

HARVARD
Kenneth C. Griffin



**GRADUATE SCHOOL
OF ARTS AND SCIENCES**

THESIS ACCEPTANCE CERTIFICATE

The undersigned, appointed by the
Department of Physics
have examined a dissertation entitled


Digital quantum simulation and error correction with atomic processors

presented by Simon James Evered

candidate for the degree of Doctor of Philosophy and hereby
certify that it is worthy of acceptance.

Signature 

Typed name: Prof. Mikhail Lukin

Signature 

Typed name: Prof. Susanne Yelin

Signature 

Typed name: Prof. Vladan Vuletić (MIT)

Date: January 28, 2026

Digital quantum simulation and error correction with atomic processors

A DISSERTATION PRESENTED
BY
SIMON JAMES EVERED
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
JANUARY 2026

©2026 – SIMON JAMES EVERED
ALL RIGHTS RESERVED.

Digital quantum simulation and error correction with atomic processors

ABSTRACT

Quantum processors have the potential to significantly advance our understanding of quantum systems. In particular, the programmability of digital quantum devices enables access to highly tunable quantum dynamics and observables. The central challenge, however, is suppressing errors, making quantum error correction essential for large-scale algorithms. In this thesis, I present work on quantum processing with reconfigurable atom arrays, with experiments featuring up to 448 neutral atom qubits, high two-qubit gate fidelities, arbitrary connectivity, and mid-circuit qubit readout and reuse. Using this system, I describe how we realize gate-based quantum simulations. In particular, we experimentally study Kitaev’s honeycomb model and explore hardware-efficient fermion-to-qubit encodings. We then leverage unique opportunities with atoms to study the key building blocks of scalable quantum error correction. These include below-threshold performance, deep computation at constant entropy, and universal logical gates. Together, these experiments highlight unique near-term opportunities with atom arrays and chart a path toward future large-scale atomic processors.

Contents

TITLE PAGE	i
COPYRIGHT	ii
ABSTRACT	iii
TABLE OF CONTENTS	iv
LIST OF FIGURES	viii
DEDICATION	lxvi
ACKNOWLEDGMENTS	lxviii
AUTHOR CONTRIBUTIONS	lxxii
o INTRODUCTION	i
o.1 The promise and challenge of quantum information processing	1
o.2 Quantum computing based on atom arrays	3
o.3 High-fidelity, programmable quantum circuits with atoms	4
o.4 Digital quantum simulation	4
o.5 Quantum error correction	6
o.6 Overview of the thesis	8
I NEUTRAL ATOM QUANTUM COMPUTING	10
I.1 Trapping and rearranging atoms	12
I.2 Rydberg excitation and entangling gates	12
I.3 Programmable single-qubit control	16
I.4 Cooling and enhanced loading	19
I.5 Midcircuit readout and feedforward	21
I.6 State-selective readout	27

1.7	Qubit reuse	30
1.8	The design of programmable high-fidelity circuits	32
2	HIGH-FIDELITY ENTANGLING GATES	42
2.1	Two-qubit gates	43
2.2	Parameterized entangling gates	47
2.3	Scaling of Rydberg properties with n	49
2.4	Design choices for Rydberg excitation	50
2.5	Dark states in two-photon Rydberg gates	50
2.6	High-fidelity Bell states	55
2.7	Benchmarking with repeated gates	58
2.8	Probing correlations in gates benchmarking data	59
2.9	Three-qubit CCZ gates	64
2.10	Tunable controlled phase gates	70
2.11	Automated calibration	71
2.12	Simulating two-qubit gate error sources	73
2.13	Gate calibration and benchmarking	75
2.14	Outlook	76
3	DIGITAL QUANTUM SIMULATIONS OF THE KITAEV HONEYCOMB MODEL	77
3.1	Digital quantum simulation	78
3.2	Introduction to experiments with Kitaev's honeycomb model	80
3.3	Realizing digital quantum simulations of Kitaev's model	81
3.4	Exploring the Kitaev model	86
3.5	Floquet evolution of fermion systems	89
3.6	Fermion exchange	93
3.7	Simulating the Fermi-Hubbard model	95
3.8	Outlook	98
4	EXPERIMENTS IN LOGICAL QUANTUM PROCESSING	99
4.1	From control of physical atoms to logical qubits	101
4.2	Logical processor based on atom arrays	102
4.3	Improving entangling gates with code distance	103
4.4	Fault-tolerant logical algorithms	107
4.5	Complex logical circuits using 3D codes	110
4.6	Quantum simulations with logical qubits	113
4.7	Outlook	115
5	INGREDIENTS FOR UNIVERSAL FAULT-TOLERANT QUANTUM COMPUTING	118
5.1	Introduction	119
5.2	Neutral atom logical processor	122

5.3	Below-threshold performance	124
5.4	Loss detection for improved QEC	126
5.5	Stabilizer measurement during logic operations	128
5.6	Universality and synthesizing arbitrary unitaries	129
5.7	Deep circuits at constant entropy	133
5.8	Discussion and outlook	137
5.9	Error budget and path to 10x below threshold	139
6	ANALOG QUANTUM SIMULATION OF QUANTUM COARSENING AND COLLECTIVE DYNAMICS	141
6.1	Introduction	143
6.2	Ordering dynamics	145
6.3	Domain wall dynamics	147
6.4	Higgs mode	153
6.5	Discussion and outlook	155
7	ARCHITECTURE FOR DIGITAL-ANALOG QUANTUM SIMULATIONS	157
7.1	Introduction	158
7.2	Fast analog-digital quantum simulations	160
7.3	Dynamic Floquet Hamiltonian engineering	163
7.4	Non-equilibrium quantum spin liquid	164
7.5	Signatures of a Rokhsar-Kivelson state	168
7.6	Emergent universal correlations	171
7.7	Discussion and outlook	173
8	FURTHER IMPROVEMENTS TO ENTANGLING CZ GATE FIDELITY	175
8.1	Introduction	176
8.2	Neutral atom entangling gates	177
8.3	Benchmarking low-error CZ gates	179
8.4	High-fidelity circuits with atom motion	182
8.5	Chaotic nonlocal circuits	183
8.6	Discussion and outlook	186
9	CONCLUSION	188
9.1	Charting the path to large-scale atomic processors	189
9.2	Towards useful digital quantum simulations	190
9.3	Frontier in error-corrected quantum processing	190
	APPENDIX A SUPPLEMENTARY INFORMATION ABOUT KITAEV MODEL EXPERIMENTS	191
A.1	Experimental layout	191
A.2	Measurement-based preparation of topological states	192

A.3	Decoding error detection	195
A.4	Postselection based on atom loss	196
A.5	Floquet evolution circuits	196
A.6	Encoding fermions in qubits	198
A.7	Free-fermion states and Hamiltonians	201
A.8	Chern number	203
A.9	Numerical simulations with errors	206
A.10	Emergent particle number conservation	207
A.11	Fermion hopping and density-density correlations	209
A.12	Fermion exchange protocol	210
A.13	Fermi-Hubbard implementation	212

APPENDIX B SUPPLEMENTARY INFORMATION ABOUT EXPERIMENTS IN LOGICAL ALGORITHMS 214

B.1	Correlated decoding	215
B.2	Direct fidelity estimation and tomography	217
B.3	Sliding-scale error detection	218
B.4	Overview of QEC methods	219
B.5	Surface code implementation.	223
B.6	[[8,3,2]] circuit implementation	230
B.7	Classically hard circuits with [[8,3,2]] codes	233
B.8	Simulation of bitstring probabilities	236
B.9	Physical qubit circuit implementations	238
B.10	Two-copy measurements	239

APPENDIX C SUPPLEMENTARY INFORMATION ABOUT INGREDIENTS FOR UNIVERSAL FAULT-TOLERANT QUANTUM COMPUTING 245

C.1	Analysis of error correlations	246
C.2	Decoding	248
C.3	Benchmarking surface code performance	255
C.4	Additional experiment and data analysis details	258

APPENDIX D SUPPLEMENTARY INFORMATION ABOUT ANALOG QUANTUM SIMULATION OF QUANTUM COARSENING AND COLLECTIVE DYNAMICS 265

D.1	Experimental platform	266
D.2	Local control	267
D.3	Theoretical background of coarsening dynamics	269
D.4	Structure factor and correlation length	272
D.5	Analysis of domains in global sweeps	274
D.6	Classical energy analysis	274
D.7	Analysis of locally prepared domains	275

D.8	Numerical simulations of local domains	276
D.9	Amplitude/Higgs mode	277
APPENDIX E SUPPLEMENTARY INFORMATION ABOUT ARCHITECTURE FOR ANALOG-		
	DIGITAL QUANTUM SIMULATIONS	292
E.1	Experimental system	293
E.2	Details of experiment configuration	297
E.3	Coherent mapping	299
E.4	Loss in analog simulations	303
E.5	Approach to Floquet Hamiltonian engineering	304
E.6	Engineering hopping dynamics	306
E.7	Verification of the effective Hamiltonian	308
E.8	Many-body spectroscopy	309
E.9	Domain wall dynamics under Rydberg blockade	311
E.10	Engineering ring-exchange dynamics	312
E.11	Overview of the $U(1)$ quantum spin liquid	314
E.12	Closed-loop optimization of state preparation	322
E.13	Boundary engineering with local detunings	325
E.14	Probability distribution of dimer configurations	326
E.15	Calculation of correlation functions	327
E.16	Measuring many-body coherences	329
E.17	Numerical methods	330
APPENDIX F SUPPLEMENTARY INFORMATION ABOUT FURTHER IMPROVEMENTS		
	TO ENTANGLING CZ GATE FIDELITY	337
F.1	Experimental system	338
F.2	Benchmarking protocols	343
F.3	CZ gate error sources	345
F.4	High-fidelity quantum circuits	356
F.5	Theoretical analysis of nonlocal circuits	359
APPENDIX G EXAMPLE EXPERIMENT COMMAND STRINGS		361
G.1	Moving AWG	361
G.2	Raman AWG	368
G.3	Local Raman AODs AWG	373
G.4	Rydberg AWG	374
REFERENCES		407

Listing of figures

1	Key ingredients of quantum information processing. Superposition and entanglement are the core underlying ingredients of quantum information processors. A single qubit can be held in a superposition of states $ 0\rangle$ and $ 1\rangle$, and many qubits can be entangled to access the 2^n dimensional Hilbert space.	2
2	Examples of errors that can occur in quantum systems. Suppressing errors is the central challenge of useful quantum computation. Coherent errors can build up and lead to the incorrect many-body quantum evolution. Additionally, incoherent errors, loss, leakage, and atom heating can corrupt the computation.	3
3	Analog simulation, digital simulation, and error correction. In near-term devices, analog simulation can be powerful for realizing native quantum dynamics. Meanwhile, digital devices can offer highly programmable quantum dynamics and observables, as well as the implementation of algorithms. However, these devices suffer from intrinsic error rates, such that we need quantum error correction to eventually reach large-scale, precise quantum computation and simulations.	5
4	Core ingredients of quantum error correction. Quantum information is stored redundantly across a group of physical qubits. Local errors can be corrected without disturbing the logical qubit information. Such logical qubits can be improved by increasing the code size, where we expect exponential error suppression with code distance below a characteristic physical error threshold.	6
5	Example: Steane code. A well-known example of an error-correcting code, the Steane code encodes a single distance-3 logical qubit with 7 physical qubits (e.g., atoms), denoted $[[7,1,3]]$ in $[[n,k,d]]$ notation. Image taken from the Error Correction Zoo.	7
6	Surface code and implementation with atoms. a , Surface code, with corresponding X and Z type stabilizers. b , Implementation of a $d = 5$ surface code with atoms, including data and ancilla qubits.	8

1.1	Neutral atom quantum computer architecture. a , Experimental layout, featuring optical tools including static SLM and 2D moving AOD traps, global and local Raman single-qubit laser beams, 420-nm and 1013-nm Rydberg beams, and imaging system for both global and local imaging. b , Level structure for ^{87}Rb atoms, with the relevant atomic transitions employed in this work. c , Control infrastructure used for programming quantum circuits, featuring several arbitrary waveform generators (AWGs). In particular, the moving and Raman 2D AODs are each controlled by two waveforms (one for x axis and one for y axis). An additional AWG is used in first-in-first-out (FIFO) mode for rearrangement before the circuit begins, and then the moving AOD control is switched to the Moving AWG. See Ref. ¹²³ for additional SLM and pre-circuit rearrangement details, Ref. ¹²⁹ for additional Rydberg AWG details and Rydberg excitation details, Refs. ^{262,43} for additional Raman laser and microwave control infrastructure details, and Ref. ⁴³ for additional moving AWG details. All AWGs (other than rearrangement AWG) are synchronized to < 10 ns jitter. During Rydberg gates the traps are briefly pulsed off by a TTL. The FPGA processes images from the camera real-time and in this work sends control signals to the Raman 2D AOD for local single-qubit control. d , Example array layout featuring entangling, storage, and readout zones. Zones can be directly reprogrammed and repositioned for different applications, as well as specific tweezer site locations. Tweezer beams and local Raman control are projected from out-of-plane. The entire objective field-of-view is 400- μm diameter, and consequently we do not expect or observe substantial tweezer deformation near the edges of our processor. During two-qubit Rydberg gates, we place atoms $\lesssim 2 \mu\text{m}$ apart within a gate site, and gate sites are separated such that atoms in different gate sites are no closer than $10 \mu\text{m}$ during the gate. At our present $n = 53$ and two-photon Rabi frequency of 4.6 MHz, the blockade radius is roughly $4.3 \mu\text{m}$, such that adjacent atoms are well-within blockade and distant atoms are well-outside blockade.	11
1.2	Trapping and rearranging atom arrays. We trap arrays of rubidium-87 atoms into static arrangements of optical tweezers generated from a spatial light modulator (SLM), loading up to around 1000 atoms from the magneto optic trap (MOT) into this tweezer array. Subsequently, we use a set of moving optical tweezers generated by a crossed pair of acousto-optic deflectors (AODs) to rearrange the atoms into programmable geometries.	13
1.3	Reconfigurable atom array architecture. In this thesis, we utilize a reconfigurable atom array architecture, in which coherent rearrangement of atomic qubits enables nonlocal connectivity and parallel control of blocks of atoms is achieved through large laser beams sent from the side of the glass cell. Quantum information is stored in long-lived hyperfine qubit states, and then we excite the atoms to highly excited Rydberg states to generate entanglement through controllable entangling gates. Figure adapted from Refs. ¹²³ and ⁴³	13

1.4	Atomic level diagram. Level diagram showing key levels of ^{87}Rb used in our quantum circuits. The clock states, $ 0\rangle$, $ 1\rangle$ are the qubit states used in this work. Excitation to the Rydberg state between $ 1\rangle$ and $ r\rangle$ is carried out by a two-photon transition driven by 420-nm and 1013-nm lasers. Single-qubit rotations are realized with an amplitude-modulated 795-nm laser that drives Raman rotations between the $m_F = 0$ clock states. A DC magnetic field of 8.5 G is applied in most of this thesis, besides Chapter 9.	14
1.5	Single-qubit addressing with local Raman. Fully programmable single-qubit rotations are implemented using Raman excitation through a 2D AOD. Grids of atomic qubits are illuminated in parallel to apply the same gate, while the other qubits are left unaffected.	17
1.6	Details about local Raman addressing. a , $5S_{1/2}$ hyperfine level diagram illustrating the two possible implementations of local single-qubit gates: resonant $X(\theta)$ (purple) and off-resonant $Z(\theta)$ (turquoise) rotations with two-photon Rabi frequencies Ω_{Raman} . In this work, we use the Z rotation scheme and are blue-detuned by 2 MHz from the two-photon resonance. Due to Clebsch-Gordan coefficients, $\tilde{\Omega}_{\text{Raman}}^Z = -\sqrt{3}\Omega_{\text{Raman}}^Z$. b , Schematic showing the conversion of local $Z(\pi/2)$ into local $X(\pm\pi/2)$ gates, where the pulses before (after) the central $Y(\pi)$ have positive (negative) sign, while leaving non-addressed qubit states unchanged. The Gaussian-smoothed local pulses have duration $2.5\ \mu\text{s}$ for $\pi/4$ pulses and $5\ \mu\text{s}$ for $\pi/2$ pulses, and are performed on single rows at a time with a $3\ \mu\text{s}$ gap between subsequent gates to allow the RF tones in the AODs to be changed (including this, duration is $5\text{-}8\ \mu\text{s}$ per row). In this way, arbitrary patterns of qubits, such as the example drawn, can be addressed. c , Calibration procedure used to homogenize the Rabi frequency over a $220\ \mu\text{m} \times 35\ \mu\text{m}$ array. The position calibration is illustrated for 80 sites: approximate $X(\pi/2)$ gates are locally performed and the horizontal/vertical position of all tones is scanned in parallel such that a Gaussian fit returns the optimal alignment. After this, powers are iteratively calibrated until the fitted scale factors for the individual RF tones converge to unity. d , Single-qubit randomized benchmarking of local $Z(\pi/2)$ gates. The local gates are interleaved with random global single-qubit Clifford gates and the final operation C_f is chosen to return to the initial state. Each data point is the average of 100 random sets of Clifford gates, and fitting an exponential decay to the return probability quantifies the fidelity $\mathcal{F}$ per local gate. Note that we apply all 51 global Clifford gates for each data point, such that errors from the global Clifford gates (in addition to SPAM errors) do not contribute to the fitted value.	18

1.7	Enhanced loading with lambda-enhanced gray molasses. Example atom images from tests of enhanced loading of atom arrays for different size arrays, and comparison of normal loading compared to enhanced loading with lambda-gray molasses cooling light on the D1 transition. Generally we find that above 75-80% loading fractions are widely achievable across different array sizes used in this work. We find that for smaller arrays, we can better tune the trap depth (and use higher trap depths) to achieve higher filling fractions.	20
1.8	Midcircuit readout. a , Example experiment using mid-circuit readout and feedforward to perform logical teleportation. b , Mid-circuit readout and feedforward. The imaging histogram shows high-fidelity state discrimination (500 μ s imaging time, readout fidelity $\approx$ 99.8%), and the Ramsey fringe shows that qubit coherence is unaffected by measuring other qubits in the readout zone (error probability $p \sim 10^{-3}$). The FPGA performs real-time image processing, state decoding, and feedforward. . . .	22
1.9	Fast local imaging and mid-circuit feedforward. Left: Single-shot 500 μ s local image in the readout zone, where the peak corresponds to roughly 50 photons collected by the CMOS camera. Right: Diagram of components involved in mid-circuit readout and feedforward steps. Atom detection and logical state decoding occur using the FPGA, which then outputs a conditional TTL to gate local Raman pulses performed on logical qubits in the entangling zone. Bottom: Diagram of approximate timings for a mid-circuit feedforward cycle. First, $F = 2$ population is pushed out (in 10 μ s), and then the remaining $F = 1$ population is imaged locally for 500 μ s. The 24 rows of pixels covering the readout zone are read out to the FPGA in 200 μ s, after which processing is performed. Finally, a conditional TTL output based on the decoded state gates on or off local Raman pulses. The whole readout and feedforward cycle takes less than 1 ms, and can be sped up in the future by optimizing local imaging and camera readout.	23
1.10	Suppressing decoherence with relative lightshifting. Atomic transition and pulse sequence used for local imaging of ancilla qubits. The data qubit trap light shift suppresses data qubit errors, in addition to the large spatial separation between entangling and readout zones. We avoid quickly losing the readout zone atoms during local imaging by using a 5x higher trap depth, and we pulse the ancilla qubit traps and local imaging light to image directly on resonance while avoiding negative effects of large trap light shifts.	25

1.1.1	Measuring decoherence on data qubits from local imaging beam. Characterization of error probability of data qubits during local imaging. From left to right: Data qubit error probability (fraction of population depumped from $F = 2$ to $F = 1$) as a function of local imaging duration out to 20 ms to quantify the effect of the local imaging beam on data qubit coherence for very long illumination. Data qubit error probability after 20 ms of local imaging, as a function of detuning of the local imaging beam, showing suppression of error both red- or blue-detuned from the data qubit transition. Equivalently, increasing the trap depth of the data qubits enables suppression of decoherence due to the local imaging beam. Since qubits in the readout zone are imaged while their traps are pulsed off, any light shift of the data qubit transition from the traps contributes directly to the relative detuning. Right: for a long 10.5 ms local beam illumination with optimal local imaging parameters, we observe a 0.7(1)% increase in data qubit error during an XY8 dynamical decoupling sequence. This suggests a roughly 0.034(5)% error probability for the data qubits during the 500 μ s mid-circuit readout image employed in this work.	26
1.1.2	Atomic physics of hiding beam at 1529-nm. a , Upper panel: The hiding beam is aligned to the storage zone and a knife-edge used to suppress its gaussian-tail in the readout zone. Lower panel shows the relevant atomic levels in ^{87}Rb . The hiding beam couples the excited-to-excited state transition and imparts a 6 GHz light-shift on the $5P_{3/2}$ state. At this detuning, the $5S_{1/2}$ ground state polarizability is approximately 2×10^{-5} times smaller ²⁰⁷ , thus maintaining coherence in the hyperfine qubit manifold while shifting the transition far off-resonance from the imaging light. b , Additional dephasing on atoms in the storage zone due to local imaging beams in the readout zone, for various drive powers and detunings. Here we use 20-ms of illumination and closer probe detuning than in deep circuit experiments. We plot a fit to a simple model (described in the text). For deep circuit experiments, we operate at 16 GHz red-detuned from the bare $5P_{3/2} \rightarrow 4D_{5/2}$ transition (not shown). c , By scanning the detuning of the probe light, we observe an Autler-Townes (AT) splitting of probe resonance when coupled to the $4D_{5/2}$ -level. For this measurement we operate at a low drive power compared to typical values and the drive detuning is 500 MHz red-detuned from resonance, such that both peaks can be experimentally measured in the limited probe detuning range. The small peak is likely due to leaked light and is not expected. d , The measured AT-splitting scales linearly with the Rabi frequency of the drive, or the square root of the drive power. e , Level diagram illustrating the coupled three-level ladder system which leads to the emergence of the Autler-Townes feature.	27
1.1.3	Motivation for state-selective readout. In conventional readout, one of the two spin states is converted to loss before fluorescence imaging. This causes us to lose half the atoms, and prevents detection of loss events. By comparison, using spin-to-position conversion, we can image both spin states, enabling loss detection and efficient qubit reuse.	28

1.14	State-selective readout overview. a-b, Spin-to-position conversion for non-destructive, loss-resolved qubit readout is accomplished with a state-selective 1D lattice that converts the atom spin state into position. Plot shows measured Rabi oscillation. This non-destructive readout has 0.46(4)% bit-flip error and 0.24(2)% loss. c, Results of repeated rounds of $d = 5$ surface code using loss detection, showing a snapshot of the data block and multiple ancilla blocks.	28
1.15	Spin-to-position conversion. a, Level diagram showing the ^{87}Rb hyperfine levels used to engineer a spin-selective one-dimensional optical lattice. b, Trapping potentials for the bright and dark states. The dark state only experiences a lightshift from the optical tweezers - allowing the atom to be moved around freely - while the bright state is additionally confined by the optical lattice potential. By using a blue-detuned lattice, atoms are trapped in intensity antinodes and so scattering of the lattice light is reduced. c, Schematic timeline of spin-to-position conversion. The time to transfer the clock qubit to the bright and dark states for readout is typically on the order of roughly $20\ \mu\text{s}$, but for some of our measurements is several milliseconds due to using a slow global rotation of the magnetic field (panel e). See Methods text for additional information. d, Transfer of qubit state $ 1\rangle$ to the dark state via resonant optical pumping. e, Transfer of qubit state $ 0\rangle$ into $m_F = +1, +2$ states. This is achieved with either a coherent Raman transfer to $F = 2, m_F = +2$ or incoherent pumping with σ^+ -polarized repumper. Both approaches achieve the same bright state readout fidelity, but in the specific implementation we use here the Raman transfer takes several milliseconds owing to the rotation of the external magnetic field for driving σ transitions (2 and 3). f, Quadratic suppression of readout error due to scattering. The bright state transfer ensures that at least two lattice-induced scattering events are required to cause a readout error, which occurs if the AOD tweezer has not yet moved away as the bright state becomes unpinned. Scattering further causes diabatic changes in the depth of the lattice potential and may contribute to atom loss.	29

1.16 Architecture for qubit reuse and constant-entropy computation. **a**, Illustration of processes for removing entropy generated by computation. Logical teleportation is used to ensure all physical errors are removed. **b**, Rabi oscillations measured using the same atoms for 150 cycles of non-destructive measurement and re-initialization. Each curve shows a single experimental run, averaged over 200 atoms in parallel. 3D cooling methods are used in this subplot as coherence does not need to be preserved. **c**, Local cooling with 1D polarization gradient cooling (PGC) and electromagnetically induced transparency (EIT). The finite magnetic field is compensated by applying a relative detuning between the two circularly polarized counterpropagating beams, leading to a rotating frame where the effective field is zero. The atom loss and temperature is constant as a function of time and recovers to steady state within one cycle after applying a perturbation (turning off cooling on one layer) to the system (dashed line). **d**, Additional shielding of the data qubits is provided by a 1529-nm shielding laser, which rapidly suppresses decoherence induced by the imaging (resonance is at 1529.365 nm). 33

1.17 One-dimensional imaging and cooling in finite field. **a**, Atomic level structure for cooling and imaging of ^{87}Rb . The one-dimensional (1D) techniques here rely on coupling hyperfine states separated by $\delta m_F = 2$ with counterpropagating $\sigma^\pm$ beams. In a finite magnetic field B_{ext} , the level degeneracy is lifted by the Zeeman effect; here μ_B is the Bohr magneton and $g_F = 1/2$ the Landé factor. **b**, Single-shot local image in 8.6 G external magnetic field. Spin-to-position conversion is used in the readout zone and the reservoir is partially imaged by the tails of the imaging beams. **c**, Example site-averaged imaging histogram in the readout zone. **d**, Schematic timeline of mid-circuit measurement and re-initialization used for deep circuit experiments. Due to latency bottlenecks associated with desktop-computer-based processing of images and rearrangement waveforms, we do not attempt here to optimize any speeds and simply use comfortable parameters. The circuit is 13.5-ms long. Including idle times, the spin-to-position conversion takes 4-ms, imaging takes 10-ms, cooling and re-sorting takes 13.3-ms (≈ 7 -ms latency for mid-circuit data processing), and re-initialization takes 1.1-ms. **e**, 1D polarization gradient cooling (PGC) in finite magnetic field. Top left: interference between the two counterpropagating $\sigma^\pm$ probe beams generates a helix of linear polarization⁴³⁴. Detuning the two beams rotates the helix such that, in the rotating frame, a fictitious magnetic field appears that cancels the external field and restores the zero-field PGC cooling mechanism. Bottom left: Atom loss from 1D PGC light at different relative beam detunings, Δ , in varying external magnetic field. The atoms are illuminated for 10 ms under comfortable imaging parameters. A cooling resonance is observed when Δ matches two times the Zeeman splitting (bottom right). Top right: Atom temperature around the cooling resonance in 4.3 G field, obtained from drop-recapture measurements. **f**, 1D electromagnetically induced transparency (EIT) cooling. Top: A strong σ^+ pump beam is combined with a weak σ^- probe beam to drive transitions between quantum harmonic oscillator states of the optical tweezer, cooling the atom. The EIT Fano resonance ensures heating transitions are suppressed⁹². Bottom: Ramp-down measurement of atom temperature. The SLM trap depth is adiabatically ramped down to $\sim 5 \mu\text{K}$ and held for 5-ms before being ramped back up; the atom loss probability from this process probes the temperature of all three motional axes⁴²⁶. Since the EIT cooling resonance is narrow and depends on trap frequency, spatial inhomogeneity in the SLM trap depths translates to inhomogeneous cooling and broadens the temperature distribution. Even so, the coldest sites after 1D PGC imaging and EIT cooling reach lower temperatures than conventional 3D PGC, highlighting the utility of these 1D techniques. 34

1.18	Putting all the components together. a , Experimental layout illustrating key optical tools, similar to Ref. ⁴¹ with the addition of beams for local cooling, imaging and hiding to enable qubit re-use experiments. b , Control infrastructure for programming quantum circuits. The entire waveform for all AWGs (except for rearrangement) is uploaded at the start of each experimental run. For qubit re-use experiments, the Moving, Raman AOD and Rydberg AWGs loop the same memory segment each layer. The full waveform is programmed for the Raman AWG to ensure phase continuity. For mid-circuit rearrangement, waveforms are calculated on-the-fly using a desktop computer and sent to the Rearrangement AWG operated in first-in first-out (FIFO) mode. c , Level diagram of the relevant atomic transitions of ⁸⁷ Rb. d , Processor layout used for qubit re-use experiments and relevant laser beams. Atoms are arranged into storage, entangling and readout zones, with an additional reservoir for refilling lost atoms mid-computation. The 1529-nm hiding beam illuminates the storage zone to preserve coherence of active qubits during imaging in the readout zone. Parallel two-qubit gates are performed in the entangling zone with global Rydberg beams, and local detunings are optionally applied to selected gate sites using an SLM. The readout zone is illuminated with local beams for 1D PGC imaging and EIT cooling, as well as two counter-propagating lattice beams to form a spin-dependent potential for readout via spin-to-position conversion. The entire array is addressed with global Raman control for dynamical decoupling. The same Raman light is directed through a pair of crossed AODs for local single-qubit gates. Global imaging and lambda-enhanced gray molasses cooling light are used for the initial loading.	35
1.19	Programming circuits with atoms. Example of a quantum circuit and how it is compiled into the low-level hardware components described in this chapter, including fully programmable single-qubit gates, high-fidelity Rydberg entangling gates with non-local connectivity achieved through atomic motion, and readout capabilities.	36

- 1.2.0 Neutral atom quantum computer architecture.** **a**, Experimental layout, featuring optical tools including static SLM and 2D moving AOD traps, global and local Raman single-qubit laser beams, 420-nm and 1013-nm Rydberg beams, and imaging system for both global and local imaging. **b**, Level structure for ^{87}Rb atoms, with the relevant atomic transitions employed in this work. **c**, Control infrastructure used for programming quantum circuits, featuring several arbitrary waveform generators (AWGs). In particular, the moving and Raman 2D AODs are each controlled by two waveforms (one for x axis and one for y axis). An additional AWG is used in first-in-first-out (FIFO) mode for rearrangement before the circuit begins, and then the moving AOD control is switched to the Moving AWG. See Ref.¹²³ for additional SLM and pre-circuit rearrangement details, Ref.¹²⁹ for additional Rydberg AWG details and Rydberg excitation details, Refs.^{262,43} for additional Raman laser and microwave control infrastructure details, and Ref.⁴³ for additional moving AWG details. All AWGs (other than rearrangement AWG) are synchronized to < 10 ns jitter. During Rydberg gates the traps are briefly pulsed off by a TTL. The FPGA processes images from the camera real-time and in this work sends control signals to the Raman 2D AOD for local single-qubit control. **d**, Example array layout featuring entangling, storage, and readout zones. Zones can be directly reprogrammed and repositioned for different applications, as well as specific tweezer site locations. Tweezer beams and local Raman control are projected from out-of-plane. The entire objective field-of-view is 400- μm diameter, and consequently we do not expect or observe substantial tweezer deformation near the edges of our processor. During two-qubit Rydberg gates, we place atoms $\lesssim 2\ \mu\text{m}$ apart within a gate site, and gate sites are separated such that atoms in different gate sites are no closer than 10 μm during the gate. 37
- 2.1 Level diagram showing error channels for gates.** Level diagram shows the eight states assumed in the simulation. We assume a 88 μs Rydberg state lifetime (based on measured T_1 with 1013-nm scattering lifetime subtracted), and a 110 ns lifetime for the intermediate states. We assume the following branching ratios for the intermediate states³¹⁴: $\eta_{e \rightarrow L} = 0.6142$, $\eta_{e \rightarrow 1} = 0.2504$, $\eta_{e \rightarrow 0} = 0.1354$, and the following ones for the Rydberg states²⁶³: $\eta_{r \rightarrow L} = 0.894$, $\eta_{r \rightarrow 1} = 0.053$, $\eta_{r \rightarrow 0} = 0.053$. We use the branching ratios between different channels of intermediate-state scattering as reported in Refs.³¹⁴, and we additionally assume a simplified model in which all indirect paths (through 4D and 6S) populate the ground-state manifold uniformly. The Rydberg lifetime has both radiative decay (170 μs) and black-body decay (128 μs) components, which we obtain by re-scaling the values in Ref.²⁶³ to $n = 53$. The microwave component results purely in atom loss and we assume that the radiative decay populates the ground-state manifold uniformly. We note that a more accurate treatment of the decay channels⁹⁸ could increase error modeling precision in future work. 46

- 2.2 **Parallel implementation of high-fidelity entangling gates on a neutral atom quantum computer.** **a**, Entangling gates are implemented by arranging atoms into designated gate sites where they interact via Rydberg blockade interactions when pulsing global Rydberg lasers. Two-qubit or three-qubit gates are performed by modulating the Rabi frequency $\Omega(t)$ and phase $\varphi(t)$ profiles of the laser driving the first leg of the two-photon Rydberg excitation. Inset: atoms in the qubit state $|1\rangle$ are excited to the Rydberg state $|r\rangle$ through an intermediate excited state $|e\rangle$, while atoms in $|0\rangle$ are not excited. Main gate error sources include Rydberg state decay γ_r , intermediate-state scattering γ_e , and Rydberg dephasing T_2^* . **b**, Numerical comparison of average bright and dark state populations during the Rydberg gate. Choosing opposite intermediate-state (Δ) and two-photon (δ) detuning signs at the beginning of the gate maximizes population in the dark state, which minimizes the intermediate-state scattering error. **c**, Entangling gates are implemented with a single Rydberg laser pulse with smooth phase modulation $\varphi(t)$, whose slope corresponds to a two-photon detuning $\delta(t)$. A few global parameters characterize the gate, allowing for a family of possible gate implementations, including a parameterized version of the time-optimal CZ gate²²⁵ and a smooth-amplitude CZ gate. **d**, Example gate calibration sequence. Tuning individual parameters of the parameterized time-optimal gate phase profile $\varphi(t) = A \cos(\omega t - \phi) + \delta_0 t$ allows for fast and simple global calibration (see Fig. 2.9 for additional experimental data). Error bars represent 68% confidence intervals. 51
- 2.3 **Simulated error budget for two-qubit CZ gates from Ref. ¹²⁹.** Simulations were performed at an intermediate-state detuning of $\Delta/2\pi = 7.8$ GHz and a two-photon Rabi frequency $\Omega/2\pi = 4.6$ MHz. (*) The scattering error for two detuning signs (bright/dark). The total fidelity values assume the correct (dark) detuning choice. (**) LG — leakage out of the qubit manifold to other m_F states. AL — atom loss from population left in the Rydberg state. See Chapter 8 and Ref. ¹³⁰ for a more up-to-date error modeling accounting for additional error sources we recently understood in more detail. 51
- 2.4 **Dark state physics in two-photon Rydberg gates.** **a**, In the far-detuned limit the three-level system can be effectively described as a two-level system with a “dark” state $|D\rangle$ and a “bright” state $|B\rangle$. The population in the excited state $|E\rangle$ is suppressed by a factor $\propto \Delta^{-4}$ and does not participate in system dynamics. **b**, The Hamiltonian of the bare atomic system and the effective two-level system, where α is the (time-dependent) ratio between the blue and red Rabi frequencies and $\dot{\alpha}$ is its time derivative. **c**, Intermediate-state population during the parameterized time-optimal gate in the dark and bright configurations together with their Bloch sphere trajectories in the $\{|D\rangle, |B\rangle\}$ basis. **d**, Intermediate-state population during the smooth-amplitude gate and its Bloch sphere trajectory in the instantaneous $\{|D\rangle, |B\rangle\}$ basis. 55

- 2.5 **High-fidelity CZ gates characterized with Bell states.** **a**, Parameterized time-optimal CZ gates are implemented on 20 atoms in parallel. **b**, Raw Bell state measurements upon application of a single CZ gate, with raw Bell state fidelity of 98.0(2)%. We estimate a SPAM-corrected fidelity of 99.4(4)% (not plotted). **c**, Circuit used to benchmark two-qubit gate fidelity by making a Bell state after an odd number of CZ gates interleaved with single-qubit X gates. **d**, Decay of Bell state fidelity after applying a variable number of CZ gates. A fidelity of 99.52(4)% per CZ gate is extracted from an exponential fit. The plot y-axis on (d) is rescaled by a SPAM correction factor which does not affect the extracted gate fidelity. Note that the same number of single-qubit gates are performed for each point, so that no normalization of single-qubit errors is necessary, which thereby results in a reduced y-axis intercept. Error bars represent 68% confidence intervals. 57
- 2.6 **Comparing repeated gates benchmarking sequences.** **a**, Benchmarking of the CZ- π -CZ sequence with global random rotations, which is insensitive to the single-particle phase. **b**, Benchmarking a standalone CZ gate with global random rotations, which enables separate calibration of the single-particle phase. **c**, The usual interleaved randomized benchmarking method utilizing random two-qubit Clifford gates (not performed in this work). **d**, Numerical simulation of the presented benchmarking methods, and the Bell-state-preparation method, utilizing the realistic error model developed in this work. All approaches give consistent results, with the Bell-state fidelity measurement lower-bounding the other curves. 59
- 2.7 **Global randomized benchmarking of the CZ gate.** **a**, Circuit used to benchmark two-qubit gate fidelity with global randomized benchmarking, in which random rotations R_{rand} are sampled from a Haar-random distribution and a final rotation R_f is precomputed to return population to the initial product state $|00\rangle$ in the absence of gate errors. All 21 single-qubit rotations are applied for each data point, such that the number of single-qubit operations is independent of the number of two-qubit gates. **b**, Benchmarking CZ gates on 20 atoms in parallel. We plot the probability of return to the initial product state as a function of the number of applied CZ gates for both the parameterized time-optimal gate and the smooth-amplitude gate, extracting fidelities of 99.54(2)% and 99.55(3)%, respectively. **c**, The extracted time-optimal gate fidelity is consistent within statistical error across the 10 gate sites. **d**, Scaling to larger systems, we measure a comparable fidelity of 99.48(2)% for the time-optimal gate implemented on 60 qubits in parallel. **e**, Site-by-site analysis reveals homogeneous, high-fidelity gate performance across the two rows. 60
- 2.8 **Robust calibration of gates as a function of rise time.** Robustness to experimental offsets. For some systematic experimental offsets, such as finite rise time of the 420-nm Rydberg laser pulse, an exact CZ gate can still be found. The relative values of gate parameters for the time-optimal gate are plotted as a function of pulse rise time. For no rise time, the parameter values used here are: $A/2\pi = 0.0988$, $\omega/\Omega = 1.3629$, $\phi_0 = -2.6082$, and $\partial_0/\Omega = -0.0187$ 60

2.9	Empirical optimization of two-qubit gate fidelity. a , Calibrating gate parameters for the time-optimal gate, indicating the chronological sequence of calibrations performed prior to measurement in Fig. 3b. We scan individual gate parameters and measure the probability of return to the initial state after application of 20 entangling gates. b , Analogous calibration of gate parameters for the smooth-amplitude gate, which we performed in the sequence shown prior to the smooth-amplitude measurement. Additional calibration of other gate parameters was performed before these measurements.	61
2.10	CZ gate single-particle phase calibration and benchmarking. a , Digital circuit for global randomized benchmarking method used to calibrate single-particle phase, where a Z rotation can be performed after each CZ gate to compensate for the acquired phase. R_{rand} are single-qubit rotations sampled from a Haar-random distribution, and the 5 rotations R_f are computed to return the atom pair to the initial product state in the absence of gate errors. For the o CZ gates point, 20 random rotations are applied, as well as a final rotation precomputed to return population to the initial state. b , Experimental data used for calibrating the single-particle phase by optimizing the return probability $P_{ 00\rangle}$ after 20 CZ gates (inset). For the optimal choice of ϕ , we extract a 99.48(2)% CZ gate fidelity, fitting to an exponential decay.	61
2.11	Correlated errors in experimental shots. a , Distribution of errors in each experimental shot as a function of the number of CZ gates applied for the 20-atom data, showing qualitative agreement with a Poisson distribution centered at the experimental mean for number of errors in a shot. b , Histogram of the number of errors in a shot, averaging over all numbers of gates for the 20-atom data. We compare to one model assuming a Poissonian distribution of errors about the mean, finding some deviation from our data. In a second model, we consider that in each shot there is a slightly different gate fidelity, sampled from a Gaussian distribution with a mean of 99.54% and standard deviation of 0.3%. This second model appears to better capture our data. c , Repeating the analysis for the 60-atom data, we notably find no shot (out of the 5053 total repetitions) with 18 or more errors out of the 30 gate sites. Again, the data is qualitatively similar to a Poisson distribution model. d , Averaging over the 60-atom data for all numbers of gates, we find again a small quantitative deviation between the data and a model with a Poisson distribution of errors in each shot. The data appears to be better described by a model where in each shot, we sample fidelities from a Gaussian distribution with a mean of 99.5% and a standard deviation of 0.3%.	65

2.12	Correlations between gate sites. a , Covariance matrices for the 20-atom data after 0 gates and 20 gates, where local correlations appear after 20 gates. b , Covariance averaged over all neighbors, next nearest neighbors, and next next nearest neighbors, as a function of number of CZ gates applied. As a guide to the eye, data is fit to quadratic curves. c , Covariance matrices for 60-atom data for 0 gates and 20 gates. d , Plotting covariance for neighboring sites for the 60-atom data. Once again, the covariance between nearby sites exhibits small growth throughout the 20 CZ gates applied, in particular for the nearest neighbor sites.	66
2.13	Effect of P state Rydberg decay on measured CZ gate fidelity. CZ gate fidelity measured via randomized benchmarking ¹²⁹ . Infidelity due to Rydberg P states is removed by leaving sufficient time or distance between gates. Sparse corresponds to 2x larger separation between gate sites.	66
2.14	Time-optimal pulses for multi-qubit controlled phase gates and GHZ state data. a , The execution time of a C_NZ blockade gate as a function of the number of qubits. The N -qubit gates are realized by applying a phase flip to the $ 0\rangle^{\otimes N}$ state, which is not equivalent to the method of Ref. ²²⁵ for $N > 2$. For the CCZ gate, applying a π phase to the $ 111\rangle$ state ($ 111\rangle \rightarrow - 111\rangle$), while leaving all other basis states invariant, is related by a global bit-flip to applying a relative π phase to the $ 000\rangle$ state; however, the two implementations are not equivalent up to a global Z rotation, contrary to the two-qubit case. The time-optimal CCZ gate using the second approach realizes the CCZ gate $\approx 34\%$ faster with $(\Omega T/2\pi) = 1.72$, as compared to $(\Omega T/2\pi) = 2.61$ from the first approach. The two approaches are different because the relative phase of π is accumulated between different basis states, which have different rates of phase accumulation. In the case of applying $ 111\rangle \rightarrow - 111\rangle$, the states with the slowest relative rate are $ 111\rangle$ and $ 011\rangle$ which are driven with the Rabi frequency of $\sqrt{3}\Omega$ and $\sqrt{2}\Omega$, respectively, resulting in the phase accumulation rate proportional to $(\sqrt{3}-\sqrt{2})\Omega \approx 0.32\Omega$. In contrast, when the relative phase is applied on the state $ 000\rangle$, the smallest accumulation rate is given by the $ 001\rangle$ state which is driven with the Rabi frequency Ω . In general, an arbitrary global single-qubit rotation at the end of the gate can be included to incorporate all of the above approaches in the optimization procedure. b , Time-optimal phase profiles (without analytic parameterization) for the C_NZ gates up to 6 qubits realized by applying a phase flip to the $ 0\rangle^{\otimes N}$ state. c , Circuit used to generate the GHZ state $(000\rangle + 111\rangle)/\sqrt{2}$ after two CCZ gates. d , GHZ states measured experimentally upon applying this circuit to 7 three-qubit groups in parallel, with populations in $ 000\rangle$ and $ 111\rangle$ of 92.9(3)% and a parity contrast of 89(1)%, giving a raw GHZ state fidelity of 90.9(6)%.	68

- 2.15 **Realizing fast CCZ gates.** **a**, Entangling zone comprising three-qubit gate sites, with 21 qubits total. The triangular configuration enables strong, symmetric interactions between three qubits. **b**, Circuit used to calibrate the CCZ gate, by generating GHZ states $(|000\rangle + |111\rangle)/\sqrt{2}$ after applying 2, 6, and 10 CCZ gates. **c**, Phase profile used for implementing a parameterized CCZ gate with a single, fixed-amplitude pulse. **d**, Raw GHZ state fidelity as a function of CCZ gates applied, with an exponential decay fit $\propto 0.979(2)^{N_{CCZ}}$ (see Fig. 2.14d for example GHZ state characterization). **e**, Theoretical scaling of gate duration as a function of number of qubits for native Rydberg blockade multi-qubit gates (see Fig. 2.14b for corresponding phase profiles). 69
- 2.16 **Design of controlled-phase gates.** **a**, Example phase profiles for different CPHASE gates. **b**, Gate parameter values as a function of the entangling phase of the CPHASE gate. **c**, Duration of a controlled-phase gate CPHASE(θ). The CZ gate ($\theta = \pi$) is the longest in this gate family. Since faster gates are expected to have higher fidelity, an average CPHASE gate should perform even better than the CZ reported in this work, which is an exciting perspective for near-term digital simulation. **d**, Plotting on a log-log plot, we see that for small angles θ , the gate time decreases with an approximate power law $\Omega T(\theta)/2\pi = 1.6 \times (\theta/2\pi)^{0.23}$. This suggests that applying very small phases can be costly and should be taken into account when designing digital simulation schemes. While these small-angle gates are time-optimal when applying a single, fixed-amplitude pulse, different approaches could perform better. Exploring other exact and approximate gate schemes, such as Rydberg dressing and a detuned 2π -pulse, in the small-angle regime is an interesting direction for future work. 71
- 2.17 **CPHASE gate benchmarking.** **a**, Circuit used to benchmark and calibrate the entangling phase gates, adapted from the approach in ²⁹⁷. One atom is prepared in $|0\rangle$ and the other in $|+\rangle_y = (|0\rangle + i|1\rangle)/\sqrt{2}$ with a local Raman gate. Then a series of CPHASE gates are applied to the atom pair, which effectively rotates the phase of the initial $|+\rangle_y$ state. A precalculated final $Z(\varphi)$ gate ensures that both atoms return to $|0\rangle$ in the absence of errors. Errors, for example coming from a miscalibrated entangling phase or qubit loss/leakage during the gate, will reduce the probability of finding both atoms in $|0\rangle$ at the end of the circuit (the return probability). **b**, Data used to extract the entangling phase in Fig. ??d, shown for two different CPHASE gates. For different numbers of gates applied, we scan the $Z(\varphi)$ gate before the final local $\pi/2$ pulse in order to extract the amount of phase accumulated. The trend shows how much phase has been accumulated over the 10 gates, and the reduction in the peak return probability is a result of gate errors. Note that for a smaller number of gates, we use the same Raman pulse sequence as for the 10 gate sequence and only reduce the number of CPHASE gates. 72

2.18	Automated Rydberg beam calibration. a , Example calculated Rydberg beam profile and the effect of adding two separate peak-correction holograms. The local peak corrections are added to the base holograms with variable weights on the Rydberg beam SLMs, in addition to correcting for more global Zernike aberrations. b , Example automatic calibration routine showing the homogenization of our Rydberg beam across an array. (top) The peak-to-peak variations in row intensity across the array during calibration, as measured by the differential light shift on the hyperfine qubit. (bottom) Example of how the defocus and one of the peak corrections change during this automated calibration procedure. c , Examples of the measured light shift across the rows of our atom array for the two Rydberg beams, comparing the uniformity before and after the automated calibration procedure. We also ensure that the columns are uniform, although they are naturally more homogeneous due to the beam geometry.	73
3.1	Approach to digital simulation using reconfigurable atom arrays. Positions of atoms used for performing digital quantum simulation of the Kitaev model. At each step, the atoms are moved to a certain connectivity, and a controllable amount of evolution is applied using CPHASE gates of tunable angle. This process is repeated to realize a Floquet version of the Kitaev model.	79
3.2	Kitaev's honeycomb model. The Kitaev honeycomb model is a paradigmatic example in topological and condensed-matter physics. In this model, qubits sit on the vertices of a honeycomb lattice, and interact under highly anisotropic interaction terms depending on the orientation of the link between the two qubits. Under this Hamiltonian, 6-qubit loop operators are conserved, and when these are all set to be +1, the model has a rich phase diagram featuring different Abelian and non-Abelian topological phases. There has been a decades-long quest to engineer these anisotropic interactions in Kitaev materials, but the model has yet to be convincingly realized, motivating this first experimental demonstration.	81
3.3	Encoding fermions with qubits. Encoding fermionic particles in qubit systems is challenging, owing to the different commutation relationships of fermions and qubits. As an example, upon fermionic particle exchange, the wavefunction picks up a -1 sign, independent of the distance between the two fermionic particles. In order to encode fermions with qubits, then, this non-local nature must be captured. One conventional approach is to use a Jordan-Wigner mapping, where local fermionic operators become high-weight Pauli operators. However, in practice, this means that the operators to implement even simple nearest-neighbor hopping terms becomes prohibitively difficult in near-term devices. An alternative approach is to leverage long-range entanglement to encode the non-local nature of fermionic operators. For example, in the Kitaev model, local fermionic terms can be realized with only local qubit operations. We will leverage this encoding throughout this Chapter.	82

- 3.4 **Fermion encoding description.** Details of the fermion-to-qubit mapping and its origins. **a**, An example Jordan-Wigner (JW) ordering spanning a honeycomb lattice on a cylinder. In our case the string is constructed as a product of link operators (see text). **b**, The operators which are fixed to $+1$ (stabilizers) on the topological state include the hexagonal plaquettes and the loops around the cylinder. **c**, A local fermion hopping term between two sites results in a macroscopic JW operator. The underlying long-range entanglement enables reducing this operator to a local one by multiplying the string with the relevant stabilizers. **d**, Construction of complex-fermion hopping in terms of spin operators. The length-4 Pauli string cancels the particle-creation terms resulting from the length-2 hopping operator. **e**, Table describing operator mapping between qubits and fermions. Sites i, k belong to the even sublattice and j belongs to the odd sublattice. Sites i', j' complement i, j along the ZZ links. 83
- 3.5 **Digital quantum simulations with reconfigurable atom arrays.** **a**, The honeycomb lattice used in this work with 104 total atomic qubits and periodic boundary conditions along the shorter direction, forming a cylinder. The qubits are encoded in ^{87}Rb atoms and entangling gates are realized through excitation to interacting Rydberg states. To encode fermion statistics, we prepare a long-range entangled state characterized by hexagonal plaquette operators $W_p = X_1 Z_2 Y_3 X_4 Z_5 Y_6$ that commute with the unitary evolution and are therefore conserved. The encoded Majorana fermions live on vertices, at the ends of operator strings, and conventional (complex) fermions are formed by combining two Majoranas along a chosen link orientation. **b**, The long-range entangled state is prepared using mid-circuit measurement of ancilla qubits, and plaquettes are deterministically flipped to be $+1$ using conditional single-qubit gates (red circle). **c**, The Floquet evolution cycle consists of atom reconfiguration interspersed with tunable entangling gates and global basis changes. **d**, Accumulated entangling phase upon repeated application of $\exp(i\pi\theta[Z\otimes Z]/4)$ gates with a variable angle θ realized through fast, parameterized laser pulses. The values show extracted values of θ , and error bars represent one standard deviation. 84

- 3.6 **Measurement-based preparation of topological order.** **a**, A long-range entangled state of the toric-code type is prepared using a depth-3 circuit, independent of the system size, accompanied by mid-circuit readout of the ancilla qubits (I). The measurement results are random up to parity constraints along the periodic direction. A feedforward step realized through FPGA-triggered single-qubit Z rotations pairs the -1 outcomes (white hexagons). Finally, a parallel controlled-Y operation creates the weight-6 plaquettes (hexagons) and initializes the ZZ-link operators (ovals) to all be +1 (II). **b**, Expectation values of the weight-6 plaquettes and weight-2 ZZ-links across the array of 72 data and 32 ancilla qubits with cylindrical boundary conditions. The inset shows values averaged across the system. **c**, Parity expectation values of increasingly large loops, including loops that enclose 1, 2, 3, and 4 hexagons within a column and the loop around the cylinder (plotted for maximum decoding postselection). The largest operator is equivalent to the product of two loops enclosing the cylinder. Error bars represent 68% confidence intervals. **d**, The product of plaquette operators in each column is equivalent to loops enclosing the cylinder, which are +1 in the absence of errors. Postselection based on the parity of the measured ancilla values within each column improves the plaquette expectation values and the quality of the resulting fermion encoding. 87
- 3.7 **Measurement-based topological state preparation.** **a**, Experimental layout, with gate orientation and an example plaquette shown, and ancilla qubit motions for preparing the ZXXZ surface code state on one half of the data qubits. We perform the periodic ancilla measurement step first, such that all other qubits are still in $|0\rangle$ and do not experience most of the Rydberg gate errors. **b**, Atom positions for mid-circuit readout and feedforward, with an example feedforward pattern applied (indicated by orange circles). The mid-circuit image is analyzed, and then based on our decoding algorithm, an FPGA outputs a series of digital TTLs which determine which atoms have a local Raman gate applied. **c**, Data qubit positions for the parallel controlled-Y (CY) operation, completing the state preparation sequence. The sublattice denoted by black circles gets picked up by AOD traps and shifted to the left to perform the parallel CY. This sublattice then remains in AOD traps for the Floquet circuit. 88

- 3.8 **Kitaev model on a honeycomb lattice.** **a**, The honeycomb model, consisting of anisotropic spin interactions along the three directions of the honeycomb lattice ($J_X + J_Y + J_Z = 1$), exhibits topological order with three Abelian (A) phases and a single non-Abelian phase B. The toric code state is the fixed-point state A_Z^1 of the Abelian A_Z phase ($J_Z = 1$). **b**, Starting from this initial state, we prepare different states on the phase diagram. Top: numerically optimized sequence of two-qubit gates used to prepare the non-Abelian phase B. Each circuit layer includes CPHASE (CP) gates and global single-qubit rotations. Bottom: plaquette parity during the evolution, with postselection based on atom loss and decoding. The plaquette parity is lower than due to a longer sequence of gates and atom movements. **c**, Pauli strings of different lengths measured on the three studied states and averaged over the bulk of the system. Error bars represent 68% confidence intervals. **d**, In the fermion representation, the link operators are proportional to nearest-neighbor Majorana hoppings, $K_{ij}^{X/Y/Z} \propto i c_i c_j$. Longer Pauli strings constructed from their products result in longer-range hopping operators. **e**, The free-fermion parent Hamiltonian H_k can be reconstructed from measured two-point Majorana correlations. **f**, The Chern number C is evaluated using the learned parent Hamiltonians of the string distributions in **c**, resulting in $C=0$ in phase A and $C=1$ in phase B. The non-Abelian phase B is characterized by the underlying topological order and an odd Chern number²⁴³. 90
- 3.9 **Probing fermion dynamics.** **a**, System initialized in a two-particle state with complex fermions defined along the ZZ links. An open Pauli string is used to create a fermion pair out of the initial vacuum state. **b**, Digital evolution under two-qubit interactions in Eq. (3.1), realizing Majorana fermion hopping between connected vertices. The hopping strength along the XX and YY links is fixed, $J_X = J_Y = J$, while the value of J_Z/J is tuned. **c**, Total number of complex fermions in the system during a quench under two different Floquet drives, for $J_Z/J=1$ and $J_Z/J=8$. The K_{ij}^Z term corresponds to the fermion particle number and increasing it enforces particle number conservation. The initial background fermion number resulting from state preparation errors is subtracted for clarity and the solid lines are the result of noisy numerical simulations. **d**, System initialized with two fermions at distances $d=1$ or $d=2$ is evolved under the Floquet Hamiltonian with $J_Z/J=8$. **e**, A horizontal cut of the two-point density-density correlation relative to the marked site (green dot) at depth-11. The directional asymmetry in hopping is consistent with Pauli exclusion, which prevents fermions from occupying the same site. Error bars represent 68% confidence intervals. 91

- 3.10 **Fermion hopping dynamics.** **a**, Spatial distribution of complex-fermion density, n_i , after evolution for depth-11 from the initial two-fermion state. The density values are obtained as $n_i = (1 - Z_i Z_{i'} / \bar{z}z) / 2$, where $\bar{z}z$ is the background magnitude of the ZZ link operators in the initial state, $\bar{z}z = 0.728(3)$. **b**, The particle creation of a depth-6 Floquet Hamiltonian $H_{F,6}$, with the evolution time of individual unitaries corresponding to that of the quench experiment, $Jt = 0.125$ (see Ref¹³⁰ for more details). The particle creation is quantified here as a matrix 2-norm of the commutator between the bulk Floquet Hamiltonian and the particle number operator, $\|[H_{F,6}, N_{\text{tot}}]\|_2$. **c**, A spatial plot of density-density correlations normalized by the density of the reference site (green dot), $G_{ij} / \langle n_i \rangle$, in a small neighborhood of the reference site. 91
- 3.11 **Fermion exchange statistics.** **a**, Demonstration of the hopping gate H_1 for the two-fermion basis state (created by applying R twice) and the fermionic vacuum state. **b**, Additional data for the exchange protocol, showing additional intermediate steps of the process. During the middle three steps, the state of the system is described by a superposition of zero fermions and two fermions at different positions in the lattice. 93
- 3.12 **Fermion exchange statistics.** **a**, Direct measurement of the fermion exchange phase. A superposition of vacuum and fermion-pair states is created with a sequence of local entangling gates, which realizes a partial creation unitary R . The system is evolved under two different particle-hopping protocols: (i) hopping to an intermediate site and hopping back and (ii) exchanging the two fermions. The evolution is followed by a second application of R , which maps the exchange phase to fermion presence. Color plots show the mean fermion density at each site in a 4×4 region embedded in the full lattice. **b**, Contrast between protocols (i) and (ii) as a function of decoding postselection. Error bars represent 68% confidence intervals. 94
- 3.13 **Two-component strongly interacting fermions.** **a**, Two independent species of fermions, labeled “spin-up” and “spin-down”, are realized by dividing the system into two halves. The XX and YY links across the boundary are turned off to facilitate independent hopping dynamics; this system realizes a 4×4 lattice with two fermion species. **b**, Floquet evolution consists of independent hopping within the two halves followed by density-density interactions between the two spin states. **c**, The interactions are realized through a parallel entangling operation connecting the two halves of the system. Two parallel CNOT gates propagate a $ZZ(\theta)$ operator, applied to the spin-up half, into the interaction term. **d**, The initial state is a checkerboard pattern, staggered between the two fermion spin states with mean magnetization denoted by the colorbar. **e**, Staggered magnetization dynamics during a quench with and without the interactions. **f**, Final staggered magnetization as a function of interaction strength. Increasing interactions suppress the spin mixing dynamics. The solid lines are the result of noisy numerical simulations of the circuit. Error bars represent 68% confidence intervals. 96

