-08-28-1092.

[52] David Gross, Sepehr Nezami, and Michael Walter. "Schur–Weyl Duality for the Clifford Group with Applications: Property Testing, a Robust Hudson Theorem, and de Finetti Representations". In: *Commun. Math. Phys.* 385.3 (2021), pp. 1325–1393. URL: https://doi.org/10.1007/s00220-021-04118-7.

[53] Christopher David Niroula Pradeep andWhite et al. "Phase transition in magic with random quantum circuits". In: *Nature Physics* 20 (2024), pp. 1786–1792. DOI: 10.1038/s41567-024-02637-3. URL: https://doi.org/10.1038/s41567-024-02637-3.

[54] Poetri Sonya Tarabunga et al. "Many-Body Magic Via Pauli-Markov Chains—From Criticality to Gauge Theories". In: *PRX Quantum* 4 (4 Oct. 2023), p. 040317. DOI: 10.1103/PRXQuantum.4.040317. URL: https://link.aps.org/doi/10.1103/PRXQuantum.4.040317.

[55] Piotr Sierant, Jofre Vallès-Muns, and Artur Garcia-Saez. *Computing quantum magic of state vectors*. 2026. arXiv: 2601.07824 [quant-ph]. URL: https://arxiv.org/abs/2601.07824.

[56] Jorge Nocedal and Stephen J. Wright. "Numerical Optimization". In: Springer Series in Operations Research and Financial Engineering (2006).

[57] Gary Kochenberger et al. "The unconstrained binary quadratic programming problem: a survey". In: *Journal of combinatorial optimization* 28.1 (2014), pp. 58–81.

[58] Jules Tilly et al. "The Variational Quantum Eigensolver: A review of methods and best practices". In: *Physics Reports* 986 (2022), pp. 1–128. ISSN: 0370-1573. DOI: https://doi.org/10.1016/j.physrep.2022.08.003. URL: https://www.sciencedirect.com/science/article/pii/S0370157322003118.

[59] Alberto Peruzzo et al. "A Variational Eigenvalue Solver on a Photonic Quantum Processor". In: *Nature Communications* 5.1 (July 2014), p. 4213. ISSN: 2041-1723. DOI: 10.1038/ncomms5213. URL: https://www.nature.com/articles/ncomms5213.

[60] Kostas Blekos et al. "A review on Quantum Approximate Optimization Algorithm and its variants". In: *Physics Reports* 1068 (2024). A review on Quantum Approximate Optimization Algorithm and its variants, pp. 1–66. ISSN: 0370-1573. DOI: https://doi.org/10.1016/j.physrep.2024.03.002. URL: https://www.sciencedirect.com/science/article/pii/S0370157324001078.

[61] Edward Farhi and Aram W Harrow. "Quantum Supremacy through the Quantum Approximate Optimization Algorithm". In: *arXiv:1602.07674* (2019). DOI: 10.48550/arXiv.1602.07674. URL: https://arxiv.org/abs/1602.07674.

[62] Tadashi Kadowaki and Hidetoshi Nishimori. "Quantum annealing in the transverse Ising model". In: *Phys. Rev. E* 58 (5 Nov. 1998), pp. 5355–5363. DOI: 10.1103/PhysRevE.58.5355. URL: https://link.aps.org/doi/10.1103/PhysRevE.58.5355.

[63] Tameem Albash and Daniel A. Lidar. "Adiabatic quantum computation". In: *Rev. Mod. Phys.* 90 (1 Jan. 2018), p. 015002. DOI: 10.1103/RevModPhys.90.015002. URL: https://link.aps.org/doi/10.1103/RevModPhys.90.015002.

[64] Philipp Hauke et al. "Perspectives of quantum annealing: methods and implementations". In: *Reports on Progress in Physics* 83.5 (May 2020), p. 054401. DOI: 10.1088/1361-6633/ab85b8. URL: https://doi.org/10.1088/1361-6633/ab85b8.

[65] Andrew Lucas. "Ising formulations of many NP problems". In: *Frontiers in Physics* 2 (2014), p. 5. ISSN: 2296-424X. DOI: 10.3389/fphy.2014.00005. URL: http://journal.frontiersin.org/article/10.3389/fphy.2014.00005/abstract.

[66] Sheir Yarkoni et al. "Quantum Annealing for Industry Applications: Introduction and Review". In: *Reports on Progress in Physics* 85.10 (2022), p. 104001. URL: https://iopscience.iop.org/article/10.1088/1361-6633/ac8c54.

[67] Valentin Kasper et al. "Universal quantum computation and quantum error correction with ultracold atomic mixtures". In: *Quantum Science and Technology* 7.1 (2022), p. 015008. ISSN: 2058-9565. DOI: 10.1088/2058-9565/ac2d39. URL: https://iopscience.iop.org/article/10.1088/2058-9565/ac2d39.

[68] Paolo Giorda, Paolo Zanardi, and Seth Lloyd. "Universal quantum control in irreducible state-space sectors: Application to bosonic and spin-boson systems". In: *Physical Review A* 68.6 (2003), p. 062320. DOI: 10.1103/PhysRevA.68.062320. URL: https://link.aps.org/doi/10.1103/PhysRevA.68.062320 (visited on 02/15/2023).

[69] David J Gross. "Gauge theory-past, present, and future?" In: *Chinese Journal of Physics* 30.7 (1992), pp. 955–972.

[70] Michael E Peskin. *An Introduction to quantum field theory*. CRC press, 2018.

[71] Mary K. Gaillard, Paul D. Grannis, and Frank J. Sciulli. "The standard model of particle physics". In: *Rev. Mod. Phys.* 71 (2 Mar. 1999), S96–S111. DOI: 10.1103/RevModPhys.71.S96. URL: https://link.aps.org/doi/10.1103/RevModPhys.71.S96.

[72] John B Kogut. "An introduction to lattice gauge theory and spin systems". In: *Rev. Mod. Phys.* 51.4 (4 Oct. 1979), p. 659. DOI: 10.1103/RevModPhys.51.659. URL: https://link.aps.org/doi/10.1103/RevModPhys.51.659.

[73] J. Greensite. "The confinement problem in lattice gauge theory". In: *Progr. Part. Nucl. Phys.* 51.1 (2003), pp. 1–83. ISSN: 0146-6410. DOI: https://doi.org/10.1016/S0146-6410(03)90012-3. URL: https://www.sciencedirect.com/science/article/pii/S0146641003900123.

[74] Laurence G. Yaffe. "Confinement in SU(N) lattice gauge theories". In: *Phys. Rev. D* 21 (6 Mar. 1980), pp. 1574–1590. DOI: 10.1103/PhysRevD.21.1574. URL: https://link.aps.org/doi/10.1103/PhysRevD.21.1574.

[75] Torsten V. Zache, Daniel González-Cuadra, and Peter Zoller. "Fermion-qudit quantum processors for simulating lattice gauge theories with matter". In: *Quantum* 7 (Oct. 2023), p. 1140. ISSN: 2521-327X. DOI: 10.22331/q-2023-10-16-1140. URL: http://dx.doi.org/10.22331/q-2023-10-16-1140.

[76] U.-J. Wiese. "Ultracold quantum gases and lattice systems: quantum simulation of lattice gauge theories". In: *Annalen der Physik* 525.10-11 (2013), pp. 777–796. DOI: https://doi.org/10.1002/andp.201300104. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.201300104. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.201300104.

