

©Copyright 2026

Henry Froland

Quantum Simulation in the Era of Practical Quantum Computers

Henry Froland

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Martin J. Savage, Chair

David B. Kaplan

Silas Beane

Program Authorized to Offer Degree:
UW Department of Physics

University of Washington

Abstract

Quantum Simulation in the Era of Practical Quantum Computers

Henry Froland

Chair of the Supervisory Committee:
Professor and Senior INT Fellow Martin J. Savage
Department of Physics

Understanding the computational power of quantum devices has become a central focus of modern physics, driven in part by the promise of new discoveries in fundamental science. The properties of quantum many-body systems are encoded in a Hilbert space whose dimension grows exponentially with the number of constituent degrees of freedom, and the intricate patterns of quantum correlations that such systems develop resist the compression on which classical methods rely. As a result, many quantities of central physical interest, most notably real-time dynamics and non-equilibrium phenomena, lie beyond the reach of even the most powerful classical supercomputers. Quantum computers evade this barrier by construction: because the device is itself a controllable quantum system, it can store and evolve many-body states with resources that scale only polynomially in the system size. The simulation of physical systems is therefore among the primary applications of quantum computation, and arguably its original motivation. Although errors and noise still limit the coherence times of current hardware, quantum simulators have now reached the stage of being functional scientific instruments, capable of practical simulations that challenge their classical counterparts.

Extracting scientific value from these devices, however, requires more than raw hardware capability. Generic, hardware-agnostic algorithms are often prohibitively expensive on near-term machines, and so

specialized methods must be developed that are informed by the physics of the problem itself. Insights from quantum many-body theory, the symmetries and gauge invariance of the target system, the structure of its correlations and entanglement, and the way errors propagate through its dynamics, can be built directly into simulation, measurement, and error-suppression protocols, dramatically reducing the resources they demand. Developing such physics-informed methods is the unifying theme of this dissertation.

The research presented here contributes to both the practical and conceptual formulations of quantum simulation, for systems spanning a range of subfields from condensed matter to high-energy physics. The first part develops time evolution algorithms in both analog and digital settings. A framework for quantifying and optimizing algorithmic errors in Rydberg-atom simulators is presented, revealing that at a given level of device noise there exists an optimal time step. A hardware-efficient digitization of $SU(2)$ lattice gauge theory (LGT) is then presented that remains resource efficient at all values of the gauge coupling; quantum circuits implementing time evolution are constructed and simulations are performed on superconducting quantum hardware.

The second part of this thesis focuses on the efficient measurement of observables relevant for quantum simulation. Building on the investigation of the non-equilibrium dynamics of an $SU(2)$ LGT, its quantum complexity is measured experimentally with the aid of a novel error mitigation technique. Classical shadows protocols are then developed that leverage gauge invariance to achieve an exponential reduction in sample complexity compared to standard local schemes. Finally, a tomographic protocol based on a novel ansatz for non-Gaussian states is introduced and used to efficiently measure the build-up of level repulsion in a system of weakly interacting fermions, opening new avenues for studying non-equilibrium dynamics and thermalization.

The dissertation concludes with new approaches for combining fault tolerance with near-term simulation techniques, demonstrated through experiments on quantum devices. By balancing circuit depth against the protection afforded by fault-tolerant gadgets, beyond break-even performance is achieved for the calculation of local observables relevant to quantum simulation.

TABLE OF CONTENTS

	Page
List of Figures	iv
List of Tables	xviii
Acknowledgements	xxi
Dedication	xxiv
Chapter 1: Introduction	1
1.1 Nature as a Computer	1
1.2 Primer on Quantum Simulation	4
1.3 Model Systems	9
1.4 Outline of this Thesis	15
Chapter 2: Optimization of Algorithmic Errors in Analog Quantum Simulations	18
2.1 Introduction	18
2.2 Hamiltonian Engineering Methods	21
2.3 Characterization of Error Type and Error Rate	25
2.4 Results	29
2.5 Discussion	33
Chapter 3: Simulating Fully Gauge-Fixed SU(2) Hamiltonian Dynamics on Digital Quantum Computers	37
3.1 Introduction	37
3.2 Two Plaquette Hamiltonian	39
3.3 Digitization of the Mixed Basis	50
3.4 Circuits for Time Evolution	57
3.5 Real-Time Dynamics in Two Plaquettes	68
3.6 Discussions and Conclusion	71

Chapter 4:	Measuring Non-Stabilizerness in an SU(2) Lattice Gauge Theory	74
4.1	Introduction	74
4.2	Non-Stabilizerness of Two Plaquettes	77
4.3	Walsh Based Error Mitigation	82
4.4	Measurement of Non-Stabilizerness on a Quantum Device	88
4.5	Conclusion	90
Chapter 5:	Classical shadows for sample-efficient measurements of gauge-invariant observables	93
5.1	Introduction	93
5.2	Lattice Gauge Theory	97
5.3	Dual Pairs Protocols	99
5.4	Dual Product Protocol	108
5.5	Numerical Demonstrations	112
5.6	Protocol Comparisons	117
5.7	Discussion and Outlook	120
Chapter 6:	Measuring the Entanglement Structure of Non-Gaussian States	124
6.1	Introduction	124
6.2	Non-Gaussian States	126
6.3	Derivation of Parameterization	127
6.4	Quench Setup	137
6.5	Early-time Dynamics	137
6.6	Early Time Growth of Non-Gaussianity	141
6.7	Thermalization Dynamics	143
6.8	Discussion of Symmetries and Choice of the Initial State	144
6.9	Experimental Application	147
6.10	Cost, efficiency, and noise	148
6.11	Conclusions	150
Chapter 7:	Realizing Error Suppression in Partially Fault-Tolerant Quantum Simulations with IBM Quantum Computers	152
7.1	Introduction	152
7.2	Fault Tolerance in the $[[4, 2, 2]]$ Code	154
7.3	Error-Detected Quantum Simulations of The Ising Model	159
7.4	Discussion	165

Chapter 8:	Concluding Remarks	167
Appendix A:	Appendix to Chapter 2	171
A.1	Constant-Field Method Error Rate Derivation	171
A.2	Derivation of the $\frac{\pi}{2}$ -Pulse Error	172
A.3	First-Order Product Formula Error	173
A.4	Analysis of the Interaction Sums	177
Appendix B:	Appendix to Chapter 3	180
B.1	Explicit Form of Quantum Circuits	180
B.2	Explicit Matrix Elements of Hamiltonian Operators	180
B.3	Table of Lowest Energy Eigenstates in the character Irrep Basis	188
Appendix C:	Appendix to Chapter 4	190
C.1	The Non-Stabilizerness of Indicator States	190
C.2	Estimators for Stabilizer Entropies	193
C.3	Walsh-Paley Based Reconstruction of Wavefunction Probabilities	194
Appendix D:	Appendix to Chapter 5	198
D.1	Details of Sample Complexity Bound	198
D.2	Protocols with Fixed Boundary Conditions	199
Appendix E:	Appendix to Chapter 7	203
E.1	ORP Details	203
E.2	Block Placement Selection	207
E.3	Circuit Scheduling	209
E.4	Quantum Simulation Details	211
E.5	Classical Simulation Details	214
E.6	Details on Filtering Metrics	214
E.7	Fault Tolerant Benchmarking of Circuit Elements	225
E.8	Additional Analysis of Encoding Improvement	231
E.9	Classical Simulations of False-Vacuum Decay Dynamics	237
E.10	Other Layouts	243
E.11	Tables of Results	245
Bibliography		253

LIST OF FIGURES

Figure Number		Page
2.1	<i>Mapping hardware interactions to a target Hamiltonian through pulse engineering</i> The device Hamiltonian H_{dev} can be used to emulate a target system of interest with Hamiltonian H_{tar} using Trotter-like methods (top path) or continuous driving methods (bottom path). In the limit of perfect pulses (equivalently, infinite magnetic field), the dashed curves implement the target time evolution exactly. The solid curves represent the mappings that are approximate due to algorithmic errors. These errors increase the overall error when mapping between Hamiltonians due to idle errors and when implementing the time evolution operator $U(t)$ (Trotter errors). Algorithmic errors affect constant-field time evolution engineering methods as well. The constant-field method C_1 and one of the Trotter methods S_1 are labeled as resulting time evolution methods within the space of possible $U(t)$ operators.	21
2.2	<i>Pulse sequences for generating Heisenberg evolution from ZZ interactions, corresponding to Eqs. (2.6), (2.8) - (2.11)</i> For each pulse sequence, the magnetic-field strength (in units of Ω) is given as a function of time (in units of $\epsilon = \frac{\pi}{2\Omega}$). The length of each sequence U is denoted by τ_U . The idling times t_i are shown in the graphics for each of the Trotter sequences.	24
2.3	<i>Graphical description of leading order error</i> Different leading-order error terms contributing to the error rate analysis of the four Trotter methods considered in this chapter. The optimal scaling is given by the negative slope of each line at $k_1 = 1$ and the corresponding error rate scaling is given by the $k_1 = 1$ intercept.	28
2.4	<i>Error rates and bounds</i> Comparison of analytic bounds for the error rate for a single Trotter step from Eqs. (2.22) and (2.23) to numerical results as a function of the device parameter ϵ for 2×2 systems. The analytic bounds consistently exceed the numerical results by factors of $O(1) - O(10)$, which is an expected consequence of using the triangle inequality to evaluate the bounds.	31
2.5	<i>Effects of different step sizes</i> The effect of multiple Trotter steps on errors for 2×2 systems with $\epsilon \sim 10^{-5}$ and pulse sequences $S_{1/2}$ and S_1 . The quantity $\delta_1 - \delta_n = U_{\text{trot}}(T) - e^{-iH_{XXX}T} - \prod_n U_{\text{trot}}(T/n) - e^{-iH_{XXX}T} $ is positive when taking many small steps of $U_{\text{trot}}(T/n)$ replicates $e^{-iH_{XXX}T}$ better than taking one large step $U_{\text{trot}}(T)$. Here, $T = n\tau_U$ and τ_U is chosen optimally according to the Eq. (2.21). Taking many steps to reach a total time T is beneficial for $S_{1/2}$ and S_1 compared to taking a single large step when the step size is chosen optimally. This demonstrates the importance of choosing the step size correctly to take advantage of the Trotter error scaling.	32

2.6	<i>Error rates for different pulse sequences</i> Numerical error rates corresponding to Eq. (2.14) for a single Trotter step (Floquet period for $\mathcal{C}_1$) for 2×2 systems. The slopes of lines on these log-log plots indicate the scaling of the error rate for each method (see Table 2.1) and the y-intercepts give the scaling prefactor. The lines for $\mathcal{S}_{1/2}$, $\mathcal{S}_1$, and $\mathcal{C}_1$ are the same as the numerical ones in Fig. 2.4. Current device capabilities report pulse control to the ns level, meaning pulse sequences with $\epsilon \sim O(10^{-5})$ can be applied, indicated by the dashed vertical line.	34
3.1	<i>Spectrum of the Two Plaquette System</i> a) Two plaquette system with open boundary conditions. The system contains seven links, numbered 1 through 7, with the physical links denoted as $\kappa = 1, 2$. The magnetic degrees of freedom $\hat{X}_1, \hat{X}_2$ are associated with each plaquette. The dashed lines that run around each plaquette denote the loops used in the max tree gauge fixing procedure. b) Spectrum of the Hamiltonian for different bare couplings. The solid lines show the solution obtained by FEM. The black dashed line on the bottom shows the energy bound of Eq. (3.35). The gray lines denote the strong coupling limit of the character irrep spectrum. This limit is found by ignoring the contribution of the magnetic Hamiltonian, which is only valid when the gauge coupling is large. c) Low-lying spectrum of the character irrep Hamiltonian in the strong coupling limit. The first and fourth excited state are two-fold degenerate and each state is a linear superposition of the stacked basis states. See Table B.4 for further details.	40
3.2	<i>Digitization of the Wavefunction</i> The continuous range of ω is digitized on a discrete grid. As such the wave function will only be sampled at a discrete set of points. For the closest approximation to the undigitized theory, the maximum value of $\omega_{\max} \leq 2\pi$ must be chosen.	54
3.3	<i>Convergence properties of the two plaquette wavefunction</i> a) A slice of the ground state wave function for different values of the gauge coupling. As the gauge coupling decreases, the wave function becomes more concentrated around $\omega_1 = 0$, although it must go to 0 at both boundaries due to the Dirichlet boundary conditions. b) Convergence of the ground state energy to the undigitized value as the digitization resolution N_ω is increased. The convergence is plotted for a variety of $\omega_{\max}$. After quick convergence to $\epsilon_0 \sim 10^{-3}$ a slower exponential approach occurs, followed by a plateau. c) (<i>top</i>) Convergence to the ground state energy as $\omega_{\max}$ is varied. The dashed lines correspond to the full two plaquette system while the solid lines correspond to the system without any first-derivative interactions. The effect of the first derivative reduces the precision that can be achieved. Even though this term limits precision, digitization errors are still below those due to time-evolution algorithms. (<i>bottom</i>) The absolute value of the derivative of the energy w.r.t the truncation. It was heuristically found that the truncation at which this derivative takes a minimum corresponds to the optimal truncation for the analytically solvable case. This method is still applicable even when the exact solution is not known. d) The solution of the stationary condition for a range of gauge couplings. The agreement is very good with the analytical solution Eq. (3.53).	55

3.4	<i>Accuracy of the ν Truncation</i> a) Calculation of the time dependent magnetic energy for the low-energy state $ 5\rangle$, defined in Eq. (3.54), as the maximum ν value is increased for $g = 0.5$ and $N_\omega = 10$. This choice of digitization resolution ensures that any errors due to the ω_1, ω_2 truncation will be at the percent level. b) The probability $\langle \Pi_{\nu>1} \rangle$ of leaking out of the $\nu = 0, 1$ sector as a function of time for $g = 0.5, N_\omega = 10$ and $\nu_{\max} = 5$	57
3.5	<i>Circuits for the digitization presented in Sec. 3.4.1, which constructs sub-components for each class of differential operators</i> a) Implementation of the exact exponential of the second derivative, making use of Discrete Sine Transforms. This subcircuit only entangles qubits within each ω register and does not touch the ν register. The ancillas a_1, a_2 associated with each ω register are the outermost qubits. b) An important circuit identity used in this work: The exponential of an N -body Pauli string P is implemented by the staircase of CX s connecting all the qubits participating in the gate, followed by an $R_Z(\theta)$ rotation on one of the qubits and then another CX staircase. The green single qubit gates denote a local change of bases to change P into a string of Z s on all qubits e.g. if there is an X on qubit i then this will be a Hadamard H and if it is a Y then it will be a Hadamard followed by a phase gate HS . The different shades denote that the local basis change can depend on the particular Pauli string P being implemented. c) Circuits for implementing first order derivatives and mixed derivatives using the finite difference representation. <i>(top)</i> Implements the exponential of a single first derivative. Each element in this circuit implements one of the terms in the sum Eq. (3.65). The bottom qubit of each step is the target qubit for both the V_k and CRZ gates. <i>(bottom)</i> An example of one of circuit elements $e^{-it \frac{\partial}{\partial \omega_1} \frac{\partial}{\partial \omega_2}}$ that interleaves the top circuit.	59
3.6	<i>Trotter error for the circuits presented in Sec. 3.4.1, which constructs sub-components for each class of differential operators</i> Noiseless simulator results for the expectation value of the magnetic Hamiltonian at a different number of Trotter steps.	60
3.7	<i>Approximation for time evolution circuits presented in Sec. 3.4.2</i> a) The expectation value of the magnetic part of the Hamiltonian for differing digitization resolutions and percentage of terms dropped. Rapid convergence is observed for $n_q \geq 3$ and for truncation levels $\delta < .6$. b) Different stages of the time evolution operator for a second-order Trotter step. The silver circuit denotes the entire Trotter step. The blue sub-circuits act only on the ω registers, typically involving the largest angles, and is identical for each. The green sub-circuits couple each ω register to the ν register and the circuit structure is again identical for both ω . The yellow sub-circuit couples everything and generally contains the smallest rotations, therefore being the elements most affected by small-angle truncations. It also involves the highest weight terms, and so focusing on truncating these elements allows the circuit to be compressed at only a small cost in algorithmic error. c) The accuracy of first- and second-order Trotter on magnetic observables. Repeating the blue sub-circuit in the second step allows the circuit to take much larger time steps while only needing a small amount of extra gates to implement. .	65

3.8	<i>Upper bounds on resource estimates for the Pauli Truncation circuits</i> The cumulative weight of the Pauli decomposition of $H_{\square\square}$ is shown as a function of the truncation percentage δ . The color of the curves denotes the magnitude of the coefficient in the decomposition. The blue point demonstrates the truncation level and actual gate depth used in the quantum simulation in this work, highlighting the effect of circuit compilation.	68
3.9	<i>Quantum simulation of $\langle H_B \rangle$ dynamics</i> a) The time evolution due to $H_{\square\square}$ at gauge coupling $g = 0.5$. The dashed line corresponds to the solution obtained by directly taking the matrix exponential of the digitized Hamiltonian and the solid line is the Trotterized evolution. The points are the mitigated results from ibm_fez . b) The physical layout of the Heron chip. The highlighted regions are an example of the parallel layouts used to generate the data point at $t = 0.6$	69
4.1	<i>Digital quantum simulation of Two Plaquettes</i> a) The two-plaquette system with open boundary conditions is gauge-fixed using max tree gauge fixing. The system contains seven links, numbered 1 through 7, with the physical links denoted as $\kappa = 1, 2$. The magnetic degrees of freedom $\hat{X}_1, \hat{X}_2$ are associated with each plaquette. The dashed lines that run around each plaquette denote the loops used in the max tree gauge-fixing procedure. b) The digitized Hamiltonian is compiled into digital quantum gates via Trotterization that approximate e^{-iHt} . Shown below the circuit diagram is an example of initial state amplitudes given by a smooth gaussian profile in ω_1, ω_2 and the final amplitudes after time evolution by the circuit. To measure the non-stabilizerness, many copies of the time evolved state are prepared and then a set of single qubit Clifford gates are sampled from the ensemble $Cl(2)^{\otimes n}$ and appended to the end of each time-evolved state. These copies are subsequently measured and the noisy bitstring distributions are recorded. c) From the set of random measurements, error mitigation can be applied to the expectation values of different Pauli Z strings $\langle Z_k \rangle$ to estimate the noiseless values. These mitigated expectation values can be used to reconstruct the error-mitigated bitstring distribution $p(s)$, which can then be used to calculate non-linear functions of the state, like the Stabilizer Rényi Entropy or the purity.	76
4.2	<i>Non-Stabilizerness of low-energy eigenstates</i> a) Non-stabilizerness of ground state $ e_0\rangle$ (top) and first two excited eigenstates $ e_1\rangle, e_2\rangle$ (bottom). For large gauge coupling, the non-stabilizerness plateaus to a constant value for $ e_1\rangle, e_2\rangle$, which are degenerate in this limit. For a fixed resolution, once a certain coupling is reached, the $\omega_{\max}$ truncation begins to decrease from 2π , causing the effective support of the wavefunction over $\omega_1 - \omega_2$ to change at a different rate with g , which is signified by the transparent lines (these lines are not shown for the bottom plot). b) Non-stabilizerness in degenerate subspace $\mathcal{S}_{12}$. The states that have the minimum non-stabilizerness in this subspace are separable, while the states that have the maximum non-stabilizerness can be characterized by their properties under $\omega_1 - \omega_2$ exchange. c) Convergence of non-stabilizerness with increasing digitization resolution for first ten even eigenstates. d) Time evolution of $\mathcal{M}_2$ for different values of gauge coupling (top) and different levels of Pauli truncation ϵ in Hamiltonian (bottom). In the top plot, multiple digitization resolutions ranging from $n_q = 3$ to $n_q = 6$ are shown for each value of the coupling.	79

- 4.3 *Walsh hierarchy of smooth wavefunctions* **a)** Different Walsh bands for two-dimensional function of ω_1, ω_2 (top) with Walsh indices and associated coefficients (below). The lowest, most significant, bands are highlighted in blue, and the next most significant are shown in red. The strategy to reconstruct the wavefunction involves starting from the shells closest to the origin and working out in each direction. Once a shell is filled, the next shell starts being filled. **b)** Reconstructed wavefunction probability density for fixed ω_2 at different Walsh cutoffs k^* . The solid black lines correspond to exact reconstruction where the cutoff is removed. Both a strong coupling wavefunction at $g = 2$ and a weak coupling wavefunction at $g = 0.6$ are shown. **c)** Number of Walsh bands needed to reconstruct distribution $p(\omega_1, \omega_2)$ up to fixed error $\delta = .1/.04$ (top/bottom) as function of n_q . The number of bands in each case increases linearly, corresponding to an exponential growth in the number of required Walsh coefficients until a plateau is reached. The different points in the plot denote the values of g for which the ground state wavefunction was calculated. 87
- 4.4 *Quantum measurement of non-stabilizerness* **a)** Bitstring distributions for specific instance of random Cliffords. The raw data (left) shows small populations of many bitstrings whose populations are spuriously amplified by the noise of the quantum device. The bitstring distribution after application of BSDR (right) removes these spurious populations and restores much of the probability mass to the true bitstring. **b)** Values of non-stabilizerness over time. The solid lines correspond to evolution due to the exact propagator and the Trotterized circuit. The circles represent the calculations from the raw data and the diamonds represent the mitigated data. The raw data is almost constant, and the features of the dynamics are highly suppressed. The mitigated values are able to accurately track the real time dynamics for all times studied. 91
- 5.1 *Sample-efficient measurements of gauge-invariant observables in $\mathbb{Z}_2$ LGT.* Illustration of $\mathbb{Z}_2$ LGT in $(2 + 1)$ D and its Ising dual. For the LGT qubits reside on links l of the lattice and Gauss law constraints G_s are associated with sites s . The Hamiltonian consists of plaquette (magnetic) terms $W_{\square} = \prod_{l \in \square} \sigma_l^z$ and on-site (electric) terms σ_l^x . In the dual description as a $(2 + 1)$ D Ising model the qubit degrees of freedom are associated with the plaquette. There is a one-to-one correspondence between operators and states within the physical Hilbert space of the LGT, independent of boundary conditions: the magnetic terms map to Pauli- Z operators on the Ising qubits and the electric terms map to two-body Pauli- X terms acting on the qubits across the relevant link. In the inverse map from the Ising model to the LGT, single X operators on the Ising side of the duality map to a string of σ^x operators along some path γ to an arbitrary reference plaquette $\square_0$. For PBC, the dual Ising model obeys a global parity constraint and such single-body X operators are unphysical. 94

5.2	<i>Symmetry respecting randomization protocols</i> We introduce three symmetry-aware random measurement protocols—Global and Local Dual Pairs and Dual Product—designed to estimate gauge-invariant observables efficiently, leveraging prior knowledge that expectation values of gauge-variant operators vanish. For Global Dual Pairs, random pairs of Ising qubits are chosen and random two-body symmetry-respecting unitaries associated with these degrees of freedom are mapped back to the LGT side of the duality as the randomizing operations prior to measurement. Local Dual Pairs is similar but leverages a promise of geometric locality in the observables of interest to limit the choice of pairs to local patches. Finally, the Dual Product protocol is constructed from the standard Product Protocol applied to the dual Ising model. This protocol slightly breaks symmetry when we have periodic boundary conditions (PBC) since such operations are not parity-respecting. A trade-off exists between sample complexity and circuit depth.	95
5.3	<i>Schematic overview of the Global Dual Pairs protocol with periodic boundary conditions (PBC).</i> Step 1. A pairing, π , of dual lattice sites is chosen uniformly at random. For each pair, a random unitary $U_{[ij]}$ is constructed in the Ising formulation. Step 2. For each pair $[ij] \in \pi$, the dual unitary $\mathcal{U}_{[ij]} = \Phi^{-1}[U_{[ij]}]$ is implemented on the LGT side of the duality and a computational basis measurement is performed. Step 3. The output bit string is mapped back to the dual Ising theory, where efficient shadow channel inversion is performed to estimate expectation values of gauge-invariant observables O	99
5.4	<i>LGT–Ising duality for mapping a (parity-respecting) entangling operation $\Phi^{-1}(Y_i X_j)$.</i> While this operation is two-body on the Ising side (left panel), it corresponds to an extended operation in the LGT (right panel) acting between plaquettes i and j along a path γ_{ij} , which is arbitrary except for its fixed endpoints.	103
5.5	<i>Overview of the Local Dual Pairs protocol.</i> The Local Dual Pairs protocol is a variant of our scheme for geometrically local gauge-invariant observables that have support on a patch of at most $L \times L$. Step 1 consists of picking a random tiling $\tau \in \mathcal{T}$ and, then, picking a random pairing repeated within each patch of the tiling in order to construct an associated random unitary U_π . Step 2 consists of applying the associated unitary $\mathcal{U}_\pi$ on the LGT side of the duality, followed by measurement in the computational basis. Step 3 is the post-processing step, which proceeds essentially identically to the Global Dual Pairs protocol restricted to a given patch of a tiling, where the data is filtered to a single choice of tiling according to the support of a given target variable.	107
5.6	<i>An ancilla enables mixing PBC (even parity) and tPBC (odd parity) sectors.</i> a) Adding an ancilla a to a reference link r allows us to enforce either PBC or tPBC in the $\mathbb{Z}_2$ LGT. b) In the Dual Product protocol, the z -component of the state of the link r is “copied” to the ancilla a via a CNOT. The randomizing unitaries, chosen according to the usual Product Protocol, but applied to the Ising degrees of freedom, then mix between the PBC and tPBC sectors on the LGT side of the duality. In the right panel, we show how the plaquette and Gauss law operators change with the ancilla a : The plaquette operators no longer share a link, while the Gauss law operator is effectively extended by a fifth leg, $\sigma_r^x \rightarrow \sigma_r^x \sigma_a^x$	110

5.7	<i>Schematic overview of the Dual Product protocol with periodic boundary conditions (PBC).</i> Step 1. Randomly choose a unitary U_{prod} consisting of a single qubit Clifford gate on each Ising qubit. Step 2. The dual unitary $\mathcal{U}_{\text{prod}} = \Phi^{-1}[U_{\text{prod}}]$ is implemented on the LGT side of the duality and a computational basis measurement is performed. Step 3. The output bit string is mapped back to the dual Ising theory, where efficient shadow channel inversion is performed to estimate expectation values of gauge-invariant observables O	111
5.8	<i>Numerical demonstration.</i> a) Examples of a ribbon (blue) and loop (red) operator on both the $\mathbb{Z}_2$ LGT and Ising side of the duality. b) Estimates of ribbon (of maximal length) and loop operator (two plaquettes) expectation values on a minimal lattice of size $N_x = 3, N_y = 2$ with PBC via Global Dual Pairs, Dual Product, and Product protocols, compared to the exact solution for the ground state of the $\mathbb{Z}_2$ LGT as a function of the coupling g . Error bars represent the standard deviation of the median of means estimator over 50 experimental repetitions of $\nu = 1000$ shots each. c) Average relative error ϵ_{avg} (see Eq. (5.32)) versus number of samples ν for the same loop and ribbon operator as in panel (b) for the Product, Dual Product and Global Dual Pairs protocols for the ground state corresponding to coupling $g = 0.5$. d) Average relative error ϵ_{avg} versus volume for loop and ribbon operators of fixed weight and operators with extensive weight for the Global Dual Pairs and Dual Product protocols. In particular, for the ribbon operators the extensive weights are 2,4,6, and 8 and for the loop operators they are 2,4,4, and 6, for systems sizes $V = 4, 6, 7, 10$. The average is taken over 50 experimental repetitions of $\nu = 1000$ shots each using the ground state of the $\mathbb{Z}_2$ LGT with coupling $g = 0.5$. The sample cost of the Global Dual Pairs protocol depends both on the operator weight in the dual Ising theory and the size of the lattice, whereas the Dual Product protocol depends only on operator weight.	116
5.9	a) An illustration of an example loop operator of the sort considered in the discussion in Section 5.6.2. b) An illustration of an example (set of) ribbon operators of the sort considered in the discussion in Section 5.6.2. The way in which one picks paths to connect the end points is arbitrary, as all choices are equivalent up to the action of Gauss law operators.	118
6.1	<i>Non-Gaussian state reconstruction.</i> a) Few-body correlation measurements (e.g. four-point correlations $C^{(4)}$) on a subsystem A of a fermionic quantum simulator are used to reconstruct the reduced state ρ_A . Besides single-site resolved imaging, this requires pair-wise tunneling operations $\mathcal{U}_{kl}$ between sites k and l . After classical post-processing, the protocol gives access to the entanglement spectrum ε_i . b) The state reconstruction is based on an expansion around Gaussian states ρ_g , which form a subset of integrable states in the Hilbert space $\mathcal{H}$. To probe the onset of quantum chaos, we go to non-Gaussian states $\tilde{\rho}$ by including higher-order correlations. c) Four-point correlations lead to non-Gaussian matrix elements $\langle \mathbf{m} \rho_A \mathbf{n} \rangle$ with Hamming distance $d_H(\mathbf{m}, \mathbf{n}) = 0, 2, 4$ (top, middle, bottom).	125
6.2	<i>Elemental Diagrams.</i> a) Propagator diagrams carrying index p . b) Source terms, represented by diamonds carrying index p with the corresponding occupations number m_p, n_p . c) k -point vertex carrying indices $i_1, \dots, i_k$	133

6.3	<i>Construction of I_1 and I_2 Interaction Diagrams.</i> Starting from the Gaussian diagram as the generating functional (top), 4-point vertices are introduced. The diagram for I_1 is constructed on the left and the diagram for I_2 is constructed on the right following the rules discussed in this section.	134
6.4	<i>Diagrams contributing at $O(\tilde{C}^{(4)})$.</i> a) I_1 : Propagators connect repeated indices of the vertex, all sources are paired up. b) I_2 : One propagator connects one set of repeated indices of the vertex, the other two end in sources. Some sources can be paired, others remain unconnected. c) I_3 : All propagators connect to sources.	135
6.5	<i>Non-equilibrium evolution of reconstructed state.</i> a,b) Non-Gaussianities $\theta(\sigma)$, see Eq. (6.30), build up analogously for the exact state ρ_A [(a)] and the reconstructed state $\tilde{\rho}$ [(b)], as indicated by the dashed line which shows a linear fit to the early time behavior $\theta(\rho_A)$. Here we show numerical data for a system of size $N = 10$, $N_A = 5$ with 10 different initial states (shaded lines) and their averages (solid line). c) We test the accuracy of the state reconstruction through the spectral error of levels i of the entanglement spectra $\{\tilde{\epsilon}\}$, $\{\epsilon_g\}$ for different interaction strengths U (at a fixed time). The non-Gaussian parameterization reduces the error to $\mathcal{O}(U^2)$ relative to the Gaussian parameterization with error $\mathcal{O}(U)$	138
6.6	<i>Evolution of level statistics of non-Gaussian state parameterization</i> a,b) The gap ratio distributions of both, the parameterization $\tilde{\rho}$ (filled, purple) and the exact state ρ_A (unfilled, gold), evolve from Poisson (pink dashed) at time $\tau = 0.015$ to GUE (black dashed) at time $\tau = 0.47$, thus signaling the onset of quantum ergodic evolution. c) The average gap ratio $\langle r \rangle$ grows with time and saturates to the value ~ 0.6 indicating level repulsion (GUE). Data is shown for a system of size $N = 10$, $N_A = 5$, interaction strength $U = 5.6 \times 10^{-3}$; averages are taken over 10 different initial states, where shaded regions signify 95% confidence intervals. The dashed gray lines signify the times displayed in (a) and (b).	139
6.7	<i>Average gap ratio of exact state and non-Gaussian parameterization</i> a) The average gap ratio of the exact state over a range of times Jt and interaction strengths U/J . The black dashed line shows the crossover from a Poisson distribution to GUE, indicating quantum chaos through level repulsion. b) The average gap ratio for the parameterization qualitatively reproduces that of the exact state over the times and interaction strengths studied.	140
6.8	<i>Self-similarity of non-Gaussian correlations.</i> a) Example of a “density-density” correlation, $\tilde{C}_{ijji}^{(4)}$ with $i \neq j$, for the same interaction strengths as shown in Fig. 6.5. For $U = 10^{-3}, 10^{-2}$, the correlations collapse onto the same curve when rescaled by the interaction strength. As $U \sim J$, J' oscillations begin to appear on top of this curve that breaks this behavior. b) Example of a “pair-density” correlation function $\tilde{C}_{ijki}^{(4)}$ with $i \neq j \neq k$. c) Example of a “pair-pair” correlation function ($i \neq j \neq k \neq l$).	140
6.9	<i>Rank of the state.</i> Growth of rank for a system of $N = 10$ during the free evolution averaged over 30 random initial states. The pink region shows the range of growth rates over the averaged states.	145
6.10	<i>Position space initial state.</i> a) The average gap ratio of the exact state with a position space initial state. b) The average gap ratio $\langle r \rangle(t)$ predicted by the parameterization.	146

- 6.11 *Shot noise analysis.* **a)** The build-up of level repulsion for $N = 8$, $N_A = 4$ averaged across 4 initial states. The gold line represents the exact result, the purple line represents the ansatz with exact four-point functions, and the blue dots denote the ansatz measured with finite shots; the shading indicates the number of samples used. **b)** An example of four-point function (points) compared to the exact four-point function (continuous curve). The error bars are given by the standard deviation from the mean and the number of shots are $N_S = 5 \cdot 10^4$. . 150
- 7.1 *Error Detection in Quantum Simulations on Heavy-Hex Topology.* **a)** The stabilizers and logical operators that act on data qubits d_0 - d_3 in the $[[4, 2, 2]]$ error-detecting code. The Z -stabilizer is shown in red and the X -stabilizer is shown in blue. Logical operators are grouped into commuting sets represented by the color. **b)** The placement of code blocks (grey) onto the `ibm_boston` quantum processor. Each code block has four data qubits (black) and two ancillas for Z - and X -stabilizer extraction (red and blue respectively). Bridge qubits between blocks used for routing are shown in green. **c)** Circuit operation for encoded quantum simulations. Dashed outlines indicate nFT operations, while solid outlines indicate FT gadgets. Each code block starts in the “compute” configuration (shown in the left box in the center). Within each code block, the logical $|\overline{00}\rangle$ state is prepared with a nFT circuit. Non-FT logical operations (orange boxes) are performed in the compute configuration, example circuits of which are shown at the bottom. The blue boxes represent FT syndrome extraction done in the “syndrome” configuration within each block (right two boxes), which is implemented with the circuit shown in the bottom. Applications of logical operations and syndrome extraction are repeated R times in any given computation. 158
- 7.2 *Encoded Simulations of Dynamics in the 1D Ising Model.* **a)** The realized logical connectivity grid by the block placement in Fig. 7.1b). Logical qubits (b, l_b) (grey) are joined by intra-block gates (hatched orange) and inter-block gates (orange). The black outline traces the gates used to simulate a 1D chain. **b)** The energy density $\langle \overline{H}_n \rangle$ computed from MPS simulations as a function of time t and logical qubit $n = (b, l_b)$ for $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). A size-three true-vacuum bubble located at the center is evolved with $n_T = 500$ Trotter steps of size $\delta t = 0.1$. The black rectangle shows the t for which device results are compared in c). **c)** Comparison of encoded (blue) and unencoded (grey) magnetization $\langle \overline{Z}_n \rangle$ from `ibm_boston` to MPS simulations (black) at $t = 8$, corresponding to 16 Trotter steps of $\delta t = 0.5$ for $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). The encoded simulations shown have $R = 1$ syndrome extraction rounds. **d)** The cumulative distribution function (CDF) of the per-qubit, per-time absolute error $|\langle \overline{Z} \rangle_{\text{meas}} - \langle \overline{Z} \rangle_{\text{MPS}}|$ for each value of g_z . The median error (dotted vertical lines) and CDFs corresponding to one layer ($R = 1$) of syndrome measurements are shown compared to unencoded results (solid and dashed lines respectively). Results up to $t = 4$ are used, where an encoding improvement is observed. **e)** The signal-survival factor α , defined in Eq. (7.7), as a function of t for various R (colored lines) and each value of g_z compared to unencoded results (dashed grey line). The shading for $t = 4$ -8 represents times excluded for the CDFs shown in d). The insets show the residual error σ in the fit. The statistical uncertainty is determined through bootstrap resampling. . . 160

7.3	<i>Encoded Simulations of the 2D Ising Model.</i> a) The magnetization $\langle \bar{Z}_n \rangle$ of a 3×3 true-vacuum bubble as a function of time t for $g_z = 0.5$ (top row) and $g_z = 1.5$ (bottom row), computed from MPS simulations with $\delta t = 0.1$. The black solid and dashed rectangles show cross-sections of the grid at $t = 4$ for which data is plotted in b). b) Encoded (blue) and unencoded (grey) results from <code>ibm_boston</code> compared to MPS results (black) for the horizontal (solid) and vertical (dashed) cross-sections outlined in a), as a function of logical qubit number $n = (b, l_b)$. The encoded data has $R = 1$ syndrome measurement rounds. c) Cumulative distribution functions (CDFs) of the per-qubit, per-time absolute error $ \langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}} $ for both values of g_z . The median error and CDFs corresponding to encoded runs with $R = 1$ syndrome measurements and unencoded runs are shown compared to unencoded results (solid and dashed lines respectively). d) The signal-survival factor α as a function of t for various R (colored lines) and each value of g_z , compared to unencoded results (dashed grey line). The insets show the residual error σ in the fit. The statistical uncertainty is determined through bootstrap resampling.	161
7.4	<i>Encoding Improvement in α.</i> Percent improvement in α of encoded over unencoded results as a function of time t for 1D and 2D simulations. Results with $R = 1$ syndrome extraction rounds for the melting regime ($g_z = 0.5$) are shown in shades of red, and the localized regime ($g_z = 1.5$) is shown in shades of blue. The shading represents statistical uncertainty determined through bootstrap resampling.	164
A.1	<i>Error scalings</i> Scaling of the contributions $\mathcal{E}$ to the error from different sources in a single Trotter step of the pulse sequences $\mathcal{S}_{1/2}$ and $\mathcal{S}_1$ as a function of system geometry. N_y is varied on the x axis and each line represents the contribution to the error with a different value of N_x . Here the analytic bounds of Eqs. (A.38) - (A.41) are compared to values obtained by numerically evaluating the norm of Eq. (2.15) for each type of contribution. Note that in the latter three panels, the $N_x = 1, N_y = 1$ points are much lower than the rest of the $N_x = 1$ lines because of the fact that there are three-body contributions to the error for these terms; when the system has less than three sites, these contributions are absent, leading to a much smaller error.	176
A.2	J_1, J_2, J_3 sums evaluated for explicit array configurations with fits. The extent of the array in the x direction is $N_x = 20$ atoms. While there is a deviation for smaller array sizes, a linear scaling is still seen.	178
D.1	<i>Summary of LGT-Ising duality for PBC and FBC.</i> a) Hilbert space structure and operators for $\mathbb{Z}_2$ LGT with PBC and b) for the corresponding Ising dual which is subject to a parity constraint. Spheres mark one local spin-1/2 d.o.f. c) For FBC, the LGT includes bulk and boundary Gauss laws, as well as plaquette and electric operators. d) The dual Ising theory for FBC is no longer parity constraint but electric operators dualize differently in bulk and on the boundary.	200

D.2	<i>Implementation of Global Dual Pairs with Fixed Boundary Conditions</i> Example of a) loop and b) ribbon operators for $N_x = 3$ and $N_y = 3$ with FBC via the modified Global Dual Pairs protocol, compared to the exact solution of the $\mathbb{Z}_2$ LGT ground state. Error bars represent standard deviation of the mean over a single experiment of $\nu = 5000$ shots.	202
E.1	<i>Convergence of ORP on ibm_boston.</i> Four Trotter steps of size $\delta t = 0.4$ are applied to four code blocks, each initialized in the $ 00\rangle$ state, with couplings $g_x = g_z = 1$. The solid line shows $\langle \bar{Z}_{(0,0)} \rangle$ computed through ORP as a function of the number of detectors k used in the postselection criterion, with the noiseless MPS prediction given by the horizontal dashed line. The corresponding acceptance fraction is shown by the dotted line (right axis). The vertical red dotted line marks the value of k^* chosen by the plateau-finding algorithm.	206
E.2	<i>Example $[[4, 2, 2]]$ Code Block Placement onto Heavy-Hex Connectivity.</i> Several blocks (grey) are shown placed onto a heavy-hex connected quantum processor. Data qubits are given in black, Z - and X -type ancillas are shown in red and blue respectively, and a bridge qubit is shown in green.	208
E.3	<i>Encoded Trotter Step Circuits.</i> a) An example placement of two blocks in the compute configuration. Data qubits are given in black, Z - and X -type ancillas are shown in red and blue respectively, and a bridge qubit is shown in green. b) The CNOT fanout necessary for implementing a $\bar{R}_{ZZ}$ between logical qubits $(17, 0)$ and $(20, 1)$ (dark orange), with the associated SWAP gate converted to CNOTs. c) A Trotter step involving logical qubits in blocks 17 and 20, with the inter-block gate between qubits $(17, 0)$ and $(20, 1)$. Single-logical-qubit gates (yellow) and intra-block $\bar{R}_{ZZ}$ (light orange) are folded into the CNOT fanout required for the inter-block $\bar{R}_{ZZ}$ gate, the last part of which is shown in dark orange. Block 20 exits the Trotter step in frame $f = 1$	210
E.4	<i>Circuits for Switching Between Compute and Syndrome Block Configurations.</i> a) The compute and syndrome block configurations as shown in the center of Fig. 7.1c). Data qubits are shown in black, Z -type ancillas are shown in red, and X -type ancillas are shown in blue. b) The circuit to switch from the compute to the syndrome configuration. c) The corresponding circuit to switch from the syndrome to the compute configuration. The syndrome extraction circuit applies a $\text{SWAP}(d_2, d_3)$, which is taken into account in the configuration switching circuit.	213
E.5	<i>Comparison of Δ_j and β_j Predictions on ibm_boston.</i> The metrics Δ_j and β_j are calculated using Eq. (E.8) and Eq. (E.11) respectively. The two solid lines show the predicted value of $\langle \bar{Z} \rangle$ as a function of the number k of detectors used in postselection for Δ_j (dark) and β_j (light). The horizontal dashed line is the noiseless MPS prediction. The vertical dashed lines correspond to the value of k^* returned from the plateau finding algorithm using Δ_j (dark) and β_j (light). The expectation value is from four Trotter steps on four code blocks.	223

E.6	<i>Depiction of Most Important Detectors</i> Calculated for the logical observable $\bar{O} = \bar{Z}_{(b=0,0)}$, as determined by the rankings calculated in Table E.4. Highlighted in red are the locations of the seven most important detectors, as determined by both ORP and Pauli-propagation. The grey region shows the observable's circuit backwards lightcone, containing all of the detectors that are connected to the observable through gates. The dark blue region captures the physical reverse-lightcone of $\bar{O}$, showing the detector locations where the back-propagated observable has the largest Pauli weights.	225
E.7	<i>Logical Error Rate of Different Circuit Components</i> a) The different FT and nFT circuit elements whose logical error rates are considered in App. E.7.1. The green elements denote initialization gadgets, labeled FT (solid) and nFT (dashed). The bottom two circuits $\bar{R}_{\text{FT}}/\bar{R}_{\text{nFT}}$ are FT/nFT initialization gadgets followed by a nFT logical rotation $\bar{R}_{ZZ}$. b) The LER of the two different initialization gadgets. The grey region is where nFT crosses to a lower LER than the FT procedure. The grey and black dashed lines represent $\mathcal{O}(p)$ and $\mathcal{O}(p^2)$ scaling for the logical error rate p_L . c) The LER for the FT and nFT initialization followed by a logical rotation. d) The LER of the syndrome extraction circuits for multiple rounds. e) The LER of syndrome extraction as a function of number of rounds for different p . 227	227
E.8	<i>Memory Experiment Performance of Multiple Rounds on ibm_boston</i> a) The Total Variation Distance (TVD) from the logical $ \bar{00}\rangle$ state over 60 rounds of syndrome extraction. b) The acceptance fraction over 60 rounds. c) Suppression of errors at leading order as a function of number of syndromes used for postselection.	230
E.9	<i>Ancilla vs. No-Ancilla Rotations.</i> a) The top row shows ancilla-free circuits G_2 and G_4 implementing $e^{-i\pi/8\bar{Z}_{(b,0)}\bar{Z}_{(b,1)}}$ within a single code block (left) b and $e^{-i\pi/8\bar{Z}_{(a,0)}\bar{Z}_{(b,0)}}$ between two code blocks a and b (right). The bottom row shows circuits G'_2, G'_4 with ancilla qubits that implement the same rotations. b) The LER p_L as a function of the physical error rate p for the circuits in a). The logical error rate is computed by taking the average over Haar-random states under repeated rotations with depolarizing noise.	230
E.10	<i>Postselection Performance and Signal Survival in Encoded Simulations.</i> a) The signal-survival factor α as a function of logical qubit over all t for $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). Encoded results with $R = 1$ are given in blue, and unencoded results are shown in grey. Results for the 1D chain are shown in the left column, and 2D grid results are shown in the right column. The shading represents one standard deviation of the α distribution over qubits. Error bars on individual points represent uncertainty determined through bootstrap resampling. b) Convergence of an example observable on qubit $(14, 1)$ at $t = 5$ as a function of the number of detectors kept for postselection k . The shading represents uncertainty from bootstrap resampling. Results from the unencoded runs and from MPS are given by grey and black dashed horizontal lines. The vertical black dotted line shows the k^* selected by the plateau finding algorithm. Acceptance rates for each k are shown by the dotted grey line. . .	232

E.11	<i>Comparison of α for Corner and Bulk Qubits in 2D Simulations.</i> The signal-survival factor α is shown for encoded runs as a function of time t for R syndrome extraction rounds (colors) compared to unencoded runs (grey). The left panel shows α for corners of the grid (logical qubits with two neighbors), and the right panel shows α for bulk qubits (those with more than two neighbors). The insets show the error in the α fits, σ . Results from the localized regime ($g_z = 1.5$) are used.	233
E.12	<i>Group Velocity in the 1D and 2D Ising Model.</i> a) The group velocity $v(k)$ as a function of momentum k for $g_z = 0.5$ (blue) and $g_z = 1.5$ (red) calculated in a 20-qubit 1D chain with periodic boundary conditions. b) The group velocity $ v(k_x, k_y) $ for a 5×5 2D lattice with $g_z = 0.5$ (left) and $g_z = 1.5$ (right). Dispersion relations are calculated using exact diagonalization, and a spectral derivative is used to compute $v(k_x, k_y)$	234
E.13	<i>Effective Inter-Block Coupling Under Residual Ancilla Sign Flips on ibm_boston.</i> a) The effective inter-block coupling $\langle J_{\text{eff}}/J \rangle$ as a function of time t averaged over all inter-block connections in the 1D simulations shown in Fig. 7.2 for $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). Different R values indicate amount of syndrome measurements. b) The same as a) but for 2D results presented in Fig. 7.3.	236
E.14	<i>Encoding Improvement Without Postselection in 2D Simulations.</i> a) Cumulative distribution functions (CDFs) of the per-qubit, per-time absolute error $ \langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}} $ for both values of g_z . The median error and CDFs corresponding to $R = 1$ syndrome measurements are shown compared to unencoded results (solid and dashed lines respectively). b) The signal-survival factor α as a function of t for various R (colored lines) and each value of g_z , compared to unencoded results (dashed grey line). The insets show the residual error σ in the fit.	237
E.15	<i>Classical Simulations of False-Vacuum Decay Dynamics in 1D.</i> An initial true-vacuum bubble of size one, two, three (columns) at the center of the 42-qubit chain is evolved with $g_z = 0.5$ (top row) and $g_z = 1.5$ (bottom row). The color in the heatmap shows the energy density $\langle \bar{H}_n \rangle$ as a function of logical qubit position $n = (b, l_b)$ and time t . The time evolution uses $n_T = 500$ Trotter steps of size $\delta t = 0.1$ and is computed with MPS.	240
E.16	<i>Classical Simulations of False-Vacuum Decay Dynamics on a 2D Lattice with Missing Edges.</i> An initial true-vacuum bubble of size 1×1 , 2×2 , 3×3 (rows) at the center of the 42-qubit lattice is evolved with $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). The qubit colors show the magnetization $\langle \bar{Z}_n \rangle$ as a function of logical qubit position $n = (b, l_b)$ and time t (columns). The time evolution uses $n_T = 100$ Trotter steps of size $\delta t = 0.1$ calculated with MPS. The logical connectivity grid corresponds to Fig. 7.2a) with all available edges used.	241
E.17	<i>Classical Simulations of False-Vacuum Decay Dynamics on a Uniform 2D Lattice.</i> An initial true-vacuum bubble of size 1×1 , 2×2 , 3×3 (rows) at the center of the 5×5 lattice is evolved with $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). The qubit colors show the magnetization $\langle \bar{Z}_n \rangle$ as a function of logical qubit position $n = (b, l_b)$ and time t (columns). The time evolution uses $n_T = 500$ Trotter steps of size $\delta t = 0.1$ and is calculated with statevector simulations.	242

E.18	<i>Block Placement and Logical Connectivity Graph in the Absence of Faulty Qubits and Gates.</i>	
	a) An example of a block placement (grey) of 21 blocks assuming no faulty qubits and gates on a heavy-hex quantum computer. Within each block, the data qubits are shown in black, Z -ancillas are shown in red, X -ancillas are shown in blue, and bridge qubits are shown in green. b) The resulting connectivity between logical qubits (grey) optimized for similarity to a square lattice. Thick orange hatched edges represent intra-block connections and thin orange edges show inter-block connections.	244
E.19	<i>Overlapping Block Placement.</i> Code blocks are split into two groups (red and blue) for serialized syndrome extraction. Data qubits within each block are shown in black and shared ancilla qubits are shown in green.	245

LIST OF TABLES

Table Number		Page
2.1	The error rates, step size, and step parameters for each method to simulate H_{XXX} evolution. The three rightmost columns show the breakdown of τ_U into the evolution due to the rotations t_{rot} , the evolution to compensate for the rotations t_{comp} , and the optimal optional evolution time t_* . t_* is given by Eqs. (2.17) and (2.21), and following similar analysis for the other Trotter sequences. For $\mathcal{C}_1$ the step size corresponds to the Floquet period $\tau_{\mathcal{C}_1} = 4\epsilon$. The essentially identical properties of $\mathcal{S}_1$ and $\tilde{\mathcal{S}}_1$ and the better error rate scaling of $\mathcal{S}_2$ at the cost of longer device time can be seen.	35
3.1	Estimated and exact values of $\langle H_B \rangle$ as a function of time.	71
4.1	Raw, mitigated and Trotter values of $\mathcal{M}_2$ as a function of time. The raw values are determined by using the bitstring distributions that are read directly off of the device without any post-processing.	90
5.1	<i>Summary of costs for estimating the expectation value of a gauge-invariant observable O with $\ O\ _\infty = 1$ to an additive error ϵ.</i> We assume that O is k -local on the LGT side of the duality or k_{dual} -local on the Ising side and, also, assume that $k_{\text{dual}} \in \Theta(1)$. Recall that $V = N_x N_y$ is the size of the (dual) lattice and, for the Local Dual Pairs protocol, L is the linear size of the geometrically local patches considered. Circuit depths for implementing the random rotations in the Dual Pairs and Dual Product protocols are based on the standard decomposition of Pauli rotations into two qubit gates and assume no parallelization of the randomizing unitaries.	115
B.1	Terms in the Pauli string decomposition of $H_{\square\square}$ that are retained in the circuits. The colored circuit elements correspond to those in the main text. The Pauli string column denotes the rotations $e^{-ic_p P}$ in the circuit for each string P with corresponding coefficient c_p . The truncation percentage δ denotes the fraction of Pauli strings that would need to be truncated for this term to have the smallest coefficient. As an example, all terms in this table were included in the circuit which was run at $\delta = .66$ and so all δ values in this table are greater than .66.	185
B.2	The columns are the same as in the Table B.1. This table (continued in Table B.3) contains all of the Pauli string in the decomposition $\mathcal{D}(H_{\square\square})$ that were dropped in the quantum simulations. As such all entries in the δ column are less than $\delta = .66$. Note that all of the terms dropped are included in the yellow circuit element. They all have either an X or a Y located on the ν register. Additionally, they are generally much higher weight and have support on a greater number of qubits.	186

B.3	Continuation of Table B.2.	187
B.4	Energies and explicit eigenstate constructions for the seven lowest-lying eigenstates of the two plaquette Hamiltonian in the strong coupling limit. For the states that are degenerate in this limit, the magnetic Hamiltonian breaks the degeneracy and splits the states.	189
E.1	The inter-block R_{ZZ} CNOT fanout depth for pairs of logical qubits (l_a, l_b) between pairs of code blocks (a, b) corresponding to the block placement shown in Fig. 7.1b). The depth includes CNOTs used for SWAPs and bridge qubits where necessary.	211
E.2	<i>Details of 1D Chain Simulations on ibm.boston at $t = 8$.</i> The columns compare unencoded counts to counts in encoded runs with various numbers of syndrome extraction rounds R . The second row shows the total number of measurements (both mid-circuit and terminal). The maximum coherence time used by measurements on any qubit in units of $dt = 4 \times 10^{-9}$ s is given in the third row. The fourth row gives the total duration of the circuit. Depths, total numbers of two-qubit gates, and numbers of shots executed are given in the fifth, sixth and seventh rows respectively	212
E.3	<i>Details of 2D Grid Simulations on ibm.boston at $t = 8$.</i> The same rows and columns as Table E.2 but for the 2D simulations presented in Sec. 7.3.	212
E.4	<i>Comparison Between Detector Ranking Using Pauli-propagation and ORP.</i> The top row shows the position in the ranking, where 1 is the highest ranking and 10 is the lowest ranking among the detectors shown. The middle row shows the detectors $v_{(b,r)}$ ranked by their Δ_j value and the bottom row shows them ranked by their γ_j value. The rankings are computed for the observable $\bar{Z}_{(b=0,0)}$ using a four-block circuit consisting of $n_T = 4$ Trotter steps with step size $\delta t = 0.1$, and one syndrome extraction round per Trotter step, starting from an all $ 00\rangle$ initial state. The rankings are computed for $g_x = g_z = 1$	224
E.5	<i>Magnetization in 1D Simulations at $t = 8$.</i> For each logical qubit (b, l_b) (first column), $\langle \bar{Z}_n \rangle$ from MPS results (second and fifth columns) is compared to unencoded results (third and sixth columns) and encoded results with $R = 1$ syndrome extraction rounds (fourth and seventh columns). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.	246
E.6	<i>Absolute Error in 1D Simulations.</i> The absolute error $ \langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}} $ by decile (first column) is compared between unencoded runs (second and seventh columns) and encoded runs for various numbers of syndrome extraction rounds R (columns 3-6 and 8-11). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.	247
E.7	<i>Signal-Survival Factor for 1D Simulations.</i> The signal-survival factor α calculated as a function of time t (first column) in unencoded results (second and seventh columns) is compared to encoded results for various numbers of syndrome extraction rounds R (columns 3-6 and 8-11). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is the residual error σ in the α fits.	248

E.8	<i>Magnetization in 2D Simulations at $t = 4$, first half.</i> For each logical qubit (b, l_b) (first column), $\langle \overline{Z}_n \rangle$ from MPS results (second and fifth columns) is compared to unencoded results (third and sixth columns) and encoded results with $R = 1$ syndrome extraction rounds (fourth and seventh columns). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.	249
E.9	<i>Magnetization in 2D Simulations at $t = 4$, Continued.</i> The remaining logical qubits from Table E.8, following the same column structure. For each logical qubit (b, l_b) (first column), $\langle \overline{Z}_n \rangle$ from MPS results (second and fifth columns) is compared to unencoded results (third and sixth columns) and encoded results with $R = 1$ syndrome extraction rounds (fourth and seventh columns). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.	250
E.10	<i>Absolute Error in 2D Simulations.</i> The absolute error $ \langle \overline{Z} \rangle_{\text{meas}} - \langle \overline{Z} \rangle_{\text{MPS}} $ by decile (first column) is compared between unencoded runs (second and seventh columns) and encoded runs for various numbers of syndrome extraction rounds R (columns 3-6 and 8-11). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.	251
E.11	<i>Signal-Survival Factor for 2D Simulations.</i> The signal-survival factor α calculated as a function of time t (first column) in unencoded results (second and seventh columns) is compared to encoded results for various numbers of syndrome extraction rounds R (columns 3-6 and 8-11). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is the residual error σ in the α fits.	251
E.12	<i>Acceptance Rates of $R = 1$ Encoded Simulations.</i> The acceptance rates are shown for both 1D simulations (second and third columns) and 2D simulations (fourth and fifth columns) as a function of time t (first column). These encoded simulations use $R = 1$ syndrome extraction rounds, corresponding to the data reported in the main text. The uncertainty represents the standard deviation over all qubits for a given t . The uncertainty is the standard deviation over all qubits for each t	252

ACKNOWLEDGEMENTS

I must first acknowledge the senior researchers who played important roles in my development as a scientist. I want to thank my advisor, Martin Savage, for introducing me to the world of quantum simulation and quantum computing. He took a risk in bringing an unseasoned first-year graduate student into IQuS, and I hope I have made that substantial financial investment worthwhile. For his encouragement to start the Quantum Computing Club, for his wisdom on navigating the politics of the academy and the world more generally, and for his emphasis to work on things that are truly important, my perspectives have been indelibly changed. Next, I thank Niklas Mueller, who was the first to sit down with me, look through code, and work at a whiteboard. He taught me that science is not just about the results, but also about the story that communicates them. He was the first person who caused me to critically question my own research ideas, to slow down and ask “Why does this matter?”. Finally, I thank Dorota Grabowska, who is always willing to listen to my half-baked ideas and help turn them into something more fully fledged. It is rare and powerful to have someone who is willing to listen to everything you have to say, and I am grateful for that. Further, they are always willing to give advice not only about research, but about life more generally. I would also like to thank David Kaplan and Silas Beane for agreeing to serve on my reading committee after last-minute changes had to be made to its composition.

Next, I must acknowledge the members of IQuS, all who enriched my experience: Alessandro Roggero, Fabio Anza, Xiaojun Yao, Sebastian Grieninger, Stephan Caspar, Hersh Singh, Marc Illa Subina, Saurabh Kadam, Francesco Turro, Anthony Ciavarella, Ramya Bhaskar, Roland Farrell, and Sarah Powell - I learned so much from all of you. Marc Illa was a postdoc for the bulk of my time at IQuS, and I could always look forward to his smiling face when I came into the office. Anthony Ciavarella, the senior graduate student during my first few years, was always willing to answer my silly questions. Roland Farrell, another senior graduate student, was a wonderful role model and friend for us younger students. Katie Hennessy was an angel disguised as a group administrator and I’m incredibly fortunate that she decided to retire as I graduated and not before. To Nikita Zemlevskiy, Ivan Chernyshev, Jeremy Hartse, and Zhiyao Li: we all joined IQuS not

knowing what to expect, and having you alongside me to navigate the challenges and uncertainties of research in a new group made everything a little bit easier. Special mention must be made of Nikita Zemlevskiy - roommate and officemate, collaborator and commiserator. I've learned a lot from working alongside him and watching him grow as a researcher, and I'm glad we were able to navigate through grad school together.

I would like to acknowledge the financial support for the research enclosed herein from the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, InQubator for Quantum Simulation (IQUS) under Award Number DOE (NP) Award DE-SC0020970 via the program on Quantum Horizons: QIS Research and Innovation for Nuclear Science, and from the Department of Physics and the College of Arts and Sciences at the University of Washington.

To my quantum friends outside of IQUS. Christopher Kang - I had no idea that the undergrad asking for deadlift tips, who also happened to share research interests, would become a long-time friend. You are as precocious as you are refined, and though you are younger than me, I value your advice and perspectives. I still hold out hope that we will either write a paper together or coalesce our loose ideas into a successful venture pitch. Jordi Arnau Montaña López - despite your brief time in Seattle, it was always wonderful hanging out and learning about the mathematically rigorous side of quantum; I hope to visit you in Munich soon and always look forward to catching up over Zoom. I also benefited greatly from collaborations with Robert Ott, Torsten Zache, Jacob Bringewatt, and Andreas Elben. I learned a great deal from all four and I value the insights we were able to discover and share together. To the friends I made at the QSC Summer School in 2024 - Sohan, Adrian, Abdel, Mary, and Kevin - it was great getting to know you all at ORNL, and I am always happy when we reconnect at conferences.

Transitioning to my non-quantum friends, I must thank Murali Saravanan and Will Marshall (and Nikita). We shared many memories and experiences spanning multiple years and addresses in Seattle. I will always look back fondly on our adventures, and appreciate the collective patience with me for often forgetting to close the shower curtain after use and for using too much oil to cook. To Kent Wilson - one of my first roommates in Seattle, I'm glad we have been able to stay in touch despite moving down separate paths as the years move on.

In the summer of 2025, I interned at HRL Laboratories, working on noise modeling and tensor network simulations of their TN3 spin qubit device. It was an amazing experience that completely reframed my

perspective on quantum, due largely to the efforts of Daniel Volya, Chris Swank, Stephen Carr, Nic Ezzell, Liam O'Brien, Thaddeus Ladd, and the rest of the Quantum Theory and Simulation (QTS) group. I am grateful for how welcoming the group was, for the value they placed on learning and exploration, and for allowing me to fall back in love with research. While the actions of certain government agencies may have complicated my immediate return, I do not believe our story is over yet.

Finally, to my parents, who have been my emotional support system for much of my life and have helped me become the person I am today. My mother is kind, caring, and almost always willing to put up with whatever mood I am in. Knowing that she has my back regardless of the situation, but will also call me out and tell me when I am out of line, gives me great comfort. My father is equally caring, fiercely rational, and very good at grounding me in reality. His advice in life is always sound though somewhat aphoristic, except when it comes to picking stocks (to be clear: he is very good at it, but for some reason, any time I have listened, I have only picked losers). I measure my love for them as an ocean: impossible to quantify exactly, but vast, stretching as far as the eye can see.

...C, I'm on my way.

DEDICATION

To Garrett Odell, for igniting my curiosity about the world at a young age.

Chapter 1

INTRODUCTION

1.1 *Nature as a Computer*

One of the most exciting frontiers in fundamental science is that of complexity. It has many faces, from macroscopic, non-linear dynamical systems obeying the laws of classical physics, all the way down to microscopic systems that are governed by the laws of quantum mechanics. In every case it is what gives rise to the weird and wonderful phenomena observed in nature, and it is the same richness that makes those phenomena so challenging to calculate. Pen and paper calculations must usually give way to exascale computing clusters and multi-threaded GPUs to gain even the slightest degree of predictive power. There is no clearer example than the properties of quantum systems that contain more than a few interacting degrees of freedom, which in this work will be referred to as quantum many-body systems (QMBS).

Quantum many-body systems are ubiquitous in nature and have the capacity to describe systems ranging from materials science and chemistry [1,2], to the structure of the most extreme astrophysical environments like black holes [3]. In particular, interacting quantum fields, one of the most fundamental representations of nature, have natural representations as QMBS when regularized by a lattice. Quantum fields are the natural language for describing particle and nuclear physics, and many of the outstanding problems in these areas can only be addressed by solving Quantum Chromodynamics (QCD), the theory of the strong force [4–6]. QCD describes the interactions of quarks and gluons, determining the properties of nuclear matter. One of the theory’s most important features is that it exhibits *asymptotic freedom* [7], which means that at high energies the interaction strength is weak and the theory describes freely propagating quarks. In this regime, many properties of interest can be calculated through perturbation theory [8]. However, as the energy scale is lowered, the interaction strength increases and perturbation theory is no longer applicable. Quarks and gluons form bound states known as hadrons and require non perturbative methods to understand their structure. At this stage, the only *ab initio* way to make predictions about QCD is through Lattice QCD, which makes

use of Monte Carlo importance sampling of the QCD path integral [9–13]. Lattice QCD has been able to successfully predict quantities like decay rates, scattering amplitudes, masses and more [14–17]. However, for cases where one is interested in finite-density behavior or real-time dynamics, the field configurations that are sampled in Monte Carlo become complex and one encounters a sign problem [18].

It has been recognized for a long time that quantum many-body systems are challenging to simulate. While the reasons for this are manifold, one of the central issues is that the size of the quantum mechanical state space, known as the *Hilbert Space* and denoted by $\mathcal{H}$, grows exponentially with the number of constituents of the system. These constituents are generally comprised of *qudits*, which are quantum mechanical d level systems. For a system of N qudits, the dimension of the Hilbert Space is $\dim(\mathcal{H}) = d^N$. This thesis deals mostly with the case $d = 2$, known as a “qubit”, and d will be reserved to represent other quantities. For $N = 1$, all quantities of interest can easily be calculated by hand. As N grows, systems up to about $N \sim 30$ can be simulated on a conventional laptop computer. After this point, holding the entire statevector in memory requires petabytes of storage, and even the most powerful supercomputers can only reach $N = 48$ using exact matrix methods [19].

Having observed this scaling of the space requirements for simulating quantum systems, it is natural to wonder whether it is really necessary to model the entirety of the Hilbert space to make accurate predictions of physical observables. Sometimes it is not, and there exist sophisticated methods that can compress the relevant information of a quantum state, thereby giving access to larger system sizes. In many cases of interest, however, the answer is yes, and the reason lies in the fundamental structure of correlations enabled by quantum mechanics. One source of these correlations is a property known as *entanglement*, a feature that arises from the collective description of multiple quantum systems. As a simple example, take two Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, for systems A and B respectively. The joint Hilbert space is obtained by taking a tensor product $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$ such that the basis states of the joint space have the form $|\psi\rangle = |\psi_A\rangle \otimes |\psi_B\rangle$, where $|\psi_A\rangle \in \mathcal{H}_A$ and likewise for B . Despite the fact that basis states of $\mathcal{H}_{AB}$ decompose as tensor products, forming superpositions can lead to states in the joint space $\mathcal{H}_{AB}$ that cannot be decomposed as a product of the two subsystems. If A and B are both single qubits, an example of such an entangled state is the Bell state, defined as

$$|\Phi\rangle = \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle) \quad (1.1)$$

The correlations between the two subsystems carry uniquely quantum mechanical signatures, and have

been tested to high accuracy in experiments [20]. As the number of qubits grows, the set of entanglement structures that admissible states can express grows with it, and it is precisely this proliferation that resists the compression mentioned above. In these cases the exponentially large Hilbert space is truly needed to describe the many-body quantum mechanical phenomena at hand.

To understand why quantum computers are useful in this regard, one must step back to see how computational difficulty is quantified. **NP** (nondeterministic polynomial time) is the class of decision problems whose proposed solutions can be verified in polynomial time by a classical computer, even when finding a solution may take far longer. The traveling salesman problem is a classic example: checking whether a given route is short enough is easy, whereas discovering the best route may be hard. The **NP**-complete problems are the hardest in **NP**, in the sense that any **NP** problem can be efficiently transformed into them, so an efficient solution to one would solve all of **NP**. Turning to the quantum world, **BQP** (bounded-error quantum polynomial time) is the class of problems a quantum computer can solve efficiently with high probability, the quantum analogue of the classical class P, and it contains tasks such as integer factoring through Shor's algorithm. **QMA** (quantum Merlin-Arthur) is in turn the quantum analogue of **NP**, the class of problems whose quantum proof can be efficiently verified by a quantum computer. A central **QMA**-complete problem is estimating the ground-state energy of a quantum system, the local Hamiltonian problem, so that **BQP** stands to **QMA** much as P stands to NP, solving against verifying with quantum resources on both sides. It is no coincidence that this canonical hard problem is a physics problem: the difficulty of characterizing quantum systems is exactly what makes their simulation a natural target for quantum hardware.

While it presents a challenge on one hand, the apparent complexity of quantum mechanics can also serve as a boon. In 1980, Paul Benioff devised a method of computing based on evolution under the Schrödinger equation, uniting ideas from discrete classical computing with the unitary evolution present in quantum mechanics [21]. The first true description of a quantum computer, however, came from Richard Feynman [22, 23]. Feynman's view considers Nature as a computation, proceeding in step according to some predetermined set of local rules. Utilizing classical hardware, representing quantities like entanglement requires computing resources scaling exponentially with volume. However, using a manifestly quantum mechanical substrate to do the computation allows one to represent the necessary correlations using resources that scale proportionally to the volume of the system. In this regard, quantum computers are able to perform computations with an exponential separation in efficiency compared to their classical counterparts.

From this vision, enormous effort has gone into harnessing large quantum systems to perform calculations

beyond the computational reach of classical processors, in search of a so-called *quantum advantage*. While the status of current quantum hardware keeps a concrete demonstration of advantage out of reach, one can still implement practical simulations. Practical simulation is a heuristic term, defined only for the purposes of this dissertation. It means that while deep circuits cannot be run and physical error rates remain relatively high, quantum programs can be executed at a scale that requires more specialized classical methods to match (e.g. tensor networks or Clifford simulators). Additionally, there are techniques that can be used to counteract the effects of hardware errors, allowing results from real quantum devices that contain a meaningful quantum mechanical signature. One of the true values in practical simulation is that they provide a new avenue for development of both theory and methods that will accelerate the arrival of the computational capabilities forecast by Feynman.

Among the many schemes that may hold the promise of a quantum computational advantage, or at least a whisper of practical utility, the one that stands at the forefront is the simulation of fundamental physics. With problem classes ranging from thermal and ground state preparation to non-equilibrium dynamics, simulating physics is one of the most exciting near-term applications of quantum hardware. This thesis focuses on the last of these, the non-equilibrium time evolution of a quantum system, which reduces to the following: given the Hamiltonian H of a system, implement the unitary time evolution operator given by

$$U = e^{-iHt} \tag{1.2}$$

where t is the evolution time and i is the imaginary unit. Apart from capturing the dynamics of many interesting quantum systems, this problem appears in many other quantum applications like quantum phase estimation (QPE) [24, 25] and the variational quantum eigensolver (VQE) [26].

1.2 Primer on Quantum Simulation

This section provides an abstract example of what a general quantum simulation entails. Along the way, concrete examples will be provided that demonstrate standard techniques that are used to accomplish this goal. The problem of Hamiltonian simulation can be broken into five main steps: (i) encode the problem into a register of qubits, (ii) prepare the initial state, (iii) evolve the state according to U , (iv) measure the state, (v) classical post-processing on the measurement data.

1.2.1 Encoding

The importance of the first step cannot be overstated, as it can be likened to choosing a good basis for performing the calculation. For systems found in condensed matter, like spin systems, the mapping into qubits is natural. However for more exotic systems like those found in high-energy physics (HEP), there are many different bases to choose from, all with different convergence properties and different potential hardware implementations. In this context, the particular choice of basis is referred to as a *digitization*. The most foundational example of this is due to Jordan, Lee, and Preskill (JLP) [27, 28], work that is considered the gold standard in quantum simulation. In these works it is rigorously proven that scattering in scalar field theory lies in **BQP**. Part of this result hinges on the choice of basis, which among other things allows the quantum field to be represented on very few qubits.

1.2.2 State Preparation

Once the representation of the system is chosen, the initial state must be prepared. Generally one begins with the fiducial state $|0\rangle^{\otimes N}$ and applies a state preparation unitary U_{prep} to arrive at $|\psi_{init}\rangle = U_{prep} |0\rangle^{\otimes N}$. The type of states one would like to prepare depends strongly on the phenomena of interest. In the context of quantum field theories, generally it is the low energy eigenstates of the Hamiltonian that are of interest. Taking the lattice scalar field as an example, each site can be represented as a Harmonic oscillator. In the absence of an interaction, e.g. a $\lambda\phi^4$ term, the ground state can be represented as a Gaussian in field space, for which efficient state preparation procedures have been developed [29]. As it turns out, including the interaction does not fundamentally change this picture. Excitations can be placed on top of this, such as wavepackets, allowing one to study scattering events relevant for collider physics. For states that have exponentially decaying correlations, it is possible to design very scalable, shallow circuits that efficiently prepare the wavefunction to high fidelity. There are some cases where state preparation is not required. Studies of thermalization and quantum chaos provide a central example of this, where one is interested in the time evolution of a system due to a chaotic Hamiltonian, which for the purposes of this work is defined as one whose spectrum exhibits Wigner-Dyson statistics [30–34] and can be treated using random matrix theory (RMT). Such systems display a property known as *universality*, meaning that after sufficiently long evolution times, many of the initial details of the dynamics get erased and only a small number of conserved quantities, like energy, determine the properties of the state. To see these dynamics, the initial state must have overlap with eigenstates that are

far away from the edges of the spectrum, and so for most chaotic Hamiltonians, any product state will be sufficient.

1.2.3 Time Evolution

Once the desired initial state is prepared, the time evolution operator must be compiled into quantum gates. There are a vast number of ways to implement this operator, among which is by use of the product formulas, also known as a Trotter-Suzuki decomposition or “Trotterization” [35, 36]. This approach stems from the Baker-Campbell-Hausdorff (BCH) formula, which states that for two real or complex matrices A and B , their exponential may be expressed as

$$e^A e^B = e^C, \quad C = A + B + \frac{1}{2}[A, B] + \frac{1}{12}[A, [A, B]] - \frac{1}{12}[B, [B, A]] + \dots \quad (1.3)$$

where C is a formal series of commutators that does not necessarily converge. While this expression may seem incredibly unwieldy in the context of quantum simulation, its implications are far reaching. Consider a Hamiltonian acting on a system of spins which is a sum of terms $H = \sum_{i \in \Gamma} H_i$ where Γ denotes a set of sets of spins and H_i is a Hermitian operator supported on spins in i . Expressing the full exponential e^{-iHt} in terms of a native quantum gate set may not be possible. Expressing the exponential of each of these terms individually simplifies the operator structure such that the entire evolution may be approximated as a product. The first-order Trotter decomposition can be expressed as

$$\mathcal{J}_1(t) = \prod_{i \in \Gamma} e^{-iH_i t} \quad (1.4)$$

Generally such a formula will not be a good approximation of the true propagator U , as it amounts to keeping only the first term in the BCH series. By inspecting the form of higher order commutators, the overall norm of these terms will be controlled by any prefactors that dress the operators in the sum. If the evolution time t is sufficiently small, then the higher order terms will be suppressed by powers of t and can be dropped with error $\mathcal{O}(t^2)$. While this limits the evolution time achievable with a single step, a sequence of many first order steps can be taken to achieve a desired evolution time, increasing the error by a factor of t . As the number of steps is taken to infinity the following limit holds

$$\lim_{n \rightarrow \infty} \left(\prod_i e^{-iH_i t/n} \right)^n = e^{-iHt} \quad (1.5)$$

By careful rearrangement of the terms in a product formula, even better convergence can be achieved. For example, symmetrizing the first-order product formula constructs a second-order product formula defined as

$$\mathcal{J}_2(t) = \mathcal{J}_1(-t/2)^\dagger \mathcal{J}_1(t/2) \quad (1.6)$$

which for a given time step has error of order $\mathcal{O}(t^3)$. Higher-order product formulas can be recursively constructed in a similar manner, pushing the approximation error to higher orders in t . Beyond its use in quantum simulation, Trotterization has also found widespread applicability in many areas of computational physics, from numerically evaluating partition functions of quantum systems in Quantum Monte Carlo to providing a time evolution method, TEBD, in Matrix Product States.

1.2.4 Measurement

Once the quantum state has been evolved, it is necessary to perform measurements to extract the desired information. In all cases, the output from a quantum device is a bitstring s , a list whose entries are either 0 or 1, corresponding to the result for each qubit. For a register of N qubits, the probability of seeing a particular bitstring is determined by Born's rule from quantum mechanics, namely $p(s) = |\langle s | \psi \rangle|^2$. By running the circuit many times, the probability distribution p can be reconstructed. From this distribution, the expectation value and uncertainty of some observable O can be estimated using bootstrap resampling. As a concrete example, take $O = Z_i Z_j$ for qubits i, j . Since O is diagonal, given the measured distribution $p(s)$ the mean can immediately be estimated as

$$\langle O \rangle = \sum_s p(s) (-1)^{s_i + s_j} \quad (1.7)$$

where the sum is over all 2^N possible bitstrings. When O contains off-diagonal matrix elements, e.g. it contains X s or Y s on a set of qubits, then prior to measurement a local change of basis is applied that diagonalizes the observable in the readout basis. As any observable may be decomposed into a sum of Pauli strings, this method allows one to estimate general observables. For general observables, this can still be non-trivial for two reasons. Since the number of Paulis for a system grows exponentially with the number of qubits, observables that decompose as a sum of exponentially many Paulis require estimating a great number of individual observables and essentially forces one to resort to full state tomography. Second, it is possible that over the course of the computation the information that one is interested in has been smeared over the exponentially large Hilbert space and so the resources required for readout will also scale exponentially with system size, causing a signal-to-noise problem.

More powerful methods have been proposed, among which are a class of techniques known as *random measurement protocols*. These protocols have been used to estimate a variety of quantities from fidelities to Rényi entropies. The random measurement approach involves constructing an ensemble of unitaries and for each shot, randomly sampling an element of this ensemble to append to the end of the circuit and measuring. Generally this approach will still have a measurement overhead that is exponential in sample size, although in some cases it is a less severe exponential overhead than in the naive approach. One of the most well known random measurement techniques that is able to estimate a large number of relevant observables is “Classical Shadows” [37]. The sampled unitaries and associated measurement bitstrings are collected, forming the shadow table $\{U_i, s_i\}$. It is useful to view the average effect of this protocol on a state ρ as a quantum channel called the shadow channel $\mathcal{M}(\cdot)$, i.e. $\rho \rightarrow \mathcal{M}(\rho)$. From the shadow table, the list of classical snapshots of the state $\hat{\rho}_i = \mathcal{M}^{-1}[U_i^\dagger |s_i\rangle \langle s_i| U_i]$ can be used to estimate the expectation value of an arbitrary observable. While the associated sample complexity depends heavily on the ensemble of unitaries sampled from, on average very few measurements are needed for a wide variety of observables. For simulations where one is interested in the expectation value of a local observable, such as a local magnetization of spins or the chiral condensate in the Schwinger model, shadows will typically be less performant than reconstructing the bitstring distribution. However, for global observables like fidelity or entanglement, classical shadows can require far fewer measurements than other approaches. An additional advantage is that the data in the shadow table can be re-used for many different observables, spawning the phrase “measure first, ask questions later”.

1.2.5 Post-processing of Results

One ever present feature of modern day quantum simulators that has so far had no mention is the presence of noise. Understanding how to deal with noise has so far been one of the primary tasks in most practical simulations to date. In the absence of a fully fault tolerant quantum device, circuits must be designed in both a shallow and scalable way so that the device does not decohere before a meaningful signal can be extracted. Even with optimal circuits, errors will still happen, and they must be mitigated accordingly [38]. One immediate tool is to use a physics informed post-selection. This approach leverages the fact that physical systems have symmetries and conserved quantities associated with them. If one is doing ground state simulations of a system with a $U(1)$ symmetry, any state subject to the dynamics of the Hamiltonian must conserve the total charge, and so any measurement result that lives in a different charge sector heralds

an error during the computation and should be discarded. For more general cases, a variety of powerful error mitigation strategies have been developed for this purpose. Two techniques that should be included in any quantum simulation workflow are Dynamical Decoupling (DD) [39] and Pauli Twirling (PT) [40]. Dynamical Decoupling reduces idle errors, i.e. errors that affect qubits that are not actively undergoing an operation, by subjecting them to sequences of operations whose logical effect on the encoded information is the identity, but which average out errors that occur during the sequence. This sequence occurs on a shot-by-shot basis. To mitigate errors at the ensemble level, Pauli Twirling applies Pauli operators to the beginning and end of a quantum gate so that logically the same gate is applied, but coherent errors of the form $e^{i\theta P}$ for some Pauli P are transmuted into incoherent errors on average. Mitigating incoherent errors can be advantageous compared to coherent errors as their effect on expectation values is easier to model, generally taking the form of some multiplicative or additive bias. This bias factor can be canceled by running reference circuits that either systematically amplify the noise or learn its impact on observables some other way, e.g. as is done with Zero Noise Extrapolation (ZNE) [41]. One drawback of many of these methods is that they are heuristic and are unable to remove all systematic sources of uncertainty. Probabilistic Error Cancellation (PEC) [42] stands apart as a promising method that is provably free of bias, however its shot overhead grows exponentially with the depth of the circuit, making it unfeasible for overly complex calculations.

1.3 Model Systems

This section reviews various types of many-body systems that are relevant to different parts of this thesis. Each will contain one or two specific examples that serve as archetypes for models of their specified class. Before turning to specific examples, it is useful to fix the two axes along which these systems will be compared. The first is the dimension of the local Hilbert space. Spin and fermionic systems have finite dimensional local spaces, whereas the gauge fields considered at the end of this section carry continuous, group-valued degrees of freedom and therefore infinite dimensional local spaces. The second axis is the structure of the spectrum. In each class there is a solvable limit, in which the Hamiltonian reduces to a quadratic form that can be diagonalized exactly, and an interacting regime in which this structure is lost and the dynamics becomes ergodic. The archetypes below are chosen to sit on either side of this divide.

1.3.1 Fermionic Models

Consider a one dimensional lattice with N sites, each hosting a set of operators c_i and $c_i^\dagger$ that satisfy the Canonical Anti-Commutation Relations (CAR)

$$\{c_i, c_j\} = 0, \quad \{c_i^\dagger, c_j\} = \delta_{ij} \quad (1.8)$$

The Hilbert space of a single site is given by $\mathcal{H} = \text{span}(|0\rangle, |1\rangle)$ with basis states whose relation to the $c_i, c_i^\dagger$ is described with the compact notation

$$c_i |\alpha_i\rangle = \alpha_i |1 - \alpha_i\rangle \quad (1.9)$$

$$c_i^\dagger |\alpha_i\rangle = (1 - \alpha_i) |1 - \alpha_i\rangle \quad (1.10)$$

$$(1.11)$$

where $\alpha_i \in \{0, 1\}$. This also motivates the definition of the number operator $n_i = c_i^\dagger c_i$, which counts the number of excitations on the site. The Hilbert space $\mathcal{F}$ of the entire lattice is described by the N -fold tensor product of each of the single site Hilbert spaces. This *Fock space* decomposes into the following direct sum

$$\mathcal{F} = \bigoplus_{\gamma=0}^N \mathcal{H}_\gamma \quad (1.12)$$

where each of the $\mathcal{H}_\gamma$ is spanned by states of the form $|\alpha_1, \dots, \alpha_N\rangle$ such that $\sum_i n_i |\alpha_1, \dots, \alpha_N\rangle = \gamma |\alpha_1, \dots, \alpha_N\rangle$. A defining feature of fermionic systems is that moving one excitation through another picks up a phase. This appears at the operator level through the relations

$$c_i |\alpha_1, \dots, \alpha_N\rangle = (-1)^{\sum_{k=1}^{i-1} \alpha_k} \alpha_i |\alpha_1, \dots, 1 - \alpha_i, \dots, \alpha_N\rangle \quad (1.13)$$

$$c_i^\dagger |\alpha_1, \dots, \alpha_N\rangle = (-1)^{\sum_{k=1}^{i-1} \alpha_k} (1 - \alpha_i) |\alpha_1, \dots, 1 - \alpha_i, \dots, \alpha_N\rangle \quad (1.14)$$

where the overall phase depends on the ordering of the fermionic modes, which is arbitrary. This phase is not observable, however, since any fermionic operator corresponding to a physical observable must commute with the parity operator defined by $P = (-i)^N \prod_{i=1}^{2N} \gamma_i$, where $\gamma_{2i} = -i(c_i^\dagger - c_i)$ and $\gamma_{2i-1} = c_i^\dagger + c_i$ are Majorana fermions that satisfy $\{\gamma_i, \gamma_j\} = 2\delta_{ij}$ and are Hermitian $\gamma_i = \gamma_i^\dagger$.

A simple fermionic model that describes noninteracting particles hopping between adjacent lattice sites is

$$H = J \sum_i c_i^\dagger c_{i+1} + h.c. \quad (1.15)$$

where $h.c.$ denotes the Hermitian conjugate. Note that the only terms that appear in this Hamiltonian are bilinear in the $c_i, c_i^\dagger$. Models of this type, that are quadratic in these operators, are exactly solvable as the Hamiltonian can be diagonalized in a plane wave basis. States that can be written as the exponential of a quadratic fermionic Hamiltonian are known as *Gaussian states* and their properties are completely determined by

$$C_{ij}^{(2)} = \text{Tr}[\rho c_i^\dagger c_j] \quad (1.16)$$

in the sense that any higher-order correlation function can be decomposed into a sum of products of the entries of $C^{(2)}$, a property known as Wick's theorem. In terms of the Majorana operators introduced above, the most general quadratic Hamiltonian, including pairing terms that do not conserve particle number, can be written compactly as $H = \frac{i}{4} \sum_{ij} A_{ij} \gamma_i \gamma_j$ for an antisymmetric matrix A .

An extremely useful expression for the density matrix of a gaussian state in terms of the correlation matrix is

$$\rho = \det(C^{(2)}) \times e^{-\sum_{ij} \left[\log \frac{C^{(2)}}{1 - C^{(2)}} \right]_{ij} c_i^\dagger c_j} \quad (1.17)$$

where the prefactor is due to normalization. From this expression, many interesting quantities like the entropy can be directly calculated from the spectrum of $C^{(2)}$ using *Wick's Theorem*, a factorization that is reminiscent of free field theories [43, 44] or of Isserlis' theorem in probability theory [45]. This last point has a sharp computational consequence. Because the entire state is fixed by the polynomially many entries of $C^{(2)}$, dynamics generated by a quadratic Hamiltonian can be simulated efficiently on a classical computer by evolving $C^{(2)}$ alone, rather than the exponentially large state vector; in the circuit picture these are the matchgate, or free fermion, circuits. It is exactly this classical tractability that is lost once interactions are switched on.

Moving from the solvable to the more general interacting case, another level of complexity can be reached by adding an extra term to the free-fermionic Hamiltonian

$$H = J \sum_{i,\sigma} c_{i,\sigma}^\dagger c_{i+1,\sigma} + h.c. \quad (1.18)$$

$$+ U \sum_i n_{i,\uparrow} n_{i,\downarrow} \quad (1.19)$$

where the σ correspond to an extra label that has been added denoting the spin of an excitation. This is known as the Fermi-Hubbard model, and is a paradigmatic model in condensed matter physics, hosting

a variety of phases of matter and considered classically intractable¹ in two and higher dimensions [47]. Physically the model encodes a competition between the hopping term, which favors delocalization, and the onsite repulsion, which favors localized, singly occupied sites; at strong coupling and half filling this competition drives the system into a Mott insulating phase. The presence of the quartic term in time evolution generates non-Gaussian correlations that take the state off the manifold of Gaussian states, causing it to become quantum chaotic and eventually thermalize. Since this term removes the Gaussian structure, no efficient classical description of the general dynamics is known, which is precisely what makes the model a natural target for quantum simulation.

1.3.2 Spin Models

Spin systems are defined by a lattice of spins that hosts a set of Pauli operators $\{X_i, Y_i, Z_i\}$ on each site. Spins do not satisfy the same set of exchange statistics that fermions do, and should instead be viewed as a system of bosons whose number levels are truncated to $\{0, 1\}$. Operators on the same site anti-commute, e.g. $\{X_i, Z_i\} = 0$ and satisfy $X_i Z_i = iY_i$. A spin- $\frac{1}{2}$ is nothing other than a qubit, so spin models map onto quantum hardware with no encoding overhead at all, a directness that makes them the default setting for many practical quantum simulations. One of the canonical systems that can be built out of such operators is the Mixed Field Ising Model (MFIM)

$$H = J \sum_i Z_i Z_{i+1} + h \sum_i X_i + g \sum_i Z_i \quad (1.20)$$

where the g, h terms correspond to external magnetic fields that couple to individual spins.

The value of g determines the classical solvability of this model. For $g = 0$, this model is integrable and can be mapped to free fermions [48]. Considering first the limit $h \gg J$ (disordered), the ground state of this system will have all the spin aligned in X basis eigenstates so that the energy is minimized. On the other hand, in the $J \gg h$ (ordered) limit, all the spins will be aligned/anti-aligned in Z basis depending on the sign of J . Thus, at comparable values of $J \sim h$ there is competition between the two non-commuting terms. The system is driven through a quantum phase transition at $J = h$, which is also the limit in which the Ising field theory arises.

¹In one dimension, this model is solvable via the Bethe ansatz [46], however adding a next-to-next-to-nearest neighbor hopping term makes it again unsolvable.

For $g \neq 0$, by contrast, the integrability is spoiled and the model is ergodic and capable of generating highly complex states. This is why the model recurs as a minimal, hardware friendly setting for studies of quantum chaos and thermalization. Despite this, it is exceedingly easy to implement at the level of quantum gates, only requiring single qubit gates and nearest-neighbor interactions, making it a powerful testbed for developing quantum simulation techniques. More general spin models, such as the Heisenberg and XXZ chains, are obtained by adding XX and YY couplings, and a similar dichotomy between integrable and chaotic regimes persists. These are minimal additions to the Ising model that bring it closer to describing realistic behavior in magnetic materials.