- 3.14 **Fermi-Hubbard implementation.** **a**, Local gates used to prepare the checkerboard pattern. In order to correct the plaquettes, we pre-compensate with the decoding algorithm such that local gates both flip the plaquettes to all be +1 and also initialize the operators on ZZ-links to give the correct complex fermion configuration. **b**, Depiction of how the system is cut in half for the Fermi-Hubbard implementation. Atom moves are shown for the XX, YY, and ZZ terms (no periodic boundary conditions are used here unlike in earlier experiments). These gates comprise the rest of the Floquet circuit in addition to the interaction terms. Note that the middle links are turned off for the XX and YY gates to realize independent hopping on the two halves. 97
- 4.1 **Logical qubit control.** At the basis of the experiments described in this Chapter is the ability to operate on encoded logical qubits efficiently. This is achieved by performing transversal gates in which the same operation is applied on a whole block of physical qubits in parallel. For example, single-qubit logical gates can be performed by shining light on a whole block of atomic qubits with standard AOD optical tools. Even more significantly, coherent transport of atomic qubits enables efficient two-qubit logical gates by directly interlacing the two logical qubits. 100
- 4.2 **A programmable logical processor based on reconfigurable atom arrays.** **a**, Schematic of the logical processor, segmented into three zones: storage, entangling, and readout. Logical single-qubit and two-qubit operations are realized transversally with efficient, parallel operations. Transversal CNOTs are realized by interlacing two logical qubit grids and performing a single global entangling pulse that excites atoms to Rydberg states. Physical qubits are encoded in hyperfine ground states of ^{87}Rb atoms trapped in optical tweezers. **b**, Fully programmable single-qubit rotations are implemented using Raman excitation through a 2D AOD; parallel grid illumination delivers the same instruction to multiple atomic qubits. **c**, Mid-circuit readout and feedforward. The imaging histogram shows high-fidelity state discrimination (500 μs imaging time, readout fidelity $\approx 99.8\%$), and the Ramsey fringe shows that qubit coherence is unaffected by measuring other qubits in the readout zone (error probability $p \sim 10^{-3}$). The FPGA performs real-time image processing, state decoding, and feedforward. . . . 101

- 4.3 **Transversal entangling gates between two surface codes.** **a**, Illustration of transversal CNOT between two $d = 7$ surface codes based on parallel atom transport. **b**, The concept of correlated decoding. Physical errors propagate between physical qubit pairs during transversal CNOT gates, creating correlations that can be utilized for improved decoding. We account for these correlations, arising from deterministic error propagation (as opposed to correlated error events), by adding edges and hyperedges that connect the decoding graphs of the two logical qubits. **c**, Populations of entangled $d = 7$ surface codes measured in the XX and ZZ basis. **d**, Measured Bell pair error as a function of code distance, for both conventional (top) and correlated (bottom) decoding. We estimate Bell error with the average of the ZZ populations and the XX parities. To reduce code distance we simply remove selected atoms the grid, as shown on the right, ensuring unchanged experimental conditions (for $d = 3$, four logical Bell pairs are generated in parallel). Error bars represent standard error of the mean. 105
- 4.4 **Fault-tolerant logical algorithms** **a**, Circuit for preparation of logical GHZ state. Ten color codes are encoded non-fault-tolerantly (nFT), and then parallel transversal CNOTs between computation and ancilla logical qubits perform FT initialization. The ancilla logical qubits are moved to storage, and a 4-logical-qubit GHZ state is created between the computation qubits. Logical Clifford operations are applied before readout to probe the GHZ state. **b**, State-preparation-and-measurement (SPAM) infidelity of the logical qubits without (nFT) and with (FT) the transversal-CNOT-based flagged preparation, compared to physical qubit SPAM. **c**, Logical GHZ fidelity without postselecting on flags (nFT), postselecting on flags (FT), and postselecting on flags and stabilizers of the computation logical qubits, corresponding to error detection (EDFT). **d**, GHZ fidelity as a function of sliding-scale error detection threshold (converted into the probability of accepted repetitions) and of number of successful flags in the circuit. **e**, Density matrix of the 4-logical-qubit GHZ state (with at most 3 flag errors) measured via full state tomography involving all 256 logical Pauli strings. 106

4.5	Zoned logical processor: scaling and mid-circuit feedforward. a , Atoms in storage and entangling zones, and approach for creating and entangling 40 color codes with 280 physical qubits. b,c , 40 color codes are prepared with an nFT circuit, and then 20 transversal CNOTs are used to fault-tolerantly prepare 20 of the 40 codes, whose fidelity is plotted. Logical decoherence is smaller than the physical idling decoherence experienced during the encoding steps. d , Mid-circuit measurement and feedforward for logical entanglement teleportation. The middle of three logical qubits is measured in the X -basis, and by applying a mid-circuit conditional, locally pulsed logical S rotation on the other two logical qubits, the state $ 0_L 0_L\rangle + 1_L 1_L\rangle$ is prepared. e , Measured logical qubit parity with and without feedforward, showing that feedforward recovers the intended state with Bell fidelity of 77(2)% (ZZ parities of 83(4)% not plotted). No mid-circuit refers to turning off the mid-circuit readout, and postselecting on the middle logical being in state $ +_L\rangle$ in the final readout. By postselecting on perfect stabilizers of only the two computation logicals (error detection in the final measurement) the feedforward Bell fidelity is 92(2)% (not plotted). In d three of the additional blocks are flag qubits and the other four are prepared but unused for this circuit.	109
4.6	Complex logical circuits using 3D codes. a , $[[8,3,2]]$ block codes can transversally realize $\{\text{CCZ}, \text{CZ}, \text{Z}, \text{CNOT}\}$ gates within each block, and transversal CNOTs between blocks. By preparing logical qubits in $ +_L\rangle$, performing layers of $\{\text{CCZ}, \text{CZ}, \text{Z}\}$ alternated with interblock CNOTs, and measuring in the X -basis, we realize classically hard sampling circuits with logical qubits. b , Measured sampling outcomes for a circuit with 12 logical qubits, 8 logical CZs, 12 logical CNOTs, and 8 logical CCZs. By increasing error detection, the measured distribution converges toward the ideal distribution. c , Circuit involving 48 logical qubits with 228 logical CZ / CNOT gates and 48 logical CCZs. d , Classical simulation runtime for calculating an individual bit-string probability; bottom plot is estimated based on matrix multiplication complexity. e , Measured normalized XEB as a function of sliding-scale error-detection for 3, 6, 12, 24, and 48 logical qubits. For all sizes we observe a finite XEB score which improves with increased error detection. Diagram shows 48-logical connectivity, with logical triplets entangled on a 4D hypercube. f , Scaling of raw (red) and fully error-detected (black) XEB from e . Physical upper-bound fidelity (blue) is calculated using best measured physical gate fidelities. Diagrams show physical connectivity. $[[8,3,2]]$ cubes are entangled on 4D hypercubes, realizing physical connectivity of 7D hypercubes.	111

4.7	Logical two-copy measurement. a , Identical scrambling circuits are performed on two copies of 12 logical qubits, and then measured in the Bell basis to extract information about the state. Z basis measurements are corrected with an $[[8,3,2]]$ decoder (when full error detection is not applied). b , Measured entanglement entropy as a function of subsystem size, showing expected Page-curve behavior ²⁵⁶ for the highly scrambled state, improving with increased error detection. c , Measured and simulated magic (associated with non-Clifford operations) as a function of number of CCZ gates applied, performed on two copies of scrambled 6-logical-qubit systems. d , Pauli string measurement and zero-noise extrapolation using logical qubits. Plot shows the absolute values of all 4^{12} Pauli string expectation values, which only have five discrete values for our digital circuit; Pauli strings with the same theory value are grouped. By analyzing with sliding-scale error detection we improve toward the theoretical expectation values (squares) while also improving toward a purity of 1. By extrapolating to perfect purity, we extrapolate the expectation values and better-approximate the ideal values (shaded regions are statistical fit uncertainty).	114
5.1	The challenge of designing a QEC architecture. Quantum error correction provides a path to reach very low error rates, but comes at the cost of complexity constraints, associated with qubit overhead, restricted logical operations, and long run times. This motivates consideration of the hardware, logical encoding, and algorithm design to find the most efficient ways of realizing error-corrected quantum processing. In this chapter, we study the core underlying mechanisms of error correction in order to guide this design.	120
5.2	Architectures and mechanisms for fault-tolerant quantum computation. a , We study the key building blocks of fault-tolerant processing. We utilize an architecture based on reconfigurable atom arrays trapped in optical tweezers, where the logical processor is segmented into storage, entangling, readout, and reservoir zones. Underlying physical mechanisms are identified and characterized. b , Spin-to-position conversion for non-destructive, loss-resolved qubit readout is accomplished with a state-selective 1D lattice that converts the atom spin state into position. Plot shows measured Rabi oscillation. This non-destructive readout has 0.46(4)% bit-flip error and 0.24(2)% loss. c,d , Stabilizer measurement on a $d = 5$ surface code is interspersed with global coherent errors injected on the data qubits. Each CZ layer corresponds to one time-step. Repeated correction reduces error build-up through both the Zeno effect and error tracking (right plot is at a fixed $\theta/2\pi = 0.016$). For visual clarity, an acceptance fraction of 50% is used in this plot.	121

- 5.3 **Below-threshold repeated quantum error correction leveraging loss detection.** **a**, Results of repeated rounds of $d = 5$ surface code using loss detection, showing a snapshot of the data block and multiple ancilla blocks. **b, c**, Products of stabilizer measurement results between rounds are used to detect qubit errors, which we refer to as ‘detectors’. ‘Bare’ counts loss as state $|0\rangle$, ‘detect loss’ does not make this erroneous assignment, ‘supercheck’ multiplies detectors around lost atoms, and ‘post.’ post-selects on all atoms of a detector being present. **b**, Detector error probability as a function of data qubit loss in each shot, analyzed by partitioning the total dataset. **c**, Average over all data. **d**, Logical error per round for a surface code after 4 QEC rounds in both bases, decoded using most-likely error methods (‘bare MLE’), machine learning (‘bare ML’), MLE with loss information, and ML with loss information. ML with loss renders the error per round for $d = 5$ as $2.14(13)\times$ lower than $d = 3$. No post-selection is used. Small points are the four $d=3$ quadrants. Results are averaged between $|+_L\rangle$ and $|0_L\rangle$ initialization bases. **e**, Logical error per round as a function of the mean qubit loss, plotted as a cumulative density function. **f**, Relative physical error contribution to overall error budget. **g**, Distribution of detector errors per shot, suggesting the absence of large-scale correlated errors. 123
- 5.4 **Exploring the interplay of logic gates and entropy removal.** **a**, Atom images illustrating two-qubit logic gates and stabilizer measurement. Lattice surgery is realized using ancillas to measure the logical product $Z_L^1 Z_L^2$, and transversal gates are realized via atomic motion. **b**, Quantum circuits for realizing transversal gates and lattice surgery operations. CZ gates are realized as transversal CNOTs and Hadamards. **c**, Dependence of error of the logic operation on the ancilla measurement error. Lattice surgery errors rapidly worsen with increasing ancilla measurement errors (injected in post-processing). **d**, N repeated logic operations are interspersed with rounds of QEC stabilizer measurement. Transversal CNOT has a lower error than lattice surgery (left), and has an optimum of roughly 3 CNOTs per round, as seen most clearly when modest post-selection, characterized by acceptance fractions (AF) are used (right). Error detection is used for lattice surgery in **d** to compensate for having $< d$ rounds. 125
- 5.5 **Characterizing effects of loss and leakage in repeated QEC.** **a**, Superchecks. A lost data qubit results in anti-commuting stabilizers and a flickering error pattern. Producting stabilizers surrounding the lost qubit recovers commuting superchecks. **b**, Detection correlation p_{ij} matrix for five rounds of QEC on a $d = 5$ surface code. Data qubit errors appear as space-like correlations and ancilla measurement errors as time-like correlations between adjacent layers. Leakage results in additional persistent correlations between QEC rounds. By selecting shots with the fewest lost data qubits, these correlations are suppressed, indicating that leakage is dominated by loss which can be detected. Similarly, selecting shots with the most loss enhances the correlations. . . . 127

- 5.6 **Synthesizing arbitrary-angle rotations with a universal fault-tolerant gate set.** **a**, Subjecting codes of various dimensionality to global rotations. Plateaus are seen at robust angles where the stabilizers also revive. 3D Reed-Muller codes, only with all positive stabilizer signs, have an additional plateau at 45-degrees corresponding to non-Clifford T gates. **b**, Programmable angles $R(\theta, \varphi)$ are realized by an alternating sequence of H and T gates. The circuit shows an implementation of such a sequence, with T implemented via state preparation and H implemented via quantum teleportation. **c**, Polar angle plots show generated angles using entangled Reed-Muller codes, measured by tomography, for different maximum numbers of T gates. Experimental results are consistent with theory within statistical error. **d**, Experimental results show that the minimum separation between generated angles decreases exponentially with length of sequence. Inset shows a rescaling where ε is the angular separation to the target angle. Bloch sphere shows all measured angles. For visual clarity, variable degrees of acceptance fraction are used. 130
- 5.7 **Universality with 3D codes.** **a**, Codes of various dimensions subject to global rotation, here plotted without any normalization applied to the logical operator or stabilizer expectation values. **b**, Two-copy measurement of a $[[15,1,3]]$ code after a global rotation¹⁹⁶. The additive Bell magic has a plateau at one unit of magic under a global T . The same magic is obtained by analyzing either the underlying physical 15-qubit system or the single logical qubit with error detection. **c**, The encoded logical state is connected via a unitary Clifford decoding circuit to a product state of the encoded state and Pauli inputs on the physical qubits. The Clifford circuit conserves magic and therefore the total physical state and encoded logical state must have the same total magic. This implies that, while 15 physical T 's are applied to the system, only 1 physical T is produced, which can only happen with entanglement. **d**, Atom image of a register of 2D and 3D color codes. A transversal CNOT can be performed between the face of the $[[15,1,3]]$ code (control) and the $[[7,1,3]]$ (target). **e**, Code switching. $|T_L\rangle$ is prepared transversally on the 3D $[[15,1,3]]$ code and then measurement and in-software feedforward teleports the logical T onto a 2D $[[7,1,3]]$ code (here also with a H); this illustrates the equivalence between code switching protocols and logical teleportation. Error detection is used in both plots. **f**, Error-corrected test of Bell's inequality³¹. We measure $S = E(X, T) + E(X, T^\dagger) + E(Y, T) - E(Y, T^\dagger)$, where $E(A,B)$ is the expectation value of the Steane and Reed-Muller codes in the A and B bases, respectively, and obtain $S = 1.99(3) \times \sqrt{2}$ with error detection. 131

5.8 **Deep logical circuits at constant entropy.** **a**, Schematic of a 2D cluster state created from Steane codes in space and time. The same atomic qubits are re-used every other time layer for up to 27 layers. **b**, Results of repeated logical state preparation of Steane codes. Stabilizer error is constant as a function of cycle. **c,d**, Logical and physical correlations in 1D (**c**) and 2D (**d**) cluster states. Logical correlations persist while stabilizer error correlations rapidly decay. Here and below, various degree of acceptance fraction are applied. **e**, Circuit diagram illustrating that logical-level evolution is unitary and allows the logical operator to propagate throughout the algorithm whereas physical-level evolution is dissipative and does not let physical errors propagate. **f**, High-rate $[[16,6,4]]$ codes are entangled in space and time direction. **g**, Logical evolution and physical error removal with $[[16,6,4]]$ codes (1D cluster state in time direction). Permutation CNOTs within the block (black curve) extend the correlation length in comparison to non-permuted qubits. **h**, Entanglement structure in space and time, with up to $96 d = 4$ logical qubits active simultaneously. **i**, Logical 2D cluster expectation values (averaged across space and the first 9 teleportation layers) as a function of the number of co-propagating logical operators which agree (i.e., cluster state stabilizers have the same outcome for each logical qubit within the block). 134

6.1 **Application of local detunings.** **a**, In addition to the SLM used for atom trapping, a second SLM generates a set of superimposed tweezers that can apply arbitrary detuning profiles to the atoms. The pair of crossed AODs for local hyperfine control⁴¹ are not used in this work. **b**, A spin-echo sequence is used to measure and calibrate the differential AC Stark shift of the local detuning beams on the ground-state clock transition. Darker curves indicate a higher target detuning, weighted between 1 and 4, and the grey curves for unaddressed sites demonstrate negligible crosstalk between sites. The pattern here corresponds to a rotated version of the domains in the main text. **c**, Measured differential AC Stark shift of the $5S_{1/2} \rightarrow 5P_{3/2}$ transition. Unless stated otherwise, the laser amplitude used is 1, corresponding to $\approx 160 \mu\text{W}$ per site and imparting an approximate light shift on the differential $|g\rangle \rightarrow |r\rangle$ transition of $\delta_0 = 12(2)\text{MHz}$. The amplitude is the square root of power relative to $160 \mu\text{W}$. When sites are weighted, the total power imparted on the array remains constant but is proportionally redistributed. **d**, The local detunings are used to prepare deterministic checkerboard orderings, both for a single global order (left) or for domains with different orders (right). A site polarization of unity, measured at the start of the hold time for varying Δ/Ω , corresponds to $|g\rangle$ ($|r\rangle$) for addressed (unaddressed) sites, while -1 corresponds to the flipped state. Note that at finite Δ/Ω , the polarization of $|r\rangle$ sites $\neq 1$; at large Δ/Ω , mean-field shifts reduce the spin polarization of both $|g\rangle$ and $|r\rangle$ along the domain wall. **e**, Dependence of the site polarization on Δ/Ω in the bulk of the domain. As the local light shift becomes weaker relative to the global detuning, the probability of preparing $|g\rangle$ sites decreases slightly. 142

- 6.2 **Ordering via coarsening.** **a**, Atoms arranged in a square lattice are prepared in the ground state $|g\rangle$ and driven from the disordered phase, shown here as a fluorescence image (atoms in $|r\rangle$ are detected as loss and are indicated with red circles), into an antiferromagnetic ordered phase, across a quantum phase transition. The subsequent coarsening dynamics lead to the progressive development of long-range order (shown here using the real-space correlation function $G(x, y)$ and experimental snapshots). The snapshots show the local single-spin-flip-corrected staggered magnetization, with white (dark grey) shading indicating AF_1 (AF_2) orders. **b**, The detuning $\Delta(t)$ is swept linearly, with a fixed Rabi frequency Ω , to a final value Δ in the ordered phase, where it is held constant for the duration of the hold time. **c**, The correlation length ξ grows over the course of the hold time. The solid line is a fit to $(c_0 + c_1 t)^\alpha + c \cos(\omega t + \varphi)$, representing sinusoidal oscillations on top of a power-law growth. Errorbars are smaller than markers. **d**, The early-time growth of the squared correlation length is consistent with a linear increase with time, as expected for coarsening dynamics with a non-conserved order parameter. The dynamics are faster at lower Δ/Ω , closer to the critical point. **e**, The structure factor $S(k, t)$ of the staggered magnetization, for different hold times, at $\Delta/\Omega = 2.0$. Inset: the structure factors collapse onto a single curve when the time axis is rescaled by the correlation length $\xi(t)$ and the magnitude of the structure factor is rescaled by a time-dependent amplitude $\xi^2(t)b(t)$. For the data shown here, $\Omega/2\pi = 3.8$ MHz. 144
- 6.3 **Domains and energy transfer.** **a**, Probability distribution of an atom belonging to an ordered domain of a certain size. At late hold times, atoms are more likely to participate in large domains and very small domains. **b**, The area of the largest domain grows, while that of the second-largest domain decreases as the system is held in the ordered phase. **c**, Classical energy of each snapshot during the hold time, calculated using Eq. (6.2). The total classical energy of the system is conserved while the bulk (domain wall) energy increases (decreases). Inset: separation of domains into bulk and domain walls for a single snapshot. The domain walls are identified as regions of the array where neither AF ordering is observed (light and dark blue). White and grey indicate AF_1 and AF_2 orderings respectively. The data presented here are for $\Delta/\Omega = 3.0$, with $\Omega/2\pi = 3.8$ MHz for **a, b** and $\Omega/2\pi = 6.0$ MHz for **c**. 148

- 6.4 **Dynamics of seeded domains.** **a,** We deterministically create a centered domain of AF_2 order surrounded by a backdrop of AF_1 order, by “pinning” each ground-state atom in the target ordering using a local light shift δ_i , while sweeping $\Delta(t)$ and maintaining a constant Rabi frequency Ω . Upper left: target amplitude of the local detuning pattern; $|\delta_i|$ is inversely proportional to the number of neighboring Rydberg atoms surrounding each ground-state atom. Center: Rydberg population of a single shot immediately after the local detuning is quenched off. Right: local staggered magnetization, demarcating the prepared domains. **b,** Evolution of the average local staggered magnetization with time, as the prepared state evolves under different values of Δ/Ω . **c,** The squared radius of the central domain, r^2 , decreases linearly with time. **d,** The rate of change of the area, r^2 , increases with proximity to the critical point, $(\Delta - \Delta_c)/\Omega$. The dashed line is a guide to the eye for the theoretically expected scaling relationship. **e,** Change in the radially averaged local staggered magnetization with Manhattan distances $d_m = 0, 2, 8$ from the center of the injected square domain. The order initially changes near the domain wall and begins to change in the center of the domain only at later times. Atoms in the bulk of the dominant ordering far from the domain wall, in gray, remain close to their initially prepared states. The dynamics are plotted for $\Delta/\Omega = 2.5$ and $\Omega/2\pi = 6.0$ MHz. 149
- 6.5 **Curvature-driven dynamics.** **a,** We create two separate AF orderings with a zigzag domain wall between them. During the hold time, the domain wall straightens into a vertical line. **b,** Horizontal cuts of the staggered magnetization at points where the domain wall is locally curved (red, blue) versus straight (green, yellow). **c,** The domain wall’s horizontal position for each row (as indicated by the color) as a function of the hold time. In rows where the domain is locally curved, the domain wall moves towards the center. All data shown are with $\Omega/2\pi = 6.0$ MHz; **b** is measured at $\Delta/\Omega = 3.0$ and **a,c** at $\Delta/\Omega = 2.5$ 151

- 6.6 **Higgs mode oscillations.** **a**, Long-lived oscillations of the staggered magnetization observed when one sublattice is initially pinned and then released. **b**, Numerical calculation of the energy difference, in a 10×10 system with periodic boundary conditions, of the initial pinned state (blue) and the $\mathbb{Z}_2$ state (green) relative to the ground state, as well as the gap to the first excited state (red). The pinned state is much lower in energy than the $\mathbb{Z}_2$ state at finite detuning. **c**, Schematic diagram of the effective potential for the amplitude mode. **d**, Oscillation parameters extracted from fitting $\varphi(t) \approx \varphi_0 + A \cos(\omega t + \theta_0)e^{-\gamma t}$ to data shown in **a**. Left: Measured oscillation frequencies (red points), plotted along with numerically determined values and calculated ground-state energy gaps (both for a 10×10 lattice). A lower local detuning $|\delta|$ leads to oscillations of lower amplitude and frequency (orange). Oscillation frequencies of the magnetization are indicated in purple. Upper right: the measured oscillations decay more rapidly as the phase transition is approached. Lower right: in the disordered phase, oscillations occur around $m_s = 0$, while in the ordered phase, a finite offset emerges. **e**, Oscillations of the correlation length. For final detunings close to the quantum critical point, these oscillations are superposed with substantial growth of the correlation lengths. The correlation length oscillates at double the frequency of the order parameter in the disordered phase, and matches the frequency of the order parameter oscillations in the ordered phase. All data shown are for $\Omega/2\pi = 3.1$ MHz. 152
- 7.1 **Analog-digital atomic processor for probing many-body quantum matter.** **a**, Analog evolution is performed in the interacting Rydberg qubit manifold and the hyperfine spin qubit is used for digital processing. Non-destructive readout of each atom obtains spin and loss information, after which lost sites are refilled from the reservoir and the spin state is reinitialized. The circuit repeats until the reservoir is depleted. The geometry of the simulation zone is programmed for each experiment, shown here for a 271-site kagome lattice. **b**, Two π -pulses are used to map between the qubit manifolds, shown for mapping down from the Rydberg to the hyperfine qubit. **c**, Rabi oscillations on the Rydberg qubit, measured after mapping down. The contrast is 97.8(1)% and loss is $\approx 1.8\%$ (not shown), dominated by loss when turning off the tweezer traps and rearrangement infidelity as the reservoir depletes. Here and later, all data are locally postselected on no atom loss. **d**, Ramsey fringe showing coherence after mapping down $|+\rangle$. Atoms are located at next-nearest-neighbor sites for the 1D chain and 2D kagome lattice, emulating a blockaded many-body state. The contrast is reduced in 2D where effects of long-range Rydberg interaction tails are larger. **e**, Many-body states and dynamics can be engineered and probed using this zoned analog-digital architecture. 159

- 7.2 **Hamiltonian engineering via periodic global detuning drives.** **a**, Example detuning profiles per Floquet cycle. Gray profile implements a many-body echo via π phase jumps of the Rydberg drive, realized as smooth detuning pulses. Perturbations are added on top of this echo to engineer an effective Hamiltonian (red profile). **b**, Schematic of a single Rydberg excitation initialized at the center of a 31-site chain, used as the initial state in **c-e**. **c**, Dynamics of microstate populations, grouped by Hamming distance from the initial state in the Z -basis and computed over the central 19 sites of the chain. Under the many-body echo, the state periodically revives, here twice per Floquet cycle (left). The Floquet drive perturbs the dynamics, illustrated by the modified evolution of the populations, to yield an effective Floquet Hamiltonian (right). **e**, At stroboscopic times, the number of spin excitations is approximately conserved. **f**, Single-spin quantum walk under an effective blockaded XX Hamiltonian. The hopping rate is tuned via the drive parameter ε . **g**, Detuning profile per Floquet cycle used to engineer simultaneous two-body hopping and ring-exchange dynamics, and resulting dynamics on a six-site hexagonal plaquette starting from **h**, a single spin and **i**, a Néel state. Data are averaged over 15 plaquettes evolved in parallel. The decay in contrast is primarily attributed to the sensitivity to static detuning offsets. Explicit parameterizations for **a** and **g**. 162
- 7.3 **Dynamical probes of interactions and excitations.** **a**, Two-spin quantum walk under an effective blockaded XX Hamiltonian. **b**, Taking the difference with the summed density profiles of the two individual spins reveals an effective spin-spin repulsion. For clarity, we subtract the background evolution of the vacuum initial state. **c**, Protocol for many-body spectroscopy of a single spin excitation on the central site, $r = 0$. In this case, the ancilla in $|+\rangle$ is replaced by preparation of the central site in four bases, $\alpha \in \{\pm X, \pm Y\}$, and classical post-processing to obtain $\mathcal{G}(r, t)$. **d**, The Fourier transform, $\mathcal{G}(k, \omega)$, as a function of momentum, k , and frequency, ω , across the 1D Brillouin zone. Left schematic illustrates the spin wave excitation. The imaginary component is the spectral function, whose peak is fit to obtain the single-particle dispersion. From this we extract $(\mu + 4U)/2\pi = 0.407(2)$ MHz and $J/2\pi = -0.140(1)$ MHz, and independently verify these coefficients using Hamiltonian learning of the dynamics of different initial states. Background evolution, appearing at $k = 0$, is subtracted. 165

- 7.4 **Dynamical preparation of $U(1)$ quantum spin lakes.** **a**, Potential and kinetic energy per hexagonal plaquette in the Rokhsar-Kivelson Hamiltonian. Potential energy counts the number of flippable plaquettes and kinetic energy captures the plaquette resonances. **b**, Mapping of Rydberg atoms on the vertices of the kagome lattice to a QDM on the honeycomb lattice. The main data in this experiment uses $R_b/a \approx 1.32$. **c**, A multi-objective Bayesian optimizer is used in a closed loop with experiment. Each experimental run returns X and Z basis snapshots of the prepared state from which observables are computed; the RK energies and mean Rydberg density, n , are maximized, whilst atom loss is minimized. The preparation protocol is updated automatically and a new run is started. **d**, (Top) Best-performing detuning sweeps obtained from three overnight optimization runs with eight free parameters. The total time was allowed to vary to produce the longest sweep and was fixed for the shorter sweeps. Dashed line marks the ground state phase transition. (Bottom) Corresponding evolution of loss-postselected mean Rydberg density. Oscillations are due to imperfect adiabaticity starting from all qubits in $|1\rangle$. Dotted line indicates $1/3$ filling, corresponding to maximal packing. **e**, Final monomer, dimer, and blockade-violating double-dimer populations per vertex for the $0.8 \mu\text{s}$ drive. **f**, Evolution of the kinetic and potential RK energies. The reduction of $\langle T \rangle$ relative to $\langle V \rangle$ is larger at later times as the Rydberg populations, and thus vdW tail energies, increase. **g**, (Top) Schematic of the three kagome sublattices (A,B,C) per vertex. (Bottom) Local light shifts are applied to the 74 boundary atoms to compensate for their reduced mean-field vdW interaction energy. **h**, The minimum in the nematic order parameter, Φ , corresponds to uniform bulk Rydberg population on the three sublattices. Dashed line marks the detuning strength used for all subsequent data. **i**, Angular distribution of Φ reveals the emergent global rotational symmetry. Three sharp peaks at $0, \pm 2\pi/3$ are expected in a nematic phase. 166
- 7.5 **Superposition of dimer coverings.** **a**, Probability distribution of dimer configurations in isolated subsystems. The probabilities within each subsystem are nearly equal after projecting surrounding bonds onto no dimer ($|0\rangle$ in the hyperfine qubit) and restricting to perfect dimer coverings without defects, i.e., monomers and double-dimers. Inset shows a non-uniform distribution, prepared with a weakly modified preparation sweep, that is fit to a Boltzmann distribution with inverse temperature β . **b**, Schematic of the RK state. Two local configurations are isolated by projecting bonds around each single-plaquette subsystem onto no dimer. **c**, Bare and projected potential and kinetic energies, averaged over all plaquettes in the lattice. The RK prediction is $\langle T \rangle = \langle V \rangle \approx 0.29$ from classical Monte Carlo simulations. Defects reduce this expectation value in experiment, and the relative reduction of the kinetic energy is largely attributed to measurement error. **d**, Local measurement bases for the projected kinetic energy. **e**, Parity oscillations on single plaquettes using local Raman to apply Z rotations on every other site after mapping down. $U(t)$ represents the analog state preparation. **f**, Bare and projected X loops of increasing perimeter. Solid points are averaged over different measured geometries, shown individually as small points. 169

- 7.6 **Correlations in the emergent $U(1)$ gauge theory.** **a**, Dimer-dimer connected correlator. Solid points show the absolute value of correlations restricted to dimers within plaquettes in the same row. Correlations decay then plateau, indicative of weak long-range order, with additional non-universal structure arising from the lattice geometry. Small points are radially averaged over the lattice after sign rectification. **b**, Spatial plot showing the signed correlations with respect to the central site. **c**, Definition of the dimer-string correlator between plaquette centers with FM normalization. **d**, Dimer-string correlator, restricted to plaquettes in the same row (inset). The corresponding expectation value for the RK state is shown for two lattice dimensions (in units of plaquettes) obtained from Monte Carlo simulations. Monte Carlo simulations use periodic boundary conditions and the results are averaged over all four topological sectors. **e**, Relation between magnitudes of rectified dimer-dimer and dimer-string correlators. Field theory predicts a ratio of 4 on logarithmic axes. Each marker corresponds to the correlators at a given vertex separation. 171
- 8.1 **CZ entangling gates in a neutral atom processor with fast cycle rates.** **(A)** Rydberg gate realized through a smooth-amplitude gate pulse at roughly $\Omega = 2\pi \times 17$ MHz peak Rabi frequency, with a detuning profile $\delta(t)$. Two-photon excitation scheme consists of transitions at 420 nm and 1015 nm, exciting from qubit state $|1\rangle$ to the Rydberg state $|r\rangle = 53S_{1/2}, m_J = -1/2$ through the intermediate state manifold $|e\rangle = 6P_{3/2}$ with intermediate state detuning $\Delta = 2\pi \times 7.8$ GHz. **(B)** Images of the atoms separated into processor zones, with gate sites used for benchmarking and high-fidelity quantum circuits, and state-selective readout to perform loss detection. Fast calibration of CZ gates is enabled by qubit reuse and refilling lost atoms from a large atom reservoir. **(C)** Randomized benchmarking is performed over 40-80 repetitions, with a 20-30 Hz cycle rate. Example calibration of the CZ gate from starting parameters near the theoretical prediction, showing both raw and loss-postselected CZ error. Each point corresponds to the fitted maximum of a parabola scan over one of the gate parameters, with error bars given by the shaded region. 178

8.2	Stable low-error CZ gate performance. (A) Benchmarking methods based on repeated CZ gate applications: global echo randomized benchmarking (“echo RB”) ¹²⁹ and symmetric stabilizer benchmarking (“SSB”) ⁴²⁵ . (B) Return probability as a function of number of CZ gates, fitted to an exponential decay, giving a raw fidelity of 99.854(4)% and a loss-postselected fidelity of 99.941(3)%. (C) Comparison of CZ gate fidelity measured with different benchmarking protocols: echo RB with and without an extra single-qubit π pulse before readout, resonant pushout of $ 1\rangle$ versus state-selective readout (SSR), and SSB versus echo RB. (D) Stability of the CZ fidelity measured for over 10 hours with echo RB, with no recalibration of the gate. Each data point consists of a fit to the decay of return probabilities for 0, 8, 12, and 20 CZ gates. (E) Number of errors per repetition (raw and loss-postselected), where a repetition consists of benchmarking on two rows of 16 atoms brought into the entangling zone, for the data shown in D, consisting of 127,680 repetitions. Solid curves are sums of individual Poissonian fits for each number of CZ gates. (F) Sensitivity of the raw and loss-postselected CZ gate fidelity under perturbation in the 1015-nm Rydberg beam power, keeping the gate pulse parameters unchanged.	179
8.3	Error breakdown and path to further improvement. (A) Illustration of error sources for Rydberg entangling gates, including scattering from the intermediate state γ_e , Rydberg state lifetime γ_r , ground-Rydberg coherence T_2^* , and coupling to the other Rydberg state $ r'\rangle = 53S_{1/2}, m_J = +1/2$. (B) Numerical error budget showing contributions of individual errors, total theoretical errors, and experimental data. Error breakdowns are shown for the current configuration, as well as projected values attainable with technical improvements and higher laser intensities (Appendix F). . . .	181
8.4	High-fidelity quantum circuits with atom motion. (A) Raw and loss-postselected CZ fidelity measured with echo RB, with one atom in a static trap generated by a spatial light modulator (SLM) and one atom in a moving trap generated by an acousto-optic deflector (AOD). (B) Protocol for making and unmaking one-dimensional cluster states to study CZ gate performance in a quantum circuit. (C) Decay of global return probability P_{return} as a function of the number of total CZ gates, fit to an exponential decay, for $N = 20$ qubits. (D) Stability of the fitted decay over time, without recalibration.	182

8.5 High-fidelity scrambling circuits with nonlocal connectivity. (A) Nonlocal quantum circuit implementation. CZ gates of increasing nonlocal connectivity (with labels given below the circuit) are interspersed with single-qubit gates, $R = X(\pi/4)$. The resulting connectivity for the 20-qubit circuit is illustrated on the right. (B) Theoretical values for the half-system-size von Neumann entanglement entropy and $\text{XEB}_{\text{ideal}}$ as a function of the degree of nonlocal CZ gates added to the circuit (“nonlocal connectivity”). Dashed lines indicate the Page entropy value and $\text{XEB}_{\text{ideal}} = 1$. (C) Loss per qubit per CZ gate compared for different circuits in this work: randomized benchmarking (RB), the cluster state circuit in Fig. 8.4b, and the nonlocal quantum circuit in Fig. 8.5a. (D) Sampled outcome distribution for the 10-qubit nonlocal quantum circuit, run with the depth shown in A, comparing the experimental data to the ideal theory expectation and a Porter-Thomas distribution. (E) Sampled outcome distribution for the 20-qubit nonlocal circuit, with the outcomes binned in 100 groups by their probability. Comparison to numerical simulations with noise is shown. (F) Experimentally measured XEB for the 20-qubit circuit, for different degrees of nonlocal connectivity, compared to numerical simulations with noise. Solid bars are raw data, and lighter bars correspond to postselection on no atom loss. 184

A.1 Honeycomb lattice layout and experiment sequence. a, Mapping from honeycomb geometry into experimental array geometry used, indicating link orientations, an example column, and an example plaquette orientation. We use this mapping in order to shrink the number of rows needed, as well as put the atoms into a rectangular grid which is convenient for atom motion and local Raman operations. b, Loss radius definition for error detection based on atom loss. The central (white) atom is the reference point, and for loss radius of 0, we postselect on this atom being present at the end of the circuit. For a larger loss radius, we postselect on atoms within a certain local region being present. For string observables, we perform this procedure for all atoms within the string. c, High-level overview of the experimental sequence used in these experiments, including feedforward topological state preparation, Floquet evolution, and measurement steps. The Fermi-Hubbard quantum simulations have an additional part of the Floquet circuits for engineering the onsite density-density interactions. . . 193

- A.2 **Measurement-based topological state preparation.** **a**, Experimental protocol used for measurement-based preparation of the topological state. First, a $ZXXZ$ toric code state is prepared on one data qubit sublattice (black circles) using ancilla qubits (blue circles). After mid-circuit measurement and feedforward, this prepares the $ZXXZ$ toric code state. The weight-four $ZXXZ$ operators are then extended to weight-6 operators with parallel controlled-Y (CY) operations with a fresh data qubit sublattice originally in $|0\rangle$ (green circles). In the case that the ZZ link (where the CY gates were applied) are $+1$, this weight-6 operator is equal to the plaquette operators $\mathcal{W}_p$. This sequence indeed ensures that these ZZ operators are all $+1$, owing to the propagation of the single-qubit Z stabilizer on the green qubit sublattice through the CY operation, which results in the weight-2 ZZ stabilizers. **b**, Implementation of the encoding steps in parallel across the full array. The top row of ancillas also performs gates with the bottom row of data qubits. 194
- A.3 **Midcircuit decoding and error detection method.** **a**, Example of decoding and feedforward pairing based on ancilla results. White (blue) plaquettes have the value of -1 ($+1$). After the pairing procedure, all plaquettes are prepared with $+1$ value in the absence of errors. Note that the local Z gates flip plaquettes within the column to perform the 1D matching, but do not flip plaquettes in adjacent columns. **b**, Illustration of decoding postselection method based on parity of ancilla measurement results within a column, owing to the cylindrical geometry used in this work. The product of plaquettes within a column is equal to two Z loops enclosing the cylinder, which must be $+1$ in the absence of errors. **c**, Larger loops measured for the initial state, as a function of decoding postselection. We attribute the fact that the loop enclosing four hexagons is higher parity than the loop enclosing three hexagons due to certain differences in the measurement, for example the loops enclosing the cylinder are measured in the Z basis which doesn't require local Raman gates to rotate individual qubits. **d**, Comparison of different state preparation methods. The method we use in this work is based on the $ZXXZ$ toric code prepared with 32 ancilla qubits. This method can also be adapted to require only 16 ancilla qubits. A third method ("Hexagons") consists of directly measuring the weight-6 plaquette operators with 32 ancilla qubits, requiring a slightly deeper state preparation circuit. **e**, Numerical simulation of plaquette parity for these different state preparation methods, for the phenomenological noise model which matches our experimental data (depolarizing noise of strength 0.027 on all qubits after every two-qubit gate layer). The method that we use in this work ($ZXXZ$) performs the best in numerics, both for no postselection as well as with decoding postselection. . . . 197

- A.4 **Chern number analysis.** **a**, The central 12 columns are chosen as the bulk region of the system. **b**, The measured open Pauli strings. **c**, Majorana correlations in the unit-cell basis, obtained from the open strings. The correlation matrix is truncated at distance ± 1 (length-6 strings) to avoid system boundaries when starting from a reference site in the bulk. The (o, e) , (o, o) matrices are evaluated similarly, and related to (e, o) , (e, e) , respectively, through symmetry relations. **d**, The total variation distance of the measured string distribution (in phase B) from the ideal theory values with increasing error detection. The value is normalized by the point with least postselection and improves by $\sim 10\%$ for optimal error detection. **e**, Discretized Berry curvature $F_{\mathbf{k}}$ whose sum over the 1st Brillouin zone is the Chern number. The red arrows represent the discretized phase potential obtained from normalized eigenstate overlaps, with the arrow vector given by $\arg \tilde{U}_{\mathbf{k}} = (\arg U_{\mathbf{k}}^{\hat{x}}, \arg U_{\mathbf{k}}^{\hat{y}})$. **f**, The Chern number for different values of loss radius (LR) postselection, evaluated through bootstrapping with 300 trials as a function of a single batch size. The samples are drawn uniformly from the data set with replacement, and the samples from strings identified with each other are afterwards further averaged. **g**, Change in the Chern number between postselection at loss radius 3 and 0, as a function of batch size. **h**, Phase diagram of the Chern number as a function of per-layer error and initial plaquette parity values, obtained through noisy numerical simulations. The stars denote points corresponding to the phenomenological noise model inferred from experimental measurements (red with maximal postselection). 203
- A.5 **Fermion exchange protocol.** **a**, Four complex-fermion sites involved in the fermion exchange experiments, and where they are embedded in the full lattice. **b**, Full gate sequence for the exchange protocol (partial pair creation, first hop, second hop, and read-out of the exchange phase by repeating the partial pair creation). Four-body interaction terms for the hopping sections are realized by a sequence of two-qubit gate applications. **c**, Depiction of qubits undergoing the two-qubit gates at each time step, denoted by red lines. The black qubits are the atoms that are picked up by AOD traps to perform the gates, and the layout is designed such that gates can be applied in parallel at each time step, as illustrated in the diagrams. 211

- B.1 Additional surface code data.** **a**, Depiction of Bell state circuit and $d = 7$ surface codes. **b**, Diagram showing the transversal CNOT and physical error propagation rules. **c**, Covariance of the 48 measured stabilizers in both bases. The correlations near the diagonal corresponds to adjacent stabilizers within each block. Strong correlations are also observed with the stabilizers of the other block due to the error propagation in the transversal CNOT. **d**, Bell pair infidelity upper bound (as opposed to estimated Bell pair error in Fig. 2d of Main Text, see Methods text), showing improvement with increasing code distance. **e**, Probability of no detected error for each of the 96 measured stabilizers, showing agreement when compared to the theoretical values from empirically chosen error rates (experiment average = 77%, theory average = 82%). Note that X basis logical 1 and Z basis logical 2 have higher stabilizer error probability due to the error propagation in the transversal CNOT (reducing expectation values relative to if the transversal CNOT is not performed). **f**, Using the empirical error rates that correspond to data-theory agreement for the measured stabilizers in **e**, our simulations for improvement in Bell pair error, as a function of code distance, are in good agreement with experiments. The empirical error rates used are consistent with the 99.3% two-qubit gate fidelity, measured for this larger array, as well as the roughly 4% data qubit decoherence error (integrated over the entire circuit and measured by Ramsey method). These dephasing error rates are dominated by a complex moving sequence as we prepare the two surface codes in a serial fashion (see Supplementary Movie), and would be significantly smaller for a repetitive error correction experiment. 216
- B.2 $[[8,3,2]]$ and hypercube encoding.** **a**, State preparation circuit for the $[[8,3,2]]$ code, in which two 4-qubit GHZ states are simultaneously prepared and subsequently entangled. This initializes an $[[8,3,2]]$ code with logical states $|-_{L1}, +_{L2}, -_{L3}\rangle$. **b**, 4D hypercube circuit performed on 48 logical qubits (128 physical qubits). The circuit is drawn on the block-level, where each block consists of 3 logical qubits and 8 physical qubits. The first in-block gate layer is performed with a global $T^\dagger$. The local gate patterns, and the corresponding logical gates they execute within each code block, are illustrated in the inset. **c**, Diagram illustrating the code block movements and use of the processor's zoned architecture throughout the circuit. Initially, eight $[[8,3,2]]$ code blocks are prepared in the entangling zone and atoms for later state preparation of eight additional code blocks are loaded in the storage zone. The code blocks in the entangling zone are then picked up and interlaced with adjacent blocks to perform three transversal CNOT layers. The two groups of eight code blocks are then swapped and the same procedure is repeated with the second group of code blocks. The first group of code blocks are then moved back into the entangling zone and interleaved with the atoms of the first group to perform a final parallel transversal CNOT. The layers of CNOT gates connect the code blocks such that a 4D hypercube on 16 blocks of $[[8,3,2]]$ codes is constructed. 232

B.3 Additional $[[8,3,2]]$ circuit sampling data. **a**, Overlap of error-detected 12-qubit sampling data with the theoretical distribution (same data as fully error-detected case in Fig. 5b). Progressive zoom-in's show the agreement between theory and experiment, down to the level of 10^{-4} probability per bit string. This error-detected data set is composed of 23,545 shots (raw data set is 138,626 shots). Note that we simultaneously measure on two groups of 12 logical qubits; plotted here is only one of the two 12-logical groups with an XEB of 0.69(1), while in plots Figs. 5e,f and Fig. B.3b we average the two logical groups, which gives a measured XEB of 0.616(7). **b**, Same data as Fig. 5f of Main Text but with purity (orange), as measured by two-copy measurement, additionally plotted. The measured XEB is slightly below the measured purity, providing evidence that XEB is a faithful fidelity proxy. We further note that under error detection, the logical XEB for these IQP circuits should be a good fidelity proxy. Interestingly, the behavior can be different for the raw, uncorrected data, as the circuit we apply on the physical level is not IQP. Without applying error detection, not all errors are logical errors and therefore the circuit differs from IQP behavior and can lend itself to a different scaling. For systems of 3, 6, and 12 logical qubits, multiple systems are measured in parallel and their results are averaged together. We note that although our preparation of $[[8,3,2]]$ code states makes these states on a cube, it does not have CNOTs between two pairs of qubits in the first step and, therefore, does not have the full gate connectivity of a cube. Instead, one can interpret these CNOTs as having been included but then compiled away as they commute with the state. We neglect this in plotting our physical qubit connectivity, which is derived from entangling 3D cubes on a 4D hypercube connectivity, realizing a 7D hypercube. **c**, 48-qubit XEB sliding-scale error-detection data. The point with full postselection on all stabilizers being perfect returned only 8 samples, so we omit this point from the Main Text plot for clarity. 234

B.4 Additional Bell basis measurement results. **a**, Histogram of $|\text{tr}(P\rho)|^2$ for all 4^6 Pauli strings P in the 6 logical qubit circuit, as a function of stabilizer postselection threshold (i.e. number of correct stabilizers across the 6×2 logical qubits). Blue (red) indicate Pauli strings that are expected to have $|\text{tr}(P\rho)|^2 = 0.0625$ (o). The separation between the histograms improves as more postselection is applied. **b**, Signal-to-noise (purity divided by statistical uncertainty of purity) as a function of sliding-scale error detection (converted into accepted fraction), for the 12 logical qubit two-copy measurements, where subsystem size 1 indicates a single logical qubit in one copy, and subsystem size 12 indicates all logical qubits. For subsystem size 1, the signal-to-noise ratio gets worse as data is discarded, since the signal does not change (maximally mixed) but the number of repetitions decreases. In contrast, for the global purity, the signal-to-noise increases as near-unity purities are faster to measure⁶⁷. **c,d**, Entanglement entropy when analyzing the circuit as a physical Bell basis measurement as opposed to a logical Bell basis measurement. For logical entanglement entropy calculations we average over all possible subsystems of that given subsystem size, which we find behaves very similarly to e.g. contiguous subsystems due to the high-dimensional hypercube connectivity. In the physical qubit entanglement entropy calculations, we randomly choose from the possible subsystems as there are many. **c**, 6 logical (16 physical) qubits per copy; **d**, 12 logical (32 physical) qubits per copy. The finite sampling imposes a noise floor for very high entanglement entropy values. **e**, Entanglement entropy measurements as a function of logical subsystem size. **f**, Logical circuits used for benchmarking magic. For 1 CCZ, we include U_1 and omit U_0 ; for 2 CCZ, we include U_0 and omit U_1 ; for the 3 CCZ, we include both U_0 and U_1 241

B.5 Surface code preparation and decoding data. **a**, Surface code stabilizers for the two independent $d = 7$ codes following state preparation. The entire movement circuit corresponding to the transversal CNOT is implemented, and the transversal entangling gate pulse is simply turned off. The mean stabilizer probability of success across the 96 total stabilizers is 83%. The high probability of stabilizer success of the two independent codes in both the X and Z bases shows that topological surface codes were prepared (and Fig. B.1 shows that they were preserved during the transversal CNOT). We note that physical fidelities were slightly lower during this measurement due to calibration drift and so these results slightly underestimate performance relative to the data in Fig. 2 and Fig. B.1. **b**, Logical Bell pair error while optimizing the decoder by (inversely) scaling the weights of the inter-logical edges and hyperedges that connect the stabilizers of the two logical qubits (higher values correspond to lower pairing weights). More concretely, the probability p of the error mechanism corresponding to the inter-logical edges / hyperedges is scaled and the weights are calculated as $\log((1-p)/p)$. Qualitatively, optimizing this scaling value optimizes with respect to the probability that errors are before or after the transversal CNOT, as errors before the CNOT will lead to correlations between the two logical qubits, corresponding to the inter-logical edges. As the decoder is optimized by tuning the inter-logical scaling factor, the performance for all three code distances improves, and the larger code distances improve faster when approaching the optimal decoding configuration, as expected. This data is consistent with the decoder being properly optimized for all three code distances, consistent with the fact that our improvement with code size does not originate from suboptimal decoder performance for low distance. Note the y-axis is log-scale. **c**, Logical Bell pair error when using (black) and not using (gray) the ancilla stabilizer measurement values, as a function of the scaling of the inter-logical edges and hyperedges that connect the stabilizers of the two logical qubits. The ancilla measurements contribute to the correction procedure, and contribute more for smaller values of the inter-logical scaling as they correspond to errors that happen before the transversal CNOT. **ox** inter-logical scaling corresponds to conventional decoding within the two independent surface codes. For the **1x** inter-logical scaling plotted here, the $d = 7$ inter-logical scaling parameter is chosen slightly different than in Fig. 2d in order to have consistency across the three code distances (which produces measured values within error-bars). 243

B.6 Theoretical exploration of hypercube IQP circuits. **a,** Anti-concentration property of our circuits. The circuit is said to be anti-concentrated if its output distribution is spread almost uniformly amongst all outcomes, without the probability being concentrated on a subset of bitstrings. This property is crucial for many proofs of classical hardness^{59,52} and, thus, it is desired for our sampling circuits to anti-concentrate. The plot shows that the output distribution of random hypercube circuits (randomized in-block operations and randomized control/target in out-block CNOT layers) anti-concentrates as the dimension of the hypercube is increased and the XEB (which captures the output collision probability) converges to the uniform-IQP value of 2 (here using Clifford circuits; i.e., circuits composed only of random CZ and Z gates)⁵⁹. This suggests that sampling from the ideal output distribution can be classically hard. In general, the hypercube IQP circuit ensemble converges to the uniform IQP ensemble in total variation distance as the depth and hypercube dimension are increased²³¹. The specific circuit instances implemented in the experiment also anticentrate quickly with increasing hypercube dimension. **b,** A single layer of the hypercube circuit admits an efficient tensor-network contraction scheme, which allows us to evaluate the ideal and experimental XEB values. The final out-block CNOT layer is immediately followed by the measurement, which can be incorporated into a non-unitary tensor that is contracted between the two halves of the system (controls and targets of the final CNOT layer). This contraction scheme reduces the memory requirements to half the system size, which enables bitstring amplitude evaluation for the 48-qubit experiment. This simulation approach can be made significantly more expensive by applying additional out-block operations within the two subsystems, forcing the blocking of the intra-partition tensors, which increases the memory and runtime requirements (Fig. 8.5d). **c,** To understand the effects of finite XEB on required classical simulation time, we explore if our circuit families can be “spoofed” with cheaper, approximate simulation that achieves moderately high XEB scores¹⁵⁰, studied here for a 24-qubit system with full state-vector simulation. The spoofing algorithm works by independently sampling from the two halves of the system (two groups of 12 qubits), effectively removing the final layer of CNOTs. This further reduces the simulation complexity, as each of the halves can, in principle, be independently simulated with the efficient approach from **b**. The plot shows that the spoofed XEB for the 24-qubit non-Clifford circuit can be exponentially reduced by extending the circuit with additional gate layers (similarly to the approach used to decrease the performance of the efficient hypercube contraction), for a particular extension of our circuit. This result shows that future work can consider adding additional CNOT layers into these circuits to demonstrate quantum advantage (in the presence of finite experimental noise). 244

- C.1 **Entropy in deep circuits.** **a**, General hypercube encoding. The same entangling circuit structure is combined with programmable input physical states to prepare $[[7,1,3]]$, $[[15,1,3]]$ and $[[16,6,4]]$ codes (members of the family of quantum Reed-Muller codes¹⁶⁶). For the punctured codes, either the top or bottom qubit is removed. **b**, Turning off entropy removal mechanisms in the 2D $[[7,1,3]]$ cluster state circuit. In the absence of re-cooling or refilling of loss, physical stabilizer correlations persist in time. Lost atoms are assigned as qubit state 0 for this plot, and only correlations between codes constructed from the same atomic qubits in every other layer are shown. **c,d** Circuit structure and physical error correlations. By replacing CZ gates by CNOT gates in the 1D $[[7,1,3]]$ cluster state circuit, physical errors can propagate beyond a single layer, extending the stabilizer correlations. The product of adjacent stabilizers commutes with these propagated errors, recovering the rapid decay in physical correlations. 248
- C.2 **Additional data for repeated QEC characterization circuit.** **a**, Processor layout for repeated QEC on a $d = 5$ surface code. The data qubits and one ancilla block are located in the entangling zone (top) and four additional ancilla blocks are in storage (bottom); one ancilla block is unused in the four-round circuit. **b**, Processor layout for repeated QEC on a $d = 3$ code in one of four possible quadrants. **c**, Circuit for four rounds of QEC. For the XZZX rotated surface code⁵¹, $Y(\pi/2)$ gates are applied to one data qubit sublattice (A or B) for preparing and measuring in the X or Z basis. The equivalent circuit is obtained for stabilizer measurement of the CSS rotated surface code upon compiling $Y(\pi/2)$ gates. **d**, Stabilizer gate ordering. The same pattern is used globally in each round. **e**, Detector error probability for $d = 5$. Faint dashed curves correspond to the 24 individual detectors (12 for the first and last rounds), and solid curves are the mean of all deterministic detectors. In this subfigure only, to illustrate the supercheck error distribution, we assign the error of each supercheck to all detectors from which it is composed, and plot the resulting effective detectors; the overall mean is the average of these individual detector values. The first round has lower error due to the neighboring transversal state preparation. The best detector over the central three rounds after loss postselection has 27% lower error than the overall mean; one atom has anomalously high error in the Z basis. **f**, Comparison to simulation. Left: detector error probability, converting loss to qubit state 0. Right: detection correlation matrix p_{ij} ¹⁷⁰. Both metrics show good agreement between simulation and experiment in the structure and magnitude of the errors. **g**, CZ gate fidelity measured via randomized benchmarking¹²⁹. Infidelity due to Rydberg P states is removed by leaving sufficient time or distance between gates. Sparse corresponds to 2x larger separation between gate sites. 256

C.3 Exploring entropy removal during single- and two-logical-qubit operations. **a**, Additional data for entropy removal via stabilizer measurement. As studied in Fig. 1d, the final logical error depends on the balance between the injected error rate (corresponding to the injected $\theta/2\pi$ per time-step) and the entropy removal rate (number of QEC rounds in the fixed total time window). For small injected error, there is an optimal number of QEC rounds (here, 3-4) since stabilizer measurement is imperfect and introduces entropy of its own. The plot uses MLE decoding and an acceptance fraction of 66% to enhance salient features. **b**, Absence of coherent logical error. Using the same error-injection protocol as in (a), the final logical state is measured in both the X and Z basis. With no QEC, the global coherent error results in a coherent logical rotation. With one round of QEC, or more, this coherent rotation vanishes. The stabilizer measurement is performed immediately after transversal state preparation, before the majority of the error is injected, such that the lack of logical rotation compared to no QEC can be attributed to the non-deterministic X(Z) stabilizer signs randomizing the response of the Z(X) logical operator to coherent error. Here we ignore the ancillas in the MLE decoding and use a 50% acceptance fraction. All curves are the average of both bases; for left plot, the measurement is the same basis as preparation, and for the right plot it is the orthogonal basis. **c-e** Analysis of logical gate performance of two logical qubits undergoing repeated transversal CNOTs and QEC, with the circuit studied in Fig. 3 of main text. **c**, The measured detector error probability increases linearly as a function of number of CNOTs applied in each round. **d**, Logical error probability as a function of number of CNOTs per round. Fits are to a functional form of $A \cdot (p_{\text{qec}} + N\Delta p_{\text{det}})^3$, where A , p_{qec} , and Δp_{det} are fitted parameters (blue). Using the fitted values of p_{qec} and Δp_{det} in **c** produces the predicted logical error probability $P_L(p_{\text{det}})$ (red). **e**, Results of **d** divided by total number of CNOT gates. Logical gate fidelity $F_L(p_{\text{det}})$ is $1 - P_L(p_{\text{det}})/(3N)$, where 3 is the number of rounds and N the number of gates per round. 263