[77] Minh C. Tran et al. "Faster Digital Quantum Simulation by Symmetry Protection". In: *PRX Quantum* 2 (1 Feb. 2021), p. 010323. DOI: 10.1103/PRXQuantum.2.010323. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.010323.

[78] Masahito Hayashi. "Group Representation Theory". In: *Group Representation for Quantum Theory*. Cham: Springer International Publishing, 2017, pp. 21–67. ISBN: 978-3-319-44906-7. DOI: 10.1007/978-3-319-44906-7_2. URL: https://doi.org/10.1007/978-3-319-44906-7_2.

[79] David J. Gross and Frank Wilczek. "Ultraviolet Behavior of Non-Abelian Gauge Theories". In: *Phys. Rev. Lett.* 30 (26 June 1973), pp. 1343–1346. DOI: 10.1103/PhysRevLett.30.1343. URL: https://link.aps.org/doi/10.1103/PhysRevLett.30.1343.

[80] Jeff Greensite. *An introduction to the confinement problem*. Vol. 821. Springer, 2011.

[81] Pietro Silvi et al. "Finite-density phase diagram of a $(1+1)-d$ non-abelian lattice gauge theory with tensor networks". In: *Quantum* 1 (Apr. 2017), p. 9. ISSN: 2521-327X. DOI: 10.22331/q-2017-04-25-9. URL: https://doi.org/10.22331/q-2017-04-25-9.

[82] Philippe De Forcrand. "Simulating QCD at finite density". In: *arXiv preprint arXiv:1005.0539* (2010).

[83] Massimo D'Elia and Maria-Paola Lombardo. "Finite density QCD via an imaginary chemical potential". In: *Phys. Rev. D* 67 (1 Jan. 2003), p. 014505. DOI: 10.1103/PhysRevD.67.014505. URL: https://link.aps.org/doi/10.1103/PhysRevD.67.014505.

[84] John Kogut and Leonard Susskind. "Hamiltonian formulation of Wilson's lattice gauge theories". In: *Phys. Rev. D* 11 (2 Jan. 1975), pp. 395–408. DOI: 10.1103/PhysRevD.11.395. URL: https://link.aps.org/doi/10.1103/PhysRevD.11.395.

[85] Chetan Nayak et al. "Non-Abelian anyons and topological quantum computation". In: *Rev. Mod. Phys.* 80 (3 Sept. 2008), pp. 1083–1159. DOI: 10.1103/RevModPhys.80.1083. URL: https://link.aps.org/doi/10.1103/RevModPhys.80.1083.

[86] Pierpaolo Fontana, Marc Miranda-Riaza, and Alessio Celi. "Efficient Finite-Resource Formulation of Non-Abelian Lattice Gauge Theories beyond One Dimension". In: *Phys. Rev. X* 15 (3 Sept. 2025), p. 031065. DOI: 10.1103/k9p6-c649. URL: https://link.aps.org/doi/10.1103/k9p6-c649.

[87] Marc Miranda-Riaza, Pierpaolo Fontana, and Alessio Celi. *Renormalized dual basis for scalable simulations of 2+1D compact quantum electrodynamics*. 2025. arXiv: 2510.18594 [quant-ph]. URL: https://arxiv.org/abs/2510.18594.

[88] Pietro Silvi et al. "Tensor network simulation of an SU(3) lattice gauge theory in 1D". In: *Phys. Rev. D* 100 (7 Oct. 2019), p. 074512. DOI: 10.1103/PhysRevD.100.074512. URL: https://link.aps.org/doi/10.1103/PhysRevD.100.074512.

[89] Giovanni Cataldi et al. "Simulating (2+1)D SU(2) Yang-Mills lattice gauge theory at finite density with tensor networks". In: *Phys. Rev. Res.* 6 (3 July 2024), p. 033057. DOI: 10.1103/PhysRevResearch.6.033057. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.6.033057.

[90] Edoardo Ballini et al. "Symmetry verification for noisy quantum simulations of non-Abelian lattice gauge theories". In: *Quantum* 9 (2025), p. 1802. URL: https://quantum-journal.org/papers/q-2025-07-22-1802/.

[91] Erez Zohar et al. "Digital Quantum Simulation of $\mathbb{Z}_2$ Lattice Gauge Theories with Dynamical Fermionic Matter". In: *Phys. Rev. Lett.* 118 (7 Feb. 2017), p. 070501. DOI: 10.1103/PhysRevLett.118.070501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.118.070501.

[92] Luca Lumia et al. "Two-Dimensional $\mathbb{Z}_2$ Lattice Gauge Theory on a Near-Term Quantum Simulator: Variational Quantum Optimization, Confinement, and Topological Order". In: *PRX Quantum* 3 (2 Apr. 2022), p. 020320. DOI: 10.1103/PRXQuantum.3.020320. URL: https://link.aps.org/doi/10.1103/PRXQuantum.3.020320.

[93] Morten I. K. Munk, Asbjørn Rasmussen, and Michele Burrello. "Dyonic zero-energy modes". In: *Phys. Rev. B* 98 (24 Dec. 2018), p. 245135. DOI: 10.1103/PhysRevB.98.245135. URL: https://link.aps.org/doi/10.1103/PhysRevB.98.245135.

[94] Hristo N. Djidjev. "Quantum Annealing with Inequality Constraints: The Set Cover Problem". In: *Advanced Quantum Technologies* 6 (2023), p. 2300104. DOI: 10.1002/qute.202300104.

[95] Kohdai Kuroiwa and Yuya O. Nakagawa. "Penalty methods for a variational quantum eigensolver". In: *Phys. Rev. Res.* 3 (1 Feb. 2021), p. 013197. DOI: 10.1103/PhysRevResearch.3.013197. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.3.013197.

[96] Carlos A Coello Coello. "Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: a survey of the state of the art". In: *Computer Methods in Applied Mechanics and Engineering* 191.11 (2002), pp. 1245–1287. ISSN: 0045-7825. DOI: https://doi.org/10.1016/S0045-7825(01)00323-1. URL: https://www.sciencedirect.com/science/article/pii/S0045782501003231.

[97] Jonathan Welch et al. "Efficient quantum circuits for diagonal unitaries without ancillas". In: *New Journal of Physics* 16.3 (2014), p. 033040. URL: https://iopscience.iop.org/article/10.1088/1367-2630/16/3/033040/meta.

[98] Stuart Hadfield. "On the representation of Boolean and real functions as Hamiltonians for quantum computing". In: *ACM Transactions on Quantum Computing* 2.4 (2021), pp. 1–21. URL: https://dl.acm.org/doi/full/10.1145/3478519.

[99] Nick Sauerwein et al. "Engineering random spin models with atoms in a high-finesse cavity". In: *Nature Physics* 19.8 (2023), pp. 1128–1134. URL: https://www.nature.com/articles/s41567-023-02033-3.

[100] Ons Sassi and Ammar Oulamara. "Electric vehicle scheduling and optimal charging problem: complexity, exact and heuristic approaches". In: *International Journal of Production Research* 55.2 (2017), pp. 519–535. URL: https://www.tandfonline.com/doi/abs/10.1080/00207543.2016.1192695.

[101] Pauli Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[102] A Yu Kitaev. "Fault-tolerant quantum computation by anyons". In: *Annals of physics* 303.1 (2003), pp. 2–30. URL: https://www.sciencedirect.com/science/article/pii/S0003491602000180?casa_token=JbcLxiDH1KgAAAAA:Z78w7v1JWlxwgaP62OqiqS-6Z-B5U5MQV07XBJZIbgqJmkdGbCxFazCStqaFSENG2j9jPF07.