1.3.3 Lattice Gauge Theories

Gauge fields diverge from the former two types of models considered as, in many cases, they have infinite dimensional local Hilbert spaces. This is particularly true for non-abelian Gauge theories, whose local degrees of freedom are group-valued. The most common formulation in regards to quantum simulation is the Hamiltonian formulation, where time is kept continuous and space is discretized on a lattice. Recalling that the role of a gauge field is to transport color charge, the “pure gauge” Hamiltonian (i.e. no matter) in D spatial dimensions can be taken to live on a hypercubic lattice with group degrees of freedom on the links. The Hilbert space of a single link is $L^2(G, d\mathbf{g})$, which is the space of square-integrable wavefunctions over the gauge group G with Haar measure $d\mathbf{g}$. A generic wavefunction is expressed in the group element basis as

$$|\psi\rangle = \int d\mathbf{g} \psi(\mathbf{g}) |\mathbf{g}\rangle \quad (1.21)$$

where distinct group elements are orthogonal. The basic operator degree of freedom that lives on a link l is the link operator U_l , which is group element valued. This can be viewed as the lattice version of a Wilson line. Gauge transformations are generated by exponentiation of the elements of the Lie algebra $E_{L,R}^a$, which are known as the electric operators. Here a denotes the color charge and L/R denote which side of the link the operator acts on.

A crucial feature of Gauge theories is *Gauss’ Law* which enforces local charge conservation, stating that the theory must be invariant under arbitrary gauge transformations. Defining the Gauss Law operator $G^a(n)$ on site n

$$\hat{G}^a(n) = \sum_{i=1}^d \left[\hat{E}_R^a(n - e_i, e_i) - \hat{E}_L^a(n, e_i) \right] \quad (1.22)$$

then the condition that physical states must satisfy is

$$\hat{G}^a(n) |\Psi_{\text{phys}}\rangle = 0 \quad (1.23)$$

Gauss' law is therefore not an equation of motion but a constraint that carves a physical subspace out of a much larger kinematical Hilbert space. Faithfully simulating the theory requires either working directly within this subspace or energetically penalizing states that violate the constraint, a complication with no counterpart in the spin and fermionic models above.

Therefore, the Hamiltonian that defines this theory must be gauge invariant, which means in practical terms it must be constructed out of simple objects that have no free color indices. For the electric operators, this can be done by contracting the color indices to get the electric part of the Hamiltonian $E_L^a E_L^a = E_R^a E_R^a = E^2$. This object is known as the quadratic Casimir and it commutes with all elements of the group Lie algebra. The magnetic part of the Hamiltonian is formed by noting that the product of link operators around any closed loop will be gauge invariant. On a hypercubic lattice, the smallest possible closed loop is composed of four links in a single plane, an object known as a “plaquette”. Noting that the link operators are group-valued, the product must be traced over. Summing both the electric and magnetic parts over the entire lattice results in the famous Kogut-Susskind Hamiltonian [49]

$$H = \frac{g^2}{2a} \sum_l E_l^2 + \frac{1}{2ag^2} \sum_P \text{Tr } 2\mathbb{1} - P - h.c. \quad (1.24)$$

where P denotes a plaquette, the sum is over all plaquettes, and the $h.c.$ is added to make the sum over plaquettes Hermitian.

The Hamiltonian above describes the gauge field on its own. In most theories of interest it is coupled to dynamical matter, introduced as fermionic fields on the lattice sites whose hopping is dressed by the link operators so as to stay gauge invariant, schematically $\psi_n^\dagger U_n \psi_{n+1}$. Gauss' law then ties the divergence of the electric field at a site to the charge of the matter sitting there, making concrete the statement that gauge fields transport charge. The simplest realization is the Schwinger model, Quantum Electrodynamics (QED) in one spatial dimension, which despite its simplicity already displays confinement and a nonzero chiral condensate, and for that reason has become the standard benchmark for quantum simulation of gauge theories.

As the link Hilbert space is infinite dimensional, simulating this theory on qubits requires a *digitization*, which means truncating the group-valued degrees of freedom to a finite local basis. The two terms of the Hamiltonian favor opposite bases: the electric term is diagonal in the representation, or “electric”, basis

natural at strong coupling, whereas the magnetic term is diagonal in the group element basis natural at weak coupling. Any finite qubit register must truncate one of these, so the quality of a chosen digitization depends on which regime of the coupling g one wishes to reach. Physical predictions are ultimately recovered only in the continuum limit $a \rightarrow 0$, where by asymptotic freedom the bare coupling itself must be driven toward zero, connecting this lattice construction back to the high energy behavior of QCD. However, understanding how to take the continuum limit of the Kogut-Susskind Hamiltonian in a way that allows for extrapolation and RG flows is highly non-trivial and very much an area of active research.

1.4 Outline of this Thesis

The chapters of this thesis cover the projects for which I was a primary author, ordered in such a way to reflect the quantum simulation workflow described above. This thesis is organized as follows:

Chapter 2 begins with analog quantum simulation, where the model of interest is emulated through Hamiltonian engineering rather than gate-based compilation. Focusing on simulations of Heisenberg-model dynamics built from native Ising interactions and strong drive fields, we quantify the algorithmic errors introduced by Trotterization and by the continuous “spiral” drive protocol, derive bounds on the associated error rates, and show how the free parameters of these protocols can be optimized to minimize algorithmic error at fixed experimental resources. This analysis is a step toward fully quantified uncertainties in analog simulations of quantum spin systems. This work was done in collaboration with Nikita Zemlevskiy and Stephan Caspar and was published in *Physical Review A* 109 (5), 052425 [50].

Chapter 3 turns from analog to digital simulation, and from spin models to non-Abelian gauge theories. Using a Fully Gauge-Fixed formulation of $SU(2)$ lattice gauge theory, in which all redundant degrees of freedom are removed and Gauss’ law is automatically satisfied, we develop and benchmark quantum circuits for simulating the real-time dynamics of a two-plaquette system in a mixed-basis formulation that remains accurate across all values of the gauge coupling. This work charts a course toward simulations of larger two-dimensional non-Abelian systems. This work was done in collaboration with Dorota Grabowska and Zhiyao Li and is drawn from arXiv:2512.22782 [51].

Chapter 4 examines what makes such simulations genuinely hard, studying the quantum complexity of the same class of theories through the lens of non-stabilizerness, or “magic” – a resource beyond entanglement that quantifies the difficulty of classically simulating a quantum state. We measure the non-stabilizerness

generated by the dynamics of an $SU(2)$ lattice gauge theory using random measurements, connecting the resource requirements of gauge theory simulations to fundamental notions of classical simulability. A central development of this work is an approach to error mitigation of bitstring distributions, which has applicability in a variety of other settings. This work was done in collaboration with Dorota Grabowska and is drawn from arXiv:2606.14842 [52].

The next two chapters address the complementary challenge of extracting physics from a quantum simulator once a state has been prepared. Chapter 5 develops the measurement side for lattice gauge theories, introducing three randomized measurement protocols, based on the framework of classical shadows, that are tailored to systems with local (gauge) symmetries. Exploiting the duality between the $\mathbb{Z}_2$ lattice gauge theory in $(2 + 1)$ dimensions and the Ising model, these symmetry-aware protocols achieve up to exponential reductions in sample complexity for estimating gauge-invariant observables relative to symmetry-agnostic schemes, at the cost of modestly deeper randomization circuits. This work was done in collaboration with Jacob Bringewatt, Andreas Elben, and Niklas Mueller and was published in *Quantum* 10, 2127 [53].

Chapter 6 extends the measurement program to analog platforms with limited control, developing a method for reconstructing the entanglement structure of quantum states from the few-body correlation functions accessible in such experiments. Based on a systematic expansion of the many-body Wigner functional around Gaussian states, the approach uses measured two- and four-point functions to reconstruct entanglement Hamiltonians and their spectra, and is demonstrated on the post-quench thermalization dynamics of an extended Fermi-Hubbard chain, where it captures signatures of quantum chaos that Gaussian-state descriptions miss entirely. This work was done in collaboration with Robert Ott, Torsten Zache, and Niklas Mueller and was published in *Physical Review Letters* 135 (4), 040201 [54].

Finally, Chapter 7 confronts the errors of present-day hardware directly through lightweight quantum error detection. We report Hamiltonian simulation experiments on IBM’s heavy-hexagonal architecture using many interacting blocks of the $[[4, 2, 2]]$ error-detecting code, enabled by shallow syndrome extraction circuits tailored to the native qubit connectivity and a selective-filtering strategy that postselects on the detection events most relevant to a given logical observable, avoiding the exponential shot loss of full postselection. Encoded simulations of the mixed-field Ising model in one dimension outperform their unencoded counterparts despite a several-fold increase in circuit depth, achieving beyond break-even estimation of local observables without any error mitigation, while the delocalized nature of the encoding relaxes connectivity constraints enough to enable two-dimensional false-vacuum decay simulations at lower overall depth than direct implementation.

This work was done in collaboration with Nikita Zemlevskiy, Sarah Powell, Sebastian Grieninger, Xiaojun Yao, Dorota Grabowska, Zhiyao Li, Anne Lashbrook, Ziyuan Li, and Martin Savage and is drawn from arXiv:2607.24947 [55].

Chapter 2

OPTIMIZATION OF ALGORITHMIC ERRORS IN ANALOG QUANTUM SIMULATIONS

2.1 Introduction

Simulation of quantum many-body systems has been an outstanding goal of the scientific community for decades. Despite the enormous success of classical methods in determining static properties of quantum systems, they suffer from superpolynomial resource scaling for many real-time and finite-density phenomena [28, 56–60]. Quantum computing (QC) has the potential to perform such calculations with realistic resource costs [21, 22]. A variety of noisy intermediate-scale quantum (NISQ) devices are becoming available for science, and while the applicability of these devices is limited by various factors, interesting results can already be obtained for modest calculations [61–66]. In contrast to digital QC, which utilizes discrete logic gates to process information, analog QC operates by continuously modulating parameters of the experimental apparatus, e.g. a magnetic field. Analog simulation is possible in cases where the problem of interest can be mapped to a simulator’s native Hamiltonian exactly, or with controllable error [67]. Hybrid digital-analog quantum simulation protocols have also been proposed [68, 69]. Although the analog approach is more limited than digital methods, current device capabilities show that beyond-classical results may be possible with analog simulation due to larger possible system sizes [70–73].

Systems of many interacting spins are a common starting point for a number of the quantum simulations to date [74–76], due to the ease of mapping physical problems to the effective Hamiltonians describing many NISQ devices. One of the most highly studied spin models is the Heisenberg model, which serves as a testing ground for condensed matter systems [77, 78], as well as a universal model for quantum computing [79]. Simulations of the Heisenberg model constitute a proposed step towards achieving quantum advantage [80]; these simulations can be mapped to processes relevant for high-energy particle physics [81–84], quantum gravity [85], QCD processes [86, 87], and other applications. Many of the analog platforms that are currently available can be naturally described by the Ising model, which has proven to be a flexible testbed for studying a broad variety of physics [88], as well as a system where Heisenberg interactions can be realized [89–101].

Systems supporting Ising interactions include trapped ions [74, 102, 103], neutral atoms [72, 95, 104, 105], nuclear spins [106, 107], and superconducting qubits [75, 108, 109].

Major limitations of all these devices are the numerous sources of error that accompany any particular simulation. Several error correction schemes have been proposed for analog quantum computation [110, 111], but they have not yet been implemented on devices available today. Because of this, the ability to quantify these errors is crucial to make precise statements about a simulation. The three main sources of error that are important to consider for the purposes of simulation are the encoding error, the algorithmic error, and the hardware error [112]. The encoding and algorithmic error come from the approximations that are made in translating physics to a form best suited for a particular device (e.g., Hilbert space truncations, commutator error for product formulas, etc.). The hardware error comes from experimental imperfections and noise on a device, which is usually uncontrolled and needs to be dealt with using error mitigation. Much work has been done examining Trotter errors [38, 113–115] and errors induced by analog devices [116–119] individually, but little attention has been paid to the interplay of these different sources of error. As it turns out, there are particular choices of parameters that will minimize the cumulative error within and across these different categories.

With the advent of analog quantum simulation, an understanding of uncertainties stemming from mapping physical problems and implementations of simulation algorithms is now necessary. This chapter investigates the interplay of two types of algorithmic errors occurring in digital quantum computations on analog systems, Trotter errors and idle errors. Trotter errors arise due to the decomposition of the evolution operator into a finite product formula of non-commuting unitaries. Idle errors, on the other hand, stem from the inability of analog systems to turn off interactions while applying local pulses (gates). Continuously driven fields have also recently been shown to implement Heisenberg-type evolution in Ising systems [101]; such methods incur errors similar to idle errors as well. The decomposition of the algorithmic error into Trotter and idle errors is shown in Fig. 2.1.

On a perfect device, improving precision is straightforward if these two error sources are analyzed independently. In the absence of realistic device constraints, decreasing the Trotter step size is guaranteed to decrease the Trotter error, and decreasing the pulse width on the device lowers the idle error contribution. However, these error sources play against each other on real-world machines. For example, taking many short Trotter steps will incur a large error if the idle errors are high in the device. Furthermore, taking very small Trotter steps may be impossible due to limited pulse width. Therefore, it is necessary to analyze these

effects collectively to find the parameters that maximize the precision of the output from simulations on a real device. In this chapter, a general framework for analyzing the error scaling of engineered time evolution methods is presented and is used to compare the performance of the aforementioned methods. We compare with numerical results for an analog simulator with (quasi-) local Ising interactions and global longitudinal and transverse-field addressing. This analysis gives the optimal choice of simulation parameters as a function of device parameters. In Section 2.2, the general Hamiltonian engineering methods are described, Section 2.3 introduces the types of algorithmic error studied, Section 2.4 presents the results of the error analysis, and Section 2.5 discusses the implications of this chapter's findings.




Figure 2.1: *Mapping hardware interactions to a target Hamiltonian through pulse engineering* The device Hamiltonian H_{dev} can be used to emulate a target system of interest with Hamiltonian H_{tar} using Trotter-like methods (top path) or continuous driving methods (bottom path). In the limit of perfect pulses (equivalently, infinite magnetic field), the dashed curves implement the target time evolution exactly. The solid curves represent the mappings that are approximate due to algorithmic errors. These errors increase the overall error when mapping between Hamiltonians due to idle errors and when implementing the time evolution operator $U(t)$ (Trotter errors). Algorithmic errors affect constant-field time evolution engineering methods as well. The constant-field method C_1 and one of the Trotter methods S_1 are labeled as resulting time evolution methods within the space of possible $U(t)$ operators.

2.2 Hamiltonian Engineering Methods

Consider a general Hamiltonian that may be implemented in various experimental systems:

$$\begin{aligned}
 H_{\text{dev}}(t) &= H_{\text{idle}} + H_{\text{drive}}(t), \\
 H_{\text{idle}} &= \sum_{i < j} J_{ij} (c_x X_i X_j + c_y Y_i Y_j + c_z Z_i Z_j), \\
 H_{\text{drive}}(t) &= \vec{B}(t) \cdot \left(\sum_i \vec{S}_i \right).
 \end{aligned} \tag{2.1}$$

Using carefully engineered pulse sequences of the global magnetic field $\vec{B}(t)$, the coefficients $c_{x,y,z}$ in H_{idle} can be modified into any target Hamiltonian $H_{\text{tar}} = H'_{\text{idle}}$ in the interaction picture, obeying $c_x + c_y + c_z = c'_x + c'_y + c'_z$ and $\min c_{x,y,z} \leq c'_{x,y,z} \leq \max c_{x,y,z}$ [120]. The targeted Hamiltonian H_{tar} arises through dynamical decoupling of the strong magnetic fields $|\vec{B}(t)| \leq \Omega \gg \max J_{ij}$

$$\begin{aligned} H_{\text{tar}} &= \frac{1}{\tau} \int_0^\tau dt U_{\text{drive}}(t)^\dagger H_{\text{idle}} U_{\text{drive}}(t), \\ i\partial_t U_{\text{drive}}(t) &= H_{\text{drive}}(t) U_{\text{drive}}(t), \\ U_{\text{drive}}(0) &= \mathbb{I}, \end{aligned} \quad (2.2)$$

where an average is taken over a period τ of the pulse sequence with $\vec{B}(t + \tau) = \vec{B}(t)$ and Ω is the maximum field strength that can be implemented on the device. The rest of the chapter focuses on the concrete case of emulating the isotropic Heisenberg $c'_{x,y,z} = \frac{1}{3}$ evolution from Ising interactions $c_z = 1, c_{x,y} = 0$

$$H_{\text{idle}} = H_{\text{Ising}} = \sum_{i < j} J_{ij} Z_i Z_j \quad \rightarrow \quad H_{\text{tar}} = H_{XXX} = \sum_{i < j} \frac{J_{ij}}{3} (X_i X_j + Y_i Y_j + Z_i Z_j). \quad (2.3)$$

Consider two types of pulse engineering methods, intermittent and continuous driving. Intermittent driving is inspired by Trotterized time-evolution on digital quantum computers. The strong magnetic field allows approximate global X , Y , or Z $\frac{\pi}{2}$ -rotations through short pulses $\epsilon = \frac{\pi}{2\Omega} \ll J_{ij}^{-1}$ at the maximal field strength Ω

$$\begin{aligned} R_X^\pm(\epsilon) &= \exp \left\{ \left(-i\epsilon \sum_{i < j} J_{i,j} \hat{Z}_i \hat{Z}_j \mp \frac{i\pi}{4} \sum_j \hat{X}_j \right) \right\}, \\ R_Y^\pm(\epsilon) &= \exp \left\{ \left(-i\epsilon \sum_{i < j} J_{i,j} \hat{Z}_i \hat{Z}_j \mp \frac{i\pi}{4} \sum_j \hat{Y}_j \right) \right\}, \\ R_Z^\pm(\epsilon) &= \exp \left\{ \left(-i\epsilon \sum_{i < j} J_{i,j} \hat{Z}_i \hat{Z}_j \mp \frac{i\pi}{4} \sum_j \hat{Z}_j \right) \right\}. \end{aligned} \quad (2.4)$$

Note that at infinite Ω these are exact rotation gates. These pulses effectively permute the coefficients $c_{x,y,z}$, yielding a set of global gates which can be composed in the usual Trotter-like fashion. There have been several recent developments showing Trotter-like methods of Hamiltonian engineering where similar target systems are recovered [89–100]. For the native Ising interactions of Eq. (2.3), there are global evolutions generated by native XX , YY and ZZ gates

$$\begin{aligned}
R_{ZZ}(t) &= \exp\{-itH_{\text{Ising}}\}, \\
R_{XX}^\pm(t, \epsilon) &= R_Y^\mp(\epsilon)R_{ZZ}(t)R_Y^\pm(\epsilon), \\
R_{YY}^\pm(t, \epsilon) &= R_X^\mp(\epsilon)R_{ZZ}(t)R_X^\pm(\epsilon).
\end{aligned} \tag{2.5}$$

Taking the product of these three gates defines a sequence which is referred to as “pseudo-first order”

$$\mathcal{S}_{1/2}(\vec{t}, \epsilon) = R_{ZZ}(t_z)R_{YY}^+(t_y, \epsilon)R_{XX}^+(t_x, \epsilon). \tag{2.6}$$

$\mathcal{S}$ stands for the pulse sequence and the subscript indicates the error scaling with ϵ . As in digital product formulas, order is defined by the scaling of the error with the parameters of the sequence. For instance, a first-order sequence has errors that only scale like $O(t^2, \epsilon t, \epsilon^2)$. The “pseudo”-modifier indicates that this sequence generates unwanted errors at order $O(\epsilon)$ which contribute to H_{tar} in Eq. (2.2). Such errors are avoidable if the alternative gate set

$$\begin{aligned}
\tilde{R}_{XX}^\pm(t, \epsilon) &= R_Y^\pm(\epsilon)R_{ZZ}(t)R_Y^\pm(\epsilon), \\
\tilde{R}_{YY}^\pm(t, \epsilon) &= R_X^\pm(\epsilon)R_{ZZ}(t)R_X^\pm(\epsilon),
\end{aligned} \tag{2.7}$$

is used to define a “true-first order” method

$$\mathcal{S}_1(\vec{t}, \epsilon) = R_{ZZ}(t_z)\tilde{R}_{YY}^+(t_y, \epsilon)\tilde{R}_{XX}^+(t_x, \epsilon). \tag{2.8}$$

Applying $\frac{\pi}{2}$ -pulses in the same direction twice leads to destructive interference of the error terms plaguing $\mathcal{S}_{1/2}$. This comes at the cost of a residual global Z rotation by π . This overall rotation is easy to keep track of and correct for in subsequent stages of a simulation. To extend this Trotter-like approach, we present a “pseudo-second order” sequence, $\tilde{\mathcal{S}}_1$, which carefully rearranges the gates in $\mathcal{S}_1$ to remove errors at order $O(t^2)$

$$\tilde{\mathcal{S}}_1(\vec{t}, \epsilon) = R_{ZZ}(t_z/2)R_X^+(\epsilon)R_{ZZ}(t_y/2, \epsilon)\tilde{R}_{XX}^+(t_x, \epsilon)R_{ZZ}(t_y/2)R_X^-(\epsilon)R_{ZZ}(t_z/2). \tag{2.9}$$

Finally, there is the fully symmetrized version of $\mathcal{S}_1$ according to the formula $\mathcal{S}_2(\vec{t}, \epsilon) = \mathcal{S}_1(\vec{t}, \epsilon)\mathcal{S}_1(-\vec{t}, -\epsilon)^\dagger$ [121]. As one would expect, all errors $O(\epsilon^2, \epsilon t, t^2)$ vanish due to the symmetry, making this a “true-second order” sequence

$$\mathcal{S}_2(\vec{t}, \epsilon) = R_{ZZ}(t_z/2)\tilde{R}_{YY}^+(t_y/2, \epsilon)\tilde{R}_{XX}^+(t_x/2, \epsilon)\tilde{R}_{XX}^-(t_x/2, \epsilon)\tilde{R}_{YY}^-(t_y/2, \epsilon)R_{ZZ}(t_z/2). \tag{2.10}$$




Figure 2.2: *Pulse sequences for generating Heisenberg evolution from ZZ interactions, corresponding to Eqs. (2.6), (2.8) - (2.11)* For each pulse sequence, the magnetic-field strength (in units of Ω) is given as a function of time (in units of $\epsilon = \frac{\pi}{2\Omega}$). The length of each sequence U is denoted by τ_U . The idling times t_i are shown in the graphics for each of the Trotter sequences.

These pulse sequences are illustrated in Fig. 2.2. The parameter $\vec{t} = (t_x, t_y, t_z)$ specifies the amount of time to evolve without a magnetic field in the respective frames XX, YY, ZZ . They obey $\tau c'_i = t_i + O(\epsilon)$ since the finite pulse width also contributes to these coefficients, and the t_i need to be adjusted accordingly. It is important to reiterate that the sequences $\mathcal{S}_1$ and $\tilde{\mathcal{S}}_1$ reproduce evolution under H_{XXX} up to a global Z π -rotation, which is accounted for in this analysis.

The constant drive field method of Refs. [101, 122] is an example of continuous driving. Similar methods have been explored in various contexts within simulations [123]. This approach involves the constant driving

field

$$\vec{B}(t) = \frac{\Omega}{\sqrt{3}} \begin{pmatrix} \sqrt{2} \\ 0 \\ 1 \end{pmatrix}. \quad (2.11)$$

In contrast with finite pulses, this method has the effect of gradually rotating the interaction term around the Bloch sphere. It is straightforward to show that after a full period $\tau_{C_1} = 4\epsilon = 2\pi/\Omega$, the generated H_{tar} reproduces H_{XXX} at leading-order in ϵ . The error analysis for the constant field method is less involved than for the pulse sequences, since there is only one parameter ϵ . The engineered Hamiltonian in Eq. (2.2) is nothing but the leading-order term in the Magnus expansion

$$U_{\text{int}}(t) = U_{\text{drive}}(t)^\dagger U(t) = e^{\sum_k \Omega_k(t)}, \quad \tau_{C_1} H_{\text{tar}} = i \Omega_1(\tau_{C_1}). \quad (2.12)$$

Due to the scaling $\Omega_k(\tau) = O(\tau^k)$, the expected error is of order $O(\tau^2) = O(\epsilon^2)$, which is equivalent to a first-order Trotter formula. This method is referred to as C_1 with C standing for constant-field, and the subscript again indicating the error scaling with ϵ . C_1 is shown graphically in the bottom panel of Fig. 2.2.

The methods presented in this chapter for approximating desired Hamiltonians using pulse engineering, as well as the methods described for analyzing the errors incurred, are general and can be applied to any case of Eq. (2.1). The next section describes how to analyze and quantify the errors associated with different time evolution methods.

2.3 Characterization of Error Type and Error Rate

Consider the problem of simulating evolution for a total time T , which is broken up into steps of length τ in physical time. As seen in Fig 2.2, the time it takes to implement a single step on the device varies for each sequence; these are labeled by τ_U for a pulse sequence U . This section analyzes the error accrued during a single step of the sequences described and shows how to optimize the length of this step to minimize the error. The contributions of different sources of error are isolated by considering a time evolution operator U that is a smooth function of two parameters t and ϵ that admits the following Taylor expansion:

$$U(t, \epsilon) = \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} t^{k_1} \epsilon^{k_2} U_{t^{k_1} \epsilon^{k_2}}, \quad U_{t^{k_1} \epsilon^{k_2}} = \frac{1}{k_1! k_2!} \partial_t^{k_1} \partial_\epsilon^{k_2} U(t, \epsilon) \big|_{t, \epsilon=0}. \quad (2.13)$$

The figure of merit that will be used to compare the error scaling of the different evolution methods is the error rate

$$R_U(t, \epsilon) = \frac{1}{\tau_U(t, \epsilon)} \left\| U(t, \epsilon) - e^{-i\tau_U(t, \epsilon)H_{tar}} \right\| \leq \frac{1}{\tau_U(t, \epsilon)} \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} t^{k_1} \epsilon^{k_2} \left\| \mathcal{E}_{U; t^{k_1} \epsilon^{k_2}} \right\|, \quad (2.14)$$

where

$$\mathcal{E}_{U; t^{k_1} \epsilon^{k_2}} = U_{t^{k_1} \epsilon^{k_2}} - \frac{1}{k_1! k_2!} \partial_t^{k_1} \partial_\epsilon^{k_2} e^{-i\tau_U(t, \epsilon)H_{tar}} \Big|_{t, \epsilon=0} \quad (2.15)$$

are defined to be the error terms that contribute at $O(t^{k_1}, \epsilon^{k_2})$. As before, H_{tar} describes the system to be simulated (here it is H_{XXX}), and $\tau_U(t, \epsilon)$ is the optimal step size for a given time evolution U and parameters t and ϵ . In principle, any norm can be used in Eq. (2.14); throughout this chapter the spectral norm induced by the Hilbert space two-norm is used, which is defined by: $\|A\| = \sup_v \frac{\|Av\|_2}{\|v\|_2}$. While the error for a single application of U will always decrease with step size τ_U , multiple applications may cause this error to grow unfavorably. The error accrued per step, the error rate R_U , is used to account for varying step size.

The different terms $\mathcal{E}_{U; t^{k_1} \epsilon^{k_2}}$ can be characterized as follows. Terms with $k_2 = 0$ are the standard Trotter error, as these come strictly from the algorithm used to simulate time evolution (i.e. commutator error from Trotter-Suzuki formulas). The magnitude of this type of error is controlled by t . Terms with $k_1 = 0$ are called “idle errors” and come strictly from the inability to implement ideal single qubit rotations. For the device considered in this chapter, the source of idle error is due to the persistent Ising interaction during $\frac{\pi}{2}$ -pulses. The magnitude of this type of error is controlled by ϵ , which is inversely related to the maximum magnetic-field strength. Terms that correspond to $k_1, k_2 \neq 0$ are referred to as “mixed errors”, and these can be thought of as $O(\epsilon^{k_2})$ idle error that is propagated through the simulation by the $O(t^{k_1})$ action of the time evolution algorithm.

In the Trotter-like digital methods there is freedom in the choice of τ_U . The aforementioned sources of error can balance against one another to give overly pessimistic simulation results. This means that not all choices of τ_U are equally good, and some choices will give more error than others. To be precise, for a given time evolution method U , the optimal step size τ_U is split into

$$\tau_U(t, \epsilon) = 3(t_{\text{rot}} + t_{\text{comp}} + t), \quad (2.16)$$

where t_{rot} corresponds to the evolution due to the rotation gates, t_{comp} is the compensation for unwanted evolution due to those gates, and t is optional evolution time. The factor of 3 comes from mapping

$H_{\text{Ising}} \rightarrow H_{XXX}$ in Eq. (2.3). The optimization of τ_U proceeds as follows. t_{rot} is fixed by the sequence, and so t_{comp} must be applied to match H_{tar} evolution at leading-order in the expansion of Eq. (2.13). t is then chosen to minimize the contribution of terms not canceled by t_{comp} to R_U ; the optimal choice of t is labeled t_* . Here no assumptions are made about the relative smallness of t and ϵ , only that they are small enough to give systematically improvable expansions for each type of error such that Eq. (2.13) may be truncated at finite-order.

The freedom in choosing τ_U for the Trotter sequences lies in the choice of $\vec{t}$ appearing in Eqs. (2.6)-(2.10). Specifically, $\vec{t}$ is chosen so that $t_i = t_{\text{comp},i} + t$ where t is the optional evolution time of Eq. (2.16) (this is expanded upon in Appendix A.3). An important upshot of this relation is that all product formulas have a minimum value of τ_U where $t_i = 0$ for some i (e.g., corresponding to $\tau_{\mathcal{S}_1} = 6\epsilon$). In other words, for given experimental parameters and a given sequence, the step size τ_U cannot be made arbitrarily small. Therefore the optimization of t must be constrained such that τ_U is always larger than its minimal value.

The optimal choice of t (assuming this is greater than the minimal simulation time) is then given by

$$t_{U*} = \arg \min_{t_i > 0} [R_U(t)]. \quad (2.17)$$

For the low-order product formulas considered in this chapter, the solution becomes analytically tractable. Explicit expressions for these bounds are presented in subsequent sections, with derivations in Appendix A.3. For the constant-field method $\mathcal{C}_1$, the step size is completely set by the Floquet period $\tau_{\mathcal{C}_1} = 4\epsilon$ and so this method only incurs (algorithmic) error as a function of ϵ . There is a useful graphical view of the relevant error terms that makes the optimal scaling of t and error rate with ϵ manifest. The error rate in Eq. (2.14) is expanded using Eq. (2.13) and truncated based on the leading-order Trotter and idle errors (i.e. the lowest powers of strictly k_2 or k_1 respectively). This leaves a finite number of terms contributing to the error rate in unique powers of t and ϵ , which can be represented as a curve $k_2 = \mathcal{B}(k_1)$ in the $k_1 - k_2$ plane, as in Fig. 2.3. Specifically, the error rate series becomes

$$R_U(t) \leq \frac{1}{\tau_U(t, \epsilon)} \left(t^{n+1} \|\mathcal{E}_{U;t^{n+1}}\| + \epsilon^m \|\mathcal{E}_{U;\epsilon^m}\| + \sum_{k=1}^{m-1} t^\kappa \epsilon^k \|\mathcal{E}_{U;t^\kappa \epsilon^k}\| \right), \quad (2.18)$$

where κ is defined such that

$$\|\mathcal{E}_{U;t^\kappa \epsilon^k}\| \neq 0, \quad \inf(\|\mathcal{E}_{U;t^{\kappa-1} \epsilon^k}\|, \|\mathcal{E}_{U;t^\kappa \epsilon^{k-1}}\|) = 0 \quad (2.19)$$




Figure 2.3: *Graphical description of leading order error* Different leading-order error terms contributing to the error rate analysis of the four Trotter methods considered in this chapter. The optimal scaling is given by the negative slope of each line at $k_1 = 1$ and the corresponding error rate scaling is given by the $k_1 = 1$ intercept.

for some fixed k . By working in the asymptotic limit, which is defined as $\epsilon \ll 1$ with $t \propto \epsilon^\alpha$ for some number α , the error rate scaling may be analyzed in terms of the individual powers of each of the contributing error terms. The leading-order k_1 error is $\alpha(k_1 - 1) + \mathcal{B}(k_1)$, where 1 is subtracted because $\tau_U(t, \epsilon) = O(t, \epsilon)$. The optimal error rate, for a given choice of α , then scales with ϵ as

$$\sup_{k_1} \alpha(k_1 - 1) + \mathcal{B}(k_1). \quad (2.20)$$

This quantity then needs to be minimized over α . This is a modified Legendre transform of $\mathcal{B}$, where α and k_1 form a conjugate pair. Because of this, the optimal scaling of the error rate is given by $\mathcal{B}(k_1 = 1)$ and the optimal scaling α is given by $-\left. \frac{d\mathcal{B}}{dk_1} \right|_{k_1=1}$. Comparing Fig. 2.3 to the results below demonstrates these properties. An upshot of this analysis is that for curves whose slope at $k_1 = 1$ is undefined (e.g. $\tilde{S}_1$), there are actually a range of optimal scalings between the slopes from the left and right side that give rise to the same optimal error rate.

2.4 Results

2.4.1 Analytic Results

This section gives the results of applying the analysis presented in the previous section to the Hamiltonian engineering methods of Section 2.2. Specifically, $U = \mathcal{S}_{1/2}, \mathcal{S}_1, \mathcal{C}_1$ are used to demonstrate the analysis described above, with details of the calculations contained in Appendices A.1 and A.3. To isolate the contributions of different sources of error $\mathcal{E}$ to the error rate, the expansion of Eq. (2.13) is applied to the two Trotter sequences. A bound on these contributions is then placed by separating the part of the expression depending on the geometry of the system from the operator content. The leading-order contributions to the error are found to be $O(N)$. These are given in Eqs. (A.38) - (A.41).

As described in the previous section, the optimal step size τ_U is found by first choosing t_{comp} to cancel leading-order error terms and optimizing t to minimize the remaining terms at leading or next-leading order, the result of which is shown in Eqs. (A.22) and (A.23). τ_U is expected to vary for different sequences since the error made depends on the sequence used. The magnitudes of the different leading-order contributions to the error are used to find t_* . Subject to the constraints described in the previous section, the optimal t_* values are (up to $O(1)$ factors determined by system geometry, see Eqs. (A.42) and (A.43))

$$t_{\mathcal{S}_{1/2}*} = O(\epsilon^{\frac{1}{2}}), \quad t_{\mathcal{S}_1*} = 0. \quad (2.21)$$

The step size for $\mathcal{C}_1$ is fixed by $\tau_{\mathcal{C}_1}$, so there is no optimization to be done for the constant-field method.

The error rate (Eq. (2.14)) is the main point of comparison between different pulse sequences. It is computed by using the individual error source contributions $\mathcal{E}_{U;\{\dots\}}$ and the optimal step sizes t_* to find the error rate scaling as a function of ϵ for the sequences $\mathcal{S}_{1/2}$ and $\mathcal{S}_1$

$$\begin{aligned} R_{\mathcal{S}_{1/2}} &\lesssim \epsilon^{\frac{1}{2}} \sqrt{\left(\sum_{i \neq j \neq k} J_{ik} J_{kj} \right) \left(\sum_{i \neq j} J_{ij} \right)} = O(\epsilon^{\frac{1}{2}} N), \\ R_{\mathcal{S}_1} &\lesssim \epsilon \left(\sum_{i \neq j \neq k} J_{ik} J_{kj} + \sum_{i \neq j} J_{ij}^2 \right) = O(\epsilon N). \end{aligned} \quad (2.22)$$

Here $O(1)$ factors are again neglected for clarity. The full expressions can be found in Eqs. (A.44) and (A.45). To find the error rate for the constant-field method, the evolution implemented on the device is expanded in a Magnus series and the Heisenberg terms are subtracted off at leading-order to get a bound for $R_{\mathcal{C}_1}$. Dropping

$O(1)$ factors, the error rate for this method is given by:

$$R_{\mathcal{C}_1} \lesssim \epsilon \left(\sum_{i \neq j \neq k} J_{ij} J_{jk} + \sum_{i \neq j} J_{ij}^2 \right) = O(\epsilon N). \quad (2.23)$$

The error rates are found to be extensive quantities, which is not surprising: the number of terms in the Hamiltonian increases with the system size. For comparing systems of different sizes and analyzing local observables, the error rate density is useful: $\frac{R_U}{N}$.

The methods described in this chapter can be used to map between device and target Hamiltonians obeying a simple relation described in Section 2.2. As an example of this, the setup of Ref. [95] can accommodate time evolution methods similar to those presented in previous sections. Their work assumes an experimental system with dipole $XX + YY$ interactions, and tunable X and Y magnetic fields. With this setup, the pulse sequences of Eqs. (2.6), (2.8) - (2.10) implementing Trotterized time evolution carry over without change to emulate an isotropic Heisenberg system. In addition, if a tunable Z magnetic field is assumed, the constant-field method of Eq. (2.11) works as well. Following the analysis for R_U , the method proposed in Ref. [95] has a $O(\sqrt{\epsilon})$ error rate scaling, similar to $\mathcal{S}_{1/2}$, with a slightly different prefactor.

There exist many sequences that have the same error scaling. Examples of this are the sequences $\mathcal{S}_1$ and $\tilde{\mathcal{S}}_1$, which are essentially reordered versions of each other. The method of Ref. [95] and $\mathcal{S}_{1/2}$ also both have $ER_U = O(\sqrt{\epsilon})$. This non-uniqueness comes from the fact that different combinations of $\pm \frac{\pi}{2}$ -rotations may yield the same evolution and similar scaling but with different prefactors due to the combination of the errors incurred (see Appendix A.2). This chapter shows low-lying sequences up to $O(\epsilon^2)$ in the error scaling. It is possible to extend these using the standard Trotter methods [121] to stitch them together to create higher-order sequences. Furthermore, as the subscript on $\mathcal{C}_1$ suggests, it is possible to create higher-order variants of constant-field methods by stitching together lower-order methods in the appropriate manner, e.g., $\mathcal{C}_2(2\epsilon) = \mathcal{C}_1(\epsilon)\mathcal{C}_1(-\epsilon)$.

2.4.2 Numerical Results

The reliance of this analysis on inequalities rooted in the triangle inequality is a source of inflation of the derived upper bounds. In practice, cancellations between different terms give rise to a smaller error overall, compared to the worst-case bounds. Realistic values for the error rate can be found by evaluating Eq. (2.14) numerically to understand the extent to which our bounds ignore these cancellations.




Figure 2.4: *Error rates and bounds* Comparison of analytic bounds for the error rate for a single Trotter step from Eqs. (2.22) and (2.23) to numerical results as a function of the device parameter ϵ for 2×2 systems. The analytic bounds consistently exceed the numerical results by factors of $O(1) - O(10)$, which is an expected consequence of using the triangle inequality to evaluate the bounds.

For concreteness, the physical Hamiltonian underlying experimental systems of Rydberg atoms [124, 125] is used for the rest of the paper, corresponding to $c_x = c_y = 0, c_z = 1$ in Eq. (2.1). Rydberg atoms have an all-to-all coupling mediated by a Van der Waals interaction (i.e. $J_{ij} = \frac{C_6}{|\vec{r}_i - \vec{r}_j|^6}$). This interaction only depends on the geometry of the Rydberg array which is assumed to be fixed for the duration of the simulation. In this chapter, the values of $C_6 = O(10^6)$ MHz μm^6 and the lattice spacing $a = O(1)$ μm are fixed to realistic values based on current hardware parameters [126]. Changing these parameters is equivalent to rescaling the maximum field strength Ω , so numerical calculations are performed with these constants fixed. In reality, devices require some time to change the magnetic fields; this time is determined by the slew rate. This analysis assumes the ideal case of an infinite slew rate for simplicity, although including realistic slew rates does not change the results qualitatively.

The individual contributions $\mathcal{E}$ are compared to the numerical values (obtained by directly evaluating the $\mathcal{E}$ terms in Eq. (2.15)) in Fig. A.1 as a function of system configuration. The results for the error rates of Eqs. (2.22) and (2.23), along with their numerical counterparts, can be seen in Fig. 2.4. In both Fig. 2.4 and Fig. A.1, we see the analytic bounds exceed the numerical values. This happens because the inequalities applied to derive the bounds effectively ignore the cancellations between various terms that happen in practice.




Figure 2.5: *Effects of different step sizes* The effect of multiple Trotter steps on errors for 2×2 systems with $\epsilon \sim 10^{-5}$ and pulse sequences $S_{1/2}$ and S_1 . The quantity $\delta_1 - \delta_n = ||U_{\text{trot}}(T) - e^{-iH_{XXX}T}|| - ||\prod_n U_{\text{trot}}(T/n) - e^{-iH_{XXX}T}||$ is positive when taking many small steps of $U_{\text{trot}}(T/n)$ replicates $e^{-iH_{XXX}T}$ better than taking one large step $U_{\text{trot}}(T)$. Here, $T = n\tau_U$ and τ_U is chosen optimally according to the Eq. (2.21). Taking many steps to reach a total time T is beneficial for $S_{1/2}$ and S_1 compared to taking a single large step when the step size is chosen optimally. This demonstrates the importance of choosing the step size correctly to take advantage of the Trotter error scaling.

So far, the error for a single Trotter step has been considered. Fig. 2.5 investigates whether it is better to take many Trotter steps or a single large step by considering the quantity $\delta_1 - \delta_n = ||U_{\text{trot}}(T) - e^{-iH_{XXX}T}|| - ||\prod_n U_{\text{trot}}(T/n) - e^{-iH_{XXX}T}||$. For a given Trotter time evolution method U_{trot} , this quantity measures how much closer U_{trot} with n steps is to the true time evolution than U_{trot} with one large step. Considering this figure, the effect of the different error contributions on the overall error for each pulse sequence becomes evident: while decreasing ϵ will always decrease the error, the choice of the step size τ_U is of crucial importance to be able to take advantage of the scaling properties of Trotter formulas. In other words, a poorly chosen τ_U may inhibit successive application of the sequence and incur more error than a single large application as a result of constructive interference of different $\mathcal{E}$ contributions. Figs. 2.4 and 2.5 also demonstrate the order-of-magnitude difference between the rate at which errors accrue for $S_{1/2}$ and S_1 .

2.5 Discussion

While simulation of real-life physical systems remains an outstanding goal, the onset of devices capable of supporting analog quantum simulation provides a path forward in the near-term. Because of the level of control and the large number of qubits compared to digital quantum computers, these devices are attractive platforms for attempts at simulating classically inaccessible physics. Control and understanding of errors are important for any quantitative study. As demonstrated in this chapter for the case of analog simulations, the error on real-life devices is not always minimized by trivially minimizing the sources of error individually. Instead, parameters of the simulation must be chosen specifically to maximize the precision of the simulation metrics of interest given the physical limitations of the device. This chapter has outlined and shown examples of the analysis necessary to quantify the error incurred by methods of simulating Hamiltonians using analog quantum devices. Different sources of error and the interplay between them have been investigated, and the growth of these contributions with system parameters has been described. The analytic results provide worst-case guarantees for the methods presented; in practice, numerical simulations of the methods give more realistic reflections of the error. The error rate as a function of minimal $\frac{\pi}{2}$ -pulse duration ϵ is used to compare the effectiveness of several time evolution methods. For methods emulating time evolution under H_{XXX} starting from H_{Ising} , error rates are shown in Fig. 2.6, and relevant features and parameters are given in Table 2.1. Extensions of this analysis method may include using different norms, since typically only a small subset of the Hilbert space is of interest, such as low-energy effective field theory spaces [127]. Furthermore, physically-relevant observables may have additional protection from errors due to their locality [128], and properties of the physical system being simulated may combine with algorithmic and device errors to produce unexpectedly different results [129]. Overall, the considerations presented in here are important for many analog quantum simulations on devices whose evolution is naturally described by Ising-type Hamiltonians such as H_{dev} . The discussion presented is an important step toward a complete quantification of uncertainties for quantum simulation on analog devices.

Current device precision capabilities may limit the ability to implement the time evolution methods that have been presented. To implement a sequence whose error grows as $\mathcal{O}(\epsilon)$, the device should have control of the pulse width at $\mathcal{O}(\epsilon^2)$ or better precision. If such control is not possible, then the errors will wash out any algorithmic scaling assumed by this analysis. Because of this, it is important to consider the level of control that current devices have. Recently, nanosecond level control has been demonstrated on a Rydberg




Figure 2.6: *Error rates for different pulse sequences* Numerical error rates corresponding to Eq. (2.14) for a single Trotter step (Floquet period for C_1) for 2×2 systems. The slopes of lines on these log-log plots indicate the scaling of the error rate for each method (see Table 2.1) and the y-intercepts give the scaling prefactor. The lines for $S_{1/2}$, S_1 , and C_1 are the same as the numerical ones in Fig. 2.4. Current device capabilities report pulse control to the ns level, meaning pulse sequences with $\epsilon \sim O(10^{-5})$ can be applied, indicated by the dashed vertical line.

system [126]. This corresponds to $\epsilon^2 \sim O(10^{-9})$, meaning sequences using $\epsilon \sim O(10^{-5})$ can be supported. This value of ϵ is shown by the dashed vertical line in Fig. 2.6. Although constant-field methods such as C_1 do not require rapid switching of the magnetic field, decreasing ϵ shortens the Floquet period, which makes the engineered time evolution more accurate. In other words, the level of control of the pulse width ϵ limits both types of time evolution engineering methods considered in this work. From Fig. 2.6, it can be seen that at the level of control available today, S_2 outperforms all other considered engineering methods. Unless the minimal pulse width $\epsilon \sim O(10^{-1})$, there is benefit to using higher-order formulas at the cost of rapid magnetic-field switching.

Many devices that implement Hamiltonians of the form of H_{dev} are also able to modify the geometry of the system, such as allowing arbitrary placement of sites in the experimental plane of the apparatus using optical tweezers [130]. The ability to run experiments with such configurations has been proposed to simulate interesting physical phenomena such as nontrivial θ dependence in non-linear σ -models [83], dynamical

Comparison of Time Evolution Methods					
Method	Error rate	τ_U	$\vec{t}_{\text{rot}}$	$\vec{t}_{\text{comp}}$	t_*
$\mathcal{S}_{1/2}$	$O(\sqrt{\epsilon})$	$O(\sqrt{\epsilon})$	$(\epsilon, \epsilon, 2\epsilon)$	$(\epsilon, \epsilon, 0)$	$O(\sqrt{\epsilon})$
$\mathcal{S}_1$	$O(\epsilon)$	6ϵ	$(\epsilon, \epsilon, 2\epsilon)$	$(\epsilon, \epsilon, 0)$	0
$\tilde{\mathcal{S}}_1$	$O(\epsilon)$	6ϵ	$(\epsilon, 2\epsilon, \epsilon)$	$(\epsilon, 0, \epsilon)$	0
$\mathcal{S}_2$	$O(\epsilon^2)$	12ϵ	$(2\epsilon, 2\epsilon, 4\epsilon)$	$(2\epsilon, 2\epsilon, 0)$	0
$\mathcal{C}_1$	$O(\epsilon)$	4ϵ	-	-	-

Table 2.1: The error rates, step size, and step parameters for each method to simulate H_{XXX} evolution. The three rightmost columns show the breakdown of τ_U into the evolution due to the rotations t_{rot} , the evolution to compensate for the rotations t_{comp} , and the optimal optional evolution time t_* . t_* is given by Eqs. (2.17) and (2.21), and following similar analysis for the other Trotter sequences. For $\mathcal{C}_1$ the step size corresponds to the Floquet period $\tau_{\mathcal{C}_1} = 4\epsilon$. The essentially identical properties of $\mathcal{S}_1$ and $\tilde{\mathcal{S}}_1$ and the better error rate scaling of $\mathcal{S}_2$ at the cost of longer device time can be seen.

quantum phase transitions [131], systems with topological order [132], and optimizations for 2+1-dimensional gauge theory simulations [133]. These changes to the geometry would enter the error analysis in the form of modifying the sums in the bounds of Eqs. (2.22) and (2.23).

The interactions that can natively be supported on devices described by H_{dev} are limited. Because of this, Hamiltonian engineering is an attractive way to emulate other models. While this chapter constitutes an important step towards quantifying the errors involved in mapping a physical problem to an analog device, a collective study including other sources of error such as device imperfections and other limitations would be necessary for complete control of uncertainties of an analog quantum simulation. The dynamics of Heisenberg evolution are of broad interest in physics, and analog simulations are enabling their study today. Therefore, a full understanding of the impact of errors and approximations on analog quantum simulations of the Heisenberg model is vital in a path toward scientifically practical results. With the errors under control, data from simulations of systems beyond the reach of classical computers may be used to make predictions

about nature.

Chapter 3

SIMULATING FULLY GAUGE-FIXED SU(2) HAMILTONIAN DYNAMICS ON DIGITAL QUANTUM COMPUTERS

3.1 Introduction

Gauge theories describe a wide array of physical systems from condensed matter and nuclear physics [134–138] to the fundamental interactions of the Standard Model [4–8, 139, 140]. A compelling example is Quantum Chromodynamics (QCD), an SU(3) gauge theory which provides a precise and quantitative description of the strong nuclear force over a broad range of energies. Due to its strongly-coupled nature at low energies, it gives rise to a complex array of emergent phenomena that cannot be identified from the underlying degrees of freedom. Additionally, *ab-initio* calculations of its properties are crucial for comparing theoretical predictions of the Standard Model to experimental results. A fully-systematic approach to carry out computations of QCD in the low-energy regime is the numerical estimation of the lattice regulated path integral via Monte Carlo importance sampling. This approach is computationally expensive, requiring high-performance computing, but has been remarkably successful in computing certain properties of QCD.¹ For example, many single-hadron observables, such as the hadron vacuum polarization contribution to the muon anomalous magnetic moment and the hadron spectrum with QED and isospin breaking effects, can be calculated with sub-percent accuracy [142]. Additionally, several two-hadron observables, such as the decay amplitude for $K \rightarrow \pi\pi$ and direct CP violation can be reliably extracted from lattice QCD simulations [143]. However, these classical simulation methods encounter severe sign problems when applied to real-time dynamics or systems with finite chemical potential.

In recent years, there has been growing interest in utilizing quantum computers to study physical processes that appear out of reach to classical approaches. Quantum simulation provides a powerful alternative framework for exploring large many-body quantum systems, enabling first-principle investigations of out-of-equilibrium dynamics [21–23, 63, 73, 81, 144–164] and the discovery of novel phases of quantum matter [71,

¹For an in-depth review, see Ref. [141]

165–173]. These methods have already been applied to a variety of gauge groups including $\mathbb{Z}_2$ [174, 175], $U(1)$ [156–158, 161, 163, 176–200], $SU(2)$ [62, 164, 201–218], and $SU(3)$ [64–66, 219–227]. Remarkable progress on both the algorithmic [57, 228–231] and hardware [232–237] fronts has pushed the frontier of near-term quantum simulations to systems exceeding 100 qubits [238–242]. These advances have opened the door to studying a broad range of phenomena in strongly interacting theories, such as scattering in high-energy collisions [193, 238, 243–249], processes in nuclear physics [65, 198, 199, 250–255], string-breaking [149, 150, 230, 256, 257] and hadronization and energy loss [258–260].

A key step toward realizing large-scale lattice gauge theory simulations on quantum hardware is the development of efficient, physically-faithful representations of the formally infinite Hilbert space. The fundamental task is encoding the continuous gauge group into discrete degrees of freedom that can be implemented on a quantum device. There exist a variety of approaches that utilize group-theoretic constructions to address this task, including group irreducible representation (irrep) bases [66, 205–207, 209, 219, 261–268], group-element bases [204, 216, 218, 269, 270], discrete subgroups [271–278], deformed gauge groups [214, 224], and dual bases [186, 189, 279–282]. Each of these choices, however, presents its own particular limitations, ranging from inefficiency at weak coupling, complications with systematic improvement, and difficulties in making physically-accurate predictions. Therefore, it is of vital importance to compare these various formulations, with a particular focus on the resources necessary for extracting experimentally-relevant measurements.

A recently introduced basis, the *mixed basis* [216, 218], serves as a middle ground between the irrep and group element representations of the gauge group. With the removal of gauge-redundant degrees of freedom, as well as the use of a coupling-dependent digitization scheme, this basis is believed to provide a resource-efficient formulation of $SU(2)$ for all values of the gauge coupling. This is of particular importance when extracting physical parameters from numerical simulations, due to the necessary tunings of bare couplings to reach the appropriate large volume continuum limits.

The present work develops algorithms for simulating the real-time dynamics of the minimal non-trivial example of a fully gauge-fixed $SU(2)$ Hamiltonian: two plaquettes with open boundary conditions. Sec. 3.2 provides a short review of gauge theory Hamiltonians and gauge fixing. It also details the group element, mixed and character irrep bases, each of which is optimized for distinct coupling regimes. Sec. 3.3 analyzes the mixed basis in detail, presenting the corresponding digitization scheme and mapping of the Hamiltonian into quantum circuits for time evolution. It also provides resource estimates for larger-scale simulations. Sec. 3.4 gives two different methods for translating the digitization into quantum circuits that can be implemented

onto Noisy Intermediate-Scale Quantum (NISQ)-era devices. Finally, Sec. 3.5 presents the results of a real-time simulation on **ibm_fez**, a 156-qubit superconducting chip. An in-depth explanation of the various error-mitigation strategies that are utilized in order to extract results from the noisy quantum simulation is also presented. Finally, Sec. 3.6 discusses the implications of this work and proposes directions for future work and extensions to larger system sizes.

3.2 Two Plaquette Hamiltonian

This section provides a concise overview of how to formulate lattice gauge theories using maximal-tree gauge fixing [283], as well as how to remove the remaining global gauge transformation. It also introduces several different bases that can be used to derive analytic results for a wide range of (bare) coupling values. For a pedagogical review about formulating gauge theories on the lattice and understanding gauge symmetries in the max-tree construction, see Ref. [216]. Operators on the Hilbert space will generally be denoted with a hat, however when the meaning is unambiguous, the hat will be omitted for readability.

3.2.1 Gauge Fixing Lattice Gauge Theories

The Kogut-Susskind Hamiltonian [49] is the starting point of this work. Written in terms of electric field link operators and magnetic plaquette operators, the general form of a lattice gauge Hamiltonian is

$$\hat{H} = \frac{g^2}{2a} \sum_{\ell \in \text{links}} \hat{E}_\ell^2 + \frac{1}{2g^2 a} \sum_{p \in \text{plaq.}} \text{Tr} \left[2 - \hat{P}_p - \hat{P}_p^\dagger \right] \quad (3.1)$$

where each plaquette is constructed out of gauge link operators, g is the gauge coupling, and a is the lattice spacing, usually set to the identity. Before introducing the commutation relations of the operators that appear in this Hamiltonian, it is beneficial to understand how these operators are related to the mathematics of Lie groups.

Typically, a gauge link operator is defined as the path-ordered Wilson line of the gauge field,

$$U_{rs}(\ell) \equiv \left[\mathcal{P} \exp \left(i \int_{\mathbf{n}}^{\mathbf{n}+\mathbf{e}_i} dx^\mu A_\mu(x) \right) \right]_{rs} \quad (3.2)$$

where the gauge link $\ell(\mathbf{n}, \mathbf{e}_i)$ originates at point $\mathbf{n}$ and extends in the direction $\mathbf{e}_i$ and r, s are color indices. The Hilbert space of this operator is the space of all square-integrable wavefunctions $L^2(G, \mathfrak{g})$ over the Lie group G with respect to the Haar measure $d\mathfrak{g}$. Therefore, any gauge link operator can be written in the group




Figure 3.1: *Spectrum of the Two Plaquette System* **a)** Two plaquette system with open boundary conditions. The system contains seven links, numbered 1 through 7, with the physical links denoted as $\kappa = 1, 2$. The magnetic degrees of freedom $\hat{X}_1, \hat{X}_2$ are associated with each plaquette. The dashed lines that run around each plaquette denote the loops used in the max tree gauge fixing procedure. **b)** Spectrum of the Hamiltonian for different bare couplings. The solid lines show the solution obtained by FEM. The black dashed line on the bottom shows the energy bound of Eq. (3.35). The gray lines denote the strong coupling limit of the character irrep spectrum. This limit is found by ignoring the contribution of the magnetic Hamiltonian, which is only valid when the gauge coupling is large. **c)** Low-lying spectrum of the character irrep Hamiltonian in the strong coupling limit. The first and fourth excited state are two-fold degenerate and each state is a linear superposition of the stacked basis states. See Table B.4 for further details.

element basis as

$$\hat{U}_{rs}(\ell) \equiv \int d\mathbf{g}_\ell u_{rs}(\mathbf{g}_\ell) |\mathbf{g}_\ell\rangle \langle \mathbf{g}_\ell| \quad (3.3)$$

where u_{rs} are the matrices that define the group elements in the fundamental representation. The full Hilbert space is simply the tensor product of the local Hilbert spaces of each link.

The electric operators are also defined using the group element Hilbert space. In particular, the electric operators are responsible for translations in the space of group elements,

$$\hat{\Theta}_{L\mathfrak{h}} = e^{i\phi^a(\mathfrak{h})\hat{E}_L^a} \quad \hat{\Theta}_{R\mathfrak{h}} = e^{i\phi^a(\mathfrak{h})\hat{E}_R^a} \quad (3.4)$$

where $\phi^a(\mathfrak{h})$ are parameters typically called normal coordinates, where $a \in \{1, \dots, \dim(G)\}$. The operators $\hat{\Theta}$ act on the group element basis states as

$$\hat{\Theta}_{L\mathfrak{h}} |\mathfrak{g}\rangle = |\mathfrak{h}^{-1}\mathfrak{g}\rangle \quad \hat{\Theta}_{R\mathfrak{h}} |\mathfrak{g}\rangle = |\mathfrak{g}\mathfrak{h}^{-1}\rangle ; \quad (3.5)$$

since the left and right electric field operators act on states differently, both need to be present for non-Abelian gauge theories.

The electric field operators satisfy the commutation relations

$$\begin{aligned} [\hat{E}_L^a, \hat{E}_L^b] &= -if^{abc}\hat{E}_L^c & [\hat{E}_R^a, \hat{E}_R^b] &= if^{abc}\hat{E}_R^c \\ [\hat{E}_L^a, \hat{E}_R^b] &= 0 \end{aligned} \quad (3.6)$$

and given that the left and right electric operators commute, they each furnish an independent Lie algebra. Additionally, each electric operator has a commutation relation with the gauge link operator

$$[\hat{E}_L^a, \hat{U}_{mn}^{(r)}] = T_{mm'}^{(r)a} \hat{U}_{m'n}^{(r)} \quad [\hat{E}_R^a, \hat{U}_{mn}^{(r)}] = \hat{U}_{mn'}^{(r)} T_{n'n}^{(r)a} \quad (3.7)$$

where r indicates the representation and is *not* summed over.

So far, there has been no mention of gauge invariance in this discussion, though all the ingredients are now in place. For a theory that is gauge invariant, there exist states that are equivalent to each other; these states are related via gauge transformations. In deriving the Kogut-Susskind Hamiltonian from the Wilson action, Weyl gauge is imposed. This is an incomplete gauge-fixing condition and purely spatial gauge transformations still remain. A gauge transformation Ω at the lattice site $\mathbf{n}$ is performed by applying the operator

$$\hat{\Theta}_\Omega(\mathbf{n}) \equiv \exp\left(i\phi^a(\Omega)\hat{G}^a(\mathbf{n})\right) \quad (3.8)$$

where $\hat{G}^a$ is the Hermitian operator

$$\hat{G}^a(\mathbf{n}) = \sum_{i=1}^d \left[\hat{E}_R^a(\mathbf{n} - \mathbf{e}_i, \mathbf{e}_i) - \hat{E}_L^a(\mathbf{n}, \mathbf{e}_i) \right] \quad (3.9)$$

and is commonly called the Gauss Law operator. When this operator acts on any state in the physical subspace,

$$\hat{G}(\mathbf{n}) |\Psi_{\text{phys}}\rangle = 0 \quad \forall \mathbf{n} \quad (3.10)$$

is satisfied. It is important that these equivalent states are not treated as distinct. One possible approach for avoiding this is gauge-fixing, a method for choosing a single representative state from each set of equivalent states.

Gauge fixing on a spatial lattice can be done systematically, without the need to solve the Gauss law equation at each site. Instead, the residual spatial gauge transformations that are still present after Weyl gauge-fixing are used to set certain gauge links to the identity. Every time a gauge transformation is used to fix a specific gauge link, neighboring links are also transformed. In two or more spatial dimensions, it is not possible to set all the gauge links to the identity, which is consistent with gauge bosons being dynamical in these dimensions. Ref. [283] showed that the largest set of gauge links that can be set to the identity are those that make up a maximal tree.² This convention removes all of the residual spatial gauge transformations except for those at a single lattice site, which will be taken to be the origin. This last spatial gauge transformation, which is a global gauge transformation, must be treated differently. Once all the spatial gauge redundancy is fully removed, the resultant Hamiltonian only spans the physical subspace and there are no more equivalent states.

This process has been completed for SU(2) lattice gauge Hamiltonians in two- and three-dimensions in Ref. [216, 218]. In the first work, the Kogut-Susskind Hamiltonian is gauge-fixed using the maximal tree gauge fixing procedure and rewritten in terms of new canonically-conjugate operators, $\hat{\mathcal{E}}_{L,R}^a(\kappa)$ and $\hat{X}(\kappa)$. The operators $\hat{\mathcal{E}}_{L,R}^a(\kappa)$ are related to parallel-transported electric link operators and $\hat{X}(\kappa)$ are related to parallel-transported gauge links; both sets of operators are transported along a path $P(\kappa)$. In terms of these

²A maximal tree is a set of all gauge links such that if any additional link was added to the set, the lattice would now contain a closed loop. See Ref. [216] for more details

variables, the Hamiltonian is given by

$$\begin{aligned}\hat{H}_E &= \frac{g^2}{2a} \sum_{\ell} \left(\sum_{\kappa \in t_+(\ell)} \hat{\mathcal{E}}_L^a(\kappa) - \sum_{\kappa \in t_-(\ell)} \hat{\mathcal{E}}_R^a(\kappa) \right)^2 \\ \hat{H}_B &= \frac{1}{2g^2a} \sum_p \text{Tr} \left(\hat{I} - \prod_{\kappa \in p} \hat{X}(\kappa)^{\sigma(\kappa)} \right) + \text{h.c.}\end{aligned}\quad (3.11)$$

where $\sum_p$ and $\sum_{\ell}$ denotes sums over all plaquettes and all links of the lattice, $t_{\pm}(\ell)$ is the set of all physical links such that ℓ is contained in the path P as a positive/negative link and $\sigma = \pm 1$, depending on whether the link is traversed in the positive or negative orientation. The precise form of the Hamiltonian depends on the chosen max-tree and path convention; physical results will be convention-independent.

In addition, Ref. [216] provides a group-element basis for constructing the Hilbert space of the Hamiltonian of Eq. (3.11). Using the axis-angle representation of rotations in three-dimensional space, each loop variable is represented by a two-dimensional matrix

$$\hat{X} = \begin{pmatrix} \cos \frac{\omega}{2} - i \cos \vartheta \sin \frac{\omega}{2} & -ie^{-i\varphi} \sin \vartheta \sin \frac{\omega}{2} \\ -ie^{i\varphi} \sin \vartheta \sin \frac{\omega}{2} & \cos \frac{\omega}{2} + i \cos \vartheta \sin \frac{\omega}{2} \end{pmatrix} \quad (3.12)$$

where the variables have ranges

$$\omega \in [0, 2\pi], \quad \vartheta \in [0, \pi], \quad \phi \in [0, 2\pi] . \quad (3.13)$$

The corresponding canonically-conjugate variables are given by

$$\hat{\mathcal{E}}_{R/L} = \frac{1}{2} \left(\hat{\Sigma} \pm \hat{\mathbf{L}} \right) . \quad (3.14)$$

In axis-angle coordinates, $\hat{\mathbf{L}}$ is given by

$$\begin{aligned}\hat{L}^x &= i \left(\sin \varphi \frac{\partial}{\partial \vartheta} + \cot \vartheta \cos \varphi \frac{\partial}{\partial \varphi} \right) \\ \hat{L}^y &= i \left(-\cos \varphi \frac{\partial}{\partial \vartheta} + \cot \vartheta \sin \varphi \frac{\partial}{\partial \varphi} \right) \\ \hat{L}^z &= -i \frac{\partial}{\partial \varphi} ,\end{aligned}\quad (3.15)$$

which is simply the (conventional) angular momentum operator. The operator $\hat{\Sigma}$ can be written in terms of $\hat{\mathbf{L}}$ and the position vector $\hat{\boldsymbol{\eta}}$ as

$$\hat{\Sigma} = 2i\hat{\boldsymbol{\eta}} \frac{\partial}{\partial \omega} + \cot \frac{\omega}{2} \left(\hat{\boldsymbol{\eta}} \times \hat{\mathbf{L}} \right) , \quad (3.16)$$

where $\hat{\eta}$ is the axis of rotation in three-dimensional space,

$$\hat{\eta} \equiv (\sin \vartheta \cos \varphi, \sin \vartheta \sin \varphi, \cos \vartheta) . \quad (3.17)$$

The operator $\hat{\Sigma}$ is analogous to the boost operator of the Lorentz group. The commutation relations of these two operators are

$$\begin{aligned} [\hat{L}^a, \hat{L}^b] &= i\epsilon^{abc} \hat{L}^c \\ [\hat{\Sigma}^a, \hat{\Sigma}^b] &= i\epsilon^{abc} \hat{L}^c \\ [\hat{\Sigma}^a, \hat{L}^b] &= i\epsilon^{abc} \hat{\Sigma}^c , \end{aligned} \quad (3.18)$$

whereas the commutation relations for the Lorentz group operators have a relative minus sign for the second and third commutation relation. The Haar measure for each link in this basis is

$$d\mathbf{g} = 4 \sin^2 \frac{\omega}{2} \sin \vartheta d\omega d\vartheta d\varphi \quad (3.19)$$

The fixing of the last (global) gauge transformation is carried out in Ref. [218]. The key insight of that work is that the Hamiltonian can be conceptualized as a system of rods all connected together at the origin. The κ^{th} rod has length ω_κ and its orientation is given by the polar angle ϑ_κ and azimuthal angle φ_κ . Such a system can be written either in a lab-frame coordinate system or in a body-frame coordinate system. An Euler rotation relates these two coordinate systems. In the lab-frame coordinate system, the location of each rod is parameterized as in Eq. (3.17). This same configuration of rods can be written using body-frame variables, coupled with an Euler rotation. The definition of the body frame is arbitrary, but Ref. [218] chooses to define the $\hat{z}$ axis of the body frame by aligning it with one rod and the $\hat{x} - \hat{y}$ plane via the location of another rod. In this coordinate system, the location of each rod is given by

$$\begin{aligned} \boldsymbol{\eta}_1 &= \mathbf{R} \cdot (0, 0, 1) \\ \boldsymbol{\eta}_2 &= \mathbf{R} \cdot (\sin \Theta, 0, \cos \Theta) \\ \boldsymbol{\eta}_\mu &= \mathbf{R} \cdot (\sin \theta_\mu \cos \phi_\mu, \sin \theta_\mu \sin \phi_\mu, \cos \theta_\mu) \quad \mu \geq 3 \end{aligned} \quad (3.20)$$

where $\mathbf{R}$ is the rotation matrix written in terms of Euler angles. Rotation of the full system leaves the lengths of each rod invariant. In carrying out the change of basis to go from the lab-frame coordinate system to the body-frame coordinate system, the remaining global gauge transformation is made visible. The Hamiltonian

can then be trivially fixed to the total gauge-singlet sector and only depends on the variables $\{\omega_\kappa, \Theta, \theta_\mu, \phi_\mu\}$; all dependence on the Euler angles is eliminated from the Hamiltonian. Ref. [218] provides the explicit forms of all possible electric bilinears as well as magnetic loop operators in this new coordinate system. The new basis is called the ‘sequestered basis’, as all the dependence on the total charge of the system is *sequestered* into the Euler angles. The focus of this chapter is to simulate the two plaquette version of this system, on a NISQ-era device.

3.2.2 Two Plaquette Hamiltonian

Using Eq. (3.11) and a loop convention where each loop only encompasses one plaquette, the Hamiltonian for the two plaquette system with open boundary conditions, shown in Fig. 3.1a, is given by

$$\begin{aligned} \hat{H}_{\square\square} = & \frac{g^2}{2} \left(4\hat{\mathcal{E}}_1^2 + 6\hat{\mathcal{E}}_2^2 - 2\hat{\mathcal{E}}_{R2} \cdot \hat{\mathcal{E}}_{L2} + 2\hat{\mathcal{E}}_{R1} \cdot \hat{\mathcal{E}}_{R2} - 4\hat{\mathcal{E}}_{R1} \cdot \hat{\mathcal{E}}_{L2} \right) \\ & + \frac{1}{g^2} \left(4 - \text{tr } \hat{X}_1 - \text{tr } \hat{X}_2 \right). \end{aligned} \quad (3.21)$$

The lattice spacing has been set to the identity, as is conventionally done, and

$$\hat{\mathcal{E}}_{L\kappa} \cdot \hat{\mathcal{E}}_{L\kappa} = \hat{\mathcal{E}}_{R\kappa} \cdot \hat{\mathcal{E}}_{R\kappa} \equiv \hat{\mathcal{E}}_\kappa^2, \quad (3.22)$$

where κ is not summed over. A basis that spans the Hilbert space of $\hat{H}_{\square\square}$ must be defined to analyze the real-time behavior of this system. Additionally, both $\hat{\mathcal{E}}_\kappa$ and $\hat{X}$ are infinite-dimensional operators and the basis needs to be truncated. The calculation of any physical observable is basis-independent but truncation will affect this universality. The effects should disappear as the truncation is removed. This chapter utilizes three different bases to measure both static and dynamic properties.

Group Element Basis

In the group element basis, the eigenstates $|\mathbf{g}\rangle$ of the gauge link operators (or the loop operators) are the basis states. Any parameterization of $|\mathbf{g}\rangle$ can be used for defining the action of the Hamiltonian on the Hilbert space but some parameterizations are advantageous from the perspective of digitization and global gauge fixing. For $\text{SU}(2)$, two natural parameterizations are Euler angles or axis-angle coordinates. The axis-angle parameterization is particularly convenient [216, 218]. Any wavefunction that spans only the physical Hilbert space *i.e.* any gauge-singlet wavefunction is given by

$$\langle \omega_1 \omega_2 \Theta | \Psi = \Psi(\omega_1, \omega_2, \Theta)$$

$$\equiv \frac{u(\omega_1, \omega_2, \Theta)}{4 \sin \frac{\omega_1}{2} \sin \frac{\omega_2}{2}} \quad (3.23)$$

where the rescaled wavefunction $u(\omega_1, \omega_2, \Theta)$ is introduced to simplify later calculations. The magnetic component of the fully gauge-fixed Hamiltonian is diagonal and only depends on the radial coordinates ω_1 and ω_2 ,

$$H_{\square\square}^{(B)} = \frac{2}{g^2} \left(2 - \cos \frac{\omega_1}{2} - \cos \frac{\omega_2}{2} \right). \quad (3.24)$$

The electric component of the Hamiltonian can be determined by evaluating each of the electric bilinears that appear in Eq. (3.21). The explicit expressions for the electric bilinears are relegated to Appendix B.2.1.

Combining the various elements together, the electric component of the Hamiltonian is given by

$$\begin{aligned} H_{\square\square}^{(E)} = & -\frac{g^2}{2} \left[\left\{ 4 \left(\frac{\partial^2}{\partial \omega_1^2} + \cot \frac{\omega_1}{2} \frac{\partial}{\partial \omega_1} \right) - 2 \csc^2 \frac{\omega_1}{2} \mathcal{N} + \left(\cot \frac{\omega_2}{2} \frac{\partial}{\partial \omega_1} \right) \left(\sin \Theta \frac{\partial}{\partial \Theta} \right) + (\omega_1 \leftrightarrow \omega_2) \right\} \right. \\ & \left. + \left(\frac{1}{2} \cot \frac{\omega_1}{2} \cot \frac{\omega_2}{2} \right) \left(\sin \Theta \frac{\partial}{\partial \Theta} \right) - \left(\frac{1}{2} \cos \Theta \cot \frac{\omega_1}{2} \cot \frac{\omega_2}{2} - \frac{1}{2} \right) \mathcal{N} \right] \end{aligned} \quad (3.25)$$

where $\mathcal{N}$ is the differential operator,

$$\mathcal{N} = -\frac{\partial^2}{\partial \theta^2} - \cot \theta \frac{\partial}{\partial \theta}. \quad (3.26)$$

In this basis, $H_{\square\square}$ is a partial differential operator and its spectrum can be found numerically using Finite Element Methods (FEM). FEM is a general numerical approach for solving low-dimensional partial differential equations that utilizes a discrete mesh of the computational domain.³ In order to ensure that the eigenfunctions of $H_{\square\square}$ are normalizable and that the mesh does not encounter any singularities, FEM is applied to solve for the rescaled wavefunctions, $u(\omega_1, \omega_2, \Theta)$. In this case, the radial coordinates, $\omega_{1,2}$ have a range of $[0, 2\pi]$ with Dirichlet boundary conditions while the angular coordinate Θ has range $[0, \pi]$ with Neumann boundary conditions. Fig. 3.1b plots the low-energy spectrum as a function of the bare gauge coupling. These solutions provide a beneficial ‘exact’ solution to which the mixed basis and character irrep basis can be compared. This is particularly helpful, as while these bases can be used for larger system sizes, the FEM approach faces insurmountable computational barriers when applied to larger systems. Without fully gauge-fixing, FEM could not even be applied to the two-plaquette system as, for example, the (naively parameterized) two-plaquette version of the Kogut-Susskind Hamiltonian involves twenty-one parameters (three quantum numbers for each of the seven links).

³This is in contrast to Finite Difference Methods (FDM), which approximate differential operators as finite differences on discrete grid.

Mixed Basis

In the mixed basis, the states that span the Hilbert space are labeled by the continuous variables ω_1, ω_2 and the discrete quantum number ν . This discrete quantum number is related to the angular variable Θ via

$$\langle \Theta \rangle \nu = P_\nu(\Theta) \quad (3.27)$$

with $P_\nu(\Theta)$ the (rescaled) Legendre polynomial of degree ν .⁴ Legendre polynomials are used to span the Hilbert space because they are the eigenfunctions of the operator $\mathcal{N}$,

$$\mathcal{N}P_\nu(\theta) = \nu(\nu + 1)P_\nu(\theta). \quad (3.29)$$

The matrix elements in this basis are found by calculating

$$\begin{aligned} \langle \omega'_1 \omega'_2 \nu' | \hat{H}_{\square\square} | \omega_1 \omega_2 \nu \rangle &= \int d(\cos \Theta) d(\cos \Theta') \langle \nu' \rangle \Theta' \\ &\times \langle \omega'_1 \omega'_2 \nu' | \hat{H}_{\square\square} | \omega_1 \omega_2 \nu \rangle \langle \Theta \rangle \nu \\ &= \int d(\cos \Theta) P_{\nu'}^*(\Theta) H_{\square\square} P_\nu(\Theta) \end{aligned} \quad (3.30)$$

where $H_{\square\square}$ is the sum of the magnetic and electric Hamiltonians, written in the group element basis in Eq. (3.24) and Eq. (3.25).

To convert between the group element basis and the mixed basis, only four integrals need to be calculated:

$$\int d(\cos \Theta) P_{\nu'}(\Theta) \left\{ \begin{array}{c} 1 \\ \mathcal{N} \\ \cos \Theta \\ \sin \Theta \frac{\partial}{\partial \Theta} \end{array} \right\} P_\nu(\Theta). \quad (3.31)$$

⁴The conventional normalization of the Legendre polynomials is

$$\int_0^\theta d(\cos \theta) P_{\nu'}(\theta) P_\nu(\theta) = \frac{2}{2\nu + 1} \delta_{\nu'\nu} \quad (3.28)$$

but due to the requirement that the basis states of the Hilbert space be orthonormal, the Legendre polynomials in this work are rescaled.

The results of these integrals, tabulated in Appendix B.2.2, are used to recast expressions in $H_{\square\square}$. The necessary digitization, particularly of the components that depend on ω_κ , will be discussed below in Sec. 3.3.

Using the mixed basis, it is possible to derive a strong bound on the ground state energy of the two plaquette system. Recall that the ground state energy of a Hamiltonian, E_0 , can be bounded from above by the expectation value of that Hamiltonian in any state Ψ ,

$$E_\Psi \equiv \langle \Psi | \hat{H} | \Psi \rangle \geq E_0. \quad (3.32)$$

In general, the goodness of this bound depends on how well the wavefunction Ψ matches the true ground state. A good ansatz for the two plaquette system is found by assuming that the ground state does not depend strongly on Θ and therefore the functional dependence of the ground state on ω_κ is encapsulated in the differential operator

$$-2g^2 \left(\frac{\partial^2}{\partial \omega_\kappa^2} + \cot \frac{\omega_\kappa}{2} \frac{\partial}{\partial \omega_\kappa} \right) + \frac{2}{g^2} \left(1 - \cos \frac{\omega_\kappa}{2} \right), \quad (3.33)$$

which is the component of the Hamiltonian that is diagonal in ν for each ω . This differential equation is related to that of the Mathieu equation. Therefore, the ansatz for the ground state is

$$u(\omega_1, \omega_2, \Theta) = \frac{\text{se}_2(-8/g^4, \omega_1/4) \text{se}_2(-8/g^4, \omega_2/4)}{\sqrt{2}} \quad (3.34)$$

where Ψ and u are related via Eq. (3.23). The parity odd Mathieu function, $\text{se}_n(q, x)$ has period π when n is even or 2π when n is odd. For this system, n has to be even for ψ to be finite on the full range of ω_κ . With this ansatz, E_Ψ is

$$\begin{aligned} E_\Psi &= \frac{g^4 b_2(-8/g^4) - 4g^4 + 16}{4g^2} \\ &= \begin{cases} 3\sqrt{2} & g \ll 1 \\ 4/g^2 & g \gtrsim 1 \end{cases} \end{aligned} \quad (3.35)$$

where $b_n(q)$ is the characteristic number that guarantees that $\text{se}_n(q, x)$ has period π or 2π . Comparing this ansatz to the FEM solution demonstrates that there is very good agreement between the two for a wide range of gauge couplings.

Character Irrep Basis

The Kogut-Susskind Hamiltonian is formulated in the group irrep basis, also called the electric basis. The eigenstates of the single link Hilbert space are $|JM_L M_R\rangle$ where J is the total angular momentum and M_L

(M_R) is the angular momentum projected onto the $\hat{z}$ axis of the lab (body) frame; J is also commonly referred to as the irrep of $SU(2)$ group. In analogy, the basis of the Hilbert space of the fully gauge-fixed two-plaquette Hamiltonian can be chosen to be $|J_1 J_2 \nu\rangle$ where $J_{1,2}$ are again irreps of $SU(2)$. The relationship between the group element basis and this basis is

$$\langle \omega_1 \omega_2 \Theta | J_1 J_2 \nu \rangle = \chi_\nu^{J_1}(\omega_1) \chi_\nu^{J_2}(\omega_2) P_\nu(\Theta) \quad (3.36)$$

where χ_ν^J are (rescaled) generalized character functions of the irreps of the rotation group. These functions satisfy the differential equation,

$$\left[\frac{\partial^2}{\partial \omega^2} + \cot \frac{\omega}{2} + J(J+1) - \frac{\nu(\nu+1)}{4 \sin^2 \frac{\omega}{2}} \right] \chi_\nu^J(\omega) = 0 \quad (3.37)$$

where J are integers and half-integers and ν are integers, bounded by $2J \geq \nu \geq 0$. An explicit construction of these character functions is done via Clebsch-Gordan coefficients,

$$\chi_\nu^J(\omega) = \frac{i^\lambda}{2\sqrt{\pi}} \sum_M e^{-iM\omega} C_{JM\nu 0}^{JM} \quad (3.38)$$

where the prefactor of $1/(2\sqrt{\pi})$ arises due to the requirement that the basis states are orthonormal.

Relevant matrix elements can be calculated for generic J_1, J_2 and ν by making use of recurrence relations for the character functions [284]. The complete expressions are presented in Appendix B.2.3. A key result is that the magnetic component of the Hamiltonian only changes the J_κ quantum numbers by at most $\pm 1/2$. In contrast, the electric component of the Hamiltonian only changes the ν quantum number by at most ± 1 . With these matrix elements, it is possible to construct a finite-dimensional Hamiltonian, with cutoff $J_\kappa^{\max}$. Note that ν is naturally truncated due to the requirement that

$$2 \text{Min}(J_1, J_2) \geq \nu \geq 0 \quad (3.39)$$

and since the Hamiltonian is symmetric under $J_1 \leftrightarrow J_2$ exchange, as is evident from the differential form of the Hamiltonian, both J_1 and J_2 have the same cut-off.

Solving for the spectrum in the strong coupling limit is simple, as the electric component preserves the quantum numbers J_1, J_2 and the contributions from the magnetic component can be ignored. Additionally, since the quantum number ν is bounded from above by $2 \text{Min}(J_1, J_2)$, for a given J_1, J_2 , the relevant sub-block of the Hamiltonian is finite-dimensional. In short, in the strong coupling limit, at leading order, the Hamiltonian is block-diagonal. The results for the lowest states are shown in Fig. 3.1c and in Table B.4.

This basis can be used for other values of the gauge coupling if the magnetic component of the Hamiltonian is included. In this case, the Hamiltonian is no longer diagonal in $J_{1,2}$ and its spectrum cannot be found analytically. Using exact diagonalization for different values of $J_{\kappa}^{\max}$, it can be shown that at small values of the gauge coupling, the energy eigenstates diverge as $1/g^2$. However, using FEM, it is shown that the low-lying eigenstates become independent of g in the limit of small g . These large truncation errors at weak coupling are why electric bases are disfavored for this regime.

3.3 Digitization of the Mixed Basis

The Hilbert space degrees of freedom that may be represented by a register of qubits is finite and so a finite representation of the continuous gauge field must be determined. This is commonly referred to as a ‘digitization’ (distinct from ‘discretization’, which corresponds to representing a continuous space-time on a lattice). This section focuses on the manner in which the mixed basis can be digitized, in particular the derivatives with respect to ω . Of key importance is how to efficiently implement derivative operators onto a quantum device, which will be the focus of the next section.

3.3.1 Conventions

In the following sections, various ideas and tools will appear in slightly different contexts. Therefore, to ease notation and presentation, the following conventions are made.

For a single qubit, one may define the set of Pauli operators

$$\sigma_{\alpha} = \{I, X, Y, Z\} . \quad (3.40)$$

For N -qubits, the set of all Pauli strings (modulo phase) is the Pauli group, and is defined as

$$\mathcal{P}_N = \{ \otimes_{l=1}^N \sigma_l^{\alpha_l} \mid \alpha_l = 0, 1, 2, 3 \} . \quad (3.41)$$

This group defines a complete basis for the set of all Hermitian operators O (e.g. a Hamiltonian or density matrix) on N qubits, and is normalized as

$$\text{Tr} PP' = 2^N \delta_{P,P'} \quad \forall P, P' \in \mathcal{P} . \quad (3.42)$$

Given such a Hermitian operator O on N qubits, its (normalized) Pauli decomposition is denoted as

$$\mathcal{D}(O) = \left\{ \left(P, \frac{1}{2^N} \text{Tr} PO \right) \mid P \in \mathcal{P}_N \right\} . \quad (3.43)$$

$\mathcal{D}(O)$ will be used to either denote the set of pairs as shown above or the sum itself, where the meaning will be clear based on context.

3.3.2 Formulating Gauge Field Operators

An important feature of a good digitization is that observables (e.g. expectation values of Hermitian operators) converge quickly to the continuum limit as the resolution is increased. For fast convergence, the range of support of the wavefunctions of interest must be fully covered in both field and conjugate momentum space. In particular, if the support of the wave function is not sufficiently covered then predictions for observables will not converge to the correct solution. If the region of field space covered is much larger than the support of the wave function, then the digitization will sample regions of the wave function that are exponentially small and the convergence will be slow. The following digitization achieves good convergence properties by optimally spanning the regions of wave function support.

The continuous gauge field degrees of freedom represented by $\{\omega_1, \omega_2\} \in [0, 2\pi)^{\otimes 2}$ must be embedded into a register of $2n_q$ qubits with Hilbert space dimension 2^{2n_q} . This is done by letting each ω_i take values on a grid of size 2^{n_q} with the range $[0, \omega_{\max})$. The discretized variables are denoted as ω_i , for $i = 1, 2$ and assigned values

$$\omega_i = \delta\omega(n + 1/2) \quad n \in \{0, \dots, 2^{n_q} - 1\} \quad \delta\omega = \omega_{\max}/2^{n_q} . \quad (3.44)$$

When performing classical simulations when the system is not embedded into qubits, the total grid size is denoted as N_ω . In addition to the two ω -registers, there is an additional quantum number ν that takes integer values and requires an additional n_ν qubits, giving a total simulation Hilbert space of size $2^{2n_q + n_\nu}$. This layout is graphically illustrated in Fig. 3.2.

There are three types of terms that appear in the Hamiltonian, namely the first-order derivatives of ω_i , the second-order derivatives of ω_i and functions of ω_i . Their action on the qubit register are treated separately. The field variable ω_i is defined by the operator

$$\hat{\omega} = \frac{\delta\omega}{2} \left[2^{n_q} I - \sum_{j=0}^{n_q-1} 2^j \hat{Z}_j \right] \quad (3.45)$$

where $\hat{Z}_j$ is the Pauli $\hat{Z}$ operator acting on the j th qubit. Functions of $\hat{\omega}$ can then be constructed as $\mathcal{D}(f(\hat{\omega}))$,

explicitly

$$f(\hat{\omega}) = \sum_{P \in \{I, Z\}^{\otimes n_q}} f_P P$$

$$f_P = \frac{1}{2^{n_q}} \text{Tr}[f(\hat{\omega}) P] \quad (3.46)$$

where the sum above is over all strings of Pauli Z operators. The second-order derivatives can be implemented using the exact lattice Laplacian with Dirichlet boundary conditions, which is given by

$$\frac{\partial^2}{\partial \omega^2} \mapsto \text{DST}_{II}^{-1}(-\hat{k}^2) \text{DST}_{II} \quad (3.47)$$

where DST_{II} is a variation of the Discrete Sine Transform (DST), reflecting the boundary conditions in field space:

$$(\text{DST}_{II})_{kn} = \begin{cases} \sqrt{\frac{2}{N}} \sin \frac{\pi(n+\frac{1}{2})(k+1)}{N}, & \text{for } k \neq N-1 \\ \sqrt{\frac{1}{N}} \sin \pi(n+\frac{1}{2}), & \text{for } k = N-1 \end{cases}. \quad (3.48)$$

Note that $\hat{k}$ is the lattice momentum operator, which is diagonal in Fourier space with eigenvalues $\pi k / \omega_{\max}$ with $k \in \{1, \dots, N\}$. It has been shown that when using this exact formulation, the matrix elements of this term converge faster than polynomial to the continuum value as the grid spacing $\delta\omega$ goes to zero [27, 28, 112, 285]. Therefore it is advantageous to use compared to a finite difference approach.

The first-order derivative is, however, formulated using the finite central difference

$$\left(\frac{\partial}{\partial \omega} \right)_{kn} = \begin{cases} \frac{1}{2\delta\omega}, & \text{for } k - n = 1, \\ -\frac{1}{2\delta\omega}, & \text{for } k - n = -1, \\ 0 & \text{else} \end{cases} \quad (3.49)$$

The consequence of using the finite difference is that the digitized form of $H_{\square\square}$ will only have polynomial convergence towards the continuum Hamiltonian as the spacing in field space is decreased. However, in the context of quantum simulation, even for large $\delta\omega$ this error will generally be subleading to algorithmic errors stemming from approximations to continuous time evolution. While it is desirable to use spectral methods for the first-order derivative, as was done for the Laplacian, the Dirichlet boundary conditions complicate such an approach. This can be understood intuitively because the DST acts on data with odd symmetry, e.g. a sine wave. A first derivative maps functions with odd symmetry to functions with even symmetry, and thus is

incompatible with the use of a DST. Other types of spectral methods for first order derivatives with Dirichlet boundary conditions have been formulated, but often require non-uniform grid spacings, for example using Chebyshev nodes [286]. This approach is undesirable as such grid choices have denser coverage towards the boundaries of the field region. This causes both oversampling of the singularities of the Hamiltonian as well as asymmetric sampling of the wavefunctions for low g .

3.3.3 Optimal Gauge Field Truncation

In analogy with digitizations of scalar fields, the gauge field is truncated at some value $\omega_{\max} \leq 2\pi$. Fig. 3.3a shows the amplitude of the ground state wave function

$$u_0(\omega_1, \omega_2) = |\langle \omega_1, \omega_2 | \psi_{\text{gs}} \rangle|^2 \quad (3.50)$$

along the ω_1 axis ($\omega_2 = \pi/2$) for different coupling strengths. As the gauge coupling is decreased, u_0 becomes increasingly localized at small ω_1 , decaying exponentially for large values. Therefore, the optimal $\omega_{\max}$ must be characterized as a function of gauge coupling. This is done by considering the precision

$$\epsilon_0 = \left| 1 - \frac{E'_0}{E_0} \right| \quad (3.51)$$

where E'_0 is the ground state energy predicted by the digitization and E_0 is the ‘true’ ground state energy. Similar results are observed for low-lying eigenstates. The solutions for the digitization are compared to those found via FEM methods in Fig. 3.1b, the strong coupling solutions of Sec. 3.2.2 and the ansatz of Eq. (3.34), and are in good agreement. The solutions found via FEM are of particular use, and they have a relative precision of approximately 10^{-4} .