C.4 Theoretical characterization of logical-error-per-round ratio. **a**, Numerical simulations of rotated surface code initialized in $|0_L\rangle$ undergoing repeated QEC rounds with stabilizer measurement gate ordering in Ref.¹⁷⁰. As a simple model, we apply a theory error channel with error probability $p = 0.6\%$, where qubit resets and measurements experience uniform single-qubit depolarizing noise, CNOT gates are followed by uniform two-qubit depolarizing noise, and data qubits experience an idling single-qubit uniform depolarizing channel during ancilla qubit resets. We plot different LEPR ratios ($r_{d/(d+2)} = \text{LEPR}(d)/\text{LEPR}(d+2)$) using the PyMatching decoder²⁰¹, observing a change of at most 17% as the number of rounds is increased from four. The LEPR for a circuit with N logical qubits and n QEC rounds is defined as $\text{LEPR} = P_{L,\text{max}}(1 - (1 - P_L/P_{L,\text{max}})^n)$ where $P_{L,\text{max}} = 1 - \frac{1}{2^N}$ is the logical error rate of a fully mixed state. **b**, Numerical simulations of repeated QEC on a single $|+_L\rangle$ surface code using the same circuit as the $d = 5$ surface code experiment in Fig. 3 of the main text. We use a simplified error model where we take the experimental error model and then turn all loss rates to 0 and double the Pauli error rates for idle errors, reset, measurement, and gate errors (single and two-qubit gates) on both data and ancilla qubits. We plot LEPR ratios for varying numbers of QEC rounds using an MLE decoder⁷⁰, observing changes of at most 9%, indicating that performance remains stable even for deeper circuits. **c**, Numerical simulations for circuits with 25 QEC rounds on four logical qubits, with interleaved transversal gate layers using the approximate experimental error model described above. Half of the qubits, selected uniformly at random, are initialized $|+_L\rangle$, while the other half are initialized in $|0_L\rangle$. Each gate layer consists of random pairing of transversal CNOT gates followed by logical single-qubit Pauli gates selected uniformly at random. After applying the random gate sequence U , the inverse $U^\dagger$ is applied, followed by transversal measurement in the same basis as initialization. By varying the number of gate layers interspersed between rounds, the LEPR ratio averaged over ≈ 100 randomly sampled circuits varies by at most 2%. 264

- D.1 Application of local detunings.** **a**, In addition to the SLM used for atom trapping, a second SLM generates a set of superimposed tweezers that can apply arbitrary detuning profiles to the atoms. The pair of crossed AODs for local hyperfine control⁴¹ are not used in this work. **b**, A spin-echo sequence is used to measure and calibrate the differential AC Stark shift of the local detuning beams on the ground-state clock transition. Darker curves indicate a higher target detuning, weighted between 1 and 4, and the grey curves for unaddressed sites demonstrate negligible crosstalk between sites. The pattern here corresponds to a rotated version of the domains in the main text. **c**, Measured differential AC Stark shift of the $5S_{1/2} \rightarrow 5P_{3/2}$ transition. Unless stated otherwise (e.g., in Extended Data Fig. D.6), the laser amplitude used is 1, corresponding to $\approx 160 \mu\text{W}$ per site and imparting an approximate light shift on the differential $|g\rangle \rightarrow |r\rangle$ transition of $\delta_0 = -12(2)\text{MHz}$. The amplitude is the square root of power relative to $160 \mu\text{W}$. When sites are weighted, the total power imparted on the array remains constant but is proportionally redistributed. **d**, The local detunings are used to prepare deterministic checkerboard orderings, both for a single global order (left) or for domains with different orders (right). A site polarization of unity, measured at the start of the hold time for varying Δ/Ω , corresponds to $|g\rangle$ ($|r\rangle$) for addressed (unaddressed) sites, while -1 corresponds to the flipped state. Note that at finite Δ/Ω , the polarization of $|r\rangle$ sites $\neq 1$; at large Δ/Ω , mean-field shifts reduce the spin polarization of both $|g\rangle$ and $|r\rangle$ along the domain wall. **e**, Dependence of the site polarization on Δ/Ω in the bulk of the domain. As the local light shift becomes weaker relative to the global detuning, the probability of preparing $|g\rangle$ sites decreases slightly. 283
- D.2 Correlation lengths and structure factor.** **a**, Structure factor $S(\mathbf{k})$ calculated at early (left) and late (right) times. **b**, Radially averaged structure factor $S(k)$ at early (left) and late (right) times for $\Delta/\Omega = 2.0$. The correlation length ξ is extracted by fitting $S(k)$ to $b\xi^2/\pi(k^2\xi^2 + 1)^{3/2}$. **c**, At various values of Δ/Ω , the structure factors collapse onto a single curve. In this case, we show data for $\Delta/\Omega = 3.0$. **d**, We find a clear dependence of amplitude b on the nonequilibrium correlation length ξ , and that their relation is Δ -dependent. **e**, For early times, the squared correlation length ξ^2 evolves linearly with time for all values of Δ/Ω tested. **f**, Dependence of $\partial_t \xi^2$ on the distance from the quantum critical point. We find that the dynamics accelerate as the system is held closer to the phase transition. The departure from the theoretically expected scaling (dashed line) stands in contrast to the case with locally seeded domains and points to the possible role of finite-size effects due to boundaries for these global sweeps. 284

- D.3 **Analysis details on domain wall dynamics.** **a**, Convolution to calculate the coarse-grained-local staggered magnetization for each single shot. This is used to distinguish the bulk checkerboard orderings from the domain walls. Extremal values of 0 and 4 indicate the ideal coarse-grained local staggered magnetization of either order while values of 1–3 indicate deviations. **b**, Demarcation of domain wall boundaries and the bulk for the same shot as shown in **a**. Red (white) indicates the domain walls (bulk). **c**, Staggered magnetization radially averaged at each Manhattan distance from the center of the prepared square. The radius is extracted as the d_m at which the linearly interpolated m_s crosses 0 (red lines). **d**, Sweeps to various end detunings for the motion of the zigzag domain wall. The acceleration of the dynamics near the QPT is also clearly supported for this shape of the domain wall. For both **c** and **d**, the left and right plots show data for $\Delta/\Omega = 2.0$ and $\Delta/\Omega = 3.5$, respectively. All data shown is for $\Omega/2\pi = 6$ MHz. 285
- D.4 **Numerical simulations of locally prepared domains.** **a**, TDVP simulations of coarsening dynamics observed when the system is initialized with a single square domain in the center (performed on a 15×15 lattice). The initial state approximates the experimental state preparation. **b**, **c**, The change in the radially averaged local staggered magnetization at Manhattan distances $d_m = 0, 2, 8$ from the center of the injected square. The layer at $d_m = 2$, closer to the domain boundary, changes order before $d_m = 0$. We also find that similarly, the atoms in the bulk ($d_m = 8$) remain in their initially prepared order. In these simulations, a slight acceleration in the dynamics is also observed on going closer to the phase transition. **d**, Simulation of the zigzag initial state on a 16×15 lattice. **e**, As in the experimental results, the locally curved points of the domain wall move towards the center of the domain while points with no domain-wall curvature remain stable. 286
- D.5 **Numerical simulations of the amplitude mode.** **a**, Bond-dimension scaling of the energy gaps $\Delta E_1, \Delta E_2$ of the first two excited states above the ground state as a function of Δ/Ω , obtained via DMRG on a 10×10 lattice with (i) open boundary conditions (OBC), and (ii) periodic boundary conditions (PBC). **b**, Order-parameter dynamics at $\Delta/\Omega = 2$, of an initial state prepared with pinning fields $\delta_l/\Omega = -0.7$, on a 10×10 lattice with OBC, simulated via TDVP with $\chi = 64, 128$. **c**, Order-parameter dynamics at (i) $\Delta/\Omega = 0$, (ii) $\Delta/\Omega = 1.15$, (iii) $\Delta/\Omega = 2$, of an initial state prepared with pinning fields $\delta_l/\Omega = -0.7$, on a 10×10 lattice with PBC, simulated via TDVP with $\chi = 64, 128, 256$. The order-parameter dynamics are well modeled as a damped harmonic oscillator, as shown by the fitting. **d**, Oscillation parameters (blue): (i) frequencies, (ii) damping rates, (iii) offsets, and (iv) amplitudes as a function of Δ/Ω , obtained from functional fits to the numerical simulations on a 10×10 PBC lattice with $\chi = 128$ shown in **c**. The frequencies agree with the bulk (PBC) energy gaps for a 10×10 PBC lattice obtained in panel **a**(ii), and with the oscillation frequencies for a 8×8 PBC lattice, for comparison. 287

- D.6 Effect of applied local detunings on the Amplitude mode a**, Amplitude mode oscillations in the disordered phase at $\Delta/\Omega = 0$. The magnitude of the applied local detunings is varied as a fraction of $\delta_0 = -2\pi \times 12(2)\text{MHz} \approx -4\Omega$ (see Extended Data Fig. D.1c). **b**, Higgs-mode oscillations for varying magnitudes of the local detunings at $\Delta/\Omega = 1.1$. Note that for **a** and **b**, the frequency extracted here for $\delta = 0.06\delta_0 \approx -\Omega/4$ is plotted in yellow for comparison. **c**, With decreasing magnitude of δ , we observe a decrease in the oscillation frequency. The dependence is more pronounced closer to the phase transition. This behaviour can be qualitatively captured by considering the anharmonicity of the Landau mean-field potential in Eq. (D.11) as the phase transition is approached. **e**, The damping is stronger near the phase transition for both detunings. **f**, The oscillations at $\Delta/\Omega = 1.1$ are centered around finite φ_0 suggesting a transition into the ordered phase at this value. The data shown here is taken for the same Ω as all other Higgs-mode oscillations, $\Omega/2\pi = 3.1\text{ MHz}$. 288
- D.7 Amplitude mode after a sweep through the QPT**. Higgs-mode oscillations are apparent in parallel to coarsening in multiple observables. Note that we fit the the oscillation after an initial decay on the time-scales less than a Rabi-cycle. Extracted frequencies are consistent across observables. **a**, Oscillations of the magnetization $\langle n_i \rangle$. **b**, **c**, **d**, Dynamics of the two-point correlation function in position space for various end detunings. 289
- D.8 Ordered-phase quenches and frequency doubling**. **a**, Sweep profile for quenches into the ordered phase. Under a constant Rabi drive at $\Omega/2\pi = 6\text{ MHz}$, the global detuning Δ is ramped from large negative values to large positive values deep into the ordered phase ($\Delta/\Omega = 3.3$). The local detuning δ is quenched off quickly over 50 ns. The global Δ is then also quenched down to varying final detuning values closer to the QPT. **b**, Oscillations of the staggered magnetization in the ordered phase following this preparation sequence. In agreement with the Landau mean-field picture, these oscillations occur around a nonzero value of the order parameter m_s . Note that we fit the the oscillation after an initial decay on the time-scale less than a Rabi-cycle. **c**, Oscillations in the correlation length on top of a growing background. **d**, Ratio between oscillation frequencies extracted from the staggered magnetization (ω_{m_s}) and the correlation length (ω_ξ). In the disordered phase, $\omega_\xi/\omega_{m_s} \approx 2$ while in the ordered phase, this changes to $\omega_\xi/\omega_{m_s} \approx 1$ 290
- D.9 Ratio between oscillation frequencies**. **a**, Experimentally extracted oscillation frequency of the amplitude mode. The frequency is plotted with respect to the distance from the critical point, $\Delta_c = 1.12$ measured in ¹²³. We see an approximate agreement with the theoretically predicted factor, $\omega(-|q|)/\omega(|q|) = 1.9^{61,75,72}$. Deviations close to the phase transition are apparent. **b**, Numerical simulations using MPS methods on 10×10 sites with periodic boundary conditions described in Extended Data Fig. D.5. We again see approximate agreement with the expected ratio, as well as deviations near the numerically determined critical point $\Delta_c = 1.0$ 291

E.1 Analog-digital quantum simulations. **a**, Level diagram showing the relevant atomic transitions in ^{87}Rb . Raman transitions are used for single-qubit gates, the spin-dependent lattice is used for spin-to-position conversion, and the local SLM is used for local light shifts. The clock qubit $|1\rangle = |F=2, m_F=0\rangle$ is excited to the lower-energy $m_J = -1/2$ Rydberg state via a two-photon transition. All Rydberg pair states lie higher in energy than the blockaded subspace. We suppress coupling to the nearby $m_J = +1/2$ Rydberg state, which has $3 \times$ smaller matrix element, by increasing our DC magnetic field to 11.3 G. The intermediate state detuning of +1 GHz and $n = 70$ Rydberg state is used for all 2D experiments. For the 1D chain in Figs. 2,3 we instead use -6.3 GHz and $n = 53$. **b**, Schematic experiment timeline. The total loop takes 16.4 ms for 1D experiments and 24 ms for 2D experiments; in 2D, spin-to-position conversion is performed in two groups due to limited laser power for the AOD tweezers, and image processing and rearrangement are slower for the larger lattice. **c**, Example pulse sequence used for analog-digital quantum simulation. Not all steps are used in all experiments. For the QSL experiments, the first Rydberg π -pulse is omitted when starting from the all- $|0\rangle$ initial state before mapping up. Gaussian Rydberg π -pulses are used to minimize coupling to the other m_J state. The Rydberg Rabi frequency is reduced for the analog dynamics via the 420-nm intensity. For 1D experiments at $n = 53$ and performed in AOD tweezers, the atom spacing is reconfigured from $a = 6.0 \mu\text{m}$ for local Raman rotations to $3.6 \mu\text{m}$ for Rydberg dynamics. In 2D, $a = 6.0 \mu\text{m}$ is used throughout. **d**, Characterization of coherent mapping. As in Fig. 1d, single qubits are arranged at next-nearest neighbor sites in 1D and 2D to emulate the vdW tails in a many-body state. Microstate-dependent phase accumulation from vdW tails occurs during the gap for the Raman π -pulse after the analog dynamics (here, a single-qubit $\pi/2$ -pulse). We observe the decay in Ramsey contrast for increasing vdW tail energy, E_{vdW} , by varying the lattice spacing in 1D (180 ns gap time) and also for increasing gap time in 2D ($2\pi \times 670$ kHz tail energy at $\sqrt{3}a$). Analytical formula includes only phase from vdW tails during the gap time t , given as $\cos^N(\phi/2)$ for N neighboring qubits and $\phi = E_{\text{vdW}}t$. Numerics simulate the entire measurement where the effect of vdW tails during the Rydberg pulses leads to a further reduction in contrast. Both curves are rescaled by the measured contrast of the hyperfine qubit without mapping and no other sources of error are included in the numerics. 295

E.2 **Loss detection in analog quantum simulations.** **a**, Distribution of losses per shot in QSL experiments. Baseline losses exclude the many-body dynamics and follow a Poisson distribution (solid line). The loss increases when performing analog dynamics and a non-Poissonian tail from correlated losses appears. There are 271 atoms in total. **b**, Spatial correlations of loss events. The two-point connected correlator follows a dipolar $1/r^3$ decay, suggesting the presence of Rydberg S - P interactions resulting from an atom, originally in the S state, undergoing blackbody decay to a nearby P state. These long-range interactions facilitate so-called avalanche events^{136,165,53} in which nearby atoms come into resonance with unwanted transitions and are also lost, explaining the tail in **a**. A background offset of $\approx 2.5 \times 10^{-5}$ is subtracted. Data in **a,b** are postselected on the atom being present in the prior image to remove correlations from mid-circuit rearrangement which are created when the reservoir starts to deplete. **c,d** Comparison of local and global loss postselection methods. Data are binned according to the number of missing atoms in each snapshot, including both losses and imperfect initial lattice filling. Postselection on no lost qubits within the local operator approximates the expectation value obtained from shots with a vanishing number of missing atoms across the entire 271-site lattice (dark curves). Randomly converting loss to qubit state $|0\rangle$ or $|1\rangle$ in post-processing strongly affects the calculated expectation value (light curves). We illustrate this for the **c**, three-body dimer populations and **d**, six-body potential energy from the QSL experiments. Dashed lines show the means of the entire dataset with local postselection. There are 285,948 shots in total and 6 shots have zero missing atoms. 298

E.3 Engineering hopping dynamics via Floquet driving. **a**, Interleaved analog and digital operation. Moveable tweezers are used to reconfigure the chain between non-interacting and interacting geometries. Local Raman is used for state preparation and non-destructive readout obtains spin and loss information. **b**, Comparison of Floquet and perturbative dynamics at large static detuning. Starting from a single excitation in a 31-site chain, a quantum walk is realized via Floquet driving (red) or perturbative second-order hopping (gray), tuned by ε or Ω/Δ , respectively. For the same hopping energy (left), Floquet driving exhibits better number conservation (right). Floquet drive parameters are $\gamma + 2\varepsilon = -0.15$ and $\theta = 0$, optimized for number conservation. The effective dynamics at $\varepsilon = 0.1$ has anomalously bad number conservation, possibly from an improperly calibrated resonant detuning for this data. **c**, Verification of the effective Hamiltonian. The measured dynamics, starting from several initial states and measuring in different bases, are fit to numerical simulation of the ansatz blockaded XX Hamiltonian. Plot shows a two-dimensional slice of the cost function, where the dashed lines mark the learned coefficients. **d**, Extracted coefficients of the effective Hamiltonian. The spin-flip terms $g_{X,Y}$ are small, consistent with approximate number conservation. **e**, Domain wall excitations. Increasing $g_{X,Y}$ allows for the creation of domain wall pairs when a spin is flipped in the $\mathbb{Z}_2$ state, and subsequent hopping dynamics of these seeded domain walls washes out the initial order. We use an odd-length chain with $|r\rangle$ initialized at the ends such that domain walls do not propagate from the boundaries. **f**, This process occurs when the chemical potential, tuned via γ , brings the $\mathbb{Z}_2$ state into resonance with the band of domain wall excitations. To a good approximation, γ does not couple to the other terms in the effective Hamiltonian²⁴⁹. **g**, A resonance appears when scanning the chemical potential across the band. This approximately coincides with the presence of a Luttinger liquid phase in the ground state phase diagram^{249,8,430}.

..... 300

- E.4 Engineering ring exchange via Floquet driving.** **a**, Atom image of the array of 15 isolated hexagonal plaquettes. **b**, Continuously tuneable six-body interactions. The speed of the ring-exchange dynamics is varied via the parameter A of the Floquet drive, shown here via the populations on the two sublattices. $A = -1.75$ uses the same data as in Fig. 2h. **c**, Homogeneity of the engineered Hamiltonian. Sublattice dynamics are shown for individual plaquettes, where the several-percent-scale inhomogeneity of the 420-nm and 1013-nm light shifts result in the spatial variations of the static detuning offset. **d**, Numerical simulations of the engineered Floquet dynamics. Good agreement is found when including the measured $T_2^* \approx 1 \mu\text{s}$ in numerics and further multiplying by an empirical exponential decay, $\exp[-\langle n \rangle t / T_1]$, where $\langle n \rangle \approx 2.5$ is the mean Rydberg density and the Rydberg lifetime is $T_1 \approx 30 \mu\text{s}$, dominated by 1013-nm scattering. The T_2^* is dominated by inhomogeneity between plaquettes and shot-to-shot fluctuations of the 1013-nm light shift and therefore is modelled as a Gaussian-distributed detuning shift with zero mean and $2\pi \times 200 \text{ kHz}$ standard deviation. **e**, Stabilization of entangled states. The coherence time of a quasi-adiabatically prepared GHZ state, $\approx (|010101\rangle + |101010\rangle) / \sqrt{2}$, can be extended by Floquet driving, as compared to holding the state at large static detuning, $\Delta/\Omega = 5.9$. We attribute this to the combined effect of an energy gap generated by the ring-exchange interaction and the reduced incoherent error owing to the lower average Rydberg population between stroboscopic times. The engineered Hamiltonian contains non-number-conserving three-body interactions which admix the state $|100100\rangle$ and its cyclic permutations, resulting in oscillations in both the populations and coherences. Decay of populations at static detuning is consistent with Rydberg T_1 313
- E.5 Experimental optimization of QSL state preparation.** **a**, Parameterization of $\Omega(t)$ and $\Delta(t)$ profiles. The detuning sweep is a cubic interpolation between five points (red markers) defined by seven parameters. The Rabi frequency is constant with a 20 ns turn-off time. **b**, Summary table of parameters used in the main text, found using a Bayesian optimizer. Detuning sweep is later clipped between $\pm 20 \text{ MHz}$. Final row is a by-hand modification of optimized profile in row 3. **c**, Results of one round of closed-loop optimization. Metrics used in the multi-objective optimization are shown as a function of trial number. Each trial takes roughly 7 minutes to obtain 2000 Z -basis and 4000 X -basis shots. At the end of the round, we repeat the best 2-4 trials to select the final sweep. For some rounds (not shown), we also included minimization of the magnitude of open length-5 X strings but found this had little effect and was within error bars for all reasonable sweeps. **d**, The measured RK potential energy is positively correlated with both the mean density and RK kinetic energy across different preparation sweeps. **e**, Sensitivity to sweep parameters. By rescaling the time of the $0.8 \mu\text{s}$ sweep (row 3 in **b**), we find the RK energies can change significantly in this out-of-equilibrium regime. The local maxima of $\langle T \rangle$ and $\langle V \rangle$ coincide. For slower sweeps, $\langle T \rangle$ decreases faster than $\langle V \rangle$ which we explore numerically in ED Figure 6. 319

E.6 Preparation of quantum spin lakes in numerical simulations. **a**, Schematic of the energy spectrum of the Rydberg Hamiltonian on the kagome lattice as a function of detuning. For $\Delta \gg \Omega$, monomer excitations are strongly gapped whereas vdW tails and perturbative couplings lead to a small splitting between dimer coverings. Hemidiabatic preparation^{369,162} corresponds to sweeping through the phase transition at a rate (dashed line) adiabatic relative to monomer excitations but sudden relative to the low-energy splitting. **b**, To confirm that the final sweep rate in Fig. 8.4d corresponds to the hemidiabatic regime, we modify the $0.8\mu\text{s}$ detuning profile by accelerating the sweep through the phase transition (gray dashed line) by a multiplicative rate factor after the inflection point (dotted line). **c**, Evidence for hemidiabatic preparation of quantum spin lakes. We use exact diagonalization to simulate a 3×3 unit cell lattice with periodic boundary conditions along each lattice vector. Starting from the ground state at $t = 0$, we evolve under the modified sweep from **b** and plot the final potential and kinetic energies for both the PXP approximation and the full simulation of the long-range vdW interaction. For large rate factors, we see that accelerating the sweep decreases these RK energies due to the large number of monomers in the final state. For small rate factors, although the PXP model approaches the RK state ($\langle T \rangle = \langle V \rangle \approx 0.19$ for this system size), the inclusion of vdW tails changes the relative phases and amplitudes between dimer configurations which reduces both T and V . Rather, the optimal rate with tails corresponds to the hemidiabatic regime where the sweep is slow enough to remove monomers but fast enough that vdW tails dephasing is limited. As expected, the experimentally optimized sweep corresponds to rates near the maxima of T and V . **d**, Effect of disorder on unprojected energies. To account for sources of disorder in the experiment, we simulate the $0.8\mu\text{s}$ sweep starting from the empty product state $|1\rangle^{\otimes N}$ 100 times, randomly sampling on-site detuning shifts, global detuning shifts, in-plane positional shifts, and out-of-plane positional shifts from Gaussians with widths of $2\pi \times 0.043\text{ MHz}$, $2\pi \times 0.2\text{ MHz}$, $0.02a$, and $0.08a$, respectively. We plot V , T , and T after 100 ns of idle evolution under the vdW interaction to simulate the gap for the Raman pulse during mapping down. We attribute the dip in these energies at the end of the sweep to finite-size effects which differ from the experiment. **e**, Effect of same disorder on projected energies, where bonds around the loop are projected onto no dimer (see Fig. 5c). 323

E.7 **Dimer configurations in the Rokhsar-Kivelson state.** **a**, Comparison between the standard and weakly modified detuning sweeps used in the main text. Zoom-in highlights the scale of the modification. **b**, Probability distribution of dimer configurations in isolated subsystems using the modified sweep. The experimental distribution follows a thermal distribution with a fitted inverse temperature $\beta = 9.0(2)/V_0$, where V_0 is the nearest-neighbor interaction strength (light gray bars). The same data is used for the inset of Fig. 5a. The dimer configurations are explicitly shown for two subsystem sizes. The vdW tail energy is larger for configurations with more flippable plaquettes within the subsystem, shaded in red, which explains the reduced probability of these configurations. **c**, Probability distribution within a 2×1 subsystem for varying number of dimers surrounding the subsystem, using the standard $0.8 \mu\text{s}$ sweep. The uniform distribution in Fig. 5a corresponds to 0 outer dimers and the single configuration with 8 outer dimers is omitted. The outer dimers restrict the state of the lattice surrounding the subsystem to be in one of a finite number of allowed dimer configuration which satisfy the boundary conditions, such that even the perfect RK state has a non-uniform probability distribution within subsystems in this analysis. The RK distribution is shown for a 12×6 lattice with periodic boundary conditions obtained from Monte Carlo simulations; we find a qualitatively similar distribution for other lattice dimensions although the precise amplitudes differ slightly (not shown). We find experiment is in good agreement with this qualitative distribution. 324

E.8 Loops and strings in the quantum spin liquid. **a**, Time-dependent matrix product state (MPS) simulations of $\langle V \rangle$ and $\langle T \rangle$, compared to experiment. For the final state, MPS predicts $\langle T \rangle / \langle V \rangle = 0.92$. The lattice used for the numerics is three plaquettes tall and infinite in the horizontal direction, with the van der Waals interaction truncated at $\sqrt{7}a$ and a bond dimension of $\chi = 256$. **b**, Including the phase accumulation from vdW tails during the 100 ns mapping down gap reduces the expectation value in MPS numerics to $\langle T \rangle / \langle V \rangle = 0.76$, compared to $\langle T \rangle / \langle V \rangle = 0.47(1)$ in experiment; numerics cannot account for the Rydberg π -pulse in the three-level system. In experiment, plaquettes on the boundary (bdy) have higher $\langle T \rangle$ than in the bulk, where van der Waals tails are larger, but similar $\langle V \rangle$. For this plot only, the four corner plaquettes are excluded. **c**, Spatial plot shows the energy ratio in experiment. **d**, Evolution of experimental RK energies after state preparation, where the Rabi drive is turned off during the hold time. Decay of $\langle V \rangle$ is consistent with Rydberg T_1 error, while $\langle T \rangle$ decays faster as the different dimer configurations accumulate relative phases but the populations remain approximately unchanged. **e**, Perimeter-law decay of closed X and Z loops. Data for X loops are the same as in Fig. 5f. **f**, Normalization of open X strings. See Fig. 6c for the corresponding definition for Z strings. **g**, Open X strings are small and decay to zero with increasing string length, consistent with the state being described by a deconfined field theory^{386,429}. **h**, Dynamics of open strings. X strings, which create monomers at their endpoints, are finite when there is a high density of monomers. Z strings are trivially +1 for the initial product state. We show open strings with length 3 which are normalized by the length-6 closed loops that enclose a single plaquette. **i**, Effect of defects on the dimer-string correlator. Postselection on no monomers or double-dimers inside the closed loop (such that it is +1) slightly improves agreement between experiment and the RK prediction from Monte Carlo simulations. . . . 328

- E.9 Diagonal correlators in the emergent $U(1)$ gauge theory.** **a**, Schematic of the rectified dimer-dimer correlator between two vertices. The sum involves nine (unconnected) dimer-dimer correlation functions, as highlighted, and the colors indicate the relevant phases, determined by the sublattices (a,b,c) of the neighboring plaquette pair. Here, $\omega = e^{2\pi i/3}$. **b**, Rectified dimer-dimer correlator in experiment. **c**, Comparison to MPS simulations of the experimental preparation protocol and to Monte Carlo simulations of the RK state on a 36×6 lattice with periodic boundary conditions. Inset shows real and imaginary components in MPS. MPS numerics use the PXP approximation, a bond dimension of $\chi = 350$, and a four-plaquette-tall lattice that is infinite in the horizontal direction. **d**, Schematic of rectified dimer-string correlator between two vertices. The sum involves nine unrectified dimer-string correlators between plaquette centers and the colors indicate the plaquette sublattices that define the phases. **e**, Rectified dimer-string correlator in experiment. While $\tilde{C}_s$ is a real quantity for a translationally invariant system, here we find an appreciable imaginary component (inset). **f**, Comparison of the dimer-string correlator to MPS and Monte Carlo simulations. Inset shows real and imaginary parts in MPS numerics, with the experimental data shown as faint markers. In **c** and **f** the experimental correlations are shown along the horizontal direction only to compare to numerics, while in **b** and **e** they are radially averaged over the lattice, excluding vertices on the boundary. 332
- E.10 Finite-size scaling in numerics and experiment.** **a,b** Finite-size scaling of rectified correlators in Monte Carlo simulations of the RK state. The correlators converge towards the predicted $1/r^2$ and $1/r^{1/2}$ decays at large system sizes. Dashed lines are guides to the eye showing exponential decays with correlation length $\xi(L_y) \propto L_y$. Experiment is well-captured by the 36×6 data (ED Fig. 9). **c**, Relation between the magnitudes of the rectified dimer-dimer and dimer-string correlators in Monte Carlo simulations. The ratio of 4 from $U(1)$ gauge theory holds for all system sizes until the dimer-dimer correlator plateaus. **d**, Effect of topological sector in Monte Carlo with periodic boundary conditions. Within the $(+1, +1)$ topological sector, defined as $+1$ eigenvalues of the two non-contractible Z -strings, both correlators are symmetric about half-length (dashed line), as enforced by translation and reflection symmetries. Averaging over all four topological sectors suppresses the correlators overall, and leads to a distance-dependent decay of the dimer-string correlator beyond half-length. In all other plots, correlators are averaged over sectors. Data are for a 36×36 lattice. **e**, Finite-size scaling in experiment. Upon varying the lattice height from $L_y = 3$ to $L_y = 6$, at fixed width of $L_x = 13$, we find the rectified dimer-string correlator is largely independent of system size while the rectified dimer-dimer decays until the correlations plateau. This plateau arises from explicit breaking of the lattice's $\mathbb{Z}_3$ symmetry and decreases in magnitude as the lattice height increases. For all other data in this work, $L_y = 6$. **f**, The field theory prediction (solid line) is in good agreement with the data at short distances for all system sizes until the plateau in the dimer-dimer correlations is reached. . . 336

F.1	Architecture for high-fidelity gates and circuits. (A) Atom array layout utilized in this work, with a large reservoir zone for fast cycle rates, an entangling zone with gate sites demarcated (dark sites are the subset used for randomized benchmarking), and additional sites in the entangling zone used for state-selective readout ⁴² . The top five rows of sites are used for initial loading and rearrangement of atoms into the reservoir. (B) Level diagram for Rydberg excitation to $n = 53$ Rydberg state $ r\rangle$. The lasers used for other aspects of the architecture, including for Raman single-qubit gates and state-selective readout, are described in our previous work ⁴²	339
F.2	Readout methods and spatial homogeneity for echo RB. Echo RB curves for the data plotted in Fig. 8.2c. Measured return probability as a function of number of CZ gates (A) with and without a single-qubit π pulse before state-selective readout (B) with pushout versus state-selective readout (SSR). (C) CZ fidelity across the gate sites used randomized benchmarking.	343
F.3	Symmetric stabilizer benchmarking. (A) Circuit used for symmetric stabilizer benchmarking (SSB); figure adapted from Ref. ⁴²⁵ . (B) Optimization of single-particle phase ϕ after the CZ gate, here looking at the return probability after 12 CZ gates. (C) Comparison of SSB and echo RB approaches, taken interspersed with each other to avoid any potential bias from experimental drift. Measured fidelities are within error bars of each other.	345
F.4	Understanding leakage in CZ gates. Measurement of the buildup of population in $F = 1$ during the echo RB sequence. At the end of the RB sequence, population in $F = 2$ is expelled from the trap with near-resonant pushout light. Then, a single-qubit π pulse transfers $ 0\rangle$ to $F = 2$, and we perform state-selective readout. Plotted is the residual population in $F = 1$ as a function of the number of CZ gates, consisting of buildup of leakage population. The green curve is for the raw data, and the tan curve is for postselecting on the other atom being in the correct state at the end of the sequence.	348
F.5	Understanding sensitivity of CZ fidelity to drifts. Sensitivity of the measured CZ gate fidelity to Rydberg beam (A-B) powers and (C-D) positions. Red points are loss-postselected, and blue points are raw CZ fidelities. Faint colored curves are quadratic fits for beam powers and quartic fits for beam positions, and the black curves are from numerical simulations with offsets added as the only free parameter (see text for more details). (E) Positions of Rydberg beams on camera used for tracking and active feedback, averaged between each experimental repetition.	353

F.6	Path to higher CZ fidelity. Numerical projection of the CZ fidelity for improved parameters. (A) CZ error as a function of Rabi frequency (relative to current Rabi frequency Ω_0), with and without contributions of the other Rydberg m_J state $ r'\rangle$. (B) CZ error as a function of detuning, relative to the current detuning Δ_0 , where the optimum balances contributions from Rydberg T_2^* and intermediate state scattering. The other m_J level is not included for this plot. (C) CZ error as a function of Δ and Ω , compared to our current Rabi frequency Ω_0 and intermediate state detuning Δ_0 (color-bar is log-scale). Dashed lines denote equal beam intensity, for $1\times$, $2\times$, and $3\times$ current intensity I_0 used in this work. The effect of coupling to $ r'\rangle$ prevents going to much higher Rabi frequencies. (D) Projection of CZ fidelity improvements achievable with higher laser intensities, as a function of Δ and Ω . For this plot, we consider a configuration with $2\times$ improved 1015-nm T_2^* and no coupling to $ r'\rangle$. The white stars mark the optimal configuration that minimizes CZ error for each intensity, with the associated CZ error labeled.	354
F.7	Numerical simulation of fidelity in cluster state circuit. Comparison between fitted fidelity of the decay in global return probability and the true CZ fidelity put into the numerical simulation, for the cluster state circuit in Fig. 8.4. Black lines denote $y = x$	356
F.8	Nonlocal circuit setup and supplementary data. (A) Atom images for the CZ gate locations for the nonlocal circuit (cluster state circuit repeats the first two nearest-neighbor moves). (B) Sampling from 10-qubit circuit, showing first 30 unordered outcomes. (C) Distribution of atoms lost per shot (out of 20 atoms total), for the 670,890 shots for the cluster state circuit and the 26,140 shots for the nonlocal circuit. The solid curves are a sum of Poissonian fits of the loss distributions for each different number of CZ gates.	358
F.9	Theoretical analysis of nonlocal circuits. (A) Entanglement entropy as a function of circuit depth, comparing our nonlocal circuit to applying only nearest-neighbor gates at all depths. (B) Entanglement entropy as a function of depth, for different rotation angles θ for the single-qubit gates in the circuit, where we use $\theta = 0.25$ corresponding to a 45-degree X rotation. (C) Distribution of ordered outcome probabilities as a function of nonlocal connectivity, compared to the Porter-Thomas distribution (PT). (D) KL divergence from $PT(x)$ as a function of nonlocal connectivity. (E) Distribution of eigenvalues of the reduced density matrices for subsystems on the boundary, compared to the Marchenko-Pastur distribution (black curve) expected for Haar-random pure states, for nonlocal connectivity of 11.	360

TO MY FAMILY AND FRIENDS WHO SUPPORTED ME ON THIS JOURNEY.

Acknowledgments

MY PHD EXPERIENCE was made very special by the people I got to work with, and I feel extremely lucky to have been able to work with many thoughtful, talented, and enthusiastic scientists.

I would first like to thank my advisor, Prof. Misha Lukin. Working with Misha is a one-of-a-kind experience. He is extremely skilled at bringing together a great group of people, and then pioneering a wide range of research directions. Misha always pushed us to think deeply about our results, and reflect on how to make our research even more impactful. His ability to give his students freedom to pursue new ideas while keeping things on track was also very essential to our research. I will always remember how we were planning to change out our vacuum chamber in the middle of my PhD, and Misha suggested we first just try to improve the gate fidelity, and this decision then snowballed into a series of very cool science projects. This insight was just one of many that would change the trajectory of our research for the better. Misha has really shaped how I think about scientific research, and being a part of his research group has been transformative.

I would also like to thank our two other PIs, Prof. Vladan Vuletić and Prof. Markus Greiner, who brought such complementary skill sets and research vision. Vladan was very kind during my first year helping me figure out how to get into research during the Covid-19 pandemic, including letting me research some Rydberg physics with his group. Vladan has been an invaluable resource whenever we have a question about laser cooling, atomic physics, or many other topics. My discussions with Vladan and his group about different atomic physics techniques have been a lot of fun and very scientifically stimulating. Markus is always very enthusiastic to come up with and discuss different new optics ideas, and has shown me the real value of making time to brainstorm creative ideas. I always look forward to hearing his latest thoughts on combining different optical elements in new ways. Additionally, I would like to thank Prof. Susanne Yelin for her valuable insights, feedback, and collaboration on projects during my PhD, as well as for being a great member of my committee.

I was fortunate to work with a truly stellar, large group of experimental and theoretical collaborators during my PhD:

Sepehr Ebadi's passion for science was infectious, and he taught me a lot about the experiment early on. He brought an especially deep love for the technical aspects of our work, and always ensured that our posts were well-tightened (!). He left a lasting mark on how I think about AMO

physics, as I often think back to different lessons he taught me years ago, from high-power optics and choices of lasers to his approaches for leading science projects. We also had fun playing chess and joking around in lab.

Dolev Bluvstein's enthusiasm and passion has been a real driving force for our science. His leadership ability has always been very inspiring to me, as he managed to forge new research directions across a large team of experimental and theoretical collaborators. Even more important was the energy and intention he put into caring about the wellbeing and scientific development of the community around him, ranging from encouraging fun boba breaks to providing invaluable mentorship to many people. We had a blast in lab together learning how to run complex quantum circuits with atoms and developing as scientists and people. Our exceptional teamwork and joy in working together really propelled our experiments in unexpected ways, and I couldn't have wished for a better friend and colleague to take this journey with.

Tom Manovitz was really a joy for me to work with, both because he has a lot of technical experience and breadth of scientific knowledge and, more importantly, because he is a very thoughtful and caring person. I had a lot of fun working with him closely on CCZ gates, experiment upgrades, and digital simulation experiments, and I cherish many memories of learning from him about life and science. His research group is very lucky to have such a kind and talented leader, and I look forward to hearing updates on what they get up to!

Marcin Kalinowski was a close collaborator of mine on most of my PhD projects, and it was really an amazing experience to work closely with him. It's very unique how Marcin is both a master at numerics and theory, and is equally, if not more, passionate about aligning optics. During our collaborations, I would ask him a ton of theory questions that he would patiently answer, and in turn he enthusiastically asked me what all the different tools were doing in the lab. On a personal note, our weekend brunches also reminded me of the bigger picture as we went through our PhDs. I'm very excited to continue our friendship into our next career endeavors!

Harry Zhou was initially my "Lukin group mentor" at the start of my PhD, and I was excited when I found out that I would get to work with him in the middle of my PhD. His scientific wisdom and passion for error correction were a great influence on our team. He also built the high-power 1013-nm setup which has served our Rydberg gates well ever since, and has been one of the most stable setups in the lab! Whenever I talk to Harry about anything from his ideas about science to his advice on my life decisions, I've always come away with a deeper and more nuanced understanding of the world, and I really value his continued mentorship and friendship.

Sasha Geim is an exceptional scientist, and an equally incredible teammate. Whenever I'm stuck on a problem, I know that talking to Sasha will help, whether it be how to debug the experiment or lead a science project. I've learned a lot seeing her become an expert on many topics over the last few years, ranging from optics to many-body Rydberg physics to error correction to running circuits with atoms. It's recently been inspiring to see how she led the Floquet project and turned it into something next level, and how she is defining the next frontier in quantum error correction experiments with atoms. I'm looking forward to continued fun science brainstorming with Sasha in the coming years and seeing where she takes the experiment!

Sophie Li is an exceptionally talented experimentalist, with a great attitude toward jumping in

and making progress in the lab, and she has been such a positive part of our team. I was really impressed with how she led the coarsening project to success through sheer willpower and in-depth analysis of the domains. Recently I've had a great time improving the Rydberg gate with her, from seeing how her working enclosure of the beam paths immediately improved our Rydberg coherence to making the fidelity continue to increase with each new upgrade to designing and running some fun high-fidelity circuits together. I'm very excited to stay in close touch and see the science she does over the next couple of years.

Muqing Xu joined the lab later on in my PhD and has been an amazing postdoc to have in our group. Muqing's style of work complements mine well; he often takes a more methodical approach to problems. It's been inspiring seeing how he quickly develops a deep understanding of new topics, and has the persistence to figure out the answer even when it's very challenging. This last year of working together closely has been a blast, debugging the experiment and improving the two-qubit gate fidelity. I hope to continue to learn from and collaborate with Muqing in the future.

I had a lot of fun times in lab together with this group, from blasting Boiler Room sets and ABBA while taking high-quality data to brainstorming wacky atomic physics ideas. We also had plenty of good times outside of lab, and it was always great to talk about life and physics together. Highlights include horoscopes at Grendel's, Black Sheep bagels, contemplative walks by the river, Atom Family BBQs and lunches, and brainstorming acronyms for atom array architectures.

In the Lukin Group and CUA, I found a collaborative community of passionate scientists. Being able to work with and learn from so many smart and thoughtful theorists – including Maddie Cain, Nishad Maskara, Pablo Bonilla Ataides, Rahul Sahay, Nazlı Uğur Köylüoğlu, Katherine van Kirk, Christian Kokail, Johannes Feldmeier, Jordan Cotler, and many others – was a great privilege. Some of my favorite times were spent getting lunch together and learning about entirely new areas of science.

I also got to learn from a large group of very talented experimentalists throughout my PhD, including Atom Array 2 (Jinen Guo, Tout Wang, Pasha Stroganov, Elias Trapp, Luke Stewart, Allen Chiu, Mohamed Abobeih, Lisa Peters, Nicholas Lyu, Renhao Tao, Sebastian Geier), Atom Cavity (Paloma Ocola, Polnop Samutpraphoot, Elmer Guardado-Sanchez, Brandon Grinkemeyer, Tamara Đorđević, Ivana Dimitrova, Offek Tziperman, Eirini Mandopoulou, Matthew Bilotta, Michel Tào, Danilo Shchepanovich, Andrei Ruskuc, Shivam Mundhra), the Vuletić group (Wenchao Xu, Tamara Sumarac, Emily Qiu, Matthew Peters, Sergio Cantu, Aditya Venkatramani), the previous generation of the Atom Array experiment (Harry Levine, Alex Keesling, Giulia Semeghini, Ahmed Omran), and many other members of the community. I'm truly very thankful to everyone.

I'm also grateful for the many kind administrative staff members I've gotten help from during my PhD (Lisa, Sam, Elaine, Adam, Jacob, Silke), the custodial staff especially for managing all the trash and boxes from our lab, and the building staff for help with many things including the temperature stability of our lab. I also gratefully acknowledge financial support from the NDSEG Fellowship.

I'm very thankful to the many people who got me interested in physics research, starting with my

amazing high-school physics teacher, Dr. Alvaro Salcedo, who taught me that there is no such thing as a free lunch in physics! At Stanford, Prof. Monika Schleier-Smith was a key reason I fell in love with AMO physics and quantum science, and I had many wonderful mentors in her lab (Jacob Hines, Ogi Markovic, Emily Davis, Greg Bentsen, Tracy Li, Avikar Periwal, Eric Cooper). I was also fortunate to spend some time in Munich, Germany at MPQ before starting my PhD, and am thankful to Immanuel Bloch, Christian Gross, Simon Hollerith, David Wei, Kritsana Srakaew, Jun Rui, and Toni Rubio-Abadal for making me feel welcome and for teaching me a lot about building and operating quantum gas microscope experiments.

Finally, I would like to thank my family and friends for their support throughout my physics journey, and for bringing me a lot of joy and welcome distractions. Too many people were supportive during my PhD to list everyone here, but a few highlights were the legendary Lexingtonians and Montrosians who distracted me from PhD stress with pickles and pop the balloon (Bridget, Giacomo, Sara, Michele, Deepali, Daniele, and of course their pets Paisley and Devi for keeping me on task when I was working at home!), undergraduate physics friends (Deepti, Sam, Josie, Nicel, Scott, Manisha, Carlos), our Lukin/Geiner/Semeghini lab qual-exam study group, CUA pizza talk committee, catching up with high school friends in Boston, playing cello in the GSAS Student Center Orchestra with Yi, and visiting with Olivia in New York and Boston.

Doing a PhD is by no means easy, and my family's unending support and love was essential to making it through the challenging times, as well as celebrating the successes. My parents have been endlessly supportive in my pursuit of science, even if I still haven't quite managed to explain what I actually do. Hanging out with my two sisters over holidays, as well as our sibling calls, was very restorative and helped remind me about my broader life outside of work. I'm also thankful for my family's patience with me when I was working hard and laser focused in the PhD – during one holiday break, my mom asked if I was dating someone new because I kept sneaking out at night for long phone calls, when really I was just brainstorming with Dolev as we were finishing up a paper. I'd also like to thank my extended family – my many uncles, aunts, and cousins – for their love and support. I don't know if I could've made it to the end of the PhD without so much love from friends and family.

Author Contributions

The work presented in this thesis have been previously published and co-authored, as outlined below.

Chapter 2 is based on:

Simon J. Evered*, Dolev Bluvstein*, Marcin Kalinowski*, Sepehr Ebadi, Tom Manovitz, Hengyun Zhou, Sophie H. Li, Alexandra A. Geim, Tout T. Wang, Nishad Maskara, Harry Levine, Giulia Semeghini, Markus Greiner, Vladan Vuletić, Mikhail D. Lukin. “High-fidelity parallel entangling gates on a neutral atom quantum computer,” *Nature* 622, 268–272 (2023).

Chapter 3 and Appendix A are based on:

Simon J. Evered*, Marcin Kalinowski*, Alexandra A. Geim, Tom Manovitz, Dolev Bluvstein, Sophie H. Li, Nishad Maskara, Hengyun Zhou, Sepehr Ebadi, Muqing Xu, Joseph Campo, Madelyn Cain, Stefan Ostermann, Susanne F. Yelin, Subir Sachdev, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. “Probing the Kitaev honeycomb model on a neutral-atom quantum computer,” *Nature* 645, 341–347 (2025).

Chapter 4 and Appendix B are based on:

Dolev Bluvstein, **Simon J. Evered**, Alexandra A. Geim, Sophie H. Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J. Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J. Gullans, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. “Logical quantum processor based on reconfigurable atom arrays,” *Nature* 626, 58–65 (2023).

Chapter 5 and Appendices C and G are based on:

Dolev Bluvstein*, Alexandra A. Geim*, Sophie H. Li, **Simon J. Evered**, J. Pablo Bonilla Ataides, Gefen Baranes, Andi Gu, Tom Manovitz, Muqing Xu, Marcin Kalinowski, Shayan Majidy, Christian Kokail, Nishad Maskara, Elias C. Trapp, Luke M. Stewart, Simon Hollerith, Hengyun Zhou, Michael J. Gullans, Susanne F. Yelin, Markus Greiner, Vladan Vuletić, Madelyn Cain, and Mikhail D. Lukin. “A fault-tolerant neutral-atom architecture for universal quantum computation,” *Nature* 649, 39–46 (2025).

Chapter 6 and Appendix D are based on:

Tom Manovitz*, Sophie H. Li*, Sepehr Ebadi*, Rhine Samajdar, Alexandra A. Geim, **Simon J. Ev-**

ered, Dolev Bluvstein, Hengyun Zhou, Nazli Ugur Koçluoglu, Johannes Feldmeier, Pavel E. Dolgirev, Nishad Maskara, Marcin Kalinowski, Subir Sachdev, David A. Huse, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. “Quantum coarsening and collective dynamics on a programmable simulator,” *Nature* 638, 86–92 (2025).

Chapter 7 and Appendix E are based on:

Alexandra A. Geim, Nazli Ugur Koçluoglu, **Simon J. Evered**, Rahul Sahay, Sophie H. Li, Muqing Xu, Dolev Bluvstein, Nik O. Gjonbalaj, Nishad Maskara, Marcin Kalinowski, Tom Manovitz, Ruben Verresen, Susanne F. Yelin, Johannes Feldmeier, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. “Engineering quantum criticality and dynamics on an analog-digital simulator,” arXiv:2602.18555 (2026).

Chapter 8 and Appendix F are based on:

Simon J. Evered^{*}, Muqing Xu^{*}, Sophie H. Li^{*}, Alexandra A. Geim, J. Pablo Bonilla Ataides, Marcin Kalinowski, Dolev Bluvstein, Nishad Maskara, Christian Kokail, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. “High-fidelity entangling gates and nonlocal circuits with neutral atoms,” arXiv:2604.25987 (2026).

^{*}These authors contributed equally to this work.



Introduction

0.1 THE PROMISE AND CHALLENGE OF QUANTUM INFORMATION PROCESSING

The realization of large-scale, fault-tolerant quantum computers will have transformative impacts on science and society³³⁸. As quantum simulators, these devices can help us learn more about nature, probing complex highly-entangled quantum systems relevant to a wide variety of physical and chemical systems, including strongly-correlated materials and models for gravity^{198,367}. Ad-




Figure 1: Key ingredients of quantum information processing. Superposition and entanglement are the core underlying ingredients of quantum information processors. A single qubit can be held in a superposition of states $|0\rangle$ and $|1\rangle$, and many qubits can be entangled to access the 2^n dimensional Hilbert space.

ditionally, quantum information processors can be used for quantum sensing and metrology, such as sensing astrophysical signals more accurately leveraging long-distance entanglement⁴⁰⁹, or performing algorithms to improve sensing accuracy¹¹. Beyond their utility as instruments for scientific discovery, such machines have strong implications for society, for example in cryptography^{400,158}.

Motivated by the promise of quantum computers, there has been a significant effort to build large-scale quantum computers across a variety of physical platforms. Central to the promise of quantum computers is the exponentially growing Hilbert space. For a single qubit, the system is described by a two-dimensional Hilbert space, accessed through superpositions of the basis states $|0\rangle$ and $|1\rangle$. For n qubits, the Hilbert space becomes 2^n -dimensional, and can be fully accessed by entangling the qubits with each other (Fig. 1). However, quantum computations face the fundamental issue that all quantum systems are susceptible to sources of decoherence. In particular, there are many types of errors that can corrupt quantum computations, as shown in Fig. 2. Coherent errors, for example under-rotation in a single-qubit gate, can cause evolution to become more and more imprecise over longer durations. Additionally, there are incoherent errors, for example scattering, which can convert one spin state into the other. At the hardware level, there can also be issues with qubit loss or leakage into other states, as well as heating of the qubits throughout the computation which can lead to higher error rates. For many large-scale applications requiring very precise circuit execution, error rates must also be incredibly low, for example at the $< 10^{-10}$ level. Despite




Figure 2: Examples of errors that can occur in quantum systems. Suppressing errors is the central challenge of useful quantum computation. Coherent errors can build up and lead to the incorrect many-body quantum evolution. Additionally, incoherent errors, loss, leakage, and atom heating can corrupt the computation.

tremendous progress across several platforms, physical qubit error rates will likely never reach these low levels due to fundamental sources of errors, requiring the use of quantum error correction.

0.2 QUANTUM COMPUTING BASED ON ATOM ARRAYS

Developing a quantum computing architecture involves many complexities and constraints. One must think about scalability to large qubit numbers, while maintaining high-fidelity and programmable control. Meanwhile, additional ingredients such as non-local connectivity can reduce the requirements for large-scale algorithms.

In this thesis, we will use arrays of atoms, which have emerged as a versatile platform for quantum computation^{342,41,288}, simulation^{63,123}, networking^{100,194,181}, and metrology^{455,138,73}. In this approach, atoms are cooled to microkelvin temperatures and individually trapped in tightly focused beams of light called optical tweezers. These optical traps can be dynamically rearranged, allowing the atoms to be reconfigured into arbitrary two-dimensional geometries. Quantum information is stored in long-lived ground-state hyperfine qubits, in this work using rubidium-87 atoms. Atoms are excited to highly excited atomic states called Rydberg states in order to interact. These techniques have rapidly advanced from trapping a single atom³⁷⁸ to the precise control of up to thousands of individually-trapped atomic qubits^{288,90}. Several features make this platform uniquely powerful for quantum information processing: atomic qubits are inherently identical, non-local

connectivity can be achieved by physically moving atoms during a computation, and large numbers of atoms can be efficiently trapped and manipulated.

0.3 HIGH-FIDELITY, PROGRAMMABLE QUANTUM CIRCUITS WITH ATOMS

In this thesis, we demonstrate methods for reducing the error per operation in a neutral-atom quantum computer. Until very recently, neutral-atom quantum computing faced very high error rates compared to other quantum architectures. In 2019, there was a breakthrough that realized two-qubit gates with 97.5% fidelity through the Levine-Pichler gate. In Chapter 2, we will discuss how we used an improved gate scheme and implemented technical upgrades to achieve a gate fidelity of 99.5%. This made the error rates of neutral atoms comparable to those of other platforms and brought them below the error-correction thresholds necessary for scalable quantum computation. In Chapter 8 we will describe further improvements to achieve raw fidelities of 99.85% which improve to 99.94% upon postselection on neither atom being lost.

Meanwhile, we will demonstrate fully programmable quantum circuits by implementing a variety of new tools for local single-qubit gates, tunable entangling gates, non-local connectivity, mid-circuit readout and feedforward, state-selective readout, and qubit reuse. In doing so, we will show how to realize high-fidelity and programmable digital quantum computations and simulations with atom arrays.

0.4 DIGITAL QUANTUM SIMULATION

Until recently, quantum simulations with atom arrays were realized through an “analog” modality. An alternative method is digital or gate-based quantum simulation, in which quantum dynamics are decomposed into discrete operations such as single-qubit and two-qubit gates. While analog simulations are powerful for quantum dynamics that are inherent to the platform, such as PXP or




Figure 3: Analog simulation, digital simulation, and error correction. In near-term devices, analog simulation can be powerful for realizing native quantum dynamics. Meanwhile, digital devices can offer highly programmable quantum dynamics and observables, as well as the implementation of algorithms. However, these devices suffer from intrinsic error rates, such that we need quantum error correction to eventually reach large-scale, precise quantum computation and simulations.

Ising dynamics in the case of Rydberg arrays, digital approaches could offer new avenues for enhancing the programmability for (i) state preparation, (ii) Hamiltonian evolution, and (iii) the types of observables that are accessible. Additionally, for applications requiring very low error rates (e.g., $< 10^{-10}$), such as a very precise quantum chemistry calculation, future fault-tolerant quantum processors can extend digital quantum simulations to incorporate quantum error correction. In this way, studying digital quantum simulation on current noisy devices allows us to begin exploring algorithms relevant for these future low-error regimes.

Digital quantum simulations of many different types of models have been actively explored in superconducting qubit and trapped ion processors. Recent progress in digital quantum processors based on atom arrays has pushed the scale, level of control, and fidelity such that realizing such digital quantum simulations has become possible. Atom arrays enable several new avenues. For one, parallel coherent transport of atoms opens opportunities for realizing dynamics involving non-local connectivity. Additionally, multi-qubit gates can be realized through Rydberg excitation of several atoms placed within the Rydberg blockade radius.




Figure 4: Core ingredients of quantum error correction. Quantum information is stored redundantly across a group of physical qubits. Local errors can be corrected without disturbing the logical qubit information. Such logical qubits can be improved by increasing the code size, where we expect exponential error suppression with code distance below a characteristic physical error threshold.

0.5 QUANTUM ERROR CORRECTION

In this thesis, we will additionally explore how to use atom arrays to perform quantum error correction, with the goal of charting a path toward large-scale fault-tolerant quantum processors. The central challenge to realizing precise quantum computation and simulation is suppressing errors³³⁸. Remarkably, low-error quantum computation can be achieved by delocalizing information across many qubits to form “logical qubits”, in which a local error will not corrupt the encoded information^{399,410}. Such quantum error correction provides a path to arbitrarily precise quantum simulations, but introduces significant complexity and overheads. For instance, large-scale applications are typically thought to require tens of thousands to millions of qubits and days-long computations³⁹, constituting orders-of-magnitude improvements from the current state-of-the-art. Therefore, developing a scalable quantum architecture with efficient ways to control many qubits is fundamental to useful quantum processing.

Fig. 5 shows a classic example of a quantum error-correcting code: the $[[7,1,3]]$ Steane code (also




Figure 5: Example: Steane code. A well-known example of an error-correcting code, the Steane code encodes a single distance-3 logical qubit with 7 physical qubits (e.g., atoms), denoted $[[7,1,3]]$ in $[[n,k,d]]$ notation. Image taken from the [Error Correction Zoo](#).

known as the color code for its colored stabilizers). That is, if we consider any single error on a given qubit in the code, then based on how the stabilizers “light up”, we are able to deduce what error occurred. For example, consider a bit-flip X-type error on qubit 5. Then the stabilizers whose sign is flipped will be the red and blue Z stabilizers. You can work out for all other X or Z errors on a given qubit which corresponding stabilizers will flip in sign, and you will see that they are all unique. This tells us that there is a one-to-one mapping from the stabilizer measurements to the corresponding single error. However, if we now have two errors in the qubit array, we won’t have this nice one-to-one mapping, but we will at least know that an error occurred. These observations tell us that the code distance for the Steane code is $d = 3$, meaning that the code can detect up to $d - 1 = 2$ errors at a time, but can correct up to 1 error. In order to measure the stabilizers of an error-correcting code, we will use additional “ancilla qubits”, as shown in Fig. 6b, which we entangle with the data qubits forming the code and subsequently measure to extract the stabilizer information.

Whereas the ultimate goal is to build a large-scale error-corrected quantum computer, exploring applications of quantum computers in the near term and implementing the first error-corrected




Figure 6: Surface code and implementation with atoms. **a**, Surface code, with corresponding X and Z type stabilizers. **b**, Implementation of a $d = 5$ surface code with atoms, including data and ancilla qubits.

algorithms can give insights that guide this path. Therefore, in this thesis, we explore a wide range of settings in quantum simulation and computation, including digital quantum simulations with physical qubits, analog-digital quantum simulations, first studies of error-corrected algorithms and quantum simulations, and probing the core ingredients for fault-tolerant quantum computing architectures. We find such wide-ranging explorations to be extremely fruitful, enabling many scientific explorations, highlighting new directions for algorithmic design, elucidating aspects for how to think about atomic processing architectures, and helping us chart a path toward large-scale fault-tolerant atomic processors.