[103] Masafumi Udagawa, Ludovic Jaubert, et al. "Spin Ice". In: *Spin Ice* (2021). URL: https://link.springer.com/book/10.1007/978-3-030-70860-3.

[104] Cristian Micheletti, Philipp Hauke, and Pietro Faccioli. "Polymer physics by quantum computing". In: *Physical Review Letters* 127.8 (2021), p. 080501. URL: https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.127.080501.

[105] Francesco Slongo et al. "Quantum-inspired encoding enhances stochastic sampling of soft matter systems". In: *Science Advances* 9.43 (2023), eadi0204. URL: https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.127.080501.

[106] Franz Georg Fuchs et al. "Constraint preserving mixers for the quantum approximate optimization algorithm". In: *Algorithms* 15.6 (2022), p. 202. URL: https://www.mdpi.com/1999-4893/15/6/202.

[107] Stuart Hadfield et al. "From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz". In: *Algorithms* 12.2 (2019), p. 34. DOI: https://doi.org/10.3390/a12020034.

[108] Andreas Bärtschi and Stephan Eidenbenz. "Grover Mixers for QAOA: Shifting Complexity from Mixer Design to State Preparation". In: *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2020), pp. 72–82. DOI: 10.1109/QCE49297.2020.00020. URL: https://ieeexplore.ieee.org/abstract/document/9259965?casa_token=K-97Mmf157YAAAAA:z02ptP_fUJhfZeRwwRHxjkTqK5X3MHWbtrni7EouW
q1YrQx8P7.

[109] Yannick Deller et al. "Quantum approximate optimization algorithm for qudit systems". In: *Phys. Rev. A* 107.6 (2023), p. 062410. DOI: 10.1103/PhysRevA.107.062410. URL: https://link.aps.org/doi/10.1103/PhysRevA.107.062410.

[110] Jonathan Wurtz and Peter J Love. "Counterdiabaticity and the quantum approximate optimization algorithm". In: *Quantum* 6 (Jan. 2022), p. 635. DOI: 10.22331/q-2022-01-27-635. URL: https://quantum-journal.org/papers/q-2022-01-27-635/.

[111] Pieter W. Claeys et al. "Floquet-Engineering Counterdiabatic Protocols in Quantum Many-Body Systems". In: *Phys. Rev. Lett.* 123 (9 Aug. 2019), p. 090602. DOI: 10.1103/PhysRevLett.123.090602. URL: https://link.aps.org/doi/10.1103/PhysRevLett.123.090602.

[112] Dries Sels and Anatoli Polkovnikov. "Minimizing irreversible losses in quantum systems by local counterdiabatic driving". In: *Proc. Natl. Acad. Sci. U.S.A.* 114.20 (2017), E3909–E3916. DOI: 10.1073/pnas.1619826114.

[113] Dan Sun et al. "Optimizing counterdiabaticity by variational quantum circuits". In: *Philos. Trans. R. Soc. A* 380.2239 (2022), p. 20210282. DOI: 10.1098/rsta.2021.0282.

[114] Jiahao Yao, Lin Lin, and Marin Bukov. "Reinforcement Learning for Many-Body Ground-State Preparation Inspired by Counterdiabatic Driving". In: *Phys. Rev. X* 11 (3 Sept. 2021), p. 031070. DOI: 10.1103/PhysRevX.11.031070. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.031070.

[115] Antonio Ferrer-Sánchez et al. "Physics-informed neural networks for an optimal counterdiabatic quantum computation". In: *Machine Learning: Science and Technology* 5.2 (May 2024), p. 025035. DOI: 10.1088/2632-2153/ad450f. URL: https://dx.doi.org/10.1088/2632-2153/ad450f.

[116] Kazutaka Takahashi and Adolfo del Campo. "Shortcuts to Adiabaticity in Krylov Space". In: *Phys. Rev. X* 14 (1 Feb. 2024), p. 011032. DOI: 10.1103/PhysRevX.14.011032. URL: https://link.aps.org/doi/10.1103/PhysRevX.14.011032.

[117] P. Chandarana et al. "Digitized-counterdiabatic quantum approximate optimization algorithm". In: *Phys. Rev. Res.* 4 (1 Feb. 2022), p. 013141. DOI: 10.1103/PhysRevResearch.4.013141. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.4.013141.

[118] Pranav Chandarana et al. "Digitized Counterdiabatic Quantum Algorithm for Protein Folding". In: *Phys. Rev. Appl.* 20 (1 July 2023), p. 014024. DOI: 10.1103/PhysRevApplied.20.014024. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.20.014024.

[119] Alejandro Gomez Cadavid et al. "Efficient DCQO algorithm within the impulse regime for portfolio optimization". In: *arXiv:2308.15475* (2023). URL: https://arxiv.org/abs/2308.15475.

[120] Robert Zeier and Thomas Schulte-Herbrüggen. "Symmetry principles in quantum systems theory". In: *J. Math. Phys.* 52.11 (Nov. 2011), p. 113510. ISSN: 0022-2488. DOI: 10.1063/1.3657939.

[121] Kris Coolsaet, Sven D'hondt, and Jan Goedgebeur. "House of Graphs 2.0: A database of interesting graphs and more". In: *Discrete Appl. Math.* 325 (2023). Available at https://houseofgraphs.org, pp. 97–107. ISSN: 0166-218X. DOI: 10.1016/j.dam.2022.10.013. URL: https://www.sciencedirect.com/science/article/pii/S0166218X22004036.

[122] Johannes Siemons. "Automorphism groups of graphs". In: *Arch. Math.* 41.4 (Sept. 1983), pp. 379–384. ISSN: 1420-8938. DOI: 10.1007/BF01371410. URL: https://doi.org/10.1007/BF01371410.

[123] Frédéric Sauvage et al. "Building spatial symmetries into parameterized quantum circuits for faster training". In: *Quantum Sci. Technol.* 9.1 (2024), p. 015029. DOI: 10.1088/2058-9565/ad152e. URL: https://doi.org/10.1088/2058-9565/ad152e.

[124] Stephane Foldes. "Symmetries of directed graphs and the Chinese remainder theorem". In: *J. Comb. Theory B* 28.1 (1980), pp. 18–25. ISSN: 0095-8956. DOI: 10.1016/0095-8956(80)90053-2. URL: https://www.sciencedirect.com/science/article/pii/0095895680900532.

[125] C. Hubig et al. "Strictly single-site DMRG algorithm with subspace expansion". In: *Phys. Rev. B* 91 (15 Apr. 2015), p. 155115. DOI: 10.1103/PhysRevB.91.155115. URL: https://link.aps.org/doi/10.1103/PhysRevB.91.155115.

[126] Claudius Hubig. "Symmetry-protected tensor networks". PhD thesis. Ludwig Maximilian University, 2017.

[127] Christos H. Papadimitriou and Mihalis Yannakakis. "Optimization, approximation, and complexity classes". In: *J. Comput. Syst. Sci.* 43.3 (1991), pp. 425–440. ISSN: 0022-0000. DOI: 10.1016/0022-0000(91)90023-X. URL: https://www.sciencedirect.com/science/article/pii/002200009190023X.

[128] F. Hadlock. "Finding a Maximum Cut of a Planar Graph in Polynomial Time". In: *SIAM J. Comput.* 4.3 (1975), pp. 221–225. DOI: 10.1137/0204019. URL: https://doi.org/10.1137/0204019.