Fig. 3.3b shows the precision for different choices of $\omega_{\max}$ truncation as a function of digitization resolution N_ω . The convergence is characterized by three stages. The first stage is rapid convergence to $\epsilon_0 \sim 10^{-3}$ level for small N_ω . Following this, the precision decays more slowly, though still exponentially, as N_ω is increased. Eventually, a plateau is reached. The precision value of this plateau decreases as $\omega_{\max}$ increases. There are sharp spikes observed before the plateau is reached, which are attributed to the digitized solution crossing the FEM solution, as opposed to approaching it. Such crossings are purely incidental and should not be interpreted as signs of optimality in representation of the wavefunction. Since the precision of the FEM solution is limited, it is not possible to study convergence below precision levels of 10^{-4} . Despite this it is clear that even for grid sizes as low as $N_\omega = 8$ ($n_q = 3$) it is possible to reach $\epsilon_0 \sim 10^{-3}$ levels of




Figure 3.2: *Digitization of the Wavefunction* The continuous range of ω is digitized on a discrete grid. As such the wave function will only be sampled at a discrete set of points. For the closest approximation to the undigitized theory, the maximum value of $\omega_{\max} \leq 2\pi$ must be chosen.

precision. This is at least an order of magnitude below algorithmic error due to time-evolution and on par with device noise-levels for near-term devices.

The optimal truncation is found by scanning over a range of $\omega_{\max}$ values for a fixed g , N_ω and calculating the corresponding ϵ_0 . The top plot of Fig. 3.3c shows the precision ϵ_0 for $H_{\square\square}$. Also considered is the ‘free’ effective theory of $H_{\square\square}$, where terms containing a first derivative are subtracted off. The theory is not completely free, as there is still a collective ν quantum number, but the interactions due to the first derivatives are no longer present. This theory is studied to understand how ϵ_0 behaves without the error induced by first derivatives. The ground state of the latter Hamiltonian can be found exactly and is given by the ansatz in Eq. (3.34). The corresponding energy is Eq. (3.35). For the system without first derivatives, the precision decays exponentially in $\omega_{\max}$ until the optimal value is reached, and then increases due to over-sampling regions of field space where the wave function has no support. However, once first derivatives are included, the precision plateaus at a value orders of magnitude higher, and can only be increased by increasing N_ω . Thus there is a range of choices of $\omega_{\max}$. The sharp dips are again due to crossings and should be ignored when considering the optimality of $\omega_{\max}$.




Figure 3.3: *Convergence properties of the two plaquette wavefunction* **a)** A slice of the ground state wave function for different values of the gauge coupling. As the gauge coupling decreases, the wave function becomes more concentrated around $\omega_1 = 0$, although it must go to 0 at both boundaries due to the Dirichlet boundary conditions. **b)** Convergence of the ground state energy to the undigitized value as the digitization resolution N_ω is increased. The convergence is plotted for a variety of $\omega_{\max}$. After quick convergence to $\epsilon_0 \sim 10^{-3}$ a slower exponential approach occurs, followed by a plateau. **c)** (*top*) Convergence to the ground state energy as $\omega_{\max}$ is varied. The dashed lines correspond to the full two plaquette system while the solid lines correspond to the system without any first-derivative interactions. The effect of the first derivative reduces the precision that can be achieved. Even though this term limits precision, digitization errors are still below those due to time-evolution algorithms. (*bottom*) The absolute value of the derivative of the energy w.r.t the truncation. It was heuristically found that the truncation at which this derivative takes a minimum corresponds to the optimal truncation for the analytically solvable case. This method is still applicable even when the exact solution is not known. **d)** The solution of the stationary condition for a range of gauge couplings. The agreement is very good with the analytical solution Eq. (3.53).

The gauge field truncation can be optimized without knowing the exact value of the ground state energy by looking at the derivative of the energy with respect to $\omega_{\max}$. It is empirically observed that the optimal $\omega_{\max}$ can equivalently be defined by the stationary condition

$$\omega_{\max} = \arg \min_{\omega} \left| \frac{\partial E_i(\omega)}{\partial \omega} \right|. \quad (3.52)$$

This condition can also be observed for the digitization of the scalar field. Fig. 3.3d (bottom) shows the derivative of the ground state energy predicted by the digitization. For the system without first derivatives, the minimum of these curves align with the points where the ground state energy most closely approaches the expression Eq. (3.35). When first derivatives are included, the precision plateaus at low values for a range of $\omega_{\max}$. The optimal truncation in the absence of first derivatives is found by scanning over a range of gauge couplings and finding the solution of stationary condition, Eq. (3.52). Fig. 3.3d shows that the numerical solution found by the scan closely matches the analytical solution [216]

$$\omega_{\max} = \min \left(g(N-1) \sqrt{\frac{\sqrt{8}\pi}{N}}, 2\pi \right) \quad (3.53)$$

which was found by minimizing $\langle [\partial/\partial\omega, \omega] - 1 \rangle$ for the ground state of a single plaquette.

In addition to the truncation of the gauge field, the discrete quantum number ν must also be truncated. An advantage of the mixed basis is that the low-lying eigenstates are supported solely in the $\nu = 0, 1$ sector, and so only a single qubit is required to represent this quantum number. Furthermore, $H_{\square\square}$ only connects states that have $|\nu - \nu'| \leq 1$, and so for low-energy simulations heavy truncations on this degree of freedom will have little effect on the precision. To study the effect of different ν truncations on the time evolution of observables, the system is initialized in the state

$$|5\rangle = \frac{1}{\sqrt{5}}(|E_0\rangle + |E_1\rangle + |E_2\rangle + |E_3\rangle + |E_4\rangle) \quad (3.54)$$

where $|E_n\rangle$ is the n^{th} eigenstate of the digitized Hamiltonian and then time evolved according to different maximum ν values $\nu_{\max}$. Fig. 3.4a shows the time dependence of the expectation value for the magnetic part of the Hamiltonian H_B . As $\nu_{\max}$ is increased, the values of $\langle H_B \rangle$ rapidly converge to the same curve, suggesting that $\nu_{\max} = 2$ is already sufficient for extracting high-precision estimates of observables.

The leakage out of the $\nu = 0, 1$ sectors is also studied. Defining the projector onto states with $\nu > 1$ as

$$\Pi_{\nu>1} = \mathbb{I} - \sum_{J_1, J_2} (|J_1 J_2 0\rangle \langle J_1 J_2 0| + |J_1 J_2 1\rangle \langle J_1 J_2 1|), \quad (3.55)$$

the expectation value $\langle \Pi_{\nu>1} \rangle$ encodes the probability that the state has leaked out of the low-energy $\nu = 0, 1$ subspace. Fig. 3.4b shows the leakage probability as a function of time. As the state evolves, the leakage does not grow noticeably over time, showing that a low-energy state with little support on values of $\nu > 1$ will not evolve to states with larger support over a range of times.




Figure 3.4: *Accuracy of the ν Truncation* **a)** Calculation of the time dependent magnetic energy for the low-energy state $|5\rangle$, defined in Eq. (3.54), as the maximum ν value is increased for $g = 0.5$ and $N_\omega = 10$. This choice of digitization resolution ensures that any errors due to the ω_1, ω_2 truncation will be at the percent level. **b)** The probability $\langle \Pi_{\nu>1} \rangle$ of leaking out of the $\nu = 0, 1$ sector as a function of time for $g = 0.5$, $N_\omega = 10$ and $\nu_{\max} = 5$.

3.4 Circuits for Time Evolution

In this section the digitization presented above is translated into quantum circuits that can be implemented on current quantum devices. The circuit architecture for implementing time evolution is not unique, and two distinct approaches are developed that have different depths based on the level of digitization that is required for the simulation task.

In the first approach to circuit design, a single Trotter step is formed from component sub-circuits. This has the advantage of avoiding any pre-computation of circuits as the form of the components are already known. Additionally, most of these components have an asymptotic scaling that is no more than quadratic in

n_q . For the component implementing a diagonal unitary that encodes various transcendental functions, exact implementations scale exponentially in the number of elementary gates with n_q , the number of qubits used to digitize ω_1 and ω_2 . While this is seen as an unavoidable feature of this system, implementing such gates up to a desired level of precision scales more favorably than the worst-case exponential [287].

This approach gives the greatest benefit when high resolution in field space is required. When the number of qubits per plaquette is of order one, it is advantageous to employ a different approach. This begins by directly decomposing the Hamiltonian into a sum of Pauli strings and implementing the time evolution operator by a series of rotations. This approach scales exponentially in n_q , although dropping strings with weights below some threshold is found to give circuit-depth reductions that outweigh the losses in accuracy. The second approach is more suited to simulation tasks where only low resolution in field space is necessary. The following sections give explicit circuit decompositions for each of these approaches and studies the accuracy and different approximations that go into each. The relative resource counts are then compared.

3.4.1 Circuit from Differential Operators

This approach proceeds by organizing the terms in the two-plaquette Hamiltonian according to the highest order differential operator each term contains. In particular, the terms in the Hamiltonian are grouped in the following way

$$H_{\square\square} = H_{\partial^2} + H_{\partial} + H_{\omega} \quad (3.56)$$

where H_{∂^2} collects all the terms that have only a second derivative, H_{∂} contains any terms that have a first order derivative, and H_{ω} is all remaining terms, *i.e.* those that are only functions of ω with ν -dependent coefficients. The total time evolution operator is Trotterized according to these groupings, $e^{-itH_{\square\square}} \approx e^{-itH_{\partial^2}} e^{-itH_{\partial}} e^{-itH_{\omega}}$.

The circuits for $e^{-itH_{\partial^2}}$ and $e^{-itH_{\omega}}$ are straightforward, and mirror the circuits involved in simulating the scalar field. The Laplacian term is implemented by switching to the Fourier basis with a DST_{II} , which is based off of the Quantum Fourier Transform (QFT) [288]. In this basis each second derivative is diagonal and proportional to the square of the conjugate momentum operator $\hat{k}^2$,

$$\hat{k}^2 = \frac{\pi}{\omega_{\max}} \left(\frac{2^{n_q} - 1}{2} I - \sum_j^{n_q-1} 2^{j-1} Z_j \right)^2 \quad (3.57)$$




Figure 3.5: Circuits for the digitization presented in Sec. 3.4.1, which constructs sub-components for each class of differential operators **a)** Implementation of the exact exponential of the second derivative, making use of Discrete Sine Transforms. This subcircuit only entangles qubits within each ω register and does not touch the ν register. The ancillas a_1, a_2 associated with each ω register are the outermost qubits. **b)** An important circuit identity used in this work: The exponential of an N -body Pauli string P is implemented by the staircase of CX s connecting all the qubits participating in the gate, followed by an $R_Z(\theta)$ rotation on one of the qubits and then another CX staircase. The green single qubit gates denote a local change of bases to change P into a string of Z s on all qubits e.g. if there is an X on qubit i then this will be a Hadamard H and if it is a Y then it will be a Hadamard followed by a phase gate HS . The different shades denote that the local basis change can depend on the particular Pauli string P being implemented. **c)** Circuits for implementing first order derivatives and mixed derivatives using the finite difference representation. (*top*) Implements the exponential of a single first derivative. Each element in this circuit implements one of the terms in the sum Eq. (3.65). The bottom qubit of each step is the target qubit for both the V_k and CRZ gates. (*bottom*) An example of one of circuit elements $e^{-it\frac{\partial}{\partial\omega_1}\frac{\partial}{\partial\omega_2}}$ that interleaves the top circuit.




Figure 3.6: Trotter error for the circuits presented in Sec. 3.4.1, which constructs sub-components for each class of differential operators. Noiseless simulator results for the expectation value of the magnetic Hamiltonian at a different number of Trotter steps.

The circuit implementing a DST_{II} is shown in Fig 3.5a. While using the QFT as an algorithmic primitive, extra elements (not shown) are required to enforce the boundary conditions. The DST circuit requires an additional ancilla qubit changing the qubit count per gauge variable to be $n_q + 1$. H_ω is completely diagonal in the mixed basis and so it can be represented as sum of Pauli Z strings. It is implemented using the primitive shown in Fig. 3.5b with ν -dependent rotation angles.

The construction of the circuits for e^{-itH_∂} is more involved. Recalling the results from Sec. 3.2.2 and Sec. 3.2.2, the matrix representation of the Hamiltonian is organized into (ν, ν') sectors, where nontrivial sectors satisfy the selection rule $|\nu - \nu'| \leq 1$. The $\nu = \nu'$ case only involves the circuits already discussed. For $\nu \neq \nu'$, the Hamiltonian H_∂ is expressed as

$$H_\partial = \sum_{\nu} H_{\nu}^{+} |\nu\rangle \langle \nu + 1| + \sum_{\nu} H_{\nu}^{-} |\nu\rangle \langle \nu - 1| \quad (3.58)$$

where the expressions for $H_{\nu}^{\pm}$ are given in Appendix B.2.2.

The coupling of different ν sectors is reminiscent of a finite difference matrix in ν space. As already noted, implementing such an operator exactly requires a prohibitive number of gates. This is where the expressive

power of the mixed basis shines through. Truncating to $\nu_{\max} = 1$ gives the off-diagonal Hamiltonian the form

$$H_{\partial} = \sigma^+ H_0^+ + \sigma^- H_1^- \quad (3.59)$$

where $\sigma^{\pm}$, defined as

$$\sigma^{\pm} = \frac{1}{2}(X \pm iY), \quad (3.60)$$

acts on the single-qubit ν -register. Applying $HS^{\dagger}$ to the ν -register turns the projectors into

$$\sigma^{\pm} \mapsto \frac{1}{2}(Z \pm iX), \quad (3.61)$$

where H is a Hadamard and S is a $\pi/2$ phase gate. Thus under this transformation, H_{∂} maps to

$$H_{\partial} \mapsto \frac{1}{2}(Z \otimes H_{\text{sym}} - iX \otimes H_{\text{anti}}) \quad (3.62)$$

where $H_{\text{sym}}, H_{\text{anti}}$ are the symmetric and anti-symmetric parts of H_0^+ ,⁵

$$\begin{aligned} H_{\text{sym}} &= H_0^+ + H_1^- \\ H_{\text{anti}} &= H_0^+ - H_1^-. \end{aligned} \quad (3.63)$$

Explicitly, these components are given by

$$\begin{aligned} H_{\text{sym}} &= \frac{1}{\sqrt{3}} \frac{\partial}{\partial \omega_1} \frac{\partial}{\partial \omega_2} + \frac{1}{4\sqrt{3}} \cot \frac{\omega_1}{2} \cot \frac{\omega_2}{2} \\ H_{\text{anti}} &= -\frac{1}{2\sqrt{3}} \left(\cot \frac{\omega_1}{2} \frac{\partial}{\partial \omega_2} + \cot \frac{\omega_2}{2} \frac{\partial}{\partial \omega_1} \right). \end{aligned}$$

The complete time evolution induced by the first derivatives may be approximated by

$$\begin{aligned} e^{-itH_{\partial}} &\approx e^{-itg^2 Z \otimes H_{\text{sym}}} e^{-itg^2 X \otimes H_{\text{anti}}} \\ &\approx e^{-i \frac{g^2 t}{\sqrt{3}} Z \otimes \frac{\partial}{\partial \omega_1} \frac{\partial}{\partial \omega_2}} e^{-i \frac{g^2 t}{4\sqrt{3}} Z \otimes \cot \frac{\omega_1}{2} \cot \frac{\omega_2}{2}} \\ &\quad e^{-\frac{g^2 t}{2\sqrt{3}} X \otimes \cot \frac{\omega_1}{2} \frac{\partial}{\partial \omega_2}} e^{-\frac{g^2 t}{2\sqrt{3}} X \otimes \cot \frac{\omega_2}{2} \frac{\partial}{\partial \omega_1}} \end{aligned} \quad (3.64)$$

where the tensor product separates the ν register from the ω registers and the dependence of gauge coupling has been reinserted.

⁵ H_0^+ and H_1^- are Hermitian conjugates

Therefore, the problem of implementing the time evolution due to the H_{∂} terms is reduced to exponentiation of the finite difference operator. It has been recognized that this operator can be recast as a sum of Matrix Product Operators (MPO) [289] with an induced Trotter error of $t^2(n_q - 1)/2$. To see this, one starts from the identity

$$\frac{\partial}{\partial \omega} \mapsto \frac{-i}{2\delta\omega} \sum_k I^{\otimes n_q - k} \otimes V_k \left(-\frac{\pi}{2} \right) \left(Z_k \otimes |1\rangle \langle 1|^{\otimes k-1} \right) V_k^\dagger \left(-\frac{\pi}{2} \right). \quad (3.65)$$

where $\hat{Z}_k$ is the Pauli $\hat{Z}$ operator on the k^{th} qubit in the corresponding register and $V_k(\phi)$ is a unitary change of basis defined by

$$V_k(\phi) = \left(\prod_m^{k-1} \text{CX}_k^m \right) P_k(\phi) H_k \quad (3.66)$$

where CX_k^m is a CNOT gate with control on qubit m and target on qubit k , $P_k(\phi) = \text{diag}(1, e^{i\phi})$ is a phase gate on qubit k , and H_k is a Hadamard on qubit k , not to be confused with a Hamiltonian. After taking the exponential and Trotterizing the sum over k in Eq. (3.65), the operator $e^{-t\partial/\partial\omega}$ can be implemented by a sequence of controlled $\text{CRZ}_k^{1,\dots,k-1}(t/\delta\omega)$ rotations for each qubit in the register k , where $\text{CRZ}_k^{1,\dots,k-1}(\theta)$ is a controlled $R_Z(\theta)$ with target k and controls $1, \dots, k-1$.

Inspection of Eq. (3.64) shows that it is insufficient to implement terms just of the form $e^{-t\partial/\partial\omega}$, as there are entangling operations between the two gauge variable registers and the ν register. Therefore, it is necessary to modify the above circuits by interleaving the extra operations between each set of V_k and $V_k^\dagger$. As an example, the first term will need to implement an operator which is the product of a first derivative on each ω . Expanding the MPO expression for the first derivative gives

$$\frac{\partial}{\partial \omega_1} \frac{\partial}{\partial \omega_2} \mapsto -\frac{1}{4\delta\omega^2} \sum_{k,k'} I^{\otimes n_q - k} I'^{\otimes n_q - k'} \otimes V_k V_k' \left(Z_k Z_k' \otimes |1\rangle \langle 1|^{\otimes k-1} |1'\rangle \langle 1'|^{\otimes k'-1} \right) V_k^\dagger V_k'^\dagger$$

where the (un)primed operators denote operators acting on the $(\omega_1)\omega_2$ register and the argument of the V_k, V_k' have been suppressed. A similar interleaving can be made for the remaining terms in Eq. (3.64), and for concreteness, the circuit structures for the case just discussed are shown in Fig. 3.5c.

Fig. 3.6 shows different numbers of first-order Trotter steps of these circuits at various times t . If the noise profile of the device is partially or completely characterized, there exists an optimal step size for a target

total simulation time [50, 290]. For even a single first-order step, this method remains accurate at up to a step size of approximately $t = 0.5$, at which point dephasing is observed and deviations from the exact curve begin to grow. Modifying these circuits to perform a second-order step, with error that falls off as $\mathcal{O}(t^3)$ is straightforward, and requires modifications of the V_k circuits as outlined in [289].

3.4.2 Circuits by Pauli Truncation

Simulating time evolution requires converting the digitized time evolution operator $U = e^{-itH_{\square\square}}$ into a sequence of two-qubit gates directly implementable on a quantum device. The most straightforward way to do this, while not the most efficient, is to calculate $\mathcal{D}(H_{\square\square})$. Once this decomposition is performed, the sum in the exponent of U can be split into a product over Paulis using a Trotter-Suzuki formula, which approximates the exact time evolution operator. This time evolution operator $U_T(t) = \prod_P e^{-itc_P P}$, where $(P, c_P) \in \mathcal{D}(H_{\square\square})$, can be turned into circuit-level operations using the identity displayed in Fig. 3.5b. Generally this procedure will have a two-qubit gate count that scales as $\mathcal{O}(4^{2n_q+n_\nu})$. However, many of the rotations, particularly those corresponding to high weight strings, will have small angles that do not greatly impact the expectation value of observables over some range of evolution times. The rotations with small angles less than $\theta_{\min}$ can be dropped, which gives a compressed approximation of the exact time evolution operator.

It is necessary to understand the impact of dropping small angles for different levels of gauge field digitization and at different values of the gauge coupling. By removing these terms, the system is evolved instead by the Hamiltonian

$$H_\delta = H_{\square\square} + \theta_{\min} V$$

$$V = - \sum_{c_P \leq \theta_{\min}} \left(\frac{c_P}{\theta_{\min}} \right) P \quad (3.67)$$

where δ corresponds to the percentage of terms in the Pauli decomposition that have coefficient below $\theta_{\min}$.

The time evolution implemented on the device is given by the unitary

$$U_\delta(t) = U U_\theta(t) \quad (3.68)$$

where the U_θ (and its component) matrix, given by

$$U_\theta(t) = e^{-i\theta_{\min} \int_0^t d\tau V_I(\tau)}$$

$$V_I(\tau) = e^{-i\tau H_{\square\square}} V e^{i\tau H_{\square\square}}, \quad (3.69)$$

is the error term in the interaction picture. Therefore the extent to which the perturbation V grows under the dynamics of the two plaquette system directly controls the appropriate level of truncation. The error in evolution on some observable O is quantified by

$$\langle O(t) \rangle_\delta = \langle U_\theta^\dagger O(t) U_\theta \rangle. \quad (3.70)$$

which is the time-evolved value of the observable. This is compared to the untruncated value $\langle O(t) \rangle$ ($\delta = 0$).

To understand the effect of Pauli truncation error, the system is initialized in the state $|5\rangle$ defined in Eq. (3.54) and the expectation value of the magnetic part of the Hamiltonian H_B is calculated. Fig 3.7a shows this expectation value for various levels of truncation and gauge field digitization. For a coarse field digitization of $n_q = 2$, the expectation value is generally more sensitive to truncations. As the digitization resolution is increased, a larger percentage of the terms in the Pauli decomposition are able to be dropped, showing that an increasingly small fraction of the terms in the decomposition are responsible for capturing the relevant dynamics. For diagonal matrices, this can be understood via the Sequency Hierarchy Truncation [287]. The extension to non-diagonal matrices has not yet been developed, but similar arguments should hold.

Having truncated the Hamiltonian to retain only the Pauli strings that are most relevant for implementing time evolution, the time evolution operator U_δ must be converted into a quantum circuit. To minimize the number of CNOTs each Trotter step contains, the Pauli rotations are grouped to model the terms that appear in the Hamiltonian. This structure is presented in Fig. 3.7b. The first group includes strings that only act on individual ω registers, the second group involves strings that couple a single ω register to the ν register and have a Z acting on the ν register, and the third group includes strings that act on all of the qubits. The reason for classifying the rotations in this way is two-fold. First, as already mentioned, this grouping respects how the digitized Hamiltonian couples different quantum numbers, and so this structure establishes a hierarchy in the significance of how each term contributes to time evolution. The terms in group three generally have much smaller rotation angles and are the ones to be dropped first. Second, clustering terms that act on smaller subsets of qubits makes compilation easier, as only CNOT cancellations over a subset of qubits need to be considered. For compilation the open source software package **bqskit** is used [291]. The circuit that implements the i^{th} group is defined to be U_i , and so the first-order product formula is $U_T = U_3 U_2 U_1$. As group one is restricted to weight- n_q Paulis, it will generally involve fewer CNOTs to implement, and therefore




Figure 3.7: *Approximation for time evolution circuits presented in Sec. 3.4.2* **a)** The expectation value of the magnetic part of the Hamiltonian for differing digitization resolutions and percentage of terms dropped. Rapid convergence is observed for $n_q \geq 3$ and for truncation levels $\delta < .6$. **b)** Different stages of the time evolution operator for a second-order Trotter step. The silver circuit denotes the entire Trotter step. The blue sub-circuits act only on the ω registers, typically involving the largest angles, and is identical for each. The green sub-circuits couple each ω register to the ν register and the circuit structure is again identical for both ω . The yellow sub-circuit couples everything and generally contains the smallest rotations, therefore being the elements most affected by small-angle truncations. It also involves the highest weight terms, and so focusing on truncating these elements allows the circuit to be compressed at only a small cost in algorithmic error. **c)** The accuracy of first- and second-order Trotter on magnetic observables. Repeating the blue sub-circuit in the second step allows the circuit to take much larger time steps while only needing a small amount of extra gates to implement.

it is advantageous to use the second-order Trotter formula $U'_T = U_1^- U_3 U_2 U_1^+$, where $U_1^\pm$ involve the same set of operations as U_1 with the rotation angle c_{pt} replaced by $c_{pt}/2$ and $\pm$ denotes the same/reversed ordering of operations as compared to U_1 . The explicit rotations for each term U_1, U_2, U_3 are recorded in Appendix B.1.

Fig. 3.7c shows the evolution of the magnetic energy using both first- and second-order formulas for a variable number of Trotter steps. The second-order method allows for longer steps, giving similar levels of accuracy as three steps of first-order Trotter but a notably smaller number of CNOTs. Also shown for comparison is the exact evolution at $\delta = 0.66$ levels of truncation. Even with over half of the Pauli strings dropped, for long enough times, the error will be subleading to algorithmic error due to Trotterization.

3.4.3 Resource Estimates

To count the gates needed to implement a Trotter step using the differential operator method, the three stages are considered individually. For the Laplacian term, the quantum DST circuits can be implemented using $\mathcal{O}(n_q^2)$ elementary quantum gates, where the quadratic scaling is a direct result of the use of the Quantum Fourier Transform. The diagonal momentum operator requires an additional n_q^2 gates, so this entire stage requires only $\mathcal{O}(n_q^2)$.

Implementing the time evolution $e^{-t\partial/\partial\omega}$ requires at most $9n_q^2 - 33n_q + 34$ CNOT gates for $n_q \geq 3$ [289]. Using the interleaving method previously discussed, the two qubit gate count is the product of the gate counts for the two types of terms that are interleaved, e.g. $e^{-it\frac{\partial}{\partial\omega_1}\frac{\partial}{\partial\omega_2}}$ will require $(9n_q^2 - 33n_q + 34)^2 \sim \mathcal{O}(n_q^4)$. The decomposition of the other terms involving functions of ω_1, ω_2 in Eq. (3.64) will asymptotically require $\mathcal{O}(2^{2n_q+1})$ Z strings that must be classically pre-computed. This situation mirrors the counts estimated for compact $U(1)$ gauge theories [197]. In this case it was found that using the Walsh basis to encode functions of the gauge variables and truncating interactions with small angles was sufficient to break the exponential scaling to polynomial, suggesting an efficient path forward to completely reducing the scaling to polynomial in n_q . In absence of these more specialized methods the gate count for a single Trotter step using the differential operator method will scale as $\mathcal{O}(2^{2n_q+1})$. It is once more emphasized that the only place where the exponential scaling enters is in exactly implementing the diagonal functions of ω_1, ω_2 .

Turning now to the gate counts for the Pauli truncation method, even though a large percentage of the gates can be dropped to calculate time-evolved quantities up to some specified level of error, the scaling will generally be exponential in the number of qubits per plaquette. However, for large-scale simulations it may be more efficient to utilize the qubit budget for increasing the lattice volumes as opposed to increasing the resolution in field space. For instance, on a two-dimensional $(L+1) \times (L+1)$ lattice, the dominant interaction is between pairs of plaquettes. The number of such interactions is $\binom{L^2}{2}$, which is quartic in the

linear dimension L . This suggests that the only exponential scaling is due to resolving the Hilbert space at each physical link. As already seen in the previous section, ground state properties can be resolved with precision of 10^{-3} already at $n_q = 3$. Understanding the resource requirements for the Pauli truncation circuits up to per-mille precision is done by considering the number of two-qubit gates needed for time-evolution.

Analytically decomposing $H_{\square\square}$ into Pauli strings for arbitrary n_q is challenging, and so an upper-bound on the number of two-qubit gates is estimated by observing that every Pauli string of weight $w \geq 2$ in the Hamiltonian will require $2(w - 1)$ CNOTs, as shown in Fig. 3.5b. The estimate

$$\#(\text{Two Qubit}) \leq \sum_{P \in \mathcal{D}(H_{\square\square})} 2(w(P) - 1) \quad (3.71)$$

is numerically calculated for digitization resolutions of $H_{\square\square}$ up to $n_q = 6$. Fig. 3.8 shows the upper bound as a function of the percentage of truncated strings from the Hamiltonian. This upper bound still scales exponentially in the number of qubits per plaquette, although this scaling can be improved by a factor of $2 - 3$ with appropriate truncation. In addition, careful compilation of the circuits can bring this number down by another factor of ~ 2 . However, compilation is also generally a hard problem, so to compare the raw estimate to the number of post-compilation gates, the number of gates used on the quantum device in this work is also shown.




Figure 3.8: *Upper bounds on resource estimates for the Pauli Truncation circuits* The cumulative weight of the Pauli decomposition of $H_{\square\square}$ is shown as a function of the truncation percentage δ . The color of the curves denotes the magnitude of the coefficient in the decomposition. The blue point demonstrates the truncation level and actual gate depth used in the quantum simulation in this work, highlighting the effect of circuit compilation.

3.5 Real-Time Dynamics in Two Plaquettes

This section presents the results of a real-time simulation on the 156 qubit **ibm_fez** superconducting quantum chip. The circuit architecture used is the Pauli truncation approach presented in Sec. 3.4.2 at a truncation level of $\delta = 0.66$ on five qubits ($n_q = 2, n_\nu = 1$). Preparing the example low-energy state $|5\rangle$ in Eq. (3.54) requires additional circuits for state preparation. To isolate the performance of the time evolution circuits, the system is initialized in the product state $|00000\rangle$. Note that all of the states in the Hilbert space that $H_{\square\square}$ acts upon satisfy Gauss' law, and so this product state is still a valid physical state. The system is then time

evolved at gauge coupling $g = 0.5$ and the expectation values of the magnetic part of the Hamiltonian $\langle H_B \rangle$ are measured. For the gauge coupling considered here, $\mathcal{D}(H_B)$ has the following explicit decomposition into Pauli observables:

$$\begin{aligned} H_B = & 0.386 I \otimes I \otimes I \otimes I \otimes I - 0.0701 I \otimes I \otimes I \otimes I \otimes Z - 0.1445 I \otimes I \otimes I \otimes Z \otimes I \\ & + 0.0326 I \otimes I \otimes I \otimes Z \otimes Z - 0.0701 I \otimes I \otimes Z \otimes I \otimes I - 0.1448 I \otimes Z \otimes I \otimes I \otimes I \\ & + 0.0325 I \otimes Z \otimes Z \otimes I \otimes I. \end{aligned} \quad (3.72)$$

Fig. 3.9a shows the evolution of this observable from time $t = 0.1$ to $t = 0.6$, in increments of 0.1. To make maximum use of the device, up to 20 copies of the circuit were run in parallel at different locations on the chip, with 100 active qubits for every job. To limit the effect of cross-talk error between different parallel circuits, the layout routing was chosen such that at least one inactive qubit separated any active qubits in different subcircuits. An example of the layout used to generate the $t = 0.6$ data point is shown in Fig. 3.9b.




Figure 3.9: *Quantum simulation of $\langle H_B \rangle$ dynamics* **a)** The time evolution due to $H_{\square}$ at gauge coupling $g = 0.5$. The dashed line corresponds to the solution obtained by directly taking the matrix exponential of the digitized Hamiltonian and the solid line is the Trotterized evolution. The points are the mitigated results from **ibm.fez**. **b)** The physical layout of the Heron chip. The highlighted regions are an example of the parallel layouts used to generate the data point at $t = 0.6$.

Due to the inherent noisiness of current machines, utilizing various layers of error mitigation was crucial

in obtaining these results. Once the bare circuits are created and layouts chosen, Dynamical Decoupling (DD) [39, 292] was added to suppress idle errors and further mitigate the effects of cross-talk between qubits. In addition, Pauli Twirling (PT) [40] and measurement twirling [293] are applied to each circuit. In many cases, these methods allow the dominant noise channels of the device to be approximated as a depolarizing channel. The benefit is that the impact of a depolarizing channel on observable estimates is effectively dealt with using the final mitigation technique, Operator Decoherence Renormalization (ODR) [199, 213, 225, 294].

ODR is a very powerful error mitigation technique that allows the raw data from the device to be ‘renormalized’ in order to extract accurate physical results. For this method, two types of circuits are run on the quantum device. The first is the circuit for the physical process of interest, referred to as the physics circuit. The second, called the mitigation circuit, is a circuit that has the same pattern of single- and two-qubit gates, but is classically efficient to simulate. Classically simulating the mitigation circuit provides a noise-free and therefore numerically exact result to which the quantum simulation can be compared. Assuming the dominant error channel can be approximated by sufficiently weak depolarizing noise, the quantum estimate of the observable O will deviate from the true value by a multiplicative factor p_O , *i.e.* $\langle O \rangle_{\text{meas}} = p_O \langle O \rangle_{\text{true}}$. If the physics circuit is affected by the noise in the same way as the mitigation circuit, the factors of p_O between both are the same. Therefore, the renormalized estimate of the physical observable is given by

$$\langle O_{\text{phys}} \rangle_{\text{true}} = \langle O_{\text{phys}} \rangle_{\text{meas}} \frac{\langle O_{\text{mit}} \rangle_{\text{true}}}{\langle O_{\text{mit}} \rangle_{\text{meas}}} \quad (3.73)$$

where the subscripts *meas/true* denote the estimate using the raw/renormalized quantum data and the subscripts *phys/mit* denote the physics/mitigation circuits.

For this technique to work, the choice of mitigation circuit is of central importance. For the smaller time steps, the physics circuits were ‘Cliffordized’, meaning that the rotation angles of all R_Z gates were rounded to the nearest $\pi/2$. This transforms the mitigation circuit into a Clifford circuit with the same gate instructions as the physics circuit. When the time step grows large enough, the Cliffordization procedure can return a mitigation circuit that has vanishing expectation values, causing Eq. (3.73) to become inapplicable. In this case the Trotter structure of the circuit is utilized for designing the mitigation circuit. For all but $t = 0.1$, two steps of the second order formula were used, each with time step $t/2$, to reach time t . The gate sequence of this circuit can be rearranged so that when the second step has time step $-t/2$, the first step will be undone and the resultant state will be the initial state. As such all expectation values of Pauli Z strings will be 1 and so ODR can be used at minimal addition to circuit depth.

To generate the points in Fig. 3.9a, between 20 and 60 Pauli twirls are used with eight measurement twirls per Pauli twirl and 100 shots per circuit. The measurement twirls are averaged together and then ODR is applied to each Pauli twirl. The resultant renormalized Pauli twirls are then bootstrap resampled to get the mean and standard error for each run. These values are reported in Table 3.1, which also includes the depth of the circuit for each time step as well as the exact value of the Trotterized evolution. The exact value is calculated by Exact Diagonalization.

Time t	depth	ibm_fez	Trotter
0.1	25	0.0637 (16)	0.0683
0.2	48	0.2013 (41)	0.2015
0.3	48	0.3834 (41)	0.3846
0.4	51	0.5509 (54)	0.5510
0.5	51	0.6269 (82)	0.6137
0.6	51	0.549 (11)	0.5471

Table 3.1: Estimated and exact values of $\langle H_B \rangle$ as a function of time.

3.6 Discussions and Conclusion

Gauge theories are ubiquitous in descriptions of nature, appearing in the Standard Model, many condensed matter systems and theories of gravity. *Ab initio* calculations of their properties are vital for comparing theoretical predictions to experimental results. This is particularly true for QCD, whose strongly-coupled nature at low energies gives rise to a complex array of emergent phenomena as well as important contributions to high-energy processes. While massive theoretical and algorithmic developments in classical computing have allowed us to probe many aspects of gauge theories, there remain a plethora of open questions that do not seem amenable to these methods. For example, the determination of the phase diagram of finite-density QCD

and direct, nonperturbative calculations of real-time dynamics of hadrons appear inaccessible via classical techniques due to sign and signal-to-noise problems. Even our knowledge of whether the electroweak gauge interactions in the Standard Model can be given a precise, nonperturbatively regularized definition using lattice field theory faces similar difficulties, as sign problems obstruct the verification of candidate regularizations of chiral gauge theories.

Quantum simulations of Hamiltonian formulations of lattice gauge theories provide a powerful alternative framework for exploring strongly correlated quantum systems. There is an increased need for theoretical tools that make target systems amenable for such quantum simulations. For lattice gauge theories, a necessary step is constructing a finite-dimensional representation of the Hamiltonian, so that the continuous gauge groups of the Standard Model can be encoded into a discrete register of qubits. Ideally, this encoding can be efficiently simulated at all values of the (bare) gauge coupling and is also systematically improvable. In the last decade, many Hamiltonians for both Abelian and non-Abelian gauge theories have been constructed. The proposal put forth in Refs. [216, 218], combining gauge fixing and a basis inspired by the axis-angle representation of $SU(2)$, is one such promising approach for simulating real-time dynamics that have hitherto been inaccessible on classical devices.

The focus of this work was constructing the full algorithmic pipeline for simulating the real-time dynamics of the fully gauge-fixed two plaquette system with open boundary conditions, working in the mixed basis. While the explicit focus was on the two plaquette system, many of the developments are applicable to larger systems in two- and three-dimensions. The workflow started from the Hamiltonian written in the group element basis and ended with the measurement of an observable on a quantum processor. Throughout the development of this workflow, several key points were addressed.

First, this work validated that previously introduced mixed-basis truncation schemes work well for larger system sizes; an avenue for additional study is whether this scheme, particularly the optimal value of $\omega_{\max}$ given in Eq. (3.53), will continue to provide small truncation errors as system size is increased further. An important development in this work was demonstrating how the inclusion of first-order derivatives, only present in systems with two or more plaquettes, inhibits the exponential convergence predicted by the Nyquist-Shannon theorem. In particular, approximating first derivatives with finite-differences introduces digitization errors that scale polynomially with the truncation scale. This work demonstrated that with a qubit allowance of three qubits per plaquette, the low-energy properties of the two-plaquette system could be recovered with a precision of 10^{-3} . However, it is not clear how this precision depends on the lattice

volume and therefore an important question that must be addressed in future work is whether the inclusion of first-derivative terms dramatically limits the precision achievable in this formulation. Such first derivative terms appear in many other quantum many-body systems beyond $SU(2)$, and developing formulations of such terms compatible with quantum simulation is an area where more developments are needed.

Second, this work presented two distinct algorithms for implementing time-evolution on a digital quantum computer. The two algorithms had different quantum resource scalings and thus had different regimes of applicability. The first, described in Sec. 3.4.1, sorted the Hamiltonian according to the highest order differential operator that each term contained. A sub-circuit was constructed for each type of term. The benefit of this algorithm is that it scales more favorably in circuit depth for a large number of qubits and gives a Trotterization scheme with less error. Additionally, it can be easily adapted for systems with a larger number of plaquettes. However, it could not be implemented on currently-available quantum processors due to its depth. The second algorithm, presented in Sec. 3.4.2, decomposed the Hamiltonian directly into Pauli strings, truncating the Pauli strings that had a coefficient below some *a posteriori* determined value. This resulted in a quantum circuit that was sufficiently small as to be simulated on **ibm_fez**.

Lastly, this work determined an error mitigation strategy for simulating the real-time dynamics of the two-plaquette system on NISQ-era hardware. Previously-developed methods, namely Dynamical Decoupling, Pauli Twirling and Operator Decoherence Renormalization, were all necessary to achieve percent-level agreement between the quantum and classical results. It is not yet clear how the necessary quantum resource will scale with lattice volume and gauge coupling. Furthermore, different low-energy observables may require different error mitigation strategies.

The fully gauge-fixed two-plaquette $SU(2)$ Hamiltonian is a very fruitful system for understanding some of the hurdles that lie ahead on the path to simulating QCD on quantum processors. It is a relatively simple system, exactly solvable using FEM, but contains both non-Abelian gauge dynamics and non-trivial interacting degrees of freedom. This work charts a course towards simulations of larger two- and three-dimensional systems while also demonstrating the viability of the mixed-basis formulation for studying the properties of $SU(2)$ gauge theories at all values of the gauge coupling.

Chapter 4

MEASURING NON-STABILIZERNESS IN AN $SU(2)$ LATTICE GAUGE THEORY

4.1 Introduction

Applications of quantum computers to simulating particle and nuclear physics have attracted considerable interest in recent years, enabling *ab initio* studies of key phenomena such as particle collisions, confinement, and thermalization [21–23, 59, 81, 144, 148, 154, 184, 199, 225, 230, 238, 243, 245, 247, 253, 295, 296]. Gauge theories play a central role in the Standard Model and so their lattice regularizations, known as Lattice Gauge Theories (LGT), have been a primary area of focus [62, 64–66, 149, 151, 157, 158, 161, 176, 181, 182, 192, 194, 207, 211, 220–222, 226, 227, 246, 266, 268, 281, 282, 297, 298]. These are paradigmatic models of strongly-coupled quantum many-body systems and describe a wide array of physical systems, making them ideal targets for future quantum simulators.

Much of the difficulty in understanding the properties of strongly-coupled systems stems from their underlying quantum complexity properties. Initially, quantum complexity was understood through the lens of entanglement, a mechanism that allows for correlations in systems beyond what is possible classically [147, 299–302]. However, it has been realized that entanglement is not the only property responsible for the observed complexity in quantum systems, with non-stabilizerlessness [303–308], or “magic”, being another essential resource. Non-stabilizerlessness quantifies the amount of non-Clifford resources needed to prepare a particular state. While first viewed as a quantifier of computational complexity, the time evolution of non-stabilizerlessness has been linked to physical phenomena like quantum chaos and thermalization [309–311]. More recently, its relevance to high-energy physics has been explored in contexts like particle interactions [312, 313], the AdS-CFT correspondence [314] and string-breaking [315, 316]. However, systematic studies of the role of non-stabilizerlessness in the dynamical properties of non-abelian LGTs remain relatively unexplored.

One additional subtlety when studying non-stabilizerlessness in the context of LGTs is that the Hilbert space of the gauge field is formally infinite dimensional. Non-stabilizerlessness is naturally defined on qudit degrees of freedom, which have finite local dimension. Therefore to both simulate such theories and make contact with

standard measures of non-stabilizerness, it is necessary to “digitize” the Hilbert space. This renders the gauge field finite so that it may be mapped onto a register of qudits. There is no unique way to perform this mapping, and a variety of different bases have been developed [204, 205, 209, 214, 216, 261, 271, 273, 280, 317]. Among these bases, the *mixed basis* has emerged as one that can faithfully represent the gauge field across a wide array of different gauge couplings by combining both continuous magnetic variables and discrete electric variables to describe the wavefunction.

This work studies the dynamics of non-stabilizerness in non-abelian LGTs by considering pure gauge $SU(2)$ in the mixed basis. The smallest non-trivial system in this formulation is that of two interacting plaquettes, whose low-energy spectrum is described by smooth functions of two magnetic variables and a single electric variable. Central to this work is the development of a novel approach for mitigating errors in bitstring distributions, called Bit String Decoherence Renormalization (BSDR). This approach takes advantage of the fact that the amplitudes of such smooth wavefunctions admit a hierarchical expansion in terms of operators whose errors can be efficiently mitigated. Similar hierarchies have been utilized in the context of state preparation [287, 318, 319] and error correction [320] for a variety of different quantum applications, like options pricing, machine learning, and solving differential equations [321, 322]. This mitigation method, in conjunction with randomized measurement protocols, can additionally allow for the extraction of a number of non-trivial state properties beyond just the non-stabilizerness, like entanglement and the Inverse Participation Ratio [323].

This work is organized as follows. Section 4.2 introduces the relevant measure of non-stabilizerness and numerically examines its properties in both low-energy eigenstates as well as the dynamical setting of a quantum quench. Section 4.3 first presents the measurement protocol used to extract non-stabilizerness from the quantum device and introduces an improved estimator that significantly increases the efficiency of the calculation. Then, the novel mitigation technique is developed and its range of applicability is reported. Section 4.4 reports the non-stabilizerness measurements on an IBM quantum computer, **ibm_marrakesh**. Finally, Section 4.5 offers concluding remarks and outlines directions for future work.




Figure 4.1: *Digital quantum simulation of Two Plaquettes* **a)** The two-plaquette system with open boundary conditions is gauge-fixed using max tree gauge fixing. The system contains seven links, numbered 1 through 7, with the physical links denoted as $\kappa = 1, 2$. The magnetic degrees of freedom $\hat{X}_1, \hat{X}_2$ are associated with each plaquette. The dashed lines that run around each plaquette denote the loops used in the max tree gauge-fixing procedure. **b)** The digitized Hamiltonian is compiled into digital quantum gates via Trotterization that approximate e^{-iHt} . Shown below the circuit diagram is an example of initial state amplitudes given by a smooth gaussian profile in ω_1, ω_2 and the final amplitudes after time evolution by the circuit. To measure the non-stabilizerness, many copies of the time evolved state are prepared and then a set of single qubit Clifford gates are sampled from the ensemble $Cl(2)^{\otimes n}$ and appended to the end of each time-evolved state. These copies are subsequently measured and the noisy bitstring distributions are recorded. **c)** From the set of random measurements, error mitigation can be applied to the expectation values of different Pauli Z strings $\langle Z_k \rangle$ to estimate the noiseless values. These mitigated expectation values can be used to reconstruct the error-mitigated bitstring distribution $p(s)$, which can then be used to calculate non-linear functions of the state, like the Stabilizer Rényi Entropy or the purity.

4.2 Non-Stabilizerness of Two Plaquettes

Having mapped the continuous Hilbert space of the gauge field into a finite-dimensional quantum system, its quantum complexity can be quantified by the non-stabilizerness. On a system of n qubits the Pauli group¹ is defined as $\mathcal{P} = \{\otimes_i^n \sigma_i \mid \sigma_i \in I, X, Y, Z\}$. The n qubit Clifford group $Cl(n)$ is defined as the group of unitaries that normalize $\mathcal{P}$, and the stabilizer states are the set of states that are fixed by the action of $Cl(n)$. Computations involving only gates from $Cl(n)$ acting on stabilizer states can be performed efficiently classically [324]. The non-stabilizerness quantifies the distance of a state from the set of stabilizer states, and as such captures one notion of the quantum complexity of a state.

A measure of non-stabilizerness [305] that is computable for larger systems of qubits is the Stabilizer Rényi Entropy (SRE). A pure state $\rho = |\psi\rangle\langle\psi|$ is decomposed as

$$\rho = \frac{1}{d} \sum_{P \in \mathcal{P}} c_P P, \quad c_P = \text{tr} \rho P \quad (4.1)$$

where $d = 2^n$. The purity of the state enforces the condition $\sum_P |c_P|^2/d = 1$, implying that the set of squared coefficients $|c_P|^2/d$ form a probability distribution. The entropies of this distribution are the α -Stabilizer Rényi Entropies and are defined by

$$\mathcal{M}_\alpha = \frac{1}{1-\alpha} \log \frac{1}{d} \sum_P |c_P|^{2\alpha} \quad (4.2)$$

This quantity vanishes for stabilizer states and is non-increasing under application of non-Clifford unitaries for $\alpha \geq 2$, making it a good measure of non-stabilizerness from a resource theory perspective [325]. Consequently, this work uses the specific case of $\mathcal{M}_2$. A system can be said to have a large amount of non-stabilizerness if, for example, $\mathcal{M}_2$ grows extensively with system size. For mixed states the definition of the SRE must be modified since the set of Pauli expectations no longer forms a distribution. In this case, these expectations must be normalized by the purity, which is expressed as

$$S_2 = \frac{1}{d} \sum_P |c_P|^2 \quad (4.3)$$

and allows the SRE to be defined for mixed states. This definition only reliably detects stabilizer states when the purification of the state is itself a stabilizer [326].

¹To form a group, $\mathcal{P}$ must also include the possible phases $\pm 1, \pm i$. It is standard convention in the literature to refer to the set of all n qubit Pauli operators without phases as the Pauli group.

4.2.1 Non-Stabilizerness of Low-Energy Eigenstates

The low-energy spectrum of the digitized mixed-basis is comprised of smooth functions encoded into wavefunction amplitudes. Understanding the complexity of preparing such states will inevitably involve describing the stabilizer properties of such wavefunctions. Previous studies of ground state non-stabilizerness in the electric basis [327] show markedly different behavior than in the mixed basis, emphasizing the importance of basis choice when studying the theory at different values of the gauge coupling. The electric basis diagonalizes the Hamiltonian in the strong coupling limit, and so the eigenstates in this limit are trivially stabilizers. The mixed basis serves as an intermediate choice between the electric and magnetic representations and the eigenstates in the mixed basis exhibit non-vanishing non-stabilizerness at all values of the gauge coupling.

Figure 4.2a) (top) shows $\mathcal{M}_2$ for the ground state as a function of the gauge coupling. For large values of the coupling, the non-stabilizerness approaches the value calculated using the strong-coupling ground state wavefunction. In this limit, the wavefunction has the functional form $u_0(\omega_1, \omega_2) = \sin \frac{\omega_1}{2} \sin \frac{\omega_2}{2}$ and the non-stabilizerness is $\mathcal{M}_2 \sim 1.4$. As the gauge coupling is decreased, the support of the wavefunction, *i.e.* where it is not exponentially small, starts to decrease and becomes concentrated around ω_1, ω_2 close to zero. At intermediate values of the gauge coupling, $g \sim \mathcal{O}(1)$, the non-stabilizerness oscillates log-periodically, *i.e.* with minima that occur regularly with respect to $\log g$. Once g reaches the value such that $\omega_{\max}$, as given by Eq. (3.53), is less than 2π , the range of ω decreases with the support of the wavefunction and qualitatively different behavior is observed.

To understand the origin of these oscillations, it is instructive to consider a simpler wavefunction that is constant for $\omega_1, \omega_2 < C$ for $0 \leq C \leq 2\pi$ and vanishing outside this region. When such a state is encoded into a system of $2n$ qubits, the cutoff C will correspond to an integer between 0 and $2^n - 1$, as in Eq. (3.44). As shown in Appendix C.1, the non-stabilizerness of such constant wavefunctions will have minima when the integer representation of C coincides with a power of two. When the wavefunction varies over its range of support, as is the case with the ground state of the two-plaquette system, there will instead be a minimum near a power of two. As the support of the ground state wavefunction decreases with g , the observed log-periodic oscillatory behavior occurs.

The behavior of non-stabilizerness in the first and second excited eigenstates $|e_1\rangle$ and $|e_2\rangle$, shown in Fig. 4.2a) (bottom), closely mirrors that of the ground state, with saturation at large gauge coupling followed




Figure 4.2: *Non-Stabilizerness of low-energy eigenstates* **a)** Non-stabilizerness of ground state $|e_0\rangle$ (top) and first two excited eigenstates $|e_1\rangle$, $|e_2\rangle$ (bottom). For large gauge coupling, the non-stabilizerness plateaus to a constant value for $|e_1\rangle$, $|e_2\rangle$, which are degenerate in this limit. For a fixed resolution, once a certain coupling is reached, the $\omega_{\max}$ truncation begins to decrease from 2π , causing the effective support of the wavefunction over $\omega_1 - \omega_2$ to change at a different rate with g , which is signified by the transparent lines (these lines are not shown for the bottom plot). **b)** Non-stabilizerness in degenerate subspace $\mathcal{S}_{12}$. The states that have the minimum non-stabilizerness in this subspace are separable, while the states that have the maximum non-stabilizerness can be characterized by their properties under $\omega_1 - \omega_2$ exchange. **c)** Convergence of non-stabilizerness with increasing digitization resolution for first ten even eigenstates. **d)** Time evolution of $\mathcal{M}_2$ for different values of gauge coupling (top) and different levels of Pauli truncation ϵ in Hamiltonian (bottom). In the top plot, multiple digitization resolutions ranging from $n_q = 3$ to $n_q = 6$ are shown for each value of the coupling.

by subsequent oscillatory behavior as the gauge coupling decreases. The most striking feature is that the non-stabilizerness of both eigenstates approaches the same value in the strong coupling limit, owing to the degeneracy of $|e_1\rangle$ and $|e_2\rangle$ in the same limit. This degeneracy is related to both the separability of H_E and the $\omega_1 \leftrightarrow \omega_2$ swap symmetry implemented by the operator $\mathbb{T}_{12} = \frac{1}{d} \sum_{P \in \mathcal{P}_{n_q}} P_1 \otimes P_2$, where $d = 2^{n_q}$, $\mathcal{P}_{n_q}$ is the set of 4^{n_q} n_q -qubit Pauli strings (including the identity), and P_j denotes the string P acting on the ω_j

register. In the strong coupling limit, $|e_1\rangle$ and $|e_2\rangle$ form a two-dimensional subspace $\mathcal{S}_{12} = \text{span}(|e_1\rangle, |e_2\rangle)$ that is composed of states of the form $|\psi(\vartheta, \varphi)\rangle = \cos \frac{\vartheta}{2} |e_1\rangle + e^{i\varphi} \sin \frac{\vartheta}{2} |e_2\rangle$ and is parameterized by the two angles (ϑ, φ) . In this subspace, H_E is diagonal whereas H_B is not. Treating H_B as a perturbation for large but finite g using degenerate perturbation theory, it is seen that the basis states of $\mathcal{S}_{12}$ that diagonalize this perturbation have the same non-stabilizerness.

This equivalence may naively suggest that from the perspective of Hamiltonian simulation, the best basis to choose is the one that is aware of the structure of the perturbation, namely the one that splits the degeneracy. However, considering the range of values of the non-stabilizerness over the entire subspace $\mathcal{S}_{12}$ reveals a different picture. Figure 4.2b) shows the non-stabilizerness over $\mathcal{S}_{12}$ as a function of ϑ and φ . The states that maximize and minimize non-stabilizerness over $\mathcal{S}_{12}$ can be characterized by how they transform under $\mathbb{T}_{12}$. The basis is chosen such that the coordinate $(0, 0)$ corresponds to the state that is symmetric under ω exchange while the coordinate $(\pi, 0)$ corresponds to one that is anti-symmetric. The states that minimize the non-stabilizerness lie at the coordinates $(\pi/2, 0)$ and $(\pi/2, \pi)$, corresponding to separable states of the form $|\psi\rangle = |f(\omega_1)\rangle \otimes |g(\omega_2)\rangle$ with $f(\omega) = \sin \frac{\omega}{2}$ and $g(\omega) = \sin \omega$. The states with the highest non-stabilizerness are defined by the coordinates $(m\pi/4, n\pi/2)$ with $m, n \in \{1, 3\}$. These states can be defined by the property that exchange of ω_1, ω_2 is equivalent to complex conjugation of the amplitudes, *i.e.* $|\psi\rangle^* = \mathbb{T}_{12} |\psi\rangle$. This demonstrates that the most advantageous basis for quantum simulation may not always be the one that respects the structure of the Hamiltonian, and instead is the one that minimizes complexity.

Finally, the convergence of the non-stabilizerness with increasing digitization resolution is studied. Low digitization resolutions have the effect of coarse-graining small-scale variations in the wavefunction amplitudes, lowering the perceived non-stabilizerness. As the resolution is increased, these variations can have a non-trivial contribution to the resource requirements, causing the non-stabilizerness to grow extensively with the number of qubits. On the other hand, for smooth wavefunctions whose frequency spectrum decays quickly, increasing the resolution does not noticeably change the overall structure of the state and the non-stabilizerness will plateau below the expectation for typical, Haar random, quantum states. For the purposes of simulating these states in the continuum limit, this suggests that it is only necessary to simulate the state up to a minimum resolution, aligning with the expectation that the mixed basis requires a low number of quantum resources for different values of g .

For low-energy states $|e_i\rangle$, the convergence to the infinite resolution limit ($n_q \rightarrow \infty$) is exponential. To understand how quickly this convergence happens, the non-stabilizerness for the first ten eigenstates is

calculated for a fixed coupling of $g = 1$. The non-stabilizerness at infinite resolution $\mathcal{M}_2^\infty$ is extracted by fitting its values at finite n_q to the curve $\mathcal{M}_2^{n_q}(|e_i\rangle) = ae^{-bn_q} + \mathcal{M}_2^\infty$, with the finite n_q deviation being shown in Fig. 4.2c). While there is modest variance in the rates of convergence, all but one of the states studied reaches 10^{-2} precision with respect to the continuum value once $n_q = 6$. This stands in contrast to the required digitization levels for extracting values for local observables such as the energy density, for which it has been shown that as few as $n_q = 3$ qubits are required to get within per-mille precision of the continuum value [51]. This difference demonstrates that capturing the full complexity of a quantum state is not necessary for calculating the value of local observables in quantum simulations.

4.2.2 Time Evolution and Truncation Effects

The convergence of non-stabilizerness with digitization in the low-energy subspace becomes especially relevant when considering time evolution. This behavior is studied by considering the following quench-like scenario: The system is prepared in a superposition of low-lying eigenstates of the Hamiltonian at some fixed gauge coupling g_0 then evolved with the Hamiltonian at a different coupling g . The exact superposition is not important, only that the states in it are sufficiently far from the middle of the spectrum so that over the ranges of times considered the time-evolved state does not mix with energy bands that contain a number of states that grow extensively with system size. For concreteness, the initial state is an equal superposition of the first three eigenstates, $|\psi(0)\rangle = (|e_0\rangle + |e_1\rangle + |e_2\rangle)/\sqrt{3}$, which has energy on the order of the mass gap of the system.

Fig. 4.2d) (top) shows the time-evolved non-stabilizerness for $g_0 = 10$, very close to the strong-coupling limit, for $g = \{0.6, 1, 2\}$. For $g = 2$ the time evolution does not rapidly mix eigenstates of the g_0 Hamiltonian and therefore the non-stabilizerness exhibits small oscillations between the initial states. As g is decreased, the time evolution induces mixing between the initial state and other eigenstates in increasingly higher energy bands, causing the oscillations to grow and then subsequently diminish as the dynamics become dominated by ramp-like behavior. The same type of growth is observed when initializing the system in eigenstates corresponding to small g and then time evolving at values of larger g . This behavior can be understood by recognizing that the value of the gauge coupling changes the eigenbasis of the Hamiltonian and therefore time-evolving the eigenstate of one Hamiltonian with another Hamiltonian has the effect of introducing more components into the functional decomposition of the wavefunction. The number of components that

contribute non-trivially will depend on how closely the two eigenbases overlap. This suggests, that for wavefunctions representing smooth functions, a greater number of components will give a higher amount of non-stabilizerness.

When simulating real-time dynamics on a quantum device, it is necessary to also consider how approximations to the time evolution operator modify the complexity of the prepared state. The circuits in this work are constructed by decomposing the Hamiltonian into a sum of Pauli strings and dropping strings that have sufficiently small prefactors from the decomposition. When viewed from the perspective of Trotterization, this has the effect of removing many small R_z rotations (up to Cliffords) and so it is expected that as long as the prefactor is small enough the complexity of the state will not be affected. Figure 4.2d) (bottom) shows the time evolution of $|\psi(0)\rangle$, where ϵ percent of Pauli string with the smallest coefficients are removed from the Hamiltonian. Surprisingly, as many as $\epsilon \approx 0.6$ of the Paulis, corresponding to 60% of the total operator content of the Hamiltonian, can be removed with no visual change to the non-stabilizerness, allowing for the complexity dynamics to be studied on the quantum device for deeper circuit depths.

4.3 Walsh Based Error Mitigation

This section outlines the techniques used for measuring the non-stabilizerness on a quantum device. The distribution of measured bit strings is the primary observable used to calculate $\mathcal{M}_2$. Obtaining accurate estimates from quantum hardware necessitates the development of techniques that allow one to account for the effects of noise present on a device. Since the non-stabilizerness depends on the entire bitstring distribution, without *a priori* knowledge of this distribution, naive scaling arguments would imply a mitigation overhead that grows proportionally to the number of bitstrings, *i.e.* exponentially with system size. An approach tailored for mitigating the noise in bit string distributions, that in certain cases ameliorates the dependence on system size, is presented. It is shown that for wavefunctions that are the digitizations of smooth functions, e.g. eigenstates of the two-plaquette system, it is sufficient to mitigate a constant number of observables.

4.3.1 Random Measurements and Improved Estimator

Measuring the non-stabilizerness of a state on a quantum device is made possible through the use of randomized measurements [306]. The protocol begins by randomly sampling N_U single-qubit Clifford unitaries, forming an ensemble of circuits that are averaged over. Once the desired state is prepared, the

random Clifford rotations are applied to each qubit, measured N_M times and the outcome bitstrings are recorded. This setup is illustrated in Fig. 4.1b). The collection of N_M outcomes, per random Clifford, gives an estimate for a probability distribution in the computational basis $p(s)$, where the dependence on the Clifford unitary has been suppressed. For the measured distribution, the following estimators are computed

$$O_j^{(2)} = \sum_{s_1, s_2} 2^{-w(s_1 \oplus s_2)} p(s_1) p(s_2) \quad (4.4)$$

$$O_j^{(4)} = \sum_{s_1, s_2, s_3, s_4} 2^{-w(s_1 \oplus s_2 \oplus s_3 \oplus s_4)} p(s_1) p(s_2) p(s_3) p(s_4) \quad (4.5)$$

where $w(s)$ denotes the Hamming weight of the bitstring s and the subscript $j \in \{1, \dots, N_U\}$ indexes the circuit in the measurement ensemble. Averaging $O_j^{(2)}, O_j^{(4)}$ over the ensemble of random Cliffords provides experimental estimates of the Stabilizer Rényi Entropy $\mathcal{M}_2$ and the purity S_2 , which are given by

$$\begin{aligned} \mathcal{M}_2 &= -\log X^{(4)} + \log d X^{(2)} - \log d \\ S_2 &= d X^{(2)} \end{aligned} \quad (4.6)$$

in the limit $N_U \rightarrow \infty$, where $X^{(\ell)} = \frac{1}{N_U} \sum_{j=1}^{N_U} O_j^{(\ell)}$.

The expressions for the estimators $O_j^{(2)}$ and $O_j^{(4)}$ involve summing over 16^n terms and so the classical post-processing cost of the calculation can be prohibitive for large n . It is beneficial to introduce an improved estimator, which is simply a re-writing of Eq. (4.4). The derivation fundamentally relies on the Walsh-Hadamard (WH) transform of a Boolean function $f(s)$, denoted $\hat{f}(k) = \sum_s f(s) (-1)^{k \cdot s}$. For notational convenience the WH transform will be denoted as $c_k = \hat{f}(k)$ and referred to as the *Walsh coefficients*. While the use of the WH transform has recently been recognized to give advantages in similar contexts [328–330], the fact that the estimators $O_j^{(\ell)}$ are written in terms of bitstrings allows the equations to be written in closed form. Specifically for $\ell = 2, 4$, the estimators are given by

$$O_j^{(\ell)} = \frac{1}{d^2} \sum_k 3^{w(k)} \hat{p}_j^{(\ell)}(k); \quad (4.7)$$

details of this derivation are in Appendix C.2. The key observation is that while the original expression for $\tilde{\mathcal{M}}_2$ involved summing over 16^n different terms, the improved expression only sums over 2^n . Implementing the WH transform requires $\mathcal{O}(n2^n)$ operations, but it is also possible to put single-qubit Hadamards after

the sequence of random Cliffords and then make computational basis measurements. This reduces post-processing to strictly $\mathcal{O}(2^n)$ at the potential cost of extra hardware error. Single-qubit error rates are generally an order of magnitude below that of two-qubit gates and so this is not expected to be a large contribution to the error budget on current-day NISQ-devices. This reduction in complexity enables the estimators to be easily evaluated on a laptop for the same system sizes that are accessible via exact diagonalization ($n \sim 24$).

4.3.2 Error Mitigation Technique

Error mitigation refers to a host of techniques that aim to reduce the effects of device noise on measured quantities, often involving application of compensating factors in post-processing. Many powerful techniques have already been developed that correct noisy expectation values of operators [38,40,115,293]. However, the observables $O^{(2),(4)}$ are functions of $p(s)$ and require getting mitigated estimates of the bitstrings themselves. Different strategies for accomplishing this have been developed in different contexts, including matrix inversion [331], maximum likelihood estimation [332–334], and Monte-Carlo sampling [335]. This work introduces a novel approach to error mitigation called Bit String Decoherence Renormalization (BSDR), aimed at mitigating the bitstring distributions $p(s)$ relevant for calculating quantities like non-stabilizerness or Rényi entropies.

A quantum state $|\psi\rangle$ written in the computational basis is given by $|\psi\rangle = \sum_s e^{i\phi_s} \sqrt{p(s)} |s\rangle$. Similarly, projectors onto the computational basis states can be written as

$$|s\rangle\langle s| = \sum_k (-1)^{k \cdot s} Z_k \quad (4.8)$$

where Z_k denotes a string of Pauli Z operators and k is a length n bitstring. Here, Z_k means there is an identity on qubit j if the j^{th} entry of k is 0 and Z where the j^{th} entry is 1. The desired bitstring distribution can therefore be formulated as an expectation value

$$p(s) = \text{Tr} \rho |s\rangle\langle s| \quad (4.9)$$

with $\rho = |\psi\rangle\langle\psi|$. It then follows straightforwardly that the bitstring distribution $p(s)$ is just the WH transform of the Walsh coefficients c_k , which coincide with the set of expectation values for Z strings. This suggests that to mitigate the distribution $p(s)$, one should mitigate the Z expectation values and then reconstruct $p(s)$ from the set of mitigated observables. A secondary implication of the equivalence between $p(s)$ and c_k is that the estimators Eq. (4.7) for non-stabilizerness and purity can be expressed directly in terms of the Clifford

averaged Z expectation values. Using this approach, the noisy estimate for the bitstring probability is given by

$$\tilde{p}(s) = \sum_k (-1)^{k \cdot s} \tilde{c}_k \quad (4.10)$$

where the tilde denotes the mitigated estimate for both $p(s)$ and c_k . Defining the bias of the estimate to be $\text{Bias}(p(s)) = \tilde{p}(s) - p(s)$, then $\text{Bias}(\tilde{p}(s)) = \sum_k (-1)^{k \cdot s} \text{Bias}(c_k)$, which implies that when the mitigation method for the observables is unbiased, the estimator $\tilde{p}(s)$ is also unbiased.

To understand why inferring $p(s)$ from mitigated observables is advantageous, it is important to note that the noise affecting hardware can generally be modeled as Pauli noise

$$\mathcal{E}(\rho) = \sum_P a_P P \rho P \quad (4.11)$$

where the sum runs over all Pauli strings and $0 \leq a_P \leq 1 \ \forall P$. The effect on an expectation value is simply rescaling by a coefficient r , *i.e.* $\text{Tr} \mathcal{E}(\rho) O = r \text{Tr} \rho O$, with r defined as

$$r = \sum_{P \in O^+} a_P - \sum_{P \in O^-} a_P \quad (4.12)$$

where $O^{+/-}$ are the set of Pauli strings that commute/anti-commute with O . While Pauli noise acts in a predictable way for Pauli string observables, the effect on bitstrings is generally more complicated and not just a rescaling. Observable mitigation methods like Zero-Noise Extrapolation (ZNE) and Operator Decoherence Renormalization (ODR) [199, 213, 225, 294] take advantage of this structure, assuming that the noise strength is weak enough so that r can be canceled. While there is likely a regime of noise strength for which a simple relation between the bitstring probabilities and noise strength holds, the fact that $p(s)$ is a combination of many Pauli strings obfuscates the point at which this limit becomes appropriate. Therefore, applying the wealth of methods for mitigating observables O and using these mitigated values to infer the bitstring distribution is found to be much more effective.

4.3.3 Walsh Based Reconstruction

It is undesirable to require mitigating every c_k , as their number scales exponentially with the number of qubits. However, for states that encode smooth functions such as the low-energy states of the two-plaquette system, the Walsh coefficients admit a hierarchy that allows for efficient reconstruction. This idea has been used before in the context of compression for digital signals. To understand how this hierarchy appears, it is

necessary to define some notation. The *sequency* of a bitstring is the digital analog of frequency, defined as the number of sign changes in the row of the n -qubit Hadamard matrix $H_2 \otimes \cdots \otimes H_2$, with $H_2 = (X + Z)/\sqrt{2}$ the single-qubit Hadamard, indexed by that bitstring under natural (binary) ordering. The sequency ordering is then obtained by re-indexing rows according to the decimal value of their bit-reversed Gray code: for each bitstring k , one computes its Gray code representation, reverses the order of its bits, *i.e.* reflects the bitstring left-to-right, and ranks rows by the resulting decimal values in ascending order. The ν^{th} sequency band is defined by the collection of sequency ordered bitstrings that lie in the range $2^{\nu-1}$ to $2^\nu - 1$. For multivariate functions defined on a D -dimensional interval ($D = 2$ for the two-plaquette eigenstates), each band is characterized by a tuple of band coefficients. As shown explicitly in Appendix C.3, for sufficiently smooth functions the coefficients in each band can generally be bounded as

$$c_\nu \sim \prod_{i=1}^D \frac{1}{2^{\nu_i}} \quad (4.13)$$

which gives a prescription for how reconstruction should proceed.

Fig. 4.3a) describes the strategy for reconstruction, taking the ground state of the two-plaquette system at couplings $g = 0.6$ and $g = 2$ as concrete examples. The bands with greater influence on the reconstruction will have the lowest band indices (ν_1, ν_2) , and so starting from the origin the index of each band is incremented individually, alternating between ν_1 and ν_2 . The bottom figure explicitly shows the entire sequency spectra, with the bands containing the lowest, most dominant coefficients highlighted. It is actually unnecessary to calculate all of the coefficients within each band, as only including the most dominant coefficient generally gives the largest improvement. However, how these maximum coefficients are distributed within bands is non-trivial, and without further knowledge of the specific function, one must mitigate the entire band to account for its contribution. Fig. 4.3b) shows the reconstruction quality at various cutoffs compared to the exact wavefunction for fixed ω_2 . For low cutoffs only the shape of the distribution is recovered, although as the cutoff is increased the profile of the reconstructed distribution becomes indistinguishable from the exact curve.

To measure the quality of the reconstruction, denote the reconstructed distribution up to sequency number k^* as $p(s; k^*)$. The error between the exact distribution $p(s)$ and $p(s; k^*)$ is quantified by the Total Variational Distance (TVD), δ , defined by

$$\delta(p_1, p_2) = \frac{1}{2} \sum_s |p_1(s) - p_2(s)| \quad (4.14)$$




Figure 4.3: *Walsh hierarchy of smooth wavefunctions* **a)** Different Walsh bands for two-dimensional function of ω_1, ω_2 (top) with Walsh indices and associated coefficients (below). The lowest, most significant, bands are highlighted in blue, and the next most significant are shown in red. The strategy to reconstruct the wavefunction involves starting from the shells closest to the origin and working out in each direction. Once a shell is filled, the next shell starts being filled. **b)** Reconstructed wavefunction probability density for fixed ω_2 at different Walsh cutoffs k^* . The solid black lines correspond to exact reconstruction where the cutoff is removed. Both a strong coupling wavefunction at $g = 2$ and a weak coupling wavefunction at $g = 0.6$ are shown. **c)** Number of Walsh bands needed to reconstruct distribution $p(\omega_1, \omega_2)$ up to fixed error $\delta = .1/.04$ (top/bottom) as function of n_q . The number of bands in each case increases linearly, corresponding to an exponential growth in the number of required Walsh coefficients until a plateau is reached. The different points in the plot denote the values of g for which the ground state wavefunction was calculated.

and which satisfies the properties of a distance over distributions. Assuming the hierarchy in Walsh coefficients satisfies the condition Eq. (4.13), it is shown in Appendix C.3 that for a D -dimensional function the TVD is bounded as

$$\delta \leq \sqrt{\frac{DC^D}{6}} \Delta \sum_j 2^{-m_j} \quad (4.15)$$

where $m_j = \lfloor \log_2 k_j^* \rfloor$ denotes the most significant bit of the cutoff k_j^* in the j^{th} direction, $C \approx 1.356$ is a constant, and Δ denotes the maximum derivative of the function over its whole domain when considered on a continuous interval, *i.e.* an infinitely-resolved digitization. Crucially, assuming that the underlying distribution is sufficiently smooth, the bound becomes independent of n_q . Fig. 4.3c) shows the number of bands required to reach a constant TVD of $\delta = 0.04$ and $\delta = 0.1$ as a function of system size n , for a range of gauge couplings $g = \{0.6, 1, 2\}$. As the size is increased, the number of bands increases linearly (corresponding to an exponential increase in coefficients), until the critical threshold is hit, at which point the curves plateau, demonstrating an independence from system size.

4.4 Measurement of Non-Stabilizerness on a Quantum Device

Using the circuits approximating time evolution due to the Pauli-truncated Hamiltonian approach outlined in Sec. 4.2, a second-order Trotter step at varying time steps is simulated on IBM's Heron r2 device, **ibm_marrakesh**. The non-parallelizable CNOT depth is 51, which to the authors' knowledge is the deepest circuit where the non-stabilizerness has been measured on a quantum device. The number of Clifford circuits sampled ranged from $N_U = 100$ to $N_U = 200$. Running many circuits in parallel on a single chip was crucial for allowing economic use of the quantum resources. This required a carefully chosen layout, prioritizing high performing qubits. To choose the layout, the calibration parameters of the chip were queried before each run and the qubits were filtered based on having readout error less than 0.05 and having connections to other qubits with two-qubit gate error rate less than 0.01. Once the filtered set of qubits were chosen, contiguous patches of five qubits were found such that every patch was separated by at least one inactive qubit on the device. This approach minimized cross-talk between resonators.

The error mitigation methods used were Pauli Twirling (PT), measurement twirling, dynamical decoupling, and BSDR. For each individual circuit, the measurement twirls were averaged together, and then BSDR was applied to individual Pauli twirls, allowing for effective renormalization of $p(s)$. Each shot of the physics circuit was Pauli twirled 28 times and then measurement twirled 5 times, where the number of measurements

N_M for each Pauli twirl were split evenly among the measurement twirls. The random Cliffords were omitted on the mitigation circuits so that a single set of mitigation circuit runs could be applied to all the physics circuits, allowing most of the quantum resources used to be applied to the physics circuits.

Figure 4.4a) (top) shows an example bitstring distribution from one particular random Clifford with the raw results and the mitigated results (where all the twirls are averaged over). The primary effect of the noise is to transfer probability from the most significant strings and distribute it among the less likely strings. This redistribution of probability can have a severe impact on the predicted value of non-stabilizerness from the device, and so applying the mitigation technique developed in this work is crucial. By mitigating the $\langle Z_k \rangle$ and inferring the distribution $p(s)$ from the mitigated operators, the effects of the small but many erroneous bitstrings are suppressed and the most likely ones are amplified, restoring their significance.

The predictions of the non-stabilizerness are shown in Fig. 4.4b), and the numerical values are presented in Table 4.1. It is immediately apparent that without mitigation, the experimentally measured non-stabilizerness is higher than that predicted by classical simulation of the Trotter evolution, and the variation over time is suppressed. It is counterintuitive that the unmitigated non-stabilizerness is higher than that of the noiseless circuit in light of results on the complexity of noisy quantum circuits [336]. This effect can be understood by considering the Pauli spectrum of a state, *i.e.* the set of all expectation values of Pauli strings. For stabilizer states on n qubits, there will be 2^n Paulis that have expectation value plus or minus one, and the rest will be exactly zero. On the other hand, states with high non-stabilizerness generally have many nonzero Pauli expectation values with small magnitude. As previously observed, Pauli noise will reduce the magnitude of large expectation values and with finite measurement statistics, expectation values that should vanish exactly will be estimated to have a small but nonzero value. These effects taken together will cause the otherwise highly structured Pauli spectrum of a low magic state to resemble the more generic spectrum of a high magic state.

Time t	Cliffords N_U	raw	mitigated	Trotter
0.1	100	1.893 (33)	1.13 (23)	1.271
0.2	150	1.921 (37)	1.58 (15)	1.623
0.3	150	1.923 (37)	1.57 (15)	1.487
0.4	200	1.736 (28)	1.29 (18)	1.168
0.5	100	2.030 (38)	2.06 (14)	1.93
0.6	100	2.056 (42)	1.85 (21)	2

Table 4.1: Raw, mitigated and Trotter values of $\mathcal{M}_2$ as a function of time. The raw values are determined by using the bitstring distributions that are read directly off of the device without any post-processing.

4.5 Conclusion

Quantum simulation promises to answer questions about strongly-interacting many-body systems beyond the reach of current approaches. Non-abelian gauge theories are an appealing target, as their symmetries hinder analytic progress and large Hilbert spaces create computational barriers. This work improves the simulation of these theories in two ways: by investigating the quantum complexity of an $SU(2)$ LGT in the mixed basis and by developing methods for efficient measurement of quantum information quantities on a quantum device. First, the Stabilizer Rényi Entropy was calculated for a system of two interacting plaquettes. The non-stabilizerness of low-lying eigenstates was determined and its convergence with local Hilbert space dimension studied. Extrapolating from increasingly resolved digitization levels, $n_q = 6$ qubits per plaquette were sufficient to reach one percent of the SRE value for low-energy eigenstates, which is 2-3 times the number of qubits needed for comparable precision in local observables like the energy [51]. It was also shown that different combinations of degenerate eigenstates can alter the required resources for simulating low-energy physics. Second, randomized measurement schemes were improved by rewriting conventional




Figure 4.4: *Quantum measurement of non-stabilizerness* **a)** Bitstring distributions for specific instance of random Cliffords. The raw data (left) shows small populations of many bitstrings whose populations are spuriously amplified by the noise of the quantum device. The bitstring distribution after application of BSDR (right) removes these spurious populations and restores much of the probability mass to the true bitstring. **b)** Values of non-stabilizerness over time. The solid lines correspond to evolution due to the exact propagator and the Trotterized circuit. The circles represent the calculations from the raw data and the diamonds represent the mitigated data. The raw data is almost constant, and the features of the dynamics are highly suppressed. The mitigated values are able to accurately track the real time dynamics for all times studied.

estimators as a weighted sum over Clifford-averaged expectation values, reducing classical post-processing from $\mathcal{O}(16^n)$ to $\mathcal{O}(n2^n)$ via the Fast Walsh-Hadamard Transform. This work is the first application of this method to non-stabilizerness measurement protocols.

This connection between the non-stabilizerness estimator and operator expectation values motivated the central contribution of this work: a novel error mitigation scheme for bitstring distributions, called BSDR. The key insight is that common noise channels on current devices can be expressed as Pauli noise, which is more effectively corrected at the operator level. For states admitting a hierarchy in their Walsh spectrum, such as the low-energy subspace of field theories like the scalar field and $SU(2)$ in the mixed basis, as well as states arising in applications such as PDEs and machine learning, only a constant number of observables is needed to reach a fixed error level in the reconstructed distribution.

BSDR was applied to real data from randomized measurements on IBM quantum devices, where raw

results failed to capture time variations of non-stabilizerness. The mitigated results more accurately show the time dependency and the buildup of non-stabilizerness after a quench from a stabilizer state. This timescale is connected to quantum chaos and ergodicity [310, 311], making BSDR suitable for studying non-equilibrium quantum dynamics on real devices through the lens of bitstring distributions [337].

Extending this work to numerical studies of more plaquettes is an important future step. It has been recognized that the non-stabilizerness, in particular a basis-independent measure like the non-local magic, can serve as signals for the internal structure of bound particles. Understanding the implications of this for gauge fields is a promising direction [338–340]. Furthermore, systematic studies comparing the non-stabilizerness in other bases like the electric representation and connecting this to the practical gate complexity and non-Clifford resources required in simulation are essential for reaching larger lattice volumes.

Another vital direction is determining how the Walsh spectrum of a state is related to other measures of complexity like non-stabilizerness and non-Gaussianity. It has been shown that for certain quantum states, a hierarchy in the non-Gaussian correlation structure allows for efficient representations [54]; extending this concept to the realm of sequency truncation may allow for experimentally-friendly probes of otherwise complex observables. Additionally, understanding the underlying structure in the Walsh spectrum between states that are considered complex by different classes of resource theories may elucidate new phenomena in the context of quantum chaos and non-equilibrium dynamics.

Finally, it is expected that the BSDR error mitigation strategy presented in this work will find great use in many other simulation tasks relevant for high energy physics. As an example, many of the primitives that are used to construct wavepackets in scattering experiments, namely plane waves and Gaussian profiles, constitute smooth functions and so BSDR should prove to be beneficial [238, 243, 332]. Using BSDR will be advantageous for studying complex observables like the non-Stabilizerness, Rényi entropy, Inverse Participation Ratio, and other quantities that are strongly sensitive to device noise.

Chapter 5

CLASSICAL SHADOWS FOR SAMPLE-EFFICIENT MEASUREMENTS OF GAUGE-INVARIANT OBSERVABLES

5.1 Introduction

Programmable quantum simulators and computers enable the investigation of large many-body systems, for instance, opening new routes for studying systems out-of-equilibrium [73, 145–150, 230, 341] or discover emergent phases of matter [71, 165, 166] that elude classical algorithms. A key point is that one can not only simulate such complex quantum states, but also characterize them at scale. This includes both the extraction of observables to study the properties of these systems [37, 342], but also verification and benchmarking of the computational devices themselves [156, 343–349].

A broad range of methods have been developed to characterize quantum states, from tomography protocols that aim to reconstruct the full state [345, 350–356] to randomized measurement schemes [357–362], including classical shadows [37, 363–374], which bypass full tomography and often allow efficient estimation of many properties directly from measurement data. The latter are especially popular in experiments [361, 374–381], since many variants can be realized on current devices using only simple operations and computational-basis measurements.

Randomized measurement protocols are designed to be general and state-agnostic, following the philosophy of “measure first, ask questions later” [382]. This flexibility, however, can increase the learning cost unnecessarily. In practice, prior knowledge about the structure of unknown states or relevant observables can guide measurement pre-processing [383, 384]. Key examples include symmetries—such as particle number or permutation symmetry—arising from the target state, the observables, or hardware constraints on measurement operations [358, 365, 385–389]. A common feature is that exploiting prior knowledge reduces the measurement cost but can increase circuit complexity.

Systems with *local* symmetries, such as Lattice Gauge Theories (LGTs), exhibit distinctive Hilbert space structures and state properties. Quantum simulation of LGTs is a key application for quantum computers and simulators [151], owing to their central role in describing high-energy and nuclear physics [59, 82,




Figure 5.1: *Sample-efficient measurements of gauge-invariant observables in $\mathbb{Z}_2$ LGT.* Illustration of $\mathbb{Z}_2$ LGT in $(2 + 1)$ D and its Ising dual. For the LGT qubits reside on links l of the lattice and Gauss law constraints G_s are associated with sites s . The Hamiltonian consists of plaquette (magnetic) terms $W_{\square} = \prod_{l \in \square} \sigma_l^z$ and on-site (electric) terms σ_l^x . In the dual description as a $(2 + 1)$ D Ising model the qubit degrees of freedom are associated with the plaquette. There is a one-to-one correspondence between operators and states within the physical Hilbert space of the LGT, independent of boundary conditions: the magnetic terms map to Pauli- Z operators on the Ising qubits and the electric terms map to two-body Pauli- X terms acting on the qubits across the relevant link. In the inverse map from the Ising model to the LGT, single X operators on the Ising side of the duality map to a string of σ^x operators along some path γ to an arbitrary reference plaquette $\square_0$. For PBC, the dual Ising model obeys a global parity constraint and such single-body X operators are unphysical.

295, 296, 390], condensed matter systems [391–394], fermion-to-qubit mappings [395–397], and because of connections to quantum error correction [398–400]. We will show that incorporating prior knowledge when designing randomized measurement schemes for LGTs leads to especially pronounced trade-offs between sample complexity and circuit complexity. The common eigenspace of local symmetry constraints is exponentially smaller than the full Hilbert space used in the simulation (though still exponentially large in system size) and this can enable large, even exponential, improvements in sample complexity. However, such schemes require more complex circuit implementations, and analytical tasks—particularly the analysis and




Figure 5.2: *Symmetry respecting randomization protocols* We introduce three symmetry-aware random measurement protocols—Global and Local Dual Pairs and Dual Product—designed to estimate gauge-invariant observables efficiently, leveraging prior knowledge that expectation values of gauge-variant operators vanish. For Global Dual Pairs, random pairs of Ising qubits are chosen and random two-body symmetry-respecting unitaries associated with these degrees of freedom are mapped back to the LGT side of the duality as the randomizing operations prior to measurement. Local Dual Pairs is similar but leverages a promise of geometric locality in the observables of interest to limit the choice of pairs to local patches. Finally, the Dual Product protocol is constructed from the standard Product Protocol applied to the dual Ising model. This protocol slightly breaks symmetry when we have periodic boundary conditions (PBC) since such operations are not parity-respecting. A trade-off exists between sample complexity and circuit depth.

inversion of the associated shadow channels—become challenging.

In this chapter, we introduce three protocols for sample-efficient measurement of gauge-invariant observables in $\mathbb{Z}_2$ LGT in (2+1)D, for which a dual Ising formulation allows rigorous analysis. We analyze the trade-offs between sampling and circuit complexity, provide an analytic inversion of the quantum channel underlying the shadow protocol, and establish rigorous performance guarantees. The three variants are:

- a. **Dual Pairs Protocols.** These are inspired by the “All-Pairs” approach of Ref. [387], where shot-wise random connectivity between qubit pairs is used to implement particle-number-preserving randomization. Our protocols additionally leverage the LGT–Ising duality: most components are analyzed ‘virtually’ via the duality, while physical operations are carried out on the LGT state realized on the device. Unlike Ref. [387], where particle-number-preserving randomization was introduced primarily for hardware reasons, here it directly reduces sampling cost. We develop two variants in Section 5.3:
 - i. The **Global Dual Pairs Protocol**, applicable to arbitrary gauge-invariant observables, exponentially improves sampling efficiency over standard protocols.
 - ii. The **Local Dual Pairs Protocol** is tailored for geometrically local gauge-invariant observables, yields further improvements in sampling efficiency and reduces both sampling and circuit complexity compared to the Global Dual Pairs Protocol.
- b. The **Dual Product Protocol** exploits the Ising–LGT duality without relying on an all-pairs structure. It achieves the best sampling efficiency among our protocols but is also the most demanding in terms of circuit resources, requiring depth $\mathcal{O}(V^2)$ (with V the system volume) and an additional ancilla qubit for periodic boundary conditions; it is discussed in Section 5.4.

All protocols are benchmarked against the ‘standard’ classical shadow protocol, the **Product Protocol**, which applies single-qubit Clifford randomization [37], and which uses no prior information about the state or the observables to be predicted.

The paper is organized as follows: In Section 5.2, we introduce the $\mathbb{Z}_2$ LGT and its duality to the Ising model. Section 5.3 presents the two variants of the Dual Pairs Protocol—the Global Dual Pairs and Local Dual Pairs protocols. In Section 5.4, we describe the Dual Product Protocol. Numerical demonstrations of all protocols are given in Section 5.5. Finally, in Section 5.7, we discuss the implications of our results and outline future directions, including extensions to other LGTs, potentially beyond the Abelian case.

5.2 Lattice Gauge Theory

Lattice Gauge Theories (LGTs), are models in which the local degrees of freedom, on links (i.e. edges of the lattice graph), are associated with elements of a group, subject to a set of local symmetry operators imposing constraints on the Hilbert space that involves their neighbors. These operators, Gauss laws, may not mutually commute (they do for LGTs based on Abelian groups), but commute with the Hamiltonian ensuring invariance of eigenstates under gauge transformations, see e.g., Refs. [49, 401, 402] for reviews.

A simple example is the $\mathbb{Z}_2$ LGT [403, 404] in 2+1 spacetime dimensions. The underlying Abelian group $\mathbb{Z}_2$ has two elements, corresponding to local states encoded in spin- $\frac{1}{2}$ degrees of freedom (qubits) residing on the links of a rectangular $N_x \times N_y$ lattice; see Fig. 5.1. The Hamiltonian is given by

$$H = - \sum_{\square} W_{\square} - g \sum_l \sigma_l^x, \quad (5.1)$$

where $W_{\square} := \prod_{l \in \square} \sigma_l^z$ is the plaquette operator and $\sigma_l^{x,z}$ are Pauli operators acting on link l . Gauss law operators are defined at each site s (node of the lattice graph) as

$$G_s = \prod_{l \in s} \sigma_l^x, \quad (5.2)$$

where $l \in s$ denotes links emanating from a site. These operators commute with the Hamiltonian, $[H, G_s] = 0$, and thus label superselection sectors; that with $G_s |\psi\rangle = |\psi\rangle \forall s$ is usually of interest (called the “physical” subspace). Gauge-invariant operators are invariant under local symmetry transformations

$$U O_S U^\dagger = O_S, \quad U := \prod_{s \in S} (G_s)^{\alpha(s)}, \quad (5.3)$$

and their expectation value can be non-zero in a physical state. Here, $\alpha(s)$ is an integer for cyclic groups such as $\mathbb{Z}_2$ ($\alpha(s) \in \{0, 1\}$ for $\mathbb{Z}_2$), and S denotes the set of sites (or links between sites) on which O_S has support. For LGTs with continuous groups, G_s are replaced by exponentiated Gauss law *generators*, $G_s \rightarrow e^{iG_s}$ with real, continuous $\alpha(s)$; multiple non-commuting Gauss law operators exist per site in the non-Abelian case [49, 401, 402]. A lattice can feature different boundary conditions, which can lead to additional superselection sectors. For periodic boundary conditions (PBCs), discussed in the main text, two additional operators $V_x \equiv \prod_{\ell \in \mathcal{P}_x} \sigma_\ell^x$, and $V_y \equiv \prod_{\ell \in \mathcal{P}_y} \sigma_\ell^x$ define superselection sectors; the periodic paths $\mathcal{P}_{x/y}$ are shown in the right panel of Fig. 5.1 in blue and red, respectively. Other boundary conditions, such as fixed boundary conditions (FBC) are discussed in Appendix D.2.