0.6 OVERVIEW OF THE THESIS

This thesis will focus on the development of methods for high-fidelity neutral-atom quantum computing and their application to digital quantum simulation and error correction. Chapter 1 will present an overview of the platform, including the key techniques that we developed in my PhD. Along the way, we will discuss how to realize high-fidelity, programmable quantum circuits with atoms, including the best practices for mitigating errors. In Chapter 2, we will discuss high-fidelity entangling gates, showing how we improved the state-of-the-art fidelity of CZ gates, as well as implementation of a wide variety of CCZ and CPHASE gates. Chapter 3 will detail our approach to

realizing digital quantum simulations in this platform, which enable experimental studies of the Kitaev honeycomb model and fermion-to-qubit encodings. Chapters 4 and 5 will focus on quantum error correction in this platform, including the realization of the first error-corrected algorithms and the exploration of ingredients for universal fault-tolerant quantum computing. Chapters 6 and 7 will provide an overview of several analog and analog-digital quantum simulation studies that we have performed with our neutral-atom processor. Chapter 8 will describe recent further improvements to the Rydberg entangling gate performance. Finally, we will conclude with an outlook of exciting scientific directions at the frontier of neutral atom quantum computation.

1

Neutral atom quantum computing

This section will focus on giving a high-level overview of the components we use for neutral atom quantum processing: trapping and rearranging atoms, generating entanglement through Rydberg excitation, programmable single-qubit gates, cooling techniques, mid-circuit readout and feedforward, state-selective readout, and qubit reuse. Additional details are also covered in the PhD theses of Sepehr Ebadi, Dolev Bluvstein, Harry Levine, and Alex Keesling.




Figure 1.1: Neutral atom quantum computer architecture. **a**, Experimental layout, featuring optical tools including static SLM and 2D moving AOD traps, global and local Raman single-qubit laser beams, 420-nm and 1013-nm Rydberg beams, and imaging system for both global and local imaging. **b**, Level structure for ^{87}Rb atoms, with the relevant atomic transitions employed in this work. **c**, Control infrastructure used for programming quantum circuits, featuring several arbitrary waveform generators (AWGs). In particular, the moving and Raman 2D AODs are each controlled by two waveforms (one for x axis and one for y axis). An additional AWG is used in first-in-first-out (FIFO) mode for rearrangement before the circuit begins, and then the moving AOD control is switched to the Moving AWG. See Ref. ¹²³ for additional SLM and pre-circuit rearrangement details, Ref. ¹²⁹ for additional Rydberg AWG details and Rydberg excitation details, Refs. ^{262,43} for additional Raman laser and microwave control infrastructure details, and Ref. ⁴³ for additional moving AWG details. All AWGs (other than rearrangement AWG) are synchronized to < 10 ns jitter. During Rydberg gates the traps are briefly pulsed off by a TTL. The FPGA processes images from the camera real-time and in this work sends control signals to the Raman 2D AOD for local single-qubit control. **d**, Example array layout featuring entangling, storage, and readout zones. Zones can be directly reprogrammed and repositioned for different applications, as well as specific tweezer site locations. Tweezer beams and local Raman control are projected from out-of-plane. The entire objective field-of-view is $400\text{-}\mu\text{m}$ diameter, and consequently we do not expect or observe substantial tweezer deformation near the edges of our processor. During two-qubit Rydberg gates, we place atoms $\lesssim 2\text{ }\mu\text{m}$ apart within a gate site, and gate sites are separated such that atoms in different gate sites are no closer than $10\text{ }\mu\text{m}$ during the gate. At our present $n = 53$ and two-photon Rabi frequency of 4.6 MHz , the blockade radius is roughly $4.3\text{ }\mu\text{m}$, such that adjacent atoms are well-within blockade and distant atoms are well-outside blockade.

1.1 TRAPPING AND REARRANGING ATOMS

Our quantum computing architecture is built from arrays of individually trapped neutral atoms. In this thesis, we will use rubidium-87, one of the simplest and widely-used atoms in AMO physics. Rubidium is a hydrogen-like atom, in that it has a single valence electron, in contrast to more complicated atoms with multiple valence electrons. Atoms in their ground state only interact at nanometer scales, but the strength of van der Waals interactions scales remarkably quickly with the principal quantum number as n^{11} , such that for Rydberg states with high values of n (53 and 70 are the two values used in this thesis), atoms can interact at micrometer scales. In later sections, we will show how these strong interactions enable high-fidelity entanglement between the atoms.

Our experiments begin by generating a cloud containing millions of cold ^{87}Rb atoms in a magneto-optical trap inside of a glass vacuum cell. These atoms are then loaded stochastically into programmable, static arrangements of 852-nm traps generated with a spatial light modulator (SLM). We use an additional set of 852-nm moving traps generated by a pair of crossed acousto-optic deflectors (AODs, DTSX-400, AA Opto-Electronic) to rearrange the atoms into defect-free arrays^{28,380,123}. Atoms are imaged and cooled with polarization gradient cooling light (PGC), with light collected by a 0.65-NA objective (Special Optics) onto a CMOS camera (Hamamatsu ORCA-Quest C15550-20UP). This camera was chosen for fast electronic readout times, which will be useful for mid-circuit readout and feedforward.

1.2 RYDBERG EXCITATION AND ENTANGLING GATES

Entangling gates (270-ns duration) between clock qubits are performed with fast two-photon excitation using 420-nm and 1013-nm Rydberg beams to $n=53$ Rydberg states, utilizing a time-optimal two-qubit gate pulse²²⁵ detailed in Chapter 2¹²⁹. Figure 1.4 overviews the atomic level structure used and an example pulse sequence for running a quantum circuit. The atoms are excited from




Figure 1.2: Trapping and rearranging atom arrays. We trap arrays of rubidium-87 atoms into static arrangements of optical tweezers generated from a spatial light modulator (SLM), loading up to around 1000 atoms from the magneto optic trap (MOT) into this tweezer array. Subsequently, we use a set of moving optical tweezers generated by a crossed pair of acousto-optic deflectors (AODs) to rearrange the atoms into programmable geometries.




Figure 1.3: Reconfigurable atom array architecture. In this thesis, we utilize a reconfigurable atom array architecture, in which coherent rearrangement of atomic qubits enables nonlocal connectivity and parallel control of blocks of atoms is achieved through large laser beams sent from the side of the glass cell. Quantum information is stored in long-lived hyperfine qubit states, and then we excite the atoms to highly excited Rydberg states to generate entanglement through controllable entangling gates. Figure adapted from Refs. ¹²³ and ⁴³.




Figure 1.4: Atomic level diagram. Level diagram showing key levels of ^{87}Rb used in our quantum circuits. The clock states, $|0\rangle$, $|1\rangle$ are the qubit states used in this work. Excitation to the Rydberg state between $|1\rangle$ and $|r\rangle$ is carried out by a two-photon transition driven by 420-nm and 1013-nm lasers. Single-qubit rotations are realized with an amplitude-modulated 795-nm laser that drives Raman rotations between the $m_F = 0$ clock states. A DC magnetic field of 8.5 G is applied in most of this thesis, besides Chapter 9.

the $|F = 2, m_F = 0\rangle$ state to the $53S_{1/2}$ Rydberg state in a two-photon scheme with a 420-nm σ^+ -polarized and a 1013-nm σ^- -polarized light. We are able to increase the intermediate-state detuning from previous work, while maintaining similar or higher two-photon Rabi frequencies, in two ways. First, by operating at $n = 53$, we benefit from 50% increase in the Rabi frequency compared to the previously used $n = 70$. Additionally, we upgraded our 10W 1013-nm fiber amplifier laser to a 100W 1013-nm laser (IPG Photonics). Combined, this allows us to operate at the intermediate-state detuning of $\Delta/2\pi = 7.8$ GHz with a Rabi frequency of $\Omega/2\pi = 4.6$ MHz in most of this thesis, and even higher Rabi frequencies in Chapter 8.

The 1013-nm laser seed originates from a Toptica DL Pro external-cavity diode laser which is then locked to and filtered through a Stable Laser Systems ultra-low-expansion cavity (Finesse of 50,000 at 1013 nm) that then injection locks another laser diode, which then seeds our 100W laser. Our 420-nm laser is an 840-nm TiSapph (MSquared) which is locked (but not filtered) to the same

cavity (Finesse of 30,000 at 840 nm) and then frequency doubled with an MSquared ECD-X. For realizing the pulses and waveform shaping for quantum gates, we use an arbitrary waveform generator (Spectrum M4i.663 1-x8) that allows for arbitrary amplitude, frequency, and phase control which simultaneously drives two 420-nm AOMs (MQ240-B40Ao,2-UV, AA Opto-Electronic) in a tandem configuration and maps the RF waveform onto the 420-nm light. The 1013-nm AOM (M1377-aQ8oL-1 coated for 1013 nm, Isomet) is pulsed on for several hundred microseconds and intensity stabilized during the duration of the entire circuit.

The two Rydberg beams are shaped into flat intensity tophat profiles by SLMs in order to maximize intensity while maintaining homogeneity across the gate region¹²³. We tune beam parameters, including X and Y positions, focus, and aberrations, to optimize homogeneity as measured by the differential lightshift on the hyperfine qubit states. We stabilize the beam positions using a reference camera and motorized mirror mounts. To compensate for relative drift between the beam position and the atom array, we re-calibrate the position by stepping the beam positions to maximize the intensity at the atoms as measured by the differential light shift on the qubit transition, which takes ~ 5 minutes. We find that keeping the beams well-centered on the atoms is important to ensure homogeneity and reduce sensitivity to relative beam drifts, and further find that gate parameters are highly reproducible as long as the beams are properly positioned.

We can extract the single-photon Rabi frequencies of the Rydberg excitation using measurements of the two-photon Rabi frequency and the 420-nm lightshift on the ground-Rydberg transition, taking note of the appropriate Clebsch-Gordan coefficients for the multiple intermediate states. We extract $\Omega_{420} = 2\pi \times 237$ MHz and $\Omega_{1013} = 2\pi \times 303$ MHz, where we adopt a convention such that the two-photon Rabi frequency is given by $\Omega = \Omega_{420}\Omega_{1013}/2\Delta$ (with higher values in Chapter 8 as described there). Note that due to the presence of multiple intermediate states, the first-order scattering estimate is proportional to $\frac{4}{3} \times (\Omega_{420}/2\Delta)^2$ and thus, from the scattering perspective, the effective single-photon Rabi frequencies are well-balanced (273 MHz vs 303 MHz).

In Chapter 2, we will describe our journey in high-fidelity Rydberg gates, including how our understanding developed over the course of my PhD, and in Chapter 8, we will describe recent further improvements to the fidelity that we achieved at the end of my PhD.

1.3 PROGRAMMABLE SINGLE-QUBIT CONTROL

Beyond generating entanglement, we need to be able to apply single-qubit control in order to realize universal quantum circuits. To drive arbitrary global single-qubit rotations we use a Raman laser system²⁶² which globally illuminates the atoms with a Raman Rabi frequency of around 1 MHz. Single-qubit rotations are implemented using robust BB1 pulses^{444,427}, while Z rotations are implemented by adjusting the phase in the control software.

To enable individual single-qubit gates, we use the same Raman laser system as our global rotation scheme and illuminate only chosen atoms using a pair of crossed AODs. The focused beam waist in the plane of the atoms is $1.9\ \mu\text{m}$, which is large enough to be robust to fluctuations in atomic positions, and small enough to prevent cross-talk to neighboring atoms separated by $\gtrsim 6\ \mu\text{m}$. For Raman excitation, polarization needs to be carefully considered. Unlike the global path, the local beam propagation direction is perpendicular to the atom quantization axis (set by the external magnetic field). Therefore the fictitious magnetic field $\vec{B}_{\text{fict}}$ responsible for driving the transitions, as described in Ref. ²⁶², preferentially drives $\sigma^\pm$ hyperfine transitions rather than the desired π clock transition²⁶⁰. There exist two possible approaches to single-qubit gates, as illustrated in Fig. 1.6a. First, off-resonant $\sigma^\pm$ dressing generates differential light shifts between qubit states enabling fast local $Z(\theta)$ gates. Global $\pi/2$ rotations convert these to local $X(\theta)$ gates. Second, one can directly apply local $X(\theta)$ gates with direct π transitions by slightly rotating the quantization axis towards the local beam direction; this could be achieved with an external field but, conveniently, $\vec{B}_{\text{fict}}$ has a DC component that naturally rotates the axis. Note that, if the local beam is quickly turned on, this same




Figure 1.5: Single-qubit addressing with local Raman. Fully programmable single-qubit rotations are implemented using Raman excitation through a 2D AOD. Grids of atomic qubits are illuminated in parallel to apply the same gate, while the other qubits are left unaffected.

fictitious DC field causes leakage out of the $m_F = 0$ subspace, therefore Gaussian-smoothed pulses are used.

For most circuits we perform local rotations row-by-row (only 1 Y tone at a time); this enables arbitrary fine-tuning of X coordinates and powers at each site for homogenizing and calibrating rotations (Figure 1.6b). We calibrate using the procedure in ED Figure 1.6c and find these calibrations are stable on month timescales. To quantify the fidelity, we perform randomized benchmarking using 0, 10, 20, 30, 40, and 50 local $Z(\pi/2)$ rotations (per site) on 16 sites, obtaining $\mathcal{F} = 99.912(7)\%$ as shown in Figure 1.6d (note that the single-qubit gates we execute globally have fidelity closer to 99.99%^{43,129}). This approaches the Raman scattering limit for our $\sigma^\pm$ scheme (error of $\sim 7 \times 10^{-4}$ per $\pi/2$ pulse), but when not well-calibrated is limited by inhomogeneity, in particular, associated with distortions of the y position of the rows.

The key challenge for local X gates is ensuring polarization homogeneity, since the Rabi frequency is sensitive to the degree of circularity. We find inhomogeneity both across the array, introduced by a sharp dichroic cut-off noted in Ref.⁴¹, as well as inhomogeneity within each optical tweezer due to polarization breakdown near the tweezer focus. To reduce the first effect, we added




Figure 1.6: Details about local Raman addressing. **a**, $5S_{1/2}$ hyperfine level diagram illustrating the two possible implementations of local single-qubit gates: resonant $X(\theta)$ (purple) and off-resonant $Z(\theta)$ (turquoise) rotations with two-photon Rabi frequencies Ω_{Raman} . In this work, we use the Z rotation scheme and are blue-detuned by 2 MHz from the two-photon resonance. Due to Clebsch-Gordan coefficients, $\tilde{\Omega}_{\text{Raman}}^Z = -\sqrt{3}\Omega_{\text{Raman}}^Z$. **b**, Schematic showing the conversion of local $Z(\pi/2)$ into local $X(\pm\pi/2)$ gates, where the pulses before (after) the central $Y(\pi)$ have positive (negative) sign, while leaving non-addressed qubit states unchanged. The Gaussian-smoothed local pulses have duration $2.5 \mu\text{s}$ for $\pi/4$ pulses and $5 \mu\text{s}$ for $\pi/2$ pulses, and are performed on single rows at a time with a $3 \mu\text{s}$ gap between subsequent gates to allow the RF tones in the AODs to be changed (including this, duration is $5\text{--}8 \mu\text{s}$ per row). In this way, arbitrary patterns of qubits, such as the example drawn, can be addressed. **c**, Calibration procedure used to homogenize the Rabi frequency over a $220 \mu\text{m} \times 35 \mu\text{m}$ array. The position calibration is illustrated for 80 sites: approximate $X(\pi/2)$ gates are locally performed and the horizontal/vertical position of all tones is scanned in parallel such that a Gaussian fit returns the optimal alignment. After this, powers are iteratively calibrated until the fitted scale factors for the individual RF tones converge to unity. **d**, Single-qubit randomized benchmarking of local $Z(\pi/2)$ gates. The local gates are interleaved with random global single-qubit Clifford gates and the final operation C_f is chosen to return to the initial state. Each data point is the average of 100 random sets of Clifford gates, and fitting an exponential decay to the return probability quantifies the fidelity $\mathcal{F}$ per local gate. Note that we apply all 51 global Clifford gates for each data point, such that errors from the global Clifford gates (in addition to SPAM errors) do not contribute to the fitted value.

a second copy of the dichroic into the path with a half-waveplate between the pair, such that any angle-dependent phase shifts upon reflection from the dichroics are equally applied to both the s- and p-polarized components and the polarization remains close to circular. Second, polarization breakdown of a circularly-polarized tweezer results in an off-axis fictitious field with components both parallel and perpendicular to the external magnetic field (in the plane of the tweezer focus)⁴⁰¹; the parallel components can drive Raman transitions and result in dephasing of the clock qubit. Since the magnitude of the maximum off-axis field falls off linearly with tweezer waist, we mitigated this by increasing the waist to $2.5\ \mu\text{m}$. Finally, to increase the projection of the Rabi frequency drive along the magnetic field axis, we displace the Raman beam by roughly 1.5 mm within the back aperture of the objective whose size is 5.5 mm, so that the Raman beam comes in at an angle.

1.4 COOLING AND ENHANCED LOADING

In many of the experiments described in this thesis, for initially loading and cooling the atoms, we utilize global polarization gradient cooling (PGC). We additionally added a new cooling path for lambda-enhanced gray molasses, with the motivations to reduce Doppler error in Rydberg excitation, as well as have access to longer Rydberg evolution times for analog quantum simulations. We chose lambda-enhanced gray molasses on the D₁ line of rubidium (795 nm light) to reach colder temperatures than PGC within a few milliseconds. We estimate an average radial motional quantum number of $\bar{n} \approx 1 - 2$ from fitting a drop-recapture curve⁴²⁶ to $\sim 10\ \mu\text{K}$ in $\sim 1\ \text{mK}$ deep traps.

Following rearrangement and cooling, we prepare atoms in the qubit basis, formed from clock states $|0\rangle = |F=1, m_F=0\rangle$ and $|1\rangle = |F=2, m_F=0\rangle$. Previous work^{264?} used Raman-assisted optical pumping to pump into $|F=1, m_F=0\rangle$, in which π pulses are repeatedly applied on the $|F=1, m_F=-1\rangle \rightarrow |F=2, m_F=-1\rangle$ transition and the $|F=1, m_F=+1\rangle \rightarrow |F=2, m_F=+1\rangle$ transition, followed by resonant depumping of the $|F=2\rangle$ manifold. The main disadvantage of




Figure 1.7: Enhanced loading with lambda-enhanced gray molasses. Example atom images from tests of enhanced loading of atom arrays for different size arrays, and comparison of normal loading compared to enhanced loading with lambda-gray molasses cooling light on the D1 transition. Generally we find that above 75-80% loading fractions are widely achievable across different array sizes used in this work. We find that for smaller arrays, we can better tune the trap depth (and use higher trap depths) to achieve higher filling fractions.

this scheme is that it can require scattering many photons (we performed 40-70 pumping cycles in Refs. ²⁶⁴) to end up in $|F = 1, m_F = 0\rangle$, and empirically we find this causes significant enough heating to negate the benefits of the colder atoms following the Λ -enhanced gray molasses cooling step.

To address this challenge, instead we can first optically pump into the $|F = 2, m_F = +2\rangle$ stretched state with σ^+ -polarized 780-nm light, which scatters only several photons ^{446,176}. Then, we rotate the magnetic field by 90 degrees such that the Raman laser propagation axis has an orthogonal component and can thus drive $\sigma^\pm$ -polarized transitions in the hyperfine manifold (see Ref. ²⁶²). We apply two separate Raman pulses that transfer population first from $|F = 2, m_F = +2\rangle$ to $|F = 1, m_F = +1\rangle$ and then to $|F = 2, m_F = 0\rangle$.

We use Knill composite π pulses ⁴²⁷ for these transfers steps, and suppress unwanted transitions to adjacent m_F states by utilizing Gaussian-shaped optical pulses ⁴²⁷. We note that during this transfer process, the magnetic field direction is along the axis of the tweezer, which quadratically suppresses vector light shifts from the tweezer polarization gradient ⁴²³ that would otherwise limit co-

herence of the non-clock m_F states. Finally, we rotate the magnetic field back to its original configuration before performing quantum circuits.

Additionally, this cooling allowed us to have higher initial filling into the tweezer array from the MOT, as shown in Figure 1.7. In conventional loading into neutral atom arrays, 50-60% of the sites are stochastically loaded prior to rearrangement. However, the Λ -enhanced gray molasses cooling described above enables for higher initial loading fractions. While such enhanced loading is not strictly necessary since we rearrange the atoms after loading, it can be helpful in certain contexts in order to increase the total number of atoms. Additionally, starting from a more densely loaded array can make rearrangement work better in practice.

In the sections below, we will describe additional methods for *local* cooling and imaging.

1.5 MIDCIRCUIT READOUT AND FEEDFORWARD

For many protocols in quantum information science, the ability to measure some qubits in the system while maintaining coherence of the rest of the system is useful. For instance, mid-circuit readout and feedforward can be a powerful tool for preparing long-range entanglement, to study different types of topological order; this will be explored in Chapter 4. In quantum error correction, we would like to be able to measure and decode the state of a logical qubit mid-circuit and apply feedforward logical gates where necessary; a demonstration of such a protocol is explored in Chapter 5 for realizing logical quantum teleportation.

Several ingredients go into achieving practical midcircuit qubit detection. For example, since we want the imaging to be local, we wanted to use a one-dimensional beam geometry, compared to our typical 3D imaging beams, and find ways to preserve coherence of the data qubits. This can be achieved with a single focused beam or two counterpropagating beams. Additional imaging beams could be added to cool along other axes, but for the sake of simplicity, we began with just a single




Figure 1.8: Midcircuit readout. **a**, Example experiment using mid-circuit readout and feedforward to perform logical teleportation. **b**, Mid-circuit readout and feedforward. The imaging histogram shows high-fidelity state discrimination (500 μ s imaging time, readout fidelity $\approx 99.8\%$), and the Ramsey fringe shows that qubit coherence is unaffected by measuring other qubits in the readout zone (error probability $p \sim 10^{-3}$). The FPGA performs real-time image processing, state decoding, and feedforward.

resonant beam along one axis. In addition, we would like to make midcircuit imaging fast in order to quickly implement feedforward operations. We will discuss these various considerations and our approaches to these challenges in this section.

Specifically, to perform mid-circuit readout^{113,402,176,282,277,316} of selected qubits without affecting the others, we use a local imaging beam focused on the readout zone which is roughly 100 μ m spatially separated from the entangling zone^{43,98}. The local imaging beam consists of 780-nm circularly polarized light, with a near-resonant component from $F = 2$ to $F' = 3$ and a small repump component. This beam is sent through the side of our glass vacuum cell, co-propagating with the global Raman and 1013-nm Rydberg beams. We use cylindrical lenses to shape the beam with focused beam waists of 30 μ m in the plane of the atom array and 80 μ m out of the plane. After moving some of the atoms to this readout zone, we first perform local pushout of population in the $F = 2$ ground state manifold (by turning off the repump laser frequency), followed by local imaging of the remaining $F = 1$ population.

As depicted in Fig. 1.9, we collect an average of about 50 photons per imaged atom. To avoid losing the atoms too quickly during mid-circuit imaging (which, unlike our global imaging scheme, does not have multi-axis cooling), we use deep (roughly 5-mK) traps (helping retain the atoms), and




Figure 1.9: Fast local imaging and mid-circuit feedforward. Left: Single-shot 500 μs local image in the readout zone, where the peak corresponds to roughly 50 photons collected by the CMOS camera. Right: Diagram of components involved in mid-circuit readout and feedforward steps. Atom detection and logical state decoding occur using the FPGA, which then outputs a conditional TTL to gate local Raman pulses performed on logical qubits in the entangling zone. Bottom: Diagram of approximate timings for a mid-circuit feedforward cycle. First, $F = 2$ population is pushed out (in 10 μs), and then the remaining $F = 1$ population is imaged locally for 500 μs . The 24 rows of pixels covering the readout zone are read out to the FPGA in 200 μs , after which processing is performed. Finally, a conditional TTL output based on the decoded state gates on or off local Raman pulses. The whole readout and feedforward cycle takes less than 1 ms, and can be sped up in the future by optimizing local imaging and camera readout.

stroboscopically pulse them on and off out of phase of the local imaging light to avoid deleterious effects of the deep traps such as inhomogeneous light shifts and fluctuating dipole force heating (Fig. 1.9)²¹⁵. From a double-Gaussian fit to the two distributions, we extract an imaging fidelity of over 99.9%. Because this fit can lead to an overestimate of the imaging fidelity (for example due to atom loss during imaging), we compare the total SPAM error (measured by amplitude of Ramsey fringe) with local imaging versus with global imaging for the same state preparation sequence, extracting 0.14(5)% higher error with local imaging; with these considerations we conservatively estimate a local imaging fidelity of around 99.8%.

A number of design considerations facilitate local imaging in the readout zone while preserving coherence of the data qubits in the entangling zone⁹⁸. The main sources of decoherence are rescattering of photons from the locally imaged atoms as well as beam reflections and tails of the local imaging beam hitting the data qubits. To suppress this decoherence, there are different approaches

that we considered. The first way to suppress decoherence is by separating the local imaging region far away from the data qubit regions where we want to preserve coherence. We took special care to design the beam path to try to mitigate unwanted reflections. Still, with this care taken, we found that there could be small amounts of light hitting the data qubits for specific beam positions, and this was not always monotonically decreasing with the distance between the local imaging beam and the data qubits.

Therefore, to suppress decoherence further, we added local light shifting to separate the frequency of transitions for the atoms in the two regions relative to each other. The first way that we did this was by using the traps themselves. For the $500\text{-}\mu\text{s}$ mid-circuit imaging used in Ref. ⁴¹, we were able to achieve unchanged coherence (identical within the errorbars) of the data qubits with the local imaging light on as without it. To understand these effects more quantitatively, we measure the error probability of the data qubits in the entangling zone while the local imaging beam is on in the readout zone for up to 20 ms and with higher intensities than used for local imaging. We suppress decoherence by light shifting the data qubits' 780-nm transition to be different from that of the locally imaged qubits by several 10s of MHz, as studied in Fig. 1.11. Data qubit decoherence is further suppressed by the large spatial separation between the readout zone and the entangling zone, where intensity from the local imaging beam's Gaussian tail should theoretically fall off rapidly. Even at large separations, we find that stray beam reflections (e.g. from the glass cell window and other optical elements) can hit the data qubit region. To mitigate this effect, we displace reflections away from the atom array by angling the local imaging beam as it hits the glass cell window. The estimated effects of re-scattered photons from the imaged atoms, especially with the added relative detuning, is negligible. With all these considerations, we find that we are able to suppress data qubit decoherence rates to $\lesssim 0.1\%$ per $500\text{ }\mu\text{s}$ of local imaging exposure, as illustrated in Fig. 1.11.

The full mid-circuit readout and feedforward cycle occurs in slightly less than 1 ms, including local pushout, local imaging, readout of the camera pixels, decoding of the logical qubit state on the




Figure 1.10: Suppressing decoherence with relative lightshifting. Atomic transition and pulse sequence used for local imaging of ancilla qubits. The data qubit trap light shift suppresses data qubit errors, in addition to the large spatial separation between entangling and readout zones. We avoid quickly losing the readout zone atoms during local imaging by using a 5x higher trap depth, and we pulse the ancilla qubit traps and local imaging light to image directly on resonance while avoiding negative effects of large trap light shifts.

FPGA, and a local Raman pulse which is gated on or off by a conditional trigger (Fig. 1.9). In future work, this approach to mid-circuit readout and feedforward can be considerably improved to enable mid-circuit readout close to 100- μ s scale³⁹¹. This method can directly be extended to perform many rounds of measurement and feedforward, where groups of ancilla atoms are consecutively brought to the readout zone throughout a deep quantum circuit.

The second way that we added a relative detuning shift is by shifting the excited state of the data qubits. This has major advantages, including that imaging is not sensitive to uniformity of the light shift, and the hyperfine qubit is not greatly impacted. Initially, we were planning to try using 776nm light to shift the excited state, as this can be potentially combined with 780nm paths in a convenient way. However, based on the experience of a team in the Vuletic group at MIT, we decided to use 1530nm light. This wavelength has the added benefit of being far away from 780nm and 795nm transitions of rubidium-87.

To preserve the coherence of qubits in the storage zone, we illuminate them with a single beam of 1529-nm light (Figure 1.12). By coupling the $5P_{3/2}$ state to the $4D_{5/2}$ state, we impart a strong Stark-shift on the excited $5P_{3/2}$ state²⁰⁷. This causes probe light in the readout zone to appear off-resonant to the storage zone atoms while maintaining qubit information in the hyperfine manifold of the ground state, see Fig. 1.12. The beam is generated by a Connet CoSF-D series 10W fiber laser




Figure 1.11: Measuring decoherence on data qubits from local imaging beam. Characterization of error probability of data qubits during local imaging. From left to right: Data qubit error probability (fraction of population depumped from $F = 2$ to $F = 1$) as a function of local imaging duration out to 20 ms to quantify the effect of the local imaging beam on data qubit coherence for very long illumination. Data qubit error probability after 20 ms of local imaging, as a function of detuning of the local imaging beam, showing suppression of error both red- or blue-detuned from the data qubit transition. Equivalently, increasing the trap depth of the data qubits enables suppression of decoherence due to the local imaging beam. Since qubits in the readout zone are imaged while their traps are pulsed off, any light shift of the data qubit transition from the traps contributes directly to the relative detuning. Right: for a long 10.5 ms local beam illumination with optimal local imaging parameters, we observe a 0.7(1)% increase in data qubit error during an XY8 dynamical decoupling sequence. This suggests a roughly 0.034(5)% error probability for the data qubits during the 500 μ s mid-circuit readout image employed in this work.

and is focused down to an elliptical waist of $35 \mu\text{m} \times 65 \mu\text{m}$. The shorter waist of the beam is aligned vertically to the center of the storage zone. We image the beam in a 4f system and apply a knife-edge in the image plane, approximately 4 beam waists from its center, to suppress its Gaussian-tail. Stray 1529-nm light, even at low powers, can degrade the imaging quality in the readout zone. We find, therefore, that beam-shaping is critical for maintaining stable imaging quality and coherence on the storage-zone atoms for the layout of our array.

We measure dephasing of the storage-zone qubits, as a function of detuning from the bare transition, while readout-zone qubits are illuminated with local probe and repumper light, Fig. 1.12. We capture the key features of the spectrum with a simple model where the additional dephasing at each drive power scales as $\exp\left(-\frac{\Omega_{\text{probe}}^2 \Gamma_{\text{probe}} t}{4\Delta_{LS}^2}\right)$ where Ω_{probe} , Γ_{probe} are the Rabi frequency and scattering rate of the local imaging beams, t is the illumination time, and Δ_{LS} is the calculated light-shift of $5D_{3/2}$ due to the coupling to $4D_{5/2}$ and $4D_{3/2}$. More complex on-resonance or multi-level features are not captured by this simple model and are particularly sensitive at detunings between the resonances of the 4D-levels⁴⁷⁰. During all experiments with qubit re-use and local imaging, we




Figure 1.12: Atomic physics of hiding beam at 1529-nm. **a**, Upper panel: The hiding beam is aligned to the storage zone and a knife-edge used to suppress its gaussian-tail in the readout zone. Lower panel shows the relevant atomic levels in ^{87}Rb . The hiding beam couples the excited-to-excited state transition and imparts a 6 GHz light-shift on the $5P_{3/2}$ state. At this detuning, the $5S_{1/2}$ ground state polarizability is approximately 2×10^{-5} times smaller²⁰⁷, thus maintaining coherence in the hyperfine qubit manifold while shifting the transition far off-resonance from the imaging light. **b**, Additional dephasing on atoms in the storage zone due to local imaging beams in the readout zone, for various drive powers and detunings. Here we use 20-ms of illumination and closer probe detuning than in deep circuit experiments. We plot a fit to a simple model (described in the text). For deep circuit experiments, we operate at 16 GHz red-detuned from the bare $5P_{3/2} \rightarrow 4D_{5/2}$ transition (not shown). **c**, By scanning the detuning of the probe light, we observe an Autler-Townes (AT) splitting of probe resonance when coupled to the $4D_{5/2}$ -level. For this measurement we operate at a low drive power compared to typical values and the drive detuning is 500 MHz red-detuned from resonance, such that both peaks can be experimentally measured in the limited probe detuning range. The small peak is likely due to leaked light and is not expected. **d**, The measured AT-splitting scales linearly with the Rabi frequency of the drive, or the square root of the drive power. **e**, Level diagram illustrating the coupled three-level ladder system which leads to the emergence of the Autler-Townes feature.

address the storage zone at 1529.49 nm with ≈ 1.2 W. This corresponds to an approximate lightshift of 6 GHz on the $5P_{3/2}$ state. To further characterize the 1529-nm laser we explore varying the detuning of the local imaging light and observe a clear Autler-Townes splitting, Fig. 1.12. We find that, as expected, the separation of two fitted Lorentzian peaks scales linearly with the square-root of the drive power.

1.6 STATE-SELECTIVE READOUT

An important capability for reading out atoms is the ability to perform non-destructive qubit readout, enabling loss detection as well as qubit re-use. We realize this with a one-dimensional optical




Figure 1.13: Motivation for state-selective readout. In conventional readout, one of the two spin states is converted to loss before fluorescence imaging. This causes us to lose half the atoms, and prevents detection of loss events. By comparison, using spin-to-position conversion, we can image both spin states, enabling loss detection and efficient qubit reuse.




Figure 1.14: State-selective readout overview. **a-b**, Spin-to-position conversion for non-destructive, loss-resolved qubit readout is accomplished with a state-selective 1D lattice that converts the atom spin state into position. Plot shows measured Rabi oscillation. This non-destructive readout has 0.46(4)% bit-flip error and 0.24(2)% loss. **c**, Results of repeated rounds of $d = 5$ surface code using loss detection, showing a snapshot of the data block and multiple ancilla blocks.

lattice⁴⁴⁶ which pins one of two spin states, use optical tweezers to separate the pinned and unpinned states, and then image the atom position. To further enable mid-circuit qubit measurement and re-use on large arrays, we develop methods of low-loss, high-fidelity qubit readout and re-initialization while only needing moderate trap depths (see below). In particular, as shown in Fig. 1.14a, this state-selective readout enables qubit reuse by merging the two separated tweezers back together and then re-initializing the qubits. We use these techniques in Ref.⁴² for re-using atoms and extending the depth of error-corrected computation.

We realize non-destructive qubit readout throughout this work with spin-to-position conversion^{287,438,350,446}. A one-dimensional optical lattice is formed by two 795-nm counterpropagating local beams, both sourced from the same titanium:sapphire laser (M Squared) and operated at 50-200 GHz blue-detuned of the D1 line. Both beams are σ^- polarized such that $|F = 2; m_F = -2\rangle$




Figure 1.15: Spin-to-position conversion. **a**, Level diagram showing the ^{87}Rb hyperfine levels used to engineer a spin-selective one-dimensional optical lattice. **b**, Trapping potentials for the bright and dark states. The dark state only experiences a lightshift from the optical tweezers - allowing the atom to be moved around freely - while the bright state is additionally confined by the optical lattice potential. By using a blue-detuned lattice, atoms are trapped in intensity antinodes and so scattering of the lattice light is reduced. **c**, Schematic timeline of spin-to-position conversion. The time to transfer the clock qubit to the bright and dark states for readout is typically on the order of roughly $20\ \mu\text{s}$, but for some of our measurements is several milliseconds due to using a slow global rotation of the magnetic field (panel e). See Methods text for additional information. **d**, Transfer of qubit state $|1\rangle$ to the dark state via resonant optical pumping. **e**, Transfer of qubit state $|0\rangle$ into $m_F = +1, +2$ states. This is achieved with either a coherent Raman transfer to $F = 2, m_F = +2$ or incoherent pumping with σ^+ -polarized repumper. Both approaches achieve the same bright state readout fidelity, but in the specific implementation we use here the Raman transfer takes several milliseconds owing to the rotation of the external magnetic field for driving σ transitions (2 and 3). **f**, Quadratic suppression of readout error due to scattering. The bright state transfer ensures that at least two lattice-induced scattering events are required to cause a readout error, which occurs if the AOD tweezer has not yet moved away as the bright state becomes unpinned. Scattering further causes diabatic changes in the depth of the lattice potential and may contribute to atom loss.

is a dark state and $|F = 2; m_F = +2\rangle$ experiences a maximum lightshift of approximately 6 MHz, corresponding to approximately 300 kHz trap frequency in one axis. The close detuning is a balance between minimizing off-resonant coupling to the D2 line for the dark state and reducing scattering and heating from the lattice light. Since the clock state qubit is used for computation, for readout we first optically pump $|F = 2; m_F = 0\rangle$ into the dark state with 780-nm σ^- -polarized light resonant to $F = 2$ to $F' = 3$ which is co-propagating with one port of the lattice. To suppress the probability of scattering into the dark state during readout, we further transfer $|F = 1; m_F = 0\rangle$ to $|F = 2; m_F = +2\rangle$ (bright state), which also increases the trap depth. This is achieved via either a coherent Raman transfer or with incoherent σ^+ -polarized 780-nm repumper from $F = 1$ to $F' = 2$.

Following these state transfers, the lattice is ramped up adiabatically over approximately 100 μ s. AOD tweezers pick up atoms in the dark state and move them by approximately 2 μ m over approximately 500 μ s; during this, atoms in the bright state are pinned in place by the stronger confinement of the lattice. Finally, the lattice is ramped down and conventional camera-based readout then images the position of the atom, allowing identification of the spin state as well as loss detection. In Ref. ⁴², we measure an error probability of 0.87(7)% for the dark state, 0.05(5)% for the bright state, and a 0.24(2)% probability of loss. The asymmetric error arises from trade-offs when simultaneously optimizing for loss and readout fidelity and can be tuned to be more balanced; typically, due to the pumping fidelity, the dark state error is at least $\sim 0.3\%$ higher than the bright state.

1.7 QUBIT REUSE

One-dimensional and finite-field operation for imaging and cooling

For local cooling and imaging, we use two counterpropagating 780-nm beams with opposite circular polarization ^{104,434} (Figure 1.17). The beams are red-detuned from $F = 2$ to $F' = 3$ and have a variable relative detuning; the σ^+ -polarized beam additionally contains a small repump compo-

nent. Conventional methods based on polarization-gradient cooling (PGC) require zero magnetic field, however mid-circuit operation requires a finite magnetic field to maintain the quantum state of active qubits. To this end, we developed a scheme for one-dimensional PGC in finite magnetic field. PGC is based on a linear polarization rotating along the beam propagation direction which produces a population imbalance within the hyperfine levels⁴³⁴; in finite field, this imbalance is disturbed and the cooling mechanism breaks down⁴³⁴. By transforming to a frame where the polarization rotates in time, a fictitious field appears which cancels the external field and restores the cooling effect. This condition is achieved by detuning the two counterpropagating beams - which are (anti)parallel to the external magnetic field - by two times the Zeeman splitting of adjacent m_F levels. As shown in Fig. 1.17, this detuning method works across the full range of magnetic fields studied (up to 8.6 G). Furthermore, it is broadly applicable to finite-field implementation of any one-dimensional technique based on the same polarization configuration used here, for example gray-molasses cooling^{180,66}.

While this finite-field PGC is sufficient to image without loss, we added a second stage of EIT cooling to further reduce the atom temperature⁹². The scheme, shown in Fig. 1.17, uses the same beams as the PGC imaging and only requires changing to be blue-detuned of $F = 2$ to $F' = 2$ (by ~ 80 MHz) and reducing the power in one of the beams. Since the cooling is uniaxial, it is *a priori* unclear if all three motional degrees of freedom can be cooled with such techniques. Using both drop-recapture measurements and adiabatic ramp-down measurements of the atom temperature⁴²⁶, we probe the radial and axial atom temperature and find them both to be comparable to three-dimensional techniques (Fig. 1.16 and Fig. 1.17). Furthermore, the steady-state temperature and loss is set only by the EIT cooling fidelity and is independent of the degree of heating introduced from the prior circuit.

Repeated rearrangement from reservoir

The mid-circuit image identifies the qubit state as well as which atoms are lost. Before rearrangement, the atoms are recombined into their original tweezers, balancing the trap depth between the AOD and SLM tweezers to minimize loss and using cooling throughout. After this recombination, we fill empty sites using the reservoir.

In each round of rearrangement, target rows are refilled sequentially with one parallel step per row. All atoms in each step are sourced from a single reservoir row. We choose efficient horizontal moves and optimize the reservoir-to-target row pairings to minimize travel distance. Finally, since the local imaging beams do not cover the full extent of the reservoir, the reservoir site occupancies are stored from a global image before the circuit begins and used reservoir atoms are tracked in software. This leads to a slowly growing rearrangement infidelity.

Mid-circuit re-initialization

After qubits have been measured and atom loss refilled, the spin state is re-initialized to re-use the qubit. This local state preparation is performed in the readout zone using a Raman-assisted optical pumping scheme. Local Raman is used for the coherent π -pulses and the local probe beams are used for resonant depumping of the $F = 2$ manifold. Due to the close horizontal spacing of traps in the readout zone, we minimize crosstalk between local Raman tweezers by alternating the applied π -pulses between odd and even columns. We perform 24 cycles of pumping per atom over a few hundred microseconds.

1.8 THE DESIGN OF PROGRAMMABLE HIGH-FIDELITY CIRCUITS

Here we will describe how we combine all of the tools discussed in the prior sections to design and execute high-fidelity quantum circuits with atom arrays, including some of the design considerations.

Design of circuit layouts and zones




Figure 1.16: Architecture for qubit reuse and constant-entropy computation. **a**, Illustration of processes for removing entropy generated by computation. Logical teleportation is used to ensure all physical errors are removed. **b**, Rabi oscillations measured using the same atoms for 150 cycles of non-destructive measurement and re-initialization. Each curve shows a single experimental run, averaged over 200 atoms in parallel. 3D cooling methods are used in this subplot as coherence does not need to be preserved. **c**, Local cooling with 1D polarization gradient cooling (PGC) and electromagnetically induced transparency (EIT). The finite magnetic field is compensated by applying a relative detuning between the two circularly polarized counterpropagating beams, leading to a rotating frame where the effective field is zero. The atom loss and temperature is constant as a function of time and recovers to steady state within one cycle after applying a perturbation (turning off cooling on one layer) to the system (dashed line). **d**, Additional shielding of the data qubits is provided by a 1529-nm shielding laser, which rapidly suppresses decoherence induced by the imaging (resonance is at 1529.365 nm).




Figure 1.17: One-dimensional imaging and cooling in finite field. **a**, Atomic level structure for cooling and imaging of ^{87}Rb . The one-dimensional (1D) techniques here rely on coupling hyperfine states separated by $\delta m_F = 2$ with counterpropagating $\sigma^\pm$ beams. In a finite magnetic field B_{ext} , the level degeneracy is lifted by the Zeeman effect; here μ_B is the Bohr magneton and $g_F = 1/2$ the Landé factor. **b**, Single-shot local image in 8.6 G external magnetic field. Spin-to-position conversion is used in the readout zone and the reservoir is partially imaged by the tails of the imaging beams. **c**, Example site-averaged imaging histogram in the readout zone. **d**, Schematic timeline of mid-circuit measurement and re-initialization used for deep circuit experiments. Due to latency bottlenecks associated with desktop-computer-based processing of images and rearrangement waveforms, we do not attempt here to optimize any speeds and simply use comfortable parameters. The circuit is 13.5-ms long. Including idle times, the spin-to-position conversion takes 4-ms, imaging takes 10-ms, cooling and re-sorting takes 13.3-ms (≈ 7 -ms latency for mid-circuit data processing), and re-initialization takes 1.1-ms. **e**, 1D polarization gradient cooling (PGC) in finite magnetic field. Top left: interference between the two counterpropagating $\sigma^\pm$ probe beams generates a helix of linear polarization⁴³⁴. Detuning the two beams rotates the helix such that, in the rotating frame, a fictitious magnetic field appears that cancels the external field and restores the zero-field PGC cooling mechanism. Bottom left: Atom loss from 1D PGC light at different relative beam detunings, Δ , in varying external magnetic field. The atoms are illuminated for 10 ms under comfortable imaging parameters. A cooling resonance is observed when Δ matches two times the Zeeman splitting (bottom right). Top right: Atom temperature around the cooling resonance in 4.3 G field, obtained from drop-recapture measurements. **f**, 1D electromagnetically induced transparency (EIT) cooling. Top: A strong σ^+ pump beam is combined with a weak σ^- probe beam to drive transitions between quantum harmonic oscillator states of the optical tweezer, cooling the atom. The EIT Fano resonance ensures heating transitions are suppressed⁹². Bottom: Ramp-down measurement of atom temperature. The SLM trap depth is adiabatically ramped down to $\sim 5\text{ }\mu\text{K}$ and held for 5-ms before being ramped back up; the atom loss probability from this process probes the temperature of all three motional axes⁴²⁶. Since the EIT cooling resonance is narrow and depends on trap frequency, spatial inhomogeneity in the SLM trap depths translates to inhomogeneous cooling and broadens the temperature distribution. Even so, the coldest sites after 1D PGC imaging and EIT cooling reach lower temperatures than conventional 3D PGC, highlighting the utility of these 1D techniques.




Figure 1.18: Putting all the components together. **a**, Experimental layout illustrating key optical tools, similar to Ref. ⁴¹ with the addition of beams for local cooling, imaging and hiding to enable qubit re-use experiments. **b**, Control infrastructure for programming quantum circuits. The entire waveform for all AWGs (except for rearrangement) is uploaded at the start of each experimental run. For qubit re-use experiments, the Moving, Raman AOD and Rydberg AWGs loop the same memory segment each layer. The full waveform is programmed for the Raman AWG to ensure phase continuity. For mid-circuit rearrangement, waveforms are calculated on-the-fly using a desktop computer and sent to the Rearrangement AWG operated in first-in first-out (FIFO) mode. **c**, Level diagram of the relevant atomic transitions of ^{87}Rb . **d**, Processor layout used for qubit re-use experiments and relevant laser beams. Atoms are arranged into storage, entangling and readout zones, with an additional reservoir for refilling lost atoms mid-computation. The 1529-nm hiding beam illuminates the storage zone to preserve coherence of active qubits during imaging in the readout zone. Parallel two-qubit gates are performed in the entangling zone with global Rydberg beams, and local detunings are optionally applied to selected gate sites using an SLM. The readout zone is illuminated with local beams for 1D PGC imaging and EIT cooling, as well as two counter-propagating lattice beams to form a spin-dependent potential for readout via spin-to-position conversion. The entire array is addressed with global Raman for dynamical decoupling. The same Raman light is directed through a pair of crossed AODs for local single-qubit gates. Global imaging and lambda-enhanced gray molasses cooling light are used for the initial loading.




Figure 1.19: Programming circuits with atoms. Example of a quantum circuit and how it is compiled into the low-level hardware components described in this chapter, including fully programmable single-qubit gates, high-fidelity Rydberg entangling gates with nonlocal connectivity achieved through atomic motion, and readout capabilities.

We define our entangling zone size by adjusting the size of the 420-nm and 1013-nm Rydberg “tophat” beams generated by SLM phase profiles¹²³, to be homogeneous over a certain region. As the Rydberg beams propagate longitudinally, the entangling zone is longer than it is tall. We optimize tophats to be homogeneous over roughly 250- μm horizontal extent. Taller regions are also achievable, with a trade-off with reduced laser intensity and greater challenge in homogenization. The 250- μm width of the zones used here is set by the bandwidth of our AOD deflection efficiency. We position the readout and storage zones flexibly with respect to this entangling zone, balancing various tradeoffs, for example the need to minimize decoherence on data qubit during midcircuit readout

During two-qubit Rydberg ($n=53$) gates, we place atoms $\lesssim 2\ \mu\text{m}$ apart within a “gate site”, resulting in $\gtrsim 450\ \text{MHz}$ interaction strength between pairs, significantly larger than the Rabi frequency. Importantly, due to the use of the Rydberg blockade^{224,264}, the gate is largely independent of the exact distance between atoms. Hence, precise inter-atom positioning is not required. Gate sites are separated such that atoms in different gate sites are no closer than 10 μm during the




Figure 1.20: Neutral atom quantum computer architecture. **a**, Experimental layout, featuring optical tools including static SLM and 2D moving AOD traps, global and local Raman single-qubit laser beams, 420-nm and 1013-nm Rydberg beams, and imaging system for both global and local imaging. **b**, Level structure for ^{87}Rb atoms, with the relevant atomic transitions employed in this work. **c**, Control infrastructure used for programming quantum circuits, featuring several arbitrary waveform generators (AWGs). In particular, the moving and Raman 2D AODs are each controlled by two waveforms (one for x axis and one for y axis). An additional AWG is used in first-in-first-out (FIFO) mode for rearrangement before the circuit begins, and then the moving AOD control is switched to the Moving AWG. See Ref. ¹²³ for additional SLM and pre-circuit rearrangement details, Ref. ¹²⁹ for additional Rydberg AWG details and Rydberg excitation details, Refs. ^{262,43} for additional Raman laser and microwave control infrastructure details, and Ref. ⁴³ for additional moving AWG details. All AWGs (other than rearrangement AWG) are synchronized to < 10 ns jitter. During Rydberg gates the traps are briefly pulsed off by a TTL. The FPGA processes images from the camera real-time and in this work sends control signals to the Raman 2D AOD for local single-qubit control. **d**, Example array layout featuring entangling, storage, and readout zones. Zones can be directly reprogrammed and repositioned for different applications, as well as specific tweezer site locations. Tweezer beams and local Raman control are projected from out-of-plane. The entire objective field-of-view is 400- μm diameter, and consequently we do not expect or observe substantial tweezer deformation near the edges of our processor. During two-qubit Rydberg gates, we place atoms $\lesssim 2 \mu\text{m}$ apart within a gate site, and gate sites are separated such that atoms in different gate sites are no closer than $10 \mu\text{m}$ during the gate.

gate, resulting in negligible long-range interactions. We do not observe any error on storage-zone atoms when Rydberg gates are executed in the entangling zone. Even though the tail of the tophat Rydberg excitation beams is only suppressed to $\sim 0.1\times$ intensity, the two-photon drive is far off-resonant due to the ~ 20 MHz 1013 light shift detuning which is present for the entangling zone atoms¹²⁹. We natively realize physical CZ gates; when implementing CNOTs we add physical H gates. We find minimal two-qubit cross-talk between gate sites, as probed with long benchmarking sequences in Ref.¹²⁹. Although Ref.¹²⁹ appears to find some small cross-talk seemingly originating from decay into Rydberg P states, this is considerably suppressed in the practical operation due to the $\sim 200\ \mu\text{s}$ duration between gates, during which time Rydberg atoms should either fly away or decay back to the ground state.

Dynamical decoupling

In our circuit design, we engineer our echo sequences in order to cancel out as many deleterious aspects as possible. We ensure that in our dynamical decoupling we have an odd number of π pulses between CZ gates (whenever possible), as this echoes out both systematic and spurious contributions to the single-qubit phase^{43,129}. We apply appropriate $X(\pi)$ and $Z(\pi)$ rotations between local addressing with the local Raman to cancel out errors induced by the global $\pi/2$ pulses, as well as between pulses of the 420-nm laser (when used for entangling zone single-qubit rotations⁴³) to echo out small crosstalk experienced in the storage zone by the tail of the 420-nm beam. For our global decoupling pulses we use both BB1 pulses⁴⁴⁴ and “SCROFULOUS” pulses¹⁰². To benchmark and optimize coherence during our complex circuits, we perform a Ramsey fringe measurement encompassing the entire movement and single-qubit gate sequence and optimize the observed contrast⁴³. When performing properly, our total single-qubit error is consistent with SPAM¹²⁹, an effective coherence time of 1-2 s, and the Raman scattering error of all the Raman pulses^{262,43}. We note that these measured coherence time include the movement within and between zones; although we use

fewer pulses (typically 1 per movement) than the XY16-128 sequence used to benchmark 1.5 s coherence in Ref. ⁴³, the coherence times here are naturally longer due to further-detuned tweezers used (852 nm rather than 830 nm).

Local single-qubit gates ^{177,27} with the Raman AOD are realized in arbitrary positions in space on both AOD and SLM atoms. Targeted logical qubit blocks are addressed by a grid illumination of the logical block. Arbitrary patterns of rotations on the qubit grid (e.g., during color code preparation) can be realized with row-by-row serializing, with the targeted x coordinates in each row simultaneously illuminated. With the local single-qubit gates and entangling zone two-qubit gates calibrated, the entire circuit is simply defined by the appropriate trapping SLM phase profile, and waveforms for our several AWG channels and TTL pulse generator. These several channels then program complex, varied circuits on hundreds of physical qubits.

Coherent atom motion and transfers

The SLM tweezers can have arbitrary positions, but are static. The AOD tweezers are mobile, but have several constraints ^{415,43}. In particular, the AOD array creates rectangular grids (but not all sites need to be filled). During the atom moving operations, they are only used for stretches, compressions and translations of the AOD trap array: i.e., atoms move in rows and columns, and rows and columns never cross ^{43,415}. Arbitrary qubit motions and permutation is achieved by shuttling atoms around in AOD tweezers, and then transferring atoms between AOD and SLM tweezers as appropriate. We perform gates on pairs of atoms in both AOD-AOD traps and AOD-SLM traps, with no observed difference for gate performance as measured by randomized benchmarking ^{129,132}.

We find that free-space shuttling of atoms (i.e., no transfers) in AOD tweezers comes with very little fidelity cost (other than time overhead), consistent with Ref. ⁴³. Two additional improvements here are the use of a photodiode to calibrate and homogenize the 2D deflection efficiency of our 2D AODs to percent-level homogeneity across our used region, and engineering atomic trajectories

and echo sequences to cancel out residual path-dependent inhomogeneities. For example, we move an atom $100\ \mu\text{m}$ away to realize a distant entangling gate, and then before returning the atom, we perform a Raman π pulse, so that differential light shifts accumulated during the return trip cancel with the first trip. Motion is realized with a cubic profile as in Ref. ⁴³, the characteristic free-space movement time between gates is roughly $200\ \mu\text{s}$, and acoustic lensing effects from the AOD are estimated to be negligible. We pulse the 1013 laser off during motion to remove loss effects from the large light shifts. Note that the 1013 -induced differential light shift on the hyperfine qubit is only kHz-scale but we still ensure its effects are properly echoed out.

Transferring atoms between tweezers ³⁷ presents additional challenges. We measure the infidelity of each transfer, encompassing both dephasing and loss, to be $\lesssim 0.1\%$. To achieve this performance, in our transfer from SLM to AOD, we ramp up the AOD tones' intensity (with quadratic intensity profile when possible) corresponding to the appropriate sites over a time of $100\text{--}200\ \mu\text{s}$ to a trap depth $\sim 2\times$ larger than the SLM trap depth, and then move the AOD trap $1 - 2\ \mu\text{m}$ away over a time of $50\text{--}100\ \mu\text{s}$. These time scales can likely be shortened considerably while suppressing errors using optimal control techniques. During subsequent motion we leave the AOD trap depth at this $2\times$ value. To transfer an atom AOD to SLM we perform the reversed process. During these transfer processes, the differential light shifts on the transferred atoms are dynamically changing, and can result in large unechoed phase shifts. As such, whenever possible we engineer circuits such that pairs of transfers will echo with appropriately chosen π pulses. When echoing pairs of transfers is not possible, we perform 1 cycle of XY₄ or XY₈ dynamical decoupling during the transfer. Finally, we note that low-loss transfer is highly sensitive to alignment of the AOD and SLM grid. We fix small optical distortions between the AOD and SLM tweezer grids by fine adjustment of individual SLM grid tweezers, which can be arbitrarily positioned, to overlap with the AOD traps as seen on an image plane reference camera. It is important to adjust the SLM and not the AOD, as small adjustments of individual AOD tones deviating from a frequency comb causes beating and atom loss.

AOD intermodulation effects

We observe several intermodulation effects from the AODs that can result in degraded performance for specific AOD moves. First, it is important to ensure that the frequency tones in a given AOD axis are in an exact frequency comb, as intermodulation can lead to interference and beating near trap frequencies. Second, we observe here that the relative frequencies of the X-frequency-spacing and Y-frequency-spacing is also important, and that when beat notes of these are near trap frequencies this can also lead to heating. As such, we now primarily use the AODs for translations, avoiding compressions / expansions of the grid when possible, and choose incommensurate spacings for X and Y to avoid accidental cross-resonances.

Compiling circuits into software commands

The quantum circuits are programmed with a control infrastructure consisting of five arbitrary waveform generators (AWG) (Spectrum Instrumentation), synchronized to < 10 -ns jitter. The 2-channel rearrangement AWG is used for real-time rearrangement, the 2 channels of the Rydberg AWG are used for entangling gate pulses and for local SLM detunings, the 4 channels of the Raman AWG are used for IQ (in-phase and quadrature) control of a 6.8 GHz source^{262,43} (the global phase reference for all qubits) and pulse-shaping of the global and local Raman driving, the 2 channels of the Raman AOD AWG are used for displaying tones that create the programmable light grids for local single-qubit control, and the 2 channels of the Moving AOD AWG are used for controlling the positions of all atoms during the circuit.

In order to translate the high-level design of circuits into hardware-level waveforms, we developed over time a dictionary of text commands that perform the desired functions. **Appendix G** contains example commands for the waveform generation used for our different AWGs.

2

High-fidelity entangling gates

Additional details for this chapter can be found in Ref.¹²⁹. Further recent improvements to the Rydberg entangling gate fidelity are described in Chapter 8 and in Ref.¹³²

A central ingredient in quantum information science is the ability to entangle qubits with high-fidelity.³³⁸ Neutral atom arrays have recently emerged as a promising quantum computing platform, featuring coherent control over hundreds of qubits^{123,380} and any-to-any gate connectivity in a flexible, dynamically reconfigurable architecture⁴³. The major outstanding challenge has been to

reduce errors in entangling operations mediated through Rydberg interactions²⁶⁴. Here we report the realization of two-qubit entangling gates with 99.5% fidelity on up to 60 atoms in parallel, surpassing the surface code threshold for error correction^{245,141}. Our method employs fast single-pulse gates based on optimal control²²⁶, atomic dark states to reduce scattering⁴⁶, and improvements to Rydberg excitation and atom cooling. We benchmark fidelity using several methods based on repeated gate applications^{284,148}, characterize the physical error sources, and outline future improvements. Finally, we generalize our method to design entangling gates involving a higher number of qubits, which we demonstrate by realizing low-error three-qubit gates^{306,242}. By enabling high-fidelity operation in a scalable, highly connected system, these advances lay the groundwork for large-scale implementation of quantum algorithms⁴¹⁸, error-corrected circuits¹⁴¹, and digital simulations²³¹.

2.1 TWO-QUBIT GATES

Errors limit the computational capabilities of current quantum devices and must be made sufficiently low to permit efficient quantum error correction. In particular, two-qubit gate error rates below 1% (i.e. fidelities above 99%) are required to surpass quantum error-correcting thresholds¹⁴¹. Moreover, maintaining a combination of such low error rates, highly parallel control, and a high degree of connectivity while scaling system size is crucial to realizing large-scale quantum computers. While high-fidelity entangling operations were realized on isolated qubit pairs early on^{32,25,22,315}, only recently have these techniques been extended to larger systems. State-of-the-art examples include 99.4% fidelity on a 72-qubit superconducting chip⁵ and 99.4-99.6% fidelity on a 31-ion chain²¹⁷. Scaling these systems to even larger numbers of qubits while maintaining low error and efficient control is an exciting frontier^{230,357}, yet also presents substantial platform-specific scientific and engineering challenges.