[129] M. J. D. Powell. "A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation". In: *Advances in Optimization and Numerical Analysis* (1994), pp. 51–67. DOI: 10.1007/978-94-015-8330-5_4. URL: https://doi.org/10.1007/978-94-015-8330-5_4.

[130] Michael Ragone et al. "A unified theory of barren plateaus for deep parametrized quantum circuits". In: *Nature Communications* 15.1 (Aug. 2024), p. 7172. ISSN: 2041-1723. DOI: 10.1038/s41467-024-49909-3. URL: https://doi.org/10.1038/s41467-024-49909-3.

[131] Xuchen You, Shouvanik Chakrabarti, and Xiaodi Wu. "A convergence theory for over-parameterized variational quantum eigensolvers". In: *arXiv:2205.12481* (2022). URL: https://arxiv.org/abs/2205.12481.

[132] A. Frieze and M. Jerrum. "Improved approximation algorithms for MAXk-CUT and MAX BISECTION". In: *Algorithmica* 18.1 (May 1997), pp. 67–81. ISSN: 1432-0541. DOI: 10.1007/BF02523688. URL: https://doi.org/10.1007/BF02523688.

[133] Samson Wang et al. "Noise-induced barren plateaus in variational quantum algorithms". In: *Nature Communications* 12.1 (Nov. 2021), p. 6961. ISSN: 2041-1723. DOI: 10.1038/s41467-021-27045-6. URL: https://doi.org/10.1038/s41467-021-27045-6.

[134] Phattharaporn Singkanipa and Daniel A. Lidar. "Beyond unital noise in variational quantum algorithms: noise-induced barren plateaus and limit sets". In: *Quantum* 9 (Jan. 2025), p. 1617. ISSN: 2521-327X. DOI: 10.22331/q-2025-01-30-1617. URL: https://doi.org/10.22331/q-2025-01-30-1617.

[135] Antonio Anna Mele et al. "Noise-induced shallow circuits and absence of barren plateaus". In: *arXiv:2403.13927* (2004). DOI: 10.48550/arXiv.2403.13927.

[136] Bill Fefferman et al. "Effect of Nonunital Noise on Random-Circuit Sampling". In: *PRX Quantum* 5 (3 July 2024), p. 030317. DOI: 10.1103/PRXQuantum.5.030317. URL: https://link.aps.org/doi/10.1103/PRXQuantum.5.030317.

[137] Adrián Pérez-Salinas, Hao Wang, and Xavier Bonet-Monroig. "Analyzing Variational Quantum Landscapes with Information Content". In: *npj Quantum Information* 10.1 (Feb. 2024), p. 27. ISSN: 2056-6387. DOI: 10.1038/s41534-024-00819-8.

[138] Mario A. Muñoz, Michael Kirley, and Saman K. Halgamuge. "Exploratory Landscape Analysis of Continuous Space Optimization Problems Using Information Content". In: *IEEE Transactions on Evolutionary Computation* 19.1 (2015), pp. 74–87. DOI: 10.1109/TEVC.2014.2302006. URL: https://ieeexplore.ieee.org/document/6719480.

[139] T. Lanting et al. "Entanglement in a Quantum Annealing Processor". In: *Phys. Rev. X* 4 (2 May 2014), p. 021041. DOI: 10.1103/PhysRevX.4.021041. URL: https://link.aps.org/doi/10.1103/PhysRevX.4.021041.

[140] Philipp Hauke et al. "Probing entanglement in adiabatic quantum optimization with trapped ions". In: *Frontiers in Physics* 3 (2015), p. 21. URL: https://www.frontiersin.org/journals/physics/articles/10.3389/fphy.2015.00021/full.

[141] Pablo Díez-Valle, Diego Porras, and Juan José García-Ripoll. "Quantum variational optimization: The role of entanglement and problem hardness". In: *Phys. Rev. A* 104 (6 Dec. 2021), p. 062426. DOI: 10.1103/PhysRevA.104.062426. URL: https://link.aps.org/doi/10.1103/PhysRevA.104.062426.

[142] Maxime Dupont et al. "Calibrating the Classical Hardness of the Quantum Approximate Optimization Algorithm". In: *PRX Quantum* 3 (4 Dec. 2022), p. 040339. DOI: 10.1103/PRXQuantum.3.040339. URL: https://link.aps.org/doi/10.1103/PRXQuantum.3.040339.

[143] Rishi Sreedhar et al. "The Quantum Approximate Optimization Algorithm performance with low entanglement and high circuit depth". In: *arXiv:2207.03404* (2022). URL: https://arxiv.org/abs/2207.03404.

[144] Yanzhu Chen et al. "How Much Entanglement Do Quantum Optimization Algorithms Require?" In: *Quantum 2.0 Conference and Exhibition* (2022), QM4A.2. DOI: 10.1364/QUANTUM.2022.QM4A.2. URL: https://opg.optica.org/abstract.cfm?URI=QUANTUM-2022-QM4A.2.

[145] Philipp Hauke et al. "Measuring multipartite entanglement through dynamic susceptibilities". In: *Nat. Phys.* 12.8 (2016), pp. 778–782. URL: https://doi.org/10.1038/nphys3700.

[146] Gopal Chandra Santra et al. “Squeezing and quantum approximate optimization”. In: *Phys. Rev. A* 109 (1 Jan. 2024), p. 012413. DOI: 10.1103/PhysRevA.109.012413. URL: https://link.aps.org/doi/10.1103/PhysRevA.109.012413.

[147] Vittorio Vitale et al. “Robust Estimation of the Quantum Fisher Information on a Quantum Processor”. In: *PRX Quantum* 5 (3 Aug. 2024), p. 030338. DOI: 10.1103/PRXQuantum.5.030338. URL: https://link.aps.org/doi/10.1103/PRXQuantum.5.030338.

[148] Daniel Gottesman. “Theory of fault-tolerant quantum computation”. In: *Phys. Rev. A* 57 (1 Jan. 1998), pp. 127–137. DOI: 10.1103/PhysRevA.57.127. URL: https://link.aps.org/doi/10.1103/PhysRevA.57.127.

[149] Sergey Bravyi and Jeongwan Haah. “Magic-state distillation with low overhead”. In: *Phys. Rev. A* 86 (5 Nov. 2012), p. 052329. DOI: 10.1103/PhysRevA.86.052329. URL: https://link.aps.org/doi/10.1103/PhysRevA.86.052329.

[150] Earl T Campbell, Barbara M Terhal, and Christophe Vuillot. “Roads towards fault-tolerant universal quantum computation”. In: *Nature* 549.7671 (2017), pp. 172–179. URL: https://www.nature.com/articles/nature23460.

[151] Aram W Harrow and Ashley Montanaro. “Quantum computational supremacy”. In: *Nature* 549.7671 (2017), pp. 203–209. URL: https://www.nature.com/articles/nature23458.

[152] David Gross. “Hudson’s theorem for finite-dimensional quantum systems”. In: *Journal of mathematical physics* 47.12 (2006). URL: https://pubs.aip.org/aip/jmp/article/47/12/122107/919571/Hudson-s-theorem-for-finite-dimensional-quantum.

[153] William K Wootters. “A Wigner-function formulation of finite-state quantum mechanics”. In: *Annals of Physics* 176.1 (1987), pp. 1–21. URL: https://www.sciencedirect.com/science/article/abs/pii/000349168790176X.