The $\mathbb{Z}_2$ LGT is dual to an Ising model in the following sense: Let $\mathcal{L}(\mathcal{H})$ denote the space of bounded linear operators acting on the (physical) LGT Hilbert space $\mathcal{H}_{\text{phys}}$. A fixed (V_x, V_y) sector can be mapped exactly onto a dual Ising model [403–406], using the map, $\Phi : \mathcal{L}(\mathcal{H}_{\text{phys}}) \rightarrow \mathcal{L}(\mathcal{H}_{\text{dual}})$, with dual (qubit) d.o.f.s placed at the centers of plaquettes $\square$,

$$\Phi(W_{\square}) = Z_{\square}, \quad \Phi(\sigma_l^x) = X_{\square} X_{\square + \hat{a}}, \quad (5.4)$$

where, w.o.l.g, $V_x = V_y = 1$, $X_{\square}$ and $Z_{\square}$ are Pauli operators, and $\square + \hat{a}$ are the plaquettes adjacent to link l . The inverse map is

$$\Phi^{-1}(Z_{\square}) = W_{\square}, \quad \Phi^{-1}(X_{\square}) = \prod_{l \in \gamma_{\square}} \sigma_l^x. \quad (5.5)$$

Here $\gamma_{\square}$ is a path from a reference plaquette $\square_r$ to $\square$, as depicted in the bottom left panel of Fig. 5.1. The choice of $\square_r$ is arbitrary and no observable depends on it. The corresponding dual Hamiltonian is

$$H = - \sum_{\square} Z_{\square} - g \sum_{\square} (X_{\square} X_{\square - \hat{x}} + X_{\square} X_{\square - \hat{y}}), \quad (5.6)$$

a (2+1)D transverse field Ising model with PBC; $\square - \hat{x}$ ($\square - \hat{y}$) are the plaquette to the left (below) of $\square$. For PBC (but not for other boundary conditions), the product of plaquette operators satisfies the operator identity $\prod_{\square} W_{\square} = 1$. This maps, on the Ising side, to a parity constraint $\prod_{\square} Z_{\square} = 1$ required for the duality to hold. A counting of degrees of freedom confirms the exactness of the duality: the LGT has $2N_x N_y$ qubits (one per link), subject to $N_x N_y - 1$ independent Gauss law constraints. Including the decomposition into sectors labeled by V_x and V_y , this results in $N_x N_y - 1$ physical degrees of freedom per sector—precisely matching the number of d.o.f.s in the parity-constrained Ising model. Thus, the physical Hilbert space within a given V_x, V_y sector is of dimension $2^{N_x N_y - 1}$ whereas the full Hilbert space is exponentially larger with dimension $2^{2N_x N_y}$. The duality, which can be generalized to arbitrary boundary conditions, will allow us to design several symmetry-respecting random measurement protocols, and to derive explicit performance guarantees, in the next two sections.




Figure 5.3: Schematic overview of the Global Dual Pairs protocol with periodic boundary conditions (PBC).

Step 1. A pairing, π , of dual lattice sites is chosen uniformly at random. For each pair, a random unitary $U_{[ij]}$ is constructed in the Ising formulation. **Step 2.** For each pair $[ij] \in \pi$, the dual unitary $\mathcal{U}_{[ij]} = \Phi^{-1}[U_{[ij]}]$ is implemented on the LGT side of the duality and a computational basis measurement is performed. **Step 3.** The output bit string is mapped back to the dual Ising theory, where efficient shadow channel inversion is performed to estimate expectation values of gauge-invariant observables O .

5.3 Dual Pairs Protocols

Here, we present two protocols—the Global Dual Pairs and Local Dual Pairs schemes—for estimating gauge-invariant observables. Both protocols improve sample complexity compared to the symmetry-ignorant standard Product Protocol [37], which applies local Clifford randomization, here, to the $2N_x N_y$ qubits associated with the links of the $\mathbb{Z}_2$ LGT. The advantage arises because the Product Protocol randomizes the exponentially larger full Hilbert space, whereas our protocols randomize only over the physical subspace. The Global Dual Pairs protocol is the most general, requiring only that the target observables be gauge-invariant, and achieves an exponential reduction in sampling complexity relative to the Product Protocol. The

Local Dual Pairs protocol, designed for geometrically local gauge-invariant observables, provides additional improvements in both sampling and circuit complexity.

5.3.1 Global Dual Pairs Protocol

Overview

The Global Dual Pairs protocol for estimating gauge-invariant observables, without additional restrictions, is illustrated schematically in Fig. 5.3. For simplicity of presentation, we assume that both N_x and N_y are even, although this is not a requirement. The system realized in experiment, where measurements are performed, is the LGT; the dual Ising model serves purely as a conceptual tool for constructing the protocol and a computational tool for the classical pre- and post-processing. It is applied to a physical input state in the LGT, i.e., a simultaneous $+1$ eigenstate of all Gauss law operators. The protocol proceeds as follows:

Step 1 (Classical Pre-processing): Consider the dual Ising model consisting of $V = N_x \cdot N_y$ spins. Select, uniformly at random, a pairing π that partitions all spins into disjoint pairs. For each pair $[ij] \in \pi$, choose a random two-qubit unitary $U_{[ij]}$. For systems with PBCs, a technical, but non-essential, requirement is that $U_{[ij]}$ respects the $\mathbb{Z}_2$ parity symmetry, as detailed below. No such constraint applies e.g. for FBCs. This Ising-side random unitary factorizes as

$$U_\pi = \bigotimes_{[ij] \in \pi} U_{[ij]}. \quad (5.7)$$

Step 2 (Quantum Measurement): Perform the random measurement on the quantum state on the $\mathbb{Z}_2$ LGT side of the duality associated with the dual unitary corresponding to Eq. (5.7):

$$\mathcal{U}_\pi := \Phi^{-1}(U_\pi). \quad (5.8)$$

A circuit decomposition of $\mathcal{U}_\pi$ can be computed from U_π with classical time complexity $\mathcal{O}(V)$ by applying Eq. (5.5) to each of the Pauli rotation gates in the circuit decomposition of U_π . Following this rotation, perform computational basis measurements on all LGT qubits to obtain a bit string $\vec{s} \in \{0, 1\}^{2V}$.

Step 3 (Classical Post-processing): Map the bit string $\vec{s}$ to an associated Ising dual bit string $\vec{b} \in \{0, 1\}^V$ for classical post-processing on the Ising side of the duality. In particular,

$$|\vec{b}\rangle \langle \vec{b}| := \Phi(|\vec{s}\rangle \langle \vec{s}|). \quad (5.9)$$

Mapping from a bit string $\vec{s}$ to a bit string $\vec{b}$ reduces to determining the parity around each plaquette and can be performed with classical time complexity $\mathcal{O}(V)$. Details can be found in Section 5.3.1.

Finally, invert the shadow channel $\mathcal{M}$ defined on the Ising side of the duality as

$$\mathcal{M} \circ \Phi(\rho) := \mathbb{E}_{U, \vec{b}} \left[U^\dagger |\vec{b}\rangle \langle \vec{b}| U \right], \quad (5.10)$$

where $\circ$ denotes the composition of linear maps and U is the Ising-side unitary generated in Step 1. Here, ρ is the state of the system as on the LGT side of the duality and $\Phi(\rho)$ is the associated state in the dual picture. This channel can be easily inverted to obtain an Ising-side classical shadow

$$\widehat{\Phi(\rho)} := \mathcal{M}^{-1} \left(U^\dagger |\vec{b}\rangle \langle \vec{b}| U \right), \quad (5.11)$$

associated with each sampled $U, \vec{b}$. The details are described below. Using these shadows, expectation values of local observables can be computed once the observables of interest have also been mapped to the dual picture, again with a classical cost $\in \mathcal{O}(V)$.

At this point, it is important to remind the reader that the experiment, including the states, unitary operations, and measurements, is implemented on LGT states (in Step 2), while the Ising duality is only a conceptual and (classical) numerical tool for defining the unitaries and efficiently inverting the associated channel.

Random Symmetry-Respecting Unitaries

In this section, we detail the structure and cost of implementing the random symmetry-respecting unitaries $\mathcal{U}_\pi$ constructed in the Ising picture in **Step 1** and implemented on the LGT states in **Step 2**. For systems with PBCs, we choose each local unitary $U_{[ij]}$ to respect parity symmetry. In the absence of a symmetry constraint, $U_{[ij]}$ instead is simply chosen as a generic two-qubit Haar-random unitary. Specifically, for the case with parity symmetry, we consider a decomposition,


$$U_{[ij]} = U_{\text{odd}} \oplus U_{\text{even}}, \quad (5.12)$$

where $U_{\text{odd}}, U_{\text{even}}$ are Haar random 2×2 unitaries independently drawn from a circular unitary ensemble (CUE) acting on the relevant parity sector. Such unitaries can be parameterized, up to an irrelevant global

phase, as [361]

$$U_{\text{odd/even}} = \begin{pmatrix} e^{i(\alpha+\gamma)} \cos(\beta) & e^{-i(\alpha-\gamma)} \sin(\beta) \\ -e^{i(\alpha-\gamma)} \sin(\beta) & e^{-i(\alpha+\gamma)} \cos(\beta) \end{pmatrix} \quad (5.13)$$

for appropriate choices of α , β , and γ , with circuit realization



$$(5.14)$$

where $U_{\pm}(\beta) \equiv \exp\{i\frac{\beta}{2}(Y_i X_j \pm X_i Y_j)\}$ correspond to the even/odd parity sectors. (We work in the even sector.) To translate this circuit to the LGT side of the duality, we use

$$\Phi^{-1}(Z_{\square}) = W_{\square}, \quad \square \in \{i, j\}, \quad (5.15)$$

and

$$\begin{aligned} \Phi^{-1}(Y_i X_j \pm X_i Y_j) = \\ i \prod_{l \in \gamma_{ij}} \sigma_l^x \prod_{l' \in \square_j} \sigma_{l'}^z \pm i \prod_{l \in \gamma_{ij}} \sigma_l^x \prod_{l' \in \square_i} \sigma_{l'}^z, \end{aligned} \quad (5.16)$$

where γ_{ij} is some choice of path between plaquettes i and j , as depicted in Fig. 5.4 for the $Y_i X_j$ term of this operator.

Eq. (5.16) implies that implementing $\Phi^{-1}(U_{\pm}(\beta))$ requires a unitary rotation with Pauli weight that scales with the distance between the plaquettes i, j . As the Pauli weight of a rotation directly determines the circuit depth of a decomposition into a standard gate set, this results in a higher circuit depth when compared to the standard Product protocol; for the Global Dual Pairs protocol it is $\mathcal{O}(V)$. Another consideration is the choice of paths γ_{ij} , which determine how the unitaries $\mathcal{U}_{[ij]}$ are implemented in parallel. In practice, a reasonable, though not necessarily optimal, solution to this path selection problem can be obtained using a greedy algorithm [407] that identifies non-intersecting paths for each pair $[ij] \in \pi$.

As an upper bound, without parallelization, the circuit depth scales as $\mathcal{O}(V^2)$: one factor of V arises from the number of such pairs, and the second from the decomposition of each high-weight Pauli rotation into elementary gates. In Section 5.6.2 we shall see that the large circuit depth for the Global Dual Pairs protocol compared to the Product protocol is counteracted by a large reduction in sample complexity.




Figure 5.4: *LGT–Ising duality for mapping a (parity-respecting) entangling operation $\Phi^{-1}(Y_i X_j)$.* While this operation is two-body on the Ising side (left panel), it corresponds to an extended operation in the LGT (right panel) acting between plaquettes i and j along a path γ_{ij} , which is arbitrary except for its fixed endpoints.

Mapping of Measurement Outcomes

In this subsection, we discuss the first part of **Step 3** of the Global Dual Pairs protocol where we map measurement outcomes $\vec{s} \in \{0, 1\}^{2V}$ to dual bit strings $\vec{b} \in \{0, 1\}^V$ as described in Eq. (5.9). The mapping is relatively simple: the parity of the measurement outcomes around a given plaquette corresponds to the measurement outcome that would result if we had directly measured $Z_\square$ on that same plaquette in the dual Ising theory. Thus, computing the parity of $\vec{s}$ for each plaquette allows us to determine the bit string $\vec{b}$ that results from a computational basis measurement in the dual picture. This process requires classical time complexity $\mathcal{O}(V)$. After this mapping, we treat the inversion of the channel purely as a classical shadows scheme in the dual picture, as elaborated further in the following section.

Inverting the Channel

In this subsection, we discuss the second part of **Step 3** of the Global Dual Pairs protocol—the inversion of the channel defined in Eq. (5.10). We focus on the case with PBCs, where the Ising dual retains a global parity constraint. Our approach to analyzing the channel inversion adapts techniques originally developed

in Ref. [387] for systems with particle number conservation. Other boundary conditions, where no such symmetry remains on the Ising side, lead to somewhat different (but simpler) inversion procedures and are discussed in Section D.2.

Let P_V denote the set of all pairings of the $N_x N_y = V$ sites of the dual lattice (centers of plaquettes in the original lattice). Assuming that N_x and N_y are both even, V is also even so that all sites can be paired up. The number of pairings is $|P_V| := (V - 1)!!$. Formulated on the Ising side of the duality, the Global Dual Pairs channel is given as an average over channels defined for each of the possible pairings $\pi \in P_V$ (as each pairing is selected with uniform probability):

$$\mathcal{M}(\cdot) = \frac{1}{|P_V|} \sum_{\pi \in P_V} \mathcal{M}_\pi(\cdot). \quad (5.17)$$

Then, $\mathcal{M}_\pi(\rho)$ decomposes into a product over channels acting non-trivially on each pair $[ij] \in \pi$ yielding

$$\mathcal{M}(\cdot) = \frac{1}{|P_V|} \sum_{\pi \in P_V} \prod_{[ij] \in \pi} \mathcal{M}_{[ij]}(\cdot), \quad (5.18)$$

Finally, these pairwise channels act on the Ising-side state $\Phi(\rho)$ as

$$\mathcal{M}_{[ij]} \circ \Phi(\rho) = \mathbb{E}_{U_{[ij]}, b_i, b_j} \left[U_{[ij]}^\dagger |b_i b_j\rangle \langle b_i b_j| U_{[ij]} \right], \quad (5.19)$$

where $U_{[ij]}$ are two-qubit unitaries selected according to Eq. (5.12) and b_i, b_j are the measurement outcomes mapped to the dual qubits i, j , as described in the previous section. Because of the parity symmetry, the image $\Phi(\rho)$ block diagonalizes into even and odd parity sectors as

$$\Phi(\rho) = \Phi(\rho)_{\text{even}} \oplus \Phi(\rho)_{\text{odd}}. \quad (5.20)$$

The same decomposition occurs for the channel $\mathcal{M}(\cdot)$. Thus, we can evaluate $\mathcal{M}_{[ij]} \circ \Phi(\rho)$ as a direct sum over the even and odd parity blocks. Using that $U_{\text{odd}}, U_{\text{even}}$ are sampled from a 2-design, the combined action of channel and isomorphism on the state ρ is [408]

$$\mathcal{M}_{[ij]} \circ \Phi(\rho) = \bigoplus_{p \in \text{even, odd}} \frac{1}{3} \left(\Phi(\rho)_p + \text{Tr}[\Phi(\rho)_p] \right). \quad (5.21)$$

Any gauge-invariant observable can be written on the Ising side of the duality in the basis of Pauli strings $S \in \{I, Z, X, Y\}^{\otimes V}$ with an even number of X, Y operators. Using Eq. (5.21), we evaluate the action of $\mathcal{M}_{[ij]}$ on substrings $S_i S_j$,

$$\mathcal{M}_{[ij]}[S_i S_j] =$$

$$\begin{cases} S_i S_j, & S_i S_j \in \{II, ZZ\} \\ \frac{1}{3} S_i S_j, & S_i S_j \in \{IZ, ZI\} \cup \{X, Y\}^2 \\ 0, & \text{otherwise,} \end{cases} \quad (5.22)$$

indicating that $\mathcal{M}$ is diagonal, making inversion almost trivial: the non-trivial component of the inversion is all contained in the Ising duality. In particular, for a Pauli string S ,

$$\mathcal{M}(S) = \frac{f\alpha}{3^{w_{XY}/2}} S, \quad (5.23)$$

where $w_{XY} = w_X + w_Y$ is the total number of X and Y operators in the string S , f is the fraction of pairings $\pi \in P_V$ for which $\mathcal{M}_\pi(S) \neq 0$,

$$f = |P_{w_{XY}}| \frac{|P_{V-w_{XY}}|}{|P_V|}, \quad (5.24)$$

where $|P_m|$ is the number of pairings of m objects, and α is an amplitude associated with permutations of the Z and I operators within S :

$$\alpha = \frac{1}{|P_{w_I+w_Z}|} \sum_{\substack{m=0 \text{ s.t.} \\ m \equiv w_Z \pmod{2}}}^{w_Z} \underbrace{\left(\frac{1}{3}\right)^m \binom{w_Z}{m} \binom{w_I}{m} m!}_{\text{num. of } IZ \text{ pairings}} \underbrace{|P_{w_Z-m}| |P_{w_I-m}|}_{\text{num. of } ZZ, II \text{ pairings}}, \quad (5.25)$$

where w_Z and w_I are the number of Z and I operators in the operator string S respectively, and $V = w_{XY} + w_Z + w_I$. The sum in Eq. (5.25) is over the number of pairs m of Z operators with I operators, where the notation indicates that m increments by two for each term in the sum, starting from either 0 or 1 depending on the parity of w_Z . As seen in Eq. (5.22) these pairings are the pairings that yield a non-trivial contribution to the amplitude. If $w_Z = 0$, then $\alpha = 1$.

Upon inverting $\mathcal{M}$, the number of samples ν needed for predicting, with high probability, the expectation values of M observables $\{O_j\}_{j=1}^M$ of interest up to an additive error ϵ can be bounded as [37],

$$\nu \in \mathcal{O} \left(\frac{\log M}{\epsilon^2} \max_j \text{Var}(o_j) \right), \quad (5.26)$$

where $o_j := \text{Tr}[\widehat{\Phi[\rho]} O]$, where $\widehat{\Phi[\rho]}$ is the (Ising side) shadow as given in Eq. (5.11). The operator variance can be bounded as (see Section D.1):

$$\text{Var}(o_j) \leq \frac{3^{w_{XY}/2}}{f\alpha} \|O_j\|_\infty^2. \quad (5.27)$$

Thus, to reach constant error ϵ a number of samples scaling as $3^{w_{XY}/2}/(f\alpha)$ are needed. While the expressions for f and α in Eqs. (5.24) and (5.25) are somewhat unwieldy, it is easy to determine whether the scaling is polynomial or exponential in system size. In particular, under the assumption that, asymptotically in system size, $w_I = \Theta(V)$ and $w_{XY} + w_Z = \mathcal{O}(1)$, we can use Stirling's approximation to show that

$$f\alpha = \Theta\left(V^{(w_Z - k_{\text{dual}})/2}\right). \quad (5.28)$$

Note that $w_Z \leq k_{\text{dual}}$. Thus, for constant norm observables, a polynomial number of samples in the system size are needed to reach a constant error ϵ . The full costs of this protocol, including the $\mathcal{O}(V)$ costs of performing the mappings between the LGT and Ising models and the $\text{poly}(V)$ costs of computing the coefficients f and α , are summarized in Table 5.1.

5.3.2 Local Dual Pairs Protocol

In this section, we introduce a variant of the Global Dual Pairs protocol, referred to as Local Dual Pairs protocol, designed for the measurement of geometrically local observables rather than general k -local ones. This modification exploits the assumption that all observables of interest act non-trivially only on local patches of the lattice with constant size $L \times L$. Under this locality constraint, the sample complexity, classical cost of channel inversion, and the Pauli weight of the required quantum circuits are all significantly reduced.

For simplicity, we assume that N_x and N_y are divisible by L . We consider the set $\mathcal{T}$ of all distinct tilings of the lattice into adjacent $L \times L$ patches. As shifting a given tiling by L units in either the x or y direction on the lattice recovers the same tiling, the number of tilings is $|\mathcal{T}| = L^2$. With this setup, the Local Dual Pairs protocol proceeds as follows:

Step 1 (Classical Pre-processing): Pick a random tiling $\tau \in \mathcal{T}$. Then, pick a pairing $\pi \in P_{L^2}$ of sites in a single tile and consider that same pairing in each tile. The set of all such pairings P_{L^2} per tiling $\tau \in \mathcal{T}$ has size $|P_{L^2}| = (L^2 - 1)!!$. As in the Global Dual Pairs protocol, construct a unitary U_π associated with the pairing π .

Step 2 (Quantum Measurement): As in the Global Dual Pairs protocol, apply the unitary $\mathcal{U}_\pi = \Phi(U_\pi)$ on the $\mathbb{Z}_2$ side of the duality and make a computational basis measurement.




Figure 5.5: *Overview of the Local Dual Pairs protocol.* The Local Dual Pairs protocol is a variant of our scheme for geometrically local gauge-invariant observables that have support on a patch of at most $L \times L$. **Step 1** consists of picking a random tiling $\tau \in \mathcal{T}$ and, then, picking a random pairing repeated within each patch of the tiling in order to construct an associated random unitary U_π . **Step 2** consists of applying the associated unitary $\mathcal{U}_\pi$ on the LGT side of the duality, followed by measurement in the computational basis. **Step 3** is the post-processing step, which proceeds essentially identically to the Global Dual Pairs protocol restricted to a given patch of a tiling, where the data is filtered to a single choice of tiling according to the support of a given target variable.

Step 3 (Classical Post-processing): Again, as in the Global Dual Pairs protocol, map the output bit string to its corresponding bit string $\vec{b}$ in the Ising dual. Repeat the above sequence of steps a total of ν times. To estimate a geometrically local observable O , we filter the dataset to retain only those samples corresponding to a single $\tau \in \mathcal{T}$ for which O is supported within one of its patches. This filtering is efficient as the number of tilings is L^2 . From this subset, we restrict the corresponding dual output bit strings $\vec{b}$ and the unitaries $\bigotimes_{\text{patches}} U_\pi$ to the patch on which O acts non-trivially. Using this restricted data, we perform the post-processing as in the Global Dual Pairs protocol on the local patch. For data restricted to this patch, the associated channel $\mathcal{M}$ coincides with that of the Global Dual Pairs protocol and can therefore be inverted in the same manner.

Given the filtered data of a single relevant tiling $\tau \in \mathcal{T}$, the fraction f_{local} of pairings that contribute in the Local Dual Pairs protocol

$$f_{\text{local}} = |P_{w_{XY}}| \frac{|P_{L^2 - w_{XY}}|}{|P_{L^2}|}, \quad (5.29)$$

and the coefficient α are calculated as in Eq. (5.25), but with w_I replaced by the number of identity operators in a given basis operator string S within the relevant patch.

A diagram summarizing this protocol can be found in Fig. 5.5 and a summary of the cost of the protocol can be found in Table 5.1.

5.4 Dual Product Protocol

At this stage, the reader may have noticed that the parity symmetry, and the associated “all-pairs” structure on the Ising side, are not essential for an efficient measurement scheme (indeed, no parity symmetry exists under other boundary conditions). The improvement in sampling complexity arises solely from exploiting the LGT–Ising duality which eliminates the exponential overhead between the gauge-invariant subspace and the full Hilbert space.

This suggests an alternative approach: (virtually) applying the standard Product Protocol on the Ising side of the duality, which we refer to as the Dual Product Protocol. The scheme is conceptually straightforward, apart from one important caveat: For PBC, the protocol mixes the even and odd parity sectors of the dual Ising theory, where only the even sector is dual to the LGT. On the LGT side, the parity constraint appears as an operator identity, $\prod_{\square} W_{\square} = 1$, rather than a symmetry: there is no component of the full LGT Hilbert space corresponding to the odd sector.

To obtain a one-to-one correspondence between the full Ising theory (including both parity sectors) and an (extended) LGT, we introduce on the LGT side an ancilla qubit a to control the lattice boundary condition, as shown in Fig. 5.6. Picking an *arbitrary* reference plaquette $\square_r$ and selecting one of its four boundary links r , we designate that link as the control and the ancilla a , initially in the state $|0\rangle$, as the target of a $\text{CNOT}_{r \rightarrow a}$ operation. As illustrated in Fig. 5.6, this operation effectively “copies” the magnetic (z -basis) state of link r into the ancilla, while leaving the electric basis unchanged. The two plaquettes that originally shared link r are thereby separated: one now contains r , and the other contains the ancilla a . The plaquette identity $\prod_{\square} W_{\square} = 1$ then becomes $\prod_{\square} W_{\square} = \sigma_a^z \sigma_r^z$. Thus, we have effectively promoted the operator identity $\prod_{\square} W_{\square} = 1$ to a symmetry $\prod_{\square} W_{\square} = \sigma_a^z \sigma_r^z$ with the original theory without the ancilla corresponding to the sector where $\sigma_a^z \sigma_r^z = 1$.

Introducing the ancilla also affects electric operators (defined in the x -basis). In particular,

$$\text{CNOT}_{r \rightarrow a} \sigma_r^x \text{CNOT}_{r \rightarrow a} = \sigma_r^x \sigma_a^x \quad (5.30)$$

so that every electric operator involving the reference link r —including Gauss-law and ribbon operators—is effectively modified as $\sigma_r^x \rightarrow \sigma_r^x \sigma_a^x$. For example, the two Gauss-law operators that originally shared r now each become five-legged, see Fig. 5.6. The states of the original $\mathbb{Z}_2$ LGT with periodic boundary conditions are mapped by the $\text{CNOT}_{r \rightarrow a}$ operation to the extended system including the ancilla where the original z -state of link r is encoded in the two-qubit subspace $|00\rangle, |11\rangle$ of the pair (r, a) . Crucially, however, the ancilla allows the boundary condition to be dynamically controlled by the state of a , enabling superpositions of PBC ($\prod_{\square} W_{\square} = 1$) and twisted PBC (tPBC, $\prod_{\square} W_{\square} = -1$). In this enlarged formulation of the LGT-Ising duality the *entire* Ising space can now be represented on the LGT side: $\sigma_a^z \sigma_r^z = 1$ (PBC) corresponds to the even-parity sector, while $\sigma_a^z \sigma_r^z = -1$ (tPBC) corresponds to the odd-parity sector.¹

The Dual Product protocol now consists simply of single-qubit Clifford rotations on the Ising side, which mix the parity sectors. On the LGT side of the duality this mixes the $\prod_{\square} W_{\square} = 1$ (PBC) and $\prod_{\square} W_{\square} = -1$ (tPBC) sectors. By virtue of the initial state being in the $\prod_{\square} W_{\square} = 1$ (PBC) sector the shadow protocol that breaks this symmetry will still yield an estimate of $\prod_{\square} W_{\square} = 1$, in expectation.

In detail, the steps of the Dual Product Protocol are as follows:

¹This is analogous, e.g., to how ancillas can induce disorder-free localization without explicitly breaking lattice symmetries [409, 410].




Figure 5.6: An ancilla enables mixing PBC (even parity) and tPBC (odd parity) sectors. **a)** Adding an ancilla a to a reference link r allows us to enforce either PBC or tPBC in the $\mathbb{Z}_2$ LGT. **b)** In the Dual Product protocol, the z -component of the state of the link r is “copied” to the ancilla a via a CNOT. The randomizing unitaries, chosen according to the usual Product Protocol, but applied to the Ising degrees of freedom, then mix between the PBC and tPBC sectors on the LGT side of the duality. In the right panel, we show how the plaquette and Gauss law operators change with the ancilla a : The plaquette operators no longer share a link, while the Gauss law operator is effectively extended by a fifth leg, $\sigma_r^x \rightarrow \sigma_r^x \sigma_a^x$.

Step 1 (Classical Pre-processing): Select a unitary operator U_{prod} on the dual Ising qubits of the form $U_{\text{prod}} = \bigotimes_j U_j$ where each U_j is a random single qubit Clifford gate acting on site j (i.e. Hadamard or S gates, which may be generated by single qubit X and Z rotations).

Step 2 (Quantum Measurement): Introduce the ancilla a in the LGT state and apply $\text{CNOT}_{r \rightarrow a}$, effectively “copying” the σ^z -basis value of an arbitrary link r into a , while also transforming $\sigma_r^x \rightarrow \sigma_r^x \sigma_a^x$. The original state is preserved, now embedded in a slightly larger Hilbert space. Then, use $\mathcal{U}_{\text{prod}} = \Phi^{-1}(U_{\text{prod}})$ to apply operations on the LGT state: the generalized duality is understood to map two neighboring Z




Figure 5.7: *Schematic overview of the Dual Product protocol with periodic boundary conditions (PBC). **Step 1.** Randomly choose a unitary U_{prod} consisting of a single qubit Clifford gate on each Ising qubit. **Step 2.** The dual unitary $\mathcal{U}_{\text{prod}} = \Phi^{-1}[U_{\text{prod}}]$ is implemented on the LGT side of the duality and a computational basis measurement is performed. **Step 3.** The output bit string is mapped back to the dual Ising theory, where efficient shadow channel inversion is performed to estimate expectation values of gauge-invariant observables O .*

operators on the Ising side to separate plaquettes on the LGT side—one containing r and the other containing a . Simultaneously, every occurrence of σ_r^x is replaced by $\sigma_r^x \sigma_a^x$ when mapping electric operations (i.e., those involving X in the Ising dual). These operations mix the parity-even and parity-odd sectors of the Ising dual

and create superpositions between PBC and tPBC in the LGT (see Fig. 5.7 for an illustration). Following this we perform computational basis measurements to obtain a bit string $\vec{s}$.

Step 3 (Classical Post-processing): Map the bit string $\vec{s}$ to the corresponding Ising dual bit string $\vec{b}$ for channel inversion. Compared to the previously discussed protocols, $\vec{s}$ is now one bit longer due to the ancilla, but the mapping follows Eq. (5.9): the parity around each plaquette is determined, with the key difference that plaquette $\square_r$ no longer shares a link with its neighbor. As a result, the parities of the bit strings associated with each plaquette are now completely independent. Once $\vec{b}$ is obtained, post-processing reduces to the standard Product protocol applied to the Ising degrees of freedom, with a straightforward shadow channel inversion as outlined in [37].

We emphasize once more that the protocol effectively *estimates* the LGT (magnetic) boundary condition (i.e., the parity on the Ising side), thus effectively discarding previously known information, with periodic boundary conditions appearing only in expectation. The protocol can estimate *arbitrary gauge invariant observables* for which its sample complexity can be better than the Global Dual Pairs protocol: In particular,

$$\text{Var}(o_j) \leq 4^{k_{\text{dual}}} \|O_j\|_{\infty}^2, \quad (5.31)$$

where k_{dual} is the locality of the operator of interest in the dual theory, and $\|O\|_{\infty}$ is evaluated on the dual Ising Hilbert space [37]. Thus, via Eq. (5.26), the sample cost is constant in system size for gauge invariant observables that are k_{dual} -local on the Ising side of the duality, while for Global Dual Pairs the sample cost is polynomial in system size.²

For costs beyond sample complexity, the dominant classical computational cost is due to the mapping between the Ising theory and the LGT, which is linear in system size. Finally, the required circuit depth to implement the unitary rotations on the LGT side is quadratic in the system size due to the need for rotations generated by ribbon operators starting at the reference qubit r .

5.5 Numerical Demonstrations

To enable the use and testing of our protocols, we provide Python code to perform the analysis for Global Dual Pairs and the Dual Product protocol [411]. To demonstrate the accuracy and sampling costs of the symmetry

²Here, we use locality in the sense that a k -local operator acts non-trivially on k qubits (has Pauli weight k), in contrast to the notion of geometric locality leveraged in the Local Dual Pairs protocol.

aware protocols (Global Dual Pairs and Dual Product) compared to the standard symmetry-ignorant Product protocol for estimating gauge-invariant observables, we numerically simulate the protocols applied to the ground state of the $\mathbb{Z}_2$ -LGT Hamiltonian in Eq. (5.1) for various system sizes and couplings g . We estimate both ribbon operators O_R with $\Phi[O_R] = X_i X_j$ and loop operators $O_L = \prod_{l \in L} \sigma_l^z$, where L is a closed loop, as pictured in Fig. 5.8a. Our protocols accurately recover the expectation values of these observables, and we perform a system size scaling analysis. As predicted by the analysis in the previous section, for observables with extensive non-trivial support on the LGT side of the duality, we observe exponential (in system size) sample complexity improvements of the Global Dual Pairs versus the Product protocol.

Fig. 5.8b shows a minimal example comparing Global Dual Pairs, Dual Product, and Product estimation for both a ribbon operator and loop operator as compared to the true values as a function of the coupling strength g for a fixed system size $(N_x, N_y) = (3, 2)$. We consider (i) a ribbon operator of the form $O_R = \sigma_i^x \sigma_j^x$ and (ii) a single plaquette loop operator of the form $O_L = \prod_{j \in \square} \sigma_j^z$. The estimates for these example observables are computed using a fixed number of experimental shots $\nu = 1000$ for all protocols. The plots show the mean and standard deviation of these estimates over 50 runs of the experiment. In each experiment, we employ a median-of-means estimator with block sizes of approximately $\sqrt{\nu}$, which is the standard choice in this setting as it suppresses the effect of outliers in the shadow table data [37].³ The true expectation values of the observables are also shown, and we see that all protocols provide accurate estimates with comparable standard deviation of the mean. As expected, the error bars for a fixed number of shots for the Global Dual Pairs and Dual Product protocols compared to the Product protocol depends on the relative weight of the observables on the LGT side versus the Ising side of the duality. For instance, for the particular ribbon operator considered, on the Ising side of the duality the operator weight is $k_{\text{dual}} = 2$, whereas on the LGT side of the duality the operator weight is $k = 3$, and, consequently, the error bars are slightly worse for the Product protocol than the symmetry-aware protocols. For the loop operator considered, the difference in error between the two protocols is more pronounced, as $k_{\text{dual}} = 2$ and $k = 6$, and the error bars for the Product protocol are significantly larger for the same number of shots.

The expected behavior of the relative error for fixed samples is shown in Fig. 5.8c, where we compare the Global Dual Pairs, Dual Product and Product protocols, showing the scaling of the average relative error over

³While we used median of means estimation for easy comparison with much of the classical shadows literature, we note that, consistent with Ref. [377], empirically, we saw no real difference between median of means estimation compared to the simple sample mean.

M experimental repetitions

$$\epsilon_{\text{avg}} := \frac{1}{M} \sum_{m=1}^M \frac{|\langle O \rangle_{\text{shadow},m} - \langle O \rangle_{\text{exact}}|}{|\langle O \rangle_{\text{exact}}|}, \quad (5.32)$$

as a function of the number of samples ν (per experiment) for fixed V and fixed coupling $g = 0.5$ for the same representative loop and ribbon operator as in Fig. 5.8b. Although all protocols exhibit the expected $1/\sqrt{\nu}$ shot noise scaling with sample number ν , the Global Dual Pairs and Dual Product protocols achieve significantly lower errors than the Product protocol for any fixed ν , in agreement with our analytical predictions.

As previously noted and further analyzed below, the relative accuracy of the protocols depends on the operator weight of the observables being considered. To demonstrate how the errors of the Global Dual Pairs and Dual Product protocols depend on the operator weight of the observables being estimated in the Ising picture, we compute the relative error averaged over $M = 50$ experimental repetitions, ϵ_{avg} , as a function of system size V for two different representative loop and ribbon operators, shown in Fig. 5.8d. In particular, we show results for observables where the observable weight (on the Ising side) is constant and extensive in V . We use a fixed number of samples per experiment ($\nu = 1000$) and fixed coupling $g = 0.5$. The system sizes we can simulate numerically are too small to reliably verify the asymptotic scaling given by the error bounds of Eq. (5.27), but the qualitative behavior matches the expectations from our analytic results, namely, the cost of the Global Dual Pairs and Dual Product protocols increases with the Pauli weight of the observables on the Ising side of the duality. In the numerical examples shown, the size of systems for which we show data is limited by the classical emulation of the $\mathbb{Z}_2$ LGT input state, not by the complexity of our protocols.

Arbitrary Gauge-Invariant Observables

Protocol	Circuit Depth	Classical Cost per Sample	Number of Samples
Global Dual Pairs (this work)	$\mathcal{O}(V^2)$	$\mathcal{O}(\text{poly}(V))$	$\mathcal{O}(\frac{1}{\epsilon^2} \text{poly}(V))$
Dual Product (this work)	$\mathcal{O}(V^2)$	$\mathcal{O}(V)$	$\mathcal{O}(\frac{1}{\epsilon^2})$
Product [37]	1	$\Theta(k) \subset \mathcal{O}(V)$	$\mathcal{O}(\frac{1}{\epsilon^2} 4^k) \subset \mathcal{O}(\frac{1}{\epsilon^2} 4^V)$

Geometrically Local, Gauge-Invariant Observables

Protocol	Circuit Depth	Classical Cost per Sample	Number of Samples
Local Dual Pairs (this work)	$\mathcal{O}(L^4)$	$\mathcal{O}(\text{poly}(V))$	$\mathcal{O}(\frac{1}{\epsilon^2} \text{poly}(L))$
Dual Product (this work)	$\mathcal{O}(V^2)$	$\mathcal{O}(V)$	$\mathcal{O}(\frac{1}{\epsilon^2})$
Product [37]	1	$\Theta(k) \subset \mathcal{O}(L^2)$	$\mathcal{O}(\frac{1}{\epsilon^2} 4^k) \subset \mathcal{O}(\frac{1}{\epsilon^2} 4^L)$

Table 5.1: *Summary of costs for estimating the expectation value of a gauge-invariant observable O with $\|O\|_\infty = 1$ to an additive error ϵ . We assume that O is k -local on the LGT side of the duality or k_{dual} -local on the Ising side and, also, assume that $k_{\text{dual}} \in \Theta(1)$. Recall that $V = N_x N_y$ is the size of the (dual) lattice and, for the Local Dual Pairs protocol, L is the linear size of the geometrically local patches considered. Circuit depths for implementing the random rotations in the Dual Pairs and Dual Product protocols are based on the standard decomposition of Pauli rotations into two qubit gates and assume no parallelization of the randomizing unitaries.*




Figure 5.8: *Numerical demonstration.* **a)** Examples of a ribbon (blue) and loop (red) operator on both the $\mathbb{Z}_2$ LGT and Ising side of the duality. **b)** Estimates of ribbon (of maximal length) and loop operator (two plaquettes) expectation values on a minimal lattice of size $N_x = 3, N_y = 2$ with PBC via Global Dual Pairs, Dual Product, and Product protocols, compared to the exact solution for the ground state of the $\mathbb{Z}_2$ LGT as a function of the coupling g . Error bars represent the standard deviation of the median of means estimator over 50 experimental repetitions of $\nu = 1000$ shots each. **c)** Average relative error ϵ_{avg} (see Eq. (5.32)) versus number of samples ν for the same loop and ribbon operator as in panel (b) for the Product, Dual Product and Global Dual Pairs protocols for the ground state corresponding to coupling $g = 0.5$. **d)** Average relative error ϵ_{avg} versus volume for loop and ribbon operators of fixed weight and operators with extensive weight for the Global Dual Pairs and Dual Product protocols. In particular, for the ribbon operators the extensive weights are 2,4,6, and 8 and for the loop operators they are 2,4,4, and 6, for systems sizes $V = 4, 6, 7, 10$. The average is taken over 50 experimental repetitions of $\nu = 1000$ shots each using the ground state of the $\mathbb{Z}_2$ LGT with coupling $g = 0.5$. The sample cost of the Global Dual Pairs protocol depends both on the operator weight in the dual Ising theory and the size of the lattice, whereas the Dual Product protocol depends only on operator weight.

5.6 Protocol Comparisons

The cost of shadows protocols are characterized in terms of their asymptotic performance for generic observables, typically depending on the Pauli weight of the target observable. For standard, symmetry-ignorant schemes, observables represented by constant-length Pauli strings are independent of system size, but for arbitrary Pauli strings, such as arise for generic gauge-invariant observables, the Pauli weight grows with system size. The sampling-cost reductions achieved by our protocols stem from the fact that, for a given target observable, the effective Pauli weight can differ non-trivially between the two sides of the duality. Here, we discuss the relative performance of the various protocols in more detail for specific classes of observables.

5.6.1 Comparing symmetry-aware protocols: Dual Pairs versus Dual Product Protocols

The Global Dual Pairs protocol and the Dual Product protocol accomplish the same goal, measuring *any* gauge invariant observable. The sample cost of the Dual Product protocol for gauge invariant observables with $k_{\text{dual}} = \Theta(1)$ is significantly better, at least asymptotically, than the Global Dual Pairs protocol, i.e. constant as opposed to polynomial in system size. This asymptotic gain is largely free: the classical costs per sample are both dominated by the mapping between the dual Ising theory and the LGT, yielding a polynomial in system size cost. The circuit depths for implementing the random unitaries are also both $\mathcal{O}(V^2)$. However, the constant factors are slightly better for the Global Dual Pairs protocol since, unlike the Dual Product protocol, the unitary rotations can be parallelized.⁴ Thus, since the Dual Product protocol suffers from only a constant factor worsening of circuit depth to move from polynomial to constant sample complexity, in the asymptotic regime it is simply the better protocol when compared to Global Dual Pairs. However, this may not hold for small systems or for all observables.

When comparing the Dual Product and Local Dual Pairs protocols, and assuming geometrically local observables, the Local Dual Pairs protocol is superior. Even for geometrically local observables, the Dual Product scheme requires implementing rotations generated by ribbon operators extending to a reference plaquette, which prevents any reduction in the circuit depth of the randomizing unitaries. It is important to emphasize that these conclusions change in minor ways when we consider alternative boundary conditions; see for instance Section D.2 for a discussion of fixed boundary conditions.

⁴Parallelization is not possible for the Dual Product protocol because all rotations are generated by a ribbon operator that intersects with the reference plaquette $\square_r$ that defines the duality.




Figure 5.9: **a)** An illustration of an example loop operator of the sort considered in the discussion in Section 5.6.2. **b)** An illustration of an example (set of) ribbon operators of the sort considered in the discussion in Section 5.6.2. The way in which one picks paths to connect the end points is arbitrary, as all choices are equivalent up to the action of Gauss law operators.

5.6.2 Comparing symmetry-aware to symmetry-ignorant protocols: Dual Pairs and Dual Product versus Product Protocols

The standard approach for measuring k -local observables is the Product protocol, applied directly to the experimentally realized LGT system. Here, we compare this approach to the three symmetry-aware protocols developed in this work that make use of a virtual Ising system, without restrictions regarding the geometric locality of the observables of interest.

Compared to symmetry-aware protocols, the Product protocol offers a significantly lower circuit depth: it requires only a single layer of single-qubit unitaries, whereas, e.g., the Global Dual Pairs protocol involves $\Theta(V^2)$ high-weight Pauli rotations, each acting on $\mathcal{O}(V)$ qubits. However, this simplicity comes at the cost of exponentially increased sample complexity for the worst case observables. To see the origin of the difference in sample complexity, recall that for the Product protocol, $\text{Var}(o_j) \leq 4^k \|O\|_\infty^2$, where k is the locality of the operator in the LGT, and $\|O\|_\infty$ is evaluated on the full Hilbert space of the LGT [37] absent any gauge constraints. In contrast to the Product protocol, the Global Dual Pairs and Dual Product protocols have a sample cost that depends on the operator locality in the dual Ising theory, k_{dual} .

In the worst case (for a symmetry-ignorant protocol), gauge invariant observables have $k \gg k_{\text{dual}}$. For instance, for an observable of the form $X_i X_j$ (in the Ising picture), $k_{\text{dual}} = 2$, but the corresponding observable on the LGT side of the duality can have $k \sim V$ if the plaquettes i, j are separated by a distance of the order of the system size. Consequently, for such observables, the Product protocol requires a number of samples exponential in the system size, whereas the Global Dual Pairs protocol requires only polynomial samples in the system size and the Dual Product protocol requires constant samples. Intuitively, these improvements are due to the exponential reduction in dimension from the full to the physical Hilbert space of the LGT.

Beyond this extreme example, the difference in sample complexity between the Product protocol and the symmetry-aware protocols will depend on the particular observable(s) considered. To illustrate this, we examine how the locality of gauge-invariant observables differ on either side of the duality for several representative examples. In what follows, we let w_i for $i \in \{X, Y, Z\}$ stand for the number of σ^x , σ^y , σ^z operators, respectively, for some operator on the LGT side of the duality. Thus, $k = w_X + w_Y + w_Z$. Similarly, we let $\tilde{w}_i$ denote the same for X, Y , and Z operators, respectively, on the Ising side of the duality, such that $k_{\text{dual}} = \tilde{w}_X + \tilde{w}_Y + \tilde{w}_Z$.

Our first example is the Wilson loop illustrated in Fig. 5.9a: a closed product of magnetic variables, $\prod_{\ell \in \mathcal{P}_0} \sigma_\ell^z$, which, in any Abelian LGT, can be equivalently expressed as the product of plaquette operators enclosed by the loop. We consider two limiting cases: (i) a rectangular Wilson loop of size $N_h \times N_v$, where N_h and N_v denote the horizontal and vertical extents of the loop, respectively; and (ii) a "line" of adjacent plaquettes of size $N'_h \times 1$, as illustrated in Fig. 5.9a. This may include the worst-case scenario where $N_h N_v \propto V$, the total system volume. The corresponding Pauli weights are as follows: for case (i), the LGT-side weight is $k = w_Z = 2(N_h + N_v)$, while the Ising-side weight is $k_{\text{dual}} = \tilde{w}_Z = N_h N_v$. Thus, for these observables, the Global Dual Pairs protocol yields a sampling advantage only when $N_h N_v < 2(N_h + N_v)$, which holds when either N_h or N_v is less than or equal to 2. In contrast, for case (ii), the LGT-side weight is $k = w_Z = 2N'_h + 2$, while the Ising-side weight is $k_{\text{dual}} = \tilde{w}_Z = N'_h$. Both examples show that symmetry-aware schemes are not universally optimal for all observables.

The large advantage of the Global Dual Pairs over the Product Protocol is driven mostly by "electric" observables, i.e., those involving electric operators σ_ℓ^x in the LGT (X operators in the Ising dual). More generally, consider observables that are products of p pairs of Pauli X (and Y , which are products of electric and magnetic observables) operators on the Ising side, with $k_{\text{dual}} = \tilde{w}_X + \tilde{w}_Y = 2p$, see Fig. 5.9b for

illustration. For the case of p pairs $X_i X_j$, the mapping to an operator in the LGT is not unique—all give the same expectation value in the gauge invariant sector. In practice, on the LGT side, one would optimize by finding a collection of configurations with minimal cumulative Manhattan distance, resulting in a weight $k = \sum_{\{l_a\}_{a \leq p}} l_a$, while the Ising side weight is simply the smaller p . Considering Y operators corresponds to products of electric and magnetic operators in the LGT with an even greater Pauli weight, while the corresponding weight in the Ising dual stays the same.

This shows that, unsurprisingly, the largest *individual* gains are observed for electric gauge-invariant observables rather than magnetic ones, since the Gauss laws—responsible for reducing the physical Hilbert space—are defined in the electric basis⁵. For context, the key advantage of classical shadows is the ability to estimate *multiple* (non-commuting, i.e., with different eigenbasis) observables with the same samples. Thus, while sample complexity improvements may not be present for selected individual observables, provided we want to be able to estimate many arbitrary gauge invariant observables the upper bounds on sample complexity that we provide in Table 5.1 accurately reflect the sample complexity for estimating M arbitrary gauge-invariant observables, up to the $\log(M)/\epsilon^2$ factor in Eq. (5.26).

The classical computational cost (per sample) of inverting the shadow channel is comparable between the Product protocol, and the symmetry-aware Global Dual Pairs and Dual Product protocols. For both protocols the per sample classical cost is dominated by the $\in \mathcal{O}(\text{poly}(V))$ due to computing the mapping from the measurement outcomes to dual bit strings (i.e. the mapping $\vec{s} \rightarrow \vec{b}$). For Global Dual Pairs, one must also pre-compute f and α for a given observable, as described in Section 5.3.1. An overall comparison of the (worst case) resource costs for the Product, Dual Product, and Dual Pairs protocols is given in Table 5.1.

5.7 Discussion and Outlook

In this chapter, we introduced three variants of randomized measurement schemes tailored for systems with local (gauge) symmetries. We analyzed their use for estimating gauge-invariant observables in $\mathbb{Z}_2$ LGT, leveraging a LGT–Ising duality which enables detailed analysis and the derivation of formal guarantees. We systematically determined, for any given set of observables, a practical advantage over symmetry-agnostic alternatives, and provide explicit algorithms supported by numerical implementations [411].

⁵One could therefore argue that no advantage exists if measurements are restricted to magnetic observables only. However, in that case, any random measurement protocol would be entirely unnecessary, as all such observables commute and could simply be measured in the same fixed basis.

As with most protocols that make use of prior information—in this case, local (gauge) symmetry—a tradeoff between a reduction in sampling and an increase in circuit complexity arises. For NISQ-era devices, circuit complexity is most often quantified by CNOT depth: while the standard Product protocol [37], can be achieved in constant depth, the (Global) Dual Pairs and Dual Product protocols we propose require entangling (although still shallow) circuits. For instance, for the Global Dual Pairs protocols, the required unitary rotations for a given pairing requires a depth of $2N - 2$, where N is the number of plaquettes between any two pairs, implying a total CNOT depth of $\mathcal{O}(V^2)$, shown in Table 5.1. Similar circuit depths are required for the Dual Product protocol due to the need for ribbons going to the reference plaquette $\square_r$ for all rotations. For geometrically local observables, this can be reduced by the Local Dual Pairs protocol to $\mathcal{O}(L^4)$, while maintaining an exponential improvement (in L) in sample complexity over the standard Product protocol. In the near-to-intermediate term whether or not the polynomial increase in circuit depth associated with tailoring shadow protocols is worth the exponential gains in sample complexity will depend on the platform and the ability to implement circuits of sufficient depth without incurring large errors due to imperfect gates. However, in fault-tolerant quantum processors—assuming that the timescales for implementing logical gates are comparable to or better than those for measurement and readout—the potential exponential reduction in sample complexity constitutes an unalloyed advantage when simulating $\mathbb{Z}_2$ LGT directly.

Nonetheless, in practice, the same Ising-LGT duality that allows us to rigorously analyze the performance of the shadows protocols could be used to directly simulate the dual Ising model [148] instead of the LGT, and avoid the large circuit depth overheads required. In this case, the sample complexity for our shadows schemes would be identical, but the circuit depths would be significantly smaller. Our work, thus, should not be primarily understood as an explicit simulation protocol for $\mathbb{Z}_2$ LGT, but rather the key components are (a) a rigorous analysis of a constructive shadows protocol for an explicit model, demonstrating the circuit depth-sample complexity trade-off intrinsic to symmetry-preserving shadows [388] that is expected for LGTs; and (b) a blueprint for analyzing more general LGTs, including non-Abelian LGTs relevant to high energy and nuclear physics such as $SU(2)$ and $SU(3)$ LGTs even if dualities may be hard to find or do not exist.

Indeed, while the central insight of our work relies on the Ising-LGT duality specific to $\mathbb{Z}_2$ (and by extension $\mathbb{Z}_N$, and in the appropriate limit, $U(1)$ [404]) LGTs, many elements of our protocols generalize naturally. All steps of the protocols considered can be cast fully in the language of the LGT, rather than the dual Ising model. In this sense, the $\mathbb{Z}_2$ LGT case serves as a useful starting point, as the presence of a duality

allows for a rigorous analysis of the channel and shadow protocol. Building on this foundation, one can move beyond this simple setting by using these insights as a blueprint for more general LGTs—particularly to *non-Abelian* LGTs or models whose local Hilbert spaces do not resemble qubits (including LGTs based on compact groups)—where the corresponding gauge-invariant randomizing operations follow directly. However, what channels these randomizing operations represent and if (and how efficiently) they can be inverted remains a major challenge. We presently are not aware of a simple analogous approach for non-Abelian LGTs for which dualities, so far, are challenging to construct [412–418]. This raises the question of whether similar sampling-efficient protocols can be devised for arbitrary LGTs, and whether their existence implies an underlying Ising-like duality, or their absence indicates its nonexistence. An interesting development in this context is the recent demonstration of charge-singlet measurements in quantum simulations of $SU(N_c)$ LGT [172, 419]. In the absence of a duality, one could study the corresponding channels numerically in small-scale instances, but extracting scaling from this is a challenge. Overcoming these hurdles is a primary goal for future work.

A variety of alternative approaches to symmetry-respecting shadows have been recently proposed. One promising direction [420] employs classical shadows to learn unknown symmetries of quantum systems, potentially offering a “bootstrap” method for constructing symmetry-respecting classical shadows. Ref. [388] introduces a general framework to construct symmetry-respecting shadows under the assumption of oracle access to a unitary that block diagonalizes the system in the irrep basis of the associated symmetry. Another established approach uses symmetric approximate unitary k -designs for LGTs [386], which allow simple channel inversion once the symmetry-sector dimensions are known but require deep random circuits. Recent results on log-depth approximate unitary k -designs could extend to random circuits with symmetries [421–423]. Applying such constructions to systems with gauge symmetries is likely to be highly favorable when estimating properties of the state beyond k -local observables [424, 425]. Finally, there exist other symmetry-ignorant shadow tomography schemes that occupy an intermediate regime between the Product and Dual Pairs protocols in terms of both sampling cost and circuit complexity. Notable examples include the locally entangled schemes of Ref. [426] and the shallow shadow schemes of Refs. [368, 374]. The potentially exponential advantages observed in our setting stem from explicitly exploiting gauge symmetry and, thus, will extend to these settings.

Significant progress has been made in understanding the impact of noise on standard schemes [427–430], and the development of derandomized variants [384]. These ideas are likely extendable to the symmetry-

respecting circuits. However, a particularly interesting complication arises: noise can break the very symmetries that these protocols are designed to exploit. This symmetry breaking can lead to biased estimators and necessitates a fundamentally different error analysis. Relatedly, the impact of small gauge violations in the input state on the performance of our protocols remains an interesting open question. On the other hand, symmetry-aware measurement schemes can enable powerful error mitigation schemes, as they allow one to effectively project the measured state into a symmetric subspace [431].

Chapter 6

MEASURING THE ENTANGLEMENT STRUCTURE OF NON-GAUSSIAN STATES**6.1 Introduction**

Recent advances in quantum simulation have opened the door to understanding quantum many-body systems at unprecedented levels [59, 82, 105, 432–434]. These advances have come with a bevy of methods for characterizing quantum states from experimental measurements, ranging from full quantum state tomography [345, 352], over randomized measurement protocols [37, 361, 382] and (entanglement) Hamiltonian learning [171, 435–439], to many-body interference [147, 440–443] or higher-order correlation functions [444, 445]. An important goal of such methods is to reveal the structure of complex many-body states, including entanglement for characterizing quantum phases in equilibrium [386, 446, 447] or chaos and thermalization [302, 405, 448, 449].

Common approaches to measure entanglement structure incur costs in sample size or computational resources that grow exponentially with the (sub)system’s size, which makes them challenging to scale up. Moreover, these protocols are difficult to realize in experiments without full local or universal control, examples include analog quantum simulators based on ultra-cold atomic and molecular systems [432, 450–452], or high-energy and nuclear physics experiments [59, 82, 295, 434]. In such experiments, instead, one typically extracts information by measuring few-body correlation functions.

Motivated by these restrictions, we address the following question: *How much physics of entanglement can be captured if access is limited to lower order correlations?* Specifically, as illustrated in Fig. 6.1, we show how to utilize exact non-perturbative correlation functions accessible in experiments for constructing approximations to a quantum state based on an expansion of the many-body Wigner functional around Gaussian states. We expect our parameterization to capture accurate descriptions during early times, where higher correlations are small, and late times, particularly when the system is near thermal equilibrium.

As an example, we illustrate our approach with numerically simulated thermalization dynamics of a weakly-coupled extended Fermi-Hubbard chain after a quantum quench. Here, we focus on the reconstruction of a non-equilibrium state’s entanglement spectrum and analyze its level spacing distribution, which serves




Figure 6.1: *Non-Gaussian state reconstruction.* **a)** Few-body correlation measurements (e.g. four-point correlations $C^{(4)}$) on a subsystem A of a fermionic quantum simulator are used to reconstruct the reduced state ρ_A . Besides single-site resolved imaging, this requires pair-wise tunneling operations $\mathcal{U}_{kl}$ between sites k and l . After classical post-processing, the protocol gives access to the entanglement spectrum ε_i . **b)** The state reconstruction is based on an expansion around Gaussian states ρ_g , which form a subset of integrable states in the Hilbert space $\mathcal{H}$. To probe the onset of quantum chaos, we go to non-Gaussian states $\tilde{\rho}$ by including higher-order correlations. **c)** Four-point correlations lead to non-Gaussian matrix elements $\langle \mathbf{m} | \rho_A | \mathbf{n} \rangle$ with Hamming distance $d_H(\mathbf{m}, \mathbf{n}) = 0, 2, 4$ (top, middle, bottom).

as an indicator of quantum chaos, a prerequisite for thermalization. Our work demonstrates how going beyond Gaussian state descriptions – which are known to fail even qualitatively already for early times – can quantitatively capture aspects of the chaotic evolution of the system’s entanglement structure. Finally, we briefly discuss methods for probing such phenomena in analog fermion quantum simulators, based on

existing and forthcoming experimental capabilities [453–458].

6.2 Non-Gaussian States

We discuss an approach for representing generic quantum states by expanding their Wigner functionals in terms of correlation functions. We consider a system of N fermionic modes described by creation/annihilation operators $c_i^\dagger/c_i$ where $\{c_i^\dagger, c_j\} = \delta_{ij}$ ($i, j \in \{0, 1, \dots, N-1\}$), and focus on connected two- and four-point functions,

$$C_{ij}^{(2)} \equiv \text{Tr}[\rho c_i^\dagger c_j] \quad (6.1a)$$

$$C_{ijkl}^{(4)} \equiv \text{Tr}[\rho c_i^\dagger c_j^\dagger c_k c_l] - C_{il}^{(2)} C_{jk}^{(2)} + C_{ik}^{(2)} C_{jl}^{(2)}. \quad (6.1b)$$

For simplicity, we restrict ourselves to a U(1) symmetric state, e.g. a state with definite total particle number $N_F = \sum_i c_i^\dagger c_i$. Correlations of this form are accessible in current and near-term neutral-atom experiments, illustrated in Fig. 6.1(a) and discussed below ¹.

For the case of a subsystem defined by $N_A < N$ modes, the corresponding reduced state $\rho_A = \text{Tr}_{\bar{A}}(\rho)$ is equivalently described by a (bipartite) entanglement Hamiltonian (EH) $H_A \equiv -\log[\rho_A]$ [446], with entanglement spectrum $\{\varepsilon_i\}$. For Gaussian states, the EH is simply given by the covariance matrix, $C^{(2)}$, as $H_A = -\sum_{ij} \log[C^{(2)}(1 - C^{(2)})^{-1}] c_i^\dagger c_j + \text{const}$ [459]. This class of states falls short of capturing even the qualitative indicators of thermalization expected from the entanglement dynamics of the system, as we will argue below. Specifically, Gaussian states fall into a small “corner” of the many-body Hilbert space $\mathcal{H}$ that cannot be chaotic, whereas even small deviations away from that region will show signatures of quantum chaos consistent with random matrix theory [460, 461], see Fig. 6.1(b).

Our approach is to write the state as a sum of a Gaussian ρ_g , entirely determined by $C^{(2)}$, and non-Gaussian part $\delta\rho$,

$$\rho \equiv \rho_g + \delta\rho \quad (6.2)$$

thereby reducing the problem to finding an expression for $\delta\rho$. To proceed, we utilize the formalism of fermionic coherent states [462] where ρ may be expressed as

$$\rho = \int d^2\eta d^2\alpha \ e^{W(\eta)} e^{\sum_i \alpha_i \eta_i^* - \eta_i \alpha_i^*} |\alpha\rangle \langle -\alpha| \quad (6.3)$$

¹It is not essential to deal with fermions, extending our approach to bosonic models, and to higher order correlations, is straightforward.

with α, η being tuples of anti-commuting Grassmann numbers, $\alpha \equiv \{\alpha_i, \alpha_i^* \dots\}_N$, $|\alpha\rangle$ are the corresponding coherent states, and $W(\eta)$ is the Wigner functional. We assume that $W(\eta)$ takes the form

$$W(\eta) = \sum_{i,j} C_{ij}^{(2)} \eta_j^* \eta_i + \frac{1}{4} \sum_{i,j,k,l} C_{ijkl}^{(4)} \eta_l^* \eta_k^* \eta_j \eta_i \quad (6.4)$$

with $i, j, k, l \in \{0, \dots, N_A - 1\}$, which is justified for small $|C^{(4)}| \ll |C^{(2)}|$. We note that $C^{(2)}$ and $C^{(4)}$ are not approximated and should be obtained from experiment, or computed non-perturbatively². Working in the mode basis $|\mathbf{n}\rangle$ where ρ_g is diagonal, we derive explicit expressions for the matrix elements $\delta\rho_{mn} = \langle \mathbf{m} | \delta\rho | \mathbf{n} \rangle$,

$$\frac{\delta\rho_{mn}}{\rho_{g,nn}} = \begin{cases} \frac{1}{2} \sum_{i,j} (-1)^{n_i+n_j} \frac{\tilde{C}_{ijji}^{(4)}}{f_i^n f_j^n} & \text{if } d_H(\mathbf{m}, \mathbf{n}) = 0 \\ \sum_{i,j,k} (-1)^{n_i+\varphi(j,k;\{i\},\mathbf{n})} \frac{\tilde{C}_{ijik}^{(4)}}{f_i^n f_j^n f_k^n} \Gamma_{m_j n_k}^{n_j m_k} & \text{if } d_H(\mathbf{m}, \mathbf{n}) = 2 \\ \frac{1}{4} \sum_{i,j,k,l} (-1)^{\varphi(i,j;\{k,l\},\mathbf{n})+\varphi(k,l;\{i,j\},\mathbf{n})} \text{sgn}(i-j) \text{sgn}(k-l) \frac{\tilde{C}_{ijkl}^{(4)}}{f_i^n f_j^n f_k^n f_l^n} \Gamma_{n_i n_j m_k m_l}^{m_i m_j n_k n_l} & \text{if } d_H(\mathbf{m}, \mathbf{n}) = 4 \end{cases} \quad (6.5)$$

where $\rho_{g,nn} = \prod_p f_p^n$ are the diagonal matrix elements of the Gaussian state and $f_p^n \equiv n_p g_p + (1 - n_p)(1 - g_p)$, g_p is the p^{th} eigenvalue of $C^{(2)}$, $\tilde{C}^{(4)}$ is the four-point correlation in the basis where $C^{(2)}$ is diagonal, and $d_H(\mathbf{m}, \mathbf{n})$ is the hamming distance between the $\mathbf{m}$ and $\mathbf{n}$ occupation bitstrings. The symbol $\Gamma_{i'_1, i'_2, \dots, i'_k}^{i_1, i_2, \dots, i_k} \equiv \prod_a^k \delta_{i_a, 0} \prod_a^{k'} \delta_{i'_a, 1}$ is defined so as to highlight the matrix structure of the $\delta\rho$ and is pictorially represented in Fig. 6.1(c). The phases $\varphi(i, j; \{k, l\}, \mathbf{n}) \equiv \sum_{s=\min(i,j)+1}^{\max(i,j)-1} n_s$ count the occupation between modes i and j excluding the modes k and l . The parameterization exactly reproduces all correlations up to 1-body (2-point) and 2-body (4-point) and will be hereafter referred to as $\tilde{\rho}$. We demonstrate its application using an example of a fermionic system.

6.3 Derivation of Parameterization

We now provide a detailed derivation of the parameterization introduced above, filling in the steps between Eq. (6.4) and Eq. (6.5), and illustrating a general approach that can be extended to include higher-order correlation functions. First, we will present some basic identities and clarify the conventions used. Then, we will derive Eq. (6.5). Additionally, we introduce a diagrammatic formulation that aids intuition and can be

²Our ansatz is unrelated to time-dependent perturbation theory.

straightforwardly extended beyond the two-body correlations demonstrated here. The fermionic coherent state formalism for the Wigner functional follows the conventions of Cahill and Glauber [462]. Our derivation is built on fermionic states consisting of modes $i = 0, \dots, N - 1$, with creation and annihilation operators obeying the usual anticommutation relation $\{c_i, c_j^\dagger\} = \delta_{ij}$. Coherent states, denoted as $|\gamma\rangle \equiv D(\gamma)|0\rangle$, are defined by the action of the displacement operator $D(\gamma) \equiv \exp\left(\sum_i c_i^\dagger \gamma_i - \gamma_i^* c_i\right)$. Here, $c_n|0\rangle = 0$ defines the Fock vacuum, and $\gamma \equiv (\gamma_0, \gamma_0^*, \dots)$ represents a tuple of $2N$ Grassmann variables that anticommute. Grassmann integration uses the convention that $\int d^2\gamma \equiv \int \prod_i d\gamma_i d\gamma_i^*$, and $\int d\gamma_i \gamma_j = \int d\gamma_i^* \gamma_j^* = \delta_{ij}$. We will make use of the completeness relation $\mathbb{1} = \int d^2\gamma \gamma \gamma^* D(\gamma)$ and other standard relations found in [462]. Products of fermionic or Grassmann variables will contain an arrow $\leftarrow$ or $\rightarrow$ to denote decreasing or increasing mode ordering, unless the ordering does not matter (such as for products of bi-linear terms). Derivatives, including the product rule, for Grassmann variables follow [462].

6.3.1 Evaluation of non-Gaussian Wigner functional

Inserting Eq. (6.3) into Eq. (6.4) yields

$$\begin{aligned} \rho_A = \int d^2\eta d^2\alpha \ e^{\sum_{ij} C_{ij}^{(2)} \eta_j^* \eta_i} & \left(1 + \frac{1}{4} \sum_{i,j,k,l} C_{ijkl}^{(4)} \eta_l^* \eta_k^* \eta_j \eta_i \right) e^{\sum_i \alpha_i \eta_i^* - \eta_i \alpha_i^*} |\alpha\rangle \langle -\alpha| \\ & + \mathcal{O}(C^{(6)}) \end{aligned} \quad (6.6)$$

where only the leading order dependency in $C^{(4)}$ has been kept. This expression can be separated into two terms, $\rho_A \equiv \rho_g + \delta\rho$

$$\rho_g = \int d^2\eta d^2\alpha \ e^{\sum_{ij} C_{ij}^{(2)} \eta_j^* \eta_i} e^{\sum_i \alpha_i \eta_i^* - \eta_i \alpha_i^*} |\alpha\rangle \langle -\alpha|, \quad (6.7)$$

$$\begin{aligned} \delta\rho = \frac{1}{4} \sum_{i,j,k,l} C_{ijkl}^{(4)} \int d^2\eta d^2\alpha \ \eta_l^* \eta_k^* \eta_j \eta_i & e^{\sum_{ij} C_{ij}^{(2)} \eta_j^* \eta_i} \\ & e^{\sum_i \alpha_i \eta_i^* - \eta_i \alpha_i^*} |\alpha\rangle \langle -\alpha| \end{aligned} \quad (6.8)$$

where the first line is identified with the Gaussian part of ρ_A ³ and the second is the perturbation. We are interested in the matrix elements in the Fock basis, $\rho_{g,mn} \equiv \langle \mathbf{m} | \rho_g | \mathbf{n} \rangle$, $\delta \rho_{nm} \equiv \langle \mathbf{m} | \delta \rho | \mathbf{n} \rangle$, where we assume that $|\mathbf{n}\rangle \equiv \overleftarrow{\prod}_i (c_i^\dagger)^{n_i} |0\rangle$ where $\mathbf{n} \equiv (n_0, \dots)$ with $n_i \in 0, 1$, and the left-pointing arrow means that the modes are in descending order. To proceed, we change to the Grassmann coordinates $\xi_i = \sum_j U_{ij} \eta_j$, where U is the matrix that diagonalizes $C^{(2)}$. With this change the integrals over ξ and α factorize and the Gaussian part becomes

$$\begin{aligned} \rho_{g,mn} &= \int d^2 \xi d^2 \alpha e^{\sum_i g_i \xi_i^* \xi_i + \alpha_i \xi_i^* - \xi_i \alpha_i^*} \langle \mathbf{m} | \alpha \rangle \langle -\alpha | \mathbf{n} \rangle \\ &= \det[-C^{(2)}] \int d^2 \alpha e^{\sum_i \frac{\alpha_i^* \alpha_i}{g_i}} \langle \mathbf{m} | \alpha \rangle \langle -\alpha | \mathbf{n} \rangle \end{aligned} \quad (6.9)$$

where $\det[-C^{(2)}] \equiv \prod_i (-g_i^{(2)})$ and $g_i, i = 0, N-1$ are the eigenvalues of the correlation matrix $C^{(2)}$ and

$$\langle \mathbf{m} | \alpha \rangle = \prod_a^{\rightarrow} (\alpha_a)^{m_a} \exp\left\{-\frac{\alpha_a^* \alpha_a}{2}\right\} \quad (6.10)$$

$$\langle -\alpha | \mathbf{n} \rangle = \prod_a^{\leftarrow} \exp\left\{-\frac{\alpha_a^* \alpha_a}{2}\right\} (-\alpha_a^*)^{n_a}. \quad (6.11)$$

The Gaussian matrix elements become

$$\begin{aligned} \rho_{g,mn} &= \det[-C^{(2)}] \int d^2 \alpha e^{\sum_i [\frac{1}{g_i} - 1] \alpha_i^* \alpha_i} \\ &\quad \times \prod_a^{\rightarrow} (\alpha_a)^{m_a} \prod_a^{\leftarrow} (-\alpha_a^*)^{n_a} \\ &= \delta_{mn} \prod_i [n_i g_i + (1 - n_i)(1 - g_i)], \end{aligned} \quad (6.12)$$

where $\delta_{mn} \equiv \prod_i \delta_{m_i n_i}$. In the second equality we made use of the following identity,

$$\int d^2 \alpha_i e^{\lambda \alpha_i^* \alpha_i} \alpha_i^{m_i} (-\alpha_i^*)^{n_i} = \delta_{m_i n_i} [\lambda - n_i(1 + \lambda)]. \quad (6.13)$$

After computing the Gaussian component, we proceed deriving the quartic contribution, $\delta \rho$ [Eq. (6.8)], which incorporates two-body correlations. To simplify the derivation, we note first that a specific term in the integrand of Eq. (6.8) can be written as

$$\eta_l^* \eta_k^* \eta_j \eta_i e^{\sum_i \alpha_i \eta_i^* - \eta_i \alpha_i^*} = \frac{\partial}{\partial \alpha_l} \frac{\partial}{\partial \alpha_k} \frac{\partial}{\partial \alpha_j^*} \frac{\partial}{\partial \alpha_i^*} e^{\sum_i \alpha_i \eta_i^* - \eta_i \alpha_i^*} \quad (6.14)$$

³Note that one can bring the Gaussian part into an alternative form [459], noting that in the diagonal basis $\rho_g = \det(C^{(2)}) \exp\{-\sum_{ij} \log[C^{(2)}(1 - C^{(2)})^{-1} c_i^\dagger c_j]\} = \prod_i g_i [\frac{1-g_i}{g_i}]^{\hat{n}_i}$, where $\hat{n}_i$ are the number operators in the $C^{(2)}$ eigenbasis. Taking the matrix elements $\langle \mathbf{m} | \rho_g | \mathbf{n} \rangle$ yields the expression we work with in this chapter.

Integrating Eq. (6.8) by parts in the α coordinates moves the derivative terms to act on $\langle \mathbf{m} | \alpha \rangle \langle -\alpha | \mathbf{n} \rangle$. As with the Gaussian case, we switch to the ξ basis where $C^{(2)}$ is diagonal, denoting the four-point function in this basis as $\tilde{C}^{(4)}$ where $\tilde{C}_{ijkl}^{(4)} \equiv \sum_{abcd} U_{ia} U_{jb} C_{abcd}^{(4)} U_{ck}^\dagger U_{dl}^\dagger$. The ξ integrals may be evaluated as before, and yet another integration by parts, this time in the α variables, gives

$$\delta\rho_{mn} = \frac{1}{4} \det[-C^{(2)}] \sum_{ijkl} \tilde{C}_{ijkl}^{(4)} \int d^2\alpha \left[\frac{\partial}{\partial \alpha_l} \frac{\partial}{\partial \alpha_k} \frac{\partial}{\partial \alpha_j^*} \frac{\partial}{\partial \alpha_i^*} \prod_p e^{\frac{\alpha_p^* \alpha_p}{g_p}} \right] \langle \mathbf{m} | \alpha \rangle \langle -\alpha | \mathbf{n} \rangle \quad (6.15)$$

The derivative terms in Eq. (6.15) are evaluated using the product rule for Grassmann numbers,

$$\begin{aligned} \frac{\partial}{\partial \alpha_l} \frac{\partial}{\partial \alpha_k} \frac{\partial}{\partial \alpha_j^*} \frac{\partial}{\partial \alpha_i^*} \prod_p e^{\frac{\alpha_p^* \alpha_p}{g_p}} &= \frac{\delta_{il} \delta_{jk} - \delta_{ik} \delta_{jl}}{g_i g_j} \prod_{p \neq i, j} e^{\frac{\alpha_p^* \alpha_p}{g_p}} \\ &+ \frac{\delta_{jk} \alpha_i \alpha_l^* - \delta_{ik} \alpha_j \alpha_l^*}{g_i g_j g_l} \prod_{p \neq i, j, l} e^{\frac{\alpha_p^* \alpha_p}{g_p}} \\ &+ \frac{\delta_{il} \alpha_j \alpha_k^* - \delta_{jl} \alpha_i \alpha_k^*}{g_i g_j g_k} \prod_{p \neq i, j, k} e^{\frac{\alpha_p^* \alpha_p}{g_p}} \\ &+ \frac{\alpha_i \alpha_j \alpha_k^* \alpha_l^*}{g_i g_j g_k g_l} \prod_{p \neq i, j, k, l} e^{\frac{\alpha_p^* \alpha_p}{g_p}} \end{aligned} \quad (6.16)$$

Here, in second line we implicitly assume that $i \neq j \neq l$, and in the third $i \neq j \neq k$, and in the fourth $i \neq j \neq k \neq l$, otherwise these terms vanish. Inserting Eq. (6.16) into Eq. (6.15) gives

$$\delta\rho_{mn} \equiv 2 \sum_{i, j} I_1^{(i, j)} + 4 \sum_{i, j, k} I_2^{(i, j, k)} + \sum_{i, j, k, l} I_3^{(i, j, k, l)}, \quad (6.17)$$

where we abbreviated

$$\begin{aligned} I_1^{(i, j)} &= \frac{1}{4 g_i g_j} \det[-C^{(2)}] \tilde{C}_{ijji}^{(4)} \int d^2\alpha \prod_{p \neq i, j} e^{\frac{\alpha_p^* \alpha_p}{g_p}} \\ &\times \prod_q e^{-\alpha_q^* \alpha_q} \prod_a (\alpha_a)^{m_a} \prod_b (-\alpha_b^*)^{n_b} \end{aligned} \quad (6.18)$$

When $\mathbf{m} \neq \mathbf{n}$, there will be an unpaired Grassmann number and the entire integral will vanish, consequently Eq. (6.18) is non-zero only if $\mathbf{m} = \mathbf{n}$. To compute the remaining integrals, one combines the Grassmann variables into bilinear form $\sim \alpha_i^* \alpha_i$ and, using Eq. (6.13), one finds

$$I_1^{(i, j)} = \frac{1}{4} \tilde{C}_{ijji}^{(4)} \delta_{mn} (-1)^{n_i + n_j}$$

$$\times \prod_{p \neq i,j} [n_p g_p + (1 - n_p)(1 - g_p)]. \quad (6.19)$$

The second term of Eq. (6.17) can be written as

$$\begin{aligned} I_2^{(i,j,k)} &= \frac{1}{4g_i g_j g_k} \det[-C^{(2)}] \tilde{C}_{ijk}^{(4)} \int d^2 \alpha \alpha_j \alpha_k^* \\ &\times \prod_{p \neq i,j,k} e^{\frac{\alpha_p^* \alpha_p}{g_p}} \prod_q e^{-\alpha_q^* \alpha_q} \prod_a^{\rightarrow} (\alpha_a)^{m_a} \prod_b^{\leftarrow} (-\alpha_b^*)^{n_b}. \end{aligned} \quad (6.20)$$

The Grassmann integrals only evaluate to a non-zero result if they appear once and the integration is over bilinears of the same mode. This is only the case if $n_k = 0$, $m_j = 0$, $n_j = 1$, and $m_k = 1$, as well as $n_a = m_a$ for $a \neq i, j$. To represent this condition we introduce the abbreviation $\Gamma_{m_j n_k}^{n_j m_k}$,

$$\Gamma_{i'_1, i'_2, \dots, i'_{k'}}^{i_1, i_2, \dots, i_k} \equiv \prod_a^k \delta_{i_a, 0} \prod_a^{k'} \delta_{i'_a, 1}, \quad (6.21)$$

where an index placed in superscript indicates that the corresponding mode should be unoccupied $n_i = 0$, and an index in subscript means that the mode is occupied $n_i = 1$. A mode i can only appear in either super- or subscript. We further use the following identity,

$$\alpha_j \alpha_k^* \prod_a^{\rightarrow} (\alpha_a)^{m_a} \prod_b^{\leftarrow} (-\alpha_b^*)^{n_b} = (-1)^{\varphi(j,k;\mathbf{n})} \alpha_k^* \alpha_k \alpha_j^* \alpha_j \prod_{a \neq k}^{\rightarrow} (\alpha_a)^{m_a} \prod_{b \neq j}^{\leftarrow} (-\alpha_b^*)^{n_b} \quad (6.22)$$

where $\varphi(j, k; \{i\}, \mathbf{n}) = \sum_{j < s < k} n_s$ if $j < k$ and $\varphi(j, k; \{i\}, \mathbf{n}) = \sum_{k < s < j} n_s$ if $j > k$, counting the occupation between modes j and k , ignoring the occupation of mode i . Inserting this identity into Eq. (6.20) yields

$$\begin{aligned} I_2^{(i,j,k)} &= \frac{1}{4} \tilde{C}_{ijk}^{(4)} (-1)^{\varphi(j,k;\{i\},\mathbf{n})} (-1)^{n_i} \Gamma_{n_j m_k}^{m_j n_k} \delta_{n_i m_i} \\ &\times \prod_{p \neq i,j,k} \delta_{n_p m_p} [n_p g_p + (1 - n_p)(1 - g_p)]. \end{aligned} \quad (6.23)$$

The final contribution is given by

$$\begin{aligned} I_3^{(i,j,k,l)} &= \frac{1}{4g_i g_j g_k g_l} \det[-C^{(2)}] \tilde{C}_{ijkl}^{(4)} \int d^2 \alpha \alpha_i \alpha_j \alpha_k^* \alpha_l^* \\ &\times \prod_{p \neq i,j,k,l} e^{\frac{\alpha_p^* \alpha_p}{g_p}} \prod_q e^{-\alpha_q^* \alpha_q} \prod_a^{\rightarrow} (\alpha_a)^{m_a} \prod_b^{\leftarrow} (-\alpha_b^*)^{n_b} \end{aligned} \quad (6.24)$$

Here again Grassmann integrals evaluate to a non-zero result only if every Grassmann variable appears exactly linearly in the integrand. Consequently, $n_k, n_l, m_i, m_j = 0$, $m_k, m_l, n_i, n_j = 1$, and $n_a = m_a$ if $a \neq i, j, k, l$. To reorganize the integrand into Grassmann bilinears, we use the following identity

$$\begin{aligned}
& \Gamma_{n_i n_j m_k m_l}^{m_i m_j n_k n_l} \alpha_i \alpha_j \alpha_k^* \alpha_l^* \prod_{a \neq i, j}^{\rightarrow} (\alpha_a)^{m_a} \prod_{b \neq k, l}^{\leftarrow} (-\alpha_b^*)^{n_b} \\
&= \Gamma_{n_i n_j m_k m_l}^{m_i m_j n_k n_l} (-1)^{\varphi(i, j; \{k, l\}, \mathbf{n}) + \varphi(k, l; \{i, j\}, \mathbf{n})} \\
&\quad \times \text{sgn}(i - j) \text{sgn}(k - l) (\alpha_k^* \alpha_k) (\alpha_l^* \alpha_l) \\
&\quad \times (\alpha_i^* \alpha_i) (\alpha_j^* \alpha_j) \prod_{a \neq i, j, k, l}^{\rightarrow} (\alpha_a)^{m_a} \prod_{b \neq i, j, k, l}^{\leftarrow} (-\alpha_b^*)^{n_b}, \tag{6.25}
\end{aligned}$$

where $\text{sgn}(x)$ gives the sign of x and $\varphi(i, j; \{k, l\}, \mathbf{n})$ is defined below Eq. (6.22). Finally,

$$\begin{aligned}
I_3^{(i, j, k, l)} &= \frac{1}{4} \tilde{C}_{ijkl}^{(4)} (-1)^{\varphi(i, j; \{k, l\}, \mathbf{n}) + \varphi(k, l; \{i, j\}, \mathbf{n})} \\
&\quad \times \text{sgn}(i - j) \text{sgn}(k - l) \Gamma_{n_i n_j m_k m_l}^{m_i m_j n_k n_l} \\
&\quad \times \prod_{p \neq i, j, k, l} \delta_{n_p m_p} [n_p g_p + (1 - n_p)(1 - g_p)] \tag{6.26}
\end{aligned}$$

Combining Eqs. (6.17-6.19) yields our final result, Eq. (6.5). The derivation can be generalized to include higher order correlation functions by keeping additional terms in Eq. (6.8) and using the Grassmann identities that we presented. Such derivation is significantly simplified by a diagrammatic representation, akin to Feynman diagram expansions, that we will illustrate below.

6.3.2 Diagrammatic Rules

The matrix elements of the Gaussian part $\rho_{g, mn}$, as well as all higher-order terms in $\delta\rho_{mn}$ are obtained by an expansion of the fermionic Wigner functional. This involves manipulations that contain multiple Grassmann-valued derivatives and integrals where attention must be paid to their anti-commutative nature. We now outline a diagrammatic approach that simplifies these computations, providing intuition and allowing us to extend the approach to higher-orders without losing the rigor of the algebra that accounts for the anti-commuting nature of the variables. We will draw on analogies from Feynman diagram evaluation in Quantum Field Theory, but note that the language introduced is to build intuition only and should not be taken literally.




Figure 6.2: *Elemental Diagrams*. **a)** Propagator diagrams carrying index p . **b)** Source terms, represented by diamonds carrying index p with the corresponding occupations number m_p, n_p . **c)** k -point vertex carrying indices $i_1, \dots, i_k$.

The elemental components of our diagrammatic derivation are summarized in Fig. 6.2, showing the “propagator” in (a), “source term” in (b), and “(4-point)-vertex” in (c). We will now discuss how these elements can be used to construct the matrix elements of ρ . The lowest-order Gaussian part is given by

$$\rho_{mn}^g \equiv \prod_p \delta_{m_p n_p} f_p^n, \quad (6.27)$$

where $f_p^n \equiv n_p g_p + (1 - n_p)(1 - g_p)$. It will play a role akin to a generating functional for all higher order terms based on an expansion in terms of the k -ality of (potentially multiple) vertices.

We refer to the subsequent terms, involving $\tilde{C}^{(4)}, \tilde{C}^{(6)} \dots$ and their products, as “interaction diagrams”. Their construction is illustrated pictographically in Fig. 6.3 for the leading-order case involving only $\tilde{C}^{(4)}$. Here, the Gaussian terms and $\tilde{C}^{(4)}$ serve as the starting points for these interaction diagrams. Each vertex contributes a factor $\tilde{C}_{i_1, \dots, i_k}^{(k)} / (k/2)!^2$. To construct them, one must first distinguish between different cases of the 4-point function $\tilde{C}^{(4)}$: either two pairs of indices, one pair of indices are equal, or all indices are different. For example, the left panel of Fig. 6.3 shows the derivation of the diagram I_1 involving $\tilde{C}_{ijji}^{(4)}$. In this process, the “open” indices $i \neq j$ are removed from the product over p derived from the Gaussian terms. For each open index, one writes down a propagator, as well as one n -type (represented by an incoming arrow on the diamond symbol) and one m -type (represented by an outgoing arrow) source term for every open index. One begins by first connecting the propagators to the vertex in all possible ways so that the indices match. Some propagators may form a closed loop connecting repeated indices of $\tilde{C}^{(4)}$, in this case the diagram incurs no




Figure 6.3: *Construction of I_1 and I_2 Interaction Diagrams.* Starting from the Gaussian diagram as the generating functional (top), 4-point vertices are introduced. The diagram for I_1 is constructed on the left and the diagram for I_2 is constructed on the right following the rules discussed in this section.

additional factor and the corresponding index is simply summed over. Otherwise, a propagator's end point must be connected to an n - or m -type source. Finally, one directly connects m - and n -type source terms if their indices match, in this case they contribute a factor of $(-1)^{n_a} \delta_{m_a n_a}$. Some source terms may be left unmatched, see for instance the bottom right panel of Fig. 6.3 where m_k and n_j with $k \neq j$ are unmatched. In this case, they result in a factor of $\Gamma_{m_k n_j}^{n_k m_j}$. The left panel of Fig. 6.3 illustrates the unique way that $\tilde{C}_{ijji}^{(4)}$ can be connected, while the right panel of the same figure depicts the diagram for the term I_2 which involves




Figure 6.4: *Diagrams contributing at $O(\tilde{C}^{(4)})$.* **a)** I_1 : Propagators connect repeated indices of the vertex, all sources are paired up. **b)** I_2 : One propagator connects one set of repeated indices of the vertex, the other two end in sources. Some sources can be paired, others remain unconnected. **c)** I_3 : All propagators connect to sources.

the “pair-density” correlation function $\tilde{C}_{ijk}^{(4)}$; the Gaussian contribution is not explicitly shown.

Additionally, because of the anti-commutative nature, we must account for phases that originate from un-paired sources (with open arrows). We denote the ordered set of in-going indices as J_{in} , the ordered set of out-going indices as J_{out} , and $J = J_{\text{out}} \cup J_{\text{in}}$ as their union, where the ordering is the same as how the

indices appear in the corresponding k -point vertex $\tilde{C}_{i_1, \dots, i_k}^{(k)}$. For example, for the diagram depicting I_2 , shown in the right panel of Fig. 6.3, $J_{\text{in}} = \{k\}$ and $J_{\text{out}} = \{j\}$ while for I_3 $J_{\text{in}} = \{k, l\}$ and $J_{\text{out}} = \{i, j\}$. We then define $\tilde{\varphi} = \sum_{f, g \in J} \varphi(f, g; J/\{f, g\}, \mathbf{n})$ which sums the occupations of $\mathbf{n}$ between all pairs (f, g) of external indices regardless of whether they are in-going or out-going, ignoring the occupations of the external indices not in the pair denoted by $J/\{f, g\}$. The diagram is multiplied by a phase factor Φ which is

$$\begin{aligned} \Phi = (-1)^{\tilde{\varphi}} & \prod_{a < b < \#J_{\text{in}}} \text{sgn}(J_{\text{in},a} - J_{\text{in},b}) \\ & \times \prod_{a < b < \#J_{\text{out}}} \text{sgn}(J_{\text{out},a} - J_{\text{out},b}) \end{aligned} \quad (6.28)$$

where $J_{\text{in},a}$ and $J_{\text{out},a}$ denote the a^{th} -element of J_{in} and J_{out} respectively, $\#$ is the cardinality of a set, and $\text{sgn}(x)$ gives the sign of x .

6.3.3 Properties of the Parameterization

We now discuss the parameterization's properties, including its spectrum and matrix structure. The non-Gaussian contribution derived in this chapter is Hermitian $\delta\rho = \delta\rho^\dagger$ and has zero trace $\text{Tr}[\delta\rho] = 0$. However, it does not constitute a bona fide quantum state, as $\rho + \delta\rho$ is not always a positive operator if only $\tilde{C}^{(4)}$ is included. This is a consequence of feeding only partial information into the parameterization: Generically, if $\tilde{C}^{(4)}$ increases, so will higher-order correlations, $\tilde{C}^{(6)}, \tilde{C}^{(8)} \dots$. If they are not included, the magnitudes of negative eigenvalues increase with $\tilde{C}^{(4)}$, indicating that the approximation being used becomes invalid. Including them will lead to a positive semi-definite matrix. In practice this is a minimal numerical effect when used within the applicability of the ansatz: Even for small but non-trivial $\tilde{C}^{(4)}$, there may be a few negative eigenvalues, although with extremely small magnitudes. Nevertheless, this can cause problems when computing the entanglement Hamiltonian numerically by taking the matrix-logarithm of a state. We resolve this issue and define a proper state by diagonalizing the matrix and projecting the Schmidt spectrum onto the closest classical probability distribution under the 2-norm using the approach of [463].