Recently, arrays of neutral atoms have emerged as a promising quantum processing platform capable of coherent control of hundreds of qubits^{123,380} for analog quantum simulations. This platform also features a flexible, dynamically reconfigurable architecture[?], whereby entangling operations can be performed between neutral atom qubits with arbitrary connectivity and in a highly parallel manner. While these capabilities open unique opportunities for both large-scale digital simulations³³⁸ and computation with error-corrected qubits¹⁴¹, the major outstanding challenge in the field has been to improve the two-qubit gate fidelity significantly above the previously demonstrated $\approx 97.5\%$ ^{264,146}. In this Article, we experimentally realize two-qubit controlled phase (CZ) gates with 99.5% fidelity while operating on up to 60 neutral atom qubits in parallel, closing the gate-fidelity gap to other state-of-the-art platforms^{357,333,5}. This advance is achieved by using a family of optimal gate schemes^{225,320} relying on the Rydberg blockade mechanism that are robust to experimental imperfections and spontaneous scattering, alongside the implementation of several experimental tools to overcome previously dominant error sources. To characterize the two-qubit gates, we employ multiple complementary benchmarking methods using repeated gate applications, each giving consistent results. Finally, these techniques are generalized to entangling operations involving a higher number of qubits, allowing us to experimentally realize parallel, high-fidelity three-qubit entangling gates.

In our approach, quantum information is encoded in long-lived $m_F = 0$ hyperfine qubits¹⁷⁷, where high-fidelity ($>99.97\%$) coherent single-qubit rotations are driven by Raman laser pulses. Entangling operations are performed in parallel by positioning the atoms, trapped in individual optical tweezers, into designated gate sites, followed by state-selective excitation into highly excited atomic Rydberg states using a two-photon transition. Errors in such quantum operations can occur due to spontaneous emission from the intermediate atomic state $|e\rangle$, atomic temperature effects, Rydberg state decay during the gate, as well as miscalibrations and experimental imperfections such as laser noise or inhomogeneity.

We address these errors through the combination of gate schemes relying on optimal control and experimental improvements. Our method for gate implementation is inspired by the recently proposed time-optimal gate by Jandura and Pupillo^{22,5}, which utilizes a numerically optimized continuous phase profile^{42,2,146,22,5} for a single laser pulse (as opposed to a discrete phase jump between two laser pulses²⁶⁴). We generalize this gate scheme to a family of single-pulse gates with a small set of tunable gate parameters, including a version of the time-optimal gate consisting of a parameterized sinusoidal phase modulation, as well as a second, smooth-amplitude gate.

We calibrate these gates by tuning several global parameters, which lends robustness to experimental imperfections: an optimal set of gate parameters can be found even in the presence of systematic offsets such as finite Rydberg laser pulse rise time (Fig. 2.8). We further optimize our control pulses to suppress scattering from the short-lived intermediate state $|e\rangle$ by minimizing population in the “bright” dressed state ($|B\rangle \sim |1\rangle + \sqrt{\frac{2\Omega}{\Delta}} |e\rangle + |r\rangle$) containing $|e\rangle$ and maximizing population in the “dark” state ($|D\rangle \sim -|1\rangle + |r\rangle$) not containing $|e\rangle$ (where Ω is the two-photon Rabi frequency and Δ is the intermediate-state detuning)^{46,305}. This optimization is achieved through the appropriate selection of the relative signs of the intermediate and two-photon detunings, as well as through smooth pulse shaping for the smooth-amplitude gate (Fig. 2.4).

Our experimental realization makes use of the apparatus described previously in Refs.^{123,43}, with which we rearrange ^{87}Rb atoms into programmable, defect-free arrays. Two main experimental upgrades facilitate high-fidelity entangling gate operation. First, we suppress scattering by significantly increasing intermediate-state detuning while maintaining a high two-photon Rabi frequency ($\Omega/2\pi = 4.6$ MHz), enabled by excitation to a lower-lying ($n = 53$) Rydberg state with a tenfold higher power laser. Second, to suppress decoherence from atomic velocity and position fluctuations, we implement Λ -enhanced gray molasses cooling and an improved optical pumping technique to achieve colder temperatures (radial phonon occupation $\bar{n} \approx 1 - 2$).




Figure 2.1: Level diagram showing error channels for gates. Level diagram shows the eight states assumed in the simulation. We assume a $88 \mu\text{s}$ Rydberg state lifetime (based on measured T_1 with 1013-nm scattering lifetime subtracted), and a 110 ns lifetime for the intermediate states. We assume the following branching ratios for the intermediate states³¹⁴: $\eta_{e \rightarrow L} = 0.6142$, $\eta_{e \rightarrow 1} = 0.2504$, $\eta_{e \rightarrow 0} = 0.1354$, and the following ones for the Rydberg states²⁶³: $\eta_{r \rightarrow L} = 0.894$, $\eta_{r \rightarrow 1} = 0.053$, $\eta_{r \rightarrow 0} = 0.053$. We use the branching ratios between different channels of intermediate-state scattering as reported in Refs.³¹⁴, and we additionally assume a simplified model in which all indirect paths (through $4D$ and $6S$) populate the ground-state manifold uniformly. The Rydberg lifetime has both radiative decay ($170 \mu\text{s}$) and black-body decay ($128 \mu\text{s}$) components, which we obtain by re-scaling the values in Ref.²⁶³ to $n = 53$. The microwave component results purely in atom loss and we assume that the radiative decay populates the ground-state manifold uniformly. We note that a more accurate treatment of the decay channels⁹⁸ could increase error modeling precision in future work.

2.2 PARAMETERIZED ENTANGLING GATES

Inspired by the methods used in Ref. ²²⁵ to find the time-optimal gates, we use optimal control to design gates with phase profiles given by simple analytical formulas. We find that this approach makes the gate experimentally robust and reproducible as small systematic offsets (e.g., rise time, atom-atom separation, etc.) can often be compensated for and are captured by a slightly different set of optimal parameters.

We focus on two main gates: one with a fixed amplitude and a phase profile similar to the time-optimal gate of Ref. ²²⁵, and a second one where the amplitude is also varied. We note that in the past, many schemes were proposed to implement two-qubit gates with Rydberg interactions ^{224,264,225,62,305 310,396,36,332,459,36} engineer robustness to errors ^{343,422,395,392,145,304,83,226,266,412}, and to perform multi-qubit gates ^{445,219,389,239,458,265,328}. The fixed-amplitude gate, which we refer to as “time-optimal” since it resembles and is only 0.2% slower than the Jandura-Pupillo gate ²²⁵, has the phase profile given by

$$\varphi(t) = A \cos(\omega t - \phi_0) + \delta_0 t. \quad (2.1)$$

with example parameters

$$\begin{aligned} A &= 2\pi \times 0.1122, & \omega &= 1.0431 \Omega, \\ \phi_0 &= -0.7318, & \delta_0 &= 0 \times \Omega, \end{aligned}$$

which constitutes an exact gate with time $(\Omega T/2\pi) = 1.215$, i.e., slightly longer than a resonant single-atom 2π pulse. Note that this set of parameters is not unique and other parameter values can also realize an exact gate, for example at non-zero detuning δ_0 as used in Fig. 2.8.

The smooth-amplitude gate has a varying phase and a varying Rabi frequency of the 420-nm

laser. We used optimal control methods in a three-level atomic system to find a gate that optimally suppresses scattering, for a fixed intermediate-state detuning and 1013-nm Rabi frequency. The scattering from the intermediate state was incorporated through a non-Hermitian Hamiltonian proportional to the scattering rate ($-i\gamma_e |e\rangle \langle e|$). Finally, we inferred an analytical form of the phase and amplitude profiles which are given by

$$\begin{aligned}\Omega_{420}(t)/\Omega_{1013} &= \Omega_0 + \Omega_1 \text{sech}[\omega_\Omega \tau]^\alpha, \\ \varphi(t) &= \delta_0 \tau + B \tanh(\lambda \tau),\end{aligned}$$

where $\tau = t - T/2$; in principle, one can also add a relative phase offset similar to ϕ_0 in Eq. (2.1) for further fine-tuning. This ansatz realizes an exact CZ gate for the gate parameters,

$$\begin{aligned}\Omega_0 &= 32.7403, & B &= 2\pi \times 0.2503, \\ \Omega_1 &= -31.1404, & \lambda &= 0.9372 \Omega, \\ \omega_\Omega &= 0.2668 \Omega, & \delta_0 &= -0.9491 \Omega, \\ \alpha &= -0.1131,\end{aligned}$$

which has a duration of $(\Omega T/2\pi) = 1.207$. We set the reference point such that for $\Omega_{420} = \Omega_{1013}$ the system is at two-photon resonance and has a two-photon Rabi frequency Ω . This smooth-amplitude gate has advantages of stronger intermediate-state scattering suppression and reduced off-resonant coupling to other states. On the other hand, due to a larger peak Rabi frequency, this gate is more susceptible to finite-blockade effects. The error budget for both gates can be found in Fig. 2.3.

Extending beyond two-qubit gates, we find that a slightly more general ansatz allows us to implement a nearly time-optimal three-qubit CCZ gate with the phase profile given by

$$\varphi(t) = A_1 \sin(\omega_1 \tau) + A_2 \sin(\omega_2 \tau) + B \tanh(\lambda \tau) + \delta_0 \tau,$$

where the Rabi frequency is kept constant and the parameters are,

$$\begin{aligned} A_1 &= 2\pi \times 2.1460, & \omega_1 &= 0.2101 \, \Omega, \\ A_2 &= 2\pi \times -0.0719, & \omega_2 &= 1.8957 \, \Omega, \\ B &= 2\pi \times -0.6432, & \lambda &= 0.6941 \Omega, \\ \delta_0 &= -1.3646 \, \Omega, \end{aligned}$$

which results in an exact gate with duration $(\Omega T/2\pi) = 1.751$. Finally, we note that the same methods can be directly extended to the design of CPHASE(θ) and CCPHASE(θ) gates, which is important in the context of digital quantum simulation²³¹. In Fig. 2.16c-d, we show how the two-qubit gate duration scales with the phase θ .

2.3 SCALING OF RYDBERG PROPERTIES WITH n

There are several effects that depend on n , including the finite Rydberg state lifetime ($\propto n^3$ for radiative decay, $\propto n^2$ for black-body decay), matrix elements influencing the 1013-nm Rabi frequency ($\propto n^{-3/2}$), interaction energy ($\propto n^{11}$), and sensitivity to electric fields ($\propto n^7$). Weighing these relative benefits, we work at $n = 53$, where Rydberg lifetime effects begin to become more relevant but where the matrix element is favorable for increasing intermediate-state detuning. A technical challenge relevant to this choice is that, due to the small blockade radius, we place atoms at $2 \, \mu\text{m}$ separa-

tion in order to achieve strong interaction strength $V_{\text{Ryd}}/2\pi \approx 450$ MHz. Operating at such close spacing is enabled by using high numerical aperture objectives (NA=0.65 from Special Optics), and allows us to pack many atoms into the entangling zone.

2.4 DESIGN CHOICES FOR RYDBERG EXCITATION

The Rydberg detuning signs and polarization signs are carefully selected for various optimizations: for suppressing 420-induced differential light shift between $|0\rangle$ and $|1\rangle$, we are red-detuned of $6P_{3/2}$ transition; for utilizing dark-state physics (nominally the phase profile corresponds to a certain sign of two-photon detuning) we thereby choose positive two-photon detuning; which, in turn then suppresses coupling to $m_J = -1/2$ by being primarily on the upper-side of $m_J = +1/2$; and finally, the 1013-nm light shift is lower (by $\sim 30\%$) at this single-photon detuning sign since there is a magic wavelength ~ 1 GHz red-detuned of $6P_{3/2}$ for the $|1\rangle \rightarrow 5S_{1/2}$ transition²⁵⁷. Two downsides of this detuning choice are that this choice of 420-nm polarization and detuning causes a vector light shift in the hyperfine ground-state manifold that causes the m_F levels to be pushed closer together, as opposed to further apart, which could exacerbate effects arising from 420-induced vector light shifts coupling adjacent m_F states in the ground state manifold (although negligible here), and the other downside is that the $m_J = -1/2$ Rydberg pair states are closer detuned to the two-photon excitation, and so we require a larger interaction strength to suppress their excitation, though the matrix element to these states is smaller.

2.5 DARK STATES IN TWO-PHOTON RYDBERG GATES

In this section we describe the physics associated with the three-level system present in the two-photon transition to the Rydberg state and derive how the Rydberg population can be realized through either the dark or bright states. The basic intuition can be developed at the single-particle




Figure 2.2: Parallel implementation of high-fidelity entangling gates on a neutral atom quantum computer. **a**, Entangling gates are implemented by arranging atoms into designated gate sites where they interact via Rydberg blockade interactions when pulsing global Rydberg lasers. Two-qubit or three-qubit gates are performed by modulating the Rabi frequency $\Omega(t)$ and phase $\phi(t)$ profiles of the laser driving the first leg of the two-photon Rydberg excitation. Inset: atoms in the qubit state $|1\rangle$ are excited to the Rydberg state $|r\rangle$ through an intermediate excited state $|e\rangle$, while atoms in $|0\rangle$ are not excited. Main gate error sources include Rydberg state decay γ_r , intermediate-state scattering γ_e , and Rydberg dephasing T_2^* . **b**, Numerical comparison of average bright and dark state populations during the Rydberg gate. Choosing opposite intermediate-state (Δ) and two-photon (δ) detuning signs at the beginning of the gate maximizes population in the dark state, which minimizes the intermediate-state scattering error. **c**, Entangling gates are implemented with a single Rydberg laser pulse with smooth phase modulation $\phi(t)$, whose slope corresponds to a two-photon detuning $\delta(t)$. A few global parameters characterize the gate, allowing for a family of possible gate implementations, including a parameterized version of the time-optimal CZ gate²²⁵ and a smooth-amplitude CZ gate. **d**, Example gate calibration sequence. Tuning individual parameters of the parameterized time-optimal gate phase profile $\phi(t) = A \cos(\omega t - \phi) + \delta_0 t$ allows for fast and simple global calibration (see Fig. 2.9 for additional experimental data). Error bars represent 68% confidence intervals.

Error source	Time optimal	Smooth amplitude	Error type X, Y, Z, LG, AL**
Scattering* $ 1\rangle$	0.103% / 0.043%	0.036%	6%, 6%, 25%, 47%, 15%
Scattering $ 0\rangle$	0.019%	0.025%	7%, 7%, 14%, 62%, 10%
Rydberg $T_1 = 88 \mu\text{s}$	0.113%	0.085%	2%, 2%, 6%, 23%, 67%
Rydberg $T_2^* = 3 \mu\text{s}$	0.134%	0.089%	0%, 0%, 75%, 0%, 25%
Position fluct.	0.012%	0.054%	0%, 0%, 96%, 0%, 4%
Rydberg $m_J = -\frac{1}{2}$	0.06 - 0.15%	0.01%	
Total fidelity	99.53 - 99.62%	99.70%	2%, 2%, 41%, 17%, 38%

Figure 2.3: Simulated error budget for two-qubit CZ gates from Ref. ¹²⁹. Simulations were performed at an intermediate-state detuning of $\Delta/2\pi = 7.8 \text{ GHz}$ and a two-photon Rabi frequency $\Omega/2\pi = 4.6 \text{ MHz}$. (*) The scattering error for two detuning signs (bright/dark). The total fidelity values assume the correct (dark) detuning choice. (**) LG — leakage out of the qubit manifold to other m_F states. AL — atom loss from population left in the Rydberg state. See Chapter 8 and Ref. ¹³⁰ for a more up-to-date error modeling accounting for additional error sources we recently understood in more detail.

level, where the system is described by the three-level Hamiltonian in the $\{|1\rangle, |e\rangle, |r\rangle\}$ basis,

$$H = \begin{pmatrix} 0 & \frac{\Omega_b}{2} & 0 \\ \frac{\Omega_b}{2} & -\Delta & \frac{\Omega_r}{2} \\ 0 & \frac{\Omega_r}{2} & -\delta \end{pmatrix}, \quad (2.2)$$

where we use symbols $\Omega_b := \Omega_{420}$ and $\Omega_r := \Omega_{1013}$ in this section for clarity of expressions.

We also assume that the amplitude and phase of the red 1013-nm laser are kept constant at all times and the blue 420-nm phase is captured by the time-dependent two-photon detuning $\delta := \delta(t) \sim -\phi'(t)$.

At large intermediate detunings ($|\Delta| \gg |\delta|, |\Omega_{b/r}|$), this system is conveniently described in the dark-state basis (as summarized in Fig. 2.4a-b) formed by the eigenstates of (2.2) at the two-photon resonance ($\delta = 0$), which to the leading order in Ω_r/Δ is,

$$\begin{aligned} |D\rangle &= -\frac{1}{\sqrt{1+\alpha^2}} |1\rangle + \frac{\alpha}{\sqrt{1+\alpha^2}} |r\rangle, \\ |B\rangle &= \frac{\alpha}{\sqrt{1+\alpha^2}} |1\rangle + \frac{\sqrt{1+\alpha^2}\Omega_r}{2\Delta} |e\rangle + \frac{1}{\sqrt{1+\alpha^2}} |r\rangle, \\ |E\rangle &= -\frac{\alpha\Omega_r}{2\Delta} |1\rangle + |e\rangle - \frac{\Omega_r}{2\Delta} |r\rangle, \end{aligned}$$

where $\alpha = \Omega_b/\Omega_r$. Note that the “dark state” $|D\rangle$ has no contribution from the intermediate state, the “bright state” $|B\rangle$ populates the intermediate state $\propto 1/\Delta^2$, and $|E\rangle$ is composed essentially entirely from $|e\rangle$.

For our purposes the initial state is always $|1\rangle$, which is subsequently dressed by the blue light to $|\tilde{1}\rangle$. This is because the amplitude rise time of the blue laser to its initial value of $\Omega_b(0)$ is at the timescale of 10 ns, which is much slower than the adiabaticity limit set by Δ , and much faster than

the two-photon Rabi frequency relevant for populating the Rydberg state; thus, the initial state indeed corresponds to

$$\begin{aligned} |\tilde{1}\rangle &= |1\rangle + \frac{\alpha\Omega_r}{2\Delta} |e\rangle \\ &= \frac{\alpha}{\sqrt{1+\alpha^2}} |B\rangle - \frac{1}{\sqrt{1+\alpha^2}} |D\rangle + O(\Delta^{-3}) |E\rangle, \end{aligned}$$

which is well supported on the $\{|D\rangle, |B\rangle\}$ states alone. Moreover, the excited state $|E\rangle$ is detuned from the other two by Δ , and all direct couplings to it are on the order of Ω_r/Δ ; thus any population transfer out of the $\{|D\rangle, |B\rangle\}$ manifold will be suppressed by $(\Omega_r\delta)^2/\Delta^4$, and the subsequent evolution of state $|\tilde{1}\rangle$ is described by an effective two-level system (Fig. 2.4b). In this picture, the energy splitting is set by the AC stark shift (diagonal terms) and the effective off-diagonal coupling is given by a combination of the two-photon detuning and diabatic terms (off-diagonal terms). Crucially, the Rydberg state population can be realized in many inequivalent ways; for example, for $\alpha = 1$, states of the form $\sqrt{1-\beta}|B\rangle + \sqrt{\beta}|D\rangle$ have the same Rydberg population for β and $\beta \rightarrow 1-\beta$ (where $\beta \in [0, 1]$) despite very different intermediate state contributions.

First, consider the case of the parameterized time-optimal gate where the blue Rabi frequency is kept constant throughout the duration of the gate ($\alpha(t) = 1$), up to the finite rise and fall times. The initial state is simply $(|D\rangle - |B\rangle)/\sqrt{2}$ and the Hamiltonian is equivalent to

$$\tilde{H}_{\text{to}} = \begin{pmatrix} 0 & -\delta/2 \\ -\delta/2 & \frac{\Omega_r^2}{2\Delta} \end{pmatrix},$$

where the magnetic field sign is decided by Δ , and the phase of the Rabi frequency is given by the sign of the two-photon detuning δ . The time evolution under this Hamiltonian (which corresponds to driving the two-photon transition) can be solved exactly for fixed δ , and the population in the

dark state is

$$P_D = \frac{1}{2} - \delta\Delta \frac{\Omega_r^2 \sin^2 \left(\frac{\sqrt{4\delta^2\Delta^2 + \Omega_r^4}}{4\Delta} t \right)}{2\delta^2\Delta^2 + \Omega_r^4},$$

which can go above or below $1/2$ depending on the relative sign of the detunings. More precisely, the Rydberg population is realized predominantly via the dark state when $\delta\Delta < 0$ ($P_D > 1/2$), i.e., when the intermediate-state detuning Δ and the two-photon detuning δ have *opposite* signs; for a time-dependent detuning the relevant sign is the one in the beginning of the pulse. In Fig. 2.4c, we plot the intermediate-state population and the Bloch sphere trajectories for the time-optimal gate at two different signs of the two-photon detuning δ . As expected, one of the trajectories realizes the Rydberg population through the dark state and as a result minimizes the intermediate-state population, leading to suppressed scattering.

The remaining scattering comes mostly from the large admixture of the bright state in the initial state and can be further reduced by utilizing a smooth amplitude profile, which starts at low blue Rabi frequency ($\alpha \ll 1$) and only later increases to larger values, as is the case in gate schemes based on the adiabatic passage³⁰⁵. Note that operating at a fixed lower α does not further reduce scattering since a larger admixture of the bright state is necessary to realize the same integrated Rydberg population as before. In Fig. 2.4d, we show the intermediate-state population and the effective Bloch sphere trajectory for the smooth-amplitude gate introduced in the previous section. This gate occupies the instantaneous dark state for a majority of its execution time, admixing only as much bright state as is necessary to realize the required Rydberg population. We find numerically that the degree of scattering suppression depends on the speed relative to the time-optimal gate (assuming fixed 1013-nm Rabi frequency): scattering is suppressed by a factor of 1.2 (relative to the time-optimal gate) if the two gates take the same time, and by a factor of roughly 2.5 if the smooth-amplitude gate is twice as long as the time-optimal gate. This tunability is useful for choosing a gate based on the




Figure 2.4: Dark state physics in two-photon Rydberg gates. **a**, In the far-detuned limit the three-level system can be effectively described as a two-level system with a “dark” state $|D\rangle$ and a “bright” state $|B\rangle$. The population in the excited state $|E\rangle$ is suppressed by a factor $\propto \Delta^{-4}$ and does not participate in system dynamics. **b**, The Hamiltonian of the bare atomic system and the effective two-level system, where α is the (time-dependent) ratio between the blue and red Rabi frequencies and $\dot{\alpha}$ is its time derivative. **c**, Intermediate-state population during the parameterized time-optimal gate in the dark and bright configurations together with their Bloch sphere trajectories in the $\{|D\rangle, |B\rangle\}$ basis. **d**, Intermediate-state population during the smooth-amplitude gate and its Bloch sphere trajectory in the instantaneous $\{|D\rangle, |B\rangle\}$ basis.

dominant error source: a slower gate could be more beneficial when scattering dominates while a faster gate can be used when T_2^* or Rydberg decay is the main source of errors.

We note that despite the presence of multiple intermediate states which are at slightly different detunings and couple with different Rabi frequencies, the dark-state picture remains valid. We find this is the case numerically and note the intermediate-state population plots in Fig. 2.4c-d include contributions from all intermediate states in a numerical model realistic for ^{87}Rb .

2.6 HIGH-FIDELITY BELL STATES

To characterize the CZ gates realized with this approach, we create arrays of 10 Bell pairs by arranging qubit pairs into separated gate sites and pulsing global Rydberg and Raman lasers. Using the parameterized time-optimal gate, we create a Bell state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ ²⁶⁴, which is then characterized by measuring the populations of $|00\rangle$ and $|11\rangle$, and the oscillation amplitude of the two-atom parity $\langle \sigma_1^z \sigma_2^z \rangle$ upon applying a global single-qubit $\pi/2$ -pulse of variable phase. We ex-

tract a raw Bell state fidelity of $98.0(2)\%$ exceeding previous work by $\approx 2\%$ ²⁶⁴, already suggesting a greatly improved gate fidelity. Since this Bell state fidelity appears to be dominated by state preparation and measurement (SPAM) errors, to characterize the fidelity of the entangling gate more systematically, we apply an odd-numbered train of CZ gates to repeatedly entangle and disentangle the pairs, and then characterize the fidelity of the final resulting Bell state^{125,333}. We fit the decreasing fidelity to an exponential decay to extract a CZ gate fidelity $F_{\text{CZ}} = 99.52(4)\%$.

Here we outline the method used for the Bell state data in this work. We measure the Bell state fidelity as the average of coherences and populations²⁶⁴. The coherence is extracted by measuring the amplitude of parity oscillations. The populations are calculated as the sum of the $|00\rangle$ and $|11\rangle$ states, which we correct for additional atom loss as described below. The Bell state populations can be overestimated due to atom loss contributing to the perceived detection of state $|11\rangle$ ²⁶⁴, since loss shows up identically as $|1\rangle$ in our state detection procedure. To account for this, here we measure the atom loss probability when applying the sequence of gates (by turning off the pushout of state $|1\rangle$ for state discrimination), to find the extra contribution of atom loss to the Bell state population. To perform this loss subtraction, we subtract the observed $|11\rangle$ (i.e., observed loss of both atoms) population during the loss measurement directly from the measured populations, as well as measuring the loss-per-atom-per-gate which can also contribute to state $|11\rangle$ by converting $|01\rangle$ and $|10\rangle$ to $|11\rangle$. We emphasize that this correction strictly lowers the measured gate fidelity (without applying this loss subtraction, the measured CZ gate fidelity on the raw Bell state fidelity data is extracted to be $99.57(4)\%$).

We next evaluate a SPAM-corrected Bell state fidelity from the measured raw Bell state fidelity of $98.0(2)\%$ after a single gate. To extract a SPAM-corrected Bell state fidelity, we first measure relevant SPAM errors. In particular, we measure a population of $99.6(1)\%$ in state $|0\rangle$ when we try to prepare into $|0\rangle$, and likewise, a population of $99.4(1)\%$ in state $|1\rangle$ after state preparation into $|1\rangle$. These measurements include additional effects from loss and imaging/pushout fidelity. Specif-




Figure 2.5: High-fidelity CZ gates characterized with Bell states. **a**, Parameterized time-optimal CZ gates are implemented on 20 atoms in parallel. **b**, Raw Bell state measurements upon application of a single CZ gate, with raw Bell state fidelity of 98.0(2)%. We estimate a SPAM-corrected fidelity of 99.4(4)% (not plotted). **c**, Circuit used to benchmark two-qubit gate fidelity by making a Bell state after an odd number of CZ gates interleaved with single-qubit X gates. **d**, Decay of Bell state fidelity after applying a variable number of CZ gates. A fidelity of 99.52(4)% per CZ gate is extracted from an exponential fit. The plot y-axis on (d) is rescaled by a SPAM correction factor which does not affect the extracted gate fidelity. Note that the same number of single-qubit gates are performed for each point, so that no normalization of single-qubit errors is necessary, which thereby results in a reduced y-axis intercept. Error bars represent 68% confidence intervals.

ically, there is a 0.35(5)% probability that an atom is lost during the sequence, and the gate itself causes 0.17% additional loss on top of this baseline loss. Our pushout fidelity of 99.83(1)% affects the measurement fidelity of $|1\rangle$ which we additionally correct for. From these measurements, we estimate the amount of population leaked into other m_F levels during state preparation, as well as the probability of atom loss both before and after the circuit. From these values, we follow the method described in Ref. ²⁶⁴ to extract a SPAM-corrected Bell state fidelity of 99.4(4)%.

2.7 BENCHMARKING WITH REPEATED GATES

As a separate characterization of the gate fidelity, we apply random global single-qubit rotations between sequences of CZ entangling gates. This method averages over different states involved in the entangling operation similar to randomized benchmarking^{284,148,449,21}. In the absence of errors during the gate sequence, a precisely calculated final single-qubit operation returns the qubit pair to their initial $|00\rangle$ state. Applying a sequence of up to 20 CZ gates with random single-qubit rotations in between, we fit the decaying state fidelity as a function of CZ gate number and extract $F_{\text{CZ}} = 99.54(2)\%$, consistent with the Bell state method. Not only do these methods agree quantitatively, but in practice we optimize the gate with this global randomized benchmarking method (Fig. 2.9) and find that the same exact parameters are optimal for generating Bell states. The qubits also acquire a single-particle phase during the CZ gate which this benchmarking approach eliminates by using X gates in between CZ gate pairs. Therefore, we employ a second method of global randomized benchmarking without these X gates, which allows for calibration of the single-particle phase (used for calibrating the Bell state measurement) and additionally benchmarks a gate fidelity of $F_{\text{CZ}} = 99.48(2)\%$ (Fig. 2.10).

We next demonstrate that our gate methods are versatile, where various pulse profiles can all realize a high-fidelity CZ gate. Specifically, we also realize and benchmark the smooth-amplitude gate and achieve a similar fidelity of $F_{\text{CZ}} = 99.55(3)\%$. Different gate implementations can be tailored to specific use-cases; for example, the smooth-amplitude gate strongly suppresses scattering even with a closer-detuned excitation, which can help achieve high gate fidelities in situations where laser intensity is limited.

We also explore the scalability of our approach to larger system sizes. Despite the fact that all calibration and control is done globally and not for individual gate sites, we find that the fidelity of the time-optimal gate is constant across the 10 individual gate sites within statistical error. This observa-




Figure 2.6: Comparing repeated gates benchmarking sequences. **a**, Benchmarking of the CZ- π -CZ sequence with global random rotations, which is insensitive to the single-particle phase. **b**, Benchmarking a standalone CZ gate with global random rotations, which enables separate calibration of the single-particle phase. **c**, The usual interleaved randomized benchmarking method utilizing random two-qubit Clifford gates (not performed in this work). **d**, Numerical simulation of the presented benchmarking methods, and the Bell-state-preparation method, utilizing the realistic error model developed in this work. All approaches give consistent results, with the Bell-state fidelity measurement lower-bounding the other curves.

tion of homogeneity across the array highlights the inherent potential for scalability: more gate sites do not increase the calibration overhead. Motivated by this observation, in Fig. 3d we extend to a 60-qubit system by using larger Rydberg beams (while maintaining the same intensity), and achieve a gate fidelity of $F_{CZ} = 99.48(2)\%$ with good homogeneity across the array.

2.8 PROBING CORRELATIONS IN GATES BENCHMARKING DATA

To understand the requirements for continued scaling and realizing high-fidelity operation in even larger system sizes, we analyze the physical error sources in the system. In particular, we compare observed gate fidelities to detailed modeling which utilizes two 8-level atomic systems with quantitative decoherence rates informed by experimental measurements. This modeling accounts for the remaining CZ gate infidelity and reveals four main error sources (Fig. 2.3): Rydberg decay, coupling to the other Rydberg m_J level, intermediate-state scattering, and our measured ground-Rydberg $T_2^* = 3 \mu s$ which is dominated by laser light shift fluctuations and finite atomic temperature. We




Figure 2.7: Global randomized benchmarking of the CZ gate. **a**, Circuit used to benchmark two-qubit gate fidelity with global randomized benchmarking, in which random rotations R_{rand} are sampled from a Haar-random distribution and a final rotation R_f is precomputed to return population to the initial product state $|00\rangle$ in the absence of gate errors. All 21 single-qubit rotations are applied for each data point, such that the number of single-qubit operations is independent of the number of two-qubit gates. **b**, Benchmarking CZ gates on 20 atoms in parallel. We plot the probability of return to the initial product state as a function of the number of applied CZ gates for both the parameterized time-optimal gate and the smooth-amplitude gate, extracting fidelities of 99.54(2)% and 99.55(3)%, respectively. **c**, The extracted time-optimal gate fidelity is consistent within statistical error across the 10 gate sites. **d**, Scaling to larger systems, we measure a comparable fidelity of 99.48(2)% for the time-optimal gate implemented on 60 qubits in parallel. **e**, Site-by-site analysis reveals homogeneous, high-fidelity gate performance across the two rows.




Figure 2.8: Robust calibration of gates as a function of rise time. Robustness to experimental offsets. For some systematic experimental offsets, such as finite rise time of the 420-nm Rydberg laser pulse, an exact CZ gate can still be found. The relative values of gate parameters for the time-optimal gate are plotted as a function of pulse rise time. For no rise time, the parameter values used here are: $A/2\pi = 0.0988$, $\omega/\Omega = 1.3629$, $\phi_0 = -2.6082$, and $\delta_0/\Omega = -0.0187$.




Figure 2.9: Empirical optimization of two-qubit gate fidelity. **a**, Calibrating gate parameters for the time-optimal gate, indicating the chronological sequence of calibrations performed prior to measurement in Fig. 3b. We scan individual gate parameters and measure the probability of return to the initial state after application of 20 entangling gates. **b**, Analogous calibration of gate parameters for the smooth-amplitude gate, which we performed in the sequence shown prior to the smooth-amplitude measurement. Additional calibration of other gate parameters was performed before these measurements.




Figure 2.10: CZ gate single-particle phase calibration and benchmarking. **a**, Digital circuit for global randomized benchmarking method used to calibrate single-particle phase, where a Z rotation can be performed after each CZ gate to compensate for the acquired phase. R_{rand} are single-qubit rotations sampled from a Haar-random distribution, and the 5 rotations R_f are computed to return the atom pair to the initial product state in the absence of gate errors. For the 0 CZ gates point, 20 random rotations are applied, as well as a final rotation precomputed to return population to the initial state. **b**, Experimental data used for calibrating the single-particle phase by optimizing the return probability $P_{|00\rangle}$ after 20 CZ gates (inset). For the optimal choice of ϕ , we extract a 99.48(2)% CZ gate fidelity, fitting to an exponential decay.

further analyze error sources by studying gate-site correlations in the experimental data. We observe that high-weight correlated errors are largely absent from our data, which suggests the feasibility of stable, large-scale operation (Fig. 2.11). Careful analysis reveals small growth in the covariance between neighboring gate sites (Fig. 2.12), which can result either from correlated detuning fluctuations (corresponding to our T_2^*) or from long-range interactions caused by a Rydberg atom decaying into an adjacent Rydberg state (corresponding to our finite Rydberg lifetime).

Informed by these microscopic error sources, we conclude the dominant challenge in maintaining high fidelity at an even larger number of parallel gate sites is to continually scale laser power and maintain beam homogeneity, as the other decoherence mechanisms appear to be independent of system size. However, we emphasize that even with the present laser parameters, system sizes can be directly increased to hundreds of qubits by shuttling atoms in and out of the entangling zone during a quantum circuit^{37?}, or by redirecting the beam to dynamically redefine the position of the entangling zone.

We further analyze our data to characterize whether gate errors across the array are correlated. We study the 20-atom and 60-atom global randomized benchmarking data, and consider the distribution of the number of errors that occur in each experimental shot, where an error is defined as whenever a qubit pair does not return to the initial $|00\rangle$ product state. In Fig. 2.11a,c, we plot the number of errors occurring in each shot as a function of number of CZ gates applied, where the mean number of errors grows due to the 0.5% error per CZ gate. We compare our data (bottom) with a model consisting of a Poissonian distribution of errors centered at the experimental mean (top). We find that the Poisson distribution model approximates our data, and large-scale correlated errors are not common in our system.

To analyze more quantitatively in a single plot, we average the data for all numbers of gates and plot the resulting distribution in Fig. 2.11b,d. We find that higher weight errors for both the 20 atoms and 60 atoms data are greatly suppressed. More quantitatively, we compare our data to the

average of the Poissonian distributions plotted in Fig. 2.11a,c, and find small deviation from the Poisson distribution. We find this data is better described by a model in which the CZ gate fidelity is sampled from a Gaussian distribution in each shot (which would arise e.g. from global shot-to-shot fluctuations in detuning, captured by our T_2^*).

In Fig. 2.12a-b, we plot the covariance matrix between gate sites for the return to the initial state $P_{|00\rangle}$, after 0 CZ gates and after 20 CZ gates, qualitatively observing evidence of small positive covariance between nearby gate sites. In particular, Fig. 2.12b,d shows the growth of covariance between neighboring gate sites as a function of the number of CZ gates applied. These observations are consistent with known physical effects related to our error budget, namely Rydberg lifetime and T_2^* . For example, decay of Rydberg atoms to nearby P states during the gate can cause detuning shifts due to strong long-range interactions between S and P Rydberg states (decaying as $1/R^3$), as well as hopping of the P state to nearby gate sites. Additionally, fluctuations in Rydberg 1013-nm beam intensity can give rise to shot-to-shot fluctuations in gate fidelity with a spatial dependence. Classical Monte Carlo simulations of these two phenomena (not plotted here) reveal that both error sources can result in nonzero covariance, which appears to be described by quadratic growth for a small number of gates but linear asymptotically. We note that the crosstalk between gate sites, on the scale of 10 kHz, should be negligible.

This measurement of covariance after 20 gates is a highly sensitive probe to small correlations between gate sites building up over the course of the circuit. After one applied gate, this covariance appears to be smaller than other main error sources, so these correlations will have little effect for quantum circuits in which atoms are involved in gates not just at a single gate site but across the entire entangling zone. Moreover, when running quantum circuits using atom transport⁴³, the $\sim 100 \mu\text{s}$ delay between subsequent gates would result in leftover Rydberg states being expelled from the array, completely suppressing the effect from Rydberg decay described above.

In later work in Ref.⁴², we measured the CZ gate fidelity in a repeated benchmarking sequence as a

function of the duration between the gates, and find that the gate fidelity in this array increases from 99.3% to 99.5% by increasing duration between gates to 100 μ s, as the Rydberg atoms decay or eject during that time. This also implies the reported gate fidelity in Ref. ¹²⁹ may have been impacted by this effect and thereby underestimating the maximum gate fidelity. Analogously, we observe that this gate fidelity reduction is removed by reducing atom density. In quantum circuits based on atom motion, the duration between gates is sufficiently long for these Rydberg P states to decay or eject⁷⁴, which natively fixes this issue, and consequently we do not observe such effects during quantum circuits.

2.9 THREE-QUBIT CCZ GATES

We next explore the generalization of these methods to multi-qubit gates. Using optimal control methods, we find a time-optimal CCZ gate and corresponding ansatz phase profile that realizes a native CCZ gate between three qubits^{264,306,242} in a time only 44% longer than the time-optimal CZ and faster than other known CCZ profiles²²⁵. This CCZ gate is realized experimentally by rearranging triplets of atoms into triangular gate sites, and applying the CCZ pulse profile with our global laser pulse. We characterize our CCZ gate using a sequence to repeatedly entangle and disentangle a three-qubit GHZ state, and subsequently measure the final GHZ state fidelity (see Fig. 2.14d for example GHZ states). While this approach does not constitute a rigorous benchmarking of the CCZ gate fidelity (which can be done using e.g. randomized benchmarking²⁴²), our data indicates high-performing three-qubit entangling gates across 21 qubits in parallel, consistent with a fidelity $F = 97.9(2)\%$. These optimal control methods extend to higher-qubit-number controlled-Z gates. We numerically search for and find fast gates for up to six qubits, with gate times significantly shorter than those required to decompose an N-qubit controlled-Z gate into $2N$ CZ gates and various single-qubit gates³⁹³. Generically, with these global pulses and Rydberg blockade, one can




Figure 2.11: Correlated errors in experimental shots. **a**, Distribution of errors in each experimental shot as a function of the number of CZ gates applied for the 20-atom data, showing qualitative agreement with a Poisson distribution centered at the experimental mean for number of errors in a shot. **b**, Histogram of the number of errors in a shot, averaging over all numbers of gates for the 20-atom data. We compare to one model assuming a Poissonian distribution of errors about the mean, finding some deviation from our data. In a second model, we consider that in each shot there is a slightly different gate fidelity, sampled from a Gaussian distribution with a mean of 99.54% and standard deviation of 0.3%. This second model appears to better capture our data. **c**, Repeating the analysis for the 60-atom data, we notably find no shot (out of the 5053 total repetitions) with 18 or more errors out of the 30 gate sites. Again, the data is qualitatively similar to a Poisson distribution model. **d**, Averaging over the 60-atom data for all numbers of gates, we find again a small quantitative deviation between the data and a model with a Poisson distribution of errors in each shot. The data appears to be better described by a model where in each shot, we sample fidelities from a Gaussian distribution with a mean of 99.5% and a standard deviation of 0.3%.




Figure 2.12: Correlations between gate sites. **a**, Covariance matrices for the 20-atom data after 0 gates and 20 gates, where local correlations appear after 20 gates. **b**, Covariance averaged over all neighbors, next nearest neighbors, and next next nearest neighbors, as a function of number of CZ gates applied. As a guide to the eye, data is fit to quadratic curves. **c**, Covariance matrices for 60-atom data for 0 gates and 20 gates. **d**, Plotting covariance for neighboring sites for the 60-atom data. Once again, the covariance between nearby sites exhibits small growth throughout the 20 CZ gates applied, in particular for the nearest neighbor sites.




Figure 2.13: Effect of P state Rydberg decay on measured CZ gate fidelity. CZ gate fidelity measured via randomized benchmarking¹²⁹. Infidelity due to Rydberg P states is removed by leaving sufficient time or distance between gates. Sparse corresponds to 2x larger separation between gate sites.

natively realize symmetric, diagonal gates (e.g. CPHASE gates as illustrated in Fig. 2.16)²²⁵ which are important for efficient realization of digital quantum simulation algorithms²³¹.

To find time-optimal gates for the multi-qubit controlled phase gates, such as the CCZ gate, we utilize the optimal control methods similar to Ref.²²⁵. The gates in Ref.²²⁵ are found by looking for two-qubit diagonal gates up to global single-qubit Z rotations. However, for more than two qubits there are several distinct ways of realizing the controlled-Z gates which are not connected by a Z rotation but rather by a general single-qubit rotation, and these various gate realizations can be different. We use the approach where the controlled phase flip is applied to the $|000\rangle$ state to find multi-qubit controlled Z gates for larger number of qubits; we present the obtained times in Fig. 4e and Fig. 2.14a. In Fig. 2.14b, we show the time-optimal pulse profiles of multi-qubit control-Z gates up to the 6-qubit CCCCCZ. Finally, we note that an analytical ansatz (defined in the section dedicated to parameterized gates) similar to that used for the CZ gate allows us to parameterize the three-qubit controlled-Z gate with only a marginal decrease in speed ($\Omega T/2\pi = 1.75$). In fact, we find that this ansatz can also realize the CCCZ gate and we expect simple generalizations to be capable of realizing these gates for even larger numbers of qubits.

To characterize the CCZ gate experimentally, we create a GHZ state using the circuit shown in Fig. 2.14c. In Fig. 2.14d, we generate GHZ states after application of two CCZ gates, with populations of 92.9(3)% and parity contrast of 89(1)%, giving a raw GHZ state fidelity of 90.9(6)% (without any loss subtraction). We calibrate the gate by repeatedly applying the CCZ- π -CCZ part of the circuit such that after 6 and 10 CCZ gates we generate GHZ states with reduced fidelity, and we observe a 2.1(2)% reduction in the raw GHZ fidelity as a function of number of CCZ gates applied.




Figure 2.14: Time-optimal pulses for multi-qubit controlled phase gates and GHZ state data. **a**, The execution time of a $C_N Z$ blockade gate as a function of the number of qubits. The N -qubit gates are realized by applying a phase flip to the $|0\rangle^{\otimes N}$ state, which is not equivalent to the method of Ref. ²²⁵ for $N > 2$. For the CCZ gate, applying a π phase to the $|111\rangle$ state ($|111\rangle \rightarrow -|111\rangle$), while leaving all other basis states invariant, is related by a global bit-flip to applying a relative π phase to the $|000\rangle$ state; however, the two implementations are not equivalent up to a global Z rotation, contrary to the two-qubit case. The time-optimal CCZ gate using the second approach realizes the CCZ gate $\approx 34\%$ faster with $(\Omega T/2\pi) = 1.72$, as compared to $(\Omega T/2\pi) = 2.61$ from the first approach. The two approaches are different because the relative phase of π is accumulated between different basis states, which have different rates of phase accumulation. In the case of applying $|111\rangle \rightarrow -|111\rangle$, the states with the slowest relative rate are $|111\rangle$ and $|011\rangle$ which are driven with the Rabi frequency of $\sqrt{3}\Omega$ and $\sqrt{2}\Omega$, respectively, resulting in the phase accumulation rate proportional to $(\sqrt{3} - \sqrt{2})\Omega \approx 0.32\Omega$. In contrast, when the relative phase is applied on the state $|000\rangle$, the smallest accumulation rate is given by the $|001\rangle$ state which is driven with the Rabi frequency Ω . In general, an arbitrary global single-qubit rotation at the end of the gate can be included to incorporate all of the above approaches in the optimization procedure. **b**, Time-optimal phase profiles (without analytic parameterization) for the $C_N Z$ gates up to 6 qubits realized by applying a phase flip to the $|0\rangle^{\otimes N}$ state. **c**, Circuit used to generate the GHZ state $(|000\rangle + |111\rangle)/\sqrt{2}$ after two CCZ gates. **d**, GHZ states measured experimentally upon applying this circuit to 7 three-qubit groups in parallel, with populations in $|000\rangle$ and $|111\rangle$ of 92.9(3)% and a parity contrast of 89(1)%, giving a raw GHZ state fidelity of 90.9(6)%.




Figure 2.15: Realizing fast CCZ gates. **a**, Entangling zone comprising three-qubit gate sites, with 21 qubits total. The triangular configuration enables strong, symmetric interactions between three qubits. **b**, Circuit used to calibrate the CCZ gate, by generating GHZ states $(|000\rangle + |111\rangle)/\sqrt{2}$ after applying 2, 6, and 10 CCZ gates. **c**, Phase profile used for implementing a parameterized CCZ gate with a single, fixed-amplitude pulse. **d**, Raw GHZ state fidelity as a function of CCZ gates applied, with an exponential decay fit $\propto 0.979(2)^{N_{CCZ}}$ (see Fig. 2.14d for example GHZ state characterization). **e**, Theoretical scaling of gate duration as a function of number of qubits for native Rydberg blockade multi-qubit gates (see Fig. 2.14b for corresponding phase profiles).

2.10 TUNABLE CONTROLLED PHASE GATES

Finally, to enable efficient digital evolution in the subsequent chapters, we implement tunable entangling controlled-phase gates (denoted CPHASE or CP) characterized by the angle θ (normalized such that $\theta=1$ is the CZ gate). We realize each CP gate using a single Rydberg laser pulse with constant intensity and a phase profile given by the cosine function, $A \cos(\omega t + \phi)$, and a constant two-photon detuning δ ¹²⁹. These parameters are numerically calculated for each θ using optimal control methods^{129,225,320}. For Floquet evolution, we combine two CP gates with global single-qubit X operations to realize entangling gates of the form,

$$ZZ(\theta) = e^{i\theta \frac{\pi}{4} Z \otimes Z} = \text{CP}[\theta/2] (X \otimes X) \text{CP}[\theta/2] (X \otimes X), \quad (2.3)$$

which are related to CP gates through single-particle terms. This approach is not only a robust way to remove the single-qubit terms but also ensures that atoms in the entangling zone that are not undergoing gates do not pick up a spurious phase.

To calibrate the entangling gates, we adapt an approach from Ref.²⁹⁷ which allows for measuring the entangling phase. We first initialize a pair of atoms, one in $|+\rangle_y = (|0\rangle + i|1\rangle)/\sqrt{2}$ and the other in $|0\rangle$. Then we apply a series of gates to the atom pair, which causes the atom in $|+\rangle_y$ to acquire a phase according to the magnitude of the entangling phase. Finally, we apply a single-qubit rotation in the appropriate axis to the atom initially prepared in $|+\rangle_y$ to bring it to $|0\rangle$ for a perfect gate. After expelling atoms in $|1\rangle$ with resonant pushout light, both atoms should be present only if the gates are perfect. This benchmarking sequence is sensitive to both the entangling phase θ and loss from the gate operation. We run this circuit with a fixed number of entangling gates, and scan gate parameters to optimize the gate performance on the experiment, in a similar approach to our CZ gate calibration described previously in Ref.¹²⁹. We measure the return probability after 20




Figure 2.16: Design of controlled-phase gates. **a**, Example phase profiles for different CPHASE gates. **b**, Gate parameter values as a function of the entangling phase of the CPHASE gate. **c**, Duration of a controlled-phase gate CPHASE(θ). The CZ gate ($\theta = \pi$) is the longest in this gate family. Since faster gates are expected to have higher fidelity, an average CPHASE gate should perform even better than the CZ reported in this work, which is an exciting perspective for near-term digital simulation. **d**, Plotting on a log-log plot, we see that for small angles θ , the gate time decreases with an approximate power law $\Omega T(\theta)/2\pi = 1.6 \times (\theta/2\pi)^{0.23}$. This suggests that applying very small phases can be costly and should be taken into account when designing digital simulation schemes. While these small-angle gates are time-optimal when applying a single, fixed-amplitude pulse, different approaches could perform better. Exploring other exact and approximate gate schemes, such as Rydberg dressing and a detuned 2π -pulse, in the small-angle regime is an interesting direction for future work.

CPHASE gates, for different values of the entangling phase. The return probability is higher for smaller entangling phases, owing to the shorter gate duration and reduced average Rydberg population.

2.1.1 AUTOMATED CALIBRATION

Due to the large range of gate angles used for the experiments in Chapter 4, we employed automated calibration routines which enable convenient calibration, either using an automated version of the parabola scan method previously used¹²⁹ or a Nelder-Mead optimization algorithm³¹² adapted for noisy data²¹¹.

We also perform automated calibration of the Rydberg beams, due to the importance of beam homogeneity for realizing uniform gate angles across the array. We use a flat tophat intensity profile generated using an SLM to maximize homogeneity across the array¹²³. To address local deviations in the beam intensity, we use pre-calculated “peak correction” phase masks which, upon addition to the base hologram, realize localized intensity peaks. We use automated noisy Nelder-Mead optimiza-




Figure 2.17: CPHASE gate benchmarking. **a**, Circuit used to benchmark and calibrate the entangling phase gates, adapted from the approach in ²⁹⁷. One atom is prepared in $|0\rangle$ and the other in $|+\rangle_y = (|0\rangle + i|1\rangle)/\sqrt{2}$ with a local Raman gate. Then a series of CPHASE gates are applied to the atom pair, which effectively rotates the phase of the initial $|+\rangle_y$ state. A precalculated final $Z(\phi)$ gate ensures that both atoms return to $|0\rangle$ in the absence of errors. Errors, for example coming from a miscalibrated entangling phase or qubit loss/leakage during the gate, will reduce the probability of finding both atoms in $|0\rangle$ at the end of the circuit (the return probability). **b**, Data used to extract the entangling phase in Fig. ??d, shown for two different CPHASE gates. For different numbers of gates applied, we scan the $Z(\phi)$ gate before the final local $\pi/2$ pulse in order to extract the amount of phase accumulated. The trend shows how much phase has been accumulated over the 10 gates, and the reduction in the peak return probability is a result of gate errors. Note that for a smaller number of gates, we use the same Raman pulse sequence as for the 10 gate sequence and only reduce the number of CPHASE gates.




Figure 2.18: Automated Rydberg beam calibration. **a**, Example calculated Rydberg beam profile and the effect of adding two separate peak-correction holograms. The local peak corrections are added to the base holograms with variable weights on the Rydberg beam SLMs, in addition to correcting for more global Zernike aberrations. **b**, Example automatic calibration routine showing the homogenization of our Rydberg beam across an array. (top) The peak-to-peak variations in row intensity across the array during calibration, as measured by the differential light shift on the hyperfine qubit. (bottom) Example of how the defocus and one of the peak corrections change during this automated calibration procedure. **c**, Examples of the measured light shift across the rows of our atom array for the two Rydberg beams, comparing the uniformity before and after the automated calibration procedure. We also ensure that the columns are uniform, although they are naturally more homogeneous due to the beam geometry.

tion to sequentially adjust both Zernike and peak corrections on the SLM. This active optimization procedure significantly improves the homogeneity of the intensity profile across the rows, and can be used in combination with passive approaches for removing aberrations^{468,123,87}.

2.1.2 SIMULATING TWO-QUBIT GATE ERROR SOURCES

Here we describe our numerical modeling of two-qubit gate error sources in Ref. ¹²⁹. We modeled scattering and Rydberg lifetime by performing a full density-matrix simulation of the two atoms with 8 levels (including 3 intermediate states). Our modeling also explicitly includes the coupling to the other ($m_J = -1/2$) Rydberg state, which is 24 MHz lower in energy and is driven with a Rabi frequency suppressed by a factor of 3 (due to Clebsch-Gordan coefficients). For the gate we assume approximately 20 ns min-max rise/fall times of our AOM pulse profile (in a Blackman profile), which has significant implications for the off-resonant coupling to the adjacent Rydberg state, in terms of whether it is adiabatically “dressed” or diabatically occupied. Since the impact of the

other Rydberg state on the gate fidelity depends on the details of the pulse power profile and degree of calibration, we report a range of values that is reasonable for the assumptions mentioned above.

Finite temperature is also assumed in our error modeling, where for our given temperature we calculate the position spread of the atom in a trap, and the corresponding fluctuation in interaction strength from the distance-dependent blockade interaction $V_{\text{Ryd}} \propto r^{-6}$. We note that finite temperature can also contribute to single-qubit dephasing through both velocity spread and photon recoil³⁵¹, but these effects are already encompassed by our single-qubit ground-Rydberg T_2^* measurements. The T_2^* can also have contributions from other phenomena such as electric field fluctuations and fluctuations in the 1013-nm light shift (which has a differential light shift of ≈ 20 MHz on the ground-Rydberg transition), and so for our simulations we simply use the measured $T_2^* = 3 \mu\text{s}$ assuming a Gaussian distribution of detunings.

We can also use our detailed microscopic error modeling, which reproduces similar fidelities as the measured 99.5%, to project future performance. To reach 99.9% fidelity, the sum of the errors needs to be suppressed to below 0.1%, which can be achieved for example by going 2x further detuned, having a 3x longer Rydberg lifetime (for instance, exciting to a higher n state), 2x longer T_2^* (note that dephasing error scales as $\propto 1/(\Omega T_2^*)^2$), and suppressing coupling to the other m_J state. This suppression can be achieved by applying a larger magnetic field, using the smooth-amplitude gate, or eliminating coupling altogether by exciting from a stretched state or through the $6P_{1/2}$ excited state. An alternative approach to reaching 99.9% fidelity could be going to 3x higher Rabi frequency (again while suppressing coupling to the other m_J state) and 2x larger detuning. Other unique opportunities toward future improvements include single-photon excitation to the Rydberg state which circumvents intermediate-state scattering but has higher Doppler and recoil sensitivity^{351,145,226}, and has been explored in a variety of contexts with both alkali^{291,407} and alkaline-earth(-like) atoms^{124,381,282}.

2.13 GATE CALIBRATION AND BENCHMARKING

We calibrate the gate using global echo randomized benchmarking method shown. In Fig. 2.9, we show an example sequential optimization of CZ gate parameters for the time-optimal gate and smooth-amplitude gate. Once found, these gate parameters are empirically optimal for all the other benchmarking methods such as the Bell state measurements, and are consistent from day to day.

The qubits additionally pick up a global single-particle phase during the gate, which we cancel here using global X gates between pairs of CZ gates for simplicity. For many applications, such as quantum error correction, our gates are naturally used in this configuration (as was done for the quantum circuits implemented in Ref.²). We also further calibrate and benchmark the CZ gate without the X gate, by performing randomized benchmarking composed of repeated application of CZ gates and random single-qubit rotations, as shown in Fig. 2.10a. Here the final several CZ gates and random rotations are calculated to return the qubit pair from the resulting entangled state back to the initial product state. We perform a single-qubit Z rotation after each CZ gate to compensate for the accumulated single-particle phase, which we calibrate as simply another parameter of the gate to scan and optimize. In addition to calibrating this single-qubit Z rotation, this circuit additionally benchmarks a gate fidelity of 99.48(2)% on 20 atoms. To measure the raw Bell state data, we utilized both the CZ gate calibration via the first method of global randomized benchmarking and this single-particle phase calibration, as two independent calibration stages.

We note that all of our randomized benchmarking methods utilize only global rotations for simplicity. The symmetry introduced by global rotations makes us less sensitive to certain types of errors⁹³, SWAP being an extreme example, which we expect to be negligible. Nonetheless, since the atoms are identical and placed very close together, there is a large degree of symmetry between the two qubits and we expect this global benchmarking procedure to faithfully capture our gate fidelity, which we confirm with numerical simulations. This is further confirmed by the fact that the experi-

mentally extracted fidelity is consistent between this method and the Bell state measurement. Quantitatively, we simulate all the benchmarking methods, including the full randomized benchmarking protocol, using the microscopic error model developed in this work. We find that all methods give consistent results, with the Bell state fidelity lower-bounding the other curves.

2.14 OUTLOOK

These results enabled a new era of high-fidelity digital circuits with neutral atoms, opening the door for explorations of large-scale quantum error correction with efficient parallel control of logical qubits, as well as digital quantum simulation. These applications will be described in the subsequent Chapters. Furthermore, based on our detailed microscopic understanding of error sources, we envisioned various paths to further improve gate fidelity. Such further improvements to the two-qubit gate fidelity are described in 8.

3

Digital quantum simulations of the Kitaev honeycomb model

Additional details for this chapter can be found in Appendix A and Ref.¹³⁰. Some descriptions are also adapted from notes written in collaboration with Giulia Semeghini¹³¹.

3.1 DIGITAL QUANTUM SIMULATION

Until recently, quantum simulations with atom arrays were realized through an “analog” modality discussed in the previous sections. An alternative method is digital or gate-based quantum simulation, in which quantum dynamics are decomposed into discrete operations, for example single-qubit and two-qubit gates. While analog simulations are powerful for quantum dynamics that are inherent to the platform, such as PXP dynamics in the case of Rydberg arrays, digital approaches could offer new avenues for enhancing the programmability for (i) state preparation, (ii) Hamiltonian evolution, and (iii) the types of observables that are accessible. Additionally, for applications requiring very low error rates (e.g., $< 10^{-10}$), such as a very precise quantum chemistry calculation, future fault-tolerant quantum processors can extend such digital simulations to incorporate quantum error correction. In this way, studying digital quantum simulation on current quantum devices allows us to begin exploring algorithms relevant for these future low-error regimes.