[154] Dmitry Panchenko. *The sherrington-kirkpatrick model*. Springer Science & Business Media, 2013. URL: https://link.springer.com/book/10.1007/978-1-4614-6289-7.

[155] Maxime Dupont et al. “Entanglement perspective on the quantum approximate optimization algorithm”. In: *Phys. Rev. A* 106.2 (2022), p. 022423. URL: https://link.aps.org/doi/10.1103/PhysRevA.106.022423.

[156] Salvatore FE Oliviero et al. “Measuring magic on a quantum processor”. In: *npj Quantum Information* 8.1 (2022), p. 148. URL: https://www.nature.com/articles/s41534-022-00666-5.

[157] David Sherrington and Scott Kirkpatrick. “Solvable Model of a Spin-Glass”. In: *Phys. Rev. Lett.* 35 (26 Dec. 1975), pp. 1792–1796. DOI: 10.1103/PhysRevLett.35.1792. URL: https://link.aps.org/doi/10.1103/Phys.

[158] Davide Venturelli et al. “Quantum Optimization of Fully Connected Spin Glasses”. In: *Phys. Rev. X* 5 (3 Sept. 2015), p. 031040. DOI: 10.1103/PhysRevX.5.031040. URL: https://link.aps.org/doi/10.1103/PhysRevX.5.031040.

[159] Samuel Mugel et al. “Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks”. In: *Phys. Rev. R* 4.1 (2022), p. 013006. URL: https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.4.013006.

[160] Stefan H Sack and Maksym Serbyn. "Quantum annealing initialization of the quantum approximate optimization algorithm". In: *Quantum* 5 (2021), p. 491. URL: https://quantum-journal.org/papers/q-2021-07-01-491/.

[161] Xhek Turkeshi, Anatoly Dymarsky, and Piotr Sierant. "Pauli spectrum and nonstabilizerness of typical quantum many-body states". In: *Phys. Rev. B* 111 (5 Feb. 2025), p. 054301. DOI: 10.1103/PhysRevB.111.054301. URL: https://link.aps.org/doi/10.1103/PhysRevB.111.054301.

[162] Xhek Turkeshi, Emanuele Tirrito, and Piotr Sierant. "Magic spreading in random quantum circuits". In: *Nature Communications* 16.1 (2025), p. 2575. URL: https://www.nature.com/articles/s41467-025-57704-x.

[163] Giuseppe E. Santoro et al. "Theory of Quantum Annealing of an Ising Spin Glass". In: *Science* 295.5564 (2002), pp. 2427–2430. DOI: 10.1126/science.1068774. URL: https://www.science.org/doi/abs/10.1126/science.1068774.

[164] Satoshi Morita and Hidetoshi Nishimori. "Mathematical foundation of quantum annealing". In: *Journal of Mathematical Physics* 49.12 (2008). URL: https://pubs.aip.org/aip/jmp/article/49/12/125210/231148.

[165] Atanu Rajak et al. "Quantum annealing: An overview". In: *Philosophical Transactions of the Royal Society A* 381.2241 (2023), p. 20210417. URL: https://royalsocietypublishing.org/doi/10.1098/rsta.2021.0417.

[166] Ulrich Schollwöck. "The density-matrix renormalization group in the age of matrix product states". In: *Annals of physics* 326.1 (2011), pp. 96–192. URL: https://www.sciencedirect.com/science/article/pii/S0003491610001752.

[167] Román Orús. "A practical introduction to tensor networks: Matrix product states and projected entangled pair states". In: *Annals of physics* 349 (2014), pp. 117–158. URL: https://www.sciencedirect.com/science/article/pii/S0003491614001596.

[168] Román Orús. "Tensor networks for complex quantum systems". In: *Nature Reviews Physics* 1.9 (2019), pp. 538–550. URL: https://www.nature.com/articles/s42254-019-0086-7.

[169] Shi-Ju Ran et al. *Tensor network contractions: methods and applications to quantum many-body systems*. Springer Nature, 2020. URL: https://link.springer.com/book/10.1007/978-3-030-34489-4.

[170] Poetri Sonya Tarabunga et al. "Nonstabilizerness via matrix product states in the pauli basis". In: *Physical Review Letters* 133.1 (2024), p. 010601. URL: https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.131.180401.

[171] Joe O'Gorman and Earl T. Campbell. "Quantum computation with realistic magic-state factories". In: *Phys. Rev. A* 95 (3 Mar. 2017), p. 032338. DOI: 10.1103/PhysRevA.95.032338. URL: https://link.aps.org/doi/10.1103/PhysRevA.95.032338.

[172] Earl T Campbell. "Early fault-tolerant simulations of the Hubbard model". In: *Quantum Science and Technology* 7.1 (2021), p. 015007. URL: https://iopscience.iop.org/article/10.1088/2058-9565/ac3110/meta.

[173] Federica Maria Surace et al. "String-breaking dynamics in quantum adiabatic and diabatic processes". In: *arXiv preprint arXiv:2411.10652* (2024).

[174] Matthew Fishman, Steven R. White, and E. Miles Stoudenmire. "The ITensor Software Library for Tensor Network Calculations". In: *SciPost Phys. Codebases* (2022), p. 4. DOI: 10.21468/SciPostPhysCodeb.4. URL: https://scipost.org/10.21468/SciPostPhysCodeb.4.

[175] Claudio Rebbi. *Lattice gauge theories and Monte Carlo simulations*. World scientific, 1983.

[176] Michael Creutz, Laurence Jacobs, and Claudio Rebbi. "Monte Carlo computations in lattice gauge theories". In: *Physics Reports* 95.4 (1983), pp. 201–282. ISSN: 0370-1573. DOI: https://doi.org/10.1016/0370-1573(83)90016-9. URL: https://www.sciencedirect.com/science/article/pii/0370157383900169.

[177] N. Brambilla et al. "QCD and strongly coupled gauge theories: challenges and perspectives". In: *The European Physical Journal C* 74.10 (Oct. 2014), p. 2981. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-014-2981-5. URL: https://doi.org/10.1140/epjc/s10052-014-2981-5.

[178] Gert Aarts et al. "Phase Transitions in Particle Physics: Results and Perspectives from Lattice Quantum Chromo-Dynamics". In: *Progress in Particle and Nuclear Physics* 133 (2023), p. 104070. ISSN: 0146-6410. DOI: https://doi.org/10.1016/j.ppnp.2023.104070. URL: https://www.sciencedirect.com/science/article/pii/S0146641023000510.

[179] Henry Lamm, Scott Lawrence, and Yukari Yamauchi. "General methods for digital quantum simulation of gauge theories". In: *Phys. Rev. D* 100 (3 Aug. 2019), p. 034518. DOI: 10.1103/PhysRevD.100.034518. URL: https://link.aps.org/doi/10.1103/PhysRevD.100.034518.

[180] Mari Carmen Bañuls et al. "Simulating lattice gauge theories within quantum technologies". In: *The European Physical Journal D* 74.8 (Aug. 2020), p. 165. ISSN: 1434-6079. DOI: 10.1140/epjd/e2020-100571-8. URL: https://doi.org/10.1140/epjd/e2020-100571-8.

[181] Monika Aidelsburger et al. "Cold atoms meet lattice gauge theory". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 380.2216 (2022), p. 20210064. DOI: 10.1098/rsta.2021.0064. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2021.0064.