The non-Gaussian part $\delta\rho$ has diagonal and off-diagonal components. The diagonal part includes contributions from connected 4-point functions corresponding to “density-density” correlations, specifically $\tilde{C}_{ijji}^{(4)}$ with $i \neq j$. On the other hand, “pair-density” correlations, such as $\tilde{C}_{ijkj}^{(4)}$ with $i \neq j \neq k$, contribute to off-diagonal elements, where the occupations $\mathbf{m}$ and $\mathbf{n}$, when written as a bitstring, differ by a Hamming distance of $d_H(\mathbf{m}, \mathbf{n}) = 2$. Additionally, “pair-pair” correlations, $\tilde{C}_{ijkl}^{(4)}$ with $i \neq j \neq k \neq l$, connect quasi-

particle states with a Hamming distance of $d_H(\mathbf{m}, \mathbf{n}) = 4$. In our numerical simulations we observed that including only the diagonal two-body 'density-density' correlations is insufficient to produce level repulsion when the underlying state behaves chaotically. We attribute this behavior to uncorrelated neighboring eigenvalues (as in a Poissonian distribution), which is preserved by the diagonal two-body correlations. Thus the ability to measure off-diagonal correlations, not just simple non-Gaussian density correlations, is important to diagnose the phenomena we study in the main text.

6.4 Quench Setup

We apply the non-Gaussian state parameterization to study the non-equilibrium dynamics after a quantum quench of a Fermi-Hubbard chain of spinful fermions with Hamiltonian [464]

$$H(J, J', U) = J \sum_{l,s} c_{l+1s}^\dagger c_{ls} + J' \sum_{l,s} c_{l+2s}^\dagger c_{ls} + \text{h.c.} \\ + U \sum_l n_{l\uparrow} n_{l\downarrow} \quad (6.29)$$

where l labels the lattice site, $s = \uparrow, \downarrow$ labels the spin, and $n_{ls} = c_{ls}^\dagger c_{ls}$ is the number operator⁴. Utilizing exact diagonalization [465], we initialize a random eigenstate of $H_0 = H(1, 1/8, 0)$, which is Gaussian, and then evolve with $H(1, 1/8, U)$ for a fixed value of interaction strength U . This quench setup allows us to study the build-up and evolution of non-Gaussian correlations. We monitor the entanglement structure of the exact time evolved state and the parameterization $\tilde{\rho}$, utilizing exact correlation functions.

6.5 Early-time Dynamics

In the early-time stage of the evolution, there are two regimes distinguished by the dynamics, $\tau/N_A \ll 1$ and $\tau/N_A \geq \mathcal{O}(1)$ with $\tau = tU$. In the first regime, the four-point correlations grow quickly, with higher-order correlations being suppressed by powers of the interaction strength. For the second regime, higher-order correlations give a non-negligible contribution to ρ_A . We introduce a metric θ that measures the non-Gaussianity of a state,

$$\theta(\sigma) \equiv \arccos \sqrt{\mathcal{F}(\sigma|\sigma_g)} \quad (6.30)$$

⁴Extended by next-to-nearest neighbor tunneling, this model is not integrable even in one spatial dimension




Figure 6.5: *Non-equilibrium evolution of reconstructed state.* **a,b)** Non-Gaussianities $\theta(\sigma)$, see Eq. (6.30), build up analogously for the exact state ρ_A [(a)] and the reconstructed state $\tilde{\rho}$ [(b)], as indicated by the dashed line which shows a linear fit to the early time behavior $\theta(\rho_A)$. Here we show numerical data for a system of size $N = 10$, $N_A = 5$ with 10 different initial states (shaded lines) and their averages (solid line). **c)** We test the accuracy of the state reconstruction through the spectral error of levels i of the entanglement spectra $\{\tilde{\epsilon}\}$, $\{\epsilon_g\}$ for different interaction strengths U (at a fixed time). The non-Gaussian parameterization reduces the error to $\mathcal{O}(U^2)$ relative to the Gaussian parameterization with error $\mathcal{O}(U)$.




Figure 6.6: *Evolution of level statistics of non-Gaussian state parameterization* **a,b)** The gap ratio distributions of both, the parameterization $\tilde{\rho}$ (filled, purple) and the exact state ρ_A (unfilled, gold), evolve from Poisson (pink dashed) at time $\tau = 0.015$ to GUE (black dashed) at time $\tau = 0.47$, thus signaling the onset of quantum ergodic evolution. **c)** The average gap ratio $\langle r \rangle$ grows with time and saturates to the value ~ 0.6 indicating level repulsion (GUE). Data is shown for a system of size $N = 10$, $N_A = 5$, interaction strength $U = 5.6 \times 10^{-3}$; averages are taken over 10 different initial states, where shaded regions signify 95% confidence intervals. The dashed gray lines signify the times displayed in (a) and (b).

where $\mathcal{F}(\sigma|\sigma') = \frac{\text{Tr}[\sigma\sigma']}{\max\{\text{Tr}[\sigma^2], \text{Tr}[\sigma'^2]\}}$ is the max-fidelity [466], and σ_g is the Gaussian part of σ as in Eq. (6.2). In Fig. 6.5, we show the deviation of the exact state from a Gaussian state, $\theta(\rho_A)$, in (a) and the deviation of the parameterization $\theta(\tilde{\rho})$ in (b), for both time regimes. At first the infidelity grows quadratically in τ , i.e. $1 - \mathcal{F} = \mathcal{O}(\tau^2)$, implying that both $\theta(\rho_A)$ and $\theta(\tilde{\rho})$ will grow linearly during this stage. The dashed line in (b) represents the linear part of $\theta(\rho_A)$ overlayed on $\theta(\tilde{\rho})$, highlighting that the growth is the same. Once the second stage of evolution has been reached, the expansion restricted to 1- and 2-body correlations is no longer quantitative, indicating that higher-order correlations not included in $\tilde{\rho}$ are now contributing to the dynamics. For small couplings $U \ll J, J'$ we also observe that θ has a self-similar profile in time (solid purple, dashed pink), which is inherited from the underlying correlations of ρ_A .




Figure 6.7: Average gap ratio of exact state and non-Gaussian parameterization **a)** The average gap ratio of the exact state over a range of times Jt and interaction strengths U/J . The black dashed line shows the crossover from a Poisson distribution to GUE, indicating quantum chaos through level repulsion. **b)** The average gap ratio for the parameterization qualitatively reproduces that of the exact state over the times and interaction strengths studied.




Figure 6.8: Self-similarity of non-Gaussian correlations. **a)** Example of a “density-density” correlation, $\tilde{C}_{ijji}^{(4)}$ with $i \neq j$, for the same interaction strengths as shown in Fig. 6.5. For $U = 10^{-3}, 10^{-2}$, the correlations collapse onto the same curve when rescaled by the interaction strength. As $U \sim J, J'$ oscillations begin to appear on top of this curve that breaks this behavior. **b)** Example of a “pair-density” correlation function $\tilde{C}_{ijki}^{(4)}$ with $i \neq j \neq k$. **c)** Example of a “pair-pair” correlation function ($i \neq j \neq k \neq l$).

6.6 Early Time Growth of Non-Gaussianity

In this section, we provide additional quantitative details on the early-time dynamics of the quantum quench. Specifically, we examine the dynamics of our measure of non-Gaussianity by comparing the deviation between the exact state, $\theta(\rho_A)$, and our parameterization, $\theta(\tilde{\rho})$, from a Gaussian state. As shown in Fig. 6.5(a) and (b), we observed that for small couplings $U \ll J, J'$, θ grows linearly and exhibits self-similarity during the early times. Both $\theta(\rho_A)$ and $\theta(\tilde{\rho})$ behave similarly in this regime, so our discussion focuses on $\theta(\rho_A)$. Specifically, without loss of generality, we consider the dynamics of the full ρ instead of the reduced ρ_A because we have checked that $\theta(\rho)$ exhibits the same linear growth as $\theta(\rho_A)$.

The system is composed of N sites with interactions described by Hamiltonian $H = H_0 + H_I$, where $H_0 = J \sum_{l,s} c_{l+1s}^\dagger c_{ls} + J' \sum_{l,s} c_{l+2s}^\dagger c_{ls} + h.c.$ and $H_I = U \sum_l n_{l\uparrow} n_{l\downarrow}$. The full system is translationally invariant with periodic boundary conditions, and so we define momentum-space mode operators

$$b_{k\sigma}^\dagger = \frac{1}{\sqrt{N}} \sum_l e^{ikl} c_{l\sigma}^\dagger. \quad (6.31)$$

Because ρ is pure $\text{Tr}[\rho^2] = 1$, allowing us to write the fidelity as $\mathcal{F}(\rho|\rho_g) = \text{Tr}[\rho\rho_g]/\max\{\text{Tr}[\rho_g^2], \text{Tr}[\rho^2]\} = \text{Tr}[\rho\rho_g]$, where ρ_g is the Gaussian state with the same two-point correlations as the exact state. In this study, the quench begins from the initial plane wave state,

$$\rho(0) = |\mathbf{n}_0\rangle \langle \mathbf{n}_0|, \quad |\mathbf{n}_0\rangle = \prod_k (b_{k\sigma}^\dagger)^{(\mathbf{n}_0)_{k\sigma}} |0\rangle \quad (6.32)$$

where $\mathbf{n}_0$ is the initial occupation of momentum modes written as a bitstring and $|0\rangle$ is the Fock vacuum. The fidelity, as a function of evolution time, is

$$\begin{aligned} \mathcal{F} &= \langle \mathbf{n}_0 | e^{i(H_0+H_I)t} \rho_g(t) e^{-i(H_0+H_I)t} | \mathbf{n}_0 \rangle \\ &= \langle \mathbf{n}_0 | e^{iH_I t} \rho_g(t) e^{-iH_I t} | \mathbf{n}_0 \rangle + \mathcal{O}(JU t^2). \end{aligned} \quad (6.33)$$

In the second line, we separated the time evolution operator into its free and interacting components, using the first-order Trotter-Suzuki formula, and we used the fact that $|\mathbf{n}_0\rangle$ is an eigenstate of the free Hamiltonian H_0 . Likewise, the Gaussian state may be written as $\rho_g = [\prod_{k,\sigma} g_{k\sigma}] \exp\{\sum_{k,\sigma} \log[g_{k\sigma}(1-g_{k\sigma})^{-1}] b_{k\sigma}^\dagger b_{k\sigma}\}$ where the indices $k \in \{-\pi, \dots, \pi - 2\pi/N\}$, $\sigma \in \{\uparrow, \downarrow\}$. Here, the two-point function $\langle b_{k\sigma}^\dagger b_{k'\sigma'} \rangle = g_{k\sigma} \delta_{kk'} \delta_{\sigma\sigma'}$ is diagonal in the plane wave basis for translationally invariant systems. The early time behavior of the $g_{k\sigma}$

can be found by solving the Heisenberg equations of motion $\partial_t(b_{k\sigma}^\dagger b_{k\sigma}) = i[H_0 + H_I, b_{k\sigma}^\dagger b_{k\sigma}]$. In the plane wave basis $H_0 = \sum_{k\sigma} \varepsilon_k b_{k\sigma}^\dagger b_{k\sigma}$ and

$$H_I = \frac{U}{N} \sum_{k_1, k_2, k_3, k_4} \delta(k_1 + k_2 - k_3 - k_4) b_{k_1, \uparrow}^\dagger b_{k_2, \downarrow}^\dagger b_{k_1, \downarrow} b_{k_3, \uparrow}$$

with $\varepsilon_k = 2[J \cos(k) + J' \cos(2k)]$, and so integrating the equations of motion gives

$$\begin{aligned} g_{k\sigma}(t) &= g_{k\sigma}(0) + \frac{iU}{N} \int_0^t ds \sum_{k_1, k_2, k_3, k_4} \delta(k_1 + k_2 - k_3 - k_4) \\ &\quad \times \langle [b_{k_1, \uparrow}^\dagger b_{k_2, \downarrow}^\dagger b_{k_1, \downarrow} b_{k_3, \uparrow}, b_{k\sigma}^\dagger b_{k\sigma}] \rangle(s) \\ &\approx (\mathbf{n}_0)_{k\sigma} + \frac{iUt}{N} G_{k\sigma} \end{aligned} \quad (6.34)$$

where we have used the fact the initial state is $|\mathbf{n}_0\rangle$ and

$$\begin{aligned} G_{k\sigma} &= \sum_{k_1, k_2, k_3, k_4} \delta(k_1 + k_2 - k_3 - k_4) \\ &\quad \langle [b_{k_1, \uparrow}^\dagger b_{k_2, \downarrow}^\dagger b_{k_1, \downarrow} b_{k_3, \uparrow}, b_{k\sigma}^\dagger b_{k\sigma}] \rangle(0) . \end{aligned}$$

In defining $G_{k\sigma}$ we have treated the commutator term as stationary on the timescales considered, which is an approximation that is used in the literature [467]. With knowledge of $g_{k\sigma}$, at early times we may express the Gaussian state as $\rho_g(t) = B(t) |\mathbf{n}_0\rangle \langle \mathbf{n}_0| B^\dagger(t)$ with $B(t) = \exp\{\frac{iUt}{N} \sum_{k,\sigma} G_{k\sigma} / [(\mathbf{n}_0)_{k\sigma} (1 - (\mathbf{n}_0)_{k\sigma})] b_{k\sigma}^\dagger b_{k\sigma}\}$, which comes from expanding $\log[g_{k\sigma}(1 - g_{k\sigma})^{-1}]$ for small Ut . Thus Eq. (6.33) becomes

$$\begin{aligned} \mathcal{F} &= \langle \mathbf{n}_0 | e^{iH_I t} B(t) | \mathbf{n}_0 \rangle \langle \mathbf{n}_0 | B^\dagger(t) e^{-iH_I t} | \mathbf{n}_0 \rangle + \mathcal{O}(JUt^2) \\ &= \langle \mathbf{n}_0 | (1 + iH_I t + \frac{iUt}{N} \sum_{k\sigma} G_{k\sigma} b_{k\sigma}^\dagger b_{k\sigma}) | \mathbf{n}_0 \rangle \\ &\quad \times \langle \mathbf{n}_0 | (1 - iH_I t - \frac{iUt}{N} \sum_{k\sigma} G_{k\sigma} b_{k\sigma}^\dagger b_{k\sigma}) | \mathbf{n}_0 \rangle + \mathcal{O}(\tau^2) \\ &= 1 + \mathcal{O}(\tau^2) , \end{aligned} \quad (6.35)$$

where on the second line we have expanded the exponents for small Ut .

In conclusion, we have shown that the terms linear in $\tau = Ut$ cancel, and so the fidelity $\mathcal{F}$ will decrease from 1 at most quadratically in τ . Because we defined the distance $\theta(\rho) = \arccos \sqrt{\mathcal{F}(\rho | \rho_g)}$, we expect that $\theta(\rho)$ grows linearly at early times.

We attribute the observed self-similarity to the system's two-body (“four-point”) correlations, although the exact mechanism remains unknown to us at this point and deserves further study. To illustrate this effect, Fig. 6.8 shows selected four-point functions used in the calculation of $\theta(\tilde{\rho})$ as functions of τ for the same values of U as in Figure Fig. 6.5 (a) and (b). For $U \ll J, J'$, the correlations collapse onto the same curve. However, for $U \sim J, J'$, oscillations on this curve begin to grow until all resemblance to the small U correlations is lost.

As a measure for how accurately the parameterization captures the entanglement spectrum $\{\epsilon_i\}$ of the exact state, we consider the spectral error

$$\Delta_i = |\epsilon_i - \epsilon'_i| \quad (6.36)$$

where $\{\epsilon'_i\}$ is the entanglement spectrum of either $\tilde{\rho}$ or ρ_g . In Fig. 6.5(c), we show the spectral error between $\rho_g, \tilde{\rho}$ and ρ_A in levels $i = 0, 10, 50, 100$ as a function of τ with $t = 10/J$ fixed. The figure shows that, for small interaction strength, the spectral error between ρ_g and ρ_A (dashed, pink) behaves as $\mathcal{O}(U)$, while the spectral error between $\tilde{\rho}$ and ρ_A (dashed, purple) scales as $\mathcal{O}(U^2)$. Thus $\tilde{\rho}$ captures the leading order in U dependence of ρ_A , and furnishes a better approximation of the exact state than just ρ_g . For larger interaction strengths neither ρ_g nor $\tilde{\rho}$ accurately estimate the spectrum of ρ_A .

6.7 Thermalization Dynamics

The statistical distribution of the entanglement spectrum $\{\epsilon_i\}$ is an indicator of quantum ergodicity and chaos [460]. Specifically, we consider the gap ratio [448]

$$r_i = \frac{\min(\delta_i, \delta_{i-1})}{\max(\delta_i, \delta_{i-1})} \quad (6.37)$$

where $\delta_i = \epsilon_{i+1} - \epsilon_i$. Figs. 6.6(a),(b) show the gap ratio distributions $P_t(r)$ of ρ_A and $\tilde{\rho}$ averaged across symmetry sectors and different initial states. These are compared to both a non-repulsive Poisson distribution and a Gaussian Unitary Ensemble (GUE). At early times the level statistics of both the exact state and parameterization are consistent with Poisson behavior while subsequent times show (GUE) level repulsion. Fig. 6.6(c) shows the average gap ratio $\langle r \rangle(t)$ of the parameterization versus the exact state over time, demonstrating that the level statistics of the exact state are reproduced. Fig. 6.7 shows the average gap ratio over a range of interaction strengths U/J and time Jt , comparing both ρ_A in a) with $\tilde{\rho}$ in b). Even at large U

the parameterization captures the level statistics of the exact state, and accurately reproduces the transition line between Poisson and GUE distributions.

A Gaussian state cannot exhibit level repulsion because its EH is quadratic and parameterized solely by $C^{(2)}$. Regardless of $C^{(2)}$, the EH cannot show level repulsion, except in the single-particle sector, simply because of the inherent level degeneracies in the many-particle sectors of a quadratic Hamiltonian. The incorporation of small $C^{(4)}$ s is sufficient to break this degeneracy after short time evolution. The time scale at which this happens, nearly identical for the exact and the parametrized state, reveals a crucial insight: the emergence of quantum chaos necessitates the presence of two-body correlations of a certain magnitude.

6.8 Discussion of Symmetries and Choice of the Initial State

We now discuss the symmetries, relevant for this study, of the Fermi-Hubbard model with a next-nearest-neighbor hopping [468]. One $U(1)$ symmetry is local particle number conservation,

$$N = \sum_l n_{l,\uparrow} + n_{l,\downarrow}. \quad (6.38)$$

For a subsystem of size N_A , there are $2N_A + 1$ particle number sectors with associated quantum numbers $n \in \{0, \dots, 2N_A\}$. Additionally, the model has an $SU(2)$ symmetry, generated by the spin operators

$$S^i = \frac{1}{2} \sum_l \sum_{r,r'} c_{l,r}^\dagger (\sigma^i)_{rr'} c_{l,r'} \quad (6.39)$$

where $l \in \{0, \dots, N_A - 1\}$, $r, r' \in \{\uparrow, \downarrow\}$ and $\vec{\sigma} = \{\sigma^x, \sigma^y, \sigma^z\}$ is a vector of Pauli matrices. This symmetry corresponds to global angular momentum conservation. The quantum numbers associated with this symmetry are the eigenvalues of the total magnetization operator $\hat{S}^z = \frac{1}{2} \sum_l n_{l\uparrow} - n_{l\downarrow}$ and the Casimir operator $S^2 = (S^x)^2 + (S^y)^2 + (S^z)^2$. There are $2N_A + 1$ quantum number sectors associated with S^z , denoted by $m \in \{-N_A, \dots, N_A\}$ and $N_A + 1$ quantum numbers associated to S^2 , denoted by $s \in \{0, \dots, N_A(N_A + 1)\}$. The three operators N, S^z and S^2 form the symmetry operators whose eigensectors are important for analyzing the dynamics of this model in this chapter. In practice, one must project the state $\rho_A(t)$ into these symmetry sectors when analyzing the level distribution of the entanglement Hamiltonian, because eigenvalues from different sectors remain trivially uncorrelated when the initial state also respects these symmetries.

There is also a particle-hole like symmetry that exists at half-filling, i.e. $n = N_A$ (see [469]), however its details are largely irrelevant for our purposes. Additionally, modifications of the model may feature additional




Figure 6.9: *Rank of the state*. Growth of rank for a system of $N = 10$ during the free evolution averaged over 30 random initial states. The pink region shows the range of growth rates over the averaged states.

symmetries, such as e.g., non-Abelian η -pairing symmetry [469], however they are either not relevant for us or explicitly broken. Space-time symmetries are broken explicitly when one restricts the state to a subsystem.

In this study, we focus on initial momentum-space Fock states given by Eq. (6.32). We now discuss how to choose such an initial state for the quench experiment that allows measurement of the build-up of level repulsion. Because studying level repulsion requires block-diagonalizing the state into symmetry sectors, and because the model has so many symmetries reflected in the measured correlation functions, one option is to resolve all of them when analyzing ρ_A . Another is to break the symmetry explicitly by the initial condition. Opting for a combination of the two, we will break the S^2 and particle-hole symmetries of the model by choice of a suitable initial state. The particle-hole symmetry is broken by choosing an initial state that is not at half-filling. The S^2 symmetry may be broken by choosing a state that is not an eigenstate of the Casimir, i.e. $\mathcal{O}_0 = [\rho_0, S^2] \neq 0$. The algorithm for finding such an initial state is as follows:

1. Choose a global number sector n that is away from half-filling (i.e. $n \neq N_A$).
2. Randomly generate a trial product state in the Fock basis, Eq. (6.32), that satisfies $\sum_{k\sigma} (\mathbf{n}_0)_{k\sigma} = n$

3. Check $\text{Tr}[\mathcal{O}_0^\dagger \mathcal{O}_0] \neq 0$. If this condition is satisfied, accept this initial state and proceed with time evolution. If it is not, return to step 2.




Figure 6.10: *Position space initial state.* **a)** The average gap ratio of the exact state with a position space initial state. **b)** The average gap ratio $\langle r \rangle(t)$ predicted by the parameterization.

An alternative strategy is to start from initial *position-space* Fock states

$$|\tilde{\mathbf{n}}_0\rangle = \prod_k (c_{k\sigma}^\dagger)^{(\tilde{\mathbf{n}}_0)_{k\sigma}} |0\rangle \quad (6.40)$$

which are un-entangled product states, making their experimental preparation simpler. Symmetries can be adjusted in the exact same way as for momentum-space states. Since such a state is a product state and therefore a very specific type of Gaussian state, it has zero entanglement and an initially trivial entanglement Hamiltonian with rank $\mathcal{R}_A = 1$. Consequently, our primary observable, the entanglement gap ratio, is initially undefined. To enable comparison with a generic Gaussian state, we first evolve the state with the free Hamiltonian H_0 for a duration t_0 before turning on interactions, i.e.

$$|\psi\rangle(t) = \begin{cases} e^{-itH_0} |\tilde{\mathbf{n}}_0\rangle & \text{for } t < t_0 \\ e^{-i(t-t_0)(H_0+H_I)} e^{-it_0H_0} |\tilde{\mathbf{n}}_0\rangle & \text{for } t \geq t_0 \end{cases} \quad (6.41)$$

During the onset period, we monitor the effective rank $\mathcal{R}_A \equiv \lim_{\alpha \rightarrow 0} \exp\{\frac{1}{1-\alpha} \log(\text{Tr}[\rho^\alpha])\}$ of the bipartite reduced density matrix. It increases until it reaches full rank while remaining a Gaussian state, see Fig. 6.9. We ensure that t_0 is always chosen large enough to achieve full rank before turning on interactions. In experiment, the turning on of interactions could be realized by keeping the two fermion species in separate lattices first before bringing them together using a reconfigurable array of optical tweezers or moving a spin-selective optical lattice.

After the interactions are turned on, our analysis is identical to that of the momentum state quench. Fig. 6.10 shows the average gap ratio over a range of interaction strengths for both ρ_A and $\tilde{\rho}$. Even for large U the parameterization captures the level statistics of the exact state, and accurately maps the cross over between uncorrelated levels (Poisson) and level-repulsion (GUE).

6.9 Experimental Application

Our approach can be applied in conjunction with current ultracold-atom simulators to address entanglement structure in experiments. We consider platforms which load fermionic atoms into optical lattices to realize Hubbard models with tunable nearest-neighbor hopping and on-site interaction strengths U [432,451,470,471] similar to the model considered in Eq. (6.29).

To extract the entanglement structure using the above protocol, we require spin-resolved measurements of both diagonal fermion density correlations $\langle n_i n_j \rangle$ as well as off-diagonal fermion correlations $\langle c_i^\dagger c_j \rangle$, $\langle c_i^\dagger c_j^\dagger c_k c_l \rangle$ across the subsystem, see Fig. 6.1(a). Besides site-resolved imaging techniques, such as quantum gas microscopes [105], this requires the ability of realizing tunneling operations between arbitrary pairs of sites i and j in the subsystem. The required tunneling operations can for instance be efficiently implemented using optical tweezers [453,472].

To understand how one can use pair-wise tunneling to measure off-diagonal correlations, we write the operator $c_i^\dagger c_j = S_{ij}^x + iS_{ij}^y$, where $S_{ij}^x = \frac{1}{2}(c_i^\dagger c_j + c_j^\dagger c_i)$ and $S_{ij}^y = \frac{-i}{2}(c_i^\dagger c_j - c_j^\dagger c_i)$. By applying $\pi/2$ -tunneling operations, $\mathcal{U}_{ij}^x(\pi/2) = \exp\{-i\frac{\pi}{4}S_{ij}^x\}$ and $\mathcal{U}_{ij}^y(-\pi/2) = \exp\{i\frac{\pi}{4}S_{ij}^y\}$, we bring S_{ij}^x and S_{ij}^y , and thus $c_i^\dagger c_j$, into diagonal form where they are read from measurements of the local occupation numbers. The same technique can be applied to another set of sites in parallel to obtain the four-body operators. In general, one needs to perform $N_A(N_A - 1)$ measurements for one-body correlations, where N_A is the number of modes in the subsystem (including spin), and $N_A(N_A - 1)(N_A - 2)(N_A - 3) + O(N_A^2)$ measurements

for two-body correlations. k -body correlations require measurements in $O(N_A^{2k})$ bases where for $k \rightarrow \infty$ one performs, at exponential cost, full tomography. Furthermore, generalized measurement protocols based on random atomic beam splitter operations in optical lattices [385, 473–477] can be applied in the present context.

The advantage of our method is the use of polynomial (in system size) number of correlation functions—standard observables in experiments—as input, rather than requiring full state tomography. However, depending on the specific application, a complexity analysis must also account for the sampling cost associated with estimating correlation functions to the desired precision and, if targeting the EH’s level spacing distribution, the classical representation of the reduced density matrix, which we address in 6.10. Our analysis suggests that the spectral properties of the EH we target may be costly to resolve, potentially exponentially so—especially at early times, where higher-order correlations tend to be overestimated with finite shots. This direction has only recently begun to be addressed [478, 479].

The accuracy of approximation is dependent on the strength of correlations that the state exhibits, which one has no a priori knowledge of in a quantum simulation experiment. Convergence can be explicitly tested by measuring higher-order correlations and incorporating them into the parameterization, as described now.

6.10 Cost, efficiency, and noise

The method proposed here enables the reconstruction of an approximate quantum state using lower-order correlation functions. Applications—such as extracting the level distribution of an associated entanglement Hamiltonian—incur costs in both the sampling complexity of measuring the correlation functions and in the classical postprocessing required. These procedures are also sensitive to noise. In the following, we address these challenges and note that strict complexity-theoretic performance guarantees remain an open question warranting further investigation.

We begin by addressing the cost associated with classical postprocessing—specifically, whether our application requires storing the (exponentially large) reduced density matrix on a classical computer in order to compute its spectrum. In principle, it is not necessary to store the full matrix to access spectral information. Instead, one can apply generalized degenerate perturbation theory to an order consistent with the included correlation functions. To illustrate this approach, we combine the Gaussian part with the diagonal contribution

of $\delta\rho$ into a single diagonal matrix,

$$D_{\mathbf{mn}} = \delta_{\mathbf{mn}}(\rho_{g,\mathbf{nn}} + \delta\rho_{\mathbf{nn}}) . \quad (6.42)$$

We work in the occupation basis where ρ_g is diagonal and $\delta\rho_{\mathbf{nn}}$ are the corresponding contributions with $d_H(\mathbf{n}, \mathbf{n}) = 0$ from Eq. (6.5). The remaining contributions with $d_H(\mathbf{m}, \mathbf{n}) = 2$ or 4 form the perturbation $V_{\mathbf{mn}}$. We can thus calculate the perturbed eigenvalues $D'_{\mathbf{nn}}$ of $\rho = D + V$ as

$$D'_{\mathbf{nn}} = D_{\mathbf{nn}} + \sum_{\mathbf{m} \neq \mathbf{n}} \frac{V_{\mathbf{nm}} V_{\mathbf{mn}}}{D_{\mathbf{n}} - D_{\mathbf{m}}} . \quad (6.43)$$

The sum over all occupation numbers $\mathbf{m}$ is *not* exponentially large in the subsystem size because for any given $\mathbf{n}$ only contributions with $d_H(\mathbf{m}, \mathbf{n}) = 2$ or 4 from the four-point functions arise, which reduces the amount of terms to be computed to a number polynomial in subsystem size, specifically $O(N_A^4)$. This shows that one can have access to the entanglement spectrum without classically reconstructing the reduced state as a matrix.

Proceeding this way allows for the estimation of certain eigenvalues—for example, the lowest-lying levels—at a given order, but it is insufficient to capture level repulsion at low orders in perturbation theory, see e.g. [480]. The entire reduced state must be explicitly constructed from our parameterization and diagonalized which incurs a cost exponential in subsystem size, though importantly does not scale with the size of the full system. This holds despite the fact that two- and four-point functions carry sufficient information to discern level repulsion. Fortunately, since level repulsion is a universal quantity, it is sufficient to choose a subsystem large enough to yield meaningful level statistics, but still small enough to be classically representable.

Second, we now address the effect of shot noise on reconstructing the spectral statistics. In Fig. 6.11 (a) we emulated the measurement protocol for 4 initial states at times $\tau = 0.01, 0.02, 0.05, 0.08, 0.1, 0.2, 0.3, 0.4, 0.5$ for four runs with $N_S = 10^4, 3 \cdot 10^4, 5 \cdot 10^4, 8 \cdot 10^4$, where N_S is the number of shots for each basis in which correlation functions are extracted, for each run. After times $\tau \geq .05$, we observe that the gap ratio from the estimated four point functions is able to closely track the gap ratio from the exact four-point functions. However, for the first few time steps the convergence to the exact prediction is observed to be much slower, behavior that can be attributed to the typical size of non-Gaussian correlations at very early times.

Fig. 6.11 (b) depicts the real part of an example four-point function with the associated finite-shot estimate at $\tau = 0.01, 0.1, 0.2, 0.3, 0.4, 0.5$, which for $\tau = 0.01$ is of the order $\mathcal{O}(10^{-6})$ while the shot-noise estimated value is $\mathcal{O}(10^{-3})$, both smaller than the two-point functions which are $\mathcal{O}(10^{-1})$. At later times the four-point




Figure 6.11: *Shot noise analysis.* **a)** The build-up of level repulsion for $N = 8$, $N_A = 4$ averaged across 4 initial states. The gold line represents the exact result, the purple line represents the ansatz with exact four-point functions, and the blue dots denote the ansatz measured with finite shots; the shading indicates the number of samples used. **b)** An example of four-point function (points) compared to the exact four-point function (continuous curve). The error bars are given by the standard deviation from the mean and the number of shots are $N_S = 5 \cdot 10^4$

function has grown sufficiently such that it can be accurately measured while still being small enough for the parameterization to remain valid.

In general, experimental imperfections and shot noise tend to overestimate higher-order correlation functions relative to those of Gaussian states which are structurally special. This leads to an overestimation of level repulsion. Taken together, these observations highlight that estimating entanglement-related properties of near-Gaussian states is a subtle and nuanced task. While initial complexity-theoretic results provide valuable insights [478, 479], this merits further investigation.

6.11 Conclusions

In this section, we addressed the entanglement structure of quantum states, i.e. entanglement Hamiltonians and their spectrum, from the perspective of lower-order k -body correlation functions. Our approach is based

on a systematic expansion of the Wigner function around Gaussian states, where non-Gaussian contributions are constructed directly from higher-order correlations. While our primary example focused on $k = 1$ and 2, the method can be systematically extended to all orders. Furthermore, the approach is model-independent and applicable in higher dimensions.

Our motivation for this study was multi-fold. On one hand, our approach provides a practical avenue for measuring the entanglement structure in analog quantum simulator experiments with limited control. Especially at weak-coupling, these experiments could enable the exploration of quantum chaos and thermalization at late times, connecting entanglement dynamics to the growth of lower-order correlations. We posit that the parameterization we have developed is useful at late times near-thermal equilibrium, where states are expected to be close to Gaussian. Thus our approach offers a pathway from the early to the late stages of the thermalization dynamics [481–485].

One area that remains unexplored is the connection between correlation functions and entanglement in high-energy and nuclear physics experiments. In ultra-relativistic heavy ion collisions many-particle correlations are a key experimental observable [486,487]. Despite the lack of direct access to any off-diagonal correlations in experiment, correlation functions are central for theoretical descriptions, including effective kinetic theory [488] and hydrodynamics [489]. Within these frameworks, our approach may see use for studying physics encoded in the entanglement spectrum. Furthermore, the intersections of our approach with other theoretical tools available for studying quantum-many body dynamics, such as matrix product states [490] or time-dependent variational techniques [491,492] ought to be explored. Our technique is more broadly applicable for characterizing quantum states without resorting to full tomography, e.g. to probe phases of matter through the low-energy spectrum of the EH [446] or to evaluate measures of quantum complexity [308]. Finally, non-Gaussianity can be viewed as a quantum information theoretic resource [493–495]. Its role for measuring complexity of quantum many-body systems should be explored [305,307,313,496–499].

Chapter 7

REALIZING ERROR SUPPRESSION IN PARTIALLY FAULT-TOLERANT QUANTUM SIMULATIONS WITH IBM QUANTUM COMPUTERS

7.1 Introduction

Performing quantum simulations with sufficient precision to compute many scientifically relevant observables requires levels of error suppression attainable only through quantum error correction (QEC) [500–504]. The qubit and gate overhead of fault-tolerant (FT) quantum computation has so far kept large-scale QEC out of reach for practical quantum simulations. Leveraging high gate fidelities and flexible connectivity, recent demonstrations on trapped ions [505, 506] and neutral atoms [507–509] have achieved “beyond break-even” performance, with the encoded error rate falling below the unencoded rate [510]. Solid-state platforms have seen similar success, albeit in more restricted settings [511–518]. These results suggest that in certain cases, presently available hardware can benefit from FT implementations beyond what error mitigation alone can achieve. Realizing this benefit in quantum simulations requires carefully balancing the error-rate (ER) improvements from FT against its associated overhead.

An approach to computing with encoded qubits that is compatible with currently available hardware is to supplement FT components with non-FT (nFT) ones, allowing some physical errors to go undetected in exchange for lower circuit overhead [519–524]. In this partially FT approach, for a physical error rate p ,¹ nFT components contribute a factor of Ap to the logical error rate (LER), whereas FT components that remove $\mathcal{O}(p)$ errors provide a suppression of Bp^2 , giving a total logical error rate of $p_L(p) = Ap + Bp^2$. By carefully designing the nFT components such that $A \ll B \ll 1$, the computation can go below the “pseudothreshold”, i.e., p such that $p_L(p) < p$. The nFT components generally consist of logical rotations, for which a fully FT implementation would require magic state injection [525] or code switching [526, 527]. Although topological codes such as the surface code [528, 529] and the heavy-hex code [530] are naturally suited to limited-connectivity devices [531], the measurement and decoding overhead of their lattice-surgery

¹In practice many processes contribute to p . In this work p characterizes the combined error rate of all of these processes, which is typically dominated by two-qubit gate errors.

operations places them out of reach for present-day simulations [532].

While the full power of QEC may still be beyond reach, immediate improvements can be achieved through the adoption of lightweight, low-overhead quantum error detection (QED) encodings [533–539]. A motivation for this work is to use FT to advance quantum simulations of lattice gauge theories [540–546] describing the fundamental forces of Nature [59, 82, 295, 390, 434, 547]. The mappings and dynamics of such simulations are organized by symmetries that group the degrees of freedom into repeating units. As such, it is natural to study FT schemes whose code blocks coincide with these units.² The recently studied Iceberg codes are a family of $[[n + 2, n, 2]]$ codes that use two extra physical qubits to encode n logical qubits into a single block [502, 505, 533, 548–552]. Using multiple interacting code blocks of the $[[4, 2, 2]]$ Iceberg code, we perform Hamiltonian simulations on IBM’s heavy-hexagonal quantum computers and achieve beyond break-even performance on the calculation of local observables. This is done through the development of shallow syndrome extraction circuits designed for the device’s native qubit connectivity. We implement nFT logical rotations, minimizing their depth by selecting specific code block layouts and gate scheduling. One unavoidable feature of QED codes is the exponential loss of ensemble size as the number of syndrome extraction rounds increases. For a computation with many code blocks, this loss prohibits running even a single round of syndrome extraction if all shots record an error and are discarded. This work proposes a selective-filtering strategy, Observable-Ranked Postselection (ORP), that discards detection events based on their probability of causing a logical error, determined through correlations between syndrome flips and logical observable values.

We demonstrate the benefits of these QED methods through simulations of false-vacuum decay in the 1+1D and 2+1D Mixed-Field Ising model (MFIM). Despite a two to six times increase in circuit depth, encoded 1D simulations outperform their unencoded counterparts in estimating local observables by 2-6% in the regimes considered in this work. As the logical degrees of freedom are spread across a block rather than being localized to a single physical qubit, the encoding relaxes the connectivity constraints of a logical computation. As a result, encoded simulations on a 2D square lattice are therefore possible at a 1.5 times lower overall circuit depth than a direct unencoded implementation. There, the encoding improvement grows with circuit depth, exceeding 200% at late times. Importantly, we compare the performance of encoded

²In the Schwinger model (quantum electrodynamics in 1+1D), for example, a spatial site maps to two qubits, e.g. Refs. [182, 184, 199, 225], suggesting the use of code blocks that implement two logical qubits.

and unencoded circuits *without any* error mitigation schemes that utilize auxiliary circuits to “learn” the noise [41, 42, 52, 184, 199, 213, 294, 553–555]. This choice isolates the improvements due to QED from those introduced by mitigation schemes. In both 1D and 2D, we find that infrequent error detection in systems mapped to many Iceberg-code logical qubits occupies a crossover regime between NISQ-era heuristics and fully FT computation.

Sections 7.2–7.4 constitute the main text of this chapter and are designed to be read sequentially. Section 7.2 describes the implementation of multiple blocks of the $[[4, 2, 2]]$ code on heavy-hex connectivity, as well as the details of our detector filtering technique, ORP. Section 7.3 presents results of encoded and unencoded simulations of the Ising model on `ibm_boston`. Section 7.4 concludes with a discussion of the results and the role of encoded computations in the early FT era of quantum simulation for scientific discovery. The appendices to this chapter contain information supporting the results in the main text. App. E.1 elaborates on ORP. App. E.2 explains the procedure used to determine optimal device layouts. Logical circuit scheduling is discussed in App. E.3. App. E.4 and App. E.5 provide the details and design choices behind our quantum and classical simulations, respectively. The remaining appendices provide further information and expand on these discussions.

7.2 Fault Tolerance in the $[[4, 2, 2]]$ Code

The key consideration behind the FT design of logical circuits is that the spread of errors should be controlled and minimized. Namely, such circuits should not introduce more errors than they can detect, and they should not propagate detectable errors into undetectable ones. This work implements a version of FT that is inspired by level-1 FT (1-FT) [504] by choosing “gadgets”, i.e. subcircuits that implement an encoded operation, that satisfy the following conditions:

1. A gadget with no errors during operation detects one input error and outputs at most one error.
2. A gadget with one error during operation and no input errors outputs at most one error.

For modest-sized circuit gadgets, such as those used in this work, these conditions can be verified by classical simulations. Two-qubit gates are the operations that typically suffer from the highest levels of noise and therefore set the dominant physical error rate. While faulty two-qubit gates can in principle produce

undetectable errors at the same rate as single-qubit errors, this work finds that using these conditions as criterion for FT is sufficient for practical error suppression.

The $[[4, 2, 2]]$ code is a Calderbank-Shor-Steane (CSS) code [556] that forms two logical qubits from four physical qubits, and is the smallest code in the family of Iceberg codes. It is a distance $d = 2$ code and so is capable of detecting single-qubit errors. The stabilizers of this code, shown on the left side of Fig. 7.1a), are given by

$$S_X = X_0 X_1 X_2 X_3, \quad S_Z = Z_0 Z_1 Z_2 Z_3. \quad (7.1)$$

These operators define the codespace as the joint $+1$ -eigenspace of S_X and S_Z , i.e., the even-parity sector of the four-qubit Hilbert space. The logical $|\overline{00}\rangle$ state is a four-qubit GHZ state, $|GHZ(4)\rangle = 1/\sqrt{2}(|0000\rangle + |1111\rangle)$, and the logical operators on the codespace are given by

$$\overline{X}_i = X_{i+1} X_3, \quad \overline{Z}_i = Z_0 Z_{i+1} \quad (7.2)$$

for $i = 0, 1$ and are shown on the right side of Fig. 7.1a). Any operator that contains only an even number of physical X s and Z s preserves the codespace, so errors of this form correspond to undetectable logical errors. Further, since operators on the logical space are equivalent up to multiplication by the stabilizers, the only errors this code fails to detect must contain two X s or two Z s.

7.2.1 The $[[4, 2, 2]]$ Code on Heavy-Hex Quantum Computers

To perform encoded computations on a quantum device, it is necessary to choose a hardware embedding that allows for efficient FT syndrome extraction and shallow implementation of logical operations. Six physical qubits are used to represent each $[[4, 2, 2]]$ code block: four data qubits d_0 - d_3 and two ancilla qubits a_x , a_z used for syndrome readout and logical operations. The embedding of a single code block into the heavy-hex connectivity requires one degree-3 vertex and two degree-2 vertices. Within this single-block topology, logical operations and syndrome extraction have different optimal arrangements of qubits. We call these the “compute” and “syndrome” configurations, which are shown in Fig. 7.1c) in the center of the circuits. The FT syndrome extraction circuit shown at the bottom right of Fig. 7.1c) measures S_X and S_Z simultaneously and requires all code blocks to have the same topology. During syndrome extraction, the ancillas a_x and a_z also act as mutual error flags, controlling the proliferation of “hook” errors that propagate from ancilla to data qubits [549, 557–559]. This circuit is derived from Ref. [560] and is verified to be 1-FT through classical

simulations. The circuits required to switch between the compute and syndrome configurations (and the gate overhead) are given in App. E.4.

To maintain scalability and locality of logical operations, we represent logical qubits with multiple $[[4, 2, 2]]$ code blocks. The optimal block placement is chosen in several steps which are explained in App. E.2. Avoiding noisy two-qubit gates and qubits with high measurement error rates, the maximum number of code blocks we find is 21, corresponding to 42 logical qubits. Logical qubits are labeled by their block and qubit number within the block:

$$(b, l_b) \text{ with } l_b \in \{0, 1\} . \quad (7.3)$$

Any given placement of code blocks can realize a number of distinct logical connectivity graphs. This work selects the placement that has the most square-grid-like logical connectivity graph from all possible 21-block placements, which is relevant for the 2D simulations discussed in Sec. 7.3. See App. E.2 for more details on the layout selection process. The chosen block placement on `ibm_boston` has four missing internal edges and is shown in Fig. 7.1b). This block placement has both directly adjacent blocks, and blocks that have a “bridge” qubit between them, shown in green.

7.2.2 Observable-Ranked Postselection (ORP)

Postselection is the task of identifying and discarding shots corrupted by errors. Syndrome measurements serve as imperfect proxies for identifying these affected shots, both due to the partially FT approach used in this work and because multiple errors can incorrectly return the logical state into the codespace. As a result, errors affecting logical information will not always trigger a syndrome. Further, a logical observable will be unaffected by errors that are causally disconnected from it or otherwise commute with it. We introduce ORP to sidestep the exponential shot loss associated with complete postselection.³ This method ranks error detection events by their probability of causing a logical error and uses this ranking to iteratively clean the measured dataset.

Active reset of the ancilla qubits between rounds consumes device coherence time, and so resets are omitted in this work. Each measured outcome therefore reflects the accumulated parity of syndrome

³The form of the exponential shot loss is a constant plus an exponential in the number of rounds of stabilizer measurements. The constant is determined by the ratio of the size of the codespace to the physical-qubit space. It is typically small enough to be neglected for practical ensemble sizes.

measurements in the computation up to that point [561]. The measurement outcome for syndrome type- g on block b and round r is denoted as $s_j = s_{(g,b,r)}$. These are combined into “detectors” by comparing each stabilizer’s outcome s_j to its value in a previous round,

$$v_{j+1} = s_{j-1} \oplus s_{j+1} , \quad (7.4)$$

where $j \pm 1$ is shorthand for $(g, b, r \pm 1)$ and $\oplus$ is the XOR operation implementing modulo-2 addition. The definition (7.4) implies a detector only fires ($v_j = 1$) when this comparison is violated. Crucially, in the absence of errors, all detectors will take the deterministic value $v_j = 0$.

A detector v_j ’s impact on a given observable $\overline{O}$ is inferred by conditioning $\langle \overline{O} \rangle$ on the value of v_j by

$$\Delta_j = \left| \langle \overline{O} \rangle_{v_j=0} - \langle \overline{O} \rangle_{v_j=1} \right| , \quad (7.5)$$

which measures the correlation between detection events and logical errors that bias $\langle \overline{O} \rangle$. For the j th detector, Δ_j can be interpreted as quantifying the effective probability that, given a detection event, a logical error occurred [512, 521, 562].⁴ Detectors are then ranked by Δ_j , and ORP uses this ranking to postselect on detection events with high Δ_j . A cut at level k keeps only shots in which none of the top- k detectors fire, resulting in a sequence of increasingly strict cuts with smaller ensemble sizes and progressively cleaner estimates of $\langle \overline{O} \rangle_k$. Each time the cutoff k is increased, the proportion of shots containing harmful errors relative to the remaining ensemble size decreases, and so the expectation value of a logical observable progressively approaches the noiseless value until it plateaus at a level set by the undetectable errors. However, the ensemble size also decreases with increasing k , inflating the statistical uncertainty on $\langle \overline{O} \rangle_k$. The optimal cut level therefore balances the bias from residual device noise against the variance from shot loss.

Since the residual bias is bounded below by the undetectable errors, once the shots carrying the dominant detectable errors have been removed $\langle \overline{O} \rangle_k$ stops moving and its distribution is consistent with statistical fluctuations alone. We identify plateaus in k over which this criterion holds, indicating that the estimate has converged. The optimal cut level k^* is chosen to be at the lower edge of the most stable plateau, and ORP uses this value to compute $\langle \overline{O} \rangle_{k^*}$. Technical details of ORP and the plateau-finding algorithm are given in App. E.1 and alternatives considered are discussed in App. E.6.

⁴For a local observable $\overline{O}$, the detectors selected by ORP are found to coincide with the backwards lightcone of $\overline{O}$, such that detectors causally disconnected from $\overline{O}$ have values of Δ_j parametrically smaller than those within the lightcone. This is discussed in App. E.6.5.




Figure 7.1: *Error Detection in Quantum Simulations on Heavy-Hex Topology.* a) The stabilizers and logical operators that act on data qubits d_0 - d_3 in the $[[4, 2, 2]]$ error-detecting code. The Z -stabilizer is shown in red and the X -stabilizer is shown in blue. Logical operators are grouped into commuting sets represented by the color. b) The placement of code blocks (grey) onto the `ibm_boston` quantum processor. Each code block has four data qubits (black) and two ancillas for Z - and X -stabilizer extraction (red and blue respectively). Bridge qubits between blocks used for routing are shown in green. c) Circuit operation for encoded quantum simulations. Dashed outlines indicate nFT operations, while solid outlines indicate FT gadgets. Each code block starts in the “compute” configuration (shown in the left box in the center). Within each code block, the logical $|00\rangle$ state is prepared with a nFT circuit. Non-FT logical operations (orange boxes) are performed in the compute configuration, example circuits of which are shown at the bottom. The blue boxes represent FT syndrome extraction done in the “syndrome” configuration within each block (right two boxes), which is implemented with the circuit shown in the bottom. Applications of logical operations and syndrome extraction are repeated R times in any given computation.

7.3 Error-Detected Quantum Simulations of The Ising Model

The rich collective behavior emerging from simple nearest-neighbor spin interactions in the Ising model makes it an ideal platform for real-time quantum simulations. Recent work has used it to examine the dynamics of thermalization [563], string breaking/confinement [230,564], scattering [243,565], and Majorana edge modes [566]. This work focuses on simulations of false-vacuum decay under quench dynamics in the MFIM in 1D and 2D. False-vacuum decay has been well studied in both 1D and 2D using classical [567–577] and quantum methods [578–581]. This work extends the tensor network study of Ref. [571] to nonzero longitudinal field. The MFIM Hamiltonian with open boundary conditions (OBCs) is

$$\bar{H} = -J \sum_{\langle ij \rangle} \bar{Z}_i \bar{Z}_j - \sum_i (g_x \bar{X}_i + g_z \bar{Z}_i) . \quad (7.6)$$

The sum over $\langle ij \rangle$ denotes nearest-neighbor coupling in 1D and couples adjacent logical qubits on a 2D lattice.

False-vacuum decay phenomenology models the nucleation of true-vacuum bubbles from a false-vacuum background and their subsequent evolution [582,583]. This work studies the quench dynamics of product-state bubbles, i.e., the evolution of bubbles after their nucleation. Our simulations begin in a product state $|\psi(t=0)\rangle = \prod_{i \in B} \bar{X}_i |\bar{1}\rangle^{\otimes N}$ on N qubits, where B is the initial true-vacuum bubble region. Recent work has predicted that bubbles melt for $|g_z|/|g_x| < 1$ and remain localized for $|g_z|/|g_x| > 1$ in the limit of large J [584], respectively indicating the presence and absence of false-vacuum decay. We work with the couplings $J = 1$, $g_x = 0.75$, and identify a bubble-melting regime at $g_z = 0.5$ and a localized regime at $g_z = 1.5$, where the bubble's extent is fixed by the quench dynamics. This phenomenon is driven by the linear potential between bubble walls that causes the momentum of excitations to wrap around the Brillouin zone, resulting in bubble oscillations known as Bloch oscillations [585]. In 1D, the oscillation of the bubble width (equivalently, Stark localization of the interface [567,576]) is predicted to have an amplitude set by the ratio $r_{\max} \sim |g_x|/|g_z|$.

In quantum simulations using the encoding discussed in the previous section, logical operations between qubits $(b, 0)$ and $(b, 1)$ within a block b are straightforwardly implemented in the compute configuration, shown in Fig. 7.1c). Depending on how adjacent blocks are placed relative to one another, some logical inter-block gates require routing through ancilla and bridge qubits. These inter-block gates constitute the bulk of the depth in our time evolution circuits. The MFIM Hamiltonian (7.6) is realized by the operations




Figure 7.2: *Encoded Simulations of Dynamics in the 1D Ising Model.* a) The realized logical connectivity grid by the block placement in Fig. 7.1b). Logical qubits (b, l_b) (grey) are joined by intra-block gates (hatched orange) and inter-block gates (orange). The black outline traces the gates used to simulate a 1D chain. b) The energy density $\langle \bar{H}_n \rangle$ computed from MPS simulations as a function of time t and logical qubit $n = (b, l_b)$ for $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). A size-three true-vacuum bubble located at the center is evolved with $n_T = 500$ Trotter steps of size $\delta t = 0.1$. The black rectangle shows the t for which device results are compared in c). c) Comparison of encoded (blue) and unencoded (grey) magnetization $\langle \bar{Z}_n \rangle$ from `ibm.boston` to MPS simulations (black) at $t = 8$, corresponding to 16 Trotter steps of $\delta t = 0.5$ for $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). The encoded simulations shown have $R = 1$ syndrome extraction rounds. d) The cumulative distribution function (CDF) of the per-qubit, per-time absolute error $|\langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}}|$ for each value of g_z . The median error (dotted vertical lines) and CDFs corresponding to one layer ($R = 1$) of syndrome measurements are shown compared to unencoded results (solid and dashed lines respectively). Results up to $t = 4$ are used, where an encoding improvement is observed. e) The signal-survival factor α , defined in Eq. (7.7), as a function of t for various R (colored lines) and each value of g_z compared to unencoded results (dashed grey line). The shading for $t = 4-8$ represents times excluded for the CDFs shown in d). The insets show the residual error σ in the fit. The statistical uncertainty is determined through bootstrap resampling.




Figure 7.3: *Encoded Simulations of the 2D Ising Model.* a) The magnetization $\langle \bar{Z}_n \rangle$ of a 3×3 true-vacuum bubble as a function of time t for $g_z = 0.5$ (top row) and $g_z = 1.5$ (bottom row), computed from MPS simulations with $\delta t = 0.1$. The black solid and dashed rectangles show cross-sections of the grid at $t = 4$ for which data is plotted in b). b) Encoded (blue) and unencoded (grey) results from `ibm.boston` compared to MPS results (black) for the horizontal (solid) and vertical (dashed) cross-sections outlined in a), as a function of logical qubit number $n = (b, l_b)$. The encoded data has $R = 1$ syndrome measurement rounds. c) Cumulative distribution functions (CDFs) of the per-qubit, per-time absolute error $|\langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}}|$ for both values of g_z . The median error and CDFs corresponding to encoded runs with $R = 1$ syndrome measurements and unencoded runs are shown compared to unencoded results (solid and dashed lines respectively). d) The signal-survival factor α as a function of t for various R (colored lines) and each value of g_z , compared to unencoded results (dashed grey line). The insets show the residual error σ in the fit. The statistical uncertainty is determined through bootstrap resampling.

$\bar{Z}_i$, $\bar{X}_i$, $\bar{Z}_i \bar{Z}_j$ available on our logical connectivity graph.

7.3.1 Results from IBM's Quantum Computers

We simulate the dynamics of a 1D chain under the Hamiltonian in Eq. (7.6) using the logical qubit grid realized by the block placement shown in Fig. 7.1b). By retaining all intra-block connections and removing specific inter-block edges from the grid, we implement a chain of 42 logical qubits shown in Fig. 7.2a).

The expectation value of the energy density $\bar{H}_n = -J\bar{Z}_n\bar{Z}_{n+1} - (g_x\bar{X}_n + g_z\bar{Z}_n)$ as a function of time and logical qubit $n = (b, l_b)$ determined from MPS simulations is shown in Fig. 7.2b). For $g_z = 0.5$ the bubble oscillates about its center, while for $g_z = 1.5$ it remains localized. This behavior is expected from the Bloch-oscillation picture of App. E.9, in which $r_{\max}$ falls from well above the bubble size to comparable with it, and is consistent with the melting/localization transition reported in Ref. [584]. The magnetization $\langle\bar{Z}_n\rangle$ from encoded results determined through ORP and unencoded runs on `ibm_boston` are compared to MPS results for $t = 8$ in Fig. 7.2c). Deviations from MPS expectations are quantified through the absolute error $|\langle\bar{Z}\rangle_{\text{meas}} - \langle\bar{Z}\rangle_{\text{MPS}}|$. The cumulative distribution function (CDF) of the absolute errors over all qubits and $t \leq 4$ is shown in Fig. 7.2d) for both values of g_z . The largest encoding improvement in median absolute error is found to be $36.1\% \pm 3.3\%$ up to $t = 4$ for $g_z = 0.5$ and $R = 1$ (corresponding to a single syndrome extraction round at the end of the circuit). The median absolute error is lower for encoded runs and the improvement is consistent for both g_z . Further, their CDFs lie predominantly to the left of unencoded CDFs, overlapping only at large absolute errors. Error detection is found to concentrate the bulk of the distribution at small errors but retains a tail of poorly-performing samples. This is consistent with most qubits and gates operating just below the pseudothreshold, with several outliers.

To quantify how faithfully the encoded circuit reproduces the ideal dynamics independent of the overall signal magnitude, $\langle\bar{Z}\rangle_{\text{meas}}$ is fit to $\langle\bar{Z}\rangle_{\text{MPS}}$ with a line through the origin. The slope α is the signal-survival factor and σ is the root-mean-square of the fit residuals

$$\langle\bar{Z}\rangle_{\text{meas}} = \alpha\langle\bar{Z}\rangle_{\text{MPS}} + \epsilon, \quad \sigma = \sqrt{\frac{1}{n} \sum_i \epsilon_i^2}. \quad (7.7)$$

Here the sums run over all n samples of the observable (for a given qubit and time). The signal-survival $\alpha = 1$ indicates perfect retention, while $\alpha < 1$ implies signal loss due to noise, so higher values of α at comparable σ indicate more faithful results. Figure 7.2e) shows α and σ as a function of t for both couplings and various numbers of syndrome extraction rounds R . Here R denotes the application of up to R syndrome extraction rounds, with at most one syndrome round per Trotter step. Encoded runs for all values of R are found to outperform their unencoded counterparts up to intermediate times, after which only small- R encoded results for $g_z = 1.5$ show an improvement in α . As expected, α trends lower with t due to the accumulation of undetected errors. The typical encoding improvement in the localized regime compared to the melting regime over all t is attributed to a narrower backwards lightcone at $g_z = 1.5$, which restricts error propagation to a smaller spatial region and lets the encoding detect a greater fraction of errors. Several points in the melting

regime break the monotonic decay of α ; this is again attributed to the faster propagation of errors at $g_z = 0.5$. Since α is similar for unencoded runs at $g_z = 0.5$ and 1.5, this suggests that encoding is more effective when dynamics are slow and error propagation is limited.

An unforeseen feature of the encoding improvement is its dependence on R , represented by the colors in Fig. 7.2e). Early-time results show a modest advantage of increased syndrome extraction frequency (larger R), but beyond $t \sim 4$ this reverses and fewer extraction rounds perform better. This indicates that conserving coherence outweighs the benefit of additional syndrome extractions, and is due to the interplay between ORP and error behavior in the absence of ancilla resets. Without resets, an ancilla that detects an error remains in $|1\rangle$, and subsequent logical operations that couple to it inject effective errors that propagate through the system. ORP removes some of this contamination, but the degradation with extraction frequency reveals its limitation: error detection events that ORP does not choose leave coherent contamination in the retained shots. Notably, adding resets does not significantly improve α , indicating that the coherence-time cost of resets roughly cancels the benefit of eliminating syndrome-induced error propagation. See App. E.8 for a detailed analysis of this effect.

Two-dimensional lattice simulations are implemented by using all inter- and intra-block edges in the logical connectivity graph shown in Fig. 7.2a). Figure 7.3a) shows the magnetization $\langle \bar{Z}_n \rangle$ on the grid for several selected times determined from MPS simulations. Although the dynamics in 2D is found to be slower than in 1D, an initial 3×3 bubble is still found to spread for $g_z = 0.5$ and remain localized for $g_z = 1.5$.

Figure 7.3b) shows $\langle \bar{Z}_n \rangle$ for horizontal and vertical cross-sections on the grids in a) at $t = 4$ (solid and dashed outlines respectively). While unencoded results show significant decoherence toward $\langle \bar{Z}_n \rangle = 0$, encoded results computed with ORP are closer to MPS expectations. The encoding largely removes the geometric overhead associated with representing a square lattice on heavy-hex hardware, and only the depth overhead from Trotter step serialization is present. The retained geometric overhead is reflected in large unencoded circuit depths, and the encoding improvement is expected to be larger than in 1D simulations. The CDFs of the absolute error over all logical qubits and times are shown in Fig. 7.3c), with encoded CDFs lying entirely to the left of their unencoded counterparts. Further, the median absolute error for encoded runs is seen to be significantly smaller than unencoded runs. The largest encoding improvement in median absolute error is found to be $56.9\% \pm 0.8\%$ for $g_z = 0.5$ and $R = 1$. Both the unencoded and encoded distributions show higher median error in Fig. 7.3c) compared to Fig. 7.2d) due to increased circuit depth.

The encoded signal-survival factor α is found to exceed unencoded values for all times and all R with the




Figure 7.4: *Encoding Improvement in α* . Percent improvement in α of encoded over unencoded results as a function of time t for 1D and 2D simulations. Results with $R = 1$ syndrome extraction rounds for the melting regime ($g_z = 0.5$) are shown in shades of red, and the localized regime ($g_z = 1.5$) is shown in shades of blue. The shading represents statistical uncertainty determined through bootstrap resampling.

exception of $R = 16$ past $t \sim 4$ (Fig. 7.3d)). Similar to the CDFs, the α values are observed to decrease faster in 2D than in 1D. Further, the improvement from encoding is significantly larger than in 1D. App. E.8 shows that α drops significantly and the improvement disappears for $R = 4, 8, 16$ when postselection is removed. This indicates that the improvement in α seen in Fig. 7.3 is largely due to postselection as opposed to reduced gate depth. In contrast to 1D, the decay of α is observed to be monotonic, which we attribute to the slower proliferation of undetectable errors. Runs with $R = 1, 2$ are found to produce comparable results, indicating that the gain from more syndrome measurement rounds is balanced with the added errors from longer circuits.

Figure 7.4 shows the percentage improvement in α over time for $R = 1$ in 1D and 2D simulations. The improvement in the results from runs with encoding grows approximately linearly in 2D, with an average slope of $23.5 \pm 0.8\%$ in the melting regime and $22.8 \pm 0.5\%$ in the localized regime. Because each caught error prevents corruption over a lightcone volume scaling as $\mathcal{O}(t^2)$ in 2D compared to $\mathcal{O}(t)$ in 1D, detection has greater leverage in 2D and the encoding improvement grows with t . The greater than 200% improvement in 2D is also driven by rapidly decohering unencoded results, as is similarly seen in Fig. 7.3b). Dynamics at $g_z = 0.5$ and $g_z = 1.5$ are more similar in 2D than in 1D, explaining their close trajectories in Fig. 7.4.

Beyond $t \sim 4.5$, the melting regime in 1D shows a disadvantage from encoding because of the proliferation of errors enabled by fast dynamics. In 1D, the encoding improvement is approximately constant: for $g_z = 0.5$ an improvement is seen up to $t = 4.5$ (average improvement of $5.99 \pm 0.99\%$), and for $g_z = 1.5$ the average improvement is $2.5 \pm 0.23\%$ over all t . Additional analysis of these results is shown in App. E.8 and data, including acceptance rates, is given in App. E.11. Together, these results show that both the density of nFT gates and the dynamics being simulated affect the encoding improvement.

7.4 Discussion

The partially FT quantum simulations performed in this work using IBM’s quantum computers provide a concrete demonstration of the benefit of encoded error detection on contemporary hardware. Despite significant overhead associated with encoding and the shot loss resulting from postselection, results obtained using our scheme show improvements for both 1D and 1D simulations. These results reside in a crossover regime between NISQ heuristics and fully FT device operation [519], which shapes our circuit design. Our circuits combine FT syndrome extraction with nFT logical operations, resulting in jointly optimized circuit overhead and error resilience properties. Increased frequency of syndrome extraction is seen to improve results at early times, as expected from a fully-FT computation. Late-time performance instead benefits from the nFT reduction in overhead provided by sparser error detection. This balance reflects a central tradeoff in the crossover regime: more syndrome measurements aid filtering but consume some of the coherence time of the quantum computer. A similar tradeoff exists for the addition of ancilla resets following syndrome readout. Designed to operate inside this tradeoff, ORP postselects on a ranked subset of observable-relevant checks and recovers much of the benefit of full postselection without an exponential loss in statistics. The implementation of a QEC code would avoid the filtering machinery built to manage the postselection overhead. Our circuit design choices and error-detection scheme are tailored to the specific simulation and observables of interest. We expect this kind of customization to be a persistent feature of the crossover regime.

The potential for encoded advantage in quantum simulation is underscored by its sensitivity to device performance. The absence of an advantage on devices with slightly worse error rates ($\mathcal{O}(10^{-3})$) and coherence times ($\mathcal{O}(100\mu\text{s})$) than `ibm_boston` shows that even modest improvement in device characteristics is enough to cross the pseudothreshold. Since our figure of merit is local observable estimation, our results do not directly imply a favorable logical error-rate scaling. However, the improvement in observable estimation

suggests that state-of-the-art devices are in the vicinity of the pseudothreshold for encoding procedures with a small footprint. This reinforces the understanding that any improvement in the crossover regime is tied to specifics of code choice, overhead, and hardware operating point rather than being a generic feature of the encoding, and should become more robust and widespread as physical fidelities improve.

While the encoded improvements found in this work hold across dynamics and dimensionality, both encoded and unencoded results would likely benefit further from error mitigation. Studying FT constructions that remove leading-order errors, alongside mitigation techniques that suppress higher-order contributions, is a very promising direction and should be pursued. This includes both standard error mitigation techniques based on physical-noise learning [40–42, 52, 199, 213, 293, 294], as well as potential techniques operating directly on syndrome information. The incorporation of such techniques involves circuit and observable-dependent design choices, and the implications for the encoded advantage remain to be seen.

Beyond the direct advantage in observable estimation, encoding expands the versatility of logical computations where no shallow unencoded analog exists. It enables simulations whose direct unencoded implementation carries substantially higher overhead from connectivity constraints alone. The 2D lattice simulations in this work illustrate this: the realized connectivity graph is costly to embed directly on heavy-hex hardware. This approach extends to more complex logical connectivities, such as three-dimensional lattices required for lattice gauge theories, e.g. Refs. [59, 82, 295, 390, 434] and the Fermi-Hubbard model, e.g. Ref. [586]. As these constraints relax, we expect a more diverse range of simulations to become feasible in the crossover regime, where the dominant constraints shift to coherence time and circuit depth. The benefit of encoding is therefore not limited to solely improving the accuracy of a given simulation, but also extends to making entirely new classes of simulations accessible.

Chapter 8

CONCLUDING REMARKS

Nature, when examined at its smallest scales, is governed by quantum mechanics, and the phenomena it produces are as rich as they are difficult to predict. For some of these phenomena, perturbative methods provide systematic and controlled understanding; for a far larger set, no such expansion exists and the collective behavior of many interacting degrees of freedom must be confronted directly. It is precisely these highly-correlated regimes, the low-energy dynamics of QCD, the physics of strongly-interacting electrons, and the non-equilibrium evolution of quantum many-body systems, that stand to benefit most immediately from quantum simulation. The premise of this thesis, laid out in Chapter 1, is that such systems are most naturally studied on computational hardware that shares their underlying structure: nature as a computer and the computer as a snapshot of nature.

The work presented in this dissertation contributes to different stages of the quantum simulation pipeline. All are unified by the use of physics-informed methods, exploiting the symmetries, structure, and approximate descriptions of the physical system being simulated, rather than treating the simulation as an oracle. Chapters 2, 3, and 4 focused on implementing time evolution and quantifying the resulting errors. Chapter 2 developed a framework for optimizing algorithmic errors in analog quantum simulations, showing that the interplay between Trotterization and the native resource dynamics of an analog device can be tuned to minimize the total simulation error. Chapter 3 used a Fully Gauge-Fixed formulation of the $SU(2)$ Hamiltonian and simulated its dynamics, demonstrating its effectiveness and efficiency at a range of different gauge couplings and culminating in real-time simulations on superconducting hardware. Chapter 4 turned to the quantum complexity of these simulations themselves, studying the dynamics of non-stabilizerness. There it was shown that the low-energy subspace converges exponentially to its continuum non-stabilizerness with the number of qubits per plaquette, and that aggressive truncations of the Hamiltonian leave the complexity dynamics of a quench essentially unchanged, permitting deeper time evolution than naive circuit constructions would allow.

Chapter 4 also marks the transition from implementing time evolution to addressing the second half of the pipeline: measurement. Extracting a quantity as fragile and globally sensitive as the stabilizer Rényi

entropy from modern quantum devices is foremost a noise problem, and the chapter’s central contributions are measurement-theoretic in nature: a rewriting of randomized-measurement estimators as Clifford-averaged operator expectation values and the Bit String Decoherence Renormalization (BSDR) scheme, which mitigates errors in entire bitstring distributions by exploiting the hierarchical nature of smooth wavefunctions. These tools turned raw device data into quantitative tracking of complexity growth after a quench. Chapters 4, 5, and 6 together develop methods for efficiently extracting complicated quantities from quantum simulators. Chapter 5 introduced classical shadow protocols tailored to systems with local symmetries, using the duality between $\mathbb{Z}_2$ lattice gauge theory and the Ising model to obtain rigorous, and in some cases exponential, reductions in the sample complexity of estimating gauge-invariant observables. Chapter 6 addressed the extraction of entanglement structure itself, showing that an expansion of the many-body Wigner functional around Gaussian states allows the entanglement spectrum of a weakly non-Gaussian state to be reconstructed from experimentally accessible few-body correlation functions.

An underlying theme throughout this thesis has been complexity, a quantity that admits several distinct but complementary characterizations. The first is the computer-science perspective of complexity classes, which asks which problems can be solved or verified efficiently, and which places Hamiltonian simulation squarely inside **BQP**. The second is operational, i.e. how many elementary steps (circuit depth, gate count, measurement samples) must be performed to guarantee a solution to a given accuracy. The error bounds of Chapter 2, the resource estimates of Chapter 3, and the sample complexity guarantees of Chapter 5 all are of this flavor. The third characterization is intrinsic to the quantum states themselves, using measures that uniquely quantify the resources a state contains. Entanglement is one such measure, non-stabilizerness is another, and non-Gaussianity is a third. These are the “Big Three”, the three axes, of quantum complexity as it pertains to quantum many-body systems and quantum phenomena. The first two have received much of the spotlight, as they respectively bound the performance of tensor network and stabilizer-based classical simulation methods. It is likely, however, that non-Gaussianity will play an increasingly important role in the coming years, as fermionic quantum simulators mature and as the boundary between efficiently simulable (Gaussian) and genuinely complex dynamics is probed in experiments of the type proposed in Chapter 6.

These three notions of complexity are not independent, and the results of Chapter 6 offer a concrete bridge between them. Level repulsion in the entanglement spectrum, i.e. the emergence of Wigner-Dyson statistics in the gap ratios of the entanglement Hamiltonian, appears once the state has developed sufficient complexity across all three axes, and it is a diagnostic of quantum chaos and the onset of thermalization. In other words,

resource measures and a physical phenomenon switch on together. Counter examples to this exist [478], but are contrived and do not capture the behavior of quantum many-body systems generally. The fact that the difficulty of classically describing a state and the physics of its thermalization furnish two complementary descriptions bolsters the perspective of nature as a computer: the complexity of a physical process is not incidental to the physics, but central to it.

Quantum simulations will change substantially over the next several years, and the first steps in that transformation will come from synthesizing fault-tolerant constructions with the techniques we already have, rather than waiting for fully error-corrected machines to arrive. The final chapter of this thesis, Chapter 7, makes first steps in this direction by performing Hamiltonian simulations on logical qubits encoded in many interacting blocks of the $[[4, 2, 2]]$ error detection code. The philosophy is deliberately pragmatic: rather than demanding full fault tolerance, the computation mixes fault-tolerant components and shallow syndrome extraction circuits with carefully constructed non-fault-tolerant logical rotations, keeping the leading-order contribution to the logical error rate small. With these ingredients, encoded simulations of Ising dynamics in one dimension outperform their unencoded counterparts despite a several-fold increase in circuit depth, and without any error mitigation applied to either. These results indicate that today’s devices sit at the pseudothreshold, and that the boundary between the NISQ era and the Fault-Tolerant era will not be a sharp line but a crossover, defined by practical approaches to simulation.

The progress that the field of quantum simulation has witnessed over the course of completing this dissertation has been remarkable. At its outset, the field was coming off of Google’s quantum supremacy claim, devices offered a few tens of noisy qubits, and it was genuinely unclear how quantum technologies would find broad usage, let alone when. Today, “if” is no longer seriously in question; “When” remains debated, but the conversation is shaped largely by the community’s focus on asymptotic advantages à la Shor. These milestones are important, but they measure the field against its most distant horizon rather than its nearest opportunities.

The regime of non-asymptotic utility, practical advantage on problems of finite, scientifically relevant size, remains comparatively unexplored. There is a lack of widespread, hands-on usage of quantum devices and very few researchers routinely perform practical simulations of the type discussed in this dissertation. A useful comparison is to the recent trajectory of artificial intelligence. In many ways these two technologies, quantum and AI, find themselves in opposite positions. Large Language Models (LLMs) work in practice but lack many underlying theoretical guarantees, while quantum computing works in theory but noisy devices

have limited its practice. No one anticipated how powerful LLMs would become until practitioners simply scaled them and found that things kept improving; the decisive ingredient was not a theorem but rather empirical engagement with real systems. Quantum simulation has not yet had this moment, partially because comparatively little effort is directed at making calculations practical rather than provable. The experience gathered in this thesis suggests that such engagement is rewarded and that one often finds that the situation on real devices is considerably better than the literature would suggest. The physics-informed methods developed here are offered in that spirit: as tools for closing the gap between what is proven possible asymptotically and what can be computed, today, about the quantum many-body systems that nature has given us.

Appendix A

APPENDIX TO CHAPTER 2

A.1 Constant-Field Method Error Rate Derivation

For the choice of $\vec{B}(t)$ given by (2.11), time evolution in the interaction picture is given by

$$U_{\text{Ising}}(t, 0) = \exp_{\mathcal{T}} \left\{ -i \int_0^t dt' H_{\text{int}}(t') \right\},$$

where $H_{\text{int}}(t) = \frac{1}{2} \sum_{i \neq j} J_{ij} Z_i(t) Z_j(t)$, $Z_i(t) = 2\vec{e}(t) \cdot \vec{S}_i$, and

$$\vec{e}(t) = \frac{1}{3} \begin{pmatrix} 2\sqrt{2} \sin^2 \left(\frac{\Omega t}{2} \right) \\ -\sqrt{6} \sin \Omega t \\ 1 + 2 \cos \Omega t \end{pmatrix}. \quad (\text{A.1})$$

The error rate at Floquet periods $\tau_{\mathcal{C}_1} = \frac{2\pi}{\Omega}$ (i.e., the period of $\vec{e}(t)$) is given by

$$ER_{\mathcal{C}_1}(\tau_{\mathcal{C}_1}) = \frac{1}{\tau_{\mathcal{C}_1}} \left\| \exp_{\mathcal{T}} \left\{ -i \int_0^{\tau_{\mathcal{C}_1}} dt H_{\text{int}}(t) \right\} - \exp \left\{ -i \frac{\tau_{\mathcal{C}_1}}{3} H_{XXX} \right\} \right\|. \quad (\text{A.2})$$

To get the tightest possible bounds on $ER_{\mathcal{C}_1}$ with minimal application of triangle inequalities, the two propagators are organized in powers of Ω (equivalently powers of $\tau_{\mathcal{C}_1}$). For the target evolution,

$$\exp \left\{ -i \frac{\tau_{\mathcal{C}_1}}{3} H_{XXX} \right\} = 1 - i \frac{\tau_{\mathcal{C}_1}}{3} H_{XXX} - \frac{\tau_{\mathcal{C}_1}^2}{18} H_{XXX}^2 + O(\tau_{\mathcal{C}_1}^3).$$

Working up to second-order in a Magnus expansion, the Floquet evolution gives:

$$\exp_{\mathcal{T}} \left\{ -i \int_0^{\tau_{\mathcal{C}_1}} dt H_{\text{int}}(t) \right\} = e^{-i \sum_{k=1}^{\infty} \Omega_k(\tau_{\mathcal{C}_1})} \quad (\text{A.3})$$

$$= 1 - \left(i\Omega_1(\tau_{\mathcal{C}_1}) - i\Omega_2(\tau_{\mathcal{C}_1}) + \frac{1}{2}\Omega_1^2(\tau_{\mathcal{C}_1}) \right) + O(\tau_{\mathcal{C}_1}^3), \quad (\text{A.4})$$

where

$$\Omega_1(\tau_{\mathcal{C}_1}) = \int_0^{\tau_{\mathcal{C}_1}} dt H_{\text{int}}(t) = \frac{\tau_{\mathcal{C}_1}}{3} H_{XXX}, \quad (\text{A.5})$$

$$\Omega_2(\tau_{C_1}) = \frac{1}{2} \int_0^{\tau_{C_1}} dt_1 \int_0^{t_1} dt_2 [H_{\text{int}}(t_1), H_{\text{int}}(t_2)]. \quad (\text{A.6})$$

Evaluating Eq. (A.6) and defining $\vec{e}_i = \vec{e}(t_i)$

$$\begin{aligned} \Omega_2(\tau_{C_1}) = & -i \int_0^{\tau_{C_1}} dt_1 \int_0^{t_1} dt_2 \left(8 \sum_{i \neq j \neq k} J_{ij} J_{jk} (\vec{e}_1 \cdot \vec{S}_i) ((\vec{e}_1 \times \vec{e}_2) \cdot \vec{S}_j) (\vec{e}_2 \cdot \vec{S}_k) \right. \\ & \left. + 2 \sum_{i \neq j} J_{ij}^2 (\vec{e}_1 \cdot \vec{e}_2) (\vec{e}_1 \times \vec{e}_2) \cdot \vec{S}_j \right), \end{aligned} \quad (\text{A.7})$$

and the following definitions to simplify expressions,

$$\Lambda_{\alpha\beta\gamma}(t_1, t_2) = (\vec{e}_1)_\alpha (\vec{e}_1 \times \vec{e}_2)_\beta (\vec{e}_2)_\gamma, \quad (\text{A.8})$$

$$\tilde{\Lambda}_\alpha(t_1, t_2) = (\vec{e}_1 \cdot \vec{e}_2) (\vec{e}_1 \times \vec{e}_2)_\alpha, \quad (\text{A.9})$$

the Ω_1 terms exactly cancel the target Heisenberg evolution, giving the error rate at leading-order in τ_{C_1} as:

$$ER_{C_1}(\tau_{C_1}) = \frac{1}{\tau_{C_1}} \|\Omega_2(\tau_{C_1})\| + O(\tau_{C_1}^3) \quad (\text{A.10})$$

$$\begin{aligned} & \leq \frac{1}{2\pi\Omega} \left(\left\| \int_0^1 \int_0^{t_1} dt_1 dt_2 \Lambda(t_1, t_2) \right\| \left(\sum_{i \neq j \neq k} J_{ij} J_{jk} \right) \right. \\ & \quad \left. + \left\| \int_0^1 \int_0^{t_1} dt_1 dt_2 \tilde{\Lambda}(t_1, t_2) \right\| \left(\sum_{i \neq j} J_{ij}^2 \right) \right) \end{aligned} \quad (\text{A.11})$$

$$= \frac{2\epsilon}{\pi} \left(\frac{\sqrt{142 + 24\pi^2}}{18} \left(\sum_{i \neq j \neq k} J_{ij} J_{jk} \right) + \frac{1}{3} \left(\sum_{i \neq j} J_{ij}^2 \right) \right) \quad (\text{A.12})$$

$$= O(\epsilon N). \quad (\text{A.13})$$

A.2 Derivation of the $\frac{\pi}{2}$ -Pulse Error

As noted in the main body of the paper, the global rotation gates cannot be perfectly implemented due to the persistent Ising interaction. The deviation from the ideal behavior is found using a Magnus expansion, following Appendix A of Ref. [122]:

$$\Upsilon_X^\pm(\epsilon) = R_X \left(\mp \frac{\pi}{2} \right) R_X^\pm(\epsilon) = \exp \left(-i \sum_{k=1} \chi_k^\pm \epsilon^k \right), \quad (\text{A.14})$$

$$\Upsilon_Y^\pm(\epsilon) = R_Y \left(\mp \frac{\pi}{2} \right) R_Y^\pm(\epsilon) = \exp \left(-i \sum_{k=1} \nu_k^\pm \epsilon^k \right). \quad (\text{A.15})$$

To derive the error rate bounds, it is only necessary to work to $O(\epsilon^2)$ in $\Upsilon_{X,Y}^\pm$. The expressions for $\chi_{1,2}^\pm$ and $\nu_{1,2}^\pm$ are:

$$\begin{aligned}\chi_1^\pm &= \int_0^1 dt e^{\pm \frac{i\pi t}{4} \sum_i X_i} \left(\sum_{i < j} J_{ij} Z_i Z_j \right) e^{\mp \frac{i\pi t}{4} \sum_i X_i} \\ &= \sum_{i < j} J_{ij} \left(\frac{Y_i Y_j + Z_i Z_j}{2} \pm \frac{Y_i Z_j + Z_i Y_j}{\pi} \right),\end{aligned}\tag{A.16}$$

$$\nu_1^\pm = \sum_{i < j} J_{ij} \left(\frac{X_i X_j + Z_i Z_j}{2} \mp \frac{X_i Z_j + Z_i X_j}{\pi} \right),\tag{A.17}$$

and

$$\begin{aligned}\chi_2^\pm &= \frac{-i}{2} \int_0^1 \int_0^{t_1} dt_1 dt_2 \left[e^{\pm \frac{i\pi t_1}{4} \sum_i X_i} \left(\sum_{i < j} J_{ij} Z_i Z_j \right) e^{\mp \frac{i\pi t_1}{4} \sum_i X_i}, e^{\pm \frac{i\pi t_2}{4} \sum_i X_i} \left(\sum_{i < j} J_{ij} Z_i Z_j \right) e^{\mp \frac{i\pi t_2}{4} \sum_i X_i} \right] \\ &= -i \sum_{i < j, m < n} J_{ij} J_{mn} \left(\frac{1}{\pi^2} [Y_i Y_j, Z_m Z_n] \pm \frac{1}{8\pi} [Y_i Y_j - Z_i Z_j, Z_m Y_n + Y_m Z_n] \right) \\ &= \frac{1}{\pi} \sum_{i, m \neq j} J_{ij} J_{jm} \frac{Y_i X_j Z_m + X_j Z_m Y_i}{\pi} \pm \frac{Y_i Y_m X_j + Z_i X_j Z_m + Y_m X_j Y_i + X_j Z_m Z_i}{8},\end{aligned}\tag{A.18}$$

$$\nu_2^\pm = -\frac{1}{\pi} \sum_{i, m \neq j} J_{ij} J_{jm} \frac{X_i Y_j Z_m + Y_j Z_m X_i}{\pi} \mp \frac{X_i X_m Y_j + Z_i Y_j Z_m + X_m Y_j X_i + Y_j Z_m Z_i}{8}.\tag{A.19}$$

A.3 First-Order Product Formula Error

As an example, the analysis of Section 2.3 is applied to the product formulas $\mathcal{S}_{1/2}$ and $\mathcal{S}_1$, with the specific choice $H_{\text{tar}} = H_{XX}$. The result of this analysis will be Eqs. (2.21) and (2.22). The error rate formulas are

$$ER_{\mathcal{S}_{1/2}} \leq \frac{1}{\tau_{\mathcal{S}_{1/2}}(t, \epsilon)} \left(t^2 \left\| \mathcal{E}_{\mathcal{S}_{1/2}; t^2} \right\| + \epsilon \left\| \mathcal{E}_{\mathcal{S}_{1/2}; \epsilon} \right\| \right),\tag{A.20}$$

$$ER_{\mathcal{S}_1} \leq \frac{1}{\tau_{\mathcal{S}_1}(t, \epsilon)} \left(t^2 \left\| \mathcal{E}_{\mathcal{S}_1; t^2} \right\| + \epsilon^2 \left\| \mathcal{E}_{\mathcal{S}_1; \epsilon^2} \right\| + t\epsilon \left\| \mathcal{E}_{\mathcal{S}_1; t\epsilon} \right\| \right).\tag{A.21}$$

The choices of t that optimize the error rates are

$$t_{\mathcal{S}_{1/2}*} = \epsilon^{\frac{1}{2}} \sqrt{\frac{\left\| \mathcal{E}_{\mathcal{S}_{1/2}; \epsilon} \right\|}{\left\| \mathcal{E}_{\mathcal{S}_{1/2}; t^2} \right\|}} + O(\epsilon),\tag{A.22}$$

$$t_{\mathcal{S}_1*} = \epsilon \left(\sqrt{4 - \frac{2 \left\| \mathcal{E}_{\mathcal{S}_1; t\epsilon} \right\| - \left\| \mathcal{E}_{\mathcal{S}_1; \epsilon^2} \right\|}{\left\| \mathcal{E}_{\mathcal{S}_1; t^2} \right\|}} - 2 \right).\tag{A.23}$$

The corresponding error rate bounds (valid for $t_{\mathcal{S}_{1/2}*}, t_{\mathcal{S}_1*} > 0$) are given by

$$ER_{\mathcal{S}_{1/2}} \leq 2\sqrt{\|\mathcal{E}_{\mathcal{S}_{1/2};t^2}\| \|\mathcal{E}_{\mathcal{S}_{1/2};\epsilon}\|} \epsilon^{\frac{1}{2}} + O(\epsilon), \quad (\text{A.24})$$

$$ER_{\mathcal{S}_1} \leq \left(2\sqrt{\|\mathcal{E}_{\mathcal{S}_1;t^2}\| (4\|\mathcal{E}_{\mathcal{S}_1;t^2}\| + \|\mathcal{E}_{\mathcal{S}_1;\epsilon^2}\| - 2\|\mathcal{E}_{\mathcal{S}_1;t\epsilon}\|)} + \|\mathcal{E}_{\mathcal{S}_1;t\epsilon}\| - 4\|\mathcal{E}_{\mathcal{S}_1;t^2}\| \right) \epsilon. \quad (\text{A.25})$$

An important point to consider is that Eq. (A.23) suggests that $t_{\mathcal{S}_1*}$ may be negative, while in both Eq. (A.20) and (A.21), $t_{\mathcal{S}_{1/2}*}$ and $t_{\mathcal{S}_1*}$ were assumed to be positive numbers. This apparent contradiction simply reflects the fact that this is an optimization problem whose solution is on the boundary $t_{\mathcal{S}_1*} = 0$. It is now clear which error terms must be calculated to recover the optimal step sizes and the error rates, which is done below.

A.3.1 First-Order Error in ϵ

At leading-order in ϵ and t , Eqs. (2.6) and (2.8) take the form

$$\begin{aligned} \mathcal{S}_{1/2} = & \mathbb{1} - i \sum_{i < j} J_{ij} (t_x X_i X_j + t_y Y_i Y_j + t_z Z_i Z_j) \\ & - i\epsilon \sum_{i < j} J_{ij} \left(X_i X_j + Y_i Y_j + 2Z_i Z_j + \frac{2}{\pi} (Y_i Z_j + Z_i Y_j - X_i Z_j - Z_i X_j) \right), \end{aligned} \quad (\text{A.26})$$

$$\mathcal{S}_1 = R_Z(\pi) \left(\mathbb{1} - i \sum_{i < j} J_{ij} (t_x X_i X_j + t_y Y_i Y_j + t_z Z_i Z_j) - i\epsilon \sum_{i < j} J_{ij} (X_i X_j + Y_i Y_j + 2Z_i Z_j) \right). \quad (\text{A.27})$$

The leading-order target evolution is $\mathbb{1} - i\tau_U(t, \epsilon)H_{XXX}$, and matching these evolutions requires $t_x, t_y = \epsilon + t$ and $t_z = t$. These choices imply that $\tau_U(t, \epsilon) = 2\epsilon + t$. With these choices, the leading-order Heisenberg evolution is the same for $\mathcal{S}_{1/2}$ and $\mathcal{S}_1$, implying that $\tau_{\mathcal{S}_{1/2}} = \tau_{\mathcal{S}_1}$. The individual terms in the expansion (2.13) are

$$\mathcal{S}_{1/2;t} = -i \sum_{i < j} J_{ij} (X_i X_j + Y_i Y_j + Z_i Z_j), \quad (\text{A.28})$$

$$\mathcal{S}_{1;t} = R_Z(\pi) \left(-i \sum_{i < j} J_{ij} (X_i X_j + Y_i Y_j + Z_i Z_j) \right), \quad (\text{A.29})$$

$$\mathcal{S}_{1/2;\epsilon} = -2i \sum_{i < j} J_{ij} \left(X_i X_j + Y_i Y_j + Z_i Z_j + \frac{Y_i Z_j + Z_i Y_j - X_i Z_j - Z_i X_j}{\pi} \right), \quad (\text{A.30})$$

$$\mathcal{S}_{1;\epsilon} = R_Z(\pi) \left(-2i \sum_{i < j} J_{ij} (X_i X_j + Y_i Y_j + Z_i Z_j) \right). \quad (\text{A.31})$$

The error terms that contribute to the error rates at this order are

$$\mathcal{E}_{\mathcal{S}_1;t}, \mathcal{E}_{\mathcal{S}_1;\epsilon} = 0, \quad (\text{A.32})$$

$$\mathcal{E}_{\mathcal{S}_{1/2};t} = 0, \quad (\text{A.33})$$

$$\mathcal{E}_{\mathcal{S}_{1/2};\epsilon} = -\frac{2i}{\pi} \sum_{i < j} J_{ij} ((Y_i - X_i)Z_j + Z_i(Y_j - X_j)). \quad (\text{A.34})$$

There is $O(\epsilon)$ error appearing for $\mathcal{S}_{1/2}$ which causes the $t \propto \sqrt{\epsilon}$ scaling. To find the optimal t_* for $\mathcal{S}_1$, higher-order terms must be calculated.