Digital quantum simulations of many different types of models have been actively explored in superconducting qubit and trapped ion processors. Recent progress in digital quantum processors based on atom arrays has pushed the scale, level of control, and fidelity such that realizing such digital quantum simulations has become possible. Atom arrays enable a few specific avenues. For one, parallel coherent transport of atoms opens opportunities for realizing dynamics involving non-local connectivity. Additionally, multi-qubit gates can be realized through Rydberg excitation of several atoms placed within the Rydberg blockade radius.

One area in which near-term digital quantum simulations can offer unique opportunities is in probing long-range entangled states and topological phases leveraging mid-circuit measurement. A paradigmatic case that we study in this Chapter is Kitaev’s honeycomb model²⁴³, which features a non-Abelian chiral spin liquid with a topologically non-trivial energy band. While certain Abelian and non-Abelian states have recently been prepared in analog and digital systems^{386,376,14,218}, the




Figure 3.1: Approach to digital simulation using reconfigurable atom arrays. Positions of atoms used for performing digital quantum simulation of the Kitaev model. At each step, the atoms are moved to a certain connectivity, and a controllable amount of evolution is applied using CPHASE gates of tunable angle. This process is repeated to realize a Floquet version of the Kitaev model.

Kitaev model and its extensions^{414,121} constitute a promising platform for exploring both static and dynamic properties of such exotic phases of matter⁴⁴³, including the dynamical interplay between topological order and strongly correlated fermions. In our approach, we will use auxiliary “ancilla” qubits to project a system of “data” qubits into a specific entangled state. Whereas long-range entangled states can be prepared in analog simulators, mid-circuit measurement provides a route to generating these states in a single step. After measuring the ancilla qubits, the system is projected into a state where 6-atom loop operators are randomly ± 1 valued. By using additional local addressing of individual atoms based on the ancilla measurement results, the system can be prepared in the special state where all these operators are $+1$ -valued. This scheme for measurement-based state preparation can be extended to other types of order, and is in general can be a very efficient way of generating entanglement in large quantum systems.

3.2 INTRODUCTION TO EXPERIMENTS WITH KITAEV’S HONEYCOMB MODEL

Quantum computers have the potential to fundamentally advance our ability to simulate strongly correlated many-body quantum systems, which can be challenging to simulate classically¹³⁷. Fermionic models are especially important due to their central role in understanding phenomena across chemistry, materials science, and fundamental physics¹². Recently, analog quantum simulators¹⁸² based on fermionic atoms in optical lattices have been advancing our understanding of the Fermi-Hubbard model^{451,388,78}. However, digital quantum simulation approaches offer new possibilities with improved control and programmability, providing access to a much wider range of models and, eventually, to fault-tolerant operation. Early experiments demonstrated key building blocks^{26,18,299} of digital fermionic encodings, and recent studies have extended the work to two dimensions^{95,313,216}.

Here, we realize digital quantum simulations of Kitaev’s honeycomb model using a reconfigurable atom array processor. Leveraging measurement-based topological state preparation, tunable Floquet circuits, and error detection, we probe the phase diagram and dynamics of this model. In particular, we prepare and characterize several low-energy states, including that of the non-Abelian spin liquid phase. We further utilize this model as an efficient fermion-to-qubit encoding, investigating fermion properties through quench dynamics and engineering strong interactions to realize the Fermi-Hubbard model on a square lattice.

Encoding fermions with qubits

Implementing fermionic models on gate-based quantum computers presents unique challenges associated with encoding the non-local nature of fermions within qubit systems. As a result, fermion-to-qubit encodings require either macroscopic operators to represent local fermion hopping²²⁹ or rely on long-range entangled ancilla states to preserve the locality of operations^{431,118,85}. The latter case requires preserving topological order throughout computations to accurately represent




Figure 3.2: Kitaev’s honeycomb model. The Kitaev honeycomb model is a paradigmatic example in topological and condensed-matter physics. In this model, qubits sit on the vertices of a honeycomb lattice, and interact under highly anisotropic interaction terms depending on the orientation of the link between the two qubits. Under this Hamiltonian, 6-qubit loop operators are conserved, and when these are all set to be +1, the model has a rich phase diagram featuring different Abelian and non-Abelian topological phases. There has been a decades-long quest to engineer these anisotropic interactions in Kitaev materials, but the model has yet to be convincingly realized, motivating this first experimental demonstration.

fermionic particles, and can additionally be related to models that host exotic states of matter with topological order⁴³¹.

3.3 REALIZING DIGITAL QUANTUM SIMULATIONS OF KITAEV’S MODEL

The key idea behind our approach is the use of long-range entanglement to encode fermionic statistics⁴³¹. In particular, while conventional fermion encodings²²⁹ require high-weight qubit operators, topologically ordered states contain long-range correlations that can natively encode these properties. The Kitaev model is a paradigmatic example of such a fermion encoding based on topological order²⁴³. Here qubits are located on the vertices of the honeycomb lattice and are effectively coupled through anisotropic two-qubit interactions,

$$K_{ij}^X = -J_X X_i X_j, \quad K_{ij}^Y = -J_Y Y_i Y_j, \quad K_{ij}^Z = -J_Z Z_i Z_j, \quad (3.1)$$

Encoding fermions with qubit devices is challenging!


Example: fermion particle exchange



Need to encode this non-local nature...


One approach (Jordan-Wigner transformation)

Local fermion operators $\rightarrow$ high-weight Pauli operators




Kitaev model: local fermion operators = local qubit terms!

Topological state $\rightarrow$ resource



Efficient fermion hopping



Key idea: **fermionic statistics** can be well-captured with **long-range entanglement**

Figure 3.3: Encoding fermions with qubits. Encoding fermionic particles in qubit systems is challenging, owing to the different commutation relationships of fermions and qubits. As an example, upon fermionic particle exchange, the wavefunction picks up a -1 sign, independent of the distance between the two fermionic particles. In order to encode fermions with qubits, then, this non-local nature must be captured. One conventional approach is to use a Jordan-Wigner mapping, where local fermionic operators become high-weight Pauli operators. However, in practice, this means that the operators to implement even simple nearest-neighbor hopping terms becomes prohibitively difficult in near-term devices. An alternative approach is to leverage long-range entanglement to encode the non-local nature of fermionic operators. For example, in the Kitaev model, local fermionic terms can be realized with only local qubit operations. We will leverage this encoding throughout this Chapter.




Figure 3.4: Fermion encoding description. Details of the fermion-to-qubit mapping and its origins. **a**, An example Jordan-Wigner (JW) ordering spanning a honeycomb lattice on a cylinder. In our case the string is constructed as a product of link operators (see text). **b**, The operators which are fixed to +1 (stabilizers) on the topological state include the hexagonal plaquettes and the loops around the cylinder. **c**, A local fermion hopping term between two sites results in a macroscopic JW operator. The underlying long-range entanglement enables reducing this operator to a local one by multiplying the string with the relevant stabilizers. **d**, Construction of complex-fermion hopping in terms of spin operators. The length-4 Pauli string cancels the particle-creation terms resulting from the length-2 hopping operator. **e**, Table describing operator mapping between qubits and fermions. Sites i, k belong to the even sublattice and j belongs to the odd sublattice. Sites i', j' complement i, j along the ZZ links.

acting along the three distinct link directions with interaction strengths J_X, J_Y and J_Z . The X_i, Y_i, Z_i are the Pauli operators acting on qubit i , and we refer to the three link orientations as XX, YY, and ZZ links, respectively. In such a system, there are long-range entangled states which correspond to the common eigenstates of the conserved, mutually commuting plaquette operators, $W_p = X_1 Z_2 Y_3 X_4 Z_5 Y_6$. When acting on such states, the spin interactions in Eq. (3.1) can be described as Majorana hopping operators, $ic_i c_j$, along appropriate links. The resulting system of Majorana fermions, c_i , located on the vertices of the lattice, realizes the fermionic mapping of Kitaev's model²⁴³.

For the experimental realization of this model, we utilize a neutral-atom quantum computer based on a dynamically reconfigurable atom array composed of 72 data qubits and 32 ancilla qubits. Using the apparatus described previously in Refs.^{123,43,129,41}, qubits are encoded in the long-lived hyperfine levels of ^{87}Rb atoms, individually trapped in optical tweezers, and can be reconfigured during computation while maintaining qubit coherence⁴³. Fully programmable global and local




Figure 3.5: Digital quantum simulations with reconfigurable atom arrays. **a**, The honeycomb lattice used in this work with 104 total atomic qubits and periodic boundary conditions along the shorter direction, forming a cylinder. The qubits are encoded in ^{87}Rb atoms and entangling gates are realized through excitation to interacting Rydberg states. To encode fermion statistics, we prepare a long-range entangled state characterized by hexagonal plaquette operators $W_p = X_1 Z_2 Y_3 X_4 Z_5 Y_6$ that commute with the unitary evolution and are therefore conserved. The encoded Majorana fermions live on vertices, at the ends of operator strings, and conventional (complex) fermions are formed by combining two Majoranas along a chosen link orientation. **b**, The long-range entangled state is prepared using mid-circuit measurement of ancilla qubits, and plaquettes are deterministically flipped to be +1 using conditional single-qubit gates (red circle). **c**, The Floquet evolution cycle consists of atom reconfiguration interspersed with tunable entangling gates and global basis changes. **d**, Accumulated entangling phase upon repeated application of $\exp(i\pi\theta[Z \otimes Z]/4)$ gates with a variable angle θ realized through fast, parameterized laser pulses. The values show extracted values of θ , and error bars represent one standard deviation.

single-qubit gates are implemented via fast Raman transitions⁴¹, while high-fidelity entangling gates are achieved by exciting atoms to interacting Rydberg states¹²⁹. A key new element involves the parallel implementation of a family of gates, $ZZ(\theta) = \exp(i\pi\theta[Z \otimes Z]/4)$, with tunable angle θ . We intersperse these gates with atom motion and single-qubit rotations to realize Floquet circuits, which consist of repeated digital evolution under the two-qubit interactions in Eq. (3.1) applied along appropriate links.

To efficiently create the topological order²⁸¹, ancilla qubits are used to project data qubits into the target long-range entangled state with cylindrical boundary conditions. This is achieved by measuring commuting plaquette operators W_p in two steps. First, ancilla qubits are used to measure weight-4 operators on one sublattice of the data qubits, initially all prepared in $|0\rangle$. Such a mid-circuit measurement projects the system into a toric code state²⁴⁵. In the second step, the two sublattices are entangled with parallel controlled-Y gates, completing the weight-6 plaquette operators required for this fermion encoding. Although the ancilla measurement outcomes are random ± 1 values, we apply conditional (feedforward) local single-qubit Z gates to deterministically prepare all plaquettes to have +1 value in the absence of errors.

To characterize the resulting state, we measure the plaquette operators W_p , finding an average parity of 0.444(4). Furthermore, we measure longer closed-loop operators, enclosing up to four hexagons, and observe finite parity in all cases. We can improve the quality of this topological order by performing error detection, by noting that for the cylindrical geometry used here, the product of ancilla values in a given column must have even parity in the absence of errors. Using this decoding postselection method, we observe an improvement in the plaquette parity to 0.57(1).

3.4 EXPLORING THE KITAEV MODEL

The highly anisotropic interaction structure of the Kitaev model, Eq. (3.1), gives rise to a rich phase diagram, featuring several distinct topological phases. Of particular interest is the non-Abelian spin liquid phase, known as phase B, which emerges when the three coupling strengths J_X, J_Y and J_Z are sufficiently uniform²⁴³. The fermionic description of the model results in its low-energy states being characterized by the underlying spin-liquid order and effective fermion correlations²⁴³. The initial state after feedforward corresponds to a fixed-point state A_Z^I of the Abelian A_Z phase ($J_X/J_Y=0$), where all ZZ-link operators should be +1. We observe a mean ZZ-link parity of 0.883(8), and the measured plaquette operators characterize the spin-liquid nature of the prepared state.

To evolve from this initial state (A_Z^I) to target states of interest in the phase diagram, we apply a numerically optimized version of a Floquet unitary circuit²³¹ constructed from repeated application of the nearest-neighbor gates, with optimized entangling phases. We find that a circuit depth of 6 is sufficient at this system size to probe all points of interest on the phase diagram (as characterized by Majorana correlations), consistent with the slow gap closing with system size²³¹. We measure the plaquette parity values to verify the preservation of the topological order and observe that they decay with increasing circuit depth, as expected due to accumulating incoherent errors. To mitigate this, we utilize a new type of state-selective readout, which converts atomic state into spatial position^{446,42}, allowing us to differentiate between lost atoms and the two computational states^{381,282}. By detecting atom loss and employing the error-detecting technique based on ancilla results, we observe that we can better preserve the plaquette expectation values during computation.

To characterize states across the phase diagram, we measure two-point Majorana fermion correlations, $\langle c_i c_j \rangle$, which are constructed as products of the link operators along the path connecting the two sites. Therefore, all fermion correlations can be measured as open Pauli strings, enabled by the underlying topological order. We measure the expectation values of strings up to length 6, for three




Figure 3.6: Measurement-based preparation of topological order. **a**, A long-range entangled state of the toric-code type is prepared using a depth-3 circuit, independent of the system size, accompanied by mid-circuit readout of the ancilla qubits (I). The measurement results are random up to parity constraints along the periodic direction. A feedforward step realized through FPGA-triggered single-qubit Z rotations pairs the -1 outcomes (white hexagons). Finally, a parallel controlled-Y operation creates the weight-6 plaquettes (hexagons) and initializes the ZZ-link operators (ovals) to all be +1 (II). **b**, Expectation values of the weight-6 plaquettes and weight-2 ZZ-links across the array of 72 data and 32 ancilla qubits with cylindrical boundary conditions. The inset shows values averaged across the system. **c**, Parity expectation values of increasingly large loops, including loops that enclose 1, 2, 3, and 4 hexagons within a column and the loop around the cylinder (plotted for maximum decoding postselection). The largest operator is equivalent to the product of two loops enclosing the cylinder. Error bars represent 68% confidence intervals. **d**, The product of plaquette operators in each column is equivalent to loops enclosing the cylinder, which are +1 in the absence of errors. Postselection based on the parity of the measured ancilla values within each column improves the plaquette expectation values and the quality of the resulting fermion encoding.




Figure 3.7: Measurement-based topological state preparation. **a**, Experimental layout, with gate orientation and an example plaquette shown, and ancilla qubit motions for preparing the ZXXZ surface code state on one half of the data qubits. We perform the periodic ancilla measurement step first, such that all other qubits are still in $|0\rangle$ and do not experience most of the Rydberg gate errors. **b**, Atom positions for mid-circuit readout and feedforward, with an example feedforward pattern applied (indicated by orange circles). The mid-circuit image is analyzed, and then based on our decoding algorithm, an FPGA outputs a series of digital TTLs which determine which atoms have a local Raman gate applied. **c**, Data qubit positions for the parallel controlled-Y (CY) operation, completing the state preparation sequence. The sublattice denoted by black circles gets picked up by AOD traps and shifted to the left to perform the parallel CY. This sublattice then remains in AOD traps for the Floquet circuit.

points in the phase diagram: the initial state after feedforward state preparation A_Z^I , another point in the Abelian phase A_Z^{II} , and a point in the non-Abelian phase B. We observe a buildup of string correlations as the system is moved towards phase B and find good agreement with theory predictions, under a noise model consisting of single-qubit Pauli errors and atom loss, across all prepared states.

We probe the nature of phase B by extracting the Chern number—a topological invariant that characterizes the curvature of the system’s energy bands^{35,7,99}. Noting that the prepared state should correspond to a low-energy state of a free-fermion Hamiltonian, we infer the momentum-space parent Hamiltonian H_k from the measured real-space correlations of the state. Because the unit cell of the honeycomb lattice consists of two sites, the system possesses two Bloch bands and their Chern number can be numerically evaluated from the learned Hamiltonian. As the system transitions from the Abelian to the non-Abelian phase, we observe a change from a zero ($C=0$) to a non-zero ($C=1$) Chern number of the lowest-energy band. In phase B, an odd Chern number, combined with the underlying topological order, provides strong evidence of the phase’s non-Abelian character. In particular, it guarantees that, in the presence of a flux, a non-Abelian, unpaired Majorana mode must be present in the spectrum at zero energy^{243,10}.

3.5 FLOQUET EVOLUTION OF FERMION SYSTEMS

Having explored the topological properties in the Kitaev phase diagram, we now leverage them to simulate quench dynamics^{109,406,456} of fermionic particles. We focus on conventional fermions characterized by complex-valued operators, which can be realized by pairing two Majorana fermions, $a_i = (c_i + ic_{i'})/2$. In our realization, we pair them along the ZZ links, such that the resulting complex fermions form a square-lattice geometry, and the $Z_i Z_{i'}$ operator along an (i, i') link is related to local fermion density, $n_i = (1 - Z_i Z_{i'})/2$, with $Z_i Z_{i'} = +1$ denoting particle absence, $n_i = 0$. In this framework, the ZZ-link operators characterize the initial vacuum state, corresponding to a background




Figure 3.8: Kitaev model on a honeycomb lattice. **a**, The honeycomb model, consisting of anisotropic spin interactions along the three directions of the honeycomb lattice ($J_X + J_Y + J_Z = 1$), exhibits topological order with three Abelian (A) phases and a single non-Abelian phase B. The toric code state is the fixed-point state A_Z^I of the Abelian A_Z phase ($J_Z = 1$). **b**, Starting from this initial state, we prepare different states on the phase diagram. Top: numerically optimized sequence of two-qubit gates used to prepare the non-Abelian phase B. Each circuit layer includes CPHASE (CP) gates and global single-qubit rotations. Bottom: plaquette parity during the evolution, with postselection based on atom loss and decoding. The plaquette parity is lower than due to a longer sequence of gates and atom movements. **c**, Pauli strings of different lengths measured on the three studied states and averaged over the bulk of the system. Error bars represent 68% confidence intervals. **d**, In the fermion representation, the link operators are proportional to nearest-neighbor Majorana hoppings, $K_{ij}^{X/Y/Z} \propto i c_i c_j$. Longer Pauli strings constructed from their products result in longer-range hopping operators. **e**, The free-fermion parent Hamiltonian H_k can be reconstructed from measured two-point Majorana correlations. **f**, The Chern number C is evaluated using the learned parent Hamiltonians of the string distributions in **c**, resulting in $C=0$ in phase A and $C=1$ in phase B. The non-Abelian phase B is characterized by the underlying topological order and an odd Chern number²⁴³.

fermion density of 5.9(5)%.

To initialize real-space configurations of complex fermions, we start from the state A_Z^I and create a fermion pair by applying an open Pauli string that commutes with the plaquette operators but flips the desired $Z_i Z_{i'}$ operators such that $n_i = 1$. Individual complex fermions can, in principle, be created by connecting such strings to the open boundary.




Figure 3.9: Probing fermion dynamics. **a**, System initialized in a two-particle state with complex fermions defined along the ZZ links. An open Pauli string is used to create a fermion pair out of the initial vacuum state. **b**, Digital evolution under two-qubit interactions in Eq. (3.1), realizing Majorana fermion hopping between connected vertices. The hopping strength along the XX and YY links is fixed, $J_X = J_Y = J$, while the value of J_Z/J is tuned. **c**, Total number of complex fermions in the system during a quench under two different Floquet drives, for $J_Z/J=1$ and $J_Z/J=8$. The K_{ij}^Z term corresponds to the fermion particle number and increasing it enforces particle number conservation. The initial background fermion number resulting from state preparation errors is subtracted for clarity and the solid lines are the result of noisy numerical simulations. **d**, System initialized with two fermions at distances $d=1$ or $d=2$ is evolved under the Floquet Hamiltonian with $J_Z/J=8$. **e**, A horizontal cut of the two-point density-density correlation relative to the marked site (green dot) at depth-11. The directional asymmetry in hopping is consistent with Pauli exclusion, which prevents fermions from occupying the same site. Error bars represent 68% confidence intervals.




Figure 3.10: Fermion hopping dynamics. **a**, Spatial distribution of complex-fermion density, n_i , after evolution for depth-11 from the initial two-fermion state. The density values are obtained as $n_i = (1 - Z_i Z_{i'} / \bar{z}\bar{z})/2$, where $\bar{z}\bar{z}$ is the background magnitude of the ZZ link operators in the initial state, $\bar{z}\bar{z} = 0.728(3)$. **b**, The particle creation of a depth-6 Floquet Hamiltonian $H_{F,6}$, with the evolution time of individual unitaries corresponding to that of the quench experiment, $Jt = 0.125$ (see Ref¹³⁰ for more details). The particle creation is quantified here as a matrix 2-norm of the commutator between the bulk Floquet Hamiltonian and the particle number operator, $\|[H_{F,6}, N_{\text{tot}}]\|_2$. **c**, A spatial plot of density-density correlations normalized by the density of the reference site (green dot), $G_{ij} / \langle n_i \rangle$, in a small neighborhood of the reference site.

The dynamics of complex fermions is inherited from that of the Majorana operators, where the Kitaev interactions in Eq. (3.1) correspond to nearest-neighbor Majorana hopping terms $c_i c_j$. We first study the dynamics of the two localized fermions under two different depth-12 Floquet circuits (Fig. ??b), focusing on the role of anisotropy between the couplings. Setting two directions to be equal, $J_X = J_Y = J$, and tuning the strength of the third direction J_Z , we find that for $J_Z = J$ the number of particles grows quickly, because the Majorana hopping terms do not naturally preserve the complex-fermion particle number. By contrast, for $J_Z = 8J$, background particle creation is suppressed and the total particle number returns close to the initial value for a full effective Floquet cycle (around depths 6 and 12). This emergent particle conservation is the result of strong K_{ij}^Z terms, which corresponds to the total particle number, projecting the dynamics into approximate particle-conserving evolution^{3,127}. Indeed, unitary dynamics result in an effective Floquet Hamiltonian closely resembling that of complex fermions, and the intermediate rise in particle number can be interpreted as Floquet micromotion.

Focusing on the particle number conserving ($J_Z = 8J$) case, we evolve two initial states, with a complex fermion pair either one ($d=1$) or two ($d=2$) sites apart, and investigate the final fermion configuration at depth 11. We probe their hopping by measuring density-density correlations $G_{ij} = \langle n_i n_j \rangle - \langle n_i \rangle \langle n_j \rangle$. Such correlations capture hopping since a transported particle should be anticorrelated with its initial position. For free-fermion states, the G_{ij} correlation is equal to the (negative) magnitude of the hopping operator between sites i and j . We explore these correlations with respect to one of the initially occupied sites, normalized by its density. For the $d=1$ initial separation, we observe a strong asymmetry in the hopping direction, owing to the fermions hopping away from each other, consistent with Pauli exclusion. As expected, the asymmetry is greatly reduced for the $d=2$ separation.




Figure 3.11: Fermion exchange statistics. **a**, Demonstration of the hopping gate H_1 for the two-fermion basis state (created by applying R twice) and the fermionic vacuum state. **b**, Additional data for the exchange protocol, showing additional intermediate steps of the process. During the middle three steps, the state of the system is described by a superposition of zero fermions and two fermions at different positions in the lattice.

3.6 FERMION EXCHANGE

A defining characteristic of fermionic particles is their non-trivial exchange statistics, with the many-body wavefunction being antisymmetric under particle exchange. We directly probe the exchange statistics of the complex fermions through deterministic hopping, realized by a gate sequence designed to preserve particle number. Because the exchange statistics manifest as a global phase, we perform a Ramsey-type experiment in which we first apply a partial pair-creation operator $R \propto \sqrt{X \otimes Z \otimes Y}$ to prepare a superposition state, $(|\emptyset\rangle + |\bullet\bullet\rangle)/\sqrt{2}$, where $|\emptyset\rangle$ is the initial vacuum state and $|\bullet\bullet\rangle$ is the two-fermion state. Then, we compare two different fermion movement patterns: (i) hopping once and then hopping back to the original configuration and (ii) fully exchanging the two fermions with each other and. Finally, we apply the partial creation operator R again and measure the resulting fermion density. When the fermions are returned to their original positions, the state remains unchanged, and the second R operator completes the fermion-pair creation, yielding $|\bullet\bullet\rangle$ as expected. However, when the particles are exchanged, we observe the absence of fermions at the end, consistent with the system returning to the vacuum state, $|\emptyset\rangle$. This represents direct demonstration of the superposition state transforming into $(|\emptyset\rangle - |\bullet\bullet\rangle)/\sqrt{2}$ due to the particle exchange, which leads to destructive interference with the final creation operation.




Figure 3.12: Fermion exchange statistics. **a**, Direct measurement of the fermion exchange phase. A superposition of vacuum and fermion-pair states is created with a sequence of local entangling gates, which realizes a partial creation unitary R . The system is evolved under two different particle-hopping protocols: (i) hopping to an intermediate site and hopping back and (ii) exchanging the two fermions. The evolution is followed by a second application of R , which maps the exchange phase to fermion presence. Color plots show the mean fermion density at each site in a 4x4 region embedded in the full lattice. **b**, Contrast between protocols (i) and (ii) as a function of decoding postselection. Error bars represent 68% confidence intervals.

To quantify this process, we evaluate the difference in average density on the target sites between the outcomes of the two protocols, corresponding to the contrast resulting from the fermionic Ramsey sequence, under varying levels of error detection. We find that the accuracy of the exchange statistics improves as the quality of fermion encoding is enhanced through decoding postselection.

3.7 SIMULATING THE FERMI-HUBBARD MODEL

The Fermi-Hubbard model is one of the most important strongly interacting fermionic systems. In its simplest form it consists of spin-1/2 particles with contact interactions between different fermionic spin states,

$$H_{\text{FH}} = \sum_{\sigma \in \uparrow, \downarrow} \sum_{\langle i, j \rangle} (a_{\sigma, i}^\dagger a_{\sigma, j} + a_{\sigma, j}^\dagger a_{\sigma, i}) + U n_\uparrow n_\downarrow, \quad (3.2)$$

where σ denotes the spin state, $\langle i, j \rangle$ is the set of edges encoding the system connectivity, and U is the interaction energy. This model, and its various extensions, provides a description of real-world materials and is believed to underpin some of the most puzzling phenomena in condensed matter, such as high-temperature superconductivity²⁶¹.

To simulate this model, we split our array into two halves, representing spin up on the left side and spin down on the right side, and engineer interactions by realizing parallel gates between the two halves. After the initial encoding step, we omit the hopping terms across the boundary during evolution, enabling independent dynamics of the two spin states. Furthermore, we extend our Floquet cycle by adding a fermion interaction term. Effectively, this realizes a Fermi-Hubbard model on a 16-site square lattice.

More specifically, to implement onsite density-density interactions, $H_{\text{int}} \propto n_\uparrow n_\downarrow$, we take advantage of long-range atom motion to couple distant parts of the system. We apply two parallel CNOT gates across the two halves of the array, interleaved with evolution under the K_{ij}^Z interaction on the




Figure 3.13: Two-component strongly interacting fermions. **a**, Two independent species of fermions, labeled “spin-up” and “spin-down”, are realized by dividing the system into two halves. The XX and YY links across the boundary are turned off to facilitate independent hopping dynamics; this system realizes a 4x4 lattice with two fermion species. **b**, Floquet evolution consists of independent hopping within the two halves followed by density-density interactions between the two spin states. **c**, The interactions are realized through a parallel entangling operation connecting the two halves of the system. Two parallel CNOT gates propagate a $ZZ(\theta)$ operator, applied to the spin-up half, into the interaction term. **d**, The initial state is a checkerboard pattern, staggered between the two fermion spin states with mean magnetization denoted by the colorbar. **e**, Staggered magnetization dynamics during a quench with and without the interactions. **f**, Final staggered magnetization as a function of interaction strength. Increasing interactions suppress the spin mixing dynamics. The solid lines are the result of noisy numerical simulations of the circuit. Error bars represent 68% confidence intervals.




Figure 3.14: Fermi-Hubbard implementation. **a**, Local gates used to prepare the checkerboard pattern. In order to correct the plaquettes, we pre-compensate with the decoding algorithm such that local gates both flip the plaquettes to all be +1 and also initialize the operators on ZZ-links to give the correct complex fermion configuration. **b**, Depiction of how the system is cut in half for the Fermi-Hubbard implementation. Atom moves are shown for the XX, YY, and ZZ terms (no periodic boundary conditions are used here unlike in earlier experiments). These gates comprise the rest of the Floquet circuit in addition to the interaction terms. Note that the middle links are turned off for the XX and YY gates to realize independent hopping on the two halves.

target half only. This K_{ij}^Z term is proportional to $Z_\uparrow Z_\uparrow$, and the CNOT gates propagate it to a four-body $(Z_\uparrow Z_\uparrow) \times (Z_\downarrow Z_\downarrow)$ term across the two halves, which effectively realizes contact interactions $(4n_\downarrow n_\uparrow - 2n_\downarrow - 2n_\uparrow)$ between the two fermion species.

We initialize an antiferromagnetic checkerboard spin ordering with local single-qubit gates, and study its dynamics with varying interaction strength. In the absence of interactions ($U=0$), we observe that the staggered magnetization of the initial state quickly decreases after a few Floquet cycles, due to the independent fermion hopping. However, by increasing the interactions, we observe suppressed particle transport and significantly slower decay of the staggered magnetization. We measure the staggered magnetization after the final Floquet round for a range of interaction strengths, and find good agreement with numerical simulations of our circuit, including a non-monotonic behavior as a function of interaction strength.

3.8 OUTLOOK

Our experiments demonstrate key building blocks for digital quantum simulations of complex fermionic and topological systems. Utilizing measurement-based preparation of long-range entanglement, efficient Floquet dynamics, and error detection, they pave the way for explorations of complex and practically relevant quantum systems. These studies can be extended in a number of interesting directions. For instance, properties of the non-Abelian excitations²³¹ can be explored in the Kitaev model, including fusion and braiding properties²⁴³. Moreover, the stability of spin liquid phases can be probed in challenging-to-simulate regimes by adding different types of interactions, which can be implemented with additional atom moves. Such simulations have the potential to provide new insights into the understanding of so-called Kitaev materials⁴¹⁴. Similar methods can be extended to study a broad range of physics and chemistry models, including those relevant to materials and molecules^{294,94}, lattice gauge theories⁹⁵, and quantum gravity^{286,367}. Other fermionic encodings can also be explored, either using qubits³¹³ or hybrid analog-digital architectures with fermionic atoms^{193,168}.

While the methods used in this work can readily be scaled up to thousands of qubits^{288,275}, the achievable circuit depths in digital quantum simulations will be limited due to errors. Our experiments demonstrate that error detection approaches can be used to suppress errors in near-term simulations. Eventually, error correction will need to be integrated to further extend achievable circuit depths. The topological state-preparation circuit can be realized with encoded logical qubits using transversal operations⁴⁶⁶, while implementation of specialized error-correcting codes could further reduce resource overheads⁴⁵³. The techniques developed and demonstrated in this work provide the foundation for co-designing such simulation approaches in a hardware-efficient manner.

4


Experiments in logical quantum processing

Additional details for this chapter can be found in Appendix B and Ref.⁴¹.

Quantum computers have the potential to significantly outperform their classical counterparts for solving certain problems³³⁸. However, executing large-scale, useful algorithms on quantum processors requires very low gate error rates (generally below $\sim 10^{-10}$)¹⁵⁷, far below those that will likely ever be achievable with any physical device³⁹⁹. The landmark development of quantum error correction (QEC) theory provides a conceptual solution to this challenge^{399,410,117}. The key idea is

Efficient control of logical qubits: control logical qubits as “big atoms”

Single-qubit transversal gates: apply single-qubit operations in parallel to physical qubits



Two-qubit transversal gates (e.g. CNOT): directly interact the delocalized degrees of freedom




Figure 4.1: Logical qubit control. At the basis of the experiments described in this Chapter is the ability to operate on encoded logical qubits efficiently. This is achieved by performing transversal gates in which the same operation is applied on a whole block of physical qubits in parallel. For example, single-qubit logical gates can be performed by shining light on a whole block of atomic qubits with standard AOD optical tools. Even more significantly, coherent transport of atomic qubits enables efficient two-qubit logical gates by directly interlacing the two logical qubits.

to use entanglement to delocalize a logical qubit degree of freedom across many redundant physical qubits, such that if any given physical qubit fails, it does not corrupt the underlying logical information. In principle, with sufficiently low physical error rates and sufficiently many qubits, a logical qubit can be made to operate with extremely high fidelity, providing a path to realizing large-scale algorithms¹¹⁷. However, in practice useful QEC poses many challenges, ranging from significant overhead in physical qubit numbers¹⁵⁷ to highly complex gate operations between the delocalized logical degrees of freedom¹⁴¹. Recent experiments have achieved milestone demonstrations of two logical qubits and one entangling gate^{357,170}, and explorations of novel encodings^{385,204,435,14}.




Figure 4.2: A programmable logical processor based on reconfigurable atom arrays. **a**, Schematic of the logical processor, segmented into three zones: storage, entangling, and readout. Logical single-qubit and two-qubit operations are realized transversally with efficient, parallel operations. Transversal CNOTs are realized by interlacing two logical qubit grids and performing a single global entangling pulse that excites atoms to Rydberg states. Physical qubits are encoded in hyperfine ground states of ^{87}Rb atoms trapped in optical tweezers. **b**, Fully programmable single-qubit rotations are implemented using Raman excitation through a 2D AOD; parallel grid illumination delivers the same instruction to multiple atomic qubits. **c**, Mid-circuit readout and feedforward. The imaging histogram shows high-fidelity state discrimination ($500\ \mu\text{s}$ imaging time, readout fidelity $\approx 99.8\%$), and the Ramsey fringe shows that qubit coherence is unaffected by measuring other qubits in the readout zone (error probability $p \sim 10^{-3}$). The FPGA performs real-time image processing, state decoding, and feedforward.

4.1 FROM CONTROL OF PHYSICAL ATOMS TO LOGICAL QUBITS

One specific challenge for realizing large-scale logical processors involves efficient control. Unlike modern classical processors that can efficiently access and manipulate many bits of information³²⁵, quantum devices are typically built such that each physical qubit requires multiple classical control lines. While suitable to the implementation of physical qubit processors, this approach poses a significant obstacle to the control of logical qubits redundantly encoded over many physical qubits.

Here, we describe the realization of a programmable quantum processor based on hardware-efficient control over logical qubits in reconfigurable neutral atom arrays⁴³. We use this logical processor to demonstrate key building blocks of QEC and realize programmable logical algorithms. In particular, we explore important features of logical operations and circuits, including scaling to large

codes, fault-tolerance, and complex non-Clifford circuits.

4.2 LOGICAL PROCESSOR BASED ON ATOM ARRAYS

Our logical processor architecture is segmented into three zones. The storage zone is used for dense qubit storage, free from entangling gate errors and featuring long coherence times. The entangling zone is used for parallel logical qubit encoding, stabilizer measurements, and logical gate operations. Finally, the readout zone enables mid-circuit readout of desired logical or physical qubits, without disturbing the coherence of the computation qubits still in operation. This architecture is implemented using arrays of individual ^{87}Rb atoms trapped in optical tweezers, which can be dynamically reconfigured in the middle of the computation while preserving qubit coherence^{37,43}.

Our experiments make use of the apparatus described previously in Refs.^{123,43,129}, with key upgrades enabling universal digital operation. Physical qubits are encoded in clock states within the ground-state hyperfine manifold ($T_2 > 1\text{s}$ ⁴³), and stored in optical tweezer arrays created by a spatial light modulator (SLM)^{123,380}. We utilize systems of up to 280 atomic qubits, combining high-fidelity two-qubit gates¹²⁹, enabled by fast excitation into atomic Rydberg states interacting through robust Rydberg blockade²²⁴, with arbitrary connectivity enabled by atom transport via 2D acousto-optic deflectors (AODs)⁴³. Central to our approach of scalable control, AODs^{380,113,402,176,282,277,316,381} use frequency multiplexing to take in just two voltage waveforms (one for each axis) to create large, dynamically programmable grids of light. Fully programmable local single-qubit rotations are realized via qubit-specific, parallel Raman excitation through an additional 2D AOD¹⁷⁷. Mid-circuit readout is enabled by moving selected qubits $\sim 100\text{ }\mu\text{m}$ away to a readout zone and illuminating with a focused imaging beam^{98,43}, resulting in high-fidelity imaging as well as negligible decoherence on stored qubits. The mid-circuit^{113,402,176,282,277,316} image is collected with a CMOS camera and sent to an FPGA for real-time decoding and feedforward.

The central aspect of our logical processor is the control of individual logical qubits as the fundamental units, instead of individual physical qubits. To this end, we observe that during a vast majority of error-corrected operations, the physical qubits of a logical block are supposed to realize the same operation, and this instruction can be delivered in parallel with only a few control lines. This approach naturally multiplexes with optical techniques. For example, to realize a logical single-qubit gate³⁹⁹, we use the Raman 2D AOD to create a grid of light beams and simultaneously illuminate the physical qubits of the logical block with the same instruction. Such a gate is transversal³⁹⁹, meaning that operations act on physical qubits of the code block independently. This transversal property further implies the gate is inherently fault-tolerant³⁹⁹, meaning that errors cannot spread within the code block, thereby preventing a physical error from spreading into a logical fault. Crucially, a similar approach can realize logical entangling gates^{399,117}. Specifically, we use the grids generated by our moving 2D AOD to pick up two logical qubits, interlace them in the entangling zone, and then pulse our single global Rydberg excitation laser to realize a physical entangling gate on each twin pair of the blocks. This process realizes a high-fidelity, fault-tolerant transversal CNOT in a single parallel step.

4.3 IMPROVING ENTANGLING GATES WITH CODE DISTANCE

A key property of QEC codes is that, for error rates below some threshold, the performance should improve with system size, associated with a so-called code distance^{117,141}. Recently, this property has been experimentally verified by reducing idling errors of a code¹⁷⁰. Neutral atom qubits can be idly stored for long times with low errors, and the central challenge is to improve entangling operations with code distance. Thus motivated, we realize a transversal CNOT gate using logical qubits encoded in two surface codes. Surface codes have stabilizers that are used for detecting and correcting errors without disrupting the logical state^{117,141}. The stabilizers form a 2D lattice of 4-body

plaquettes of X and Z operators, which commute with the X_L (Z_L) logical operators that run horizontally (vertically) along the lattice. By measuring stabilizers one can detect the presence of physical qubit errors, decode (infer what error occurred), and correct the error simply by applying a software Z_L/X_L correction¹⁴¹. Such a code can detect and correct a certain number of errors determined by the linear dimension of the system (the code distance d).

To test the performance of our logical entangling gate, we first initialize the logical qubits by preparing physical qubits of two blocks in $|+\rangle$ and $|0\rangle$ states, respectively, and performing a single round of stabilizer measurements with parallel operations⁴³. While this state preparation is non-fault-tolerant (nFT) beyond $d = 3$, we are still able to probe error suppression of the transversal CNOT. Specifically, we prepare the two logicals in state $|+_L\rangle$ and $|0_L\rangle$, perform the transversal CNOT, and then projectively measure to evaluate the logical Bell state stabilizers $X_L^1 X_L^2$ and $Z_L^1 Z_L^2$. For decoding and correcting the logical state, we observe there are strong correlations between the stabilizers of the two blocks due to propagation of physical errors between the codes during the transversal CNOT³⁸. We utilize these correlations to improve performance by decoding the logical qubits jointly, realized by a joint decoding graph that includes edges and hyperedges connecting the stabilizers of the two logical qubits. Using this correlated decoding procedure, we measure ≈ 0.95 populations in the $X_L X_L$ and $Z_L Z_L$ bases, showing entanglement between the $d = 7$ logical qubits.

Studying the performance as a function of code size reveals that the logical Bell pair improves with larger code distance, demonstrating improvement of the entangling operation. In contrast, we note that when conventional decoding, i.e. independent minimum-weight perfect matching within both codes¹¹⁷, is used, the fidelity decreases with code distance. This is in part due to the nFT state preparation, whose effect is partially mitigated by the correlated decoding.

We emphasize that while these results demonstrate surpassing an effective threshold for the entire circuit (implying we surpass the threshold of the transversal CNOT), such a threshold is higher due to projective readout after the transversal CNOT. In practice, the transversal CNOT should be used




Figure 4.3: Transversal entangling gates between two surface codes. **a**, Illustration of transversal CNOT between two $d = 7$ surface codes based on parallel atom transport. **b**, The concept of correlated decoding. Physical errors propagate between physical qubit pairs during transversal CNOT gates, creating correlations that can be utilized for improved decoding. We account for these correlations, arising from deterministic error propagation (as opposed to correlated error events), by adding edges and hyperedges that connect the decoding graphs of the two logical qubits. **c**, Populations of entangled $d = 7$ surface codes measured in the XX and ZZ basis. **d**, Measured Bell pair error as a function of code distance, for both conventional (top) and correlated (bottom) decoding. We estimate Bell error with the average of the ZZ populations and the XX parities. To reduce code distance we simply remove selected atoms the grid, as shown on the right, ensuring unchanged experimental conditions (for $d = 3$, four logical Bell pairs are generated in parallel). Error bars represent standard error of the mean.




Figure 4.4: Fault-tolerant logical algorithms **a**, Circuit for preparation of logical GHZ state. Ten color codes are encoded non-fault-tolerantly (nFT), and then parallel transversal CNOTs between computation and ancilla logical qubits perform FT initialization. The ancilla logical qubits are moved to storage, and a 4-logical-qubit GHZ state is created between the computation qubits. Logical Clifford operations are applied before readout to probe the GHZ state. **b**, State-preparation-and-measurement (SPAM) infidelity of the logical qubits without (nFT) and with (FT) the transversal-CNOT-based flagged preparation, compared to physical qubit SPAM. **c**, Logical GHZ fidelity without postselecting on flags (nFT), postselecting on flags (FT), and postselecting on flags and stabilizers of the computation logical qubits, corresponding to error detection (EDFT). **d**, GHZ fidelity as a function of sliding-scale error detection threshold (converted into the probability of accepted repetitions) and of number of successful flags in the circuit. **e**, Density matrix of the 4-logical-qubit GHZ state (with at most 3 flag errors) measured via full state tomography involving all 256 logical Pauli strings.

in combination with many repeated rounds of noisy syndrome extraction¹⁷⁰, which is expected to have a somewhat lower threshold and is an important goal for future research.

4.4 FAULT-TOLERANT LOGICAL ALGORITHMS

All logical algorithms we perform in this work are built from transversal gates which are intrinsically fault-tolerant³⁹⁹. We now also use fault-tolerant state preparation to explore programmable logical algorithms. We use two-dimensional $d = 3$ color codes^{410,47}, which are topological codes akin to the surface code, but with the useful capability of transversal operations of the full Clifford group: Hadamard (H), $\pi/2$ phase (S) gate, and CNOT⁴⁷. This transversal gate set can realize any Clifford circuit fault-tolerantly. As a test case, here we create a logical GHZ state. Specifically, we implement a 10-logical-qubit algorithm, in which all ten qubits are first encoded by a nFT encoding circuit. Then, five of the codes are used as ancilla logicals, performing parallel transversal CNOTs in order to fault-tolerantly detect errors on the computation logicals¹⁷¹, and are then moved into the storage zone where they are safely kept. Subsequently four computation logicals are used to prepare the GHZ state, and logical Clifford rotations are used at the end of the circuit for direct fidelity estimation¹⁴⁰ and full logical state tomography.

We first benchmark our state initialization^{125,336,357}. Averaged over the five computation logicals, we find that by using the fault-tolerant initialization (postselecting on the ancilla logical flag not detecting errors) our $|0_L\rangle$ initialization fidelity is $99.91^{+0.04}_{-0.09}\%$, exceeding both our physical qubit $|0\rangle$ initialization fidelity ($99.32(4)\%$ ¹²⁹) and physical two-qubit gate fidelity (99.5% ¹²⁹). Then, Fig. 3c shows the resulting GHZ state fidelity obtained using the fault-tolerant algorithm is $72(2)\%$ (again using correlated decoding), demonstrating genuine multipartite entanglement. Additionally, one can postselect on all stabilizers of our computation logicals being correct; using this error detection approach, the GHZ fidelity increases to $99.85^{+0.1}_{-1.0}\%$ at the cost of postselection overhead.

Since not all nontrivial syndromes are equally likely to cause algorithmic failure, one can perform a partial postselection where syndrome events most likely to have caused algorithmic failure are discarded, given by the weight of the correlated matching in the whole algorithm. Figure 3d shows the measured GHZ fidelity as a function of this sliding threshold converted into a fraction of accepted experimental repetitions, continuously tuning the tradeoff between the success probability of the algorithm and its fidelity; e.g., discarding just 50% of the data improves GHZ fidelity to $\approx 90\%$. (As discussed below, for certain applications purifying samples can be advantageous in improving algorithmic performance.) Finally, fault-tolerantly measuring all 256 logical Pauli strings, we perform full GHZ state tomography.

The use of the zoned architecture directly allows scaling circuits to larger numbers, without increasing the number of controls, by encoding and operating on logical qubits, moving them to storage, and then accessing storage as appropriate. We make ten color codes and operate on with parallel transversal CNOTs, moved to storage, and then more qubits are accessed from storage. Repeating this process four times, we create 40 color codes with 280 physical qubits, at the cost of slow idling errors of $\sim 1\%$ logical decoherence per additional encoding step. These storage idling errors primarily originate from global Raman π pulses applied for dynamical decoupling of atoms in the entangling zone, which could be significantly reduced with zone-specific Raman controls.

Since mid-circuit readout^{113,402,176,282,277,316} is an important component of logical algorithms, we next demonstrate a fault-tolerant entanglement teleportation circuit. We first create a three-logical-qubit GHZ state $|0_L0_L0_L\rangle + |1_L1_L1_L\rangle$ from fault-tolerantly prepared color codes. Mid-circuit X -basis measurement of the middle logical creates $|0_L0_L\rangle + |1_L1_L\rangle$ if measured as $|+_L\rangle$, and $|0_L0_L\rangle - |1_L1_L\rangle$ if measured as $|-_L\rangle$. One recovers $|0_L0_L\rangle + |1_L1_L\rangle$ by applying a logical S gate to the first and third logicals conditioned real-time on the state of the middle logical, akin to the magic state teleportation circuit¹⁴¹. Measurements indicate that while $\langle X_L X_L \rangle$ and $\langle Y_L Y_L \rangle$ indeed vanish without the feedforward step, applying the feedforward correction we recover a Bell state




Figure 4.5: Zoned logical processor: scaling and mid-circuit feedforward. **a**, Atoms in storage and entangling zones, and approach for creating and entangling 40 color codes with 280 physical qubits. **b,c**, 40 color codes are prepared with an nFT circuit, and then 20 transversal CNOTs are used to fault-tolerantly prepare 20 of the 40 codes, whose fidelity is plotted. Logical decoherence is smaller than the physical idling decoherence experienced during the encoding steps. **d**, Mid-circuit measurement and feedforward for logical entanglement teleportation. The middle of three logical qubits is measured in the X -basis, and by applying a mid-circuit conditional, locally pulsed logical S rotation on the other two logical qubits, the state $|0_L 0_L\rangle + |1_L 1_L\rangle$ is prepared. **e**, Measured logical qubit parity with and without feedforward, showing that feedforward recovers the intended state with Bell fidelity of 77(2)% (ZZ parities of 83(4)% not plotted). No mid-circuit refers to turning off the mid-circuit readout, and postselecting on the middle logical being in state $|+_L\rangle$ in the final readout. By postselecting on perfect stabilizers of only the two computation logicals (error detection in the final measurement) the feedforward Bell fidelity is 92(2)% (not plotted). In **d** three of the additional blocks are flag qubits and the other four are prepared but unused for this circuit.

fidelity of 77(2)%, limited by imperfections in the original underlying GHZ state. By repeating this experiment without mid-circuit readout and instead post-selecting on the middle logical being in $|+_L\rangle$, we find a similar Bell fidelity of 75(2)%, indicating high-fidelity performance of the readout and feedforward operations.

4.5 COMPLEX LOGICAL CIRCUITS USING 3D CODES

One important challenge in realizing complex algorithms with logical qubits is that universal computation cannot be implemented transversally¹²². For instance, when using 2D codes such as the surface code, non-Clifford operations cannot be easily performed⁴⁷, and relatively expensive techniques are required for non-trivial computation^{141,64} as Clifford circuits can be easily simulated¹. In contrast, 3D codes can transversally realize non-Clifford operations, but lose the transversal H ⁴⁷. However, these constraints do not imply that classically hard or useful quantum circuits cannot be realized transversally or efficiently. Motivated by these considerations, we explore efficient realization of classically hard algorithms that are co-designed with a particular error-correcting code. Specifically, we implement fast scrambling circuits using small 3D codes, which are used for native non-Clifford operations (CCZ).

We focus on small 3-dimensional $[[8,3,2]]$ codes^{71,428,204,435}, which have various appealing features. They encode three logicals per block, feature $d = 2$ ($d = 4$) in the Z basis (X basis), implying error detection (correction) capabilities for Z (X) errors, and can realize a transversal CNOT between blocks. Most importantly, by using physical $\{T, S\}$ rotations (T is $\pi/4$ phase gate) one can realize transversal $\{CCZ, CZ, Z\}$ gates on the logical qubits encoded within each block, as well as intra-block CNOTs by physical permutation^{204,264}. Although transversal H is forbidden, initialization and measurement in either the X or Z basis effectively allows H at the beginning and end of the circuit.




Figure 4.6: Complex logical circuits using 3D codes. **a**, $[[8,3,2]]$ block codes can transversally realize $\{CCZ, CZ, Z, CNOT\}$ gates within each block, and transversal CNOTs between blocks. By preparing logical qubits in $|+_L\rangle$, performing layers of $\{CCZ, CZ, Z\}$ alternated with interblock CNOTs, and measuring in the X -basis, we realize classically hard sampling circuits with logical qubits. **b**, Measured sampling outcomes for a circuit with 12 logical qubits, 8 logical CZs, 12 logical CNOTs, and 8 logical CCZs. By increasing error detection, the measured distribution converges toward the ideal distribution. **c**, Circuit involving 48 logical qubits with 228 logical CZ / CNOT gates and 48 logical CCZs. **d**, Classical simulation runtime for calculating an individual bitstring probability; bottom plot is estimated based on matrix multiplication complexity. **e**, Measured normalized XEB as a function of sliding-scale error-detection for 3, 6, 12, 24, and 48 logical qubits. For all sizes we observe a finite XEB score which improves with increased error detection. Diagram shows 48-logical connectivity, with logical triplets entangled on a 4D hypercube. **f**, Scaling of raw (red) and fully error-detected (black) XEB from **e**. Physical upper-bound fidelity (blue) is calculated using best measured physical gate fidelities. Diagrams show physical connectivity. $[[8,3,2]]$ cubes are entangled on 4D hypercubes, realizing physical connectivity of 7D hypercubes.

We use these transversal operations to realize logical algorithms that are difficult to simulate classically^{296,322}. More specifically, these circuits can be mapped to Instantaneous Quantum Polynomial (IQP) circuits^{59,296,322}. Sampling from the output distribution of such circuits is known to be classically hard in certain instances⁵⁹, implying a quantum device can be exponentially faster than a classical computer for this task.

We prepare all logical blocks in $|+\rangle_L$, implement a scrambling circuit with 28 logical entangling gates, and then measure all logicals in the X -basis. We measure the probability of observing each of the $2^{12} = 4096$ possible logical bitstring outcomes, showing that as we progressively apply more error detection (i.e. postselection) in postprocessing, the distribution more closely reproduces the ideal theoretical distribution. To characterize the distribution overlap, we use the cross-entropy benchmark (XEB)¹⁷, which is a weighted sum between the measured probability distribution and the ideal calculated distribution, normalized such that $\text{XEB}=1$ corresponds to perfectly reproducing the ideal distribution, and $\text{XEB}=0$ corresponds to the uniform distribution which occurs when circuits are overwhelmed by noise. The 12-logical-qubit circuit XEB increases from 0.156(2) to 0.616(7) upon applying error detection. We note that XEB should be a good fidelity benchmark for IQP circuits.

We next explore scaling to larger systems and circuit depths. To ensure high complexity of our logical circuits, we use nonlocal connections to entangle the logical triplets on up to 4D hypercube graphs, which results in fast scrambling²⁵⁶. Exploring entangled systems of 3, 6, 12, 24, and 48 logical qubits, in all cases we find a finite XEB score which improves with increased error detection. The finite XEB indicates successful sampling, and the improvement with error detection shows the benefit of using logical qubits. While this improvement comes at the cost of measurement time due to error detection, improving the sample quality cannot be replaced by simply generating more samples. Thus, improving the XEB score yields significant practical gains. We obtain an XEB of ≈ 0.1 for 48 logical qubits and hundreds of nonlocal logical entangling gates, up to roughly an order of

magnitude higher than previous physical qubit implementations of similar complexity^{17,447}, showing the benefits of a logical encoding for this application.

Assuming our best measured physical fidelities, the estimated upper-bound for an optimized physical qubit implementation in our system is also significantly below the measured logical XEB. In attempting to run these complex physical circuits, in practice we find that realising non-vanishing XEB is significantly more challenging; we confirm with small physical instances that we measure values well below this upper-bound. In addition to the error-detecting benefits, it appears the logical circuit is significantly more tolerant to coherent errors, exhibiting operation that is inherently digital, just with imperfect fidelity, consistent with theoretical predictions⁵⁶. We also note that for the logical algorithms we optimize performance by optimizing the stabilizer expectation values (rather than the complex sampling output), providing further advantage for logical implementations.

Our 48-logical circuit, corresponding to a physical qubit connectivity of a 7D hypercube, contains up to 228 logical two-qubit gates and 48 logical CCZ gates. Simulation of such logical circuits is challenging due to the high connectivity (rendering tensor networks inefficient) and large numbers of non-Cliffords⁵⁴. To benchmark our circuits, we structure them such that we can leverage an efficient simulation method which takes ≈ 2 seconds to calculate the probability of each bitstring. Modeling noise in our logical circuits is even more complicated, as they are composed from 128 physical qubits and 384 T gates, thereby making experimentation with logical algorithms necessary to understand and optimize performance.

4.6 QUANTUM SIMULATIONS WITH LOGICAL QUBITS

Finally, we explore the use of logical qubits as a tool in quantum simulation, probing entanglement properties of our fast scrambling circuits, potentially related to complex systems such as black holes^{384,256}. In particular, we utilize a Bell-basis measurement made on two copies of the quan-




Figure 4.7: Logical two-copy measurement. **a**, Identical scrambling circuits are performed on two copies of 12 logical qubits, and then measured in the Bell basis to extract information about the state. Z basis measurements are corrected with an $[[8,3,2]]$ decoder (when full error detection is not applied). **b**, Measured entanglement entropy as a function of subsystem size, showing expected Page-curve behavior²⁵⁶ for the highly scrambled state, improving with increased error detection. **c**, Measured and simulated magic (associated with non-Clifford operations) as a function of number of CCZ gates applied, performed on two copies of scrambled 6-logical-qubit systems. **d**, Pauli string measurement and zero-noise extrapolation using logical qubits. Plot shows the absolute values of all 4^{12} Pauli string expectation values, which only have five discrete values for our digital circuit; Pauli strings with the same theory value are grouped. By analyzing with sliding-scale error detection we improve toward the theoretical expectation values (squares) while also improving toward a purity of 1. By extrapolating to perfect purity, we extrapolate the expectation values and better-approximate the ideal values (shaded regions are statistical fit uncertainty).

tum state, which is a powerful tool that can efficiently extract many properties of an unknown state^{103,208,190}. With this two-copy technique, we plot the measured entanglement entropy in the scrambled system. We observe a characteristic Page curve³⁸⁴ associated with a maximally entangled, highly scrambled, but globally pure state. These measurements also reveal a final state purity of $0.74(3)$, compared to the measured XEB of $0.616(7)$, consistent with XEB being a good proxy for the final state fidelity. Despite postselection overhead, we find error detection significantly improves signal-to-noise here, as near-zero entropies are exponentially faster to measure.

Two-copy measurements can also be used to simultaneously extract information about all 4^N Pauli strings²⁰⁸. Using this property and an analysis technique known as Bell difference sampling¹⁹⁶ we experimentally evaluate and directly verify the amount of additive Bell magic¹⁹⁶ in our circuits as a function of number of applied logical CCZs. This measurement of magic, associated with non-Clifford operations, quantifies the number of T gates (assuming decomposition into T) required to realize the quantum state by observing the probability that sampled Pauli strings commute with each other. Moreover, combining encoded qubits and two-copy measurement allows for additional error mitigation techniques. As an example, we measure absolute expectation values of all 4^{12} logical Pauli strings with sliding-scale error detection. Since in the two-copy measurements, for each error detection threshold we also measure the overall system purity, we can extrapolate our expectation values to the case of unit-purity (zero-noise)²⁴¹. This procedure evaluates the averaged Pauli expectation values to $\approx 10\%$ relative precision of the ideal theoretical values spanning several orders of magnitude.

4.7 OUTLOOK

These experiments demonstrate key ingredients of scalable error correction and quantum information processing with logical qubits. In addition to implementing the key elements of logical

processing, our approach demonstrates practical utility of encoding methods for improving sampling and quantum simulations of complex scrambling circuits. Future work can explore if these methods can be generalized e.g. to more robust, higher-distance codes and if such highly entangled, non-Clifford states could be utilized in practical algorithms. We note the demonstrated logical circuits are approaching the edge of exact simulation methods, and can readily be used for exploring error-corrected quantum advantage. These examples demonstrate that the use of novel encoding schemes, co-designed with efficient implementations, can allow one to implement particular logical algorithms at reduced cost.

Our observations open the door for exploration of large-scale logical qubit devices. A key future milestone would be to perform repetitive error correction¹⁷⁰ during a logical quantum algorithm to greatly extend its accessible depth. This repetitive correction can be directly realized using the tools demonstrated here by repeating the stabilizer measurement in combination with mid-circuit read-out. The use of the zoned architecture and logical-level control should allow our techniques to be readily scaled to over 10000 physical qubits by increasing laser power and optimizing control methods, while QEC efficiency can be improved by reducing two-qubit gate errors to 0.1%¹²⁹. Deep computation will further require continuous reloading of atoms from a reservoir source^{402,316}. Continued scaling will benefit from improving encoding efficiency, e.g. by using quantum low-density-parity-check (qLDPC) codes^{55,452}, utilizing erasure conversion^{448,381,282} or noise bias⁹⁸, and optimizing the choice of (possibly multiple) atomic species^{402,277}, as well as advanced optical controls¹⁷⁷. Further advances could be enabled by connecting processors together in a modular fashion using photonic links or transport^{120,113}, or more power-efficient trapping schemes such as optical lattices⁴¹⁷. Although we do not expect clock speed to limit medium-scale logical systems, approaches to speed up processing in hardware⁴⁵⁴ or with nonlocal connectivity²⁷⁸ should also be explored. We expect that such experiments with early-generation logical devices will enable experimental and theoretical advances that dramatically reduce anticipated costs of large-scale error-corrected

systems, accelerating development of practical applications of quantum computers.