[182] Erez Zohar. "Quantum simulation of lattice gauge theories in more than one space dimension-requirements, challenges and methods". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 380.2216 (2022), p. 20210069. DOI: 10.1098/rsta.2021.0069. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2021.0069.

[183] S Chandrasekharan and U.-J Wiese. "Quantum link models: A discrete approach to gauge theories". In: *Nuclear Physics B* 492.1 (1997), pp. 455–471. ISSN: 0550-3213. DOI: https://doi.org/10.1016/S0550-3213(97)80041-7. URL: https://www.sciencedirect.com/science/article/pii/S0550321397800417.

[184] L. Tagliacozzo, A. Celi, and M. Lewenstein. "Tensor Networks for Lattice Gauge Theories with Continuous Groups". In: *Phys. Rev. X* 4 (4 Nov. 2014), p. 041024. DOI: 10.1103/PhysRevX.4.041024. URL: https://link.aps.org/doi/10.1103/PhysRevX.4.041024.

[185] E. Zohar and M. Burrello. "Building projected entangled pair states with a local gauge symmetry". In: *New J. Phys.* 18.4 (Apr. 2016), p. 043008. DOI: 10.1088/1367-2630/18/4/043008. URL: https://doi.org/10.1088/1367-2630/18/4/043008.

[186] Timo Felser et al. "Two-Dimensional Quantum-Link Lattice Quantum Electrodynamics at Finite Density". In: *Phys. Rev. X* 10 (4 Nov. 2020), p. 041040. DOI: 10.1103/PhysRevX.10.041040. URL: https://link.aps.org/doi/10.1103/PhysRevX.10.041040.

[187] Manu Mathur and T. P. Sreeraj. "Canonical transformations and loop formulation of SU(N) lattice gauge theories". In: *Phys. Rev. D* 92 (12 Dec. 2015), p. 125018. DOI: 10.1103/PhysRevD.92.125018. URL: https://link.aps.org/doi/10.1103/PhysRevD.92.125018.

[188] D. Paulson et al. "Simulating 2D Effects in Lattice Gauge Theories on a Quantum Computer". In: *PRX Quantum* 2 (3 Aug. 2021), p. 030334. DOI: 10.1103/PRXQuantum.2.030334. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.030334.

[189] Jan F. Haase et al. "A resource efficient approach for quantum and classical simulations of gauge theories in particle physics". In: *Quantum* 5 (Feb. 2021), p. 393. ISSN: 2521-327X. DOI: 10.22331/q-2021-02-04-393. URL: https://doi.org/10.22331/q-2021-02-04-393.

[190] Simone Notarnicola et al. "Discrete Abelian gauge theories for quantum simulations of QED". In: *Journal of Physics A: Mathematical and Theoretical* 48.30 (July 2015), 30FT01. DOI: 10.1088/1751-8113/48/30/30FT01. URL: https://dx.doi.org/10.1088/1751-8113/48/30/30FT01.

[191] Elisa Ercolessi et al. "Phase transitions in $\mathbb{Z}_n$ gauge models: Towards quantum simulations of the Schwinger-Weyl QED". In: *Phys. Rev. D* 98 (7 Oct. 2018), p. 074503. DOI: 10.1103/PhysRevD.98.074503. URL: https://link.aps.org/doi/10.1103/PhysRevD.98.074503.

[192] Erik J. Gustafson et al. "Primitive quantum gates for an $SU(2)$ discrete subgroup: Binary tetrahedral". In: *Phys. Rev. D* 106 (11 Dec. 2022), p. 114501. DOI: 10.1103/PhysRevD.106.114501. URL: https://link.aps.org/doi/10.1103/PhysRevD.106.114501.

[193] Daniel González-Cuadra et al. "Hardware Efficient Quantum Simulation of Non-Abelian Gauge Theories with Qudits on Rydberg Platforms". In: *Phys. Rev. Lett.* 129 (16 Oct. 2022), p. 160501. DOI: 10.1103/PhysRevLett.129.160501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.129.160501.

[194] A. Mariani, S. Pradhan, and E. Ercolessi. "Hamiltonians and gauge-invariant Hilbert space for lattice Yang-Mills-like theories with finite gauge group". In: *Phys. Rev. D* 107 (11 June 2023), p. 114513. DOI: 10.1103/PhysRevD.107.114513. URL: https://link.aps.org/doi/10.1103/PhysRevD.107.114513.

[195] Gabriele Calliari et al. "Field digitization scaling in a $\mathbb{Z}_N \subset U(1)$ symmetric model". In: *arXiv:2507.22984* (2025). URL: https://arxiv.org/abs/2507.22984.

[196] G. Semeghini et al. "Probing topological spin liquids on a programmable quantum simulator". In: *Science* 374.6572 (2021), pp. 1242–1247. DOI: 10.1126/science.abi8794. URL: https://www.science.org/doi/abs/10.1126/science.abi8794.

[197] K. J. Satzinger et al. "Realizing topologically ordered states on a quantum processor". In: *Science* 374.6572 (2021), pp. 1237–1241. DOI: 10.1126/science.abi8378. URL: https://www.science.org/doi/abs/10.1126/science.abi8378.

[198] T. I. Andersen et al. "Non-Abelian braiding of graph vertices in a superconducting processor". In: *Nature* 618.7964 (June 2023), pp. 264–269. ISSN: 1476-4687. DOI: 10.1038/s41586-023-05954-4. URL: https://doi.org/10.1038/s41586-023-05954-4.

[199] Mohsin Iqbal et al. "Non-Abelian topological order and anyons on a trapped-ion processor". In: *Nature* 626.7999 (Feb. 2024), pp. 505–511. ISSN: 1476-4687. DOI: 10.1038/s41586-023-06934-4. URL: https://doi.org/10.1038/s41586-023-06934-4.

[200] Tyler A. Cochran et al. "Visualizing Dynamics of Charges and Strings in (2+1)D Lattice Gauge Theories". In: *Nature* 642 (2025), p. 315. URL: https://www.nature.com/articles/s41586-025-08999-9.

[201] Daniel González-Cuadra et al. "Observation of string breaking on a (2+ 1) D Rydberg quantum simulator". In: *Nature* 642 (2025), p. 321. URL: https://doi.org/10.1038/s41586-025-09051-6.

[202] Jesús Cobos et al. *Real-Time Dynamics in a (2+1)-D Gauge Theory: The Stringy Nature on a Superconducting Quantum Simulator*. 2025. arXiv: 2507.08088 [quant-ph]. URL: https://arxiv.org/abs/2507.08088.

[203] Julian Bender et al. "Digital quantum simulation of lattice gauge theories in three spatial dimensions". In: *New Journal of Physics* 20.9 (Sept. 2018), p. 093001. DOI: 10.1088/1367-2630/aadb71. URL: https://doi.org/10.1088/1367-2630/aadb71.

[204] Michael Fromm, Owe Philipsen, and Christopher Winterowd. "Dihedral lattice gauge theories on a quantum annealer". In: *EPJ Quantum Technology* 10.1 (Aug. 2023), p. 31. ISSN: 2196-0763. DOI: 10.1140/epjqt/s40507-023-00188-9. URL: https://doi.org/10.1140/epjqt/s40507-023-00188-9.

[205] Edoardo Ballini et al. "Quantum computation of thermal averages for a non-Abelian D_4 lattice gauge theory via quantum Metropolis sampling". In: *Phys. Rev. D* 109 (3 Feb. 2024), p. 034510. DOI: 10.1103/PhysRevD.109.034510. URL: https://link.aps.org/doi/10.1103/PhysRevD.109.034510.