A.3.2 Higher-Order Error in ϵ, t

The error terms that contribute to the error rates are (up to a factor of $R_Z(\pi)$)

$$\mathcal{E}_{\mathcal{S}_1;t^2} = 8i \sum_{i \neq j \neq k} J_{ij} J_{jk} (X_i Y_j Z_k - Y_i X_j Z_k - Y_i Z_j X_k), \quad (\text{A.35})$$

$$\begin{aligned} \mathcal{E}_{\mathcal{S}_1;\epsilon^2} = 2i \sum_{i \neq j \neq k} J_{ij} J_{jk} & \left(8Y_i Z_j X_k - 4(Y_i X_j + X_i Y_j) Z_k \right. \\ & \left. + \frac{7(Y_i X_j Y_k - X_i Y_j X_k) + Z_i(Y_j - X_j)Z_k}{\pi} \right) + \frac{12i}{\pi} \sum_{i \neq j} J_{ij}^2 (X_i - Y_i), \end{aligned} \quad (\text{A.36})$$

$$\begin{aligned} \mathcal{E}_{\mathcal{S}_1;t\epsilon} = 8i \sum_{i \neq j \neq k} J_{ij} J_{jk} & \left(2Y_i Z_j X_k - \frac{3}{2} X_i Y_j Z_k + \frac{1}{2} Y_i X_j Z_k + \frac{Y_i X_j Y_k - X_i Y_j X_k}{\pi} \right) \\ & + \frac{8i}{\pi} \sum_{i \neq j} J_{ij}^2 (X_i - Y_i). \end{aligned} \quad (\text{A.37})$$

A.3.3 Evaluation of Prefactors

The scaling prefactors in Eqs. (A.22) and (A.23) require evaluation of the following quantities:

$$\|\mathcal{E}_{\mathcal{S}_{1/2};\epsilon}\| \leq \frac{4\sqrt{2}}{\pi} J_1, \quad (\text{A.38})$$

$$\|\mathcal{E}_{\mathcal{S}_1;\epsilon^2}\| \leq 28J_3 + \frac{12\sqrt{2}}{\pi} J_2, \quad (\text{A.39})$$

$$\|\mathcal{E}_{\mathcal{S}_1;t\epsilon}\| \leq 25J_3 + \frac{8\sqrt{2}}{\pi} J_2, \quad (\text{A.40})$$




Figure A.1: *Error scalings* Scaling of the contributions $\mathcal{E}$ to the error from different sources in a single Trotter step of the pulse sequences $\mathcal{S}_{1/2}$ and $\mathcal{S}_1$ as a function of system geometry. N_y is varied on the x axis and each line represents the contribution to the error with a different value of N_x . Here the analytic bounds of Eqs. (A.38) - (A.41) are compared to values obtained by numerically evaluating the norm of Eq. (2.15) for each type of contribution. Note that in the latter three panels, the $N_x = 1, N_y = 1$ points are much lower than the rest of the $N_x = 1$ lines because of the fact that there are three-body contributions to the error for these terms; when the system has less than three sites, these contributions are absent, leading to a much smaller error.

$$\left\| \mathcal{E}_{\mathcal{S}_{1/2}; t^2} \right\| = \left\| \mathcal{E}_{\mathcal{S}_1; t^2} \right\| \leq 8\sqrt{3 + 2\sqrt{3}}J_3, \quad (\text{A.41})$$

where the factors J_1, J_2, J_3 are defined in Eqs. (A.46) - (A.48). As explained in Section 2.4, numerics are used to see the values of the different $\mathcal{E}$ terms in practice. Eqs. (A.38) - (A.41) are compared to values obtained numerically as a function of system size in Fig. A.1.

With these expressions, calculating Eqs. (2.21) and (2.22) is now straightforward.

$$t_{\mathcal{S}_{1/2}*} \approx .3\sqrt{\frac{J_1}{J_3}}\epsilon^{\frac{1}{2}} + O(\epsilon), \quad (\text{A.42})$$

$$t_{\mathcal{S}_1*} = 0. \quad (\text{A.43})$$

Note that calculating $t_{\mathcal{S}_1*}$ using the bounds derived gives a negative number, which is why Eq. (A.43) is zero; this conclusion is supported by numerics. The corresponding optimal error rate bounds are

$$ER_{\mathcal{S}_{1/2}} \leq 12.2\sqrt{J_1 J_3}\epsilon^{\frac{1}{2}} + O(\epsilon), \quad (\text{A.44})$$

$$ER_{\mathcal{S}_1} \leq (14J_3 + 2.7J_2)\epsilon. \quad (\text{A.45})$$

A.4 Analysis of the Interaction Sums

Throughout the main text, various sums over the Ising interaction terms appear. These sums encapsulate the geometry of the device used. N_x is the horizontal extent of the atoms and N_y is the vertical extent. These sums depend very cleanly on total number of atoms $N = N_x N_y$. In the following analysis it is assumed that the configuration of atoms in the experimental apparatus is two-dimensional and rectangular (i.e. $N_x, N_y \geq 2$). The sums analyzed are:

$$J_1(N) = \sum_{i \neq j} J_{ij} = \tilde{J}_{11}N - \tilde{J}_{12}, \quad (\text{A.46})$$

$$J_2(N) = \sum_{i \neq j} J_{ij}^2 = \tilde{J}_{21}N - \tilde{J}_{22}, \quad (\text{A.47})$$

$$J_3(N) = \sum_{i \neq j \neq k} J_{ij}J_{jk} = \tilde{J}_{31}N - \tilde{J}_{32}. \quad (\text{A.48})$$

By explicit numerical evaluation of the sums for the considered system sizes, the following fits are found, which are plotted in Fig. A.2:

$$\begin{aligned} \tilde{J}_{11} &= 4.46, & \tilde{J}_{12} &= -36.74, \\ \tilde{J}_{21} &= 3.91, & \tilde{J}_{22} &= 28.35, \\ \tilde{J}_{31} &= 4.05, & \tilde{J}_{32} &= -63.58. \end{aligned}$$




Figure A.2: J_1, J_2, J_3 sums evaluated for explicit array configurations with fits. The extent of the array in the x direction is $N_x = 20$ atoms. While there is a deviation for smaller array sizes, a linear scaling is still seen.

To see how this scaling arises, consider the J_1 sum, which may be written as:

$$\begin{aligned}
 J_1 &= \frac{C_6}{a^6} \sum_{\substack{x_1=1 \\ y_1=1}}^{N_x, N_y} \sum_{\substack{x_2=x_1+1 \\ y_2=y_1+1}}^{N_x, N_y} \frac{1}{((x_2 - x_1)^2 + (y_2 - y_1)^2)^3} \\
 &= \frac{C_6}{a^6} \sum_{\substack{x_1=1 \\ y_1=1}}^{N_x, N_y} \sum_{\substack{x'_2=1 \\ y'_2=1}}^{N_x-x_1, N_y-y_1} \frac{1}{((x'_2)^2 + (y'_2)^2)^3} \\
 &\leq \frac{C_6}{a^6} \sum_{\substack{x_1=1 \\ y_1=1}}^{N_x, N_y} \int_1^\infty \frac{dx'_2 dy'_2}{((x'_2)^2 + (y'_2)^2)^3}.
 \end{aligned}$$

The integral converges to some constant, so we see (for $N = N_x N_y$) that

$$J_1 = O(Na^{-6}).$$

Scalings for J_2 and J_3 may be shown in a similar fashion. This scaling may be understood intuitively in the following manner. The power law interactions are rapidly decaying and so the dominant contribution to these

sums will be given by a single sum over an effective nearest-neighbor interaction. For a square array, this implies a $\sim 4N$ scaling, which is seen in Fig. A.2. The deviation in the fit values from the ideal $4N$ scaling stem from the fact that the fits include values from geometrically thin arrays (e.g., 2×20), where boundary interactions play a larger role.

Appendix B

APPENDIX TO CHAPTER 3

B.1 Explicit Form of Quantum Circuits

This section gives explicit forms for the circuits that are constructed in Sec. 3.4.2 by specifying the Pauli decomposition $\mathcal{D}(H_{\square\square})$. The circuit used to calculate Fig. 3.9a is formed by exponentiating the terms in this decomposition. Table B.1 gives the terms in the decomposition that were retained for the circuits implemented on the quantum device, whereas Table B.2 gives the terms in the decomposition that were dropped.

The **subcircuit** column specifies the specific subcircuit, U_1 , U_2 or U_3 , represented in accordance with their graphical depiction in Fig. 3.7b. The **Pauli String** column denotes a specific Pauli string, where the left-most index corresponds to the top wire in the circuit element and the rightmost corresponds to the bottom wire. The δ column denotes at which truncation level the corresponding string would be dropped. For Fig. 3.9a, the truncation level of 0.66 is chosen. As such, all entries in Table B.1 have values $\delta > 0.66$ and all entries in Tables B.2, B.3 have values $\delta < 0.66$.

B.2 Explicit Matrix Elements of Hamiltonian Operators

This appendix contains the explicit expression for various electric bilinear operators that appear in the Hamiltonian. The matrix elements of these operators are presented in three different bases: group element, mixed and character irrep.

B.2.1 Group Element Basis

In the group element basis, the Hilbert space basis states are spanned by the three continuous variables, ω, θ, ϕ . Using this representation, the electric bilinear matrix elements are given by

$$\begin{aligned}\mathcal{E}_1^2 &= -\left(\frac{\partial^2}{\partial\omega_1^2} + \cot\frac{\omega_1}{2}\frac{\partial}{\partial\omega_1}\right) + \frac{1}{4}\csc^2\frac{\omega_1}{2}\mathcal{N} \\ \mathcal{E}_2^2 &= -\left(\frac{\partial^2}{\partial\omega_2^2} + \cot\frac{\omega_2}{2}\frac{\partial}{\partial\omega_2}\right) + \frac{1}{4}\csc^2\frac{\omega_2}{2}\mathcal{N}\end{aligned}$$

$$\begin{aligned}
\mathcal{E}_{R2} \cdot \mathcal{E}_{L2} &= \mathcal{E}_2^2 - \frac{1}{2}\mathcal{N} \\
\mathcal{E}_{1R} \cdot \mathcal{E}_{2R} &= -\cos \Theta \frac{\partial^2}{\partial \omega_1 \partial \omega_2} \\
&\quad + \frac{1}{2} \sin \Theta \left(\cot \frac{\omega_2}{2} \frac{\partial}{\partial \omega_1} + \cot \frac{\omega_1}{2} \frac{\partial}{\partial \omega_2} \right. \\
&\quad \left. + \frac{1}{2} \cot \frac{\omega_1}{2} \cot \frac{\omega_2}{2} \frac{\partial}{\partial \Theta} \right) \\
&\quad - \frac{1}{4} \left(1 + \cos \Theta \cot \frac{\omega_1}{2} \cot \frac{\omega_2}{2} \right) \mathcal{N} \\
\mathcal{E}_{1R} \cdot \mathcal{E}_{2L} &= \mathcal{E}_{1R} \cdot \mathcal{E}_{2R} + \frac{1}{2}\mathcal{N}
\end{aligned} \tag{B.1}$$

where $\mathcal{N}$ is the operator for which the Legendre polynomials are eigenfunctions, given in Eq. (3.26)

B.2.2 Mixed Basis

The mixed basis and the group element basis are related via

$$\langle \omega'_1 \omega'_2 \nu \rangle \omega_1 \omega_2 \Theta = P_\nu(\Theta) \delta(\omega_1 - \omega'_1) \delta(\omega_2 - \omega'_2) \tag{B.2}$$

To convert the Hamiltonian from the group element basis to the mixed basis, these four integrals

$$\begin{aligned}
\int d(\cos \Theta) P_{\nu'}(\Theta) P_\nu(\Theta) &= \delta_{\nu'\nu} \\
\int d(\cos \Theta) P_{\nu'}(\Theta) \mathcal{N} P_\nu(\Theta) &= \nu(\nu+1) \delta_{\nu'\nu} \\
\int d(\cos \Theta) P_{\nu'}(\Theta) \cos \Theta P_\nu(\Theta) &= \frac{(\nu+1) \delta_{\nu'\nu+1}}{\sqrt{(2\nu+1)(2\nu+3)}} \\
&\quad + \frac{\nu \delta_{\nu'\nu-1}}{\sqrt{(2\nu+1)(2\nu-1)}} \\
\int d(\cos \Theta) P_{\nu'}(\Theta) \sin \Theta \frac{\partial P_\nu(\Theta)}{\partial \Theta} &= \frac{\nu(\nu+1) \delta_{\nu'\nu+1}}{\sqrt{(2\nu+1)(2\nu+3)}} \\
&\quad - \frac{\nu(\nu+1) \delta_{\nu'\nu-1}}{\sqrt{(2\nu+1)(2\nu-1)}}
\end{aligned} \tag{B.3}$$

are used. Note that the unconventional normalization on $P_\nu(\Theta)$ is due to the requirement that the Hilbert space basis be orthonormal.

B.2.3 Character Irrep Basis

The character irrep basis and the group element basis are related via

$$\langle J_1 J_2 \nu \rangle \omega_1 \omega_2 \Theta = \chi_\nu^{J_1}(\omega_1) \chi_\nu^{J_2}(\omega_2) P_\nu(\Theta) \quad (\text{B.4})$$

where $\chi_\nu^J(\omega)$ are generalized character functions of the irreducible representations of the rotation group.

These functions are normalized such that

$$\int_0^{2\pi} 4 \sin^2 \frac{\omega}{2} \chi_\nu^{J'}(\omega) \chi_\nu^J(\omega) = \delta_{JJ'} \quad (\text{B.5})$$

where the factor of four is due to the explicit form of the Haar measure for $\text{SU}(2)$ parameterized using the axis-angle coordinates.

To convert between the group element basis and the character irrep basis, several matrix elements need to be computed. For the magnetic Hamiltonian, the only relevant matrix element is

$$\langle J'_1 J'_2 \nu' | \text{Tr } \hat{X}_1 | J_1 J_2 \nu \rangle = \frac{1}{\sqrt{2}} \delta_{J'_2 J_2} \delta_{\nu' \nu} \quad (\text{B.6})$$

$$\begin{aligned} & \times \left(\sqrt{\frac{(2J_1 - \nu + 1)(2J_1 + \nu + 2)}{(J_1 + 1)(2J_1 + 1)}} \delta_{J'_1 J_1 + \frac{1}{2}} \right. \\ & \left. + \sqrt{\frac{(2J_1 - \nu)(2J_1 + \nu + 1)}{J_1(2J_1 + 1)}} \delta_{J'_1 J_1 - \frac{1}{2}} \right), \end{aligned} \quad (\text{B.7})$$

with the matrix element for $\text{Tr } \hat{X}_2$ the same, except with the exchange of $1 \leftrightarrow 2$. For the electric component of the Hamiltonian, the relevant matrix elements are

$$\begin{aligned} \langle J'_1 J'_2 \nu' | \hat{\mathcal{E}}_\kappa^2 | J_1 J_2 \nu \rangle &= J_\kappa(J_\kappa + 1) \delta_{J'_1 J_1} \delta_{J'_2 J_2} \delta_{\nu' \nu} \\ \langle J'_1 J'_2 \nu' | \hat{\mathcal{E}}_{R2} \cdot \hat{\mathcal{E}}_{R2} | J_1 J_2 \nu \rangle &= \delta_{J'_1 J_1} \delta_{J'_2 J_2} \delta_{\nu' \nu} \\ & \times \left(J_2(J_2 + 1) - \frac{\nu(\nu + 1)}{2} \right) \\ \langle J'_1 J'_2 \nu' | \hat{\mathcal{E}}_{R1} \cdot \hat{\mathcal{E}}_{R2} | J_1 J_2 \nu \rangle &= -\frac{1}{4} \delta_{J'_1 J_1} \delta_{J'_2 J_2} \\ & \times \left(\nu(\nu + 1) \delta_{\nu' \nu} + \delta_{\nu' \nu + 1} \frac{c_{J_1 \nu}^+ c_{J_2 \nu}^+ (\nu + 1)}{\sqrt{(2\nu + 1)(2\nu + 3)}} \right. \\ & \left. + \delta_{\nu' \nu - 1} \frac{c_{J_1 \nu}^- c_{J_2 \nu}^- \nu}{\sqrt{(2\nu + 1)(2\nu - 1)}} \right) \end{aligned}$$

$$\begin{aligned}
\langle J'_1 J'_2 \nu' | \hat{\mathcal{E}}_{R1} \cdot \hat{\mathcal{E}}_{L2} | J_1 J_2 \nu \rangle &= \frac{1}{4} \delta_{J'_1 J_1} \delta_{J'_2 J_2} \\
&\times \left(\nu(\nu+1) \delta_{\nu' \nu} - \delta_{\nu' \nu+1} \frac{c_{J_1 \nu}^+ c_{J_2 \nu}^+ (\nu+1)}{\sqrt{(2\nu+1)(2\nu+3)}} \right. \\
&\quad \left. - \delta_{\nu' \nu-1} \frac{c_{J_1 \nu}^- c_{J_2 \nu}^- \nu}{\sqrt{(2\nu+1)(2\nu-1)}} \right)
\end{aligned} \tag{B.8}$$

where in the first expression $\kappa = 1, 2$ and the coefficients $c_{J\nu}^\pm$ are given by

$$\begin{aligned}
c_{J\nu}^+ &= \sqrt{(2J + \nu + 2)(2J - \nu)} \\
c_{J\nu}^- &= \sqrt{(2J + \nu + 1)(2J - \nu + 1)}
\end{aligned} \tag{B.9}$$

The full Hamiltonian will not be given here, though key components will. Of particular interest are the components that contain only first derivatives with respect to the ω variables, grouped into H_∂ . This Hamiltonian is given by

$$H_\partial = \sum_\nu |\nu\rangle \langle \nu+1| H_\nu^+ + \sum_\nu H_\nu^- |\nu\rangle \langle \nu-1| \tag{B.10}$$

where $H^\pm$ are

$$\begin{aligned}
H_\nu^+ &= A_\nu \frac{\partial}{\partial \omega_1} \frac{\partial}{\partial \omega_2} + B_\nu \left(\cot \frac{\omega_1}{2} \frac{\partial}{\partial \omega_2} + \cot \frac{\omega_2}{2} \frac{\partial}{\partial \omega_1} \right) \\
&\quad + \frac{1}{4} (\nu(\nu+1) A_\nu - 2B_\nu) \cot \frac{\omega_1}{2} \cot \frac{\omega_2}{2} \\
H_\nu^- &= D_\nu \frac{\partial}{\partial \omega_1} \frac{\partial}{\partial \omega_2} + C_\nu \left(\cot \frac{\omega_1}{2} \frac{\partial}{\partial \omega_2} + \cot \frac{\omega_2}{2} \frac{\partial}{\partial \omega_1} \right) \\
&\quad + \frac{1}{4} (\nu(\nu+1) D_\nu - 2C_\nu) \cot \frac{\omega_1}{2} \cot \frac{\omega_2}{2}
\end{aligned} \tag{B.11}$$

with the coefficients

$$\begin{aligned}
A_\nu &= \frac{(\nu+1)}{\sqrt{(2\nu+1)(2\nu+3)}} \\
B_\nu &= -\frac{(\nu+1)^2}{2\sqrt{(2\nu+1)(2\nu+3)}} \\
C_\nu &= \frac{\nu^2}{2\sqrt{(2\nu-1)(2\nu+1)}} \\
D_\nu &= \frac{\nu}{\sqrt{(2\nu-1)(2\nu+1)}}.
\end{aligned} \tag{B.12}$$

The relation that $H_\nu^+ = (H_{\nu+1}^-)^\dagger$ ensures hermiticity of the overall expression.

For the quantum simulations presented in this work, only H_0^+ and H_1^- is needed and so they are written out explicitly here:

$$\begin{aligned}
 H_0^+ &= \frac{1}{\sqrt{3}} \frac{\partial}{\partial \omega_1} \frac{\partial}{\partial \omega_2} + \frac{1}{4\sqrt{3}} \cot \frac{\omega_1}{2} \cot \frac{\omega_2}{2} \\
 &\quad - \frac{1}{2\sqrt{3}} \left(\cot \frac{\omega_1}{2} \frac{\partial}{\partial \omega_2} + \cot \frac{\omega_2}{2} \frac{\partial}{\partial \omega_1} \right) \\
 &= (H_1^-)^\dagger .
 \end{aligned} \tag{B.13}$$


Subcircuit	Pauli String	Coefficient	Truncation Percentage δ
	$I \otimes X$ $Z \otimes Z$ $X \otimes X$ $X \otimes I$ $Y \otimes Y$ $I \otimes Z$ $Z \otimes I$	-4.358185 3.049102 -2.468268 1.565659 -1.396263 0.884597 0.519954	0.9857 0.9429 0.9143 0.8571 0.7857 0.7429 0.7143
	$Z \otimes I \otimes I$ $Z \otimes Z \otimes I$ $Z \otimes I \otimes Z$ $Z \otimes Z \otimes Z$	-3.778734 -1.678690 -1.452504 -1.392269	0.9571 0.8857 0.8143 0.7571
	$X \otimes I \otimes I \otimes I \otimes I$ $Y \otimes I \otimes Y \otimes I \otimes I$	0.283820 -0.226330	0.6857 0.6714

Table B.1: Terms in the Pauli string decomposition of $H_{\square\square}$ that are retained in the circuits. The colored circuit elements correspond to those in the main text. The Pauli string column denotes the rotations $e^{-ic_p P}$ in the circuit for each string P with corresponding coefficient c_p . The truncation percentage δ denotes the fraction of Pauli strings that would need to be truncated for this term to have the smallest coefficient. As an example, all terms in this table were included in the circuit which was run at $\delta = .66$ and so all δ values in this table are greater than .66.

Subcircuit	Pauli String	Coeff.	δ
	$Y \otimes I \otimes I \otimes I \otimes Y$	-0.226330	0.6571
	$X \otimes I \otimes I \otimes Z \otimes I$	0.189503	0.6429
	$X \otimes Z \otimes I \otimes I \otimes I$	0.189503	0.6286
	$X \otimes I \otimes Y \otimes I \otimes Y$	-0.180485	0.6143
	$Y \otimes I \otimes Y \otimes Z \otimes I$	-0.151118	0.6000
	$Y \otimes Z \otimes I \otimes I \otimes Y$	-0.151118	0.5857
	$X \otimes I \otimes I \otimes I \otimes Z$	0.136188	0.5714
	$X \otimes I \otimes Z \otimes I \otimes I$	0.136188	0.5571
	$X \otimes Z \otimes I \otimes Z \otimes I$	0.126529	0.5429
	$Y \otimes X \otimes Y \otimes I \otimes I$	0.113165	0.5286
	$Y \otimes Y \otimes X \otimes I \otimes I$	-0.113165	0.5143
	$Y \otimes I \otimes I \otimes X \otimes Y$	0.113165	0.5000
	$Y \otimes I \otimes I \otimes Y \otimes X$	-0.113165	0.4857
	$X \otimes I \otimes I \otimes Z \otimes Z$	0.110025	0.4714
	$X \otimes Z \otimes Z \otimes I \otimes I$	0.110025	0.4571
	$Y \otimes I \otimes Y \otimes I \otimes Z$	-0.108602	0.4429
	$Y \otimes I \otimes Z \otimes I \otimes Y$	-0.108602	0.4286
	$X \otimes I \otimes Z \otimes Z \otimes I$	0.090931	0.4143
	$X \otimes Z \otimes I \otimes I \otimes Z$	0.090931	0.4000
	$X \otimes I \otimes Y \otimes X \otimes Y$	0.090243	0.3857
	$X \otimes I \otimes Y \otimes Y \otimes X$	-0.090243	0.3714
	$X \otimes X \otimes Y \otimes I \otimes Y$	0.090243	0.3571
	$X \otimes Y \otimes X \otimes I \otimes Y$	-0.090243	0.3429
	$Y \otimes I \otimes Y \otimes Z \otimes Z$	-0.087739	0.3286

Table B.2: The columns are the same as in the Table B.1. This table (continued in Table B.3) contains all of the Pauli string in the decomposition $\mathcal{D}(H_{\square\square})$ that were dropped in the quantum simulations. As such all entries in the δ column are less than $\delta = .66$. Note that all of the terms dropped are included in the yellow circuit element. They all have either an X or a Y located on the ν register. Additionally, they are generally much higher weight and have support on a greater number of qubits.

Subcircuit	Pauli String	Coeff.	δ
	$Y \otimes Z \otimes Z \otimes I \otimes Y$	-0.087739	0.3143
	$Y \otimes X \otimes Y \otimes Z \otimes I$	0.075559	0.3000
	$Y \otimes Y \otimes X \otimes Z \otimes I$	-0.075559	0.2857
	$Y \otimes Z \otimes I \otimes X \otimes Y$	0.075559	0.2714
	$Y \otimes Z \otimes I \otimes Y \otimes X$	-0.075559	0.2571
	$X \otimes Z \otimes I \otimes Z \otimes Z$	0.073463	0.2429
	$X \otimes Z \otimes Z \otimes Z \otimes I$	0.073463	0.2286
	$X \otimes I \otimes Z \otimes I \otimes Z$	0.065348	0.2143
	$Y \otimes I \otimes Z \otimes X \otimes Y$	0.054301	0.2000
	$Y \otimes I \otimes Z \otimes Y \otimes X$	-0.054301	0.1857
	$Y \otimes X \otimes Y \otimes I \otimes Z$	0.054301	0.1714
	$Y \otimes Y \otimes X \otimes I \otimes Z$	-0.054301	0.1571
	$X \otimes I \otimes Z \otimes Z \otimes Z$	0.052795	0.1429
	$X \otimes Z \otimes Z \otimes I \otimes Z$	0.052795	0.1286
	$X \otimes X \otimes Y \otimes X \otimes Y$	-0.045121	0.1143
	$X \otimes X \otimes Y \otimes Y \otimes X$	0.045121	0.1000
	$X \otimes Y \otimes X \otimes X \otimes Y$	0.045121	0.0857
	$X \otimes Y \otimes X \otimes Y \otimes X$	-0.045121	0.0714
	$Y \otimes X \otimes Y \otimes Z \otimes Z$	0.043870	0.0571
	$Y \otimes Y \otimes X \otimes Z \otimes Z$	-0.043870	0.0429
	$Y \otimes Z \otimes Z \otimes X \otimes Y$	0.043870	0.0286
	$Y \otimes Z \otimes Z \otimes Y \otimes X$	-0.043870	0.0143
	$X \otimes Z \otimes Z \otimes Z \otimes Z$	0.042652	0.0000

Table B.3: Continuation of Table B.2.

B.3 Table of Lowest Energy Eigenstates in the character Irrep Basis

The energy spectrum of the two plaquette Hamiltonian in the strong coupling limit can be found exactly when working in the character irrep basis. This is due to the electric component of the Hamiltonian being block diagonal, with each block of finite dimension $2 \text{Min}(J_1, J_2) + 1$; the magnetic component is negligible and can be ignored. The energies and expressions for the several lowest lying eigenstates are shown in Table [B.4](#).

Level of State	Degeneracy	In Irrep Basis: $ J_1, J_2, \nu\rangle$	Energy
Ground	None	$ 0, 0, 0\rangle$	0
First Excited	Two-Fold	$ \frac{1}{2}, 0, 0\rangle$ $ 0, \frac{1}{2}, 0\rangle$	$\frac{3}{2}g^2$
Second Excited	None	$\frac{1}{2} \frac{1}{2}, \frac{1}{2}, 0\rangle - \frac{\sqrt{3}}{2} \frac{1}{2}, \frac{1}{2}, 1\rangle$	$\frac{9}{4}g^2$
Third Excited	None	$\frac{\sqrt{3}}{2} \frac{1}{2}, \frac{1}{2}, 0\rangle + \frac{1}{2} \frac{1}{2}, \frac{1}{2}, 1\rangle$	$\frac{13}{4}g^2$
Fourth Excited	Two-Fold	$ 1, 0, 0\rangle$ $ 0, 1, 0\rangle$	$4g^2$
Fifth Excited	Two-Fold	$\frac{1}{\sqrt{3}} 1, \frac{1}{2}, 0\rangle - \sqrt{\frac{2}{3}} 1, \frac{1}{2}, 1\rangle$ $\frac{1}{\sqrt{3}} \frac{1}{2}, 1, 0\rangle - \sqrt{\frac{2}{3}} \frac{1}{2}, 1, 1\rangle$	$\frac{9}{2}g^2$
Sixth Excited	Three-Fold	$\frac{1}{\sqrt{3}} 1, \frac{1}{2}, 0\rangle - \sqrt{\frac{2}{3}} 1, \frac{1}{2}, 1\rangle$ $\sqrt{\frac{2}{3}} \frac{1}{2}, 1, 0\rangle + \frac{1}{\sqrt{3}} \frac{1}{2}, 1, 1\rangle$ $\frac{1}{3} 1, 1, 0\rangle - \frac{1}{\sqrt{3}} 1, 1, 1\rangle + \frac{\sqrt{5}}{3} 1, 1, 2\rangle$	$6g^2$

Table B.4: Energies and explicit eigenstate constructions for the seven lowest-lying eigenstates of the two plaquette Hamiltonian in the strong coupling limit. For the states that are degenerate in this limit, the magnetic Hamiltonian breaks the degeneracy and splits the states.

Appendix C

APPENDIX TO CHAPTER 4

C.1 The Non-Stabilizerness of Indicator States

This section develops a toy model for understanding the behavior of non-stabilizerness in the ground state of the two-plaquette system. For a system of dn qubits, the basis states are defined by a tuple of d variables $\vec{x} = (x_1, \dots, x_d)$, each taking values in $\{1, \dots, 2^n - 1\}$. A cutoff C is defined as $C = \sum_{j=0}^{n-1} C_j 2^j$ where the coefficients C_j denote the binary expansion of C . Given such a binary expansion of a cutoff, it is useful to define the set of bitstrings S_j as those that lie within two adjacent powers of two, j and l , *i.e.* $S_j = \{x | 2^j \leq x \leq 2^j + 2^l\}$. The d -dimensional indicator state $|\Theta_d(C)\rangle$ is defined as

$$|\Theta_d(C)\rangle = \sum_{\vec{x}} f(\vec{x}) |\vec{x}\rangle, \quad f(\vec{x}) \propto \begin{cases} A_j & \max(\vec{x}) \in S_j \\ 0 & \text{else} \end{cases} \quad (\text{C.1})$$

where the proportionality denotes up to normalization. This state encodes an indicator function on basis states such that it takes value A_j if the largest entry of $\vec{x}$ lie within two adjacent powers of two in the cutoff's binary expansion and vanishes otherwise. Solving for the non-stabilizerness of this general state can be complicated, but its structure can be constrained in a way that allows it to be exactly solvable.

First consider the one-dimensional state such that the cutoff has expansion $C = 2^a + 2^b$ with $a \geq b$, *i.e.* there are only two powers of two in the cutoff's binary expansion. For $x_1 \in S_1$ then the amplitude of the state is A and if it lies within S_2 the amplitude is B , with $|S_1| = 2^a$ and $|S_2| = 2^b$. The normalization of this state is $\langle \Theta_1 | \Theta_1 \rangle = A^2 2^a + B^2 2^b$. The non-stabilizerness of this state is found by considering the expression for $\mathcal{M}_2$ in terms of the amplitudes [587]

$$e^{-\mathcal{M}_2} = \sum_{s_1, s_2, s_3, s_4} c_{s_1} c_{s_2} c_{s_3} c_{s_4} c_{s_1}^* c_{s_2}^* c_{s_3}^* c_{s_4}^* \quad (\text{C.2})$$

where

$$\tilde{s}_1 = s_2 \oplus s_3 \oplus s_4 \quad \tilde{s}_2 = s_1 \oplus s_3 \oplus s_4 \quad \tilde{s}_3 = s_1 \oplus s_2 \oplus s_4 \quad \tilde{s}_4 = s_1 \oplus s_2 \oplus s_3. \quad (\text{C.3})$$

As the amplitudes are purely real, evaluation of the right hand side can be reduced to assessing for a given quadruplet of bitstrings (s_1, s_2, s_3, s_4) the S_j membership of the corresponding $(\tilde{s}_1, \tilde{s}_2, \tilde{s}_3, \tilde{s}_4)$. This

membership only depends on which S_j the original quadruplet is in and so for $J \leq n - 1$ possible S_j there are $\binom{J+3}{4}$ distinct cases that must be evaluated. For cutoff choice $C = 2^a + 2^b$ this reduces to 5 distinct cases, namely how many of the s_i are in S_1 and how many are in S_2 , where the number of valid quadruples for each case being denoted as t_i . The bitstring representation of any $s_i \in S_1$ can be any combination of 0 or 1 until the a^{th} bit, whereas any $s_i \in S_2$ will have a 1 at the a^{th} bit and have any combination of 0 and 1 until the b^{th} bit and 0 in all other positions. The counts t_i are enumerated as follows:

1. **None in S_2 :** For any $s_i \in S_1$, XOR-ing it with another $s_j \in S_1$ will result in another string in S_1 . There are $t_0 = 2^{4a}$ combinations of such strings.
2. **One in S_2 :** Consider the case where $s_1 \in S_2$, $s_2, s_3, s_4 \in S_1$. As with the previous case $\tilde{s}_1$ will trivially be in S_1 . For the remaining $\tilde{s}_j$, the a^{th} bit is 1 due to s_1 and so $\tilde{s}_j$ is necessarily an element of S_2 . This implies the bits in positions $b, \dots, a - 1$ must be 0, implying that the pair in S_1 must have bit patterns that cancel one another. The bits in positions $0, \dots, b - 1$ are unfixed, and so the total number of configurations is

$$\begin{aligned}
 t_1 &= 4 \times 2^b \cdot \sum_{patterns} (2^b)^3 \\
 &= 4 \cdot 2^b \cdot 2^{a-b} \cdot 2^{3b} \\
 &= 4 \times 2^{a+3b}
 \end{aligned} \tag{C.4}$$

where the 4 accounts for the different ways to place the element of S_2 and the sum is over patterns in bit positions $b, \dots, a - 1$.

3. **Two in S_2 :** Consider the case where $s_1, s_2 \in S_2$, $s_3, s_4 \in S_1$. This implies that $\tilde{s}_3, \tilde{s}_4 \in S_1$ for any choice of s_1, s_2 . By the logic of the previous case, for $\tilde{s}_1, \tilde{s}_2 \in S_2$ the constituent bitstrings s_3 and s_4 must match in positions $b, \dots, a - 1$. There are 2^{2b} choices for pairs of s_1 and s_2 , and so $t_2 = 6 \times 2^{a+3b}$, where the factor of 6 comes from all the ways of choosing two elements of the quadruplet to be in S_2 .
4. **Three in S_2 :** Consider the case where $s_1, s_2, s_3 \in S_2$, $s_4 \in S_1$. All of $\tilde{s}_1, \tilde{s}_2, \tilde{s}_3$ will trivially be in S_1 . Further, $\tilde{s}_4$ will always be in S_2 as s_1, s_2, s_3 will always have 0 in bit positions $b, \dots, a - 1$. Therefore $t_3 = 4 \times 2^{a+3b}$ given all the ways to choose the element of S_1 from the quadruplet.

5. **All in S_2 :** When all the s_i are in S_2 , the $\tilde{s}_j$ will trivially be in S_2 and so $t_4 = 2^{4b}$

Collecting all of the contributing terms from the cases above and assigning them their corresponding factors of A and B gives the final expression for the non-stabilizerness $\mathcal{M}_2$ of

$$\mathcal{M}_2 = -\log \left(\frac{A^8 2^{4a} + 14A^4 B^4 2^{a+3b} + B^8 2^{4b}}{(A^2 2^a + B^2 2^b)^4} \right) \quad (\text{C.5})$$

As a special case of the above equation, consider when $a = b$ and $A = B$, which corresponds to a state with cutoff C a power of two. The argument of the logarithm becomes unity and so the state is a stabilizer state. For $A \neq B$, these points will be minima. This suggests that the oscillations observed in the ground state of the two-plaquette system are an instance of the more generic behavior where states supported on a power of two number of basis states are generally closer to stabilizer states.

Having considered the case of one dimension, the two-dimensional case follows straightforwardly. In particular, the computational basis quadruplets (s_1, s_2, s_3, s_4) decompose into x and y directions such that the constraints act independently. Therefore, there are 25 total cases, corresponding to all products of the five cases discussed in the one dimensional case. The number of contributing terms for each of the 25 is just a product of the terms in the constituent one-dimensional cases. For cases $(2, 2), (3, 3), (4, 4), (2, 4), (4, 2)$ there are different summands, depending on whether the strings in S_2 for the x and y directions fall on the same amplitude $c_{(s_y, s_x)}$ or different amplitudes. Summing all of the 25 cases gives the numerator

$$\begin{aligned} \sum_{s_1, s_2, s_3, s_4} c_{s_1} c_{s_2} c_{s_3} c_{\tilde{s}_4} c_{\tilde{s}_1}^* c_{\tilde{s}_2}^* c_{\tilde{s}_3}^* c_{s_4}^* &\propto 2^{8a} A^8 + \left(14 \cdot 2^{2a+6b} + 28 \cdot 2^{5a+3b} \right) A^4 B^4 \\ &+ \left(120 \cdot 2^{2a+6b} + 48 \cdot 2^{5a+3b} \right) A^2 B^6 \\ &+ \left(2^{8b} + 28 \cdot 2^{a+7b} + 2^{4a+4b+1} + 14 \cdot 2^{2a+6b} \right) B^8 \quad (\text{C.6}) \end{aligned}$$

where the proportionality is up to the normalization of the state $\sqrt{4^a A^2 + (2^{a+b+1} + 4^b) B^2}$. Specializing to the case with $A = B$, the non-stabilizerness of this state can then be expressed as

$$\mathcal{M}_2 = -\log \left(\frac{2^{8a} + 76 \cdot 2^{5a+3b} + 148 \cdot 2^{2a+6b} + 28 \cdot 2^{a+7b} + 2 \cdot 2^{4a+4b} + 2^{8b}}{(2^a + 2^b)^8} \right) \quad (\text{C.7})$$

When $a = b$ the non-stabilizerness vanishes, whereas for any $b < a$ the non-stabilizerness will be non-zero. Therefore when the support of the indicator function is a power of two the state is stabilizer, and it is maximized when it is in-between a power of two, causing the oscillations in the non-stabilizerness of the ground state observed in the main body.

C.2 Estimators for Stabilizer Entropies

C.2.1 Estimator Derivation

This section gives a complete derivation for the estimators given in the main text. Begin by considering the more general case of estimating the observable

$$O^{(m)} = \sum_{\{s_1, \dots, s_m\}} A^{-w(\oplus_i s_i)} \prod_i p(s_i) \quad (\text{C.8})$$

where w is the weight of a bitstring and $\oplus_i s_i$ denotes the element-wise XOR of all the s_i . The strategy will be to rewrite the kernel $A^{-w(\oplus_i s_i)}$ as the Walsh-Hadamard transform of a set of undetermined coefficients $\{a_k\}$. Explicitly the ansatz is

$$A^{-w(\oplus_i s_i)} = \sum_k a_k (-1)^{k \cdot (\oplus_i s_i)} \quad (\text{C.9})$$

and so all that remains is to determine the values of the a_k . Plugging this in to the expression for the estimator $O^{(m)}$ gives

$$\begin{aligned} O^{(m)} &= \sum_{\{s_1, \dots, s_m\}} \sum_k a_k (-1)^{k \cdot (\oplus_i s_i)} \prod_i p(s_i) \\ &= \sum_k a_k \prod_i \sum_{s_i} (-1)^{k \cdot s_i} p(s_i) \\ &= \sum_k a_k \hat{p}_j^m(k) \end{aligned} \quad (\text{C.10})$$

where $\hat{p}(k)$ is the Walsh-Hadamard transform of the measurement distribution and in the second line we have factorized the product among the four replicas. Now all that is left is to determine the expression for the a_k . This is done by taking the Inverse Walsh-Hadamard Transform of Eq. (C.9)

$$\begin{aligned} a_k &= \frac{1}{d} \sum_s A^{-w(s)} (-1)^{k \cdot s} \\ &= \frac{1}{d} \prod_b \sum_{s_b} A^{-s_b} (-1)^{k_b s_b} \\ &= \frac{1}{d} \prod_b \left(1 + \frac{(-1)^{k_b}}{A} \right) \\ &= \frac{1}{d} \left(1 + \frac{1}{A} \right)^{n-w(k)} \left(1 - \frac{1}{A} \right)^{w(k)} \end{aligned} \quad (\text{C.11})$$

where in the second line the sum over s has been factorized into a sum over the individual bits. The observables $O^{(2)}, O^{(4)}$ correspond to the cases $m = 2, 4$ respectively and $A = -2$. This gives the final expressions in the main text Eq 4.7.

C.3 Walsh-Paley Based Reconstruction of Wavefunction Probabilities

C.3.1 One Dimension

This section first describes the general grounds by which the Total Variational Distance (TVD) between the exact bitstring distribution and the reconstructed distribution may be bounded by the sequency spectrum of the underlying distribution. Then the dominant coefficients in each sequency band of relevant classical functions are calculated to give tight bounds on the reconstruction quality. Finally, different noise models are considered to understand how different types of noise might affect the reconstruction quality of the bitstring distributions.

To bound the TVD between two distributions, it is useful to view it as the L_1 norm of the probability vector $p(s) - p(s; k^*)$. The L_2 norm of the vector is related to the Mean Squared Error (MSE)

$$\varepsilon(p_1, p_2) = \sqrt{\sum_s (p_1(s) - p_2(s))^2} \quad (\text{C.12})$$

By an application of the Cauchy-Schwarz inequality the following bound applies

$$\delta(p_1, p_2) \leq \frac{\sqrt{d}}{2} \varepsilon(p_1, p_2) \quad (\text{C.13})$$

where $d = 2^n$ is the length of the vector for n qubits. The MSE between the exact and reconstructed distributions is

$$\begin{aligned} \sum_s (p(s) - p(s; k^*))^2 &= \sum_s \left(\frac{1}{d} \sum_{k > k^*} (-1)^{k \cdot s} c_k \right)^2 \\ &= \frac{1}{d^2} \sum_s \sum_{k, k' > k^*} (-1)^{s \cdot (k+k')} c_k c_{k'} \\ &= \frac{1}{d} \sum_{k > k^*} c_k^2 \end{aligned} \quad (\text{C.14})$$

where the third line used the orthogonality of Walsh functions. Therefore to bound the MSE it is crucial to understand the asymptotic decay of the sequency spectra. For convenience k^* will start at the beginning of the band ν , thereby providing a worst case estimate for all elements of the band. In practice, starting at an

arbitrary index in the block gives a small improvement to the error which is sub-leading to the improvement by including only the first element.

Under the assumption that the bitstring distribution is encoded by a smooth function $f(x)$, it has been shown that the Paley-ordered coefficients of degree k obey the following bound [588]

$$|c_k| \leq F \times 2^{-q(k)} \quad (\text{C.15})$$

where the prefactor $F = \max_x |f^{(r)}(x)|$ is the maximum of the r^{th} derivative of f over the domain of x , $r = \sum_i k_i$ is the number of ones in k , and $q(k) = \sum_i i k_i + r(k)$. The function $q(k)$ obeys the following shift property $q(2^m + s) = (m + 1) + q(s)$ relating bitstrings between blocks. The decay of the reconstruction error Eq. (C.14) can therefore be expressed as

$$\sum_{k > k^*} \frac{1}{4^{q(k)}} = \sum_{j=m}^{n-1} \frac{1}{4^{j+1}} \sum_{k=0}^{2^j-1} \frac{1}{4^{q(k)}} \quad (\text{C.16})$$

$$= \sum_{j=m}^{n-1} \frac{1}{4^{j+1}} \prod_i^j \sum_{k_i \in \{0,1\}} \frac{1}{4^{k_i(i+1)}} \quad (\text{C.17})$$

$$= \sum_{j=m}^{n-1} \frac{1}{4^{j+1}} \prod_i^j \left(1 + \frac{1}{4^{i+1}} \right) \quad (\text{C.18})$$

$$\leq \frac{C}{3} \left(\frac{1}{4^m} - \frac{1}{4^n} \right) \quad (\text{C.19})$$

where $C = \prod_i^n (1 + 1/4^{i+1}) < 1.356$ is a constant for $n \rightarrow \infty$; the first line represents the sum over different bands. The MSE is bounded by

$$\varepsilon^2(\mathbf{p}_1, \mathbf{p}_2) \leq \frac{C \tilde{F}^2}{3d} \left(\frac{1}{4^m} - \frac{1}{4^n} \right) \quad (\text{C.20})$$

$$(\text{C.21})$$

where $\tilde{F} = \max_k F$. Assuming that n is large, the TVD dependence of the cutoff is

$$\delta(\mathbf{p}(s), \mathbf{p}(s; k^*)) \leq \max_{r,s} |\mathbf{p}^{(r)}(s)| \sqrt{\frac{C}{12}} \frac{1}{2^m} \quad (\text{C.22})$$

where m denotes the most significant bit of the cutoff k^* . Most notably, the only dependence on system size comes from the dependence on the derivatives. It is clear that for bitstring distributions that vary on the scale of individual bits, this will incur a factor exponential in system size. However for sufficiently smooth functions, namely those that have bounded derivative as $n \rightarrow \infty$, to achieve a constant TVD it is sufficient to calculate only the coefficients up to a particular band.

C.3.2 Multiple Dimensions

The case of encoding a D-dimensional function $f : \mathbb{R}^D \rightarrow \mathbb{R}$ on the interval $[0, 1]^{\otimes D}$ is now considered. In D dimensions, the Paley-ordered Walsh functions are given by $W_k(x) = \prod_{j=1}^D W_{k_j}(x_j)$ where k denotes a list of Walsh indices. It is useful to introduce the following functions

$$Y(k, x) = \prod_j y(k_j, x_j), \quad y(k_i, x_i) = \frac{1}{(r(k_i) - 1)!} \int_0^t d\tau W_{k_i}(\tau) (x_i - \tau)^{r(k_i)-1} \quad (\text{C.23})$$

where $t, \tau \in \mathbb{R}^D$ and k is a vector of bitstrings. The Walsh functions W_k and y satisfy the following properties:

1. $\int_0^1 (\prod_i t_i^{\alpha_i} W_{k_i}) dt = 0$ if $\alpha_j < r(k_j)$
2. $y(k_i, t) = \int_0^t \int_0^{t_1} \dots \int_0^{t_{r-1}} W_{k_i}(t_r) dt_r \dots dt_1$
3. $\int_0^1 y(k_i, t) dt = 2^{-q(k_i)}$

which are analogous to the single dimensional conditions given in [589].

Denoting the block index of the band vector of k as ν , there are $\prod_{j=1}^D 2^{\nu_j-1}$ terms in the same band as k , defining a hypercube with corresponding side-lengths in k space.

The bound on the Walsh coefficient c_k is found by first expanding f up to order $\rho = \sum_j r(k_j)$ around the boundary point $(1, \dots, 1)$

$$f(x) = \sum_{|\alpha| < \rho} \frac{\partial_\alpha f}{\alpha!} (x-1)^\alpha + R_\rho(x), \quad R_\rho(x) = \sum_{|\alpha|=\rho} \frac{\rho(x-1)^\alpha}{\alpha!} \int_0^1 ds (1-s)^\rho D^\alpha f(1+s(x-1)) \quad (\text{C.24})$$

where $\alpha = (\alpha_1, \dots, \alpha_D)$ is a multi-index of differentiation orders in each direction, $|\alpha| = \alpha_1 + \dots + \alpha_D$, and $D^\alpha = \prod_i \partial_{\alpha_i}$. The term $R_r(x)$ is simply the integral form of the remainder for a Taylor expansion, and t parameterizes a line from the boundary point 1 to the point x . The Walsh coefficient c_k is found by integrating the series term by term against the Walsh function, where Property 1 causes the first terms to vanish.

$$c_k = \sum_\alpha \frac{\rho}{\alpha!} \int_0^1 d\tau W_k(\tau) (x-1)^\alpha \int_0^1 ds (1-s)^r D^\alpha f(1+s(x-1)) \quad (\text{C.25})$$

The integral over s is bounded by $\Delta = \max_x f^{(\alpha)}(x)$, which is the maximum value of the directional derivatives specified by α over the integration region. Again by Property 1, the only term in the sum over α

that does not vanish is the combination of α 's that matches the ranks of the corresponding index vector ν .

Using properties 2 and 3, the Walsh coefficient is then given by

$$|c_k| \leq \Delta \rho \prod_j \frac{1}{r(k_j)!} \int_0^1 d\tau_j W_{k_j}(\tau_j) (x_j - 1)^{r(k_j)} \quad (\text{C.26})$$

$$= \Delta \frac{\rho}{\prod_j r(k_j)} \int_0^1 d\tau Y(k, \tau) \quad (\text{C.27})$$

$$\leq D \Delta \prod_j 2^{-q(\nu_j)} \quad (\text{C.28})$$

which is analogous to the bound in the one dimensional case with an extra factor of d in front. Considering the case where the cutoff in each direction is k_j^* , the error will be given by summing over all k outside this hypercube defined by $\mathcal{C} = \bigcup_j A_j$ where $A_j = \{k | k_j > k_j^*\}$. Applying a union bound to this sum gives the following

$$\begin{aligned} \sum_{k \in \mathcal{C}} \frac{1}{4^{q(k)}} &\leq \sum_{j=1}^D \sum_{k \in A_j} \frac{1}{4^{q(k)}} \\ &= \sum_j \sum_{k_j > k_j^*} 4^{-q(k_j)} \prod_{i \neq j} \sum_{k_i} 4^{-q(k_i)} \\ &= C^{D-1} \sum_j \sum_{k_j > k_j^*} 4^{-q(k_j)} \end{aligned} \quad (\text{C.29})$$

$$= \frac{C^D}{3} \sum_j \left(\frac{1}{4^{m_j}} - \frac{1}{4^n} \right) \quad (\text{C.30})$$

where the final line follows from using the single variable sum over j and m_j is the most significant bit of the cutoff k_j^* . Once again, assuming large n , the bound on the TVD becomes

$$\delta \leq \sqrt{\frac{DC^D}{6}} \Delta \sum_j 2^{-m_j} \quad (\text{C.31})$$

This bound is not tight in terms of its dependence on the dimension D . Crucially, for constant finite D and a sufficiently smooth function, the bound is independent of the number of qubits used in the digitization.

Appendix D

APPENDIX TO CHAPTER 5

D.1 Details of Sample Complexity Bound

Here, we provide details for the derivation of the bound on sample complexity for the Global Dual Pairs protocol for $\mathbb{Z}_2$ LGT with PBC. In particular, we provide a derivation of Eq. (5.27). From Lemma S1 of [37], for a traceless operator O , the variance of the corresponding shadow expectation value o_j , where $\mathbb{E}_{U,b}(o_j) = \text{Tr}(\rho O)$,

$$\text{Var}(o_j) = \text{Var} \left(\text{Tr} \left[\widehat{\Phi[\rho]} O \right] \right) \leq \|O\|_{\text{shadow}}^2, \quad (\text{D.1})$$

where ρ is the unknown quantum state and, thus, $\widehat{\Phi[\rho]}$ is the Ising-side classical shadow of this state. The so-called shadow norm is defined as

$$\|O\|_{\text{shadow}}^2 = \max_{\sigma} \mathbb{E}_U \sum_{\vec{b} \in \{0,1\}^n} \langle \vec{b} | U \sigma U^\dagger | \vec{b} \rangle \langle \vec{b} | U \mathcal{M}^{-1}[O] U^\dagger | \vec{b} \rangle^2, \quad (\text{D.2})$$

where U are the random unitaries applied before measurement, $\vec{b}$ are the measurement outcomes (i.e. bit strings), σ is some quantum state, and $\mathcal{M}$ is the shadow channel (see Eq. (5.10)). To apply these expressions, we must consider the operator O on the Ising side of the duality.

Without loss of generality, assume that O is a single Pauli string. Then, from Eq. (5.23), $\mathcal{M}$ acts proportionally to identity and is trivial to invert yielding:

$$\mathcal{M}^{-1}(O) = \frac{3^{w_{xy}/2}}{f\alpha} O, \quad (\text{D.3})$$

where we recall that w_{xy} is the total number of X and Y operators in the (Ising-side) operator string O , f is the fraction of pairings $\pi \in P_V$ for which $\mathcal{M}_\pi(O) \neq 0$ (Eq. (5.24)), and α is an amplitude associated with permutations of Z and I operators within the operator O (Eq. (5.25)). Thus, the shadow norm is bounded as

$$\|O\|_{\text{shadow}}^2 = \max_{\sigma} \mathbb{E}_{U,\vec{b}} \langle \vec{b} | U \mathcal{M}^{-1}[O] U^\dagger | \vec{b} \rangle^2 \leq \frac{3^{w_{xy}/2}}{f\alpha} \|O\|_{\infty}^2, \quad (\text{D.4})$$

where we use that $\langle \vec{b} | U O U^\dagger | \vec{b} \rangle^2 \leq \lambda_{\max}(O)^2$, where $\lambda_{\max}(O)$ is the largest eigenvalue of O . Substituting this bound on the shadow norm into Eq. (D.4) yields Eq. (5.27).

D.2 Protocols with Fixed Boundary Conditions

In this Appendix, we explicitly describe the duality with fixed boundary conditions (FBC) (Section D.2.1), i.e. where the electric flux entering the system is fixed, and explain in detail the Global Dual Pairs protocol in this formulation, including a numerical example (Section D.2.2). The Local Dual Pairs and Dual Product protocols can be generalized along similar lines.

D.2.1 Duality with FBC

We now describe how the duality is modified in the presence of fixed boundary conditions (FBC). The case of periodic boundary conditions (PBC) is discussed in the main text and, for reference, is also summarized in Fig. D.1a and b. The LGT–Ising duality under FBC is illustrated in Fig. D.1c (LGT side) and d (Ising side).

A key difference from the PBC case is the absence of global superselection operators (V_x, V_y) . With FBC, there are $N_x(N_y - 1) + N_y(N_x - 1)$ qubit degrees of freedom residing on the links of the lattice, and $N_x N_y - 1$ independent Gauss-law constraints. The subtraction by one arises from the identity $\prod_s G_s = 1$, which includes both bulk and boundary sites s of the lattice. As a result, the number of remaining physical degrees of freedom is $N_x N_y - N_x - N_y + 1$. This matches the number of degrees of freedom in the dual Ising system, which consists of $(N_x - 1)(N_y - 1)$ dual spins located on the plaquettes. The explicit operator mapping is detailed in Fig. D.1c and d. For clarity, we label operators as belonging to the bulk or the boundary, depending on their location. The electric variables in the bulk, σ_l^x , and all magnetic terms $W_\square$ map to the dual theory just as in Eq. (5.4). However, electric operators σ_l^x on the boundary map to a single $X_\square$ operator. That is,

$$\Phi(\sigma_l^x) = X_\square, \quad (\text{for } l \text{ on boundary}), \quad (\text{D.5})$$

where $\square$ is the single plaquette adjacent to the boundary link l ; see Fig. D.1b for an illustration. Equation (D.5) implies that for σ_l^x operators located at the lattice corners, the two electric fields meeting at a corner both map to the same dual operator $X_\square$. This reflects that the corner σ_l^x operators are equivalent up to the action of the local Gauss law operator. The dual Hamiltonian for $\mathbb{Z}_2$ LGT in 2+1d with FBC can be written as,

$$H_{\text{dual}}^{\text{FBC}} = - \sum_{\square} Z_\square - g \sum_{l \in \text{boundary}} X_{\eta_l} - g \sum_{\square \in \text{blk}} (X_\square X_{\square-\hat{x}} + X_\square X_{\square-\hat{y}}),$$

where η_l runs over plaquettes adjacent to l , $\square \cap l \neq \emptyset$. Unlike the PBC case, there is no parity symmetry constraint in the dual, matching with the fact that on the LGT side it is no longer true that $\prod_\square W_\square = 1$.




Figure D.1: *Summary of LGT-Ising duality for PBC and FBC. a)* Hilbert space structure and operators for $\mathbb{Z}_2$ LGT with PBC and **b)** for the corresponding Ising dual which is subject to a parity constraint. Spheres mark one local spin-1/2 d.o.f. **c)** For FBC, the LGT includes bulk and boundary Gauss laws, as well as plaquette and electric operators. **d)** The dual Ising theory for FBC is no longer parity constraint but electric operators dualize differently in bulk and on the boundary.

D.2.2 Adapting the Global Dual Pairs Protocol to FBC

The Global Dual Pairs protocol with FBC is essentially identical to the PBC case discussed in the main text, except that the parity symmetry constraint is no longer imposed. Accordingly, for each pairing π , and for pairs $[ij] \in \pi$, we apply a randomizing unitary $U_{[ij]}$ drawn from a full two-qubit (approximate) unitary

2-design, i.e., without enforcing parity symmetry. For these pairs [408],

$$\mathcal{M}_{[ij]}[\Phi[\rho]] = \frac{1}{5} \left(\Phi[\rho] + \text{Tr}[\Phi[\rho]] \right). \quad (\text{D.6})$$

Thus, the action on an operator string $S_i S_j \in \{I, X, Y, Z\}^{\otimes 2}$ is

$$\mathcal{M}_{[ij]}[S_i S_j] = \begin{cases} S_i S_j, & S_i S_j = II \\ \frac{1}{5} S_i S_j, & \text{otherwise.} \end{cases} \quad (\text{D.7})$$

Consequently, $\mathcal{M}$ is diagonal and easily invertible. That is,

$$\mathcal{M}[S] = \tilde{\alpha} S, \quad (\text{D.8})$$

where $\tilde{\alpha}$ is the amplitude associated with S , defined as

$$\tilde{\alpha} = \frac{1}{|P_V|} \sum_{m=0}^{\lfloor k_{\text{dual}}/2 \rfloor} \frac{1}{5^{k_{\text{dual}}-m}} \underbrace{\binom{k_{\text{dual}}}{2m} |P_{2m}|}_{\text{double pairings}} \underbrace{\binom{V - k_{\text{dual}}}{k_{\text{dual}} - 2m} (k_{\text{dual}} - 2m)!}_{\text{single pairings}} \quad (\text{D.9})$$

where k_{dual} is the number of non-identity operators in S (i.e. the Pauli weight of S), and where the sum is over the number of pairings m that pair non-trivial operators with one another (“double pairings”). Unlike the PBC case, all pairings contribute and, thus, there is no FBC equivalent to the fraction f in Eqs. (5.23) and (5.24). Pairings that pair a single non-trivial operator with an identity operator are denoted a “single pairing.” Assuming that $k_{\text{dual}} = \Theta(1)$, and, using Stirling’s approximation we find that $\tilde{\alpha} \sim 1/\text{poly}(V)$. Thus, following a computation nearly identical to that in Section D.1, the sample cost remains polynomial in the system size, as in the PBC case.

Figure D.2 presents a numerical demonstration of the Global Dual Pairs protocol with FBC, implemented on a $(N_x, N_y) = (3, 3)$ lattice. We estimate a loop operator O_L in (a), and a ribbon operator O_R in panel (b). The FBC adaptation accurately reproduces the expected values across a broad range of coupling strengths. In addition, the convergence to the exact result is quicker for the loop observable O_L compared to the ribbon operator O_R , which while having higher weight on the $\mathbb{Z}_2$ side of the duality, has lower weight compared to O_R on the Ising side.




Figure D.2: *Implementation of Global Dual Pairs with Fixed Boundary Conditions* Example of **a)** loop and **b)** ribbon operators for $N_x = 3$ and $N_y = 3$ with FBC via the modified Global Dual Pairs protocol, compared to the exact solution of the $\mathbb{Z}_2$ LGT ground state. Error bars represent standard deviation of the mean over a single experiment of $\nu = 5000$ shots.

Appendix E

APPENDIX TO CHAPTER 7

E.1 ORP Details

E.1.1 Detector Construction From Syndromes

Syndrome measurements record the eigenvalue of a stabilizer at the time of measurement, and connecting this value to the underlying error mechanisms is done through the construction of detectors. A detector v_j is a parity of various syndrome measurements that has a deterministic value for a noiseless circuit [590], where $j = (g, b, r)$ is a “space-time” label denoting the stabilizer g , block b , and round r that the detector measures. The same construction applies to both X -type and Z -type detectors: the generator index g runs over the X stabilizers to build the X detectors and over the Z stabilizers to build the Z detectors, with each detector formed only from syndrome bits of a single stabilizer within a single block. The two stabilizer types are therefore structurally identical, and differ only through the “temporal boundaries”, or how detectors are constructed at initialization and when terminal measurements on the data qubits are made. In this case detectors are set by the logical state in which each block is prepared and the basis in which it is measured out. *No-reset circuits.* For the no-reset circuits used in this work, the ancilla measuring a given stabilizer is not reinitialized between cycles, so its outcome records the running parity of that stabilizer’s eigenvalue history rather than its instantaneous eigenvalue. The deterministic combination is then the parity of two syndrome measurements in the same code block two cycles apart [561],

$$v_{j+1} = s_{j+1} \oplus s_{j-1} , \tag{E.1}$$

where $j \pm 1$ is shorthand for $(g, b, r \pm 1)$, so that the intervening cycle is skipped. Taking the parity across two cycles cancels the accumulated history carried by the unreset ancilla and isolates the change in the stabilizer eigenvalue, ensuring that an isolated error flips only a small, bounded set of detectors. For the first detector with $j = 1$, the syndrome measurements are padded with a reference value of 0 (so no error has occurred at the beginning of the circuit). This padding renders the first detector deterministic for every stabilizer whose eigenvalue is fixed by the state preparation. In this work, each block of the $[[4, 2, 2]]$ code is initialized in

$|\overline{00}\rangle$, which is a simultaneous $+1$ eigenstate of both S_X and S_Z . Both stabilizer eigenvalues are therefore fixed at preparation, and the padding yields a deterministic first detector for the X -type and Z -type stabilizers alike. Errors during initialization will go undetected unless the preparation circuit is itself FT, however since $|GHZ(4)\rangle$ circuits are exceptionally shallow errors during this stage do not contribute very much to the overall logical error rate. Had the block instead been prepared in a bare computational-basis product state, only the Z -type stabilizers would be fixed by the preparation. The X -type eigenvalue on the other hand would be random and its first detector would form only once two in-circuit syndrome measurements are available and their relative parities are deterministic. At the final cycle the data-qubit readout reconstructs the stabilizers compatible with the measurement basis, and this reconstructed value plays the role of s_{j+1} in the boundary detector, closing the space-time volume. For R rounds, the final detector in the measurement basis reads

$$v_{R+1} = s_{R-1} \oplus s_R \oplus \left(\bigoplus_{d \in \text{supp}(g)} d \right), \quad (\text{E.2})$$

where $\text{supp}(g)$ are the data qubits on which stabilizer g is defined; in this work it is always the S_Z stabilizer as measurements are always be made in the Z basis. For the simulations with only a single stabilizer-measurement round, the two-cycle rule has no earlier in-circuit round to reference, and the 0-padding convention makes each detector coincide with its stabilizer measurement, $v_1 = s_1$. Because the GHZ initialization fixes both eigenvalues, this holds for the X -type and Z -type stabilizers alike, supplemented by the measurement-basis boundary detector,

$$v_2 = s_1 \oplus \left(\bigoplus_{d \in \text{supp}(S_Z)} d \right). \quad (\text{E.3})$$

Reset circuits. When the ancilla is instead reset to $|0\rangle$ after each syndrome measurement, its outcome reports the instantaneous stabilizer eigenvalue rather than a running parity. The detector then reduces to the parity of syndrome measurements in consecutive cycles,

$$v_j = s_{j-1} \oplus s_j, \quad (\text{E.4})$$

where the index convention is the same as in the no-reset case and the same padding convention supplies the reference value for the first detector. Similarly, the construction of the X and Z detectors and the initialization boundary are unchanged, and both stabilizer types again acquire a deterministic first detector from the GHZ

preparation. The only modification at the terminal boundary is that the reconstructed data value is compared against a single syndrome measurement round,

$$v_{R+1} = s_R \oplus \left(\bigoplus_{d \in \text{supp}(g)} d \right), \quad (\text{E.5})$$

since s_R is already the instantaneous eigenvalue. The remaining practical difference is in which detectors are flipped by an error: a measurement fault flips a pair of *adjacent* detectors, v_j and v_{j+1} , rather than the next-nearest pair v_j and v_{j+2} produced in the no-reset case, so the time-like edges connect neighboring rather than next-neighboring cycles.

E.1.2 Detector Selection

The following plateau-finding algorithm is used with ORP to determine which detectors should be used for postselection.

1. Given an observable $\overline{O}$, detectors v_j are ranked in order of decreasing $\Delta_j = |\langle \overline{O} \rangle_{v_j=0} - \langle \overline{O} \rangle_{v_j=1}|$.¹
2. A sweep over cutoff levels k is done, where each level k only retains shots for which none of the top- k detectors record an error. Final blockwise parity postselection is always applied regardless of k . This produces a sequence of postselected expectation values $\langle \overline{O} \rangle_k$ with bootstrapped uncertainties σ_k . The acceptance fraction among all blockwise parity-even shots f_k is recorded for each k .
3. A reference value is built by combining the most heavily filtered levels using an inverse-variance weighted average,

$$O_{\text{ref}} = \frac{\sum_{k \in \{k_{\text{ref}}\}} \langle \overline{O} \rangle_k / \sigma_k^2}{\sum_{k \in \{k_{\text{ref}}\}} 1 / \sigma_k^2}, \quad (\text{E.6})$$

where $\{k_{\text{ref}}\}$ includes levels k that satisfy $f_{\text{min}} \leq f_k \leq f_{\text{ref}}$.

4. Starting from large k and moving down, a plateau is found by looking for the longest run of k where $|\langle \overline{O} \rangle_k - O_{\text{ref}}| \leq z\sigma_k$.

¹Since $\overline{O}$ is defined up to conjugation by S_X and S_Z , the calculation of Δ_j depends on which data qubits are used to compute $\langle \overline{O} \rangle$.




Figure E.1: *Convergence of ORP on ibm.boston*. Four Trotter steps of size $\delta t = 0.4$ are applied to four code blocks, each initialized in the $|\overline{00}\rangle$ state, with couplings $g_x = g_z = 1$. The solid line shows $\langle \bar{Z}_{(0,0)} \rangle$ computed through ORP as a function of the number of detectors k used in the postselection criterion, with the noiseless MPS prediction given by the horizontal dashed line. The corresponding acceptance fraction is shown by the dotted line (right axis). The vertical red dotted line marks the value of k^* chosen by the plateau-finding algorithm.

5. The selected value k^* is the most permissive k in this run. If a sufficiently long plateau isn't found, the cut level is set to be $k^* = \arg \min_k [(\langle \bar{O} \rangle_k - O_{\text{ref}})^2 + \sigma_k^2]$.

This procedure is carried out on the held-out half of the results, and the resulting k^* is applied to the complementary half to report $\langle O \rangle_{k^*}$. This work uses the parameters $f_{\text{min}} = 0.005$, $f_{\text{ref}} = 0.05$, $z = 1$, although the quality of $\langle \bar{O} \rangle_k$ is observed to be similar for a range of parameters. Observables for all times t and logical qubits (b, l_b) are calculated using this algorithm. An example of the convergence of local expectation values for four blocks on ibm.boston as a function of the number of detectors used in the postselection criteria is shown in Fig. E.1. The plateau-finding algorithm successfully chooses a point within the vicinity of the noiseless MPS prediction, predicting $\langle \bar{Z} \rangle = 0.85(2)$ compared to an ideal value of 0.88. While the associated acceptance fraction decreases exponentially, choosing to filter only on k^* detectors does not leave the total ensemble exponentially small in system size.

E.2 Block Placement Selection

The goal of layout selection is to arrange the logical qubits so that they closely resemble some target geometry, in our case a regular square lattice that we call a “logical grid”. The heavy-hex quantum processor’s physical connectivity, combined with the choice of logical connectivity, creates a large search space of candidate logical grids. This can be formulated as a rectangle placement problem, since each $[[4, 2, 2]]$ block encodes two logical qubits, we represent it as a 2×1 rectangle occupying two sites of a square lattice. Two logical qubits couple when their blocks occupy adjacent lattice sites such that they are directly connected by a physical gate. We apply this formulation to the heavy-hex connectivity, though it generalizes to arbitrary device geometries.

We first enumerate all placements P_i of blocks onto the device that maximize the number of blocks used, subject to the heavy-hex connectivity and excluding a set of qubits with poor performance. For `ibm_boston` (avoiding four noisy qubits), we find 42 different embeddings P_i supporting 21 code blocks each, i.e., 42 logical qubits per embedding.

Within a given P_i , logical qubits within two blocks are allowed to couple if their blocks are adjacent or share a single “bridge” qubit on the device. An example of this is shown in Fig. E.2, where blocks 20 and 17 are directly adjacent, while blocks 17 and 14 are considered adjacent through the green bridge qubit. Relabeling logical qubits $(b, 0) \leftrightarrow (b, 1)$ or ancillas $a_x \leftrightarrow a_z$ within a block leaves the intra-block topology unchanged, so we treat these configurations as equivalent. For a given placement P_i , the realized edge set is

$$\begin{aligned}
 E(P_i) = & \underbrace{\{(b, 0), (b, 1)\}_{b \in P_i}}_{\text{intra-block, always present}} \\
 & \cup \{ \{(a, l_a), (b, l_b)\} : \\
 & \quad (a, l_a) \text{ adjacent to } (b, l_b) \forall b \in P_i \} ,
 \end{aligned} \tag{E.7}$$

i.e., the set of intra-block edges and available adjacent inter-block edges. The degree of a given vertex v is $\deg(v) = 1 + \#\{\text{inter-block edges}\}$, and we require $2 \leq \deg(v) \leq 4$ for all v .

Finding the best logical grid is a combinatorial optimization over the space of edge configurations, which is too large to search exhaustively. Instead, we use simulated annealing [591] to find the approximate solution. Each candidate configuration is scored as $S(E) = |E| - \lambda(\kappa(E) - 1)$. Here $\kappa(E)$ is the number of connected components and λ is a large positive coefficient that penalizes disconnected graphs. This




Figure E.2: *Example $[[4, 2, 2]]$ Code Block Placement onto Heavy-Hex Connectivity.* Several blocks (grey) are shown placed onto a heavy-hex connected quantum processor. Data qubits are given in black, Z- and X-type ancillas are shown in red and blue respectively, and a bridge qubit is shown in green.

form allows the algorithm to briefly pass through disconnected configurations to reach higher-scoring connected ones. In the annealing search, modifications E' to the current layout E are accepted with the probability $p_{\text{accept}} = \min(1, e^{(S(E') - S(E))/T})$, where $T(t) = T_0 (T_{\text{end}}/T_0)^{t/(N_{\text{it}}-1)}$ is the cooling schedule which specifies the search rate. A periodic reheating every R_h steps, $T \leftarrow \max(T, T_{\text{reheat}})$ to escape local minima is included. We use the parameters $T_0 = 4$, $T_{\text{end}} = 0.01$, $T_{\text{reheat}} = 2$, $R_h = 6000$, $N_{\text{it}} = 25000$.

The set of allowed updates are single-rectangle relocations, single rectangle rotations, two-rectangle swaps, and three-rectangle cyclic swaps. In addition, with probability 0.15 an under-realized block (one that is not yet coupled to all its available neighbors) is relocated to a spot that maximizes its realized edges. This prevents the algorithm from getting stuck on grids with thin chains attached. This process is repeated 500 times for every placement P_i with random seeds. The parameters for the simulated annealing are heuristically chosen to determine a suitable logical grid. With this algorithm, we find the grid shown in Fig. 7.1b), which has 66 logical edges (45 inter-block and 21 intra-block).² App. E.10 discusses other block placements that may be more favorable on similar heavy-hex quantum processors.

²In principle, the Trotter step depth could be optimized jointly with the layout. This work does not do so and instead fixes it in advance. Trotter step optimization is discussed in App. E.3.

E.3 Circuit Scheduling

Once the optimal logical grid is determined, a Trotter step of the Hamiltonian in Eq. (7.6) must be built. Naive placement of inter-block ZZ rotations results in deep circuits and long idle times. As shown later in this section, finding the optimal Trotter step for a given grid reduces to determining the best placement of inter-block gates $U_e(\delta t) = e^{-i\delta t/2\bar{Z}_{(a,l_a)}\bar{Z}_{(b,l_b)}}$. The logical operations within a single block are defined in Eq. (7.2). It is assumed that blocks start in the compute configuration shown in the center of Fig. 7.1c). In each block b , $\bar{Z}_{(b,0)}$ and $\bar{X}_{(b,1)}$ can be implemented directly in this configuration while $\bar{Z}_{(b,1)}$ and $\bar{X}_{(b,0)}$ require $\text{SWAP}(d_1, d_2)$. A block's frame $f \in \{0, 1\}$ records whether d_1 and d_2 have been swapped relative to their starting position. Together with intra-block gates, inter-block gates on both logical qubits within a block require at least one frame toggle per Trotter step.

Given two blocks a and b , and a single connection point between them (a direct connection or through a bridge qubit), the shallowest implementation of the unitary U_e applies CNOT fanout on both blocks (to accumulate the required parity), applies a Z rotation on the central qubit, and uncomputes the accumulated parities by applying the fanouts in reverse.³ An example of such a circuit between qubits $(17, 0)$ and $(20, 1)$ is shown in Fig. E.3b). The depths of the inter-block ZZ gadgets between all neighboring pairs $(a, l_a), (b, l_b)$ are shown in Table E.1. Three choices leave the logical circuit invariant, but may lower the depth: the orientation of block $o_b \in \{0, 1\}$ (whether $a_x \leftrightarrow a_z$ is applied), $l_b \in \{0, 1\}$ which specifies whether the $(b, 0) \leftrightarrow (b, 1)$ relabeling is applied, and w_e controlling which connection to use between two blocks when multiple exist. We parameterize the schedule with $A = \{o, l, w\}$ for all blocks and edges in the logical grid. Within a block, couplings in the same frame share the same accumulated Z -parity, so this parity is computed once per frame segment via a single fanout.

Scheduling for a fixed A is itself a combinatorial optimization problem: even with the gate placements fixed, the ordering and layering of two-qubit gates into a minimum-depth circuit has a search space too large to explore exhaustively. The schedule for a given A is determined using the constrained-satisfaction solver CP-SAT from the package `ortools` [592, 593]. This is possible because of commutation between the terms of each Trotter step. Together with fixing $f = 0$ at the start of each Trotter step and placement of single qubit gates described below, this constrains the Trotter step to a single consistent ordering on all blocks, leaving

³In the case where two adjacent blocks have two connections, such as blocks 14 and 19 in Fig. E.2, the algorithm selects a schedule that may use both connections to minimize overall circuit depth.




Figure E.3: *Encoded Trotter Step Circuits*. a) An example placement of two blocks in the compute configuration. Data qubits are given in black, Z - and X -type ancillas are shown in red and blue respectively, and a bridge qubit is shown in green. b) The CNOT fanout necessary for implementing a $\bar{R}_{ZZ}$ between logical qubits $(17, 0)$ and $(20, 1)$ (dark orange), with the associated SWAP gate converted to CNOTs. c) A Trotter step involving logical qubits in blocks 17 and 20, with the inter-block gate between qubits $(17, 0)$ and $(20, 1)$. Single-logical-qubit gates (yellow) and intra-block $\bar{R}_{ZZ}$ (light orange) are folded into the CNOT fanout required for the inter-block $\bar{R}_{ZZ}$ gate, the last part of which is shown in dark orange. Block 20 exits the Trotter step in frame $f = 1$.

the depth as the only free objective. The objective minimizes the total depth with the additional degree of freedom that the exit frame need not be $f = 0$. The parameters A are determined by coordinate descent.

A full Trotter step of the Hamiltonian (7.6) can be implemented at no extra depth using the schedule of inter-block $\bar{R}_{ZZ}$ gates, provided the logical grid satisfies the conditions in App. E.2. The single-logical qubit gates and the intra-block two-logical qubit gates can be folded into the inter-block $\bar{R}_{ZZ}$ gates. These fanouts already collect the XX and ZZ parities required for single qubit $\bar{X}$ - and $\bar{Z}$ -rotations. Since each vertex v on the selected lattice has $2 \leq \deg(v) \leq 4$, all four single-logical-qubit rotations can be implemented in all blocks with the addition of only single-physical-qubit rotations. The parity required for the intra-block $\bar{R}_{ZZ}$ is also collected in each fanout at the time of the $\text{SWAP}(d_1, d_2)$ application, so the addition of this interaction also carries no extra two-qubit depth cost. The specific location at which each single-logical-qubit gate and the intra-block $\bar{R}_{ZZ}$ is inserted depends on both the block's current frame and the structure of the particular inter-block ZZ gadget being used. An example of a Trotter step spanning two adjacent blocks is shown in Fig. E.3c), with the inter-qubit and intra-qubit gates highlighted. Since a Trotter step does not necessarily return every block to $f = 0$, consecutive odd-numbered Trotter steps are appended in reverse order, so that frames of all blocks are restored to $f = 0$ after even n_T .

Blocks (a, b)	Two-qubit gate depth				Blocks (a, b)	Two-qubit gate depth			
	(l_a, l_b)					(l_a, l_b)			
	(0,0)	(0,1)	(1,0)	(1,1)		(0,0)	(0,1)	(1,0)	(1,1)
(0, 1)	6	8	8	10	(10, 13)	4	8	6	8
(0, 2)	4	8	6	8	(11, 12)	6	8	8	10
(1, 3)	6	8	8	10	(11, 14)	4	8	6	8
(2, 4)	6	8	8	10	(11, 15)	6	8	8	10
(2, 9)	6	10	8	10	(12, 15)	6	8	8	10
(3, 5)	4	8	6	8	(13, 14)	6	8	8	10
(3, 7)	4	8	6	8	(13, 16)	4	8	6	8
(4, 9)	4	8	6	8	(13, 19)	6	10	8	10
(5, 6)	6	8	8	10	(14, 15)	4	8	6	8
(5, 7)	6	8	8	10	(14, 17)	8	10	10	12
(5, 8)	6	8	8	10	(14, 19)	6	10	8	10
(6, 8)	6	8	8	10	(15, 17)	4	8	6	8
(6, 9)	4	8	6	8	(16, 18)	4	8	6	8
(6, 11)	6	8	8	10	(16, 19)	6	8	8	10
(7, 8)	6	8	8	10	(17, 19)	6	10	8	10
(8, 11)	6	8	8	10	(17, 20)	4	8	6	8
(8, 12)	4	8	6	8	(19, 20)	6	10	8	10
(9, 10)	6	8	8	10					

Table E.1: The inter-block R_{ZZ} CNOT fanout depth for pairs of logical qubits (l_a, l_b) between pairs of code blocks (a, b) corresponding to the block placement shown in Fig. 7.1b). The depth includes CNOTs used for SWAPs and bridge qubits where necessary.

E.4 Quantum Simulation Details

While more efficient time evolution techniques and Trotter step implementations are possible for unencoded circuits and MPS simulations, we execute circuits with the same Trotter order to hold observable expectation values fixed. Further, results from encoded and unencoded runs with the same number of shots are compared, resulting in biased unencoded results with small error bars. The “extra” shots in the unencoded runs could be instead used to efficiently mitigate errors in the results through noise-learning methods [40–42, 199, 213, 293, 294]. Ideally, syndromes would be measured after every operation that could spread errors to multiple qubits. In practice this is prohibitively costly due to limited device coherence time. Tables E.2 and E.3 show the overhead in terms of depth, gate counts, and coherence time for encoded and unencoded runs in 1D and 2D,

	Unencoded	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$
# measurements	42	126	168	252	420	756
Measurement duration (dt)	446	446	671	1121	2021	3821
Total duration (dt)	6017	10171	11139	13075	16947	26451
Depth	116	272	297	347	447	719
# 2-qubit gates	1126	3266	4112	5804	9188	16404
# shots	32000	32000	32000	32000	32000	32000

Table E.2: *Details of 1D Chain Simulations on ibm_boston at $t = 8$.* The columns compare unencoded counts to counts in encoded runs with various numbers of syndrome extraction rounds R . The second row shows the total number of measurements (both mid-circuit and terminal). The maximum coherence time used by measurements on any qubit in units of $dt = 4 \times 10^{-9}$ s is given in the third row. The fourth row gives the total duration of the circuit. Depths, total numbers of two-qubit gates, and numbers of shots executed are given in the fifth, sixth and seventh rows respectively

	Unencoded	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$
# measurements	42	126	168	252	420	756
Measurement duration (dt)	446	446	671	1121	2021	3821
Total duration (dt)	19135	13931	14935	16943	20959	30591
Depth	626	443	470	524	632	912
# 2-qubit gates	4286	5540	6414	8162	11658	19418
# shots	32000	32000	32000	32000	32000	32000

Table E.3: *Details of 2D Grid Simulations on ibm_boston at $t = 8$.* The same rows and columns as Table E.2 but for the 2D simulations presented in Sec. 7.3.

respectively. Given that the two-qubit gate duration is 17 dt to 22 dt on ibm_boston, both the encoded and unencoded circuits have significant idle time. The circuits in this work do not exceed the median coherence




Figure E.4: *Circuits for Switching Between Compute and Syndrome Block Configurations.* a) The compute and syndrome block configurations as shown in the center of Fig. 7.1c). Data qubits are shown in black, Z -type ancillas are shown in red, and X -type ancillas are shown in blue. b) The circuit to switch from the compute to the syndrome configuration. c) The corresponding circuit to switch from the syndrome to the compute configuration. The syndrome extraction circuit applies a $\text{SWAP}(d_2, d_3)$, which is taken into account in the configuration switching circuit.

times on `ibm_boston`.⁴ However, degradation with depth is still observed, dominated by qubits with low T_1/T_2 times, high measurement error rates, or regions with faulty two-qubit gates. The 1D encoded circuits are two to six times deeper than their unencoded counterparts. This is similarly reflected in the durations and total gate counts. Table E.3 shows the geometric overhead of embedding a square lattice onto heavy-hex connectivity for the unencoded circuits, which are ~ 1.5 times deeper than the shallowest encoded circuit. This is one of the driving factors of their degrading performance as seen in Fig. 7.3, and is also reflected in inflated gate counts and durations.

Switching between the compute and syndrome block configuration introduces a depth ~ 15 overhead at each syndrome measurement.⁵ Figure E.4 shows the SWAP circuits necessary to switch between the compute and syndrome configuration necessary for the syndrome extraction circuit, which are shown at the bottom right of Fig. 7.1c).

⁴The median T_1 and T_2 times on `ibm_boston` were $234\mu\text{s}$ and $275\mu\text{s}$ when the data in this work was taken.

⁵Some of this depth is reduced by CNOT cancellations with the Trotter steps in practice.

E.5 Classical Simulation Details

Matrix product state simulations are used as a classical comparison to the results from `ibm_boston`. Despite the quench dynamics, for the 1D chain this is efficient because of the small system size, while in 2D the cost grows sharply. Mapping an $L \times L$ grid onto a 1D MPS introduces long-range couplings, and the entanglement across a cut scales with the linear cut length. Faithfully representing a general state on a 2D lattice requires a `max_bond` growing roughly exponentially in L , with per-step cost $\mathcal{O}(L^2 \text{max_bond}^3)$. The required `max_bond` is set by the entanglement, which grows in time as the quench spreads correlations. While more efficient classical simulation techniques for 2D lattices exist, such as tree tensor networks, belief propagation, and sparse Pauli dynamics, MPS is used for simplicity.

Time evolution is simulated by applying Trotter steps with `max_bond` = 1500 and singular-value cutoff of `cutoff` = 10^{-8} . Two sets of simulations are run. Late-time simulations with $\delta t = 0.1$ with up to $n_T = 500$ (corresponding to Fig. E.16) are used to confirm the presence of physical effects absent at early times and washed out by large δt . While truncation errors in 1D simulations are negligible, significant error accumulates for a 2D lattice. The missing links of the logical connectivity grid shown in Fig. 7.2a) modify the conventional “snake” MPS ordering, which likely modifies the convergence. The error, quantified by the norm of the MPS wavefunction at $n_T = 100$ (the latest time in Fig. E.16) is $\mathcal{O}(10^{-3})$. While this value is relatively large, changing `max_bond` from 1200 to 1500 only modifies observables at $\mathcal{O}(10^{-2})$ which is sufficient for this qualitative study. The latter simulations execute the exact unencoded circuits run on `ibm_boston` with $\delta t = 0.5$ and up to $n_T = 20$ Trotter steps for a noiseless comparison to device results. The truncation errors in these simulations, both for 1D and 2D lattices, are negligible. Larger values of δt deform the dispersion relation and add an additional quench to the dynamics [243, 568, 594]. The localized regime ($g_z = 1.5$) is observed to spread entanglement less than the melting regime ($g_z = 0.5$). The results of the 5×5 simulations displayed in Fig. E.17 were computed by `qiskit` statevector simulation.

E.6 Details on Filtering Metrics

This appendix presents a detailed analysis of the filtering metric Δ_j used in ORP. It also introduces β_j , an alternative metric based on linear regression. First, both of these quantities are defined in App. E.6.1. Then, a worked example on a small system of one code block and two rounds of syndrome measurements is given in App. E.6.2 to provide a concrete implementation of these metrics. After, App. E.6.3 introduces a general,

high-level description of errors in terms of the correlations among detector measurements and correlations between detector and observable measurements. The quantities Δ_j and β_j are then expressed in terms of the general model parameters in App. E.6.4, demonstrating the types of errors that they filter. These expressions simplify substantially when considered in a weak-noise limit, and can be seen to directly connect to the probability of an error detection event propagating to a logical error on an observable. Finally, App. E.6.5 connects the detectors chosen based on the Δ_j ranking, as done in ORP, to the backwards lightcone of a target local observable using Pauli Propagation.

In what follows, the observables considered are Pauli strings and are denoted by $\bar{O}$ when considered as operators, with expectation values denoted as $\langle \bar{O} \rangle$. The result of a single measurement of such an observable will be denoted as Q (Q_i when the measurement is a member of an ensemble, where i is the shot index), which can take values $Q \in \{-1, 1\}$. Expectation values and conditional expectation values of Q over an ensemble are denoted as $\mathbb{E}[\cdot]$ and $\mathbb{E}[\cdot|\cdot]$. More general observables can be decomposed into a weighted sum of Pauli strings, and many of the following results can be extended to this context.

E.6.1 Definitions of Δ_j and β_j

Recall the definition of Δ_j from Eq. (7.5), rewritten as an expectation value over shots

$$\Delta_j = \left| \mathbb{E}[Q|v_j = 0] - \mathbb{E}[Q|v_j = 1] \right|. \quad (\text{E.8})$$

This quantity is calculated by taking the difference in the expectation values of observable Q over shots postselected on the detector v_j not detecting/detecting an error respectively. To define β_j , let X_i be a length N_d vector of detector outcomes for a single shot, so that X is an $N_s \times (N_d + 1)$ matrix where N_d is the number of detectors and N_s is the number of shots. The first column of X is all 1s, and captures detector-independent behavior. The relationship between a particular detector outcome and the value of the logical observable can be modeled as

$$Q_i = \beta_0 + \sum_{j=1}^{N_d} \beta_j X_{ij} + \varepsilon_i, \quad (\text{E.9})$$

where β_0 is a shot-independent bias term and ε_i is an error term that accounts for observable behavior that is not explainable by the detector values. This relation is represented as a matrix equation $\vec{Q} = X\vec{\beta} + \vec{\varepsilon}$ where $\vec{Q}$ collects all Q_i into a vector. To minimize the cumulative error across all shots, we seek β^* which is the

solution of the following minimization

$$\beta^* = \arg \min_{\beta} L(\vec{Q}, X, \beta), \quad (\text{E.10})$$

where $L(\vec{Q}, X, \beta) = \|\vec{Q} - X\beta\|_2^2$ is the 2-norm of the error vector ε . The solution to Eq. (E.10) satisfies the *normal equation*

$$X^T X \beta^* = X^T \vec{Q}. \quad (\text{E.11})$$

The metric β_j is then the magnitude of the j -th element of the vector $\beta^* = (X^T X)^{-1} X^T \vec{Q}$.

E.6.2 A Worked Example of Calculating Δ_j and β_j

As an explicit example, consider a circuit with one $[[4, 2, 2]]$ code block and two rounds of S_Z syndrome measurements and no S_X measurements (using one ancilla without resets, with measurements recorded to classical bits) corresponding to three detectors. The initial state is $|\overline{00}\rangle$ and the observable of interest is $\overline{O} = \overline{Z}_0$. For each shot, we consider the different times during syndrome extraction that an X error can occur. For the set of $N_s = 4$ measurements in this example, the outcomes on codeblock A are:

$$\left\{ [\text{syndrome} \mid \text{data}] \right\} = \left\{ [0\ 0 \mid 0\ 0\ 0\ 0], [1\ 0 \mid 1\ 0\ 0\ 0], [0\ 1 \mid 0\ 0\ 0\ 1], [0\ 0 \mid 1\ 0\ 0\ 0] \right\}, \quad (\text{E.12})$$

where the notation $[\text{syndrome} \mid \text{data}]$ describes a measurement outcome. In the first shot, no error occurred. In the second shot, an error occurred before the first syndrome measurement. In the third shot, an error occurred between the first and second round of extraction. In the fourth shot, an error occurred after the second round of extraction. The detectors and observable outcomes (with notation $[v_1\ v_2\ v_3 \mid Q_i]$) using the construction in App. E.1.1 are then given by,

$$[0\ 0\ 0 \mid +1], [1\ 0\ 0 \mid -1], [0\ 1\ 0 \mid +1], [0\ 0\ 1 \mid -1]. \quad (\text{E.13})$$

The values of Δ_j can then be calculated as,⁶

$$\Delta_1 = 4/3, \quad \Delta_2 = 4/3, \quad \Delta_3 = 4/3. \quad (\text{E.14})$$

⁶This example is a pure memory experiment on a single block, and so each detector heralds errors equally impacting the logical information causing all Δ_j to be equal. Adding more blocks or nFT logical rotations would introduce a hierarchy in the Δ_j , which is the case for the simulations presented in the main text.