5

Ingredients for universal fault-tolerant quantum computing

Additional details for this chapter can be found in Appendix C and Ref.⁴².

Quantum error correction (QEC)^{399,410} is believed to be essential for the realization of large-scale quantum computers^{33,338}. However, due to the complexity of operating on the encoded ‘logical’ qubits^{6,117}, understanding the physical principles for building fault-tolerant quantum devices and

combining them into efficient architectures is an outstanding scientific challenge. Here we utilize reconfigurable arrays of up to 448 neutral atoms to implement all key elements of a universal, fault-tolerant quantum processing architecture and experimentally explore their underlying working mechanisms. We first employ surface codes to study how repeated QEC suppresses errors^{117,4}, demonstrating 2.14(13)x below-threshold performance in a four-round characterization circuit by leveraging atom loss detection and machine learning decoding^{448,23}. We then investigate logical entanglement using transversal gates and lattice surgery^{205,41,69}, and extend it to universal logic through transversal teleportation with 3D $[[15,1,3]]$ codes^{344,48}, enabling arbitrary-angle synthesis with logarithmic overhead^{6,108}. Finally, we develop mid-circuit qubit reuse⁴⁴⁶, increasing experimental cycle rates by two orders of magnitude and enabling deep-circuit protocols with dozens of logical qubits and hundreds of logical teleportations^{174,247,345,366} with $[[7,1,3]]$ and high-rate $[[16,6,4]]$ codes while maintaining constant internal entropy. Our experiments reveal key principles for efficient architecture design, involving the interplay between quantum logic & entropy removal, judiciously using physical entanglement in logic gates & magic state generation, and leveraging teleportations for universality & physical qubit reset. These results establish foundations for scalable, universal error-corrected processing and its practical implementation with neutral atom systems.

5.1 INTRODUCTION

The central challenge of quantum computation is its inherent sensitivity to errors^{338,259}. Whereas classical computers are composed of intrinsically robust digital bits stabilized by dissipation²⁵⁸, quantum states are intrinsically analog objects that evolve in a continuous state space, with coherent unitary evolution that does not allow dissipation²⁵⁹. Remarkably, the discovery of quantum error correction (QEC)^{398,410,244} provides a method to realize robust quantum computation. In this approach, entanglement between many physical qubits is used to encode error-corrected ‘logi-

QEC: ability to correct errors, at the cost of *complexity* and *constraints*

Qubit overhead, restricted logic operations, long run times, ...




Figure 5.1: The challenge of designing a QEC architecture. Quantum error correction provides a path to reach very low error rates, but comes at the cost of complexity constraints, associated with qubit overhead, restricted logical operations, and long run times. This motivates consideration of the hardware, logical encoding, and algorithm design to find the most efficient ways of realizing error-corrected quantum processing. In this chapter, we study the core underlying mechanisms of error correction in order to guide this design.

cal' qubits that can have exponentially low, digitized errors while performing arbitrary, analog-like computations^{399,6}.

While recent experiments provide clear examples where QEC works in practice^{41,4,347}, large-scale fault-tolerant quantum computation (FTQC) remains a formidable challenge. In essence, it requires one to dissipatively remove errors from the physical system, while simultaneously performing arbitrary coherent manipulation of the encoded logical qubits. Such operation requires a broad range of components^{348,370}, coupled with many important hardware considerations and a wide range of QEC techniques^{404,55,69,159,339}. Implementing and combining these techniques in practical systems, and understanding the scientific principles of resulting fault-tolerant architectures, is a highly complex task at the frontier of quantum science.

Here we experimentally implement the core building blocks for scalable, universal quantum computation and identify and explore key mechanisms for efficient architectures suitable for deep-circuit FTQC. Utilizing a quantum processor based on reconfigurable neutral atom arrays and lever-




Figure 5.2: Architectures and mechanisms for fault-tolerant quantum computation. **a**, We study the key building blocks of fault-tolerant processing. We utilize an architecture based on reconfigurable atom arrays trapped in optical tweezers, where the logical processor is segmented into storage, entangling, readout, and reservoir zones. Underlying physical mechanisms are identified and characterized. **b**, Spin-to-position conversion for non-destructive, loss-resolved qubit readout is accomplished with a state-selective 1D lattice that converts the atom spin state into position. Plot shows measured Rabi oscillation. This non-destructive readout has 0.46(4)% bit-flip error and 0.24(2)% loss. **c,d**, Stabilizer measurement on a $d = 5$ surface code is interspersed with global coherent errors injected on the data qubits. Each CZ layer corresponds to one time-step. Repeated correction reduces error build-up through both the Zeno effect and error tracking (right plot is at a fixed $\theta/2\pi = 0.016$). For visual clarity, an acceptance fraction of 50% is used in this plot.

aging key hardware upgrades, we study below-threshold error correction, fault-tolerant quantum operations, arbitrary unitary synthesis, and physical error removal during deep-circuit execution, altogether achieving all ingredients for scalable (i.e., logarithmic overhead) realization of arbitrary quantum algorithms⁶. Our results further advance the quantum computing frontier along several directions: (i) advanced decoding methods leveraging erasure information^{448,23} and machine learning³⁰ are used to optimize below-threshold performance; (ii) syndrome measurements and physical entanglement are reduced to being used only where they are needed during universal logic; (iii) teleportation^{174,247,345,366} is used to achieve efficient universality with transversal operations, and enable physical qubit reset without additional overhead in the algorithm. These result in significant improvements to architecture design and orders-of-magnitude reduction to space and time overheads.

5.2 NEUTRAL ATOM LOGICAL PROCESSOR

To implement these core building blocks, our experiments utilize a logical processor⁴¹ with up to 448 atoms. Qubits are stored in the hyperfine clock states of ⁸⁷Rb atoms trapped in optical tweezers generated by a spatial light modulator (SLM) in storage, entangling, readout, and reservoir zones. Quantum circuits are programmed by shuttling qubits in the middle of the computation with a 2D acousto-optic deflector^{37,377,43}, high-fidelity entangling operations are realized via fast excitation to Rydberg states¹²⁹, and fully programmable single-qubit operations are realized via locally focused Raman beams⁴¹. Crucially, this system provides control parallelism over logical qubit blocks⁴¹: for logic gates, repeated error correction, and even universal computation, all the physical qubits within the logical block realize identical operations with parallel instructions delivered by optical controls.

A key experimental upgrade involves non-destructive, spin-resolved qubit readout using a one-dimensional optical lattice in the readout zone. Whereas conventional readout is realized via spin-




Figure 5.3: Below-threshold repeated quantum error correction leveraging loss detection. **a**, Results of repeated rounds of $d = 5$ surface code using loss detection, showing a snapshot of the data block and multiple ancilla blocks. **b, c**, Products of stabilizer measurement results between rounds are used to detect qubit errors, which we refer to as ‘detectors’. ‘Bare’ counts loss as state $|0\rangle$, ‘detect loss’ does not make this erroneous assignment, ‘supercheck’ multiplies detectors around lost atoms, and ‘post.’ postselects on all atoms of a detector being present. **b**, Detector error probability as a function of data qubit loss in each shot, analyzed by partitioning the total dataset. **c**, Average over all data. **d**, Logical error per round for a surface code after 4 QEC rounds in both bases, decoded using most-likely error methods (‘bare MLE’), machine learning (‘bare ML’), MLE with loss information, and ML with loss information. ML with loss renders the error per round for $d = 5$ as $2.14(13)\times$ lower than $d = 3$. No postselection is used. Small points are the four $d=3$ quadrants. Results are averaged between $|+_L\rangle$ and $|0_L\rangle$ initialization bases. **e**, Logical error per round as a function of the mean qubit loss, plotted as a cumulative density function. **f**, Relative physical error contribution to overall error budget. **g**, Distribution of detector errors per shot, suggesting the absence of large-scale correlated errors.

to-loss conversion followed by camera readout⁴¹, the state-selective lattice enables splitting the two qubit states into two separate tweezers⁴⁴⁶, thereby realizing spin-to-position conversion that enables both loss detection and atom retention. Combined with techniques for mid-circuit re-initialization, this enables qubit re-use for extended computation as well as a two orders of magnitude increase to the experimental cycle rate, as described below.

Another key addition involves the use of repeated error correction for removing errors produced during quantum algorithms. To probe this function of QEC, we use a distance-5 surface code logical qubit^{117,4} with up to five rounds of stabilizer measurement. Here, data qubits are held in static potentials in the entangling zone, and an ancilla block is moved over to perform parallel entangling operations required for measurements of stabilizers. After performing the stabilizer checks, the ancilla block is moved to the storage zone, and a new block is brought in to repeat the cycle. Upon injecting global coherent errors at regular intervals, we find repeated correction significantly sup-

presses logical errors, as coherent errors are projected into incoherent ones (with quadratic suppression) due to the quantum Zeno effect²²⁰ and tracked. Notably, although the coherent errors are injected as global, correlated single-qubit rotations, stabilizer measurements convert these into uncorrelated, incoherent physical errors⁵⁶, and prevent coherent errors on the logical qubit level. This highlights the general function of repeated QEC: the conversion of general errors to correctable bit-flip and phase-flip errors removes entropy by reducing the range of possible errors⁶, and detecting the incoherent errors and using them in decoding reduces entropy by keeping track of the state of the system¹¹⁷.

5.3 BELOW-THRESHOLD PERFORMANCE

We first utilize these tools to explore below-threshold performance of a surface code qubit under multiple rounds of QEC. Theoretically, logical errors can be exponentially suppressed as $(p/p_{\text{th}})^{(d+1)/2}$ by increasing code distance d if the physical error rate p is below a characteristic threshold p_{th} ^{6,117,4}. We observe that qubit loss events result in flickering stabilizer patterns⁴⁰⁸ which create time-correlations between detected error events and leads to a sharp rise in detected errors.

To correct for the effect of atom loss, we construct so-called superchecks (see Ref.⁴⁰⁸) by multiplying stabilizers around the lost atom to recover error information, which results in a suppressed detected error probability. These observations indicate that using the loss information in decoding can greatly improve QEC performance, as predicted theoretically^{448,23}. To take advantage of these features, we use most-likely error (MLE)^{70,23} and machine learning decoders³⁰ trained on simulated and experimental data, incorporating loss information into both.

We now characterize repeated QEC performance as a function of the code distance d . We find that $d=5$ has a $2.14(13)\times$ lower error per round than $d=3$, indicating below-threshold behavior for this four-round characterization circuit. In particular, we find that the loss information and ma-




Figure 5.4: Exploring the interplay of logic gates and entropy removal. **a**, Atom images illustrating two-qubit logic gates and stabilizer measurement. Lattice surgery is realized using ancillas to measure the logical product $Z_L^1 Z_L^2$, and transversal gates are realized via atomic motion. **b**, Quantum circuits for realizing transversal gates and lattice surgery operations. CZ gates are realized as transversal CNOTs and Hadamards. **c**, Dependence of error of the logic operation on the ancilla measurement error. Lattice surgery errors rapidly worsen with increasing ancilla measurement errors (injected in post-processing). **d**, N repeated logic operations are interspersed with rounds of QEC stabilizer measurement. Transversal CNOT has a lower error than lattice surgery (left), and has an optimum of roughly 3 CNOTs per round, as seen most clearly when modest postselection, characterized by acceptance fractions (AF) are used (right). Error detection is used for lattice surgery in **d** to compensate for having $< d$ rounds.

chine learning decoding together improve the QEC performance by a factor of $1.73(13)\times$ compared to conventional methods. Note that while our observations show below-threshold behavior for repeated QEC rounds, theoretically the threshold can get worse by $\approx 1.15\times$ in the limit of many repeated rounds and by $\approx 1.1\times$ when incorporating ≈ 1 transversal gate per QEC round. The error per round is $0.62(3)\%$ for $d = 5$, and reduces toward 0.1% per round on shots with no qubit loss, consistent with a p^3 scaling and roughly half of our errors being loss.

These observations are consistent with numerical simulations using a simple empirical error model based on separately characterized error rates. We further note that our observed distribution of errors-per-shot is closely consistent with these simple simulations with uncorrelated Pauli-type and loss-type errors, suggesting the absence of large-scale correlated errors in the system. We observe that time-correlations are almost fully diminished when postselecting on no qubit loss, indicating that almost all leakage ($> 80\%$) corresponds to atom loss, while leakage to other hyperfine states appears to be suppressed. This observation suggests that various scattering channels⁹⁸ naturally loss-convert, likely due to anti-trapping of metastable states.

5.4 LOSS DETECTION FOR IMPROVED QEC

Leakage types with neutral atoms. Leakage errors, which take the qubit out of the two-level computational subspace, are important to account for in error correction. The three dominant leakage errors with neutral Rubidium (or other alkali) atoms are:

- *Loss events.* Loss events are when the atom is physically lost from the optical trap. Due to the blockade nature of the gate, doing a gate with a lost atom simply turns off the gate while still applying gate error (it is identical to the atom being in state $|0\rangle$ which is also dark to the Rydberg laser).
- *Leakage to other hyperfine states in the ground-state manifold.* In the limit of a large magnetic field, these states are off-resonant and behave the same as a lost atom (turning off CZ gates). However, they are not detected through loss detection. Moreover, in the practical operating conditions of 8.6 G, the level spacings of 6 MHz (compared to Rabi frequency of 4.6 MHz) mean that adjacent hyperfine states can still off-resonantly couple to the Rydberg state and could lead to repeated errors.
- *Leakage to Rydberg states.* Population left in the Rydberg manifold can affect subsequent gates, and can lead to large error correlations. For example, many-body Rydberg evolution in dense systems observes so-called avalanche errors where a macroscopic fraction of the system has an error¹³⁶. Crucially, in our approach with a low atomic density and several hundred microseconds between gates, the Rydberg atoms (theoretically) either decay to the ground state or are expelled from the tweezer. In this way, such Rydberg leakage either convert into an error within the computational subspace, a leakage into adjacent hyperfine states, or a loss event.




Figure 5.5: Characterizing effects of loss and leakage in repeated QEC. **a**, Superchecks. A lost data qubit results in anti-commuting stabilizers and a flickering error pattern. Producing stabilizers surrounding the lost qubit recovers commuting superchecks. **b**, Detection correlation p_{ij} matrix for five rounds of QEC on a $d = 5$ surface code. Data qubit errors appear as space-like correlations and ancilla measurement errors as time-like correlations between adjacent layers. Leakage results in additional persistent correlations between QEC rounds. By selecting shots with the fewest lost data qubits, these correlations are suppressed, indicating that leakage is dominated by loss which can be detected. Similarly, selecting shots with the most loss enhances the correlations.

We observe that with several hundred microseconds between gates, effects of Rydberg leakage are not apparent, and during our repeated QEC data we observe that our leakage is at least 80% loss.

Effect of loss during repeated QEC. Although losses simply turn off subsequent gates, these lead to distinct signatures that are important to account for in the QEC design. Whereas ancilla loss will be detected in the projective measurement, losing a data qubit corresponds to unknown loss of a degree of freedom from the system^{408,29}. Without adjusting the stabilizer measurement pattern to account for such a loss, the ancilla qubits now are measuring operators which anti-commute with each other, and thereby leads to a ‘flickering’ pattern around the lost data atom. This flickering pattern is akin to the expected behavior in a subsystem code²⁵², and means that the flickering can continue for arbitrarily long times.

Erasure information and superchecks. It is useful to detect atom loss for two reasons. First, knowing about the lost atom greatly enhances the decoding performance. While bit-flip and phase-flip errors can be inferred by stabilizers, direct detection of qubit errors - or so-called erasures - means that one already has direct information about where the errors are. Such erasure information can thereby greatly improve decoding performance^{448,381,366,82,339,23}. For example, while only $(d-1)/2$ Pauli-type

errors can be corrected, up to $(d-1)$ erasure-type errors can be corrected. We do not detect erasures as soon as they occur, instead we detect them at the final qubit measurement, constituting delayed-erasure information.

Second, although lost atoms lead to anti-commuting stabilizer measurements and a flickering error pattern, these can be accounted for with the use of so-called superchecks^{408,29}. While individual stabilizer checks around a lost atom are anti-commuting, taking products of multiple checks creates superchecks which again commute with each other. We find that such superchecks are able to remove the sharp rise in detected error that occurs with increasing data loss.

5.5 STABILIZER MEASUREMENT DURING LOGIC OPERATIONS

We next extend repeated QEC to logical operations, by studying two different approaches for realizing quantum logic, where the stabilizer measurements play different roles. In the transversal method, a logical gate is realized by physically transporting and interlacing the data qubits of two surface code blocks, and then applying pairwise entangling gates^{399,117,41}. In the planar setting, logical entanglement is realized via lattice surgery²⁰⁵, where a joint logical measurement is performed by ancilla qubits added between the two codes (Figs. 3a,b). To probe these logical operations, in the transversal approach, we apply three rounds of N repeated CNOTs interleaved with two rounds of stabilizer checks. In the lattice surgery approach, we carry out a $Z_L^1 Z_L^2$ joint logical measurement using two rounds of stabilizer measurements²⁰⁵. We then measure the $X_L^1 X_L^2$ and $Z_L^1 Z_L^2$ parities of the resulting logical Bell state for both methods.

To test the role of stabilizer measurements in both cases, we first introduce additional ancilla measurement errors in postprocessing. We observe a lower logical error for transversal gates, and that the lattice surgery is significantly more sensitive to injected measurement errors. To explore the interplay of logic gates and entropy removal, we next study the transversal gate performance as a

function of the number of repeated CNOT gates per QEC round. We find optimal performance for several CNOTs per QEC round, and using postselection on decoding confidence, observe that approximately three CNOTs per round is optimal at low error rates.

These results highlight multiple key aspects of FTQC. First, the response to ancilla measurement errors highlights the key distinction between the two approaches: in the transversal gate setting, the logic is realized directly between the data qubits which store the underlying logical states, and the role of stabilizer measurements is to remove entropy, whereas in lattice surgery, the ancilla measurements directly perform the logic operation and have to be correct. The need for correct stabilizer measurements is the origin of the conventional fault-tolerance assumption of d rounds per QEC cycle^{117,205}. Conversely, with transversal operations, the observed optimum at 3 CNOTs per round corresponds to stabilizer measurements balancing the local entropy generated by the logic gate. This is consistent with theoretical predictions for universal computation with correlated decoding techniques^{70,466,69}. Second, we clearly see that the error per logical gate depends on the number of gates applied per QEC round. This highlights a key distinction between operations on physical and logical qubits: while physical gate performance can be well-characterized by a single fidelity value F , logical gate performance depends both on the decoding success probability - which reflects the physical entropy of the system, captured by the detector error probability p_{det} - and on how much that entropy increases per logical gate, quantified by Δp_{det} . We find that a simple model for logical gate fidelity F_L , where $1 - F_L \propto [(p_{\text{det}} + N\Delta p_{\text{det}})^{(d+1)/2}] / N$, with the measured Δp_{det} consistent with our physical two-qubit gate error, accurately describes the experimental data.

5.6 UNIVERSALITY AND SYNTHESIZING ARBITRARY UNITARIES

We now study how to realize arbitrary error-corrected unitary operations. The core ingredient is the Solovay-Kitaev theorem, which states that arbitrary single-qubit rotations can be approximated to




Figure 5.6: Synthesizing arbitrary-angle rotations with a universal fault-tolerant gate set. **a**, Subjecting codes of various dimensionality to global rotations. Plateaus are seen at robust angles where the stabilizers also revive. 3D Reed-Muller codes, only with all positive stabilizer signs, have an additional plateau at 45-degrees corresponding to non-Clifford T gates. **b**, Programmable angles $R(\theta, \varphi)$ are realized by an alternating sequence of H and T gates. The circuit shows an implementation of such a sequence, with T implemented via state preparation and H implemented via quantum teleportation. **c**, Polar angle plots show generated angles using entangled Reed-Muller codes, measured by tomography, for different maximum numbers of T gates. Experimental results are consistent with theory within statistical error. **d**, Experimental results show that the minimum separation between generated angles decreases exponentially with length of sequence. Inset shows a rescaling where ε is the angular separation to the target angle. Bloch sphere shows all measured angles. For visual clarity, variable degrees of acceptance fraction are used.




Figure 5.7: Universality with 3D codes. **a**, Codes of various dimensions subject to global rotation, here plotted without any normalization applied to the logical operator or stabilizer expectation values. **b**, Two-copy measurement of a $[[15,1,3]]$ code after a global rotation¹⁹⁶. The additive Bell magic has a plateau at one unit of magic under a global T . The same magic is obtained by analyzing either the underlying physical 15-qubit system or the single logical qubit with error detection. **c**, The encoded logical state is connected via a unitary Clifford decoding circuit to a product state of the encoded state and Pauli inputs on the physical qubits. The Clifford circuit conserves magic and therefore the total physical state and encoded logical state must have the same total magic. This implies that, while 15 physical T 's are applied to the system, only 1 physical T is produced, which can only happen with entanglement. **d**, Atom image of a register of 2D and 3D color codes. A transversal CNOT can be performed between the face of the $[[15,1,3]]$ code (control) and the $[[7,1,3]]$ (target). **e**, Code switching. $|T_L\rangle$ is prepared transversally on the 3D $[[15,1,3]]$ code and then measurement and in-software feedforward teleports the logical T onto a 2D $[[7,1,3]]$ code (here also with a H); this illustrates the equivalence between code switching protocols and logical teleportation. Error detection is used in both plots. **f**, Error-corrected test of Bell's inequality³¹. We measure $S = E(X, T) + E(X, T^\dagger) + E(Y, T) - E(Y, T^\dagger)$, where $E(A, B)$ is the expectation value of the Steane and Reed-Muller codes in the A and B bases, respectively, and obtain $S = 1.99(3) \times \sqrt{2}$ with error detection.

exponential precision using only digital gates such as Hadamard H and $T = e^{-i\frac{\pi}{8}Z}$ (a 45-degree rotation around the Z-axis)^{244,6,108}. We note that while the Eastin-Knill theorem prohibits realizing such a universal gate set with unitary transversal operations¹²², it can be circumvented by the introduction of logical measurement which breaks the unitarity constraint³⁴⁴.

In our approach, non-Clifford T gates and universal rotations are realized with efficient transversal circuits built out of teleportation with 3D codes. Figure 4a shows experimental measurements where we use quantum circuits to create 2D color codes (Steane codes)⁴¹⁰, and 3D color codes (Reed-Muller codes), and subject them to global phase rotations ϕ around the Z-axis³⁴⁴. For comparison, we also show the results of similar measurements for unentangled physical qubits analyzed as 3D codes, and 3D codes with incorrect values of stabilizers. The various configurations show plateaus in the logical expectation value, and revivals in the stabilizers, for multiples of 90-degree angles (multiples of 180-degree for unentangled qubits). Crucially, 3D codes, prepared in the proper entangled states, also exhibit robustness at multiples of 45-degrees, corresponding to their transversal T gate⁴⁸.

To implement unitary synthesis, transversal teleportation is utilized. Specifically, we create multiple Reed-Muller codes in the $|T_L\rangle$ state, and by applying logical CZ gates followed by X-basis measurements (with feedforwards applied in-software here), the logical information is teleported with an H gate³⁴⁴. The resulting logical states span a range of points on the Bloch sphere and match the expected angles with high precision (consistent with statistical uncertainty). We observe that the angular spacing between accessible states shrinks exponentially with the number of T gates. This enables precise synthesis of rotation angles with a logarithmic number of steps, as expected by the Solovay-Kitaev theorem^{244,1,6}.

These observations demonstrate that teleportation can be used as a powerful tool for universal processing, allowing precise analog rotations to be built from digital gates. Interestingly, although the circuit realized here is fully transversal, the measurement, decoding, and logical feedforward

ensures that the logical information propagates unitarily while physical-level dissipation serves to correct errors. We explore code-switching and transversal gates between 2D and 3D codes, again finding the underlying source of universality is the measurement of the 3D code. In addition, we note that the role of stabilizers differs fundamentally when generating logical magic. While for Clifford circuits, stabilizer measurements are used to remove entropy, for transversal T gates, correct stabilizer signs (e.g., +1 eigenvalues) are essential for realizing non-Clifford operations. This observation reveals that physical entanglement is directly required for logical magic. This can be understood by the fact that, while logical Pauli states such as $|+_L\rangle$ are eigenstates of operators $X_L = X_1 X_2 X_3 \dots$, which is a tensor product of physical operators, states such as $|T_L\rangle$ are eigenstates of $\frac{1}{\sqrt{2}}(X_1 X_2 X_3 \dots + Y_1 Y_2 Y_3 \dots)$, involving a superposition spanning the code that is necessarily entangled. Such an observation has fundamental implications: for example, an error-corrected Bell inequality test (involving T-basis measurements) requires a high degree of in-block entanglement. We perform such an error-corrected Bell inequality test and measure a CHSH inequality of $1.99(3) \times \sqrt{2}$, saturating the quantum bound³¹.

5.7 DEEP CIRCUITS AT CONSTANT ENTROPY

We now explore the ability to perform deep-circuit quantum computation on the logical level. A critical requirement is that the processor is kept at a constant entropy⁶, necessitating that physical errors do not accumulate. This is challenging because computation inevitably introduces errors in the physical qubit state while, in addition, increasing entropy in the other degrees of freedom, such as the atomic motional state. To ensure that all physical errors are removed and that computation is kept at constant entropy, we again leverage transversal teleportation^{345,247,344,366}. In this approach (detailed below) the logical information propagates throughout the circuit while physical errors are left behind. Measuring this block then enables qubit reset, re-cooling, and re-initialization of the




Figure 5.8: Deep logical circuits at constant entropy. **a**, Schematic of a 2D cluster state created from Steane codes in space and time. The same atomic qubits are re-used every other time layer for up to 27 layers. **b**, Results of repeated logical state preparation of Steane codes. Stabilizer error is constant as a function of cycle. **c,d**, Logical and physical correlations in 1D (**c**) and 2D (**d**) cluster states. Logical correlations persist while stabilizer error correlations rapidly decay. Here and below, various degree of acceptance fraction are applied. **e**, Circuit diagram illustrating that logical-level evolution is unitary and allows the logical operator to propagate throughout the algorithm whereas physical-level evolution is dissipative and does not let physical errors propagate. **f**, High-rate $[[16,6,4]]$ codes are entangled in space and time direction. **g**, Logical evolution and physical error removal with $[[16,6,4]]$ codes (1D cluster state in time direction). Permutation CNOTs within the block (black curve) extend the correlation length in comparison to non-permuted qubits. **h**, Entanglement structure in space and time, with up to $96 d = 4$ logical qubits active simultaneously. **i**, Logical 2D cluster expectation values (averaged across space and the first 9 teleportation layers) as a function of the number of co-propagating logical operators which agree (i.e., cluster state stabilizers have the same outcome for each logical qubit within the block).

physical atoms. After that the block is prepared again in a low entropy state and is then utilized for subsequent teleportation steps.

To realize constant entropy operation during deep quantum circuits, the atomic internal states, temperature, and atom filling need to be re-initialized during the computation. In order to achieve this, we combine our non-destructive internal state readout with non-destructive imaging that also re-cools the atom. While laser cooling is typically achieved using 3D beams in zero magnetic field, we implement a novel method that enables high-fidelity imaging and cooling operating with focused 1D beams in a finite B-field (required for atomic qubit control)^{434,92}. We further protect coherence of data qubits in the nearby storage zone by applying a 1529-nm shielding beam²⁰⁷. Moreover, the missing atoms in the array are refilled with atoms from the atomic reservoir. With all these methods combined, we measure the performance when subjecting the atoms to all the operations (except entangling gates) in a 27-layer circuit, explained below. For instance, by applying a perturbation on the 5th cycle, we find that the atomic filling and temperature quickly recover to a steady-state. We find that the 1D cooling methods nearly reproduce the conventional 3D performance, limited by tweezer-depth inhomogeneity (which can be straightforwardly improved).

We now use these tools to implement logical algorithms. Here, we entangle, measure, and teleport the logical blocks in alternating A and B groups. Within one layer, we bring in a fresh batch of physical qubits (group A) from the readout into the entangling zone in order to encode them into error-correcting codes, and then entangle the code blocks with each other (entangling in space direction). We then bring group B (already entangled) from storage, and perform a transversal entangling gate with the group A block (entangling in time direction). In the spirit of measurement-based quantum computing³⁴⁵, we then move group B into the readout zone, group A to the storage zone, and then measure group B. Group B is then re-initialized and the whole process is repeated in the next layer.

We first study repeated state preparation of 32 blocks of $[[7,1,3]]$ (Steane) codes (16 blocks per

alternating group) for up to 27 layers, and find that the stabilizer expectation value remains constant as a function of the cycle number, indicating a steady-state internal entropy. (We note that in these experiment the fidelity is limited by several sub-optimal choices in the circuit design structure). As an example algorithm, we entangle the $[[7,1,3]]$ codes in both the time and space directions to realize 1D and 2D cluster states and probe the resulting correlations³⁴⁵. Starting with a 1D cluster state in the time direction, we measure the correlations on both the logical-level and the physical-error-level by evaluating the correlator $\langle Z_i Z_j \rangle - \langle Z_i \rangle \langle Z_j \rangle$ between coordinates separated in time. On the logical-qubit level we observe the expected correlations between logical outcomes corresponding to successful algorithm evolution, with a decay with distance corresponding to an effective algorithmic error rate that improves with decreasing entropy (corresponding to lower acceptance fraction). Conversely, for the physical errors, we observe the correlations are rapidly suppressed. Similar behavior is found in the 2D cluster state data, where the logical cluster state stabilizers are finite across both space and time, but the correlations between stabilizer errors are rapidly suppressed. These properties persist until the reservoir begins to run out of atoms. Leveraging the regular transversal measurements, we use data qubit loss detection as well as a recurrent neural network architecture⁵⁰ for decoding these algorithms.

These observations can be understood by considering how the transversal teleportation ensures the removal of physical errors. Whereas the decoding and logical teleportation (with feedforward done in-software) maintains unitary evolution and propagates the logical information, the physical errors remain on the previous block (with errors propagated by entangling gates traveling at most one layer). This structure, where the logical information is natively teleported onto a fresh logical block in the algorithm, thereby ensures that the algorithm proceeds at constant entropy while also performing logic.

To test this teleportation method for more general encodings, we study high-rate $[[16,6,4]]$ codes³⁴⁷. Such high-rate codes⁵⁵ have a complex structure that would generally require intricate

circuit structuring to ensure leakage removal^{4,23}. In contrast, by directly using the teleportation procedure described above we find that in a temporal 1D cluster state the logical information propagates while the physical errors are suppressed with a rate similar to the simpler $[[7,1,3]]$ Steane codes. Such high-rate codes also enable new opportunities for realizing quantum algorithms. For example, the permutation CNOT operation can enable entangling logical qubits within the same block simply through re-indexing physical qubits, thereby extending the correlation length.

We also explore 2D entanglement of the $[[16,6,4]]$ blocks (with up to 16 blocks at a time). We measure the 2D logical cluster state stabilizers as a function of postselection on shots where the co-propagating logical operators agree (i.e., the cluster state stabilizers have the same outcome for each logical qubit within the block). We find that such a procedure further improves algorithm performance. This is because the logical operators - although independent degrees of freedom - are supported on the same physical qubits. Interestingly, while these easily attained in-block correlations are algorithmically useful (as illustrated using $[[8,3,2]]$ codes in Ref.⁴¹), generating such logical entanglement is only possible because the logical operators overlap, which is enabled by the underlying physical entanglement.

5.8 DISCUSSION AND OUTLOOK

We now turn to a discussion of our key observations. First, our experiments uncover an intricate interplay between quantum logic operations and entropy removal: understanding their relationship enables syndrome extraction to be applied only where necessary, and also clarifies entropy-related aspects of logical gate performance. Second, we find that physical entanglement should be deployed judiciously in fault-tolerant systems. While techniques such as transversal gates and high-rate encoding reduce the physical entanglement required for a given amount of logical operations (Figs. 3,6f-i), logical magic states demand enforcing precise entanglement structure. Third, we observe that logi-

cal teleportation can be a central mechanism for FTQC²⁴⁷. Such teleportation enables universality even with fully transversal operations and offers a native path to physical error removal in deep-circuit protocols, including those involving high-rate codes, allowing direct data qubit detection (including atom loss events) to be effectively used in the decoding algorithms (Figs. 2,6). Beyond these architectural insights, we also observe several fundamental physics aspects of QEC. While prior work has connected quantum contextuality (involving T -basis measurements) with computational hardness²⁰⁶, our results suggest additional links to the essential entanglement required for the QEC. These observations point toward a deeper understanding of how to protect algorithmic outputs, and facilitate the realization of complex, deep-circuit quantum algorithms.

While the current experiments demonstrate QEC performance at a factor of ≈ 2 below key thresholds, large-scale computation will greatly benefit from reducing physical errors. Based on our observations, we estimate that an additional three- to five-fold reduction in physical error can be achieved through direct improvements in single-qubit operations, improved calibration, and a four-fold increase in entangling laser power¹²⁹. While machine learning decoding is found to be both effective for entropy removal performance and simple algorithms, and fast ($\sim 1 \mu\text{s}$ per shot when batched on a GPU), more work is required to ensure scalability of this powerful method³⁰. Moreover, although the maximum circuit depth here is limited by a depleting atomic reservoir, in a complementary experiment conducted in a separate apparatus, coherent continuous operation on over 3000 atomic qubits⁹⁰ is demonstrated via continuous atom reloading^{185,317}, with techniques fully compatible with the methods presented here. Taken together, these techniques enable advanced experimental exploration of fault-tolerant universal algorithms. Combined with other significant progress with neutral atom systems^{227,282,288,425,311,464}, such developments demonstrate that these systems are uniquely positioned for experimental realizations of deep-circuit fault-tolerant quantum computing.

5.9 ERROR BUDGET AND PATH TO 10X BELOW THRESHOLD

To get to algorithmically relevant error rates of $\sim 10^{-10}$ ^{157,39}, a factor of 5-10x below threshold can achieve the required errors with several hundred qubits in a code block¹⁴¹. Our performance is captured by the error budget, which we now describe in further detail.

We first list our single-qubit errors and their possible improvements:

- Local single-qubit gates are approximately 99.9% fidelity, arising from a 0.05% scattering error and residual miscalibrations. Increasing the Raman detuning to 2.5 THz will further reduce scattering errors and miscalibrations from the Raman differential light shift, and improved calibration routines could thereby achieve 99.99% fidelity.
- Our coherence time in 852-nm traps is approximately 1-2 seconds depending on the dynamical decoupling sequence applied. Comparable systems have achieved coherence times of 12.6 seconds with further tweezer detunings²⁸⁸.
- We experience a total loss from movement of roughly 1% on the ancilla atoms, arising from transfers and moves between and within the zones. We previously observed performance in Ref.⁴¹ with transfer-limited loss that would correspond to 0.2% movement loss here, which we suspect in the present work arises from using an AOD RF power that is too high. Our repeated QEC sequence also experienced 0.6% background loss from vacuum that can be readily reduced to <0.01% using improved vacuum lifetime and a shorter sequence (e.g. ~ 4 ms cycle times).
- Our lattice readout currently operates with a loss rate of 0.3% and a 99.5% bit-flip fidelity. Although a new technique, similar methods in purely-lattice systems have achieved fidelities of 99.94%⁴⁴⁶.

In order to improve two-qubit gate performance, an example approach can be:

- Improve system stability, homogeneity and fast, automated calibration. Although we achieve CZ fidelities of 99.6%, drifts since the last calibration (several days in the context of the surface code benchmarking) often contribute 0.05 – 0.1% error during final data taking.
- Use the smooth-amplitude gate, higher magnetic fields, or the $6P_{1/2}$ intermediate state to suppress coupling to the adjacent $m_j = +1/2$ state, reducing the error from $\approx 0.06 - 0.15\%$ to near-zero.
- Increase both 420-nm and 1013-nm Rydberg laser powers by a factor of four. This can simultaneously allow the following improvements (numbers are from simulation; see Ref. ¹²⁹):
 - Increase Rydberg detuning from 4.8 GHz in the present work to 9.6 GHz, reducing scattering error from 0.094% to 0.052%.
 - Decrease gate time from 270 ns to 135 ns, reducing Rydberg T_1 error from 0.113% to 0.057%, and reducing dephasing error from 0.134% to 0.034%.

These can reduce the two-qubit gate error from roughly 0.5% to 0.15% through simple system improvements. The AOM pulse profile should be compensated for realizing these gate times, and the 1013-nm beam's fractional inhomogeneity will need to be improved by the corresponding increase in power.

Altogether, by reducing single-qubit gate errors by a factor of 5x and improving two-qubit errors from 0.5% to 0.15% through the improvements listed, operation at roughly 8x below threshold would be achieved. The two-order-of-magnitude increase in cycle rate demonstrated here will be instrumental to enabling these improvements. These estimates highlight that straightforward improvements can deliver the performance required for large-scale computation. We also emphasize that such performance will need to be tested and optimized in deep-circuit settings.

6

Analog quantum simulation of quantum coarsening and collective dynamics

Additional details for this section can be found in Appendix D and Ref.²⁸⁹.




Figure 6.1: Application of local detunings. **a**, In addition to the SLM used for atom trapping, a second SLM generates a set of superimposed tweezers that can apply arbitrary detuning profiles to the atoms. The pair of crossed AODs for local hyperfine control⁴¹ are not used in this work. **b**, A spin-echo sequence is used to measure and calibrate the differential AC Stark shift of the local detuning beams on the ground-state clock transition. Darker curves indicate a higher target detuning, weighted between 1 and 4, and the grey curves for unaddressed sites demonstrate negligible crosstalk between sites. The pattern here corresponds to a rotated version of the domains in the main text. **c**, Measured differential AC Stark shift of the $5S_{1/2} \rightarrow 5P_{3/2}$ transition. Unless stated otherwise, the laser amplitude used is 1, corresponding to $\approx 160 \mu\text{W}$ per site and imparting an approximate light shift on the differential $|g\rangle \rightarrow |r\rangle$ transition of $\delta_0 = 12(2)\text{MHz}$. The amplitude is the square root of power relative to $160 \mu\text{W}$. When sites are weighted, the total power imparted on the array remains constant but is proportionally redistributed. **d**, The local detunings are used to prepare deterministic checkerboard orderings, both for a single global order (left) or for domains with different orders (right). A site polarization of unity, measured at the start of the hold time for varying Δ/Ω , corresponds to $|g\rangle$ ($|r\rangle$) for addressed (unaddressed) sites, while -1 corresponds to the flipped state. Note that at finite Δ/Ω , the polarization of $|r\rangle$ sites $\neq 1$; at large Δ/Ω , mean-field shifts reduce the spin polarization of both $|g\rangle$ and $|r\rangle$ along the domain wall. **e**, Dependence of the site polarization on Δ/Ω in the bulk of the domain. As the local light shift becomes weaker relative to the global detuning, the probability of preparing $|g\rangle$ sites decreases slightly.

6.1 INTRODUCTION

Quantum phase transitions (QPTs) are transformations between states of matter that are driven by quantum fluctuations³⁶⁰. Analogously to thermal fluctuations in classical phase transitions, quantum fluctuations play a dominant role in the emergence of order in quantum systems. While classical dynamics near thermal critical points have been studied extensively over the past several decades, only recently have quantum dynamics across QPTs become experimentally accessible, due to the advent of quantum simulators^{20,236,123,380} and ultrafast spectroscopic methods in solid-state systems^{356,327,223,397}. Their universal properties have been studied in systems of varied dimensionality using the quantum Kibble-Zurek mechanism (KZM)^{340,236,123}. The KZM stipulates that a quantum system’s dynamics and correlations “freeze” in the vicinity of a QPT when the system can no longer respond adiabatically to dynamical changes. However, in many instances, other mechanisms of correlation growth beyond KZM can dominate ordering^{40,355,379,462}. In particular, when an unordered system passes through a continuous phase transition into a symmetry-broken phase, a progressive growth of long-range order, known as coarsening, is expected. These ordering dynamics are predicted to exhibit universality, manifested as self-similarity in the growth of correlations^{273,202,57}. While such phenomena are well understood in the context of thermal phase transitions^{169,151}, quantum effects in coarsening dynamics have only recently emerged as a subject of intense theoretical^{80,81,290,149,373} and experimental^{362,272,212,13} investigation.

We use a programmable quantum simulator based on Rydberg atom arrays to investigate the collective out-of-equilibrium dynamics associated with the growth of order following an Ising QPT. We observe key features of beyond-mean-field quantum coarsening processes arising from quantum fluctuations: the curvature-driven dynamics of domain walls and their acceleration when approaching the critical point. We further explore self-similarity and universality in the ordering process. Additionally, we observe long-lived coherent oscillations of the correlation length and the order pa-




Figure 6.2: Ordering via coarsening. **a**, Atoms arranged in a square lattice are prepared in the ground state $|g\rangle$ and driven from the disordered phase, shown here as a fluorescence image (atoms in $|r\rangle$ are detected as loss and are indicated with red circles), into an antiferromagnetic ordered phase, across a quantum phase transition. The subsequent coarsening dynamics lead to the progressive development of long-range order (shown here using the real-space correlation function $G(x, y)$ and experimental snapshots). The snapshots show the local single-spin-flip-corrected staggered magnetization, with white (dark grey) shading indicating AF_1 (AF_2) orders. **b**, The detuning $\Delta(t)$ is swept linearly, with a fixed Rabi frequency Ω , to a final value Δ in the ordered phase, where it is held constant for the duration of the hold time. **c**, The correlation length ξ grows over the course of the hold time. The solid line is a fit to $(c_0 + c_1 t)^\alpha + c \cos(\omega t + \varphi)$, representing sinusoidal oscillations on top of a power-law growth. Errorbars are smaller than markers. **d**, The early-time growth of the squared correlation length is consistent with a linear increase with time, as expected for coarsening dynamics with a nonconserved order parameter. The dynamics are faster at lower Δ/Ω , closer to the critical point. **e**, The structure factor $S(k, t)$ of the staggered magnetization, for different hold times, at $\Delta/\Omega = 2.0$. Inset: the structure factors collapse onto a single curve when the time axis is rescaled by the correlation length $\xi(t)$ and the magnitude of the structure factor is rescaled by a time-dependent amplitude $\xi^2(t)b(t)$. For the data shown here, $\Omega/2\pi = 3.8$ MHz.

parameter on both sides of the QPT, which we identify as an amplitude (Higgs) mode. We study these fundamental features of QPTs^{397,356,128}, demonstrating consistency with theoretical predictions³⁶⁰, and extending these studies in a regime that is difficult to simulate classically.

Our experiments are performed using a two-dimensional programmable atom array, previously described in Ref.¹²³. The measurements are conducted on a 16×16 square lattice of ^{87}Rb atoms trapped in an array of optical tweezers generated by a spatial light modulator (SLM). Atoms are initialized in the electronic ground state $|g\rangle$ and are coupled to the high-lying electronic Rydberg state

$|r\rangle$ through a two-photon excitation with time-dependent Rabi frequency $\Omega(t)$ and global detuning $\Delta(t)$. As a key upgrade, we introduce a second SLM for generating locally controlled light shifts, allowing for programmable site-dependent detunings $\delta_i(t) = \alpha_i \partial(t)$. The atoms in the $|r\rangle$ state interact strongly via a van der Waals potential, giving rise to the following Hamiltonian governing the system's dynamics:

$$\frac{H}{\hbar} = \frac{\Omega(t)}{2} \sum_i X_i - \sum_i n_i (\Delta(t) + \delta_i(t)) + \sum_{i < j} V_{ij} n_i n_j. \quad (6.1)$$

Here, $n_i \equiv |r_i\rangle \langle r_i|$ denotes the Rydberg occupation at site i , $X_i \equiv |g_i\rangle \langle r_i| + |r_i\rangle \langle g_i|$ describes the laser-induced coupling between the states at that site, and $V_{ij} \equiv V_0/|\mathbf{r}_i - \mathbf{r}_j|^6$ is the van der Waals interaction. The Rydberg interactions prevent simultaneous excitation of two atoms if they lie within a blockade radius ($R_b \equiv (V_0/\Omega)^{1/6}$) of each other. We maintain a lattice spacing a such that $R_b/a \approx 1.1$, and only nearest-neighboring sites fall within the blockade radius. For large positive values of Δ/Ω , this configuration leads to a $\mathbb{Z}_2$ -symmetry-broken checkerboard phase with two antiferromagnetically ordered ground states³⁷¹, labeled $|\text{AF}_1\rangle$ and $|\text{AF}_2\rangle$. The disordered and ordered phases are separated by a QPT belonging to the (2+1)D Ising universality class, which occurs at $\Delta/\Omega \approx 1.1$ ^{371,231,123}. The order parameter diagnosing this transition is the staggered magnetization: $m_s = \sum_{x,y} \tilde{Z}_{x,y} = \sum_{x,y} (-1)^{x+y} Z_{x,y}$, where (x,y) denotes the two-dimensional coordinates of an atom, and $Z_{x,y} \equiv |r_{x,y}\rangle \langle r_{x,y}| - |g_{x,y}\rangle \langle g_{x,y}|$.

6.2 ORDERING DYNAMICS

We first study the nonequilibrium dynamics of the atom array after crossing the QPT into the ordered phase. The state-preparation stage is similar to that used in Refs.^{123,386}. First, all atoms are initialized in $|g\rangle$, which is the ground state for $\Delta/\Omega \ll 0$. While keeping Δ negative, Ω is ramped up to its final value, remaining constant until the end of the protocol (δ_i is held at 0 for this measure-

ment). Then, $\Delta(t)$ is swept from negative to positive values, through the quantum critical point and into the ordered phase. We use a linear sweep profile for all measurements. The sweep is halted at various endpoints within the ordered phase. Subsequently, Δ and Ω are held constant for a given hold time, and finally Ω is ramped down followed by a projective readout of the atomic states.

During the hold time, we probe the dynamics of the correlation length. To quantify the growth of correlations, we evaluate the two-point correlation function $G(\mathbf{r}, t)$ and the radially averaged structure factor $S(k, t)$, from which we extract a correlation length (ξ). In contrast to the Kibble-Zurek prediction, the correlation length grows significantly with hold time, indicating the gradual establishment of long-range order. Up to a hold time of $\sim 0.4\text{--}0.5\ \mu\text{s}$, we observe that the dynamics are consistent with a linear growth of ξ^2 with time, as expected for coarsening^{57,373}. Importantly, the rate of growth increases with proximity to the QPT. Motivated by the theoretical expectation of universality of these dynamics in the thermodynamic limit^{57,24,151}, we also study the structure factor as a function of a scaling variable $k\xi(t)$. We find that the data collapse onto a single functional form $S(k, t) \sim b(\xi(t))\xi^2(t)f(k\xi(t))$ for some scaling function f and an amplitude $b(\xi)$, which is suggestive of self-similarity. In addition to the ordering, we also observe long-lived oscillations of the correlation length. In what follows, the universal aspects and origin of these oscillations are explored in detail.

In order to gain further insight into the system's dynamics, we use single-site-resolved detection to identify the domains in each individual snapshot. We measure the probability that a given atom will appear as part of a domain of area A_d , and find that over time, increasingly larger domains are formed at the expense of their smaller counterparts. This is manifested in the growth of the mean area of the largest domain, concurrent with the shrinking of the second-largest one. Due to energy conservation, the appearance of progressively larger domains has to be offset by the proliferation of very small domains and single-site spin flips. This phenomenon is known as bidirectional flow, and is generic for coarsening in closed systems^{163,151}. We quantify this flow of energy by measuring the

spatial distribution of the classical energy, as defined by the diagonal contribution to Eq. (6.1) up to a constant shift:

$$H_{\text{cl}} = -\Delta \sum_i (n_i - 1) + \sum_{i < j} V_{ij} n_i n_j. \quad (6.2)$$

For every snapshot, we identify the domain walls and the bulk, and accordingly determine the contribution of each towards $\langle H_{\text{cl}} \rangle$. While the classical energy is indeed conserved over time, it is redistributed from the domain walls into the bulk. This is consistent with a picture of coarsening which is driven by the surface tension, and elimination, of domain walls⁵⁷.

6.3 DOMAIN WALL DYNAMICS

In order to study the real-time dynamics of domains and domain walls, we deterministically prepare specific configurations of domain walls using programmable, locally controlled light shifts. We apply site-dependent negative local detunings $\delta_i < 0$, with amplitudes $|\delta| \sim 4\Omega$ (see M, on the chosen atoms prior to ramping up Ω , and then continue the state-preparation protocol as previously described. The local detuning strongly biases the chosen atoms to $|g\rangle$, and consequently, locally favors either an $|\text{AF}_1\rangle$ or an $|\text{AF}_2\rangle$ configuration. After the sweep is completed, and before the hold time begins, the local detuning is quenched off and the state is allowed to evolve freely.

We start by preparing a small square domain of one AF order within the bulk of the other³²⁶. Upon removal of the local detunings, we observe that the area of the injected domain shrinks linearly with time. This observation is in agreement with coarsening dynamics for nonconserved fields, where surface tension due to the energy cost of domain walls generates curvature-driven dynamics^{273,57}. In such a scenario, the local velocity of a domain wall is proportional to its local curvature $1/R$: $\partial_t R \propto -1/R$, and therefore, $\partial_t R^2 = -v_a$, where v_a is some positive time-independent constant. Strikingly, we find that v_a increases as one approaches the quantum critical point. This behaviour is unique for coarsening in the vicinity of a *quantum* phase transition³⁷³; in contrast, for




Figure 6.3: Domains and energy transfer. **a**, Probability distribution of an atom belonging to an ordered domain of a certain size. At late hold times, atoms are more likely to participate in large domains and very small domains. **b**, The area of the largest domain grows, while that of the second-largest domain decreases as the system is held in the ordered phase. **c**, Classical energy of each snapshot during the hold time, calculated using Eq. (6.2). The total classical energy of the system is conserved while the bulk (domain wall) energy increases (decreases). Inset: separation of domains into bulk and domain walls for a single snapshot. The domain walls are identified as regions of the array where neither AF ordering is observed (light and dark blue). White and grey indicate AF_1 and AF_2 orderings respectively. The data presented here are for $\Delta/\Omega = 3.0$, with $\Omega/2\pi = 3.8$ MHz for **a,b** and $\Omega/2\pi = 6.0$ MHz for **c**.




Figure 6.4: Dynamics of seeded domains. **a**, We deterministically create a centered domain of AF_2 order surrounded by a backdrop of AF_1 order, by “pinning” each ground-state atom in the target ordering using a local light shift δ_i , while sweeping $\Delta(t)$ and maintaining a constant Rabi frequency Ω . Upper left: target amplitude of the local detuning pattern; $|\delta_i|$ is inversely proportional to the number of neighboring Rydberg atoms surrounding each ground-state atom. Center: Rydberg population of a single shot immediately after the local detuning is quenched off. Right: local staggered magnetization, demarcating the prepared domains. **b**, Evolution of the average local staggered magnetization with time, as the prepared state evolves under different values of Δ/Ω . **c**, The squared radius of the central domain, r^2 , decreases linearly with time. **d**, The rate of change of the area, r^2 , increases with proximity to the critical point, $(\Delta - \Delta_c)/\Omega$. The dashed line is a guide to the eye for the theoretically expected scaling relationship. **e**, Change in the radially averaged local staggered magnetization with Manhattan distances $d_m = 0, 2, 8$ from the center of the injected square domain. The order initially changes near the domain wall and begins to change in the center of the domain only at later times. Atoms in the bulk of the dominant ordering far from the domain wall, in gray, remain close to their initially prepared states. The dynamics are plotted for $\Delta/\Omega = 2.5$ and $\Omega/2\pi = 6.0$ MHz.

a classical Ising transition, the dynamics should be slower near the thermal phase boundary than when deep in the ordered phase²¹³; we also observe indications of this speedup in the global sweeps. We examine the dependence of v_a on the distance to the quantum critical point, $\Delta - \Delta_c$. Near the QPT, we find that v_a is approximately consistent with a scaling $\propto (\Delta - \Delta_c)^{-\nu}$ (where $\nu \approx 0.629$ is the correlation length exponent of the (2+1)D Ising QPT³⁷¹). We also analyze the evolution of several concentric spatial layers of the system, and observe the outer layers of the central domain morphing earlier with the dynamics moving progressively inwards. This supports a picture of coarsening in which the dynamics are indeed driven by the shrinking of domain walls, as opposed to being generated within the bulk of the domains. Note that atoms in the bulk, far from the domain walls, undergo very little change over the hold period, remaining mostly in their initially prepared state.

To further explore the curvature-driven nature of the coarsening dynamics, we prepare an initial state with a zigzag domain wall. Over time, the domain wall straightens into a vertical line separating the two orders. At points of high local curvature, the domain wall moves towards the center of the domain. This is dramatically different from the dynamics at points with no local curvature, where the initially prepared domain is aligned with the diagonal of the lattice ordering. At these points, the horizontal position of the domain wall remains remarkably stationary over time. This motion is also evident in the dynamically changing domain-wall shape. As the corners of the zigzag pattern change their order, the directly adjacent rows become locally curved and the domain wall at those points begins to move towards the center of the array. These key features of curvature-driven dynamics are in agreement with numerical simulations³⁷³.




Figure 6.5: Curvature-driven dynamics. **a**, We create two separate AF orderings with a zigzag domain wall between them. During the hold time, the domain wall straightens into a vertical line. **b**, Horizontal cuts of the staggered magnetization at points where the domain wall is locally curved (red, blue) versus straight (green, yellow). **c**, The domain wall's horizontal position for each row (as indicated by the color) as a function of the hold time. In rows where the domain is locally curved, the domain wall moves towards the center. All data shown are with $\Omega/2\pi = 6.0$ MHz; **b** is measured at $\Delta/\Omega = 3.0$ and **a,c** at $\Delta/\Omega = 2.5$.




Figure 6.6: Higgs mode oscillations. **a**, Long-lived oscillations of the staggered magnetization observed when one sublattice is initially pinned and then released. **b**, Numerical calculation of the energy difference, in a 10×10 system with periodic boundary conditions, of the initial pinned state (blue) and the $\mathbb{Z}_2$ state (green) relative to the ground state, as well as the gap to the first excited state (red). The pinned state is much lower in energy than the $\mathbb{Z}_2$ state at finite detuning. **c**, Schematic diagram of the effective potential for the amplitude mode. **d**, Oscillation parameters extracted from fitting $\varphi(t) \approx \varphi_0 + A \cos(\omega t + \theta_0)e^{-\gamma t}$ to data shown in **a**. Left: Measured oscillation frequencies (red points), plotted along with numerically determined values and calculated ground-state energy gaps (both for a 10×10 lattice). A lower local detuning $|\delta|$ leads to oscillations of lower amplitude and frequency (orange). Oscillation frequencies of the magnetization are indicated in purple. Upper right: the measured oscillations decay more rapidly as the phase transition is approached. Lower right: in the disordered phase, oscillations occur around $m_s = 0$, while in the ordered phase, a finite offset emerges. **e**, Oscillations of the correlation length. For final detunings close to the quantum critical point, these oscillations are superposed with substantial growth of the correlation lengths. The correlation length oscillates at double the frequency of the order parameter in the disordered phase, and matches the frequency of the order parameter oscillations in the ordered phase. All data shown are for $\Omega / 2\pi = 3.1$ MHz.

6.4 HIGGS MODE

In addition to the curvature-driven coarsening dynamics, our experiments clearly reveal persistent long-lived oscillations of the correlation length and the order parameter across a range of experimental parameters. We explore the origin of these oscillations. First, we apply local detunings to one of the two sublattices, which biases the order parameter. We then repeat the protocol, ramping to various values of Δ on both sides of the QPT. Directly after the ramp, we quench the pinning field off and follow the dynamics. We observe large-amplitude, long-lived oscillations of the order parameter, well modeled as a damped harmonic oscillator, $m_s(t) \approx \varphi_0 + A \cos(\omega t + \theta_0)e^{-\gamma t}$, with amplitude A , frequency ω , damping γ , and offset φ_0 that strongly depend on Δ . We find that upon approaching the phase transition from both sides, ω decreases while γ and A increase; φ_0 changes from zero in the disordered phase to a nonzero value in the ordered phase.

To understand the origin of these observations, we perform numerical simulations using matrix product state (MPS) methods. Through density-matrix renormalization group (DMRG) calculations on up to 10×10 sites with periodic boundary conditions, we find that the pinned initial state corresponds to a low-energy state of the post-quench Hamiltonian even on the disordered side of the transition. This is in contrast to the high-energy $\mathbb{Z}_2$ states typically associated with oscillations due to quantum many-body scars^{34,44}. We then attempt to simulate the dynamics of the pinned initial state using the time-dependent variational principle (TDVP)¹⁸⁷ at a relatively small bond dimension $\chi = 256$ and find good qualitative agreement with the experimentally extracted frequencies. Moreover, the oscillation frequencies closely match the numerically determined ground-state gap, and are relatively robust to variations in the system size away from the critical point.

Our experimental findings, coupled with the numerical insights into their low-energy nature, suggest that the observed oscillations can be understood as an amplitude (Higgs) mode³²⁷, which is a collective excitation of the order parameter. Qualitatively, the Higgs mode can be viewed through

the lens of Landau theory. In this framework, order-parameter dynamics are determined by an effective potential describing a quartic anharmonic oscillator: $V(\varphi) = \frac{q}{2}\varphi^2 + \frac{\lambda}{4}\varphi^4 + \mathcal{O}(\varphi^6)$ (we continue to identify the oscillation on the disordered side as a "Higgs mode" for convenience). The ordered phase ($q \sim \Delta_c - \Delta < 0$) is differentiated from the disordered phase ($q > 0$) by a finite value $\varphi_0 \neq 0$ of the potential minima, which determines the offset of the oscillation. Beyond the changing offset φ_0 , this simplified picture reproduces the increase of the oscillation amplitude $\mathcal{A}$ and the decrease of the frequency ω when approaching the phase transition. To further investigate this amplitude mode, we prepare lower-energy biased states by softening the pinning field (smaller $|\delta_i|$). We find a corresponding decrease in the oscillation amplitude $\mathcal{A}$ and the frequency ω with a sharper dependence near the QPT, in agreement with Landau theory. The oscillation frequency of these lower-energy states shifts close to the many-body gap. In order to explore the Higgs mode deep in the ordered phase, where our state preparation scheme generates low-amplitude oscillations, we use an alternative experimental protocol.