[206] Emanuele Gaz et al. "Quantum simulation of non-Abelian lattice gauge theories: A variational approach to $\mathbb{D}_8$ with dynamical matter". In: *Phys. Rev. Res.* 7 (3 July 2025), p. 033012. DOI: 10.1103/l8b6-h5s5. URL: https://link.aps.org/doi/10.1103/l8b6-h5s5.

[207] Jens Nyhegn, Chia-Min Chung, and Michele Burrello. "$\mathbb{Z}_N$ lattice gauge theory in a ladder geometry". In: *Phys. Rev. Res.* 3 (1 Feb. 2021), p. 013133. DOI: 10.1103/PhysRevResearch.3.013133. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.3.013133.

[208] Victor V. Albert et al. *Spin chains, defects, and quantum wires for the quantum-double edge*. 2021. arXiv: 2111.12096 [cond-mat.str-el]. URL: https://arxiv.org/abs/2111.12096.

[209] S. Pradhan, A. Maroncelli, and E. Ercolessi. "Discrete Abelian lattice gauge theories on a ladder and their dualities with quantum clock models". In: *Phys. Rev. B* 109 (6 Feb. 2024), p. 064410. DOI: 10.1103/PhysRevB.109.064410. URL: https://link.aps.org/doi/10.1103/PhysRevB.109.064410.

[210] Matthew Fishman, Steven R. White, and E. Miles Stoudenmire. "Codebase release 0.3 for ITensor". In: *SciPost Phys. Codebases* (2022), 4–r0.3. DOI: 10.21468/SciPostPhysCodeb.4-r0.3. URL: https://scipost.org/10.21468/SciPostPhysCodeb.4-r0.3.

[211] Eduardo Fradkin and Stephen H. Shenker. "Phase diagrams of lattice gauge theories with Higgs fields". In: *Phys. Rev. D* 19 (12 June 1979), pp. 3682–3697. DOI: 10.1103/PhysRevD.19.3682. URL: https://link.aps.org/doi/10.1103/PhysRevD.19.3682.

[212] Nathanan Tantivasadakarn, Ruben Verresen, and Ashvin Vishwanath. "Shortest Route to Non-Abelian Topological Order on a Quantum Processor". In: *Phys. Rev. Lett.* 131 (6 Aug. 2023), p. 060405. DOI: 10.1103/PhysRevLett.131.060405. URL: https://link.aps.org/doi/10.1103/PhysRevLett.131.060405.

[213] Chiu Fan Bowen Lo et al. *Universal Quantum Computation with the S_3 Quantum Double: A Pedagogical Exposition*. 2025. arXiv: 2502.14974 [quant-ph]. URL: https://arxiv.org/abs/2502.14974.

[214] Giuseppe Clemente, Arianna Crippa, and Karl Jansen. "Strategies for the determination of the running coupling of $(2+1)$-dimensional QED with quantum computing". In: *Phys. Rev. D* 106 (11 Dec. 2022), p. 114511. DOI: 10.1103/PhysRevD.106.114511. URL: https://link.aps.org/doi/10.1103/PhysRevD.106.114511.

[215] Arianna Crippa et al. "Towards determining the (2+1)-dimensional quantum electrodynamics running coupling with Monte Carlo and quantum computing methods". In: *Communications Physics* 8.1 (Sept. 2025), p. 367. ISSN: 2399-3650. DOI: 10.1038/s42005-025-02243-6. URL: https://doi.org/10.1038/s42005-025-02243-6.

[216] Gunnar S. Bali and Antonio Pineda. "QCD phenomenology of static sources and gluonic excitations at short distances". In: *Phys. Rev. D* 69 (9 May 2004), p. 094001. DOI: 10.1103/PhysRevD.69.094001. URL: https://link.aps.org/doi/10.1103/PhysRevD.69.094001.

[217] Slavo Kratochvila and Philippe de Forcrand. "Observing string breaking with Wilson loops". In: *Nuclear Physics B* 671 (2003), pp. 103–132. ISSN: 0550-3213. DOI: https://doi.org/10.1016/j.nuclphysb.2003.08.014. URL: https://www.sciencedirect.com/science/article/pii/S0550321303006916.

[218] Jannis Herr, Carolin Schlosser, and Marc Wagner. "Gluelump masses and mass splittings from SU(3) lattice gauge theory". In: *Phys. Rev. D* 109 (3 Feb. 2024), p. 034516. DOI: 10.1103/PhysRevD.109.034516. URL: https://link.aps.org/doi/10.1103/PhysRevD.109.034516.

[219] Scott Aaronson and Daniel Gottesman. "Improved simulation of stabilizer circuits". In: *Phys. Rev. A* 70 (5 Nov. 2004), p. 052328. DOI: 10.1103/PhysRevA.70.052328. URL: https://link.aps.org/doi/10.1103/PhysRevA.70.052328.

[220] Austin G. Fowler et al. "Surface codes: Towards practical large-scale quantum computation". In: *Phys. Rev. A* 86 (3 Sept. 2012), p. 032324. DOI: 10.1103/PhysRevA.86.032324. URL: https://link.aps.org/doi/10.1103/PhysRevA.86.032324.

[221] Martin Hebenstreit et al. "All pure fermionic non-Gaussian states are magic states for matchgate computations". In: *Physical review letters* 123.8 (2019), p. 080503. DOI: 10.1103/PhysRevLett.123.080503. URL: https://link.aps.org/doi/10.1103/PhysRevLett.123.080503.

[222] Richard Jozsa and Akimasa Miyake. "Matchgates and classical simulation of quantum circuits". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 464.2100 (2008), pp. 3089–3106. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rspa.2008.0189.

[223] Sudipto Singha Roy, Leon Carl, and Philipp Hauke. "Genuine multipartite entanglement in a one-dimensional Bose-Hubbard model with frustrated hopping". In: *Phys. Rev. B* 106 (19 Nov. 2022), p. 195158. DOI: 10.1103/PhysRevB.106.195158. URL: https://link.aps.org/doi/10.1103/PhysRevB.106.195158.

[224] Luca Lumia et al. "Measurement-induced transitions beyond Gaussianity: A single particle description". In: *Phys. Rev. Res.* 6 (2 May 2024), p. 023176. DOI: 10.1103/PhysRevResearch.6.023176. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.6.023176.

[225] Mario Collura et al. "The quantum magic of fermionic Gaussian states". In: *arXiv:2412.05367* (2024). URL: https://arxiv.org/abs/2412.05367.

[226] Pedro R. Nicácio Falcão et al. "Nonstabilizerness in U(1) lattice gauge theory". In: *Phys. Rev. B* 111 (8 Feb. 2025), p. L081102. DOI: 10.1103/PhysRevB.111.L081102. URL: https://link.aps.org/doi/10.1103/PhysRevB.111.L081102.

[227] Piotr Sierant, Paolo Stornati, and Xhek Turkeshi. "Fermionic Magic Resources of Quantum Many-Body Systems". In: *arXiv:2506.00116* (2025). URL: https://arxiv.org/abs/2506.00116.

[228] C. E. P. Robin and M. J. Savage. *Anti-Flatness and Non-Local Magic in Two-Particle Scattering Processes*. 2025. arXiv: 2510.23426 [quant-ph]. URL: https://arxiv.org/abs/2510.23426.