To calculate the β_j , the detector matrix X and observable vector $\vec{Q}$ are constructed as,

$$X = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix}, \quad \vec{Q} = \begin{bmatrix} +1 \\ -1 \\ +1 \\ -1 \end{bmatrix}, \quad (\text{E.15})$$

and the vector of β_j is given by $\beta^* = [(X^T X)^{-1} X^T Q]_{j>0} = -[2, 0, 2]$.

E.6.3 Model of Error Mechanisms

This section presents a higher-level description of the noise on a quantum processor, inspired by the treatments in Ref. [512, 595]. This model stands in contrast to a circuit-level noise model, which describes the probabilities of a specific error occurring, where the probabilities of flipping a detector or causing a logical error are derived quantities. Instead, this model treats error mechanisms that flip detectors as the fundamental objects and directly assigns probabilities to these events, which is favorable when compared to a circuit-level noise model because the probability of a detector flip can be directly calculated from the raw device data. We will refer to the underlying causes of detector flips as ‘mechanisms’ to emphasize that detector flips are due to the aggregate effect of a number of different underlying physical errors [562]. In general there is no single underlying process that can be identified with the flip of a detector on a quantum computer.

Let $G = (V, E)$ be a graph describing the errors, where V is a set of nodes and E are the edges. Each node $v_j \in V$ corresponds to a detector and can be flipped by error mechanisms $e_j, e_{ij} \in E$ with $e_j/e_{ij} = \text{Bernoulli}(p_j/p_{ij})$ respectively. A single index j corresponds to a flip of detector v_j if $e_j = 1$ and the two indices (i, j) correspond to different “pair” processes flipping two detectors v_i, v_j if $e_{ij} = 1$. Triplet processes (i, j, k) and higher could be included to enhance the expressivity of this model, adding hyperedges to the error graph G , but analyzing these higher-order structures is beyond the scope of the current work. For brevity, both single node j and pair processes (i, j) will be denoted by a single index m . The support $\text{supp}(e_m)$ of an error mechanism e_m is defined by the set of detectors that are flipped when that mechanism occurs.

If an error mechanism occurs, it flips the value of an observable Q with probability q_m . For an individual

physical error in a Clifford circuit with Pauli noise, the observable flip probability is deterministic, $q_m \in \{0, 1\}$, corresponding to whether the resulting error flips the operator. The mechanisms of the present model are cumulative: each (p_m, q_m) aggregates all microscopic faults sharing the same detector signature, with p_m the combined rate of the class and q_m its occurrence-weighted average flip probability. An intermediate value of q_m therefore arises whenever the errors identified by a detector induce faults with different logical actions, and is expected even in Clifford experiments, for instance when a Pauli error coincides in detector signature with a leakage or crosstalk process. Intermediate values likewise arise for individual faults under non-Clifford gates.

The error mechanisms are taken to be mutually independent. The measured value of a logical observable $\overline{O}$ for a single shot is then represented as

$$Q = Q_I \prod_m (-1)^{f_m}, \quad (\text{E.16})$$

where $Q_I \in \{-1, 1\}$ is the noiseless value of Q ⁷ and $f_m \in \{0, 1\}$ correspond to parity flips of the measured observable so that $Q \in \{-1, 1\}$. The noiseless value $Q_I = -1$ with probability λ depends on the encoded wavefunction and $f_m = 1$ indicates process m has occurred and flipped the logical observable, i.e. $P(f_m = 1) = q_m p_m$. Separating the measured Q into its ideal part Q_I and associated parity flips f_m allows one to treat the effect of noise independently of the encoded information in the state. This is an approximation that is only strictly true for Clifford circuits and restricted types of noise, e.g. stochastic Pauli noise. Therefore the circuits presented in the main text will exhibit correlations between detector and observable values that are not captured by this model. However, its simplicity affords both analytical tractability and insight, and it is empirically found to capture certain features of device noise like dependence on the magnitude of $\langle \overline{O} \rangle$.

E.6.4 Expressing Δ_j and β_j in Terms of Detector Probabilities

Calculations for Δ_j

This section expresses Eq. (E.8) in terms of the probabilities p_m, q_m . A detector v_j can be written in terms of the error mechanisms e_m that it detects as

$$v_j = \bigoplus_{e_m \in E | j \in \text{supp}(e_m)} e_m, \quad (\text{E.17})$$

⁷ $\mathbb{E}[Q_I]$ approaches the noiseless value of the observable $\langle \overline{O} \rangle$ for a large ensemble size.

where the condition $e_m \in E \mid j \in \text{supp}(e_m)$ will be denoted as $m \ni j$. In words, v_j is defined the relative parity of the error mechanisms whose support contains it. First, we seek to solve for the expectation value $\mathbb{E}[Q|v_j = b]$ where $b \in \{0, 1\}$. We begin by separating the product over m in Eq. (E.16) into those whose support $\text{supp}(e_m)$ contains j and those where it does not

$$\begin{aligned} \mathbb{E}[Q|v_j = b] &= \mathbb{E} \left[Q_I \prod_{m \not\ni j} (-1)^{f_m} \right] \mathbb{E} \left[\prod_{m \ni j} (-1)^{f_m} \mid v_j = b \right] \\ &= \langle \overline{O} \rangle B_j \times \mathbb{E} \left[\prod_{m \ni j} (-1)^{f_m} \mid v_j = b \right]. \end{aligned} \quad (\text{E.18})$$

Here, $\langle \overline{O} \rangle = 1 - 2\lambda$ and $B_j = \prod_{m \not\ni j} (1 - 2p_m q_m)$ denotes the background effect of the error mechanisms that do not touch detector j . Thus, we must simplify the remaining expectation value in Eq. (E.18). Consider a single mechanism e_n whose support contains j . The expectation value is written in terms of the conditional outcomes of e_n ,

$$\begin{aligned} \mathbb{E} \left[(-1)^{\sum_{m \ni j} f_m} \mid v_j = b \right] &= \sum_{c \in \{0,1\}} P(e_n = c \mid v_j = b) (1 - 2q_n)^c \mathbb{E} \left[(-1)^{\sum_{m \ni j, m \neq n} f_m} \mid \bigoplus_{m \ni j, m \neq n} e_m = b \oplus c \right] \\ &= P_0(e_n \mid v_j = b) \mathbb{E} \left[(-1)^{\sum_{m \ni j, m \neq n} f_m} \mid \bigoplus_{m \ni j, m \neq n} e_m = b \right] \\ &\quad + P_1(e_n \mid v_j = b) (1 - 2q_n) \mathbb{E} \left[(-1)^{\sum_{m \ni j, m \neq n} f_m} \mid \bigoplus_{m \ni j, m \neq n} e_m = \bar{b} \right], \end{aligned} \quad (\text{E.19})$$

where $P_b(v) = P(v = b)$. In the first line, the total parity has been conditioned on the value $e_n = c$, which shifts the target parity to $b \oplus c$ and introduces a factor $(1 - 2q_n)^c$; here $\oplus$ denotes addition modulo 2. In the second line, $\bar{b}$ denotes the flipped value of b . The conditional probabilities are solved for by using

$$P_c(e_n | v_j = b) = \frac{P(e_n = c, v_j = b)}{P(v_j = b)} = (1 - c - p_n + 2cp_n) \frac{P_{b \oplus c}(v'_j)}{P_b(v_j)}, \quad (\text{E.20})$$

where the notation v'_j denotes the detector with mechanism n removed. Both these quantities can be solved for by observing that

$$P_b(v_j) + P_{\bar{b}}(v_j) = 1, \quad P_b(v_j) - P_{\bar{b}}(v_j) = (-1)^b \mathbb{E}[(-1)^{v_j}], \quad (\text{E.21})$$

with

$$\mathbb{E}[(-1)^{v_j}] = \prod_{m \ni j} \mathbb{E}[(-1)^{e_m}] = \prod_{m \ni j} (1 - 2p_m), \quad (\text{E.22})$$

implying that

$$P_b(v_j) = \frac{1}{2} \left(1 + (-1)^b \prod_{m \ni j} (1 - 2p_m) \right). \quad (\text{E.23})$$

Returning to evaluating Eq. (E.19), the process of removing a node is now repeated. Denote the detector with k mechanisms removed as $v_j^{(k)} = \bigoplus_{m \ni j-k} e_m$, with $v_j^{(0)} = v_j$, and the expectation value with k mechanisms removed as

$$E_b^{(k)} = \mathbb{E} \left[(-1)^{\sum_{m \ni j-k} f_m} \mid v_j^{(k)} = b \right], \quad (\text{E.24})$$

where Eq. (E.19) denotes the first step in this iteration. Generalizing to all nodes, the iterative equation reads

$$P_b(v_j^{(k)}) E_b^{(k)} = (1 - p_{n_k}) P_b(v_j^{(k+1)}) E_b^{(k+1)} + p_{n_k} (1 - 2q_{n_k}) P_b(v_j^{(k+1)}) E_b^{(k+1)} \quad \text{for } k \geq 0, \quad (\text{E.25})$$

where n_k denotes the removed mechanism. Note that the R.H.S of this equation only refers to the mechanisms that have j in their support, without the removed mechanism n_k . For both values of b , Eq. (E.25) can be written as a matrix equation,

$$V^{(k)} = M_{n_k} V^{(k+1)}, \quad V^{(k)} = \begin{bmatrix} P_0(v_j^{(k)}) E_0^{(k)} \\ P_1(v_j^{(k)}) E_1^{(k)} \end{bmatrix}, \quad M_{n_k} = \begin{bmatrix} (1 - p_{n_k}) & (1 - 2q_{n_k}) p_{n_k} \\ (1 - 2q_{n_k}) p_{n_k} & (1 - p_{n_k}) \end{bmatrix}. \quad (\text{E.26})$$

After all M mechanisms containing j have been removed, the parity is even and the sum vanishes, giving the base case $V^{(M)} = (1, 0)^T$. The matrix M_{n_k} can be diagonalized by the basis $E_{\pm}^{(k)} = P_0(v_j^{(k)}) E_0^{(k)} \pm P_1(v_j^{(k)}) E_1^{(k)}$, with $E_{\pm}^{(M)} = 1$, allowing the iteration relation to be solved exactly,

$$E_{\pm}^{(0)} = \prod_{m \ni j} \left[(1 - p_m) \pm (1 - 2q_m) p_m \right]. \quad (\text{E.27})$$

Rewriting this in terms of the original $E_0^{(0)}$ and $E_1^{(0)}$, the desired expectation value is then given by,

$$\mathbb{E}[Q|v_j = b] = \langle \overline{O} \rangle B_j \frac{\prod_{m \ni j} (1 - 2p_m q_m) + (-1)^b \prod_{m \ni j} (1 - 2p_m (1 - q_m))}{1 + (-1)^b \prod_{m \ni j} 1 - 2p_m}. \quad (\text{E.28})$$

Finally, the filtering metric is obtained by subtracting the two conditional expectation values. Denote the three products over mechanisms containing j as,

$$A_j = \prod_{m \ni j} (1 - 2p_m q_m), \quad C_j = \prod_{m \ni j} (1 - 2p_m(1 - q_m)), \quad D_j = \prod_{m \ni j} (1 - 2p_m), \quad (\text{E.29})$$

so that

$$\mathbb{E}[Q|v_j = b] = \langle \overline{O} \rangle B_j \frac{(A_j + (-1)^b C_j)}{(1 + (-1)^b D_j)}. \quad (\text{E.30})$$

Subtracting the $b = 0$ case from the $b = 1$ case gives

$$\Delta_j = \left| 2\langle \overline{O} \rangle B_j \frac{A_j D_j - C_j}{1 - D_j^2} \right|. \quad (\text{E.31})$$

The denominator is directly measurable: comparing with the parity probability derived above, $D_j = 1 - 2P_1(v_j)$, so that $1 - D_j^2 = 4P_1(v_j)(1 - P_1(v_j))$. Expanding to leading order in the mechanism probabilities, $A_j D_j - C_j \approx -4 \sum_{m \ni j} p_m q_m$ and $1 - D_j^2 \approx 4 \sum_{m \ni j} p_m$, giving

$$\Delta_j \approx \left| 2\langle \overline{O} \rangle \left(\sum_{m \ni j} p_m q_m \right) / \left(\sum_{m \ni j} p_m \right) \right|. \quad (\text{E.32})$$

The denominator can be viewed as the cumulative “firing” rate for all processes whose support contains detector v_j , while the numerator is this same quantity weighted by the probability of that process causing a logical error. Thus, Δ_j takes on the interpretation of an effective probability of all processes affecting v_j that cause a logical error, weighted by $\langle \overline{O} \rangle$. As such, detectors that take on higher values of Δ_j should be postselected on first, as these are the most detrimental to the encoded information. This forms the basis of ORP.

This analysis reveals that the scale of Δ_j is set by the value of the logical observable and the background factor B_j . If $\langle \overline{O} \rangle \ll 1$ or if the background attenuation is large, i.e., $B_j \ll 1$, and all the individual firing rates p_m are close to 1, then every value of Δ_j will take on parametrically small values. Due to finite measurement statistics, it can become challenging to resolve differences at this scale, reducing the effectiveness of Δ_j as a quality metric distinguishing different detectors.

Calculations for β_j

In the limit of many shots, $\frac{1}{N_s} X^T X$ and $\frac{1}{N_s} X^T \vec{Q}$ converge to matrices of moments. The normal equation (Eq. (E.11)) can be expressed component-wise explicitly in terms of the moments of detectors v_j and

observable Q

$$\begin{bmatrix} 1 & \mathbb{E}[v_k]^T \\ \mathbb{E}[v_j] & \mathbb{E}[v_j v_k] \end{bmatrix} \begin{bmatrix} \beta_0^* \\ \beta_k^* \end{bmatrix} = \begin{bmatrix} \mathbb{E}[Q] \\ \mathbb{E}[Q v_j] \end{bmatrix}, \quad j, k = 1, \dots, N_d. \quad (\text{E.33})$$

Eliminating $\beta_0^* = \mathbb{E}[Q] - \sum_k \mathbb{E}[v_k] \beta_k^*$ from the remaining rows reduces this to

$$G \beta^* = a, \quad (\text{E.34})$$

for the connected two point function $G_{jk} = \mathbb{E}[v_j v_k] - \mathbb{E}[v_j] \mathbb{E}[v_k]$ and $a_j = \mathbb{E}[Q v_j] - \mathbb{E}[Q] \mathbb{E}[v_j]$, with the bias recovered from the eliminated row. Each moment is evaluated using the products derived above. Rewriting the detectors in terms of parities, $v_j = \frac{1}{2}(1 - (-1)^{v_j})$, the one-point functions are

$$\mathbb{E}[v_j] = \frac{1 - D_j}{2}, \quad \mathbb{E}[Q] = \langle \overline{O} \rangle B_j A_j, \quad (\text{E.35})$$

where the second expression is independent of the choice of j . For the two-point function, the shared mechanism e_{jk} cancels from the parity $v_j \oplus v_k$, so that $\mathbb{E}[(-1)^{v_j \oplus v_k}] = D_j D_k / (1 - 2p_{jk})^2$, and

$$\mathbb{E}[v_j v_k] = \frac{1}{4} \left(1 - D_j - D_k + \frac{D_j D_k}{(1 - 2p_{jk})^2} \right), \quad \mathbb{E}[v_j^2] = \mathbb{E}[v_j] = \frac{1 - D_j}{2}, \quad (\text{E.36})$$

where the diagonal follows from $v_j^2 = v_j$ for a detector. The cross moment follows from the conditional expectation derived above, or directly from $\mathbb{E}[Q(-1)^{v_j}] = \langle \overline{O} \rangle B_j C_j$,

$$\mathbb{E}[Q v_j] = \frac{\langle \overline{O} \rangle B_j}{2} (A_j - C_j). \quad (\text{E.37})$$

The connected components are then given as,

$$G_{jk} = \frac{D_j D_k}{4} \left(\frac{1}{(1 - 2p_{jk})^2} - 1 \right) \quad (j \neq k), \quad G_{jj} = \frac{1 - D_j^2}{4}, \quad a_j = \frac{\langle \overline{O} \rangle B_j}{2} (A_j D_j - C_j). \quad (\text{E.38})$$

The limit of weak noise, $p_m \ll 1$, is used to simplify the expressions. Expanding to leading order, $D_j \approx 1 - 2 \sum_{m \ni j} p_m$ and $(1 - 2p_{jk})^{-2} - 1 \approx 4p_{jk}$, so that

$$G_{jk} \approx p_{jk} \quad (j \neq k), \quad G_{jj} \approx p_j + \sum_{k \neq j} p_{jk}, \quad a_j \approx -2 \langle \overline{O} \rangle \left(p_j q_j + \sum_{k \neq j} p_{jk} q_{jk} \right), \quad (\text{E.39})$$

where in a_j the terms carrying rates alone cancel between $A_j D_j$ and C_j , leaving the flip-weighted sums. The normal equation then reads

$$\left(p_j + \sum_{k \neq j} p_{jk} \right) \beta_j^* + \sum_{k \neq j} p_{jk} \beta_k^* = -2 \langle \overline{O} \rangle \left(p_j q_j + \sum_{k \neq j} p_{jk} q_{jk} \right). \quad (\text{E.40})$$




Figure E.5: *Comparison of Δ_j and β_j Predictions on `ibm_boston`.* The metrics Δ_j and β_j are calculated using Eq. (E.8) and Eq. (E.11) respectively. The two solid lines show the predicted value of $\langle \overline{Z} \rangle$ as a function of the number k of detectors used in postselection for Δ_j (dark) and β_j (light). The horizontal dashed line is the noiseless MPS prediction. The vertical dashed lines correspond to the value of k^* returned from the plateau finding algorithm using Δ_j (dark) and β_j (light). The expectation value is from four Trotter steps on four code blocks.

Here, the RHS is the same magnitude as the numerator of Δ_j for the same detector. Furthermore, omitting the second term on the LHS yields $\beta_j^* = \Delta_j$, showing that the correlations that cause these two metrics to differ are the pair processes p_{ij} that connect two detectors. Specifically, Δ_j attributes the logical error caused by pair mechanisms e_{ij} to *both* detectors v_i and v_j . The regression, on the other hand, splits the contribution of e_{ij} between v_i and v_j in proportion to all the other rates of mechanisms that have v_j in their support. Therefore, a ranking based on Δ_j quantifies the probability that detection events at v_j result in a logical error, whereas a ranking based on β_j quantifies the degree to which the logical error rate is uniquely attributed to the detector v_j . This distinction suggests that Δ_j is a more effective measure for identifying the noisiest shots, and thus Δ_j is used in ORP to produce the results in the main text. However, in practice, filtering based on β_j or Δ_j is found to give comparable results, see e.g. Fig. E.5.

E.6.5 The Lightcone of the Filtered Hierarchy

To connect the detector ranking chosen by Δ_j to the properties of a partially FT circuit U , we Pauli-propagate the observable $\bar{O}$ on which Δ_j is conditioned backward through U to the location of each detector v_j . The propagated observable can be represented as

$$\begin{aligned}\tilde{O} &= U\bar{O}U^\dagger \\ &= \sum_P c_P P ,\end{aligned}\tag{E.41}$$

where the sum is over all Pauli strings P . For a detector $v_{j=(g,b,r)}$, located at code block b in round r , its ranking, as determined by Pauli-propagation, is given by

$$\gamma_j = \frac{1}{\mathcal{N}} \sum_{P: \text{supp}(P) \in b} |c_P| \tag{E.42}$$

where the sum is over all Pauli strings that have support (i.e. non-identity components) contained in code block b , and $\mathcal{N} = \sum_P |c_P|$ is a normalization that ensures $0 \leq \gamma_j \leq 1$. The ranking based on Δ_j is set by the magnitude of Δ_j .

As an example, Table E.4 shows a comparison between the detector ranking determined by γ_j and the detector ranking determined by Δ_j using simulated results. The rankings are computed for the observable

Detector ranking	1	2	3	4	5	6	7	8	9	10
Δ_j	$v_{(0,1)}$	$v_{(0,2)}$	$v_{(0,4)}$	$v_{(0,3)}$	$v_{(0,5)}$	$v_{(1,1)}$	$v_{(1,2)}$	$v_{(3,1)}$	$v_{(3,3)}$	$v_{(3,4)}$
γ_j	$v_{(0,5)}$	$v_{(0,4)}$	$v_{(0,3)}$	$v_{(0,2)}$	$v_{(0,1)}$	$v_{(1,1)}$	$v_{(1,2)}$	$v_{(1,3)}$	$v_{(2,1)}$	$v_{(3,5)}$

Table E.4: *Comparison Between Detector Ranking Using Pauli-propagation and ORP.* The top row shows the position in the ranking, where 1 is the highest ranking and 10 is the lowest ranking among the detectors shown. The middle row shows the detectors $v_{(b,r)}$ ranked by their Δ_j value and the bottom row shows them ranked by their γ_j value. The rankings are computed for the observable $\bar{Z}_{(b=0,0)}$ using a four-block circuit consisting of $n_T = 4$ Trotter steps with step size $\delta t = 0.1$, and one syndrome extraction round per Trotter step, starting from an all $|00\rangle$ initial state. The rankings are computed for $g_x = g_z = 1$.

$\bar{Z}_{(b=0,0)}$ on four code blocks $b = 0, 1, 2, 3$ with $n_T = 4$ Trotter steps and a syndrome extraction rate of one




Figure E.6: *Depiction of Most Important Detectors* Calculated for the logical observable $\overline{O} = \overline{Z}_{(b=0,0)}$, as determined by the rankings calculated in Table E.4. Highlighted in red are the locations of the seven most important detectors, as determined by both ORP and Pauli-propagation. The grey region shows the observable's circuit backwards lightcone, containing all of the detectors that are connected to the observable through gates. The dark blue region captures the physical reverse-lightcone of $\overline{O}$, showing the detector locations where the back-propagated observable has the largest Pauli weights.

round per Trotter step, as illustrated in Fig. E.6. The Δ -rankings are computed using classical simulations subject to two-qubit depolarizing noise with error rate $p = 0.003$, which is comparable with error rates on current IBM hardware. These rankings are found to be consistent with those computed from hardware results. Although the two orderings do not match exactly, ORP is consistent with Pauli-propagation, since the set of the top- k detectors selected typically are the same, with the most important detectors residing in the code block containing the observable being measured. Figure E.6 shows the locations of the most important detectors chosen by the two methods and illustrates how the highest ranked detectors selected by ORP coincide with the local observable's backwards lightcone.

E.7 Fault Tolerant Benchmarking of Circuit Elements

E.7.1 Logical Error Rates of Circuit Gadgets

This appendix numerically studies the logical performance in the simulations by calculating the logical error rate (LER) as a function of the device error rate. The noise model assumes a two-qubit depolarizing channel, defined as

$$\mathcal{E}(\rho) = (1 - p)\rho + \frac{p}{15} \sum_{P \neq I} P\rho P \quad (\text{E.43})$$

where $p \in [0, 1]$ is the physical error rate and the sum is over all two-qubit Paulis P , excluding the identity. Additionally, single-qubit gates are followed by a single-qubit depolarizing channel with error rate $p/10$. In the numerical experiments, the ideal state is a logical $|\overline{00}\rangle$, which is a GHZ state on four qubits. A logical error occurs when there is a physical that keeps the state in the codespace $\mathcal{C}$ and goes undetected by the syndrome measurements. Removing shots where a nontrivial syndrome is detected is equivalent to projecting the noisy state back onto the codespace, defined by the projector

$$\Pi_{\mathcal{C}} = \frac{1}{4}(1 + S_Z)(1 + S_X), \quad (\text{E.44})$$

where undetectable errors take the state out of the simultaneous $+1$ eigenspace of S_Z and S_X . For a state ρ whose noisy evolution is given by $\tilde{\rho}$, the LER is defined by

$$p_L = 1 - F_L, \quad F_L = \text{Tr } \rho \tilde{\rho} \Pi_{\mathcal{C}}. \quad (\text{E.45})$$

The LER p_L can be expressed as a function of p . For p small, the leading order contribution is a good approximation for how p_L scales with p . For level-0 FT (no encoding) the leading order scaling is $\mathcal{O}(p)$. Gadgets that are 1-FT have leading order behavior $\mathcal{O}(p^2)$, and so mixing both types results in an LER of the form $p_L = Ap + Bp^2$, where the constants A, B encode the geometric features of a circuit. In particular, A represents nFT circuit features that can cause undetectable errors while B represents the extra circuit overhead required to detect errors. The relative magnitudes of A and B can have meaningful impacts on the error rates, and in some cases imply that a nFT circuit gadget can give superior performance to its encoded counterpart. An example of this is seen by comparing the nFT rotations used in this work to a FT approach like magic state injection.

The clearest examples of the improved performance of including nFT elements are the $|\overline{00}\rangle$ initialization circuits that prepare a four-qubit GHZ state. A simple nFT way of preparing this state requires a CNOT depth of four, whereas a 1-FT technique for preparing such a state on heavy-hex topology requires a CNOT depth of seven followed by two measurements. Both of these circuits are shown in Fig. E.7a), with their respective LERs determined from classical simulations shown in Fig. E.7b) over a range of p . When considering the LER of both of these circuit gadgets, the FT circuit has error suppression quadratic in p compared to the lack of suppression in the nFT case. For devices with long measurement times, idle errors can occur on the data qubits being measured, amplifying the FT overhead B by almost an order of magnitude and increasing the number of errors. Together with a two-qubit gate depth almost twice that of the nFT gadget, these extra




Figure E.7: *Logical Error Rate of Different Circuit Components* a) The different FT and nFT circuit elements whose logical error rates are considered in App. E.7.1. The green elements denote initialization gadgets, labeled FT (solid) and nFT (dashed). The bottom two circuits $\bar{R}_{FT}/\bar{R}_{nFT}$ are FT/nFT initialization gadgets followed by a nFT logical rotation $\bar{R}_{ZZ}$. b) The LER of the two different initialization gadgets. The grey region is where nFT crosses to a lower LER than the FT procedure. The grey and black dashed lines represent $\mathcal{O}(p)$ and $\mathcal{O}(p^2)$ scaling for the logical error rate p_L . c) The LER for the FT and nFT initialization followed by a logical rotation. d) The LER of the syndrome extraction circuits for multiple rounds. e) The LER of syndrome extraction as a function of number of rounds for different p .

circuit elements degrade the performance at large enough p . This is reflected in the LER curves, where the FT p_L passes the nFT value at a physical error rate of $p \sim 0.2$.

Considering both FT and nFT elements together in the same circuit qualitatively modifies this picture. In particular, examining the leading-order expansion for p_L reveals a crossover point at $p^* \approx A/B$ such that for $p > p^*$ there is a region of quadratic suppression of errors, which then degrades to a linear decay with no suppression for $p \lesssim p^*$. Examples of such behavior are shown in Fig. E.7c), which compares the effect of both FT and nFT state preparation, followed by a logical rotation (nFT) and syndrome extraction

(FT). The extra circuit depth pushes the LER high enough such that the nFT state prep is no longer able to outperform its FT counterpart. Further, even though the asymptotic error dependence is linear, the addition of FT components can push the LER below that of nFT rotations. The suppression at intermediate $p > p^*$ indicates that for physical error rates that are small but not yet below threshold, there are still advantages in error suppression gained by incorporating FT elements into circuit design. Finally, Fig. E.7d) analyzes the LER of the syndrome extraction circuits with FT initialization. Here, all of the circuit elements are 1-FT and so $p_L = \mathcal{O}(p^2)$. The overhead factor B increases with the number of syndrome measurement rounds, but then quickly plateaus to an $\mathcal{O}(1)$ constant as shown in Fig. E.7e).

E.7.2 Memory Experiments

To assess the effectiveness of the syndrome extraction circuits used in this work, memory experiments are performed on `ibm_boston` where both FT and nFT state prep are followed by many rounds of syndrome extraction. Assuming noiseless circuit operation, the system will remain in the logical $|\overline{00}\rangle = \frac{1}{\sqrt{2}}(|0000\rangle + |1111\rangle)$ state, and measurements in the Z -basis result in either “0000” or “1111” bitstrings. Noise populates other bitstrings with a non-zero probability, with some being detected by syndrome extraction and others causing undetectable errors. The undetected errors will appear as even parity bitstrings with Hamming-weight two in the final bitstring distribution. After postselecting all shots whose syndrome measurements flagged an error, the filtered probability distribution is compared to that of an ideal GHZ state using the Total Variation Distance (TVD) defined by

$$\text{TVD}(q, q') = \frac{1}{2} \sum_s |q(s) - q'(s)|, \quad (\text{E.46})$$

where the sum is over all bit strings, q is the measured distribution and q' is the GHZ distribution, with $q'(\text{“0000”}) = q'(\text{“1111”}) = 0.5$. Figure E.8a) shows the TVD of the postselected distribution as a function of the number of syndrome extraction rounds. The four curves compare (1) no DD and reset, (2) no DD and no reset, (3) DD and no reset and (4) DD and no reset with nFT state prep over the course of 60 rounds of syndrome extraction (all with a circuit depth of 761 CNOTs). Incrementally adding in DD and removing resets have dramatic impacts on the error, lowering it by over an order of magnitude. The nFT state preparation contains more errors for a few rounds of syndrome extraction, but remains comparable to the FT state over many rounds. Figure E.8b) shows the corresponding acceptance fraction for these cases, defined as the number of measurements that record no errors. This fraction decreases exponentially with the number of

rounds, with DD having the most profound effect on preserving shot lifetime.

The $[[4, 2, 2]]$ code is a distance-2 code, and so all events contributing to the LER at first order in the physical error rate p are flagged, and p_L will scale as $\mathcal{O}(p^2)$. Observing this scaling in quantum hardware runs requires tuning the physical error rate, which is not possible to do directly. We leverage the fact that syndrome extraction acts as the logical identity in the codespace, and so multiple rounds of extraction can serve to amplify the effective physical error rate. As such, the physical error rate can be increased by modifying which rounds of syndrome extraction are postselected against. Specifically, instead of considering the results of syndrome measurements on every round, we only consider postselecting on the final round of extraction. To account for the errors from initializing the code state $|\overline{00}\rangle$, a constant number of syndromes are additionally postselected on. As a proxy for the LER p_L , we count the fraction of even parity Hamming-weight 2 bitstrings that survive the postselection, which at leading order is proportional to the number of undetectable errors that occurred. Figure E.8c) shows the fraction of undetected errors as a function of the number of rounds with an increasing number of syndromes postselected at the beginning. This fraction grows quadratically with the number of rounds, and this number is interpreted as the scaling factor for the physical error rate. At low amplification, this fraction plateaus due to nFT state prep errors and finite measurement statistics, whereas at high amplification there is a plateau due to exceeding the pseudothreshold. Therefore, the scaling should be understood as appearing in an intermediate regime, in the plot approximately between rounds 8-30.

E.7.3 Rotations with Ancilla Qubits

A method to reduce the error spread by nFT rotations, at the cost of ancillas and measurements, was studied in Refs. [520, 596]. This method couples an ancilla qubit prepared in $|0\rangle$ to the rotation circuit in such a way that the bulk of single-qubit X errors are concentrated on the ancilla qubit. As an example, let $P = ZZ$ and $P' = ZZZ$ where the first two qubits are data and the last one is an ancilla. Then $e^{-i\theta P'} = e^{-i\theta P} |0\rangle\langle 0| + e^{i\theta P} |1\rangle\langle 1|$ and the correct rotation is recovered by postselection on ancilla state $|0\rangle$. Figure E.9a) shows such circuit pairs G_2/G'_2 and G_4/G'_4 , implementing ZZ -rotations within a block and between logical qubits of neighboring blocks. In these circuits, the only errors that spread to become high-weight errors are Z -errors indistinguishable from logical rotations. Ancilla measurements can be deferred to the end of the circuit, provided later errors do not affect them. Since X errors are accumulated




Figure E.8: *Memory Experiment Performance of Multiple Rounds on ibm_boston* a) The Total Variation Distance (TVD) from the logical $|00\rangle$ state over 60 rounds of syndrome extraction. b) The acceptance fraction over 60 rounds. c) Suppression of errors at leading order as a function of number of syndromes used for postselection.




Figure E.9: *Ancilla vs. No-Ancilla Rotations*. a) The top row shows ancilla-free circuits G_2 and G_4 implementing $e^{-i\pi/8\bar{Z}_{(b,0)}\bar{Z}_{(b,1)}}$ within a single code block (left) and $e^{-i\pi/8\bar{Z}_{(a,0)}\bar{Z}_{(b,0)}}$ between two code blocks a and b (right). The bottom row shows circuits G'_2 , G'_4 with ancilla qubits that implement the same rotations. b) The LER p_L as a function of the physical error rate p for the circuits in a). The logical error rate is computed by taking the average over Haar-random states under repeated rotations with depolarizing noise.

on the ancilla qubit, an undetectable (logical) error can occur when two physical errors cancel⁸. Note that applying this gadget with an ancilla qubit in state $|1\rangle$ implements a rotation with the opposite angle, which is a logical error that is detectable via postselection on ancilla measurement. The implications of this are

⁸This cancellation can be studied using an error model with X errors only on the ancilla qubit.

discussed in App. E.8.

While this method enables detecting more errors, its gate and measurement overhead introduces additional errors. To determine which approach is best-suited for currently available hardware, we compare implementations of $e^{-i\pi/8\bar{Z}_{(b,0)}\bar{Z}_{(b,1)}}$ within a single code block b and $e^{-i\pi/8\bar{Z}_{(a,0)}\bar{Z}_{(b,0)}}$ between two code blocks a and b . The data qubits are positioned within code blocks so that SWAP gates are not required to implement these rotations. The circuits in Fig. E.9a) are repeatedly applied such that a net $\theta = 2\pi$ rotation is implemented. The logical error rate p_L is determined from classical simulations with depolarizing noise by computing the fidelity of the initial state ρ with the noisy final state $\tilde{\rho}$, $p_L = 1 - \text{Tr}(\rho\tilde{\rho}\Pi_C)$. Since the effect of rotations depends on the initial state, we compute the Haar average of p_L over all initial states. A two-design is necessary to compute the Haar average of the fidelity [597]. The set of stabilizer states forms a three-design [598], so one- and two-qubit stabilizer states are used as initial states for the circuits in Fig. E.9a). The LER as a function of physical error rate p for eight repeated $\pi/4$ rotations of G_2 , G'_2 , G_4 , G'_4 is shown in Fig. E.9b). Both implementations exhibit $p_L = \mathcal{O}(p)$ scaling. The ancilla method G'_2 is found to have a two to three times smaller prefactor than G_2 , and G_4 is found to be roughly equivalent to G'_4 on average when accounting for errors detected by the $[[4, 2, 2]]$ code. This suggests that although the ancilla method can offer an advantage in specific cases, consistent with Refs. [520, 596], the two approaches perform similarly on average under depolarizing noise. Further, this examination assumes noiseless (and instantaneous) measurements. On realistic hardware, a measurement consumes significant coherence time and introduces measurement errors. For these reasons the simulations in this work are carried out with the standard NISQ logical rotation gadget.⁹

E.8 Additional Analysis of Encoding Improvement

E.8.1 Qubit Dependence

Figure E.10 shows additional metrics comparing encoded to unencoded runs presented in Sec. 7.3. As seen in Fig. 7.4, the encoding improvement in 1D is most pronounced at early and intermediate times, and decreases at late times. In 2D, the improvement grows with time. Figure E.10a) shows α per logical qubit across all t . For both 1D and 2D simulations, the encoded α values are on average higher than the unencoded, but are broadly distributed. Several outliers are present in both encoded and unencoded results; these can be matched to locations on the quantum processor with faulty gates or qubits with high measurement error rate

⁹Ancillas are used to implement rotations between blocks through bridge qubits where necessary, see App. E.3.




Figure E.10: *Postselection Performance and Signal Survival in Encoded Simulations.* a) The signal-survival factor α as a function of logical qubit over all t for $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). Encoded results with $R = 1$ are given in blue, and unencoded results are shown in grey. Results for the 1D chain are shown in the left column, and 2D grid results are shown in the right column. The shading represents one standard deviation of the α distribution over qubits. Error bars on individual points represent uncertainty determined through bootstrap resampling. b) Convergence of an example observable on qubit $(14, 1)$ at $t = 5$ as a function of the number of detectors kept for postselection k . The shading represents uncertainty from bootstrap resampling. Results from the unencoded runs and from MPS are given by grey and black dashed horizontal lines. The vertical black dotted line shows the k^* selected by the plateau finding algorithm. Acceptance rates for each k are shown by the dotted grey line.

or low T1/T2 times. Pauli twirling of the coherent and amplitude-damping/dephasing components of the noise could convert part of this structured error into a stochastic Pauli channel more amenable to detection and postselection. Omitting these blocks in the analysis improves results considerably, consistent with most regions of the quantum processor having an operating error rate just below the pseudothreshold. However, this exclusion is not done in the results reported in the main text. Further, the per-qubit encoded points carry visibly larger bootstrap uncertainties than the unencoded points, which is an expected feature of shot loss from postselection and increased circuit depth due to encoding. The results obtained from the simulations in 2D have a significantly larger spread than those from simulations in 1D, and are generally noisier (lower α). This again is a consequence of the deeper circuits required to implement a dense logical connectivity.

Figure E.10b) shows the effects of the plateau-finding algorithm explained in App. E.1 for qubit $(14, 1)$ at




Figure E.11: *Comparison of α for Corner and Bulk Qubits in 2D Simulations.* The signal-survival factor α is shown for encoded runs as a function of time t for R syndrome extraction rounds (colors) compared to unencoded runs (grey). The left panel shows α for corners of the grid (logical qubits with two neighbors), and the right panel shows α for bulk qubits (those with more than two neighbors). The insets show the error in the α fits, σ . Results from the localized regime ($g_z = 1.5$) are used.

$t = 5$, as an example. As detectors are added, the acceptance fraction (grey dotted line) falls steeply while the estimate settles onto a plateau in approaching the MPS value (black dashed line) and away from the biased unencoded result (grey dashed line). The plateau-finding algorithm selects k^* (blue dotted vertical line) at the onset of this stable region. If such a value is not found, the algorithm selects k^* to minimize the squared error, as explained in step five of the plateau-finding algorithm in App. E.1.2, and examples of this are shown in both 2D plots in Fig. E.10b) and in the $g_z = 1.5$ 1D plot. Across all times and qubits, most show similar convergence behavior with the exception of especially faulty device regions.

Figure E.11 shows the difference in α between corner and bulk qubits in 2D grid simulations with $g_z = 1.5$. Corners are defined as logical qubits with two connections, while bulk qubits have $\text{deg} > 2$. Since corners have fewer logical connections, fewer gates act on them and as a result they are less noisy. In addition, the backwards lightcone on corner qubits is smaller, reducing the amount of errors that can affect observables there. This is reflected in the larger α values in the left panel of Fig. E.11 compared with the right. Further, the σ values quantifying the spread in the α fit are smaller, indicating the corners more consistently reproduce the MPS expectations.




Figure E.12: *Group Velocity in the 1D and 2D Ising Model.* a) The group velocity $v(k)$ as a function of momentum k for $g_z = 0.5$ (blue) and $g_z = 1.5$ (red) calculated in a 20-qubit 1D chain with periodic boundary conditions. b) The group velocity $|v(k_x, k_y)|$ for a 5×5 2D lattice with $g_z = 0.5$ (left) and $g_z = 1.5$ (right). Dispersion relations are calculated using exact diagonalization, and a spectral derivative is used to compute $v(k_x, k_y)$.

E.8.2 Coupling Dependence

As shown in Fig. 7.4, the encoding improvement is roughly independent of g_z for 1D and 2D. However, α fits for both encoded and unencoded runs are consistently higher for $g_z = 1.5$ than 0.5, which is seen in Figs. 7.2e) and 7.3d). This difference originates in how far excitations (and errors) can propagate in each system. As shown in Fig. E.12a), the maximum group velocity $v(k) = dE(k)/dk$ is substantially larger at $g_z = 0.5$ than at $g_z = 1.5$ in 1D, where the stronger longitudinal field slows the propagation front.¹⁰ A similar, but less dramatic difference exists in 2D, shown in Fig. E.12b). The consequence of larger $v(k)$ for local observables is a larger backwards lightcone. This implies more errors have the ability to affect a given $\langle \bar{Z}_n \rangle$ in the melting regime at fixed t . At a fixed number of syndrome extraction rounds R , errors can spread more under faster dynamics before they are caught by error detection. This is the reason that α is consistently smaller in the melting regime (seen in Figs. 7.2e) and 7.3d), and degrades faster. Further, the group velocity in 2D is significantly smaller than in 1D. The melting regime in 1D has by far the largest $v(k)$, ~ 2.5 times larger than $g_z = 1.5$ in 1D and ~ 5 times larger than 2D $v(k_x, k_y)$. The fluctuations in α in Fig. 7.2e) are

¹⁰These calculations are done using exact diagonalization with periodic boundary conditions. The limits of this method prevent determining the group velocity for our lattice geometry with missing links, so a smaller uniform lattice is used instead.

attributed to this difference. Since the difference in $v(k_x, k_y)$ between the regimes is only 0.04 in 2D, errors propagate at more similar speeds, which explains why the α plots in Fig. 7.3d) are more similar than in 1D. This dependence on the backwards lightcone suggests that more frequent syndrome extraction is necessary for simulations of faster dynamics. Unfortunately, the coherence time overhead is found to suppress most of the advantage of more frequent syndrome measurement in our simulations.

E.8.3 Reset Dependence

As a result of omitting resets on ancilla qubits after syndrome measurements, an ancilla that fires remains in the state $|1\rangle$ until a subsequent error measurement flips it back to $|0\rangle$.¹¹ Any inter-block $\bar{R}_{ZZ}(\theta)$ rotation that uses this ancilla in its implementation therefore applies the opposite rotation $\bar{R}_{ZZ}(-\theta)$. Interspersing these opposite-angle inter-block rotations among the correct ones partially cancels the intended coupling J , producing an effective “slowed-down” interaction $J_{\text{eff}} = J(1 - 2p_{\text{flip}})$, where p_{flip} is the probability that the mediating ancilla is in $|1\rangle$ at the time of the gate. The effect is amplified when two ancillas are used to implement the inter-block $\bar{R}_{ZZ}(\theta)$ rotation. Figure E.13 shows the average J_{eff} over all inter-qubit couplings as a function of time for various amounts of syndrome extraction rounds R . In 1D, the reduction in J is smaller in the localized regime, where the stronger longitudinal field g_z suppresses the spread of correlations along the chain. This results in a lower average p_{flip} for $g_z = 1.5$. The differences in J_{eff} between regimes and dimensions is again attributed to the group velocity which governs the propagation speed of errors. Since errors propagate slower at $g_z = 1.5$, the effects of incorrect rotations take longer to spread. This slowdown could be further studied by running MPS simulations with the couplings J_{eff} , but this is not done in this work.

The effective reduction in J could be removed entirely with the introduction of ancilla resets after syndrome measurements, and is the primary reason that high- R simulations are found to perform worse than low- R simulations. Increasing syndrome measurement frequency adds more opportunities for stale ancillas to artificially inject errors into the simulation, and if the corresponding detection events are not selected for postselection (e.g., by ORP), they contaminate the results. With ancilla resets, this problem is absent. However, ancilla resets consume a significant amount of coherence time, and as a result errors are nevertheless introduced. We find that the implementation of resets does not significantly change the R -dependence of

¹¹The absence of resets also means that leakage on an ancilla persists across all subsequent rounds, since there is no re-initialization to return a leaked ancilla to the computational subspace.




Figure E.13: *Effective Inter-Block Coupling Under Residual Ancilla Sign Flips on ibm.boston*. a) The effective inter-block coupling $\langle J_{\text{eff}}/J \rangle$ as a function of time t averaged over all inter-block connections in the 1D simulations shown in Fig. 7.2 for $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). Different R values indicate amount of syndrome measurements. b) The same as a) but for 2D results presented in Fig. 7.3.

the results, with the exception of the highest values of R . Even for $R = 16$, where the improvement from resets is largest, the α values do not exceed those of low R , indicating that the added coherence time injects more errors than the syndrome measurement rounds are able to detect. The implementation of faster resets on the hardware will likely shift this balance. This is another example of the tradeoff between the NISQ considerations of gate depth and coherence time and FT requirements.

E.8.4 Geometric Dependence in 1D Simulations

The encoded results presented in Fig. 7.3 show a large improvement over unencoded runs, increasing with time to a greater than 200% improvement at late times. This is partially due to the geometric gate depth overhead of embedding a square lattice onto heavy-hex connectivity. Unencoded circuits are 1.5 times deeper than encoded circuits at the latest time considered (see Table E.3). To isolate the encoding improvement from the geometric overhead, we compare the encoded results without any postselection to unencoded results in Fig. E.14. The CDF plots show that encoding without postselection has a small improvement at $g_z = 0.5$ and a negligible improvement at $g_z = 1.5$. The decay of the signal-survival factor α shows the performance of




Figure E.14: *Encoding Improvement Without Postselection in 2D Simulations.* a) Cumulative distribution functions (CDFs) of the per-qubit, per-time absolute error $|\langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}}|$ for both values of g_z . The median error and CDFs corresponding to $R = 1$ syndrome measurements are shown compared to unencoded results (solid and dashed lines respectively). b) The signal-survival factor α as a function of t for various R (colored lines) and each value of g_z , compared to unencoded results (dashed grey line). The insets show the residual error σ in the fit.

the encoded circuits over time without postselection. At early times, encoded circuits without postselection perform worse than unencoded as a result of added circuit depth to perform syndrome measurements. At intermediate and late times, low- R encoded runs outperform unencoded due to the geometric overhead. Comparing to Fig. 7.3d), the maximum decrease in α from omitting postselection is ~ 0.4 . This indicates that much of the improvement seen in Fig. 7.3 is due to encoding as opposed to the geometric overhead in unencoded circuits. Thus, removing errors that can potentially spread to $\mathcal{O}(t^2)$ qubits plays a large role in the improvement observed in 2D simulations.

E.9 Classical Simulations of False-Vacuum Decay Dynamics

Phase transitions play a key role in early-universe dynamics as well as in many condensed matter systems. For high energy physics, a particularly important question is whether the Higgs field is in a true vacuum or in a false metastable vacuum. The phase transition from a false vacuum to a true vacuum is known as false-

vacuum decay and underlies a range of early-universe phenomena, from cosmological phase transitions and bubble nucleation to the possible fate of the electroweak vacuum [582, 583]. This section describes a simple 1D model for false-vacuum decay in spin systems as intuition for the oscillation and localization behavior observed in the simulations in the main text. Classical MPS simulations carried out in the same regimes considered in Sec. 7.3 are shown to support this intuition where analytical calculations are unavailable.

In false-vacuum decay phenomenology, the metastable “false” vacuum decays to the stable and energetically favored “true” vacuum by quantum tunneling [582, 583]. In the semiclassical regime, thermal or quantum fluctuations typically nucleate localized bubbles of true vacuum. Bubbles exceeding a critical size then lower the energy through expansion, driving conversion of the system to the true vacuum. The method in this work provides a lattice realization of false-vacuum decay through quench dynamics of a spin system [569, 599]. Rather than directly studying bubble nucleation, we prepare a fixed, finite system in a ferromagnetic initial state with a true-vacuum bubble in the center and follow its unitary time evolution. The presence of “decay” in this setup is indicated by the melting of the true-vacuum bubble, and decay is absent when the bubble remains localized.

Section 7.3 observes a difference in encoded quantum simulation performance between the $g_z = 0.5$ (melting regime) and $g_z = 1.5$ (localized regime). The melting or localization of true-vacuum bubbles in our simulations is due to domain-wall oscillations known as Bloch oscillations [585]. The following analysis considers two domain walls in a 1D chain. The Hamiltonian (7.6) assigns a fixed energy per flipped spin, so the bubble’s potential energy is linear in its size r

$$V(r) = -\chi r, \quad \chi = 2m|g_z|, \quad m = \left(1 - \left(\frac{g_x}{|J|}\right)^2\right)^{1/8}, \quad (\text{E.47})$$

where m is the spontaneous magnetization at $g_z = 0$ [600]. Unitary time evolution under $\bar{H}$ conserves total energy, but converts the potential energy into the kinetic energy of expanding bubble walls. As $V(r)$ changes from bubble expansion, the relative momentum of the bubble walls k changes, as given by Hamilton’s equations,

$$\frac{dk}{dt} = -\frac{dV}{dr} = \chi. \quad (\text{E.48})$$

Neglecting corrections due to nonzero g_z , the dispersion relation for a single domain wall is

$$E(k) = 2|J|\sqrt{1 + g_x^2 + 2g_x \cos k}. \quad (\text{E.49})$$

Since $E(k)$ is periodic, the group velocity $v(k) = dE(k)/dk$ changes sign as the momentum is increased in the Brillouin zone, resulting in oscillations of the bubble width

$$r(t) = \frac{4}{\chi} \left(E \left(k_0 + \frac{\chi}{2} t \right) - E(k_0) \right). \quad (\text{E.50})$$

The maximum amplitude of these oscillations is

$$r_{\max} = \frac{4}{\chi} (E(\pi) - E(0)) = \frac{8|J||g_x|}{m|g_z|}, \quad (\text{E.51})$$

and the period is

$$T_B = \frac{4\pi}{\chi} = \frac{2\pi}{m|g_z|}. \quad (\text{E.52})$$

Although strictly valid only in the analytically tractable $g_z \rightarrow 0$ limit, and not at the couplings considered in our simulations, this analysis provides intuition for how oscillations arise in the domain-wall initial state. An equivalent treatment in the large- J limit maps the system to free fermions in a linear potential [574, 584], where the extent of oscillations is also found to be $r_{\max} \sim |g_x|/|g_z|$. Reference [584] finds a transition between the bubble-melting and localized regimes at $|g_z|/|g_x| \sim 1$ for large J . Systems with $|g_z|/|g_x| < 1$ are found to thermalize, while $|g_z|/|g_x| > 1$ shows evidence of slower decay. Results from our simulations at $J = 1$ are consistent with these findings. Figure E.15 shows the energy density determined from MPS simulations through the relaxation dynamics following a quench under the Hamiltonian (7.6) in 1D. Initial bubble states $|\psi(t=0)\rangle = \prod_{i \in B} \bar{X}_i |\bar{1}\rangle^{\otimes N}$ with bubble sizes $|B| = 1, 2, 3$ are shown in the columns. The top row shows the bubble-melting regime ($g_z = 0.5$). Faint oscillations about the center of the bubble are seen for sizes two and three, consistent with Bloch oscillations with a period $T_B \sim 12$ and $r_{\max} \sim 10$. Since our simulations are far outside the small- g_z or large- J regime, we cannot compare these to predictions of T_B or $r_{\max}$ from the analysis above. The bottom row ($g_z = 1.5$) shows the bubble localized to all accessible times for sizes two and three. This is consistent with Eq. (E.51), where larger g_z makes $r_{\max}$ comparable to the bubble size and prevents oscillations. A size one bubble is seen to decay for both values of the coupling because the Bloch mechanism requires separation between the two bubble walls for the linear potential to act on, which does not exist for a single flipped spin. This confirms the localization is not due to Z -field pinning, in which a strong longitudinal field suppresses interface-moving spin flips when $|g_z| \gg |g_x|$.

We find similar behavior in 2D simulations. Figure E.16 shows $\langle \bar{Z}_n \rangle$ determined from MPS simulations throughout the lattice for a selection of times. The dynamics in 2D is found to be slower than in 1D, which is




Figure E.15: *Classical Simulations of False-Vacuum Decay Dynamics in 1D.* An initial true-vacuum bubble of size one, two, three (columns) at the center of the 42-qubit chain is evolved with $g_z = 0.5$ (top row) and $g_z = 1.5$ (bottom row). The color in the heatmap shows the energy density $\langle \overline{H}_n \rangle$ as a function of logical qubit position $n = (b, l_b)$ and time t . The time evolution uses $n_T = 500$ Trotter steps of size $\delta t = 0.1$ and is computed with MPS.

also evidenced by the group velocities shown in Fig. E.12. Further, the melting and localization dynamics is anisotropic due to the irregular geometry realized by the logical qubit connectivity. For $g_z = 0.5$ sizes 2×2 and 3×3 , the initial bubble is seen to spread more towards the right and bottom as a result of missing edges on the left and top boundaries. These missing edges also artificially localize the bubble for $g_z = 1.5$. However, physics-driven localization (as opposed to localization driven by missing links) is still seen on the right and bottom, consistent with the dynamics seen in 1D simulations. Additionally, the bottom right qubit of the 3×3 bubble for $g_z = 1.5$ melts off as a result of a missing vertical edge that would connect it to the rest of the bubble.¹²

To remove the contributions to the melting and localization dynamics from missing links, we compare

¹²The global quench of the background contributes a significant amount to the observed dynamics, which is seen in the relaxation dynamics shown in Figs. E.15 and E.16.




Figure E.16: *Classical Simulations of False-Vacuum Decay Dynamics on a 2D Lattice with Missing Edges.*

An initial true-vacuum bubble of size 1×1 , 2×2 , 3×3 (rows) at the center of the 42-qubit lattice is evolved with $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). The qubit colors show the magnetization $\langle \bar{Z}_n \rangle$ as a function of logical qubit position $n = (b, l_b)$ and time t (columns). The time evolution uses $n_T = 100$ Trotter steps of size $\delta t = 0.1$ calculated with MPS. The logical connectivity grid corresponds to Fig. 7.2a) with all available edges used.

the results of our realized grid to simulations on a uniform 5×5 lattice. The comparison of $\langle \bar{Z}_n \rangle$ is shown for several selected times in Fig. E.17. This figure confirms that melting and localization consistent with Bloch oscillations are present in the absence of grid irregularities. The top row of the $g_z = 0.5$ panel of Fig. E.17 shows a recurrence for the 1×1 bubble. This happens due to boundary effects and is absent in 1D simulations in Fig. E.15. This is additionally verified by MPS simulations of a 7×7 lattice. The 2×2 bubble




Figure E.17: *Classical Simulations of False-Vacuum Decay Dynamics on a Uniform 2D Lattice.* An initial true-vacuum bubble of size 1×1 , 2×2 , 3×3 (rows) at the center of the 5×5 lattice is evolved with $g_z = 0.5$ (top) and $g_z = 1.5$ (bottom). The qubit colors show the magnetization $\langle \bar{Z}_n \rangle$ as a function of logical qubit position $n = (b, l_b)$ and time t (columns). The time evolution uses $n_T = 500$ Trotter steps of size $\delta t = 0.1$ and is calculated with statevector simulations.

is seen to melt for $g_z = 1.5$. This melting is slower than for $g_z = 0.5$ and is also impacted anisotropically by the boundary. This is consistent with the findings of Ref. [574] where $|g_z|/|g_x| \gg 1$ significantly slows down but does not fully prevent melting.

Together, these classical simulations show that the size-three bubbles (1D) and 3×3 bubbles (2D) studied in Sec. 7.3 melt for $g_z = 0.5$ and remain localized for $g_z = 1.5$, consistent with the Bloch oscillation picture. While the dynamics in 2D is significantly slower and the contrast between the melting and localized regimes

is less pronounced, the same qualitative behavior holds, supporting the interpretation of the device results in the main text. Although our couplings lie outside the small- g_z and large- J limits where the oscillation amplitude and period can be predicted analytically, the melting and localization observed here match the Bloch oscillation mechanism qualitatively.

E.10 Other Layouts

As discussed in Sec. 7.2.1, there is significant freedom in $[[4, 2, 2]]$ code block placement onto the heavy-hex connectivity. In this work, we select a placement that avoids particularly faulty qubits and gates. The block placement options are further expanded if this constraint is lifted. While the maximum number of blocks is still 21 (corresponding to 42 logical qubits), the properties of the resulting placement may be more favorable. In particular, the resulting logical connectivity graph may be more similar to a regular square lattice. Figure E.18a) shows an example of one such placement. This placement uses 144 active physical qubits, compared to 136 in Fig. 7.1b), a result of more bridge qubits present. An example of the resulting logical connectivity graph possible with this placement is shown in Fig. E.18b). This grid has 51 inter-block edges (72 edges total), whereas the grid used in the main text (Fig. 7.2a)) has 45 inter-block edges (66 edges total). Importantly, this grid has fewer missing interior edges, which significantly contribute to the dynamics in our simulations. This suggests that more regular square lattice simulations are possible using the current generation of quantum computers with more uniform characteristics.

It is not possible to regularly tile a heavy-hex connectivity device with blocks that have the topology required for our syndrome extraction circuits (center of Fig. 7.1c)). Therefore bridge qubits are necessary, which contribute significantly to the depth of inter-block gates. A regular tiling with our block topology is possible if blocks are allowed to share ancillas. An example of such a placement is shown in Fig. E.19. This placement supports a higher encoding rate, using 126 physical qubits to encode 48 logical qubits with 24 overlapping blocks. In addition, the logical connectivity graph possible with this placement is more dense, allowing for three-body gates between triples of neighboring blocks (with ancilla shuttling). In an overlapping block placement, code blocks share ancilla qubits (shown in green) that are used only for syndrome extraction, and data qubits (black) are permanently tied to a given block. This removes the overhead of bridge qubits. Since data qubits of neighboring blocks are located physically closer, inter-block gates are considerably shallower, which is favorable for the total depth budget. The tradeoff associated with overlapping layouts




Figure E.18: *Block Placement and Logical Connectivity Graph in the Absence of Faulty Qubits and Gates.*

a) An example of a block placement (grey) of 21 blocks assuming no faulty qubits and gates on a heavy-hex quantum computer. Within each block, the data qubits are shown in black, Z -ancillas are shown in red, X -ancillas are shown in blue, and bridge qubits are shown in green. b) The resulting connectivity between logical qubits (grey) optimized for similarity to a square lattice. Thick orange hatched edges represent intra-block connections and thin orange edges show inter-block connections.

is that syndrome extraction must now be serialized as a result of ancilla sharing. Figure E.19 splits the blocks into two groups (shown in red and blue), corresponding to the serial order in which syndromes must be measured. Ancilla shuttling is also required in this scenario, which introduces an overhead larger than that of switching between the compute and syndrome layouts in Fig. 7.1c). Further, the conversion to detectors discussed in Sec. 7.2.2 and E.1 is modified due to ancilla sharing. Syndrome measurement is by far the highest overhead in coherence time, which is doubled by serialization. This work finds that the quality of local observables measured degrades significantly faster with overlapping code blocks compared to non-overlapping placements, and for this reason the non-overlapping placement is used.




Figure E.19: *Overlapping Block Placement*. Code blocks are split into two groups (red and blue) for serialized syndrome extraction. Data qubits within each block are shown in black and shared ancilla qubits are shown in green.

E.11 Tables of Results

Tables E.5, E.6 and E.7 show the data from `ibm_boston` corresponding to Fig. 7.2. A 1D 42-logical qubit chain is run with $\delta t = 0.5$ for $g_z = 0.5$ and $g_z = 1.5$. The initial state in each of these simulations is a size-three true-vacuum bubble product state. Tables E.8, E.9, E.10 and E.11 show the corresponding data for the 2D lattice simulations displayed in Fig. 7.3. The initial state on the 42-logical qubit lattice is a 3×3 true-vacuum bubble product state. Table E.12 reports the acceptance rate of $R = 1$ runs for 1D and 2D simulations for both values of g_z .

Table E.5: *Magnetization in 1D Simulations at $t = 8$* . For each logical qubit (b, l_b) (first column), $\langle \bar{Z}_n \rangle$ from MPS results (second and fifth columns) is compared to unencoded results (third and sixth columns) and encoded results with $R = 1$ syndrome extraction rounds (fourth and seventh columns). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.

	$\langle \bar{Z}_n \rangle$					
Logical qubit (b, l_b)	$g_z = 0.5$			$g_z = 1.5$		
	MPS	Unencoded	$R = 1$	MPS	Unencoded	$R = 1$
(13, 1)	0.217	0.100(6)	0.098(19)	-0.568	-0.486(5)	-0.329(27)
(13, 0)	0.142	0.131(5)	0.113(16)	-0.372	-0.250(5)	-0.385(26)
(10, 1)	0.084	0.090(5)	0.081(17)	-0.138	-0.066(5)	-0.244(36)
(10, 0)	-0.079	-0.006(6)	0.000(21)	-0.445	-0.213(5)	-0.256(11)
(9, 0)	0.003	-0.006(6)	-0.019(21)	-0.381	-0.263(5)	-0.249(11)
(9, 1)	0.082	0.009(6)	0.019(34)	-0.330	-0.260(5)	-0.299(18)
(4, 0)	0.080	-0.043(6)	0.024(32)	-0.289	-0.263(6)	-0.248(22)
(4, 1)	0.099	0.022(6)	0.071(32)	-0.381	-0.248(6)	-0.239(33)
(2, 0)	0.115	0.055(6)	0.065(18)	-0.407	-0.283(5)	-0.269(20)
(2, 1)	0.130	0.109(5)	0.057(13)	-0.342	-0.229(5)	-0.229(46)
(0, 0)	0.175	0.121(6)	0.067(46)	-0.379	-0.297(5)	-0.202(31)
(0, 1)	0.124	0.076(6)	0.056(15)	-0.351	-0.286(5)	-0.136(18)
(1, 0)	0.071	0.044(6)	0.009(32)	-0.317	-0.223(5)	-0.132(31)
(1, 1)	0.029	0.043(6)	0.030(18)	-0.397	-0.286(6)	-0.106(42)
(3, 1)	0.068	0.115(6)	-0.020(17)	-0.347	-0.166(6)	-0.152(35)
(3, 0)	0.071	0.109(6)	0.015(43)	-0.374	-0.194(5)	-0.268(18)
(7, 1)	0.029	0.071(5)	-0.010(29)	-0.228	-0.151(6)	-0.223(11)
(7, 0)	-0.090	-0.026(5)	-0.090(50)	-0.371	-0.271(5)	-0.269(34)
(5, 0)	-0.011	-0.028(6)	0.029(17)	-0.348	-0.245(5)	-0.237(40)
(5, 1)	0.145	0.096(6)	0.082(17)	-0.549	-0.462(5)	-0.450(26)
(6, 1)	0.206	0.217(6)	0.163(41)	0.623	0.549(5)	0.555(21)
(6, 0)	0.170	0.219(5)	0.149(47)	0.928	0.806(3)	0.813(10)
(11, 0)	0.245	0.223(5)	0.211(30)	0.809	0.677(4)	0.595(12)
(11, 1)	0.142	0.108(6)	0.087(17)	-0.705	-0.499(5)	-0.489(13)
(8, 0)	-0.071	-0.058(6)	-0.108(30)	-0.304	-0.231(6)	-0.179(10)
(8, 1)	-0.108	-0.070(6)	-0.089(14)	-0.329	-0.291(5)	-0.345(32)
(12, 1)	-0.001	-0.005(6)	-0.010(28)	-0.252	-0.206(6)	-0.137(32)
(12, 0)	0.084	0.038(6)	0.012(17)	-0.373	-0.249(5)	-0.309(39)
(15, 1)	0.095	0.024(6)	0.083(20)	-0.275	-0.152(5)	-0.191(13)
(15, 0)	0.100	0.052(6)	0.063(28)	-0.403	-0.318(5)	-0.243(20)

Logical qubit (b, l_b)	$g_z = 0.5$			$g_z = 1.5$		
	MPS	Unencoded	$R = 1$	MPS	Unencoded	$R = 1$
(14, 1)	0.101	0.096(5)	-0.002(11)	-0.403	-0.242(6)	-0.237(30)
(14, 0)	0.088	0.171(5)	-0.010(12)	-0.280	-0.053(6)	-0.163(39)
(17, 0)	0.080	0.104(5)	-0.012(11)	-0.333	-0.199(5)	-0.252(19)
(17, 1)	0.066	0.076(6)	0.040(18)	-0.357	-0.293(5)	-0.246(13)
(20, 0)	0.054	0.059(5)	0.032(10)	-0.425	-0.335(5)	-0.258(26)
(20, 1)	0.067	0.076(6)	0.007(15)	-0.274	-0.200(5)	-0.135(9)
(19, 0)	0.088	0.067(6)	-0.016(40)	-0.323	-0.229(5)	-0.288(22)
(19, 1)	0.013	0.012(5)	-0.048(50)	-0.383	-0.290(5)	-0.233(9)
(16, 1)	-0.077	-0.074(6)	-0.024(20)	-0.443	-0.278(5)	-0.245(26)
(16, 0)	0.083	0.012(6)	0.089(11)	-0.139	-0.125(6)	-0.247(14)
(18, 0)	0.142	0.120(6)	0.099(11)	-0.372	-0.203(5)	-0.489(35)
(18, 1)	0.217	0.166(6)	0.054(23)	-0.568	-0.301(6)	-0.468(17)

$ \langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}} $												
%	$g_z = 0.5$						$g_z = 1.5$					
	Unenc.	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$	Unenc.	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$
0%	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.001(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)
10%	0.009(0)	0.005(1)	0.007(1)	0.007(1)	0.008(1)	0.013(1)	0.014(1)	0.009(1)	0.012(1)	0.010(1)	0.011(1)	0.016(1)
20%	0.016(1)	0.014(1)	0.016(1)	0.015(1)	0.017(1)	0.023(1)	0.031(1)	0.018(1)	0.020(1)	0.021(1)	0.023(1)	0.028(1)
30%	0.022(1)	0.021(1)	0.025(1)	0.024(1)	0.027(1)	0.037(2)	0.048(1)	0.033(1)	0.033(1)	0.039(2)	0.041(2)	0.051(2)
40%	0.031(1)	0.031(1)	0.034(1)	0.035(1)	0.040(1)	0.051(1)	0.062(1)	0.047(1)	0.049(2)	0.056(2)	0.062(2)	0.083(3)
50%	0.040(1)	0.039(1)	0.043(1)	0.046(1)	0.055(1)	0.065(1)	0.074(1)	0.067(2)	0.066(2)	0.074(2)	0.084(2)	0.137(3)
60%	0.050(1)	0.049(1)	0.053(1)	0.057(2)	0.068(2)	0.081(2)	0.091(1)	0.084(2)	0.086(2)	0.095(2)	0.124(3)	0.213(4)
70%	0.062(1)	0.061(2)	0.066(2)	0.075(2)	0.082(2)	0.098(2)	0.109(1)	0.108(2)	0.111(2)	0.127(3)	0.167(3)	0.282(4)
80%	0.078(1)	0.076(2)	0.083(2)	0.091(2)	0.105(2)	0.128(2)	0.129(1)	0.137(2)	0.143(2)	0.161(3)	0.215(3)	0.349(4)
90%	0.109(1)	0.104(3)	0.115(3)	0.123(3)	0.136(3)	0.173(4)	0.165(1)	0.170(3)	0.197(3)	0.217(4)	0.283(4)	0.421(4)
100%	0.306(5)	0.398(28)	0.370(15)	0.386(22)	0.429(14)	0.582(19)	0.291(5)	0.330(17)	0.455(22)	0.478(19)	0.548(21)	0.750(23)

Table E.6: *Absolute Error in 1D Simulations.* The absolute error $|\langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}}|$ by decile (first column) is compared between unencoded runs (second and seventh columns) and encoded runs for various numbers of syndrome extraction rounds R (columns 3-6 and 8-11). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.

Signal-survival factor $\alpha(\sigma)$												
t	$g_z = 0.5$						$g_z = 1.5$					
	Unenc.	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$	Unenc.	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$
0.0	0.983(19)	0.989(12)	0.987(9)	0.987(12)	0.988(9)	0.986(11)	0.983(19)	0.988(10)	0.987(11)	0.987(10)	0.989(10)	0.984(13)
0.5	0.931(34)	0.955(44)	0.959(44)	0.947(46)	0.953(42)	0.954(46)	0.929(36)	0.954(39)	0.955(44)	0.955(44)	0.956(40)	0.959(35)
1.0	0.905(45)	0.948(44)	0.938(43)	0.941(44)	0.937(45)	0.916(89)	0.898(47)	0.944(38)	0.942(42)	0.949(43)	0.943(42)	0.924(66)
1.5	0.878(65)	0.918(79)	0.893(85)	0.897(76)	0.897(73)	0.870(110)	0.877(61)	0.913(71)	0.908(74)	0.911(74)	0.927(62)	0.922(89)
2.0	0.870(67)	0.918(66)	0.887(69)	0.883(65)	0.851(83)	0.828(136)	0.888(58)	0.927(68)	0.930(66)	0.915(67)	0.939(70)	0.912(106)
2.5	0.841(78)	0.875(80)	0.841(80)	0.871(80)	0.823(98)	0.771(129)	0.913(55)	0.960(58)	0.966(60)	0.951(71)	0.983(67)	0.953(84)
3.0	0.820(67)	0.829(90)	0.802(80)	0.768(96)	0.724(99)	0.694(116)	0.908(45)	0.922(76)	0.945(85)	0.904(69)	0.927(72)	0.839(91)
3.5	0.777(74)	0.834(53)	0.822(62)	0.802(70)	0.757(83)	0.692(95)	0.893(40)	0.918(55)	0.934(55)	0.914(65)	0.864(72)	0.713(77)
4.0	0.777(62)	0.923(42)	0.805(44)	0.767(44)	0.721(60)	0.642(74)	0.864(37)	0.886(79)	0.874(71)	0.857(74)	0.798(81)	0.636(97)
4.5	0.755(57)	0.766(47)	0.876(56)	0.723(46)	0.739(55)	0.639(83)	0.824(44)	0.859(72)	0.858(67)	0.792(82)	0.731(116)	0.539(120)
5.0	0.842(45)	0.725(46)	0.699(56)	0.682(48)	0.569(53)	0.346(58)	0.801(50)	0.800(71)	0.816(88)	0.784(98)	0.703(127)	0.480(140)
5.5	0.898(51)	0.592(49)	0.636(58)	0.568(49)	0.448(48)	0.265(49)	0.776(61)	0.810(78)	0.766(94)	0.723(107)	0.645(145)	0.430(146)
6.0	0.845(43)	0.593(52)	0.591(58)	0.449(42)	0.378(46)	0.127(42)	0.757(59)	0.771(77)	0.729(94)	0.696(102)	0.598(129)	0.371(152)
6.5	0.822(47)	0.572(53)	0.529(51)	0.327(49)	0.291(41)	0.157(38)	0.750(64)	0.753(87)	0.731(86)	0.670(113)	0.603(141)	0.345(146)
7.0	0.760(44)	0.542(63)	0.426(50)	0.223(56)	0.152(40)	0.039(33)	0.741(60)	0.757(74)	0.715(84)	0.686(100)	0.583(123)	0.285(135)
7.5	0.802(46)	0.481(45)	0.465(48)	0.259(41)	0.179(34)	-0.019(22)	0.737(58)	0.750(74)	0.701(82)	0.687(90)	0.568(126)	0.255(119)
8.0	0.779(42)	0.552(38)	0.420(39)	0.253(33)	0.110(40)	0.123(30)	0.735(56)	0.728(81)	0.679(79)	0.667(90)	0.534(117)	0.246(115)

Table E.7: *Signal-Survival Factor for 1D Simulations.* The signal-survival factor α calculated as a function of time t (first column) in unencoded results (second and seventh columns) is compared to encoded results for various numbers of syndrome extraction rounds R (columns 3-6 and 8-11). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is the residual error σ in the α fits.

	$\langle \bar{Z}_n \rangle$					
Logical qubit (b, l_b)	$g_z = 0.5$			$g_z = 1.5$		
	MPS	Unencoded	$R = 1$	MPS	Unencoded	$R = 1$
(13, 1)	-0.558	-0.113(6)	-0.257(16)	-0.093	0.094(6)	-0.147(34)
(13, 0)	-0.725	-0.159(6)	-0.362(25)	-0.116	0.047(6)	-0.094(13)
(10, 1)	-0.694	-0.134(6)	-0.357(9)	-0.111	-0.003(5)	-0.123(13)
(10, 0)	-0.538	-0.095(5)	-0.283(14)	-0.138	-0.017(6)	-0.191(10)
(9, 0)	-0.524	-0.139(5)	-0.069(24)	-0.375	-0.142(5)	-0.127(17)
(9, 1)	-0.640	-0.206(6)	-0.284(26)	-0.145	-0.092(6)	0.004(18)
(4, 0)	-0.603	-0.253(5)	-0.388(17)	0.005	-0.078(5)	-0.102(13)
(4, 1)	-0.617	-0.262(5)	-0.352(12)	-0.036	-0.058(5)	-0.115(10)
(2, 0)	-0.706	-0.210(5)	-0.418(17)	-0.134	0.017(6)	-0.223(28)
(2, 1)	-0.686	-0.216(5)	-0.408(17)	-0.279	-0.102(5)	-0.230(24)
(0, 0)	-0.627	-0.271(6)	-0.180(44)	-0.087	-0.110(6)	-0.119(14)
(0, 1)	-0.770	-0.235(5)	-0.278(27)	-0.139	-0.080(6)	-0.067(26)
(1, 0)	-0.732	-0.205(5)	-0.103(24)	-0.137	-0.025(5)	-0.085(30)
(1, 1)	-0.554	-0.166(6)	-0.098(48)	-0.087	-0.025(5)	-0.148(15)
(3, 1)	-0.432	-0.089(6)	-0.228(29)	-0.165	-0.075(5)	-0.153(32)
(3, 0)	-0.620	-0.147(5)	-0.289(23)	-0.186	-0.056(6)	-0.185(31)
(7, 1)	-0.457	-0.037(6)	-0.298(32)	-0.043	-0.081(6)	-0.115(30)
(7, 0)	-0.403	-0.071(6)	-0.296(18)	-0.208	-0.102(6)	-0.179(29)
(5, 0)	-0.249	-0.075(6)	-0.199(19)	-0.283	-0.097(5)	-0.288(10)
(5, 1)	0.182	0.166(5)	0.203(33)	0.658	0.333(5)	0.420(14)
(6, 1)	0.597	0.201(5)	0.579(22)	0.695	0.242(5)	0.652(29)

Table E.8: *Magnetization in 2D Simulations at $t = 4$, first half.* For each logical qubit (b, l_b) (first column), $\langle \bar{Z}_n \rangle$ from MPS results (second and fifth columns) is compared to unencoded results (third and sixth columns) and encoded results with $R = 1$ syndrome extraction rounds (fourth and seventh columns). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.

	$\langle \overline{Z}_n \rangle$					
Logical qubit (b, l_b)	$g_z = 0.5$			$g_z = 1.5$		
	MPS	Unencoded	$R = 1$	MPS	Unencoded	$R = 1$
(6, 0)	0.696	0.231(5)	0.597(16)	0.804	0.371(5)	0.704(32)
(11, 0)	0.762	0.343(5)	0.642(12)	0.845	0.479(5)	0.656(15)
(11, 1)	0.678	0.335(5)	0.579(23)	0.506	0.323(5)	0.466(9)
(8, 0)	0.330	0.149(5)	0.285(23)	0.486	0.267(5)	0.442(15)
(8, 1)	-0.263	-0.077(6)	-0.159(24)	-0.364	-0.135(5)	-0.304(35)
(12, 1)	-0.199	-0.129(6)	-0.182(15)	-0.268	-0.286(5)	-0.234(18)
(12, 0)	-0.125	-0.079(6)	-0.176(11)	-0.120	-0.099(5)	-0.123(31)
(15, 1)	-0.219	-0.083(5)	-0.147(49)	-0.215	-0.125(6)	-0.216(32)
(15, 0)	0.162	0.094(6)	0.164(35)	0.212	0.156(5)	0.310(17)
(14, 1)	0.520	0.248(5)	0.493(12)	0.751	0.350(5)	0.655(16)
(14, 0)	0.228	0.168(5)	0.487(38)	0.767	0.245(5)	0.649(27)
(17, 0)	-0.412	-0.095(6)	-0.288(27)	-0.212	-0.007(6)	-0.172(13)
(17, 1)	-0.107	-0.005(5)	-0.170(15)	-0.297	-0.068(5)	-0.273(46)
(20, 0)	-0.289	-0.065(6)	-0.215(11)	-0.127	-0.144(6)	-0.013(39)
(20, 1)	-0.576	-0.277(5)	-0.313(15)	-0.219	-0.191(5)	-0.194(11)
(19, 0)	-0.679	-0.374(5)	-0.434(10)	-0.102	-0.081(6)	-0.078(39)
(19, 1)	-0.709	-0.287(5)	-0.481(22)	-0.396	-0.168(5)	-0.246(26)
(16, 1)	-0.742	-0.351(5)	-0.618(25)	-0.365	-0.256(6)	-0.249(20)
(16, 0)	-0.630	-0.262(5)	-0.507(21)	-0.076	-0.023(6)	-0.109(20)
(18, 0)	-0.282	-0.092(5)	-0.227(25)	-0.261	-0.230(6)	-0.298(10)
(18, 1)	-0.279	-0.123(6)	-0.200(16)	-0.205	-0.134(6)	-0.133(31)

Table E.9: *Magnetization in 2D Simulations at $t = 4$, Continued.* The remaining logical qubits from Table E.8, following the same column structure. For each logical qubit (b, l_b) (first column), $\langle \overline{Z}_n \rangle$ from MPS results (second and fifth columns) is compared to unencoded results (third and sixth columns) and encoded results with $R = 1$ syndrome extraction rounds (fourth and seventh columns). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.

$ \langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}} $												
%	$g_z = 0.5$						$g_z = 1.5$					
	Unenc.	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$	Unenc.	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$
0%	0.000(1)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)	0.000(0)
10%	0.047(2)	0.014(1)	0.013(1)	0.016(1)	0.016(1)	0.021(2)	0.016(1)	0.011(1)	0.011(1)	0.010(1)	0.011(1)	0.015(1)
20%	0.109(1)	0.035(2)	0.033(2)	0.037(2)	0.046(2)	0.063(3)	0.042(1)	0.022(1)	0.021(1)	0.022(1)	0.024(1)	0.029(2)
30%	0.163(2)	0.063(2)	0.060(2)	0.067(2)	0.081(2)	0.097(3)	0.072(1)	0.036(2)	0.035(1)	0.036(2)	0.041(2)	0.052(2)
40%	0.215(2)	0.091(2)	0.090(2)	0.102(2)	0.118(2)	0.150(3)	0.102(1)	0.052(2)	0.049(2)	0.051(2)	0.057(2)	0.072(2)
50%	0.281(2)	0.123(2)	0.131(3)	0.137(3)	0.165(3)	0.205(4)	0.134(1)	0.067(2)	0.066(2)	0.065(2)	0.074(2)	0.101(2)
60%	0.347(2)	0.172(3)	0.173(3)	0.193(3)	0.222(4)	0.300(4)	0.177(1)	0.087(2)	0.086(2)	0.084(2)	0.096(2)	0.132(3)
70%	0.403(2)	0.255(4)	0.257(4)	0.277(4)	0.315(5)	0.421(5)	0.222(1)	0.113(2)	0.113(2)	0.110(2)	0.123(2)	0.173(3)
80%	0.456(2)	0.342(4)	0.349(5)	0.373(4)	0.425(5)	0.505(4)	0.279(2)	0.146(2)	0.142(2)	0.147(3)	0.162(3)	0.231(3)
90%	0.527(2)	0.448(4)	0.467(4)	0.495(5)	0.548(5)	0.599(4)	0.398(2)	0.200(4)	0.206(4)	0.222(4)	0.236(5)	0.341(6)
100%	0.669(5)	0.667(21)	0.672(20)	0.747(36)	0.766(25)	0.792(34)	0.763(6)	0.518(27)	0.501(15)	0.535(26)	0.540(19)	0.855(28)

Table E.10: *Absolute Error in 2D Simulations.* The absolute error $|\langle \bar{Z} \rangle_{\text{meas}} - \langle \bar{Z} \rangle_{\text{MPS}}|$ by decile (first column) is compared between unencoded runs (second and seventh columns) and encoded runs for various numbers of syndrome extraction rounds R (columns 3-6 and 8-11). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is determined through bootstrap resampling.

Signal-survival factor $\alpha(\sigma)$												
t	$g_z = 0.5$						$g_z = 1.5$					
	Unenc.	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$	Unenc.	$R = 1$	$R = 2$	$R = 4$	$R = 8$	$R = 16$
0.0	0.988(6)	0.988(10)	0.988(12)	0.984(16)	0.986(12)	0.984(18)	0.982(37)	0.987(13)	0.985(15)	0.987(11)	0.985(13)	0.981(16)
0.5	0.862(41)	0.947(37)	0.948(41)	0.943(40)	0.943(44)	0.944(44)	0.850(73)	0.944(44)	0.942(45)	0.946(42)	0.949(33)	0.940(43)
1.0	0.770(79)	0.918(48)	0.917(52)	0.912(48)	0.918(51)	0.883(78)	0.757(106)	0.907(54)	0.909(48)	0.918(48)	0.913(54)	0.879(75)
1.5	0.688(76)	0.896(63)	0.885(66)	0.899(70)	0.861(71)	0.833(104)	0.662(109)	0.892(63)	0.888(68)	0.884(65)	0.881(72)	0.864(90)
2.0	0.597(79)	0.851(84)	0.845(91)	0.867(81)	0.817(95)	0.744(135)	0.607(112)	0.879(69)	0.866(70)	0.858(68)	0.860(78)	0.810(104)
2.5	0.524(72)	0.791(112)	0.787(106)	0.773(109)	0.738(129)	0.664(151)	0.553(98)	0.860(78)	0.855(64)	0.854(77)	0.850(76)	0.762(102)
3.0	0.459(73)	0.734(115)	0.735(115)	0.710(122)	0.667(131)	0.537(159)	0.534(85)	0.854(66)	0.849(68)	0.843(78)	0.836(76)	0.734(102)
3.5	0.411(68)	0.675(126)	0.684(125)	0.642(117)	0.590(140)	0.434(170)	0.487(72)	0.836(72)	0.845(69)	0.812(63)	0.830(68)	0.725(78)
4.0	0.348(63)	0.606(133)	0.627(123)	0.555(133)	0.492(137)	0.323(153)	0.474(66)	0.834(68)	0.831(74)	0.814(71)	0.778(70)	0.620(89)
4.5	0.305(60)	0.551(124)	0.519(124)	0.506(130)	0.387(144)	0.241(144)	0.416(61)	0.788(65)	0.764(77)	0.768(73)	0.730(82)	0.547(95)
5.0	0.251(60)	0.485(121)	0.452(121)	0.412(115)	0.337(139)	0.209(127)	0.370(60)	0.722(64)	0.703(62)	0.725(69)	0.645(86)	0.417(83)
5.5	0.215(54)	0.418(116)	0.395(115)	0.374(127)	0.270(128)	0.143(119)	0.322(57)	0.660(74)	0.668(66)	0.621(71)	0.592(94)	0.321(90)
6.0	0.174(53)	0.399(113)	0.349(110)	0.296(109)	0.248(130)	0.097(74)	0.280(58)	0.618(74)	0.609(68)	0.591(67)	0.558(84)	0.192(64)
6.5	0.149(48)	0.328(99)	0.321(108)	0.267(116)	0.204(110)	0.087(80)	0.241(54)	0.561(70)	0.575(72)	0.598(81)	0.521(87)	0.123(66)
7.0	0.119(45)	0.296(100)	0.276(95)	0.229(100)	0.160(108)	0.047(47)	0.206(52)	0.535(74)	0.543(69)	0.531(81)	0.479(84)	0.117(43)
7.5	0.095(41)	0.263(82)	0.266(100)	0.191(96)	0.133(94)	0.050(48)	0.180(50)	0.505(62)	0.521(66)	0.503(79)	0.444(81)	0.072(36)
8.0	0.074(38)	0.243(87)	0.207(81)	0.167(95)	0.124(75)	0.035(33)	0.161(48)	0.481(64)	0.482(52)	0.471(67)	0.405(78)	0.055(36)

Table E.11: *Signal-Survival Factor for 2D Simulations.* The signal-survival factor α calculated as a function of time t (first column) in unencoded results (second and seventh columns) is compared to encoded results for various numbers of syndrome extraction rounds R (columns 3-6 and 8-11). The left columns show $g_z = 0.5$ and the right columns show $g_z = 1.5$. The uncertainty is the residual error σ in the α fits.

	Acceptance rate			
t	Chain		Grid	
	$g_z = 0.5.$	$g_z = 1.5$	$g_z = 0.5$	$g_z = 1.5$
0.0	0.341 ± 0.339	0.436 ± 0.347	0.414 ± 0.355	0.289 ± 0.314
0.5	0.372 ± 0.309	0.292 ± 0.310	0.312 ± 0.273	0.268 ± 0.280
1.0	0.307 ± 0.291	0.343 ± 0.286	0.339 ± 0.299	0.298 ± 0.261
1.5	0.324 ± 0.300	0.338 ± 0.264	0.303 ± 0.228	0.247 ± 0.251
2.0	0.308 ± 0.264	0.241 ± 0.246	0.229 ± 0.201	0.252 ± 0.231
2.5	0.297 ± 0.261	0.294 ± 0.253	0.226 ± 0.196	0.274 ± 0.221
3.0	0.295 ± 0.281	0.321 ± 0.258	0.208 ± 0.188	0.205 ± 0.183
3.5	0.260 ± 0.227	0.293 ± 0.214	0.200 ± 0.167	0.233 ± 0.197
4.0	0.235 ± 0.201	0.249 ± 0.241	0.176 ± 0.147	0.208 ± 0.191
4.5	0.226 ± 0.209	0.228 ± 0.177	0.255 ± 0.179	0.169 ± 0.160
5.0	0.254 ± 0.230	0.310 ± 0.224	0.196 ± 0.162	0.249 ± 0.199
5.5	0.210 ± 0.151	0.183 ± 0.148	0.255 ± 0.185	0.211 ± 0.172
6.0	0.246 ± 0.196	0.226 ± 0.203	0.200 ± 0.177	0.218 ± 0.141
6.5	0.210 ± 0.188	0.231 ± 0.197	0.229 ± 0.178	0.216 ± 0.165
7.0	0.209 ± 0.195	0.214 ± 0.197	0.139 ± 0.081	0.227 ± 0.155
7.5	0.228 ± 0.190	0.233 ± 0.207	0.173 ± 0.136	0.173 ± 0.124
8.0	0.197 ± 0.161	0.187 ± 0.179	0.201 ± 0.158	0.171 ± 0.129

Table E.12: *Acceptance Rates of $R = 1$ Encoded Simulations.* The acceptance rates are shown for both 1D simulations (second and third columns) and 2D simulations (fourth and fifth columns) as a function of time t (first column). These encoded simulations use $R = 1$ syndrome extraction rounds, corresponding to the data reported in the main text. The uncertainty represents the standard deviation over all qubits for a given t . The uncertainty is the standard deviation over all qubits for each t .

BIBLIOGRAPHY

- [1] P. Anderson. Absence of diffusion in certain random lattices. *Physical Review*, 109:1492–1505, 1958.
- [2] D. Thouless et al. Quantized hall conductance in a two-dimensional periodic potential. *Physical Review Letters*, 49(6):405, 1982.
- [3] J. Maldacena. The Large N limit of superconformal field theories and supergravity. *Adv. Theor. Math. Phys.*, 2:231–252, 1997.
- [4] S. Glashow. Partial-symmetries of weak interactions. *Nuclear Physics*, 22(4):579–588, 1961.
- [5] P. Higgs. Broken symmetries and the masses of gauge bosons. *Physical Review Letters*, 13(16):508–509, 1964.
- [6] S. Weinberg. A model of leptons. *Physical Review Letters*, 19:1264–1266, 1967.
- [7] D. Gross and F. Wilczek. Ultraviolet behavior of non-abelian gauge theories. *Physical Review Letters*, 30(26):1343–1346, 1973.
- [8] H. Politzer. Reliable perturbative results for strong interactions? *Phys. Rev. Lett.*, 30:1346–1349, 1973.
- [9] K. Wilson. Confinement of quarks. *Physical Review D*, 10:2445–2459, 1974.
- [10] M. Creutz. Monte Carlo study of quantized SU(2) gauge theory. *Physical Review D*, 21(8):2308–2315, 1980.
- [11] K. Symanzik. Continuum limit and improved action in lattice theories. 1. principles and φ^4 theory. *Nucl. Phys. B*, 226:187–204, 1983.
- [12] M. Luscher. Two particle states on a torus and their relation to the scattering matrix. *Nuclear Physics B*, 354(2-3):531–578, 1991.
- [13] D. Kaplan. A method for simulating chiral fermions on the lattice. *Physics Letters B*, 288(3-4):342–347, 1992.
- [14] S. Beane et al. Light nuclei and hypernuclei from quantum chromodynamics in the limit of SU(3) flavor symmetry. *Physical Review D*, 87(3):034506, 2013.
- [15] S. Beane and M. Savage. The quark mass dependence of two nucleon systems. *Nucl. Phys. A*, 717:91–103, 2002.