[229] Qiaofeng Liu, Ian Low, and Zhewei Yin. *Quantum Magic in Quantum Electrodynamics*. 2025. arXiv: 2503.03098 [quant-ph]. URL: https://arxiv.org/abs/2503.03098.

[230] Luca Barbiero et al. "Coupling ultracold matter to dynamical gauge fields in optical lattices: From flux attachment to $\mathbb{Z}_2$ lattice gauge theories". In: *Science Advances* 5.10 (2019), eaav7444. DOI: 10.1126/sciadv.aav7444. URL: https://www.science.org/doi/abs/10.1126/sciadv.aav7444.

[231] Daniel González-Cuadra et al. "Robust Topological Order in Fermionic $\mathbb{Z}_2$ Gauge Theories: From Aharonov-Bohm Instability to Soliton-Induced Deconfinement". In: *Phys. Rev. X* 10 (4 Oct. 2020), p. 041007. DOI: 10.1103/PhysRevX.10.041007. URL: https://link.aps.org/doi/10.1103/PhysRevX.10.041007.

[232] Sarmed A Rahman et al. "SU(2) lattice gauge theory on a quantum annealer". In: *Phys. Rev. D* 104 (3 Aug. 2021), p. 034501. DOI: 10.1103/PhysRevD.104.034501. URL: https://link.aps.org/doi/10.1103/PhysRevD.104.034501.

[233] Sarmed A Rahman et al. "Self-mitigating Trotter circuits for SU(2) lattice gauge theory on a quantum computer". In: *Phys. Rev. D* 106 (7 Oct. 2022), p. 074502. DOI: 10.1103/PhysRevD.106.074502. URL: https://link.aps.org/doi/10.1103/PhysRevD.106.074502.

[234] Dorota M. Grabowska, Christopher F. Kane, and Christian W. Bauer. "A Fully Gauge-Fixed SU(2) Hamiltonian for Quantum Simulations". In: *arXiv:2409.10610* (2024). DOI: 10.48550/arXiv.2409.10610. URL: https://arxiv.org/abs/2409.10610.

[235] Eduardo Fradkin and Leonard Susskind. "Order and disorder in gauge systems and magnets". In: *Phys. Rev. D* 17 (10 May 1978), pp. 2637–2658. DOI: 10.1103/PhysRevD.17.2637. URL: https://link.aps.org/doi/10.1103/PhysRevD.17.2637.

[236] Doga Murat Kürkçüoglu, Henry Lamm, and Andrea Maestri. *Qudit Gate Decomposition Dependence for Lattice Gauge Theories*. 2024. arXiv: 2410.16414 [quant-ph]. URL: https://arxiv.org/abs/2410.16414.

[237] A.Yu. Kitaev. "Fault-tolerant quantum computation by anyons". In: *Annals of Physics* 303.1 (2003), pp. 2–30. ISSN: 0003-4916. DOI: https://doi.org/10.1016/S0003-4916(02)00018-0. URL: https://www.sciencedirect.com/science/article/pii/S0003491602000180.

[238] G. Magnifico et al. "$\mathbb{Z}_N$ gauge theories coupled to topological fermions: QED_2 with a quantum mechanical θ angle". In: *Phys. Rev. B* 100 (11 Sept. 2019), p. 115152. DOI: 10.1103/PhysRevB.100.115152. URL: https://link.aps.org/doi/10.1103/PhysRevB.100.115152.

[239] Tomoya Hayata, Yoshimasa Hidaka, and Yuta Kikuchi. "Diagnosis of information scrambling from Hamiltonian evolution under decoherence". In: *Phys. Rev. D* 104 (7 Oct. 2021), p. 074518. DOI: 10.1103/PhysRevD.104.074518. URL: https://link.aps.org/doi/10.1103/PhysRevD.104.074518.

[240] Xiaojun Yao. "SU(2) gauge theory in $2+1$ dimensions on a plaquette chain obeys the eigenstate thermalization hypothesis". In: *Phys. Rev. D* 108 (3 Aug. 2023), p. L031504. DOI: 10.1103/PhysRevD.108.L031504. URL: https://link.aps.org/doi/10.1103/PhysRevD.108.L031504.

[241] Abner Shimony. "Degree of entanglement". In: *Annals of the New York Academy of Sciences* 755.1 (1995), pp. 675–679. URL: https://nyaspubs.onlinelibrary.wiley.com/doi/epdf/10.1111/j.1749-6632.1995.tb39008.x.

[242] Howard Barnum and Noah Linden. "Monotones and invariants for multi-particle quantum states". In: *Journal of Physics A: Mathematical and General* 34.35 (2001), p. 6787. URL: https://iopscience.iop.org/article/10.1088/0305-4470/34/35/305.

[243] Tzu-Chieh Wei and Paul M. Goldbart. "Geometric measure of entanglement and applications to bipartite and multipartite quantum states". In: *Phys. Rev. A* 68 (4 Oct. 2003), p. 042307. DOI: 10.1103/PhysRevA.68.042307. URL: https://link.aps.org/doi/10.1103/PhysRevA.68.042307.

[244] Aditi Sen(De) and Ujjwal Sen. "Channel capacities versus entanglement measures in multiparty quantum states". In: *Phys. Rev. A* 81 (1 Jan. 2010), p. 012308. DOI: 10.1103/PhysRevA.81.012308. URL: https://link.aps.org/doi/10.1103/PhysRevA.81.012308.

[245] Tamoghna Das et al. "Generalized geometric measure of entanglement for multiparty mixed states". In: *Physical Review A* 94.2 (2016), p. 022336. URL: https://journals.aps.org/pra/abstract/10.1103/PhysRevA.94.022336.

[246] Owidiusz Makuta, Błażej Kuzaka, and Remigiusz Augusiak. "Frustration graph formalism for qudit observables". In: *arXiv:2503.22400* (2025). URL: https://arxiv.org/abs/2503.22400.

[247] Anindita Bera and Sudipto Singha Roy. "Growth of genuine multipartite entanglement in random unitary circuits". In: *Phys. Rev. A* 102 (6 Dec. 2020), p. 062431. DOI: 10.1103/PhysRevA.102.062431. URL: https://link.aps.org/doi/10.1103/PhysRevA.102.062431.

[248] Pietro Smacchia et al. "Statistical mechanics of the cluster Ising model". In: *Phys. Rev. A* 84 (2 Aug. 2011), p. 022304. DOI: 10.1103/PhysRevA.84.022304. URL: https://link.aps.org/doi/10.1103/PhysRevA.84.022304.

[249] Ingemar Bengtsson and Karol Zyczkowski. "A brief introduction to multipartite entanglement". In: *arXiv:1612.07747* (2016). URL: https://arxiv.org/abs/1612.07747.

[250] Michael JD Powell. *A direct search optimization method that models the objective and constraint functions by linear interpolation*. Springer, 1994. URL: https://link.springer.com/chapter/10.1007/978-94-015-8330-5_4.

[251] D. M. Brink and G. R. Satchler. *Angular momentum*. Oxford University Press, 1993.

[252] Giuseppe Magnifico et al. "Lattice quantum electrodynamics in (3+ 1)-dimensions at finite density with tensor networks". In: *Nature communications* 12.1 (2021), p. 3600.

[253] Eduardo Fradkin. *Field Theories of Condensed Matter Physics*. 2nd ed. Cambridge University Press, 2013. DOI: 10.1017/CBO9781139015509.