- [16] S. Beane et al. $I = 2$ π - π scattering from fully-dynamical mixed-action lattice QCD. *Physical Review D*, 73:054503, 2006.
- [17] S. Beane et al. Nucleon-nucleon scattering from fully-dynamical lattice QCD. *Physical Review Letters*, 97(1):012001, 2006.
- [18] M. Troyer and U. Wiese. Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations. *Physical Review Letters*, 94:170201, 2005.
- [19] H. Raedt et al. Massively parallel quantum computer simulator, eleven years later. *Computer Physics Communications*, 237:47–61, 2018.
- [20] A. Aspect, J. Dalibard, and G. Roger. Experimental test of Bell’s inequalities using time varying analyzers. *Physical Review Letters*, 49(25):1804–1807, 1982.
- [21] P. Benioff. The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines. *Journal of Statistical Physics*, 22(5):563–591, 1980.
- [22] R. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6–7):467–488, 1982.
- [23] R. Feynman. Quantum mechanical computers. *Foundations of Physics*, 16(6):507–531, 1986.
- [24] A. Kitaev. Quantum measurements and the Abelian stabilizer problem. *Electron. Colloquium Comput. Complex.*, 1995.
- [25] A. Aspuru-Guzik et al. Simulated quantum computation of molecular energies. *Science*, 309(5741):1704–1707, 2005.
- [26] J. McClean et al. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.*, 18(2):023023, 2016.
- [27] S. Jordan et al. BQP-completeness of scattering in scalar quantum field theory. *Quantum*, 2:44, 2018.
- [28] S. Jordan, K. Lee, and J. Preskill. Quantum computation of scattering in scalar quantum field theories. *Quantum Information Computation*, 14(11–12):1014–1080, 2011.
- [29] A. Kitaev and W. Webb. Wavefunction preparation and resampling using a quantum computer. January 2008. arXiv:0801.0342 [quant-ph].
- [30] E. Wigner and P. Dirac. On the statistical distribution of the widths and spacings of nuclear resonance levels. *Mathematical Proceedings of the Cambridge Philosophical Society*, 47(4):790–798, 1951.

- [31] E. Wigner. Characteristic vectors of bordered matrices with infinite dimensions. *The Annals of Mathematics*, 62(3):548, 1955.
- [32] F. Dyson. Statistical theory of the energy levels of complex systems. III. *Journal of Mathematical Physics*, 3(1):166–175, 1962.
- [33] O. Bohigas, M. Giannoni, and C. Schmit. Characterization of chaotic quantum spectra and universality of level fluctuation laws. *Physical Review Letters*, 52(1):1–4, 1984.
- [34] M. Berry and M. Tabor. Level clustering in the regular spectrum. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 356(1686):375–394, 1977.
- [35] H. Trotter. On the product of semi-groups of operators. *Proceedings of the American Mathematical Society*, 10(4):545–551, 1959.
- [36] A. Childs et al. Theory of trotter error with commutator scaling. *Physical Review X*, 11(1):011020, 2021.
- [37] H. Huang, R. Kueng, and J. Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.
- [38] Z. Cai et al. Quantum error mitigation. *Reviews of Modern Physics*, 95(4):045005, 2023.
- [39] N. Ezzell et al. Dynamical decoupling for superconducting qubits: A performance survey. *Physical Review Applied*, 20(6):064027, 2023.
- [40] J. Wallman and J. Emerson. Noise tailoring for scalable quantum computation via randomized compiling. *Physical Review A*, 94(5):052325, 2016.
- [41] K. Temme, S. Bravyi, and J. Gambetta. Error mitigation for short-depth quantum circuits. *Physical Review Letters*, 119(18):180509, 2017.
- [42] E. van den Berg et al. Probabilistic error cancellation with sparse Pauli–Lindblad models on noisy quantum processors. *Nature Physics*, 19(8):1116–1121, 2023.
- [43] G. Wick. The evaluation of the collision matrix. *Physical Review*, 80(2):268–272, 1950.
- [44] M. Peskin and D. Schroeder. *An introduction to quantum field theory*. CRC press, 1995.
- [45] L. Isserlis. On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables. *Biometrika*, 12(1/2):134, 1918.
- [46] E. Lieb and F. Wu. Absence of Mott transition in an exact solution of the short-range, one-band model in one dimension. *Physical Review Letters*, 20:1445–1448, 1968. [Erratum: Phys.Rev.Lett. 21, 192 (1968)].

- [47] N. Schuch and F. Verstraete. Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nature Physics*, 5(10):732–735, 2009.
- [48] E. Lieb, T. Schultz, and D. Mattis. Two soluble models of an antiferromagnetic chain. *Annals of Physics*, 16(3):407–466, 1961.
- [49] J. Kogut and L. Susskind. Hamiltonian formulation of Wilson’s lattice gauge theories. *Physical Review D*, 11(2):395–408, 1975.
- [50] N. Zemlevskiy, H. Froland, and S. Caspar. Optimization of algorithmic errors in analog quantum simulations. *Physical Review A*, 109(5):052425, 2024.
- [51] H. Froland, D. Grabowska, and Z. Li. Simulating fully gauge-fixed SU(2) hamiltonian dynamics on digital quantum computers. December 2025. arXiv:2512.22782 [quant-ph].
- [52] H. Froland and D. Grabowska. Measuring non-stabilizerness in an SU(2) lattice gauge theory. June 2026. arXiv:2606.14842 [quant-ph].
- [53] J. Bringewatt et al. Classical shadows for sample-efficient measurements of gauge-invariant observables. *Quantum*, 10:2127, 2026.
- [54] H. Froland et al. Entanglement structure of non-gaussian states and how to measure it. *Physical Review Letters*, 135(4):040201, 2025.
- [55] H. Froland et al. Realizing Error Suppression in Partially Fault-Tolerant Quantum Simulations with IBM Quantum Computers. July 2026. arXiv:2607.24947 [quant-ph].
- [56] J. Preskill. Simulating quantum field theory with a quantum computer. *PoS, LATTICE2018*:024, 2018. arXiv:1811.10085 [hep-lat].
- [57] S. Jordan, K. Lee, and J. Preskill. Quantum algorithms for quantum field theories. *Science*, 336:1130–1133, 2012.
- [58] S. Jordan, K. Lee, and J. Preskill. Quantum algorithms for fermionic quantum field theories. April 2014. arXiv:1404.7115 [hep-th].
- [59] C. Bauer et al. Quantum simulation of fundamental particles and forces. *Nature Reviews Physics*, 5(7):420–432, 2023.
- [60] M. Unsal. Theta dependence, sign problems and topological interference. *Physical Review D*, 86(10):105012, 2012.
- [61] K. Bharti et al. Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, 94(1):015004, 2022.

- [62] Y. Atas et al. SU(2) hadrons on a quantum computer via a variational approach. *Nature Communications*, 12(1):6499, 2021.
- [63] Y. Kim et al. Evidence for the utility of quantum computing before fault tolerance. *Nature*, 618(7965):500–505, 2023.
- [64] R. Farrell et al. Preparations for quantum simulations of quantum chromodynamics in 1+1 dimensions. I. Axial gauge. *Physical Review D*, 107(5):054512, 2023.
- [65] R. Farrell et al. Preparations for quantum simulations of quantum chromodynamics in 1+1 dimensions. II. Single-baryon β -decay in real time. *Physical Review D*, 107(5):054513, 2023.
- [66] A. Ciavarella, N. Klco, and M. Savage. Trailhead for quantum simulation of SU(3) Yang-Mills lattice gauge theory in the local multiplet basis. *Physical Review D*, 103(9), 2021.
- [67] I. Georgescu, S. Ashhab, and F. Nori. Quantum Simulation. *Reviews of Modern Physics*, 86:153, 2014.
- [68] I. Arrazola et al. Digital-analog quantum simulation of spin models in trapped ions. *Scientific Reports*, 6(1), 2016.
- [69] L. Lamata et al. Digital-analog quantum simulations with superconducting circuits. *Advances in Physics: X*, 3(1):1457981, 2018.
- [70] S. Flannigan et al. Propagation of errors and quantitative quantum simulation with quantum advantage. *Quantum Science and Technology*, 7(4):045025, 2022.
- [71] P. Scholl et al. Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms. *Nature*, 595:233–238, 2021.
- [72] S. Ebadi et al. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature*, 595(7866):227–232, 2021.
- [73] A. Daley et al. Practical quantum advantage in quantum simulation. *Nature*, 607(7920):667–676, 2022.
- [74] C. Monroe et al. Programmable quantum simulations of spin systems with trapped ions. *Reviews of Modern Physics*, 93(2):025001, 2021.
- [75] U. Heras et al. Digital quantum simulation of spin systems in superconducting circuits. *Physical Review Letters*, 112:200501, 2014.
- [76] B. Andrade et al. Engineering an effective three-spin Hamiltonian in trapped-ion systems for applications in quantum simulation. *Quantum Science and Technology*, 7(3):034001, 2022.

- [77] S. Gong, W. Zhu, and D. Sheng. Emergent chiral spin liquid: Fractional quantum hall effect in a kagome heisenberg model. *Scientific Reports*, 4(1), 2014.
- [78] X. Ma et al. Quantum simulation of the wavefunction to probe frustrated heisenberg spin systems. *Nature Physics*, 7(5):399–405, 2011.
- [79] T. Cubitt, A. Montanaro, and S. Piddock. Universal quantum hamiltonians. *Proceedings of the National Academy of Sciences of the United States of America*, 115(38):9497–9502, 2018.
- [80] A. Childs et al. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences of the United States of America*, 115(38):9456–9461, 2018.
- [81] B. Nachman et al. Quantum Algorithm for High Energy Physics Simulations. *Physical Review Letters*, 126(6):062001, 2021.
- [82] C. Bauer et al. Quantum simulation for high-energy physics. *PRX Quantum*, 4(2):027001, 2023.
- [83] S. Caspar and H. Singh. From asymptotic freedom to θ vacua: Qubit embeddings of the O(3) nonlinear σ model. *Physical Review Letters*, 129(2):022003, 2022.
- [84] A. Ciavarella et al. Preparation for quantum simulation of the (1+1)-dimensional O(3) nonlinear σ model using cold atoms. *Physical Review A*, 107(4):042404, 2023.
- [85] J. Maldacena. A simple quantum system that describes a black hole. March 2023. arXiv:2303.11534 [hep-th].
- [86] A. Florio et al. Real-time nonperturbative dynamics of jet production in schwinger model: quantum entanglement and vacuum modification. *Physical Review Letters*, 131(2):021902, 2023.
- [87] J. Martin et al. Classical and quantum evolution in a simple coherent neutrino problem. *Physical Review D*, 105(8):083020, 2022.
- [88] R. Verresen. Everything is a quantum Ising model. January 2023. arXiv:2301.11917 [quant-ph].
- [89] L. Martin et al. Controlling local thermalization dynamics in a floquet-engineered dipolar ensemble. *Physical Review Letters*, 130(21):210403, 2023.
- [90] H. Zhou et al. Quantum metrology with strongly interacting spin systems. *Physical Review X*, 10(3):031003, 2020.
- [91] M. Tyler et al. Higher-order methods for Hamiltonian engineering pulse sequence design. *Physical Review A*, 108(6):062602, 2023.

- [92] H. Zhou et al. Robust higher-order hamiltonian engineering for quantum sensing with strongly interacting systems. *Physical Review Letters*, 131(22):220803, 2023.
- [93] H. Zhou et al. Robust hamiltonian engineering for interacting qudit systems. *Physical Review X*, 14(3):031017, 2024.
- [94] S. Geier et al. Floquet hamiltonian engineering of an isolated many-body spin system. *Science*, 374(6571):1149–1152, 2021.
- [95] P. Scholl et al. Microwave engineering of programmable XXZ hamiltonians in arrays of rydberg atoms. *PRX Quantum*, 3(2):020303, 2022.
- [96] J. Choi et al. Robust dynamic hamiltonian engineering of many-body spin systems. *Physical Review X*, 10(3):031002, 2020.
- [97] P. Richerme et al. Non-local propagation of correlations in quantum systems with long-range interactions. *Nature*, 511(7508):198–201, 2014.
- [98] P. Jurcevic et al. Quasiparticle engineering and entanglement propagation in a quantum many-body system. *Nature*, 511(7508):202–205, 2014.
- [99] M. Wall, A. Safavi-Naini, and A. Rey. Boson-mediated quantum spin simulators in transverse fields: xy model and spin-boson entanglement. *Physical Review A*, 95:013602, 2017.
- [100] T. Kiely and J. Freericks. Relationship between the transverse-field Ising model and the XY model via the rotating-wave approximation. *Physical Review A*, 97:023611, 2018.
- [101] A. Ciavarella et al. Simulating heisenberg interactions in the Ising model with strong drive fields. *Physical Review A*, 108(4):042216, 2023.
- [102] A. Bermudez et al. Long-range Heisenberg models in quasiperiodically driven crystals of trapped ions. *Physical Review B*, 95(2):024431, 2017.
- [103] T. Gonzalez-Raya et al. Digital-analog quantum simulations using the cross-resonance effect. *PRX Quantum*, 2:020328, 2021.
- [104] L. Henri et al. Quantum computing with neutral atoms. *Quantum*, 4:327, 2020.
- [105] A. Browaeys and T. Lahaye. Many-body physics with individually controlled rydberg atoms. *Nature Physics*, 16(2):132–142, 2020.
- [106] Z. Mádi et al. Time-resolved observation of spin waves in a linear chain of nuclear spins. *Chemical Physics Letters*, 268(3-4):300–305, 1997.

- [107] G. Roumpos, C. Master, and Y. Yamamoto. Quantum simulation of spin ordering with nuclear spins in a solid state lattice. *Physical Review B*, 75:094415, 2007.
- [108] Y. Salathé et al. Digital quantum simulation of spin models with circuit quantum electrodynamics. *Physical Review X*, 5(2):021027, 2015.
- [109] M. Johnson et al. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, 2011.
- [110] S. Lloyd and J. Slotine. Analog quantum error correction. *Physical Review Letters*, 80(18):4088–4091, 1998.
- [111] K. Fukui, A. Tomita, and A. Okamoto. Analog quantum error correction with encoding a qubit into an oscillator. *Physical Review Letters*, 119(18):180507, 2017.
- [112] N. Klco and M. Savage. Digitization of scalar fields for quantum computing. *Physical Review A*, 99(5):052335, 2019.
- [113] N. Wiebe et al. Higher order decompositions of ordered operator exponentials. *J. Phys. A*, 43(6):065203, 2010.
- [114] A. Childs and Y. Su. Nearly optimal lattice simulation by product formulas. *Physical Review Letters*, 123(5):050503, 2019.
- [115] S. Endo, S. Benjamin, and Y. Li. Practical quantum error mitigation for near-future applications. *Physical Review X*, 8(3):031027, 2018.
- [116] R. Shaffer et al. Practical verification protocols for analog quantum simulators. *npj Quantum Information*, 7(1):46, 2021.
- [117] R. Trivedi, A. Rubio, and J. Cirac. Quantum advantage and stability to errors in analogue quantum simulators. *Nature Communications*, 15(1):6507, 2024.
- [118] P. Poggi et al. Quantifying the sensitivity to errors in analog quantum simulation. *PRX Quantum*, 1(2):020308, 2020.
- [119] I. Cong et al. Hardware-efficient, fault-tolerant quantum computation with rydberg atoms. *Physical Review X*, 12(2):021049, 2022.
- [120] L. Masanes, G. Vidal, and J. Latorre. Time-optimal hamiltonian simulation and gate synthesis using homogeneous local unitaries. *Quant. Inf. Comput.*, 2(4):285–296, 2002.
- [121] M. Suzuki. Generalized Trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Communications in Mathematical Physics*, 51(2):183–190, 1976.

- [122] A. Ciavarella et al. State preparation in the heisenberg model through adiabatic spiraling. *Quantum*, 7:970, 2023.
- [123] T. Eckstein et al. Large-scale simulations of Floquet physics on near-term quantum computers. *npj Quantum Information*, 10(1):84, 2024.
- [124] M. Saffman, T. Walker, and K. Mølmer. Quantum information with rydberg atoms. *Reviews of Modern Physics*, 82:2313–2363, 2010.
- [125] X. Wu et al. A concise review of Rydberg atom based quantum computation and quantum simulation. *Chinese Physics B*, 30(2):020305, 2020.
- [126] J. Wurtz et al. Aquila: QuEra’s 256-qubit neutral-atom quantum computer. June 2023. arXiv:2306.11727 [quant-ph].
- [127] B. Şahinoğlu and R. Somma. Hamiltonian simulation in the low-energy subspace. *npj Quantum Information*, 7:119, 2021.
- [128] M. Heyl, P. Hauke, and P. Zoller. Quantum localization bounds Trotter errors in digital quantum simulation. *Science Advances*, 5(4):aau8342, 2019.
- [129] K. Chinni et al. Trotter errors from dynamical structural instabilities of floquet maps in quantum simulation. *PRX Quantum*, 3(1):010351, 2022.
- [130] M. Schlosser et al. Assembled arrays of rydberg-interacting atoms. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 53(14):144001, 2020.
- [131] T. Zache et al. Dynamical topological transitions in the massive schwinger model with a θ term. *Physical Review Letters*, 122(5):050403, 2019.
- [132] R. Samajdar et al. Quantum phases of Rydberg atoms on a kagome lattice. *Proceedings of the National Academy of Sciences of the United States of America*, 118(4):e2015785118, 2021.
- [133] B. Müller and X. Yao. Simple Hamiltonian for quantum simulation of strongly coupled (2+1)D SU(2) lattice gauge theory on a honeycomb lattice. *Physical Review D*, 108(9):094505, 2023.
- [134] K. Adcox. Formation of dense partonic matter in relativistic nucleus-nucleus collisions at RHIC: Experimental evaluation by the PHENIX collaboration. *Nucl. Phys. A*, 757:184–283, 2004.
- [135] I. Arsene et al. Quark gluon plasma and color glass condensate at RHIC? The Perspective from the BRAHMS experiment. *Nucl. Phys. A*, 757:1–27, 2004.
- [136] B. Back et al. The PHOBOS perspective on discoveries at RHIC. *Nucl. Phys. A*, 757:28–101, 2005.

- [137] S. Adams et al. Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR Collaboration’s critical assessment of the evidence from RHIC collisions. *Nucl. Phys. A*, 757:102–183, 2005.
- [138] M. He, H. van Hees, and R. Rapp. Heavy-quark diffusion in the quark–gluon plasma. *Prog. Part. Nucl. Phys.*, 130:104020, 2023.
- [139] A. Salam and J. Ward. Weak and electromagnetic interactions. *Il Nuovo Cimento*, 11:568–577, 1959.
- [140] E. Majorana. Teoria simmetrica dell’elettrone e del positrone. *Il Nuovo Cimento*, 14(4):171–184, 1937.
- [141] K. Nakamura et al. Review of particle physics. *Physical Review D*, 110(3):030001, 2024.
- [142] T. Blum. Lattice calculation of the lowest order hadronic contribution to the muon anomalous magnetic moment. *Physical Review Letters*, 91:052001, 2003.
- [143] Z. Bai et al. Standard model prediction for direct cp violation in $k \rightarrow \pi\pi$ decay. *Physical Review Letters*, 115(21):212001, 2015.
- [144] S. Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.
- [145] J. Schachenmayer et al. Thermalization of strongly interacting bosons after spontaneous emissions in optical lattices. *EPJ Quant. Technol.*, 2:1, 2015.
- [146] R. Nandkishore and D. Huse. Many-body localization and thermalization in quantum statistical mechanics. *Annu. Rev. Condens. Matter Phys.*, 6(1):15–38, 2014.
- [147] A. Kaufman et al. Quantum thermalization through entanglement in an isolated many-body system. *Science*, 353(6301):794–800, 2016.
- [148] N. Mueller et al. Quantum computing universal thermalization dynamics in a $(2 + 1)d$ lattice gauge theory. *Nature Communications*, 16(1):5492, 2025.
- [149] T. Cochran et al. Visualizing dynamics of charges and strings in $(2 + 1)d$ lattice gauge theories. *Nature*, 642(8067):315–320, 2025.
- [150] D. González-Cuadra et al. Observation of string breaking on a $(2 + 1)D$ Rydberg quantum simulator. *Nature*, 642(8067):321–326, 2025.
- [151] J. Halimeh et al. Quantum simulation of out-of-equilibrium dynamics in gauge theories. September 2025. arXiv:2509.03586 [quant-ph].

- [152] P. Dumitrescu et al. Dynamical topological phase realized in a trapped-ion quantum simulator. *Nature*, 607(7919):463–467, 2022.
- [153] T. Hayata, K. Seki, and A. Yamamoto. Floquet prethermalization of Z2 lattice gauge theory on superconducting qubits. *Physical Review D*, 110(11):114503, 2024.
- [154] T. Angelides et al. First-order phase transition of the Schwinger model with a quantum computer. *npj Quantum Information*, 11(1):6, 2025.
- [155] M. Meth et al. Simulating two-dimensional lattice gauge theories on a qudit quantum computer. *Nature Physics*, 21(4):570–576, 2025.
- [156] C. Kokail et al. Self-verifying variational quantum simulation of lattice models. *Nature*, 569(7756):355–360, 2019.
- [157] B. Yang et al. Observation of gauge invariance in a 71-site Bose–Hubbard quantum simulator. *Nature*, 587(7834):392–396, 2020.
- [158] A. Mil et al. A scalable realization of local U(1) gauge invariance in cold atomic mixtures. *Science*, 367(6482):1128–1130, 2020.
- [159] C. Bauer, M. Freytsis, and B. Nachman. Simulating collider physics on quantum computers using effective field theories. *Physical Review Letters*, 127(21):212001, 2021.
- [160] D. Mark et al. Observation of ballistic plasma and memory in high-energy gauge theory dynamics. October 2025. arXiv:2510.11679 [quant-ph].
- [161] Z. Zhou et al. Thermalization dynamics of a gauge theory on a quantum simulator. *Science*, 377(6603):311–314, 2022.
- [162] G. Gyawali et al. Observation of disorder-free localization using a $(2+1)$ d lattice gauge theory on a quantum processor. *Science*, 393(6806):71–75, 2026.
- [163] D. Paulson et al. Simulating 2D effects in lattice gauge theories on a quantum computer. *PRX Quantum*, 2(3):030334, 2021.
- [164] F. Turro, A. Ciavarella, and X. Yao. Classical and quantum computing of shear viscosity for $(2+1)$ D SU(2) gauge theory. *Physical Review D*, 109(11):114511, 2024.
- [165] A. Kandala et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, 2017.
- [166] G. Semeghini et al. Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572):abi8794, 2021.

- [167] S. Thompson and G. Siopsis. Quantum computation of phase transition in interacting scalar quantum field theory. *Quantum Information Processing*, 22(11):396, 2023.
- [168] D. Abanin et al. Observation of constructive interference at the edge of quantum ergodicity. *Nature*, 646(8086):825–830, 2025.
- [169] A. Morvan et al. Phase transitions in random circuit sampling. *Nature*, 634(8033):328–333, 2024.
- [170] K. Shinjo et al. Unveiling clean two-dimensional discrete time crystals on a digital quantum computer. *npj Quantum Information*, 12(1):41, 2026.
- [171] N. Mueller et al. Quantum computation of dynamical quantum phase transitions and entanglement tomography in a lattice gauge theory. *PRX Quantum*, 4(3):030323, 2023.
- [172] A. Than et al. The phase diagram of quantum chromodynamics in one dimension on a quantum computer. *Nature Communications*, 16(1):10288, 2025.
- [173] S. Schuster et al. Studying the phase diagram of the three-flavor Schwinger model in the presence of a chemical potential with measurement- and gate-based quantum computing. *Physical Review D*, 109(11):114508, 2024.
- [174] J. Mildenerger et al. Confinement in a $\mathbb{Z}_2$ lattice gauge theory on a quantum computer. *Nature Physics*, 21(2):312–317, 2025.
- [175] G. Pardo et al. Resource-efficient quantum simulation of lattice gauge theories in arbitrary dimensions: Solving for Gauss’s law and fermion elimination. *Physical Review Research*, 5(2):023077, 2023.
- [176] D. Banerjee et al. Atomic quantum simulation of dynamical gauge fields coupled to fermionic matter: From string breaking to evolution after a quench. *Physical Review Letters*, 109(17):175302, 2012.
- [177] P. Hauke et al. Quantum simulation of a lattice Schwinger model in a chain of trapped ions. *Physical Review X*, 3(4):041018, 2013.
- [178] E. Zohar, J. Cirac, and B. Reznik. Quantum simulations of gauge theories with ultracold atoms: local gauge invariance from angular momentum conservation. *Physical Review A*, 88(2):023617, 2013.
- [179] S. Kühn, J. Cirac, and M. Bañuls. Quantum simulation of the Schwinger model: A study of feasibility. *Physical Review A*, 90(4):042305, 2014.
- [180] V. Kasper et al. Schwinger pair production with ultracold atoms. *Phys. Lett. B*, 760:742–746, 2015.
- [181] E. Zohar, J. Cirac, and B. Reznik. Quantum simulations of lattice gauge theories using ultracold atoms in optical lattices. *Reports on Progress in Physics. Physical Society*, 79(1):014401, 2015.

- [182] E. Martinez et al. Real-time dynamics of lattice gauge theories with a few-qubit quantum computer. *Nature*, 534(7608):516–519, 2016.
- [183] D. Yang et al. Analog quantum simulation of (1+1)-dimensional lattice QED with trapped ions. *Physical Review A*, 94(5):052321, 2016.
- [184] N. Klco et al. Quantum-classical computation of Schwinger model dynamics using quantum computers. *Physical Review A*, 98(3):032331, 2018.
- [185] H. Lu et al. Simulations of subatomic many-body physics on a quantum frequency processor. *Physical Review A*, 100(1):012320, 2019.
- [186] D. Kaplan and J. Stryker. Gauss’s law, duality, and the Hamiltonian formulation of U(1) lattice gauge theory. *Physical Review D*, 102(9):094515, 2020.
- [187] Z. Davoudi et al. Towards analog quantum simulations of lattice gauge theories with trapped ions. *Physical Review Research*, 2(2):023015, 2020.
- [188] F. Surace et al. Lattice gauge theories and string dynamics in Rydberg atom quantum simulators. *Physical Review X*, 10(2):021041, 2020.
- [189] J. Haase et al. A resource efficient approach for quantum and classical simulations of gauge theories in particle physics. *Quantum*, 5:393, 2021.
- [190] D. Luo et al. Framework for simulating gauge theories with dipolar spin systems. *Physical Review A*, 102(3):032617, 2020.
- [191] A. Shaw et al. Quantum algorithms for simulating the lattice Schwinger model. *Quantum*, 4:306, 2020.
- [192] R. Ott et al. Scalable cold-atom quantum simulator for two-dimensional QED. *Physical Review Letters*, 127(13):130504, 2021.
- [193] N. Nguyen et al. Digital quantum simulation of the Schwinger model and symmetry protection with trapped ions. *PRX Quantum*, 3(2):020324, 2022.
- [194] H. Riechert et al. Engineering a U(1) lattice gauge theory in classical electric circuits. *Physical Review B*, 105(20):205141, 2022.
- [195] C. Bauer and D. Grabowska. Efficient representation for simulating U(1) gauge theories on digital quantum computers at all values of the coupling. *Physical Review D*, 107(3):L031503, 2023.
- [196] D. Grabowska et al. Overcoming exponential scaling with system size in Trotter-Suzuki implementations of constrained Hamiltonians: 2+1 U(1) lattice gauge theories. August 2022. arXiv:2208.03333 [quant-ph].

- [197] C. Kane et al. Efficient quantum implementation of 2+1 U(1) lattice gauge theories with Gauss law constraints. November 2022. arXiv:2211.10497 [quant-ph].
- [198] W. Zhang et al. Observation of microscopic confinement dynamics by a tunable topological θ -angle. *Nature Physics*, 21:155–160, 2025.
- [199] R. Farrell et al. Scalable circuits for preparing ground states on digital quantum computers: The schwinger model vacuum on 100 qubits. *PRX Quantum*, 5(2):020315, 2024.
- [200] L. Nagano, A. Bapat, and C. Bauer. Quench dynamics of the Schwinger model via variational quantum algorithms. *Physical Review D*, 108(3):034501, 2023.
- [201] E. Zohar, J. Cirac, and B. Reznik. Cold-atom quantum simulator for SU(2) Yang-Mills lattice gauge theory. *Physical Review Letters*, 110(12):125304, 2013.
- [202] K. Stannigel et al. Constrained dynamics via the Zeno effect in quantum simulation: Implementing non-Abelian lattice gauge theories with cold atoms. *Physical Review Letters*, 112(12):120406, 2014.
- [203] A. Mezzacapo et al. Non-Abelian $SU(2)$ lattice gauge theories in superconducting circuits. *Physical Review Letters*, 115(24):240502, 2015.
- [204] M. Mathur and T. Sreeraj. Canonical transformations and loop formulation of SU(N) lattice gauge theories. *Physical Review D*, 92(12):125018, 2015.
- [205] I. Raychowdhury and J. Stryker. Solving gauss’s law on digital quantum computers with loop-string-hadron digitization. *Physical Review Research*, 2(3):033039, 2020.
- [206] I. Raychowdhury and J. Stryker. Loop, string, and hadron dynamics in SU(2) Hamiltonian lattice gauge theories. *Physical Review D*, 101(11):114502, 2020.
- [207] N. Klco, J. Stryker, and M. Savage. SU(2) non-Abelian gauge field theory in one dimension on digital quantum computers. *Physical Review D*, 101(7):074512, 2020.
- [208] R. Dasgupta and I. Raychowdhury. Cold-atom quantum simulator for string and hadron dynamics in non-Abelian lattice gauge theory. *Physical Review A*, 105(2):023322, 2022.
- [209] Z. Davoudi, I. Raychowdhury, and A. Shaw. Search for efficient formulations for Hamiltonian simulation of non-Abelian lattice gauge theories. *Physical Review D*, 104(7):074505, 2021.
- [210] S. Rahman et al. SU(2) lattice gauge theory on a quantum annealer. *Physical Review D*, 104(3):034501, 2021.
- [211] J. Osborne et al. Large-scale 2 + 1D U(1) gauge theory with dynamical matter in a cold-atom quantum simulator. *Commun. Phys.*, 8(1):273, 2025.

- [212] J. Halimeh, H. Lang, and P. Hauke. Gauge protection in non-Abelian lattice gauge theories. *New J. Phys.*, 24(3):033015, 2022.
- [213] S. Rahman et al. Self-mitigating Trotter circuits for SU(2) lattice gauge theory on a quantum computer. *Physical Review D*, 106(7):074502, 2022.
- [214] T. Zache, D. González-Cuadra, and P. Zoller. Quantum and classical spin-network algorithms for q -deformed Kogut-Susskind gauge theories. *Physical Review Letters*, 131:171902, 2023.
- [215] A. Alexandru et al. Fuzzy gauge theory for quantum computers. *Physical Review D*, 109(9):094502, 2024.
- [216] I. D’Andrea et al. New basis for Hamiltonian SU(2) simulations. *Physical Review D*, 109(7):074501, 2024.
- [217] P. Fontana, M. Rianza, and A. Celi. Efficient finite-resource formulation of non-Abelian lattice gauge theories beyond one dimension. *Physical Review X*, 15(3):031065, 2025.
- [218] D. Grabowska, C. Kane, and C. Bauer. Fully gauge-fixed SU(2) Hamiltonian for quantum simulations. *Physical Review D*, 111(11):114516, 2025.
- [219] R. Anishetty, M. Mathur, and I. Raychowdhury. Prepotential formulation of SU(3) lattice gauge theory. *J. Phys. A*, 43:035403, 2009.
- [220] A. Alexandru et al. Gluon field digitization for quantum computers. *Physical Review D*, 100(11):114501, 2019.
- [221] Y. Atas et al. Simulating one-dimensional quantum chromodynamics on a quantum computer: Real-time evolutions of tetra- and pentaquarks. *Physical Review Research*, 5(3):033184, 2023.
- [222] A. Ciavarella and I. Chernyshev. Preparation of the SU(3) lattice Yang-Mills vacuum with variational quantum methods. *Phys. Rev. D*, 105(7):074504, 2022.
- [223] A. Ciavarella. Quantum simulation of lattice QCD with improved Hamiltonians. *Physical Review D*, 108(9), 2023.
- [224] T. Hayata and Y. Hidaka. q deformed formulation of Hamiltonian SU(3) Yang-Mills theory. *Journal of High Energy Physics*, 2023(9):123, 2023.
- [225] R. Farrell et al. Quantum simulations of hadron dynamics in the Schwinger model using 112 qubits. *Physical Review D*, 109(11), 2024.
- [226] A. Ciavarella and C. Bauer. Quantum simulation of SU(3) lattice Yang-Mills theory at leading order in large- N_c expansion. *Physical Review Letters*, 133(11):111901, 2024.

- [227] P. Balaji et al. Quantum circuits for SU(3) lattice gauge theory. *Physical Review D*, 112(5):054511, 2025.
- [228] N. Klco and M. Savage. Fixed-point quantum circuits for quantum field theories. *Physical Review A*, 102(5):052422, 2020.
- [229] M. Carena et al. Improved hamiltonians for quantum simulations of gauge theories. *Physical Review Letters*, 129(5):051601, 2022.
- [230] A. De et al. Observation of string-breaking dynamics in a quantum simulator. October 2024. arXiv:2410.13815 [quant-ph].
- [231] S. Gu, R. Somma, and B. Sahinoglu. Fast-forwarding quantum evolution. *Quantum*, 5:577, 2021.
- [232] IBM Quantum. IBM Quantum Computing — Summit 2023: The Era of Utility, 2023. <https://www.ibm.com/quantum/summit-2023>,.
- [233] Quantinuum. Quantinuum Unveils Accelerated Roadmap to Achieve Universal, Fully Fault-Tolerant Quantum Computing by 2030, 2024. <https://www.quantinuum.com/press-releases/quantinuum-unveils-accelerated-roadmap-to-achieve-universal-fault-tolerant-quantum-computing-by-2030>,.
- [234] N. Chiu et al. Continuous operation of a coherent 3,000-qubit system. *Nature*, 646(8087):1075–1080, 2025.
- [235] D. Bluvstein et al. Logical quantum processor based on reconfigurable atom arrays. *Nature*, 626(7997):58–65, 2024.
- [236] S. Evered et al. High-fidelity parallel entangling gates on a neutral-atom quantum computer. *Nature*, 622(7982):268–272, 2023.
- [237] R. Finkelstein et al. Universal quantum operations and ancilla-based read-out for tweezer clocks. *Nature*, 634(8033):321–327, 2024.
- [238] N. Zemlevskiy. Scalable quantum simulations of scattering in scalar field theory on 120 qubits. *Physical Review D*, 112:034502, 2025.
- [239] H. Yu, Y. Zhao, and T. Wei. Simulating large-size quantum spin chains on cloud-based superconducting quantum computers. *Physical Review Research*, 5(1):013183, 2023.
- [240] O. Shtanko et al. Uncovering local integrability in quantum many-body dynamics. *Nature Communications*, 16(1):2552, 2025.
- [241] E. Baumer et al. Efficient long-range entanglement using dynamic circuits. *PRX Quantum*, 5(3):030339, 2024.

- [242] E. Chen et al. Nishimori transition across the error threshold for constant-depth quantum circuits. *Nature Physics*, 21(1):161–167, 2025.
- [243] R. Farrell et al. Digital quantum simulations of scattering in quantum field theories using W states. May 2025. arXiv:2505.03111 [quant-ph].
- [244] Z. Davoudi, C. Hsieh, and S. Kadam. Scattering wave packets of hadrons in gauge theories: Preparation on a quantum computer. *Quantum*, 8:1520, 2024.
- [245] Y. Chai et al. Fermionic wave packet scattering: a quantum computing approach. *Quantum*, 9:1638, 2025.
- [246] J. Schuhmacher et al. Observation of hadron scattering in a lattice gauge theory on a quantum computer. May 2025. arXiv:2505.20387 [quant-ph].
- [247] C. Bauer. Efficient use of quantum computers for collider physics. *Journal of High Energy Physics*, 2025(11):108, 2025.
- [248] M. Illa and M. Savage. Basic elements for simulations of standard-model physics with quantum annealers: Multigrid and clock states. *Physical Review A*, 106(5), 2022.
- [249] R. Briceño et al. Toward coherent quantum computation of scattering amplitudes with a measurement-based photonic quantum processor. *Physical Review Research*, 6(4):043065, 2024.
- [250] M. Illa and M. Savage. Multi-neutrino entanglement and correlations in dense neutrino systems. *Physical Review Letters*, 130(22):221003, 2023.
- [251] V. Amitrano et al. Trapped-ion quantum simulation of collective neutrino oscillations. *Physical Review D*, 107(2):023007, 2023.
- [252] A. Baroni et al. Nuclear two point correlation functions on a quantum computer. *Physical Review D*, 105(7):074503, 2022.
- [253] I. Chernyshev et al. Pathfinding quantum simulations of neutrinoless double- β decay. *Nature Communications*, 17(1):1826, 2026.
- [254] P. Siwach, K. Harrison, and A. Balantekin. Collective neutrino oscillations on a quantum computer with hybrid quantum-classical algorithm. *Physical Review D*, 108(8):083039, 2023.
- [255] A. Kavaki and R. Lewis. From square plaquettes to triamond lattices for SU(2) gauge theory. *Communications Physics*, 7(1):208, 2024.
- [256] C. Alexandrou et al. Realizing string breaking dynamics in a Z2 lattice gauge theory on quantum hardware. *Physical Review D*, 112(11):114506, 2025.

- [257] A. Ciavarella. String breaking in the heavy quark limit with scalable circuits. *Physical Review D*, 111(5):054501, 2025.
- [258] R. Farrell, M. Illa, and M. Savage. Steps toward quantum simulations of hadronization and energy loss in dense matter. *Physical Review C*, 111(1):015202, 2025.
- [259] Z. Li, M. Illa, and M. Savage. A framework for quantum simulations of energy-loss and hadronization in non-Abelian gauge theories: SU(2) lattice gauge theory in 1+1D. December 2025. arXiv:2512.05210 [quant-ph].
- [260] J. Barata and E. Rico. Real-time simulation of jet energy loss and entropy production in high-energy scattering with matter. *Communications Physics*, 9(1):155, 2026.
- [261] M. Mathur. Harmonic oscillator prepotentials in SU(2) lattice gauge theory. *J. Phys. A*, 38:10015–10026, 2005.
- [262] M. Mathur. Loop states in lattice gauge theory. *Physics Letters B*, 640(5-6):292–296, 2005.
- [263] M. Mathur. Loop approach to lattice gauge theories. *Nuclear Physics B*, 779(1-2):32–62, 2007.
- [264] S. Kadam, I. Raychowdhury, and J. Stryker. Loop-string-hadron formulation of an SU(3) gauge theory with dynamical quarks. *Proceedings of The 39th International Symposium on Lattice Field Theory — PoS(LATTICE2022)*, 107(9):094513, 2022.
- [265] S. Kadam et al. Loop-string-hadron approach to su(3) lattice yang-mills theory: Hilbert space of a trivalent vertex. *Phys. Rev. D*, 111(7):074516, 2025.
- [266] T. Byrnes and Y. Yamamoto. Simulating lattice gauge theories on a quantum computer. *Physical Review A*, 73:022328, 2006.
- [267] M. Bañuls et al. Efficient basis formulation for 1+1 dimensional SU(2) lattice gauge theory: Spectral calculations with matrix product states. *Physical Review X*, 7(4):041046, 2017.
- [268] Z. Davoudi, A. Shaw, and J. Stryker. General quantum algorithms for Hamiltonian simulation with applications to a non-Abelian lattice gauge theory. *Quantum*, 7:1213, 2023.
- [269] P. Orland. $(2 + 1)$ -dimensional lattice QCD. *Physical Review D*, 71:054503, 2005.
- [270] S. Romiti and C. Urbach. Digitizing lattice gauge theories in the magnetic basis: reducing the breaking of the fundamental commutation relations. *The European Physical Journal C*, 84(7):708, 2023.
- [271] M. Alam et al. Primitive quantum gates for dihedral gauge theories. *Physical Review D*, 105(11):114501, 2022.

- [272] A. Alexandru et al. Spectrum of digitized QCD: Glueballs in a $S(1080)$ gauge theory. *Physical Review D*, 105(11):114508, 2022.
- [273] E. Gustafson et al. Primitive quantum gates for an $SU(2)$ discrete subgroup: Binary tetrahedral. *Physical Review D*, 106(11):114501, 2022.
- [274] E. Gustafson, H. Lamm, and F. Lovelace. Primitive quantum gates for an $SU(2)$ discrete subgroup: Binary octahedral. *Physical Review D*, 109(5):054503, 2024.
- [275] E. Gustafson et al. Primitive quantum gates for an $SU(3)$ discrete subgroup: $\Sigma(36 \times 3)$. *Physical Review D*, 110(3):034515, 2024.
- [276] B. Assi and H. Lamm. Digitization and subduction of $SU(N)$ gauge theories. *Physical Review D*, 110(7):074511, 2024.
- [277] H. Lamm et al. Block encodings of discrete subgroups on a quantum computer. *Physical Review D*, 110(5):054505, 2024.
- [278] E. Muarari et al. Highly-efficient quantum Fourier transformations for some nonabelian groups. *Physical Review D*, 110(7), 2024.
- [279] J. Bender and E. Zohar. Gauge redundancy-free formulation of compact QED with dynamical matter for quantum and classical computations. *Physical Review D*, 102(11):114517, 2020.
- [280] M. Miranda-Riaza, P. Fontana, and A. Celi. Renormalized dual basis for scalable simulations of 2+1D compact quantum electrodynamics. *Physical Review D*, 113(7):074518, 2026.
- [281] J. Halimeh et al. A universal framework for the quantum simulation of Yang–Mills theory. *Commun. Phys.*, 9(1):67, 2026.
- [282] G. Bergner et al. Toward QCD on quantum computer: orbifold lattice approach. *Journal of High Energy Physics*, 2024(5):234, 2024.
- [283] M. Creutz. Gauge fixing, the transfer matrix, and confinement on a lattice. *Physical Review D*, 15(4):1128–1136, 1977.
- [284] D. Varshalovich, A. Moskalev, and V. Khersonskii. *Quantum Theory of Angular Momentum*. World Scientific Publishing Company, Singapore, 1988.
- [285] A. Macridin et al. Electron-phonon systems on a universal quantum computer. *Physical Review Letters*, 121(11):110504, 2018.
- [286] L. Trefethen. *Approximation Theory and Approximation Practice*. Society for Industrial and Applied Mathematics, Philadelphia, PA, extended edition edition, 2013.

- [287] Z. Li, D. Grabowska, and M. Savage. Sequency hierarchy truncation (SeqHT) for adiabatic state preparation and time evolution in quantum simulations. *Quantum*, 9:1865, 2025.
- [288] A. Klappenecker and M. Roetteler. Discrete cosine transforms on quantum computers. In *ISPA 2001. Proceedings of the 2nd International Symposium on Image and Signal Processing and Analysis. In conjunction with 23rd International Conference on Information Technology Interfaces (IEEE Cat.*, pages 464–468. IEEE, Univ. Zagreb, 2001.
- [289] Y. Sato et al. Hamiltonian simulation for hyperbolic partial differential equations by scalable quantum circuits. *Physical Review Research*, 6(3):033246, 2024.
- [290] L. Clinton, J. Bausch, and T. Cubitt. Hamiltonian simulation algorithms for near-term quantum hardware. *Nature Communications*, 12(1):4989, 2021.
- [291] BQSKit Development Team. Berkeley Quantum Synthesis Toolkit (BQSKit). <https://bqskit.lbl.gov/>, 2025. Accessed: November 9, 2025.
- [292] L. Viola and S. Lloyd. Dynamical suppression of decoherence in two state quantum systems. *Physical Review A*, 58(4):2733–2744, 1998.
- [293] E. van den Berg, Z. Mineev, and K. Temme. Model-free readout-error mitigation for quantum expectation values. *Physical Review A*, 105(3):032620, 2022.
- [294] M. Urbánek et al. Mitigating depolarizing noise on quantum computers with noise-estimation circuits. *Physical Review Letters*, 127(27):270502, 2021.
- [295] A. Meglio et al. Quantum computing for high-energy physics: State of the art and challenges. *PRX Quantum*, 5(3):037001, 2024.
- [296] M. Bañuls et al. Simulating lattice gauge theories within quantum technologies. *European Physical Journal D : Atomic, Molecular and Optical Physics*, 74(8):1–42, 2019.
- [297] T. Zache et al. Quantum simulation of lattice gauge theories using Wilson fermions. *Quantum Science and Technology*, 3(3):034010, 2018.
- [298] A. Buser et al. Quantum simulation of gauge theory via orbifold lattice. *JHEP*, 09:034, 2021.
- [299] G. Vidal. Efficient classical simulation of slightly entangled quantum computations. *Physical Review Letters*, 91(14):147902, 2003.
- [300] L. Amico et al. Entanglement in many-body systems. *Reviews of Modern Physics*, 80(2):517–576, 2008.

- [301] J. Eisert, M. Cramer, and M. Plenio. Area laws for the entanglement entropy - a review. *Reviews of Modern Physics*, 82:277–306, 2010.
- [302] T. Rakovszky et al. Signatures of information scrambling in the dynamics of the entanglement spectrum. *Physical Review B*, 100(12):125115, 2019.
- [303] D. Gottesman. The Heisenberg representation of quantum computers. In *22nd International Colloquium on Group Theoretical Methods in Physics*, pages 32–43, 1998.
- [304] M. Howard and E. Campbell. Application of a resource theory for magic states to fault-tolerant quantum computing. *Physical Review Letters*, 118(9):090501, 2017.
- [305] L. Leone, S. Oliviero, and A. Hama. Stabilizer rényi entropy. *Physical Review Letters*, 128(5):050402, 2022.
- [306] S. Oliviero et al. Measuring magic on a quantum processor. *npj Quantum Information*, 8(1):148, 2022.
- [307] A. Gu, S. Oliviero, and L. Leone. Magic-induced computational separation in entanglement theory. *PRX Quantum*, 6(2):020324, 2025.
- [308] E. Tirrito et al. Quantifying nonstabilizerness through entanglement spectrum flatness. *Physical Review A*, 109(4):L040401, 2024.
- [309] L. Leone et al. Quantum chaos is quantum. *Quantum*, 5:453, 2021.
- [310] E. Tirrito, X. Turkeshi, and P. Sierant. Anticoncentration and nonstabilizerness spreading under ergodic quantum dynamics. *Physical Review Letters*, 135(22):220401, 2025.
- [311] X. Turkeshi, E. Tirrito, and P. Sierant. Magic spreading in random quantum circuits. *Nature Communications*, 16(1):2575, 2025.
- [312] C. Robin and M. Savage. Quantum complexity and new directions in nuclear physics and high-energy physics phenomenology. April 2026. arXiv:2604.26376 [quant-ph].
- [313] C. Robin and M. Savage. Quantum complexity fluctuations from nuclear and hypernuclear forces. *Physical Review C*, 112(4):044004, 2025.
- [314] C. Cao et al. Gravitational backreaction is magical. *PRX Quantum*, 6(4):040375, 2025.
- [315] S. Grieninger, M. Savage, and N. Zemlevskiy. The quantum complexity of string breaking in the Schwinger model. January 2026. arXiv:2601.08825 [hep-ph].
- [316] S. Grieninger. Nonlocal nonstabilizerness from holographic Schwinger pair production. *Physical Review D*, 114(3):L031502, 2026. arXiv:2605.04210 [hep-th].

- [317] H. Lamm, S. Lawrence, and Y. Yamauchi. General methods for digital quantum simulation of gauge theories. *Physical Review D*, 100(3):034518, 2019.
- [318] J. Welch et al. Efficient quantum circuits for diagonal unitaries without ancillas. *New Journal of Physics*, 16(3):033040, 2014.
- [319] J. Zylberman and F. Debbasch. Efficient quantum state preparation with walsh series. *Physical Review A*, 109(4):042401, 2024.
- [320] N. Klco and M. Savage. Hierarchical qubit maps and hierarchically implemented quantum error correction. *Physical Review A*, 104(6):062425, 2021.
- [321] P. Rebentrost, B. Gupt, and T. Bromley. Quantum computational finance: Monte carlo pricing of financial derivatives. *Physical Review A*, 98(2):022321, 2018.
- [322] J. Gonzalez-Conde et al. Efficient quantum amplitude encoding of polynomial functions. *Quantum*, 8:1297, 2024.
- [323] F. Alam et al. Onset of ergodicity across scales on a digital quantum processor. March 2026. arXiv:2603.12236 [quant-ph].
- [324] S. Aaronson and D. Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004.
- [325] L. Leone and L. Bittel. Stabilizer entropies are monotones for magic-state resource theory. *Physical Review A*, 110(4):L040403, 2024.
- [326] F. Trigueros and J. Guzmán. Nonstabilizerness and error resilience in noisy quantum circuits. *Phys. Rev. Lett.*, 136(24):240602, 2026.
- [327] G. Santra et al. Quantum resources in non-Abelian lattice gauge theories: Nonstabilizerness, multipartite entanglement, and fermionic non-Gaussianity. October 2025. arXiv:2510.07385 [quant-ph].
- [328] X. Huang et al. A fast and exact approach for stabilizer Rényi entropy via the XOR-FWHT algorithm. December 2025. arXiv:2512.24685 [quant-ph].
- [329] Z. Xiao and S. Ryu. Exponentially accelerated sampling of Pauli strings for nonstabilizerness. January 2026. arXiv:2601.00761 [quant-ph].
- [330] P. Sierant, J. Vallès-Muns, and A. García-Sáez. Computing quantum magic of state vectors. *Quantum*, 10:2059, 2026.
- [331] P. Nation et al. Scalable mitigation of measurement errors on quantum computers. *PRX Quantum*, 2(4):040326, 2021.

- [332] M. Lee and R. Farrell. Studying energy-resolved transport with wavepacket dynamics on quantum computers. January 2026. arXiv:2601.16180 [quant-ph].
- [333] K. Chandramouli et al. Statistical signal processing for quantum error mitigation. In *International Conference on Quantum Computing and Engineering*, pages 269–275. IEEE, 2025.
- [334] D. Baron, H. Patil, and H. Zhou. Qubit-wise majority vote: maximum likelihood quantum error mitigation for algorithms with a single correct output. In *International Conference on Quantum Computing and Engineering*, pages 124–133. IEEE, 2024.
- [335] K. Liu and Z. Cai. Quantum error mitigation for sampling algorithms. February 2025. arXiv:2502.11285 [quant-ph].
- [336] T. Schuster et al. A Polynomial-Time Classical Algorithm for Noisy Quantum Circuits. *Physical Review X*, 15(4):041018, 2025.
- [337] D. Mark et al. Maximum entropy principle in deep thermalization and in Hilbert-space ergodicity. *Physical Review X*, 14(4):041051, 2024.
- [338] K. Xu et al. Observation of glueball excitations and string breaking in a $2 + 1\text{D } \mathbb{Z}_2$ lattice gauge theory on a trapped-ion quantum computer. April 2026. arXiv:2604.07435 [hep-lat].
- [339] J. Cao et al. String breaking and glueball dynamics in $2 + 1\text{D}$ quantum link electrodynamics. January 2026. arXiv:2601.16166 [hep-lat].
- [340] R. Jha et al. Magic and entanglement in $1+1$ -dimensional $\text{SU}(2)$ lattice gauge theory. June 2026. arXiv:2606.09971 [quant-ph].
- [341] J. Eisert, M. Friesdorf, and C. Gogolin. Quantum many-body systems out of equilibrium. *Nature Physics*, 11:124, 2015.
- [342] S. Aaronson. Shadow tomography of quantum states. In *Electron. Colloquium Comput. Complex.*, pages 325–338, 2017.
- [343] J. Emerson, R. Alicki, and K. Życzkowski. Scalable noise estimation with random unitary operators. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S347–S352, 2005.
- [344] E. Knill et al. Randomized benchmarking of quantum gates. *Physical Review A*, 77(1):012307, 2008.
- [345] M. Cramer et al. Efficient quantum state tomography. *Nat Commun.*, 1(1):149, 2010.
- [346] E. Magesan, J. Gambetta, and J. Emerson. Scalable and robust randomized benchmarking of quantum processes. *Physical Review Letters*, 106:180504, 2011.

- [347] S. Boixo et al. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, 2018.
- [348] A. Erhard et al. Characterizing large-scale quantum computers via cycle benchmarking. *Nature Communications*, 10(1):5347, 2019.
- [349] R. Blume-Kohout, T. Proctor, and K. Young. Quantum characterization, verification, and validation. March 2025. arXiv:2503.16383 [quant-ph].
- [350] A. Czerwinski. Selected concepts of quantum state tomography. *Optics*, 3(3):268–286, 2022.
- [351] D. Gross et al. Quantum state tomography via compressed sensing. *Physical Review Letters*, 105:150401, 2010.
- [352] S. Flammia et al. Quantum tomography via compressed sensing: error bounds, sample complexity and efficient estimators. *New Journal of Physics*, 14(9):095022, 2012.
- [353] J. Haah et al. Sample-optimal tomography of quantum states. In *IEEE Transactions on Information Theory*, pages 913–925, 2015.
- [354] R. O’Donnell and J. Wright. Efficient quantum tomography. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 899–912. ACM, 2016.
- [355] G. Torlai et al. Neural-network quantum state tomography. *Nature Physics*, 14(5):447–450, 2018.
- [356] A. Saxena et al. Boundary measurement tomography of the bose hubbard model on general graphs. *Physical Review Research*, 6(3):033058, 2024.
- [357] S. van Enk and C. Beenakker. Measuring $\text{Tr} \rho^n$ on single copies of ρ using random measurements. *Physical Review Letters*, 108(11):110503, 2012.
- [358] A. Elben et al. Rényi entropies from random quenches in atomic hubbard and spin models. *Physical Review Letters*, 120(5):050406, 2018.
- [359] B. Vermersch et al. Unitary n -designs via random quenches in atomic Hubbard and spin models: Application to the measurement of Rényi entropies. *Physical Review A*, 97(2):023604, 2018.
- [360] A. Elben et al. Statistical correlations between locally randomized measurements: A toolbox for probing entanglement in many-body quantum states. *Physical Review A*, 99(5):052323, 2019.
- [361] T. Brydges et al. Probing rényi entanglement entropy via randomized measurements. *Science*, 364(6437):260–263, 2019.

- [362] Z. Du et al. Optimal randomized measurements for a family of nonlinear quantum properties. *PRX Quantum*, 7(1):010360, 2026.
- [363] M. Painsi and A. Kalev. An approximate description of quantum states. October 2019. arXiv:1910.10543 [quant-ph].
- [364] H. Huang et al. Demonstrating quantum advantage in learning from experiments. *Science*, 376(6598):1182–1186, 2022.
- [365] K. Van Kirk et al. Hardware-efficient learning of quantum many-body states. *Physical Review Letters*, 2025. arXiv:2212.06084 [quant-ph].
- [366] C. Hadfield et al. Measurements of quantum hamiltonians with locally-biased classical shadows. *Communications in Mathematical Physics*, 391(3):951–967, 2020.
- [367] A. Akhtar, H. Hu, and Y. You. Scalable and flexible classical shadow tomography with tensor networks. *Quantum*, 7:1026, 2023.
- [368] H. Hu, S. Choi, and Y. You. Classical shadow tomography with locally scrambled quantum dynamics. *Physical Review Res.*, 5(2):023027, 2023.
- [369] J. Kunjummen et al. Shadow process tomography of quantum channels. *Physical Review A*, 107(4):042403, 2023.
- [370] S. Gandhari et al. Precision bounds on continuous-variable state tomography using classical shadows. *PRX Quantum*, 5(1):010346, 2024.
- [371] R. Levy, D. Luo, and B. Clark. Classical shadows for quantum process tomography on near-term quantum computers. *Physical Review Res.*, 6(1):013029, 2024.
- [372] Z. Zhu, J. Lukens, and B. Kirby. On the connection between least squares, regularization, and classical shadows. *Quantum*, 8:1455, 2024.
- [373] Z. Cai et al. Biased estimator channels for classical shadows. *Physical Review A*, 111(3):L030402, 2025.
- [374] H. Hu et al. Demonstration of robust and efficient quantum property learning with shallow shadows. *Nature Communications*, 16(1):2943, 2025.
- [375] L. Joshi et al. Observing the quantum Mpemba effect in quantum simulations. *Physical Review Letters*, 133(1), 2024.
- [376] T. Zhang et al. Experimental quantum state measurement with classical shadows. *Physical Review Letters*, 127(20):200501, 2021.

- [377] G. Struchalin et al. Experimental estimation of quantum state properties from classical shadows. *PRX Quantum*, 2(1):010307, 2021.
- [378] V. Vitale et al. Robust estimation of the quantum Fisher information on a quantum processor. *PRX Quantum*, 5(3):030338, 2024.
- [379] H. Dong et al. Measuring the spectral form factor in many-body chaotic and localized phases of quantum processors. *Physical Review Letters*, 134:010402, 2025.
- [380] T. Andersen et al. Thermalization and criticality on an analogue–digital quantum simulator. *Nature*, 638(8049):79–85, 2025.
- [381] M. Votto et al. Learning mixed quantum states in large-scale experiments. *Physical Review Letters*, 136(9):090801, 2026.
- [382] A. Elben et al. The randomized measurement toolbox. *Nature Reviews Physics*, 5(1):9–24, 2023.
- [383] A. Rath et al. Importance sampling of randomized measurements for probing entanglement. *Physical Review Letters*, 127(20):200503, 2021.
- [384] H. Huang, R. Kueng, and J. Preskill. Efficient estimation of pauli observables by derandomization. *Physical Review Letters*, 127(3):030503, 2021.
- [385] G. Low. Classical shadows of fermions with particle number symmetry. August 2022. arXiv:2208.08964 [quant-ph].
- [386] J. Bringewatt, J. Kunjummen, and N. Mueller. Randomized measurement protocols for lattice gauge theories. *Quantum*, 8:1300, 2024.
- [387] S. Hearth et al. Efficient local classical shadow tomography with number conservation. *Physical Review Letters*, 133(6):060802, 2024.
- [388] F. Sauvage and M. Larocca. Classical shadows with symmetries. August 2024. arXiv:2408.05279 [quant-ph].
- [389] R. Arvind et al. Quantum tug of war between randomness and symmetries on homogeneous spaces. *Physical Review Res.*, 7(1):013105, 2025.
- [390] N. Klco, A. Roggero, and M. Savage. Standard model physics and the digital quantum revolution: thoughts about the interface. *Rept. Prog. Phys.*, 85(6):064301, 2022.
- [391] E. Fradkin. *Field theories of condensed matter physics*. Cambridge University Press, 2013.

- [392] H. Kleinert. *Gauge Fields in Condensed Matter: Vol. 1: Superflow and Vortex Lines (Disorder Fields, Phase Transitions) Vol. 2: Stresses and Defects (Differential Geometry, Crystal Melting)*. World Scientific, 1989.
- [393] X. Wen. Topological orders in rigid states. *International Journal of Modern Physics B*, 04(02):239–271, 1990.
- [394] M. Levin and X. Wen. String net condensation: A Physical mechanism for topological phases. *Physical Review B*, 71:045110, 2005.
- [395] Y. Chen, A. Kapustin, and D. Radicevic. Exact bosonization in two spatial dimensions and a new class of lattice gauge theories. *Annals of Physics*, 393:234–253, 2017.
- [396] Y. Chen. Exact bosonization in arbitrary dimensions. *Physical Review Research*, 2(3):033527, 2020.
- [397] Y. Chen and Y. Xu. Equivalence between fermion-to-qubit mappings in two spatial dimensions. *PRX Quantum*, 4(1):010326, 2023.
- [398] M. Freedman et al. Topological Quantum Computation. *Bull. Am. Math. Soc.*, 40(1):31–38, 2003. arXiv:quant-ph/0101025 [quant-ph].
- [399] C. Nayak et al. Non-abelian anyons and topological quantum computation. *Reviews of Modern Physics*, 80:1083–1159, 2008.
- [400] V. Lahtinen and J. Pachos. A Short Introduction to Topological Quantum Computation. *SciPost Physics*, 3(3):021, 2017.
- [401] J. Kogut. An introduction to lattice gauge theory and spin systems. *Reviews of Modern Physics*, 51(4):659–713, 1979.
- [402] Z. Davoudi. TASI/CERN/KITP lecture notes on “Toward quantum computing gauge theories of nature”. July 2025. arXiv:2507.15840 [hep-lat].
- [403] F. Wegner. Duality in generalized ising models and phase transitions without local order parameters. *Journal of Mathematical Physics*, 12(10):2259–2272, 1971.
- [404] D. Horn, M. Weinstein, and S. Yankielowicz. Hamiltonian approach to $\mathbb{Z}_n$ lattice gauge theories. *Physical Review D*, 19(12):3715, 1979.
- [405] N. Mueller, T. Zache, and R. Ott. Thermalization of gauge theories from their entanglement spectrum. *Physical Review Letters*, 129:011601, 2022.
- [406] J. Hartse, L. Fidkowski, and N. Mueller. Stabilizer scars. *Physical Review Letters*, 135(6):060402, 2025.

- [407] P. Black. Dictionary of algorithms and data structures. <https://xlinux.nist.gov/dads/>, 1998.
- [408] A. Mele. Introduction to haar measure tools in quantum information: A beginner’s tutorial. *Quantum*, 8:1340, 2024.
- [409] A. Smith et al. Disorder-free localization. *Physical Review Letters*, 118(26):266601, 2017.
- [410] J. Halimeh et al. Disorder-free localization for benchmarking quantum computers. October 2024. arXiv:2410.08268 [quant-ph].
- [411] H. Froland. LGTShallowShadows. <https://github.com/frolandh/LGTShallowShadows>, 2025.
- [412] V. Dotsenko. Duality transformations for discrete abelian models. simple example of duality transformation for non-abelian model. *Soviet Journal of Experimental and Theoretical Physics*, 48:546, 1978.
- [413] J. Drouffe, C. Itzykson, and J. Zuber. Lattice models with a solvable symmetry group. *Nuclear Physics B*, 147(1-2):132–134, 1979.
- [414] E. Álvarez, L. Alvarez-Gaume, and Y. Lozano. On non-abelian duality. *Nuclear Physics B*, 424(1):155–183, 1994.
- [415] V. Buchstaber and M. Monastyrsky. Generalized kramers–wannier duality for spin systems with non-commutative symmetry. *Journal of Physics A: Mathematical and General*, 36(28):7679–7692, 2003.
- [416] E. Cobanera, G. Ortiz, and E. Knill. A solution to the non-abelian duality problem. *Nucl. Phys. B*, 877(2):574–597, 2012.
- [417] W. Ji and X. Wen. Categorical symmetry and noninvertible anomaly in symmetry-breaking and topological phase transitions. *Physical Review Research*, 2(3):033417, 2020.
- [418] D. Aasen, P. Fendley, and R. Mong. Topological defects on the lattice: Dualities and degeneracies. August 2020. arXiv:2008.08598 [cond-mat.stat-mech].
- [419] A. Chakraborty, R. Lewis, and C. Muschik. Charge-singlet measurement toolbox. *Physical Review Res.*, 7(4):043162, 2025.
- [420] Y. Zhan et al. Learning conservation laws in unknown quantum dynamics. *PRX Quantum*, 5(1):010350, 2024.
- [421] Z. Ji, Y. Liu, and F. Song. Pseudorandom quantum states. In *Annual International Cryptology Conference*, pages 126–152. Springer, 2018.

- [422] F. Ma and H. Huang. How to construct random unitaries. In *IACR Cryptology ePrint Archive*, pages 806–809, 2024.
- [423] T. Schuster, J. Haferkamp, and H. Huang. Random unitaries in extremely low depth. *Science*, 389(6755):adv8590, 2025.
- [424] L. Grevink et al. Will it glue? On short-depth designs beyond the unitary group. June 2025. arXiv:2506.23925 [quant-ph].
- [425] L. Cui et al. Random unitaries from Hamiltonian dynamics. October 2025. arXiv:2510.08434 [quant-ph].
- [426] M. Ippoliti. Classical shadows based on locally-entangled measurements. *Quantum*, 8:1293, 2024.
- [427] S. Chen et al. Robust shadow estimation. *PRX Quantum*, 2(3):030348, 2021.
- [428] D. Koh and S. Grewal. Classical shadows with noise. *Quantum*, 6:776, 2022.
- [429] P. Rozon, N. Bao, and K. Agarwal. Optimal twirling depth for classical shadows in the presence of noise. *Physical Review Letters*, 133(13):130803, 2024.
- [430] R. Brieger et al. Stability of classical shadows under gate-dependent noise. *Physical Review Letters*, 134(9):090801, 2025.
- [431] H. Hu et al. Logical shadow tomography: Efficient estimation of error-mitigated observables. March 2022. arXiv:2203.07263 [quant-ph].
- [432] C. Gross and I. Bloch. Quantum simulations with ultracold atoms in optical lattices. *Science*, 357(6355):995–1001, 2017.
- [433] E. Altman et al. Quantum simulators: Architectures and opportunities. *PRX Quantum*, 2(1):017003, 2021.
- [434] D. Beck et al. Quantum Information Science and Technology for Nuclear Physics. Input into U.S. Long-Range Planning. February 2023. arXiv:2303.00113 [quant-ph].
- [435] M. Dalmonte, B. Vermersch, and P. Zoller. Quantum simulation and spectroscopy of entanglement hamiltonians. *Nature Physics*, 14(8):827–831, 2018.
- [436] E. Bairey, I. Arad, and N. Lindner. Learning a local hamiltonian from local measurements. *Physical Review Letters*, 122(2):020504, 2019.
- [437] A. Anshu et al. Sample-efficient learning of interacting quantum systems. *Nature Physics*, 17(8):931–935, 2021.

- [438] C. Kokail et al. Quantum variational learning of the entanglement hamiltonian. *Physical Review Letters*, 127(17):170501, 2021.
- [439] M. Joshi et al. Exploring large-scale entanglement in quantum simulation. *Nature*, 624(7992):539–544, 2023.
- [440] A. Daley et al. Measuring entanglement growth in quench dynamics of bosons in an optical lattice. *Physical Review Letters*, 109(2):020505, 2012.
- [441] H. Pichler et al. Thermal versus entanglement entropy: a measurement protocol for fermionic atoms with a quantum gas microscope. *New Journal of Physics*, 15(6):063003, 2013.
- [442] H. Pichler et al. Measurement protocol for the entanglement spectrum of cold atoms. *Physical Review X*, 6(4):041033, 2016.
- [443] R. Islam et al. Measuring entanglement entropy in a quantum many-body system. *Nature*, 528(7580):77–83, 2015.
- [444] T. Schweigler et al. Experimental characterization of a quantum many-body system via higher-order correlations. *Nature*, 545(7654):323–326, 2017.
- [445] M. Rispoli et al. Quantum critical behaviour at the many-body localization transition. *Nature*, 573(7774):385–389, 2019.
- [446] H. Li and F. Haldane. Entanglement spectrum as a generalization of entanglement entropy: Identification of topological order in non-abelian fractional quantum hall effect states. *Physical Review Letters*, 101(1):010504, 2008.
- [447] M. Dalmonte et al. Entanglement hamiltonians: from field theory to lattice models and experiments. *Annals of Physics*, 534(11):2200064, 2022.
- [448] V. Oganesyan and D. Huse. Localization of interacting fermions at high temperature. *Physical Review B*, 75(15):155111, 2007.
- [449] P. Chang et al. Evolution of entanglement spectra under generic quantum dynamics. *Physical Review Letters*, 123(19):190602, 2019.
- [450] I. Bloch, J. Dalibard, and W. Zwerger. Many-body physics with ultracold gases. *Reviews of Modern Physics*, 80(3):885, 2008.
- [451] A. Mazurenko et al. A cold-atom fermi–hubbard antiferromagnet. *Nature*, 545(7655):462–466, 2017.
- [452] L. Tarruell and L. Sanchez-Palencia. Quantum simulation of the hubbard model with ultracold fermions in optical lattices. *Comptes Rendus Physique*, 19(6):365–393, 2018.

- [453] D. Gonzalez-Cuadra et al. Fermionic quantum processing with programmable neutral atom arrays. *Proceedings of the National Academy of Sciences of the United States of America*, 120(35):e2304294120, 2023.
- [454] A. Impertro et al. Local readout and control of current and kinetic energy operators in optical lattices. *Physical Review Letters*, 133(6):063401, 2024.
- [455] P. Lunt et al. Realization of a Laughlin state of two rapidly rotating fermions. *Physical Review Letters*, 133(25):253401, 2024.
- [456] M. Lebrat et al. Observation of nagaoka polarons in a fermi–hubbard quantum simulator. *Nature*, 629(8011):317–322, 2024.
- [457] M. Prichard et al. Directly imaging spin polarons in a kinetically frustrated hubbard system. *Nature*, 629(8011):323–328, 2024.
- [458] T. Hartke et al. Direct observation of nonlocal fermion pairing in an attractive fermi-hubbard gas. *Science*, 381(6653):82–86, 2023.
- [459] M. Chung and I. Peschel. Density-matrix spectra of solvable fermionic systems. *Physical Review B*, 64(6):064412, 2001.
- [460] T. Guhr, A. Mueller-Groeling, and H. Weidenmueller. Random-matrix theories in quantum physics: common concepts. *Physics Reports*, 299(4-6):189–425, 1997.
- [461] M. Mehta. *Random Matrices*, volume 142. Elsevier, Amsterdam, 3rd edition, 2004.
- [462] K. Cahill and R. Glauber. Density operators for fermions. *Physical Review A*, 59(2):1538, 1999.
- [463] A. Acharya, S. Saha, and A. Sengupta. Informationally complete POVM-based shadow tomography. May 2021. arXiv:2105.05992 [quant-ph].
- [464] F. Biebl and S. Kehrein. Thermalization rates in the one-dimensional hubbard model with next-to-nearest neighbor hopping. *Physical Review B*, 95(10):104304, 2017.
- [465] P. Weinberg and M. Bukov. QuSpin: a Python Package for Dynamics and Exact Diagonalisation of Quantum Many Body Systems. Part II: bosons, fermions and higher spins. *SciPost Physics*, 2018.
- [466] R. Jozsa. Fidelity for mixed quantum states. *Journal of Modern Optics*, 41(12):2315–2323, 1994.
- [467] M. Stark and M. Kollar. Kinetic description of thermalization dynamics in weakly interacting quantum systems. August 2013. arXiv:1308.1610 [cond-mat.str-el].
- [468] F. Essler et al. *The One-Dimensional Hubbard Model*. Cambridge University Press, Cambridge, 2005.

- [469] J. Marco et al. Level statistics of the one-dimensional ionic hubbard model. *Physical Review Research*, 4(3):033119, 2022.
- [470] I. Bloch, J. Dalibard, and S. Nascimbene. Quantum simulations with ultracold quantum gases. *Nature Physics*, 8(4):267–276, 2012.
- [471] J. Léonard et al. Realization of a fractional quantum hall state with ultracold atoms. *Nature*, 619(7970):495–499, 2023.
- [472] D. Bluvstein et al. A quantum processor based on coherent transport of entangled atom arrays. *Nature*, 604(7906):451–456, 2022.
- [473] P. Naldesi et al. Fermionic correlation functions from randomized measurements in programmable atomic quantum devices. *Physical Review Letters*, 131(6):060601, 2023.
- [474] M. Gluza and J. Eisert. Recovering quantum correlations in optical lattices from interaction quenches. *Physical Review Letters*, 127(9):090503, 2021.
- [475] A. Zhao, N. Rubin, and A. Miyake. Fermionic partial tomography via classical shadows. *Physical Review Letters*, 127(11):110504, 2021.
- [476] J. Denzler et al. Learning fermionic correlations by evolving with random translationally invariant hamiltonians. *Physical Review Letters*, 133(24):240604, 2024.
- [477] M. Tran et al. Measuring arbitrary physical properties in analog quantum simulation. *Physical Review X*, 13(1):011049, 2023.
- [478] A. Gu et al. Simulating quantum chaos without chaos. October 2024. arXiv:2410.18196 [quant-ph].
- [479] L. Bittel et al. Optimal trace-distance bounds for free-fermionic states: Testing and improved tomography. *PRX Quantum*, 6(3):030341, 2025.
- [480] D. Szász-Schagrin, B. Pozsgay, and G. Takács. Weak integrability breaking and level spacing distribution. *SciPost Physics*, 11(2):037, 2021.
- [481] J. Deutsch. Quantum statistical mechanics in a closed system. *Physical Review A*, 43(4):2046–2049, 1991.
- [482] M. Srednicki. Chaos and quantum thermalization. *Physical Review E*, 50(2):888–901, 1994.
- [483] M. Rigol, V. Dunjko, and M. Olshanii. Thermalization and its mechanism for generic isolated quantum systems. *Nature*, 452(7189):854–858, 2008.

- [484] X. Yao. SU(2) gauge theory in 2+ 1 dimensions on a plaquette chain obeys the eigenstate thermalization hypothesis. *Physical Review D*, 108(3):L031504, 2023.
- [485] L. Ebner et al. Eigenstate thermalization in (2+ 1)-dimensional SU (2) lattice gauge theory. *Physical Review D*, 109(1):014504, 2024.
- [486] R. Vogt. *Ultrarelativistic heavy-ion collisions*. Elsevier, 2007.
- [487] W. Florkowski. *Phenomenology of ultra-relativistic heavy-ion collisions*. World Scientific, 2010.
- [488] P. Arnold, G. Moore, and L. Yaffe. Effective kinetic theory for high temperature gauge theories. *Journal of High Energy Physics*, 2003(01):030, 2002.
- [489] C. Shen and L. Yan. Recent development of hydrodynamic modeling in heavy-ion collisions. *Nuclear Science and Techniques*, 31(12):122, 2020.
- [490] U. Schollwoeck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, 2010.
- [491] P. Kramer and M. Saraceno. *Geometry of the time-dependent variational principle in quantum mechanics*. Springer Berlin Heidelberg, 1981.
- [492] J. Haegeman et al. Time-dependent variational principle for quantum lattices. *Physical Review Letters*, 107(7):070601, 2011.
- [493] Q. Zhuang, P. Shor, and J. Shapiro. Resource theory of non-gaussian operations. *Physical Review A*, 97(5):052317, 2018.
- [494] R. Takagi and Q. Zhuang. Convex resource theory of non-gaussianity. *Physical Review A*, 97(6):062337, 2018.
- [495] F. Albarelli et al. Resource theory of quantum non-gaussianity and wigner negativity. *Physical Review A*, 98(5):052350, 2018.
- [496] M. Hebenstreit et al. Computational power of matchgates with supplementary resources. *Physical Review A*, 102(5):052604, 2020.
- [497] D. Rattacaso et al. Stabilizer entropy dynamics after a quantum quench. *Physical Review A*, 108(4):042407, 2023.
- [498] O. Hahn, G. Ferrini, and R. Takagi. Bridging magic and non-gaussian resources via Gottesman-Kitaev-Preskill encoding. *PRX Quantum*, 6(1):010330, 2025.
- [499] F. Mele et al. Learning quantum states of continuous variable systems. *Nature Physics*, 2024.

- [500] A. Steane. Error Correcting Codes in Quantum Theory. *Physical Review Letters*, 77(5):793–797, 1996.
- [501] A. Steane. Simple quantum error correcting codes. *Physical Review A*, 54(6):4741–4751, 1996.
- [502] D. Gottesman. *Stabilizer codes and quantum error correction*. PhD thesis, California Institute of Technology, 1997.
- [503] D. Gottesman. A Theory of fault tolerant quantum computation. *Physical Review A*, 57(1):127–137, 1998.
- [504] P. Aliferis, D. Gottesman, and J. Preskill. Quantum accuracy threshold for concatenated distance-3 code. *Quantum Information Computation*, 6(2-):097–165, 2005.
- [505] S. Dasu et al. Computing with many encoded logical qubits beyond break-even. February 2026. arXiv:2602.22211 [quant-ph].
- [506] M. Perlin et al. Fault-tolerant execution of error-corrected quantum algorithms. March 2026. arXiv:2603.04584 [quant-ph].
- [507] P. Rodriguez et al. Experimental demonstration of logical magic state distillation. *Nature*, 645(8081):620–625, 2025.
- [508] B. Reichardt et al. Fault-tolerant quantum computation with a neutral atom processor. November 2024. arXiv:2411.11822 [quant-ph].
- [509] D. Bluvstein et al. A fault-tolerant neutral-atom architecture for universal quantum computation. *Nature*, 649(8095):39–46, 2026. [Erratum: *Nature* 650, E3 (2026)].
- [510] D. Gottesman. Quantum fault tolerance in small experiments. October 2016. arXiv:1610.03507 [quant-ph].
- [511] M. Urbanek, B. Nachman, and W. de Jong. Error detection on quantum computers improving the accuracy of chemical calculations. *Physical Review A*, 102(2):022427, 2020.
- [512] Z. Chen et al. Exponential suppression of bit or phase flip errors with repetitive error correction. *Nature*, 595(7867):383–387, 2021.
- [513] R. Acharya et al. Quantum error correction below the surface code threshold. *Nature*, 638(8052):920–926, 2025.
- [514] B. Hetenyi and J. Wootton. Creating entangled logical qubits in the heavy-hex lattice with topological codes. *PRX Quantum*, 5(4):040334, 2024.

- [515] L. Caune et al. Demonstrating real-time and low-latency quantum error correction with superconducting qubits. *Nature Communications*, 2024.
- [516] R. Gupta et al. Encoding a magic state with beyond break-even fidelity. *Nature*, 625(7994):259–263, 2024.
- [517] F. Vigneau et al. Quantum error detection in qubit-resonator star architecture. *PRX Quantum*, 6(4):040369, 2025.
- [518] M. Abraham et al. A digitally controlled silicon quantum processing unit. *Nature*, 655(8125):1154–1159, 2026.
- [519] J. Preskill. Beyond nisq: The megaquop machine. *ACM Trans. Quant. Comput.*, 6(3):1–7, 2025.
- [520] C. Gerhard and T. Brun. Weakly fault-tolerant computation in a quantum error-detecting code. *Phys. Rev. A*, 114(1):012429, 2026.
- [521] B. Reichardt, D. Aasen, and R. Chao. Fire and ice: Partially fault-tolerant quantum computing with selective state filtering. May 2026. arXiv:2605.15344 [quant-ph].
- [522] Y. Akahoshi et al. Partially fault-tolerant quantum computing architecture with error-corrected clifford gates and space-time efficient analog rotations. *PRX Quantum*, 5(1):010337, 2024.
- [523] K. Yamamoto et al. Quantum error-corrected computation of molecular energies. *PRX Quantum*, 7(2):020319, 2026.
- [524] D. Zhong et al. Combining error detection and mitigation: A hybrid protocol for near-term quantum simulation. October 2025. arXiv:2510.01181 [quant-ph].
- [525] S. Bravyi and A. Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A*, 71(2):022316, 2005.
- [526] A. Paetzniak and B. Reichardt. Universal fault-tolerant quantum computation with only transversal gates and error correction. *Physical Review Letters*, 111(9):090505, 2013.
- [527] L. Daguerre et al. Experimental demonstration of high-fidelity logical magic states from code switching. *Physical Review X*, 15(4):041008, 2025.
- [528] A. Kitaev. Fault tolerant quantum computation by anyons. *Annals Phys.*, 303:2–30, 1997.
- [529] E. Dennis et al. Topological quantum memory. *J. Math. Phys.*, 43:4452–4505, 2001.
- [530] C. Chamberland et al. Topological and subsystem codes on low-degree graphs with flag qubits. *Physical Review X*, 10(1):011022, 2020.

- [531] C. Benito et al. Comparative study of quantum error correction strategies for the heavy-hexagonal lattice. *Quantum*, 9:1623, 2025.
- [532] D. Horsman et al. Surface code quantum computing by lattice surgery. *New J. Phys.*, 14:123011, 2012.
- [533] N. Linke et al. Fault-tolerant quantum error detection. *Science Advances*, 3(10):1701074, 2017.
- [534] C. Vuillot. Is error detection helpful on IBM 5Q chips? *Quant. Inf. Comput.*, 18(11-12):0949–0964, 2018.
- [535] R. Harper and S. Flammia. Fault-tolerant logical gates in the ibm quantum experience. *Phys. Rev. Lett.*, 122(8):080504, 2019.
- [536] M. Takita et al. Experimental demonstration of fault-tolerant state preparation with superconducting qubits. *Phys. Rev. Lett.*, 119(18):180501, 2017.
- [537] A. Corcoles et al. Demonstration of a quantum error detection code using a square lattice of four superconducting qubits. *Nature Communications*, 6(1):6979, 2015.
- [538] K. Yamamoto et al. Demonstrating Bayesian quantum phase estimation with quantum error detection. *Phys. Rev. Res.*, 6(1):013221, 2024.
- [539] Y. Wang et al. Fault-tolerant one-bit addition with the smallest interesting color code. *Sci. Adv.*, 10(29):ado9024, 2024.
- [540] B. Pató and N. Klco. Trade-offs in Gauss’s law error correction for lattice gauge theory quantum simulations. February 2026. arXiv:2602.22121 [quant-ph].
- [541] A. Rajput, A. Roggero, and N. Wiebe. Quantum error correction with gauge symmetries. *npj Quantum Information*, 9(1):41, 2023.
- [542] L. Spagnoli, A. Roggero, and N. Wiebe. Qudit stabiliser codes for $\mathbb{Z}_N$ lattice gauge theories with matter. February 2026. arXiv:2602.20661 [quant-ph].
- [543] M. Turco, L. Spagnoli, and A. Roggero. Binary Gauss Stabilizers for Abelian Lattice Gauge Theories. July 2026. arXiv:2607.14861 [quant-ph].
- [544] L. Spagnoli, A. Roggero, and N. Wiebe. Fault-tolerant simulation of Lattice Gauge Theories with gauge covariant codes. *Quantum*, 10:1968, 2026.
- [545] X. Yao. Quantum error correction codes for truncated SU(2) lattice gauge theories. *Physical Review D*, 113(11):114512, 2026.

- [546] M. Carena et al. Quantum error thresholds for gauge-redundant digitizations of lattice field theories. *Physical Review D*, 110(5), 2024.
- [547] Z. Davoudi et al. Report of the Snowmass 2021 Topical Group on Lattice Gauge Theory. In *Snowmass 2021*, 2022. arXiv:2209.10758 [quant-ph].
- [548] J. Roffe et al. Protecting quantum memories using coherent parity check codes. *Quantum Science and Technology*, 3(3):035010, 2018.
- [549] R. Chao and B. Reichardt. Quantum error correction with only two extra qubits. *Physical Review Letters*, 121(5):050502, 2018.
- [550] C. Self, M. Benedetti, and D. Amaro. Protecting expressive circuits with a quantum error detection code. *Nature Physics*, 20(2):219–224, 2024.
- [551] E. Chertkov et al. Error detection without postselection in adaptive quantum circuits. *Physical Review Research*, 8(2):023057, 2026.
- [552] A. Vezvaei et al. Demonstration of high-fidelity entangled logical qubits using transmons. *Nature Communications*, 17(1):3281, 2026.
- [553] Y. Li and S. Benjamin. Efficient variational quantum simulator incorporating active error minimization. *Physical Review X*, 7(2):021050, 2017.
- [554] E. Dumitrescu et al. Cloud quantum computing of an atomic nucleus. *Phys. Rev. Lett.*, 120:210501, 2018.
- [555] A. Kandala et al. Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567(7749):491–495, 2019.
- [556] A. Calderbank and P. Shor. Good quantum error correcting codes exist. *Physical Review A*, 54:1098, 1996.
- [557] R. Chao and B. Reichardt. Fault-tolerant quantum computation with few qubits. *npj Quantum Information*, 4:42, 2018.
- [558] C. Chamberland and M. Beverland. Flag fault-tolerant error correction with arbitrary distance codes. *Quantum*, 2:53, 2018.
- [559] R. Chao and B. Reichardt. Flag fault-tolerant error correction for any stabilizer code. *PRX Quantum*, 1(1):010302, 2020.
- [560] B. Reichardt. Fault-tolerant quantum error correction for Steane’s seven-qubit color code with few or no extra qubits. *Quantum Science and Technology*, 6(1):015007, 2021.

- [561] G. Gehér et al. To reset, or not to reset—that is the question. *npj Quantum Information*, 11(1):39, 2025.
- [562] R. Blume-Kohout and K. Young. Estimating detector error models from syndrome data. April 2025. arXiv:2504.14643 [quant-ph].
- [563] D. Jaschke, L. Carr, and I. de Vega. Thermalization in the quantum Ising model—approximations, limits, and beyond. *Quantum Science and Technology*, 4(3):034002, 2019.
- [564] F. Surace and A. Leroise. Scattering of mesons in quantum simulators. *New Journal of Physics*, 23(6):062001, 2021.
- [565] N. Zemlevskiy. Exclusive scattering channels from entanglement structure in real-time simulations. March 2026. arXiv:2603.15621 [quant-ph].
- [566] X. Mi et al. Noise-resilient edge modes on a chain of superconducting qubits. *Science*, 378(6621):785–790, 2022.
- [567] M. Kormos et al. Real-time confinement following a quantum quench to a non-integrable model. *Nature Physics*, 13(3):246–249, 2017.
- [568] A. Milsted et al. Collisions of false-vacuum bubble walls in a quantum spin chain. *PRX Quantum*, 3(2):020316, 2022.
- [569] G. Lagnese et al. False vacuum decay in quantum spin chains. *Physical Review B*, 104(20):L201106, 2021.
- [570] C. Yin, F. Surace, and A. Lucas. Theory of metastable states in many-body quantum systems. *Physical Review X*, 15(1):011064, 2025.
- [571] L. Pavešič, D. Jaschke, and S. Montangero. Constrained dynamics and confinement in the two-dimensional quantum Ising model. *Physical Review B*, 111(14):L140305, 2025.
- [572] L. Pavešič, M. Di Liberto, and S. Montangero. Scattering and induced false vacuum decay in the two-dimensional quantum Ising model. *Nature Communications*, 2025.
- [573] U. Borla et al. Microscopic dynamics of false vacuum decay in the $2 + 1$ D quantum Ising model. January 2026. arXiv:2601.04305 [quant-ph].
- [574] F. Balducci et al. Interface dynamics in the two-dimensional quantum Ising model. *Physical Review B*, 107(2):024306, 2023. arXiv:2209.08992 [cond-mat.stat-mech].
- [575] L. Pavešić, I. Moss, and S. Montangero. False vacuum decay in a two-dimensional quantum spin system. July 2026. arXiv:2607.01994 [quant-ph].

- [576] A. Leroose et al. Quasilocalized dynamics from confinement of quantum excitations. *Physical Review B*, 102(4):041118, 2020.
- [577] G. Lagnese et al. Quenches and confinement in a Heisenberg–Ising spin ladder. *J. Phys. A*, 55(12):124003, 2022.
- [578] S. Darbha et al. False vacuum decay and nucleation dynamics in neutral atom systems. *Physical Review B*, 110(15):155103, 2024.
- [579] Y. Chao et al. Probing false vacuum decay and bubble nucleation in a rydberg atom array. *Physical Review Letters*, 136(12):120407, 2026.
- [580] D. Luo et al. Quantum simulation of bubble nucleation across a quantum phase transition. May 2025. arXiv:2505.09607 [quant-ph].
- [581] G. Humar et al. Resonant false vacuum decay in two dimensions on a 4000-qubit quantum annealer. June 2026. arXiv:2606.25889 [quant-ph].
- [582] S. Coleman. The fate of the false vacuum. 1. semiclassical theory. *Physical Review D*, 15(10):2929–2936, 1977. [Erratum: Phys.Rev.D 16, 1248 (1977)].
- [583] Jr. C. Callan and S. Coleman. The fate of the false vacuum. 2. first quantum corrections. *Physical Review D*, 16(6):1762–1768, 1977.
- [584] F. Balducci et al. Localization and melting of interfaces in the two-dimensional quantum Ising model. *Physical Review Letters*, 129(12):120601, 2022.
- [585] O. Pomponio et al. Bloch oscillations and the lack of the decay of the false vacuum in a one-dimensional quantum spin chain. *SciPost Phys.*, 12(2):061, 2021.
- [586] T. Esslinger. Fermi-hubbard physics with atoms in an optical lattice. *Annual Review of Condensed Matter Physics*, 1(1):129–152, 2010.
- [587] P. Tarabunga and C. Castelnovo. Magic in generalized Rokhsar-Kivelson wavefunctions. *Quantum*, 8:1347, 2024.
- [588] C. Yuen. Function approximation by walsh series. *IEEE Transactions on Computers*, C-24(6):590–598, 1975.
- [589] C. Yuen. Upper bounds on walsh transforms. *IEEE Transactions on Computers*, C-21(12):1273–1280, 1972.
- [590] P. Derks et al. Designing fault-tolerant circuits using detector error models. *Quantum*, 9:1905, 2025.

- [591] S. Kirkpatrick, C. Gelatt, and M. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.
- [592] L. Perron and F. Didier. CP-SAT. Google OR-Tools, version 9.12, 2025. <https://developers.google.com/optimization/cp/cp-solver/>.
- [593] L. Perron and V. Furnon. OR-Tools. Google, version 9.12, 2025. <https://developers.google.com/optimization/>.
- [594] Y. Yang et al. Simulating prethermalization using near-term quantum computers. *PRX Quantum*, 4(3):030320, 2023.
- [595] G. Kishony et al. Low-cost quantum error mitigation via auxiliary qubit return validation. May 2026. arXiv:2605.28342 [quant-ph].
- [596] D. Zhong and T. Brun. Protection of exponential operation using stabilizer codes in the early fault tolerance era. February 2026. arXiv:2602.13399 [quant-ph].
- [597] C. Dankert et al. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1):012304, 2009.
- [598] H. Zhu. Multiqubit Clifford groups are unitary 3-designs. *Physical Review A*, 96(6):062336, 2017.
- [599] G. Lagnese et al. Detecting a long-lived False vacuum with quantum quenches. *Physical Review Letters*, 133(24):240402, 2024.
- [600] P. Pfeuty. The one-dimensional Ising model with a transverse field. *Annals Phys.*, 57(1):79–90, 1970.