The Higgs-mode dynamics present a unique probe of the quantum critical point. In particular, the ratio of the oscillation frequencies on the two sides of the QPT is universal in equilibrium, and predicted to be $\omega(-|q|)/\omega(|q|) = \sqrt{2}$ ³⁵⁹ by Landau mean-field theory. However, our experimental results, in which $\omega(-|q|)/\omega(|q|) > \sqrt{2}$, indicate a significant deviation from this simplistic prediction, broadly consistent with more advanced calculations predicting $\omega(-|q|)/\omega(|q|) \approx 1.9$. The discrepancy with mean-field results emphasizes the central role of quantum fluctuations, and in particular, finite-momentum order parameter fluctuations, in the vicinity of the QPT. The dynamics of these fluctuations are also expressed in the progressively larger oscillations of the correlation length and a sharp increase in the damping term γ observed upon approaching the critical point.

6.5 DISCUSSION AND OUTLOOK

Our observations shed new light on paradigmatic collective processes in closed nonequilibrium quantum many-body systems. In particular, they highlight the important role of coarsening dynamics, and reveal their curvature-driven character in systems with a nonconserved order parameter^{57,373}. Crucially, we measure an acceleration of the ordering processes when approaching the phase transition, a signature of the intrinsically quantum nature of the dynamics³⁷³. While we observe a scaling collapse of the structure factor suggestive of self-similarity⁵⁷, the dynamically varying amplitude $b(t)$ deviates from the expected universal behavior. This deviation could originate from finite-size effects (as well as, potentially, residual disorder or decoherence), and its detailed understanding constitutes an interesting theoretical problem²⁴. Similar mechanisms may account for the slowdown of coarsening at late times observed in Figs. 6.2c,d. Further evidence for the role of finite-size effects is provided by measurements involving local control: we find that the dynamics of domains seeded away from the system’s boundaries are in much closer agreement with universal theoretical predictions.

Additionally, we observe the concurrent excitation of the Higgs mode upon crossing the phase transition. Investigation of this mode yields detailed information on important observables—such as the damping rate of the Higgs mode in the vicinity of the phase transition—which are difficult to access classically³⁵⁸. For the numerically accessible system sizes and bond dimensions, our simulations cannot capture the damping rate γ , and generally break down near the critical point. More generally, the possible interplay of coarsening with the Higgs mode presents an intriguing question that warrants further theoretical investigation.

These studies can be extended along several directions. In contrast to traditional condensed-matter systems, programmable quantum simulators can directly access correlation functions of *any* order³⁸⁶ as well as other important observables, such as the entanglement entropy^{67,420}, using,

e.g., hybrid digital-analog approaches^{43,13}. These could provide further insights into complex dynamics, particularly near the quantum critical point, where numerical calculations are prohibitively challenging. Besides the symmetry-broken ordered states probed in this work, it would also be interesting to extend our study of coarsening dynamics to the formation of topologically ordered states of matter^{386,372,429}, which cannot be characterized by local order parameters. Moreover, quantum dynamics following a first-order transition can also be explored. For instance, the local programmability of the system could allow for the preparation and biasing of metastable states^{105,106}. Starting from these states, the quantum system can order through a high-order tunneling process, known as false vacuum decay, that nucleates a critical-size domain of the true ground-state ordering^{461,96}.

7

Architecture for digital-analog quantum simulations

Additional details for this chapter can be found in Ref.¹⁵³.

Understanding emergent phenomena in out-of-equilibrium interacting many-body systems is an exciting frontier in physical science^{15,160,126}. While quantum simulators represent a promising approach to this long-standing problem³³⁸, in practice it can be challenging to directly realize the

required interactions, measure arbitrary observables, and mitigate errors. Here we use coherent mapping between the Rydberg and hyperfine qubits in a neutral atom array simulator⁴² to engineer and probe complex quantum dynamics. We combine efficient analog dynamics with fully programmable state preparation and measurement, leverage non-destructive readout for loss information and atomic qubit reuse, and use an atom reservoir for replacing lost atoms. With this analog-digital approach, we first demonstrate dynamical engineering of ring-exchange and particle-hopping dynamics via Floquet driving¹⁶⁴, and measure the spectral function of single excitations by evolving initial superposition states²⁴⁶. Extending these techniques to a 271-site kagome lattice, we employ closed-loop optimization³⁵⁴ to target an out-of-equilibrium critical quantum spin liquid of the Rokhsar-Kivelson type³⁵³. We observe the key features of such a state, including the absence of local order, many-body coherences between nearly equal-amplitude dimer configurations over up to 18 sites, and universal correlations consistent with predictions from field theory¹⁴². Together, these results pave the way for using dynamical control in analog-digital quantum simulators to study complex quantum many-body systems.

7.1 INTRODUCTION

Emergent phenomena in interacting many-body systems lie at the frontier of multiple subfields across the physical sciences^{15,160,126}. Understanding these phenomena is challenging, as numerical simulations rapidly become intractable in strongly correlated systems, while direct microscopic probes are typically inaccessible in real-world systems and materials. Quantum processors offer a route to large-scale, programmable simulations of many-body systems and their dynamics. In particular, analog quantum simulators, which aim to directly realize a quantum Hamiltonian²⁷⁹, have already shed light on diverse phenomena ranging from quantum phase transitions^{179,380} and quantum magnetism^{116,450}, to non-equilibrium dynamics¹³ and thermalization^{235,383}. To access new




Figure 7.1: Analog-digital atomic processor for probing many-body quantum matter. **a**, Analog evolution is performed in the interacting Rydberg qubit manifold and the hyperfine spin qubit is used for digital processing. Non-destructive readout of each atom obtains spin and loss information, after which lost sites are refilled from the reservoir and the spin state is reinitialized. The circuit repeats until the reservoir is depleted. The geometry of the simulation zone is programmed for each experiment, shown here for a 271-site kagome lattice. **b**, Two π -pulses are used to map between the qubit manifolds, shown for mapping down from the Rydberg to the hyperfine qubit. **c**, Rabi oscillations on the Rydberg qubit, measured after mapping down. The contrast is 97.8(1)% and loss is $\approx 1.8\%$ (not shown), dominated by loss when turning off the tweezer traps and rearrangement infidelity as the reservoir depletes. Here and later, all data are locally postselected on no atom loss. **d**, Ramsey fringe showing coherence after mapping down $|+\rangle$. Atoms are located at next-nearest-neighbor sites for the 1D chain and 2D kagome lattice, emulating a blocked many-body state. The contrast is reduced in 2D where effects of long-range Rydberg interaction tails are larger. **e**, Many-body states and dynamics can be engineered and probed using this zoned analog-digital architecture.

regimes of many-body physics, recent efforts have focused on tailoring the native dynamics^{109,167,457} or engineering the necessary interactions through periodic driving^{330,300,308}. In parallel, the advent of gate-based digital simulation is rapidly opening doors to new physical models and complex observables^{95,130,443}, although large gate counts and circuit depths pose a challenge for near-term implementations at large scales³⁹.

Hybrid approaches that combine continuous analog evolution with coherent digital computation offer a promising route to overcoming these challenges³³⁸, leveraging the efficiency of analog dynamics while retaining the universal control afforded by quantum circuits^{387,13,67}. Here we realize an analog-digital quantum simulator in a Rydberg atom array based on coherent mapping between the interacting Rydberg and digital hyperfine qubit manifolds. We use this architecture to

first explore Hamiltonian engineering via periodic driving, employing programmable initial states to study six-body ring exchange and two-body hopping as examples. These represent key ingredients for studying dynamics and excitations of complex phases, and notably underpin the realization of dynamical lattice gauge theories which are difficult to simulate classically¹⁸⁸.

We then apply our hybrid approach to two-dimensional systems, focusing on the $U(1)$ quantum spin liquid (QSL)^{143,142,353,302,346}, an exotic state whose experimental observation remains elusive. Featuring long-range entanglement and no local order, it is a paradigmatic example of a strongly-correlated system described by an emergent gauge theory similar to electromagnetism^{143,142} and is closely linked to the resonating valence bond theory introduced by Anderson¹³³. As compared with the recently studied gapped $\mathbb{Z}_2$ QSLs^{386,376}, $U(1)$ QSLs are more challenging to realize due to their unstable, gapless nature¹⁴². Enabled by the increased experimental cycle rate, we target an out-of-equilibrium $U(1)$ QSL in the honeycomb quantum dimer model using analog state preparation optimized through a closed-loop quantum-classical procedure³⁶⁹. We observe microscopic ordering characteristic of a Rokhsar-Kivelson state³⁵³, correlations consistent with field theory predictions, and further employ digital gates to directly measure many-body coherences.

7.2 FAST ANALOG-DIGITAL QUANTUM SIMULATIONS

In our experiments, programmable arrays of ^{87}Rb atoms trapped in optical tweezers are used to implement both analog and digital evolution^{130,42}. Analog evolution is implemented via excitation of the $|1\rangle \leftrightarrow |r\rangle$ Rydberg qubit with global laser pulses, such that the many-body unitary evolution $U(t)$ is governed by the Hamiltonian

$$\frac{H}{\hbar} = \frac{\Omega(t)}{2} \sum_i \sigma_i^x - \sum_i n_i (\Delta(t) - \delta_i(t)) + \sum_{i<j} V_{ij} n_i n_j \quad (7.1)$$

where Ω is the Rabi frequency, Δ is the global detuning on the Rydberg transition, δ_i are site-resolved detunings, $\sigma_i^x = |1_i\rangle\langle r_i| + |r_i\rangle\langle 1_i|$ and $n_i = |r_i\rangle\langle r_i|$ is the Rydberg occupation number. $V_{ij} = V_0/|\mathbf{x}_i - \mathbf{x}_j|^6$ denotes the van der Waals (vdW) interaction which prevents multiple Rydberg excitations within the blockade radius $R_b/a = (V_0/\Omega)^{1/6}$ for lattice spacing a ; throughout we use $R_b/a \approx 1.3 - 1.4$, forbidding nearest-neighbor excitations.

Digital processing is performed in the long-lived hyperfine clock qubit, $|0\rangle \leftrightarrow |1\rangle$, on which optical Raman transitions implement fully programmable single-qubit gates. Fast Raman and Rydberg π -pulses coherently transfer the quantum state between the two qubit manifolds. We illustrate this method for single qubits by performing Rydberg Rabi oscillations and mapping the resulting state down to the hyperfine qubit for readout. We further benchmark the preservation of phase coherence during the mapping, where the slightly reduced Ramsey contrast is due to crosstalk from the long-range tails of the vdW interaction during the mapping time.

This hybrid approach allows for two key features. First, it enables state preparation and measurement beyond the Pauli Z -basis. Second, non-destructive readout of the hyperfine qubit⁴² enables error detection based on loss information, where atom loss is a dominant source of error in Rydberg simulations that conventionally results in a measurement error. This includes correlated loss events previously observed during analog Rydberg evolution⁵³. The mapping protocol converts state initialization errors and population in blockade-violating states into these detectable loss events. Finally, since atoms are retained after readout, they can be reinitialized and reused, thereby increasing the experimental cycle rate by over an order of magnitude. Together these features give access to new observables, provide native error-mitigation techniques, and allow for measurements with high sample complexity.




Figure 7.2: Hamiltonian engineering via periodic global detuning drives. **a**, Example detuning profiles per Floquet cycle. Gray profile implements a many-body echo via π phase jumps of the Rydberg drive, realized as smooth detuning pulses. Perturbations are added on top of this echo to engineer an effective Hamiltonian (red profile). **b**, Schematic of a single Rydberg excitation initialized at the center of a 31-site chain, used as the initial state in **c-e**. **c**, Dynamics of microstate populations, grouped by Hamming distance from the initial state in the Z -basis and computed over the central 19 sites of the chain. Under the many-body echo, the state periodically revives, here twice per Floquet cycle (left). The Floquet drive perturbs the dynamics, illustrated by the modified evolution of the populations, to yield an effective Floquet Hamiltonian (right). **e**, At stroboscopic times, the number of spin excitations is approximately conserved. **f**, Single-spin quantum walk under an effective blockaded XX Hamiltonian. The hopping rate is tuned via the drive parameter ε . **g**, Detuning profile per Floquet cycle used to engineer simultaneous two-body hopping and ring-exchange dynamics, and resulting dynamics on a six-site hexagonal plaquette starting from **h**, a single spin and **i**, a Néel state. Data are averaged over 15 plaquettes evolved in parallel. The decay in contrast is primarily attributed to the sensitivity to static detuning offsets. Explicit parameterizations for **a** and **g**.

7.3 DYNAMIC FLOQUET HAMILTONIAN ENGINEERING

We first explore the engineering of interactions under nearest-neighbor Rydberg blockade in small systems using periodic driving^{249,134}. The key idea is to use a time-periodic global detuning drive to engineer a Floquet Hamiltonian, inspired by the dynamical control of quantum many-body scars⁴⁴. In the simplest case, detuning π -pulses implement a many-body echo, resulting in no stroboscopic dynamics⁴⁴. Perturbations away from this echo can then be used to control the evolution and generate the desired Floquet Hamiltonian, given by the sum of time-evolved operators over each period. Notably, although the intermediate dynamics generally involves many non-number-conserving states, we can engineer approximate number conservation at stroboscopic times.

Using this approach, we experimentally realize two types of effective Hamiltonians. First, we realize a hopping Hamiltonian with chemical potential μ , nearest-neighbor blockaded hopping energy J , and next-nearest neighbor interaction U , which can be continuously tuned by varying the parameters of the periodic detuning profile $\Delta(t)$ ²⁴⁹. We observe a single-particle quantum walk in a one-dimensional spin chain for increasing J , where the characteristic interference pattern is visible in the density profile. Importantly, this approach enables faster number-conserving hopping dynamics than can be reached at large static detuning. These results demonstrate that the Floquet protocol works in practice and is well captured by the ansatz hopping Hamiltonian, even beyond the regime of small perturbations to the many-body echo.

Next, we demonstrate the ability to engineer higher-body interactions. Specifically, on hexagonal plaquettes we generate ring exchange dynamics by additionally engineering three-body interactions such that, at large μ , the six-body ring exchange appears at second order¹³⁴. To isolate the dynamics associated with the two-body and six-body dynamics, we prepare and map up initial states of a single spin excitation and a Néel state, respectively. Crucially, we observe that the hopping and ring-exchange dynamics occur on the same timescale, quantitatively different to perturbative dy-

namics generated at large static detuning at higher orders in Ω/Δ . Such dynamics can also stabilize entangled many-body states.

Finally, we explore techniques for directly probing the resulting quantum dynamics. Focusing on the hopping Hamiltonian, we initialize and evolve two Rydberg excitations separated by a single site, and observe a clear modification of the interference pattern for the two-spin dynamics compared to that of isolated single spins. This can be understood as the blockade constraint inducing an effective repulsive spin-spin interaction, which reduces the density within the overlapping light cones of the two spins. This measurement therefore reveals the interacting nature of the Hamiltonian. To gain further insight into the low-energy properties of the engineered Hamiltonian, we employ many-body spectroscopy which directly measures the single-particle dispersion²⁴⁶. The key idea is to perform an interferometric measurement between a reference state with no excitations, $|\Psi_\downarrow\rangle = |1\rangle^{\otimes 31}$, and a probe state with a single excitation, $\sigma_0^x |\Psi_\downarrow\rangle$. We prepare the central spin of the chain in one of four superposition states in the hyperfine qubit, map this state up to Rydberg qubit, evolve under the Floquet Hamiltonian for time t , and map back down for measurement in the X -basis. Using these measurements, we extract the Green's function $\mathcal{G}(r, t) = \langle \Psi_\downarrow | \sigma_r^x(t) \sigma_0^x(0) | \Psi_\downarrow \rangle$ and plot its Fourier transform. The fitted dispersion relation $\omega(k)/2\pi = 0.407(2) - 0.281(2) \cos ka$ has the expected tight-binding form, and the real and imaginary parts of the spectral function are qualitatively consistent with Kramers-Kronig relations. Furthermore, this measurement demonstrates simultaneous phase coherence of the hyperfine and Rydberg qubits throughout the entire evolution.

7.4 NON-EQUILIBRIUM QUANTUM SPIN LIQUID

We next employ these tools for preparing and probing complex many-body states. Specifically, we study the quantum dimer model (QDM) on the honeycomb lattice, which is predicted to host a




Figure 7.3: Dynamical probes of interactions and excitations. **a**, Two-spin quantum walk under an effective blockaded XX Hamiltonian. **b**, Taking the difference with the summed density profiles of the two individual spins reveals an effective spin-spin repulsion. For clarity, we subtract the background evolution of the vacuum initial state. **c**, Protocol for many-body spectroscopy of a single spin excitation on the central site, $r = 0$. In this case, the ancilla in $|+\rangle$ is replaced by preparation of the central site in four bases, $\alpha \in \{\pm X, \pm Y\}$, and classical post-processing to obtain $\mathcal{G}(r, t)$. **d**, The Fourier transform, $\mathcal{G}(k, \omega)$, as a function of momentum, k , and frequency, ω , across the 1D Brillouin zone. Left schematic illustrates the spin wave excitation. The imaginary component is the spectral function, whose peak is fit to obtain the single-particle dispersion. From this we extract $(\mu + 4U)/2\pi = 0.407(2)$ MHz and $J/2\pi = -0.140(1)$ MHz, and independently verify these coefficients using Hamiltonian learning of the dynamics of different initial states. Background evolution, appearing at $k = 0$, is subtracted.




Figure 7.4: Dynamical preparation of U(1) quantum spin lakes. **a**, Potential and kinetic energy per hexagonal plaquette in the Rokhsar-Kivelson Hamiltonian. Potential energy counts the number of flippable plaquettes and kinetic energy captures the plaquette resonances. **b**, Mapping of Rydberg atoms on the vertices of the kagome lattice to a QDM on the honeycomb lattice. The main data in this experiment uses $R_b/a \approx 1.32$. **c**, A multi-objective Bayesian optimizer is used in a closed loop with experiment. Each experimental run returns X and Z basis snapshots of the prepared state from which observables are computed; the RK energies and mean Rydberg density, n , are maximized, whilst atom loss is minimized. The preparation protocol is updated automatically and a new run is started. **d**, (Top) Best-performing detuning sweeps obtained from three overnight optimization runs with eight free parameters. The total time was allowed to vary to produce the longest sweep and was fixed for the shorter sweeps. Dashed line marks the ground state phase transition. (Bottom) Corresponding evolution of loss-postselected mean Rydberg density. Oscillations are due to imperfect adiabaticity starting from all qubits in $|1\rangle$. Dotted line indicates $1/3$ filling, corresponding to maximal packing. **e**, Final monomer, dimer, and blockade-violating double-dimer populations per vertex for the $0.8\ \mu\text{s}$ drive. **f**, Evolution of the kinetic and potential RK energies. The reduction of $\langle T \rangle$ relative to $\langle V \rangle$ is larger at later times as the Rydberg populations, and thus vdW tail energies, increase. **g**, (Top) Schematic of the three kagome sublattices (A,B,C) per vertex. (Bottom) Local light shifts are applied to the 74 boundary atoms to compensate for their reduced mean-field vdW interaction energy. **h**, The minimum in the nematic order parameter, Φ , corresponds to uniform bulk Rydberg population on the three sublattices. Dashed line marks the detuning strength used for all subsequent data. **i**, Angular distribution of Φ reveals the emergent global rotational symmetry. Three sharp peaks at $0, \pm 2\pi/3$ are expected in a nematic phase.

deconfined QSL with an emergent $U(1)$ gauge symmetry arising from its bipartite nature^{302,142}. The QSL state is the equal-amplitude and equal-phase superposition of all perfect dimer coverings and is found at the critical point of the QDM with the Hamiltonian $H_{\text{RK}} = -tT + vV$, where T is the kinetic energy and V is the potential energy. T captures the quantum coherences of the state, favoring plaquettes which can resonate via ring exchange, while V penalizes such ‘flippable’ plaquettes. Known as the Rokhsar-Kivelson (RK) Hamiltonian³⁵³, this model features a critical point, $v = t$, known as the RK point.

To realize the QDM, atoms are placed on the vertices of the kagome lattice and encode the presence of a dimer on the bonds of the medial honeycomb lattice. Rydberg blockade enforces that each vertex is touched by at most one dimer, with vertices without a dimer referred to as monomers. While at large Δ the ground state of the Rydberg Hamiltonian, Eq. (7.1), on this lattice is a nematic phase³⁷², it was recently pointed out that, out-of-equilibrium, finite-sized regions of QSL order known as quantum spin ‘lakes’ can emerge^{369,161}. We therefore start by optimizing state preparation protocols in this out-of-equilibrium regime, further motivated by theoretical observations that tailored protocols can mitigate the impact of, for example, van der Waals tails and decoherence^{162,433,463}. More specifically, we employ a closed-loop optimization of different detuning sweeps, using the measured observables from the quantum evolution to update the sweep parameters and target the desired RK state (i.e., the ground state at the RK point). We use a Bayesian optimizer to simultaneously maximize $\langle T \rangle$, $\langle V \rangle$, and mean Rydberg density whilst minimizing atom loss, and find preparation sweeps for three evolution times that all approach the maximum filling of $1/3$.

Two characteristic features emerge from the optimized detuning sweeps. First, longer sweeps improve performance by slowing down at early times, indicating the initial evolution is quasi-adiabatic. Second, the sweeps speed up before the phase transition—in contrast to ground state preparation—and reach approximately the same final sweep rate for all protocols. Intuitively, this can be under-

stood as follows: the quasi-adiabatic evolution prepares an equal phase superposition of states and the final rapid sweep implements a dynamical projection onto the maximally-packed dimer coverings. During this final stage, the sweep must be sufficiently slow to adiabatically remove monomer excitations, while not so slow as to resolve any energy differences between dimer coverings and create phase excitations. As a result, there exists an optimum in the final sweep rate that balances these competing energy scales, thus explaining why this rate is approximately the same across the different optimized sweeps.

Focusing on the slowest sweep, we find a high final dimer population alongside finite monomer and double-dimer densities, arising from excitations as well as dressing of the low-energy maximally-packed dimer manifold at finite Δ/Ω . The RK kinetic and potential energies grow in step with these dimer populations; we explore these energies more below. Notably, we find that local loss postselection improves the state fidelity, and closely reproduces the results obtained when postselecting on the absence of loss across the entire lattice.

7.5 SIGNATURES OF A ROKHSAR-KIVELSON STATE

A key property of the RK state is the absence of local order, implying that seeding of order by the lattice boundaries must be suppressed. Such boundary effects imbalance the population on the three kagome sublattices (A,B,C), which can be quantified by the nematic order parameter, $\Phi = \sum_j \left(n_j^A + n_j^B e^{2\pi i/3} + n_j^C e^{-2\pi i/3} \right)$, where the sum runs over vertices. By applying static local light shifts to the boundary atoms, we prepare a uniform density across all three sublattices and, moreover, find that the distribution of the phase of the nematic order parameter is close to uniform, revealing the emergent rotational symmetry of the state and distinguishing it from a nematic phase.

To further probe the microscopic order, we next examine the probability distribution of dimer coverings. By choosing small subsystems and projecting the surrounding bonds onto not having a




Figure 7.5: Superposition of dimer coverings. **a**, Probability distribution of dimer configurations in isolated subsystems. The probabilities within each subsystem are nearly equal after projecting surrounding bonds onto no dimer ($|0\rangle$ in the hyperfine qubit) and restricting to perfect dimer coverings without defects, i.e., monomers and double-dimers. Inset shows a non-uniform distribution, prepared with a weakly modified preparation sweep, that is fit to a Boltzmann distribution with inverse temperature β . **b**, Schematic of the RK state. Two local configurations are isolated by projecting bonds around each single-plaquette subsystem onto no dimer. **c**, Bare and projected potential and kinetic energies, averaged over all plaquettes in the lattice. The RK prediction is $\langle T \rangle = \langle V \rangle \approx 0.29$ from classical Monte Carlo simulations. Defects reduce this expectation value in experiment, and the relative reduction of the kinetic energy is largely attributed to measurement error. **d**, Local measurement bases for the projected kinetic energy. **e**, Parity oscillations on single plaquettes using local Raman to apply Z rotations on every other site after mapping down. $U(t)$ represents the analog state preparation. **f**, Bare and projected X loops of increasing perimeter. Solid points are averaged over different measured geometries, shown individually as small points.

dimer in post-processing, we isolate the local dimer covering from the surrounding lattice. Within such subsystems of up to 27 atoms plus 12 projected bonds, we observe that the probability distribution across perfect dimer coverings is almost flat. To investigate the robustness of this uniformity, we slightly perturb the state preparation profile and find that even small modifications result in an unequal distribution. As an example, we fit such an instance to a thermal distribution at finite temperature, where the classical energy of each subsystem is equal to its internal van der Waals interaction energy, finding excellent agreement. In contrast, for the original sweep profile, the uniform distribution corresponds to an approximately infinite-temperature state that exactly coincides with the equal-amplitude prediction for the RK state. We further perform the same analysis for different projected boundary conditions around each subsystem and again find good agreement with the RK prediction.

We next turn to the many-body coherences between dimer covering. Closed X loop operators, of

which the kinetic energy operator is the smallest example, map different dimer coverings onto each other such that the expectation probes the coherences of the state. The expectation value of these loops, even in the ideal QSL, is not unity; that is, it is not a stabilizer state. However, projecting all the bonds surrounding the chosen X loop onto not having a dimer isolates the states with finite expectation value such that $\langle V \rangle_{\text{proj.}} = \langle T \rangle_{\text{proj.}} = +1$ for the RK state. For the potential energy, we find an $\approx 5 \times$ increase in the expectation value after projection. To measure the projected kinetic energy, after mapping down we apply local $Y^{1/2}$ gates to the sites in one of three plaquette sublattices. Although we find the kinetic energy is reduced relative to the potential energy, we attribute this predominately to microstate-dependent phase accumulation during the mapping down that acts as an off-diagonal measurement error.

Intuitively, the X loops are local GHZ-like states on the loop perimeter embedded within the many-body state. To test this we perform parity oscillations, again on one plaquette sublattice at a time. Using local hyperfine Z rotations on alternate sites along the loop, we measure coherent oscillations in $\langle X^6 \rangle$ with a many-body phase given by $3 \times$ the applied single-qubit phase, as expected for the state $\frac{1}{\sqrt{2}}(|101010\rangle + |010101\rangle)$. We similarly find that the loop expectation value is insensitive to the global readout phase.

Extending to larger loops of up to 33 sites in total, where the loop perimeter is always measured in the X -basis and surrounding atoms in the Z -basis for projection, we observe a slow decay in the projected expectation value with loop perimeter and a faster decay in the bare loop. This decaying ratio $\langle X^N \rangle / \langle X^N \rangle_{\text{proj.}}$ reflects the exponentially increasing number of dimer configurations over larger subsystems³⁰². The perimeter-law decay of the projected X loops, which are all $+1$ in the ideal QSL, is due to virtual ‘magnetic’ excitations, i.e., phases between dimer configurations generated at the endpoint of open Z strings, alongside incoherent errors³⁸⁶.




Figure 7.6: Correlations in the emergent $U(1)$ gauge theory. **a**, Dimer-dimer connected correlator. Solid points show the absolute value of correlations restricted to dimers within plaquettes in the same row. Correlations decay then plateau, indicative of weak long-range order, with additional non-universal structure arising from the lattice geometry. Small points are radially averaged over the lattice after sign rectification. **b**, Spatial plot showing the signed correlations with respect to the central site. **c**, Definition of the dimer-string correlator between plaquette centers with FM normalization. **d**, Dimer-string correlator, restricted to plaquettes in the same row (inset). The corresponding expectation value for the RK state is shown for two lattice dimensions (in units of plaquettes) obtained from Monte Carlo simulations. Monte Carlo simulations use periodic boundary conditions and the results are averaged over all four topological sectors. **e**, Relation between magnitudes of rectified dimer-dimer and dimer-string correlators. Field theory predicts a ratio of 4 on logarithmic axes. Each marker corresponds to the correlators at a given vertex separation.

7.6 EMERGENT UNIVERSAL CORRELATIONS

To probe the emergent $U(1)$ gauge theory, we next study the correlations of the state. We focus on two types of correlators —dimer-dimer and dimer-string— which are predicted by the field theory to decay as $1/r^2$ and $1/r^{1/2}$, respectively¹⁴². Remarkably, at the RK point, such Z -basis correlators on a lattice are known analytically¹³⁹ and exactly coincide with the predictions of continuum field theory at long distances, after accounting for the lattice-dependent substructure. We first investigate the dimer-dimer correlations $C_d(|\mathbf{x}_i - \mathbf{x}_j|) = \langle \sigma_i^z \sigma_j^z \rangle - \langle \sigma_i^z \rangle \langle \sigma_j^z \rangle$ where $\sigma_i^z = 2n_i - 1$, finding an exponential decay of the magnitude and an antiferromagnetic sign structure. The plateau at large separations suggests the presence of weak long-range order induced by the open boundary condi-

tions, which explicitly break the $\mathbb{Z}_3$ symmetry of the honeycomb lattice. We then turn to the dimer-string correlator, $C_s(|\mathbf{x}_i - \mathbf{x}_j|) = \prod_{\ell \in \Gamma_{ij}} \sigma_\ell^z$, that measures the dimer parity along an open string Γ_{ij} between two plaquettes. Similar to the definition of the Fredenhagen-Marcu (FM) order parameter for topological order¹⁴⁴, we normalize by the closed loop parity, mitigating the effect of dressing of the string by monomers and incoherent errors that both result in a perimeter-law decay of closed loops^{386,429}. The resulting expectation value displays a characteristic sign structure and magnitude decaying to zero for longer strings, once again ruling out the presence of crystalline ordering. We further measure normalized open X strings that probe for monomers created at the string endpoints and find that the expectation value is small and decays with increasing string length, consistent with a deconfined gauge theory^{178,429}.

To reveal the universal physics, we analyze rectified dimer-dimer and dimer-string correlators, where the rectification maps these both onto vertex-vertex correlators described by the field theory. We compare the logarithms of the two rectified correlators and find good agreement with the universal ratio of 4 predicted by the field theory, which is expected to hold even at finite system size. We attribute the exponential decay of the correlators to finite-size effects and the presence of monomers. Both arguments from field theory and classical Monte Carlo simulations of the RK state are consistent with such an exponential decay on a finite-size lattice, converging toward algebraic correlations with increasing size. Interestingly, we find the decay of the correlators is largely independent of lattice height, with a difference emerging only in the magnitude of the plateau in the dimer-dimer correlations. Our results are therefore consistent with the preparation of quantum spin lakes that appear locally like the RK state, with a length scale set by the monomer excitation density³⁶⁹. Furthermore, we find close agreement with unitary time-dependent matrix product state simulations of the experimental protocol, indicating that incoherent errors do not appreciably affect these results.

7.7 DISCUSSION AND OUTLOOK

These observations reveal the key signatures of a non-equilibrium $U(1)$ Rokhsar-Kivelson state and provide insights into the complex many-body physics of this state. This is especially notable given the instability of the $U(1)$ QSL in equilibrium and its gapless nature, which together render it inaccessible via either adiabatic sweeps or finite-depth digital circuits. In this out-of-equilibrium regime, we find that microscopic properties are sensitive to the specific protocol used for state preparation, highlighting the value of experimental optimization, informed by theory. Rather remarkably, the key features of the resulting state are well-described by the equilibrium theoretical predictions of the Rokhsar-Kivelson Hamiltonian at its critical point. Quantifying the deviations from these predictions at longer length scales and developing a detailed understanding of the interplay between finite monomer density, system size, and open boundary conditions constitute interesting directions for future explorations.

These results can be extended along several directions. By using local measurements, an approximate equilibrium parent Hamiltonian can be obtained using Hamiltonian learning techniques³⁴¹, potentially enabling studies of the deformations from the ideal Rokhsar-Kivelson model. By engineering such Hamiltonians using, e.g., Floquet techniques, one could then study low-energy excitations relevant to understanding emergent gauge dynamics in resonating valence bond states. This and other Hamiltonian engineering approaches could make use of variational quantum circuits^{76,234}, possibly combined with closed-loop optimization incorporating control over analog evolution, digital gates, and coherent atom rearrangement. Furthermore, such methods can be used for inverse quantum simulation²⁴⁸ to design systems or materials with desired properties.

More broadly, these experiments demonstrate the potential of a hybrid analog-digital approach to quantum simulation. Digital controls enable local operations, error detection, and fast experimental cycle rates, while analog evolution efficiently implements many-body interactions. We find

coherent mapping to be a powerful tool and envision several pathways to further improve its performance. Error mitigation based on loss information improves effective fidelities within existing technical capabilities. Using atomic motion to toggle between interacting and non-interacting geometries enables measurement of two-time correlation functions²⁸⁰ or entanglement witnesses⁴³², and could be used for circuit-based state preparation^{95,376,130} or to employ state purification techniques⁸⁹. Extending to interleaved layers of digital and analog evolution may enable new opportunities for quantum simulation or sensing²³⁴. Many of these techniques are applicable to a range of platforms^{234,13,387}, greatly expanding the breadth of accessible many-body physics in practical quantum simulators.



Further improvements to entangling CZ gate fidelity

Additional details for this chapter can be found in Ref.¹³² and Appendix F.

Creation and manipulation of entanglement with low error is essential in quantum information systems. In practice, two-qubit entangling gates constitute a dominant error source, limiting circuit depths and performance in fault-tolerant architectures. Using a neutral-atom quantum processor,

we realize entangling CZ gates with a high Rabi frequency smooth-amplitude pulse, employing state-selective readout and qubit reuse for fast calibration, and achieve state-of-the-art fidelities of 99.854(4)% which improve to 99.941(3)% upon loss postselection, with stable performance for 10 hours. We then use these low-error gates in quantum circuits with coherent atom rearrangement. We first benchmark performance by creating and disentangling cluster states, and subsequently implement scrambling circuits featuring longer-range connectivity to study non-locally entangled states generated through chaotic dynamics. These results pave the way toward deep-circuit, efficient fault-tolerant quantum computation.

8.1 INTRODUCTION

High-fidelity entanglement is at the heart of quantum science and technology, enabling wide-ranging applications in computing, simulation, metrology, and sensing³³⁸. For instance, error rates around 0.1% are generally believed to be necessary to enable efficient performance of quantum error correction (QEC), central to realizing utility-scale fault-tolerant quantum computation^{141,39,158,183}. In particular, achieving such high-fidelity performance for nonlocal quantum circuits over large numbers of qubits could enable the implementation of efficient high-rate QEC codes^{55,453} as well as significant algorithmic speedups with transversal gate architectures^{70,466}.

Neutral atom processors have recently emerged as a leading quantum computing approach, combining nonlocal qubit connectivity⁴³, scalable control of large qubit numbers^{288,417}, continuous operation^{90,267,317}, and the ability to efficiently implement quantum error correction^{41,352,311,348}. In this approach, atomic qubits are entangled using controlled-phase gates based on Rydberg blockade^{224,364}. The fidelity of these entangling operations has rapidly advanced^{264,175,129,282,329,311,342}, with recent benchmarking demonstrations reaching $\sim 99.7\%$ fidelity⁴²⁵. Moreover, methods for loss detection and erasure conversion have been shown in a range of settings, including loss-

postselected fidelities reaching $\sim 99.8 - 99.85\%$ ^{387,274,268} and improved quantum error correction with erasure information^{464,349} including below-threshold performance⁴².

Here we demonstrate a new state-of-the-art in entangling CZ gate fidelity for neutral atom qubits, with a raw fidelity of $99.854(4)\%$ that improves to $99.941(3)\%$ upon postselection on neither atom being lost. Our approach combines high Rabi frequency gates with a smooth-amplitude pulse profile based on optimal control^{225,320,129} with fast calibration based on loss-resolved qubit readout⁴², and we demonstrate how this fidelity can be maintained for over 10 hours without recalibration. We utilize these gates to realize quantum circuits based on coherent atom motion, first benchmarking similarly high-fidelity CZ gate performance in the making and unmaking of cluster states. Then, we design and implement nonlocal circuits featuring longer-range atom transport, which exhibit super-ballistic scrambling of quantum information. We study the nature of these nonlocal circuits by sampling from the output distribution, and show consistency with numerical simulations of the chaotic dynamics. These advances open the door for efficient realization of quantum algorithms with nonlocal connectivity as well as deep-circuit fault-tolerant quantum computation with high-rate QEC codes.

8.2 NEUTRAL ATOM ENTANGLING GATES

Our approach makes use of a reconfigurable neutral atom processor architecture^{41,42}, with several key upgrades to achieve and explore low error regimes. Quantum information is encoded in long-lived hyperfine clock states of ^{87}Rb atoms, and entangling gates are performed by exciting the pair of atomic qubits to $n = 53$ Rydberg states¹²⁹. We use a gate profile based on optimal control^{225,320} with a smooth amplitude and phase profile, as shown in Fig. 8.1a¹²⁹, and additional amplitude corrections to compensate for pulse imperfections²⁸²; see Appendix F. This gate is optimized to suppress scattering from the intermediate excited state $|e\rangle$ by maximizing population in the dark state



Figure 8.1: CZ entangling gates in a neutral atom processor with fast cycle rates. (A) Rydberg gate realized through a smooth-amplitude gate pulse at roughly $\Omega = 2\pi \times 17$ MHz peak Rabi frequency, with a detuning profile $\delta(t)$. Two-photon excitation scheme consists of transitions at 420 nm and 1015 nm, exciting from qubit state $|1\rangle$ to the Rydberg state $|r\rangle = 53S_{1/2}, m_J = -1/2$ through the intermediate state manifold $|e\rangle = 6P_{3/2}$ with intermediate state detuning $\Delta = 2\pi \times 7.8$ GHz. (B) Images of the atoms separated into processor zones, with gate sites used for benchmarking and high-fidelity quantum circuits, and state-selective readout to perform loss detection. Fast calibration of CZ gates is enabled by qubit reuse and refilling lost atoms from a large atom reservoir. (C) Randomized benchmarking is performed over 40-80 repetitions, with a 20-30 Hz cycle rate. Example calibration of the CZ gate from starting parameters near the theoretical prediction, showing both raw and loss-postselected CZ error. Each point corresponds to the fitted maximum of a parabola scan over one of the gate parameters, with error bars given by the shaded region.

not containing $|e\rangle$ ¹²⁹. To suppress errors from Rydberg state lifetime, we utilize high Rabi frequencies reaching a peak of roughly $2\pi \times 17$ MHz. The required high laser intensities induce substantial lightshifts on the Rydberg transition and necessitate careful stabilization of Rydberg laser positions and powers to minimize the detrimental effect on ground-Rydberg coherence T_2^* ; see Appendix F. Additionally, the coupling to the other Rydberg level $|r'\rangle$ with $m_J = +1/2$ grows as Rabi frequency is increased, which we address by applying a stronger magnetic field (larger Zeeman splitting) to suppress off-resonant excitations.

Our quantum processor utilizes up to 442 atomic qubits arranged into separated zones (Fig. 8.1b), with 17 gate sites used for high-fidelity benchmarking and quantum circuits, and a large reservoir for refilling lost atoms. We employ state-selective readout using spin-to-position conversion⁴², which enables us to directly detect atom loss errors at the end of the circuit. By reusing qubits and refilling lost atoms from the reservoir, we are able to perform fast calibration of the CZ gate, for up to 80 repetition cycles within a single round of atom loading, at a cycle rate of 20-30 Hz (Fig. 8.1c). We find that calibrating the gate from scratch, which consists of scanning each gate parameter in




Figure 8.2: Stable low-error CZ gate performance. (A) Benchmarking methods based on repeated CZ gate applications: global echo randomized benchmarking (“echo RB”) ¹²⁹ and symmetric stabilizer benchmarking (“SSB”) ⁴²⁵. (B) Return probability as a function of number of CZ gates, fitted to an exponential decay, giving a raw fidelity of 99.854(4)% and a loss-postselected fidelity of 99.941(3)%. (C) Comparison of CZ gate fidelity measured with different benchmarking protocols: echo RB with and without an extra single-qubit π pulse before readout, resonant pushout of $|1\rangle$ versus state-selective readout (SSR), and SSB versus echo RB. (D) Stability of the CZ fidelity measured for over 10 hours with echo RB, with no recalibration of the gate. Each data point consists of a fit to the decay of return probabilities for 0, 8, 12, and 20 CZ gates. (E) Number of errors per repetition (raw and loss-postselected), where a repetition consists of benchmarking on two rows of 16 atoms brought into the entangling zone, for the data shown in D, consisting of 127,680 repetitions. Solid curves are sums of individual Poissonian fits for each number of CZ gates. (F) Sensitivity of the raw and loss-postselected CZ gate fidelity under perturbation in the 1015-nm Rydberg beam power, keeping the gate pulse parameters unchanged.

series, typically takes < 40 minutes, with calibrated gate parameters remaining stable such that subsequent recalibration is faster.

8.3 BENCHMARKING LOW-ERROR CZ GATES

To characterize the resulting CZ gate performance, we utilize randomized benchmarking (RB) based on repeated application of CZ gates interspersed with single-qubit rotations (Fig. 8.2a). First, we use a global echo RB sequence using a single-qubit π pulse in between pairs of CZ gates ¹²⁹, and measure a raw fidelity of 99.854(4)% fidelity and a fidelity of 99.941(3)% when postselecting on neither atom being lost at the end of the sequence (Fig. 8.2b). Our detection of these loss errors can

be used to improve QEC thresholds^{23,42} and quantum simulation results¹³⁰. In particular, we find that the CZ error due to atom loss is 0.087(5)%, comprising $\sim 60\%$ of the total error. We measure a loss probability of 0.054(2)% per atom per CZ gate, which exceeds half of this total loss error and indicates correlated loss of both atoms during the gates; this information can be used in QEC decoding³³¹. Additionally, we explore the effects of leakage to other m_F levels⁴²⁵ in the ground state $5S_{1/2}$, $F = 1, 2$ manifolds (from Rydberg T_1 and intermediate state scattering) through measurements of $F = 1$ populations at the end of the RB sequence, which increase as 0.008(1)% per atom per CZ gate (Appendix F).

To test the robustness of the gate performance, we measure the CZ fidelity across several different benchmarking and readout methods (Fig. 8.2c). First, we compare echo RB to the symmetric stabilizer benchmarking (SSB) protocol developed in Ref. ⁴²⁵, which is theoretically insensitive to single-qubit gate errors and requires calibration of the single-qubit phase from the CZ gate. We find that both the raw and loss-postselected fidelities are equivalent for the two RB methods, consistent with no measurable effect from single-qubit gate errors or miscalibration of the single-qubit phase of the gate. We additionally compare state-selective readout with destructive spin-to-loss readout based on resonant blowout of one of the qubit states, and again find comparable results for echo RB. Finally, we measure similar fidelities with and without an extra single-qubit π pulse before readout, which suggests that any residual effects of leakage on the RB results are relatively balanced between the $F = 1$ and $F = 2$ levels.

To test the stability of our gates, we track the fidelity over long durations (Fig. 8.2d). Remarkably, we observe stable performance at $\sim 99.85\%$ fidelity for over 10 hours without any recalibration of the gate. For this large dataset, consisting of roughly 2 million runs of the RB sequence averaged over the array, we further observe that both the raw and loss-postselected return probabilities follow Poisson distributions with no observed large-scale correlated error events spanning the system (Fig. 8.2e). We find that the residual drift in the gate fidelity typically originates from drifts




Figure 8.3: Error breakdown and path to further improvement. (A) Illustration of error sources for Rydberg entangling gates, including scattering from the intermediate state γ_e , Rydberg state lifetime γ_r , ground-Rydberg coherence T_2^* , and coupling to the other Rydberg state $|r'\rangle = 53S_{1/2}, m_J = +1/2$. (B) Numerical error budget showing contributions of individual errors, total theoretical errors, and experimental data. Error breakdowns are shown for the current configuration, as well as projected values attainable with technical improvements and higher laser intensities (Appendix F).

in the Rydberg laser powers and positions, such that peak performance can generally be recovered by small adjustments in beam powers and positions, with calibrated gate pulse parameters remaining unchanged. To understand the magnitude of the effect of drifts, in Fig. 8.2f we intentionally adjust the power of 1015-nm Rydberg beam away from optimal, and observe that the gate is sensitive to $\sim 0.1\%$ -scale power changes. We further find that the sensitivity of the CZ gate to perturbations in the Rydberg beam powers and positions is well-captured by numerical error simulations that include the effects on gate detuning and T_2^* . These observations suggest that neutral-atom devices with sufficiently stable Rydberg beams can achieve high-fidelity performance over long durations passively, and over even longer timescales with automated stabilization of beam positions and powers, for example via real-time steering of QEC performance⁴⁰³.

The remaining CZ gate errors can be largely explained by a numerical error model, with the dominant contributions being decay from the Rydberg state due to finite lifetime and from the intermediate excited state in the form of off-resonant scattering (Fig. 8.3). Using our numerical error model, we project that raw CZ gate fidelities at the level of 99.9-99.95% could be enabled by improved T_2^* , suppressed coupling to $|r'\rangle$, and higher laser intensities.




Figure 8.4: High-fidelity quantum circuits with atom motion. (A) Raw and loss-postselected CZ fidelity measured with echo RB, with one atom in a static trap generated by a spatial light modulator (SLM) and one atom in a moving trap generated by an acousto-optic deflector (AOD). (B) Protocol for making and unmaking one-dimensional cluster states to study CZ gate performance in a quantum circuit. (C) Decay of global return probability P_{return} as a function of the number of total CZ gates, fit to an exponential decay, for $N = 20$ qubits. (D) Stability of the fitted decay over time, without recalibration.

8.4 HIGH-FIDELITY CIRCUITS WITH ATOM MOTION

We next extend our benchmarking of CZ gates to quantum circuits based on coherent rearrangement of atomic qubits. Several additional challenges can arise when incorporating high-fidelity entangling gates into quantum circuits. In our randomized benchmarking circuits, we place the two atoms in two traps generated by the same spatial light modulator (SLM), which allows us to put the atoms uniformly close together ($\approx 1.7\mu\text{m}$ spacing). By contrast, in quantum circuits, we perform gates between atoms in stationary SLM traps and atoms in moving traps generated by acousto-optic deflectors (AODs)⁴¹. We find that to perform these gates with high fidelity, it is important for the spacing between traps to be uniform across the array, and for the focus of the two trap arrays to be in the same plane (Appendix F). When properly calibrated, we can realize similarly high CZ gate performance, as shown in Fig. 8.4a.

To understand the CZ performance in quantum circuits with atom motion, we perform a char-

acterization circuit corresponding to repeatedly making and disentangling 20-atom cluster states (Fig. 8.4b). We measure the probability that all atoms in the array return to the initial state $|0\rangle^{\otimes 20}$ after 0, 4, 8, 12, and 16 nearest-neighbor CZ gate layers, defined as the global return probability P_{return} . By fitting P_{return} as a function of the total number of CZ gates applied across the entire qubit array, we infer the corresponding CZ gate fidelity (Fig. 8.4c). We find a raw fidelity of 99.843(6)% and a loss post-selected value of 99.956(5)%, similar to the fidelity measured in gate benchmarking sequences. While not a rigorous benchmarking method, we observe in numerical simulations that the extracted fidelity from this characterization circuit approximately agrees with the CZ gate fidelity. To test the stability of the CZ gate performance within this circuit, we leave the experiment running for 15 hours, and find that the extracted CZ fidelity remains high without any recalibration (Fig. 8.4d).

8.5 CHAOTIC NONLOCAL CIRCUITS

Finally, we explore high-fidelity entanglement in more complex circuits with longer-range qubit connectivity. Owing to the high degree of connectivity, such circuits are known to be effective in scrambling quantum information, and can be difficult to simulate classically^{154,112}. In Fig. 8.5, as an example, we design and study a class of nonlocal circuits with increasingly long-range gates as a function of depth, interspersed with $\pi/4$ single-qubit gates between layers. This circuit structure leads to super-ballistic scrambling of quantum information, in which entanglement entropy saturates at the Page entropy, the near-maximal value expected for a random pure state, in a depth which scales as $\sqrt{N}$ for system size N (Fig. 8.5b). Our circuits can also be interpreted as nonlocal kicked Ising models, where the corresponding nearest-neighbor circuit does not scramble (see Appendix F). The output random states are related to the physics of black holes¹⁹⁸, and have been used for benchmarking quantum devices^{101,91} and quantum advantage experiments¹⁷.




Figure 8.5: High-fidelity scrambling circuits with nonlocal connectivity. (A) Nonlocal quantum circuit implementation. CZ gates of increasing nonlocal connectivity (with labels given below the circuit) are interspersed with single-qubit gates, $R = X(\pi/4)$. The resulting connectivity for the 20-qubit circuit is illustrated on the right. (B) Theoretical values for the half-system-size von Neumann entanglement entropy and XEB_{ideal} as a function of the degree of nonlocal CZ gates added to the circuit ("nonlocal connectivity"). Dashed lines indicate the Page entropy value and $XEB_{\text{ideal}} = 1$. (C) Loss per qubit per CZ gate compared for different circuits in this work: randomized benchmarking (RB), the cluster state circuit in Fig. 8.4b, and the nonlocal quantum circuit in Fig. 8.5a. (D) Sampled outcome distribution for the 10-qubit nonlocal quantum circuit, run with the depth shown in A, comparing the experimental data to the ideal theory expectation and a Porter-Thomas distribution. (E) Sampled outcome distribution for the 20-qubit nonlocal circuit, with the outcomes binned in 100 groups by their probability. Comparison to numerical simulations with noise is shown. (F) Experimentally measured XEB for the 20-qubit circuit, for different degrees of nonlocal connectivity, compared to numerical simulations with noise. Solid bars are raw data, and lighter bars correspond to postselection on no atom loss.

We first assess fidelity of the CZ gate in these nonlocal quantum circuits applied on 20 qubits through measurement of atom loss, which we find to be a useful proxy for the in-situ performance of the gate in circuits. We measure the loss per qubit per CZ gate to be consistent across the randomized benchmarking, cluster state, and nonlocal circuits throughout this work (Fig. 8.5c). Additionally, we observe that the loss per shot for both circuits follows a Poisson distribution, with no evidence of large-scale loss events.

Next, we sample from the output of this circuit, by measuring all qubits in the Pauli Z basis and plotting the probability of each outcome ordered by theoretical expectation. For random states, the ideal probability distribution $p(x)$ is an exponentially distributed Porter-Thomas distribution $\text{PT}(x)$ ³³⁵. Such probability distributions can be studied using cross-entropy benchmarking (XEB), where $\text{XEB} = 2^n \sum_x q(x)p(x) - 1$ for outcomes x and system size n . The ideal value in the absence of noise is given by the second moment of the probability distribution, $\text{XEB}_{\text{ideal}} = 2^n \sum_x p^2(x) - 1$, often called the collision probability, for which the Porter-Thomas distribution $\text{PT}(x)$ has $\text{XEB}_{\text{ideal}} = 1$. We find that the theoretical output distributions of our circuits converge to $\text{PT}(x)$ for sufficient depth (Fig. 8.5b).

Experimentally, we start by studying the distribution in a 10-qubit circuit, where we can measure the probability of all 2^N outcomes. We find that the resulting distribution is similar to $\text{PT}(x)$ for outcomes x (Fig. 8.5d). Next, we extend to sampling from the 20-qubit circuit by binning the outcomes into groups (Fig. 8.5e). We find that our results are well described by numerical simulations with noise, with decoherence making the distribution closer to the uniform distribution. To characterize the 2^{20} sampled outcomes more quantitatively, in Fig. 8.5f we compare XEB for the 20-qubit circuit sampling distribution to numerical simulations, and find good agreement for the same error model across three different circuit depths; see Appendix F. While we do not directly probe the CZ gate fidelity as in Fig. 8.4, the combination of low atom loss per gate and agreement with numerical simulations suggests high-fidelity performance in these complex nonlocal circuits.

8.6 DISCUSSION AND OUTLOOK

Our results pave the way for practical fault-tolerant quantum computation with atomic processors. Corresponding to more than $3 \times$ improvement from our previous CZ gates used in below-threshold performance of a surface code logical qubit⁴², these higher fidelities, in combination with improving single-qubit operations, should enable $> 10 \times$ below-threshold QEC performance, while leveraging the ability to detect atom loss errors²³. These improved fidelities were achieved with $\approx 1.7 \times$ higher Rabi frequency, requiring suppression of coupling to $|r'\rangle$, stronger blockade with closer spacing, improved Rydberg coherence (T_2^*), as well as fast calibration of the gate pulse; see Appendix F. We find that the remaining CZ gate errors are dominated by Rydberg state lifetime and intermediate state scattering. Through technical improvements to further suppress coupling to $|r'\rangle$ and improve the Rydberg T_2^* , we expect that reaching 99.9-99.95% raw CZ fidelity will be achievable by scaling up the available laser power.

Future work will focus on implementing these high-fidelity entangling gates across larger atomic systems with up to 10,000 or more qubits, for instance leveraging efficient beam shaping¹²³, fast beam scanning^{349,437,68}, higher-power and pulsed lasers⁸⁷, and different Rydberg excitation schemes^{282,425}. Furthermore, our techniques can be extended to other gate schemes, including entangling gates at constant atomic velocity²⁶⁸ and multi-qubit gates such as CCZ^{129,73}.

The high-fidelity quantum circuits demonstrated in this work form the basis for wide-ranging explorations of highly-entangled quantum systems and fault-tolerant architectures. Extensions of our nonlocal circuits could be used to simulate nonlocal quantum models, including efficient encoding of fermions with qubits²⁹² and models for gravity³⁶⁷. Additionally, such circuits can be used to probe complex properties of quantum many-body systems, providing optimal system-size-independent sample complexity for measuring linear observables²⁰⁹. Most significantly, implementing complex nonlocal circuits with high fidelity should enable the realization of high-rate quantum

error correcting codes to achieve low error rates on large numbers of logical qubits^{55,453}.

9

Conclusion

In this thesis, we have developed a platform for high-fidelity digital quantum computation with atom arrays, and conducted experiments in quantum simulation and quantum error correction. These results highlight near-term scientific directions with atom arrays, as well as pave the way towards future large-scale error-corrected devices. Below I outline a few research directions that I think will be particularly exciting in the next few years with atom arrays.

9.1 CHARTING THE PATH TO LARGE-SCALE ATOMIC PROCESSORS

As a field, we are progressing towards building continuously-reloaded large-scale devices with $>10,000$ atoms, $>99.9\%$ entangling gate fidelity, 99.99% fidelity single-qubit gates, and flexible ways to read-out and coherently rearrange the atoms. The next frontier will be in figuring out how to combine these ingredients into a single system where everything works seamlessly in concert. Such devices will be incredibly powerful in enabling a wide range of scientific directions.

To scale beyond this level, several approaches can be considered. Trap power is one of the main limitations to scalability, with each trap typically requiring one milliwatt of power. However, once cooled to very low temperatures, atoms can be confined in orders-of-magnitude weaker traps⁶⁵. One can therefore develop methods to operate an atomic processor with most of the optical traps being very low power, only using deeper traps temporarily for loading and imaging the atoms. Such methods of iteratively assembling atomic systems have begun to be explored recently in both optical lattices¹⁸⁵ and optical tweezer arrays^{90,267}, but using this approach to scale to much larger atom arrays is an outstanding challenge. Several technical challenges will need to be overcome to fully take advantage of this power efficiency, for example suppressing atom heating during the assembly. Additionally, larger microscope objectives and methods for transporting atoms over larger distances will need to be investigated. We will also need to efficiently control and entangle these large systems, which can be accomplished by movable addressing beams or pulsed laser systems for higher intensities²⁸⁵. Scaling to larger systems will also likely require integrating new optical technologies, for example metasurfaces to generate programmable tweezer arrays²⁰³ or optical cavity arrays to achieve several-orders-of-magnitude enhancement in power³⁹⁰.

9.2 TOWARDS USEFUL DIGITAL QUANTUM SIMULATIONS

We ultimately would like to use these digital quantum processors to simulate quantum systems and make scientific discoveries. Near-term digital quantum simulations and implementations of algorithms can provide us with a sense of how quantum simulations might look in the future, once we have built large-scale devices capable of running error-corrected algorithms with very low error rates. We can use these near-term digital simulators to begin understanding how to structure quantum simulations with digital devices as well as questions such as how to verify the output of quantum simulations, what observables to study, and whether there can be ways to significantly speed up algorithms.

9.3 FRONTIER IN ERROR-CORRECTED QUANTUM PROCESSING

Finally, we will need to make progress in testing error-corrected quantum processing and algorithms to guide the design of these atomic processors. A few near-term experiments in error correction could serve as the next milestones to work toward. First, our repeated rounds of quantum error correction on a surface code logical qubit⁴² can be combined with continuous reloading methods to further extend the depth and store a qubit for long durations. For example, using 1000 atoms to create a surface code logical qubit with previously demonstrated fidelities, we theoretically would be able to store a qubit for over a day. Practically, such an experiment would enable probing noise floors at very low-error $< 10^{-7}$ rates. Second, different error-correcting codes should be explored experimentally, in particular to efficiently create large numbers of logical qubits using quantum LDPC codes⁵⁵.



Supplementary information about Kitaev model experiments

A.1 EXPERIMENTAL LAYOUT

In our experimental layout (Extended Data Fig. A.1a), the hexagonal plaquettes of the honeycomb model are embedded in a rectangular atom array with 4 rows. In these experiments, we use three

separate zones: entangling, storage, and readout⁴¹. Atoms are first sorted in the storage zone, then transported into the entangling zone before the start of the experiment. We employ an upgraded sorting algorithm, which, compared to our previous algorithm¹²³, has an additional final step to fill individual defect sites from a small atom reservoir (more details in Ref. ⁴²). For the state preparation circuit, we perform mid-circuit measurement^{113,402,176,277,316,282} of the ancilla qubits, by bringing them to a spatially separated readout zone far away ($\sim 150\mu\text{m}$) from the entangling zone and imaging them with a single, local 780-nm imaging beam⁴¹. The fidelities of individual components such as single-qubit gates and two-qubit CZ gates in this work are roughly similar to our other works^{129,41}, except for the local imaging fidelity, which was lower during data taking due to degrading trap laser power ($\sim 96\text{-}97\%$ compared to 99.8% in Ref. ⁴¹).

A.2 MEASUREMENT-BASED PREPARATION OF TOPOLOGICAL STATES

We begin our experiments by efficiently preparing a long-range entangled, topological state on a cylinder, which forms the basis for all subsequent explorations. First, we prepare a ZXXZ surface code state^{245,51} on one of the data qubit sublattices (black sites), using entanglement operations with ancilla qubits to project the data-qubit state into an eigenstate of the weight-4 ZXXZ operators. Concretely, we put the ancilla qubits in the $|0\rangle_a + |1\rangle_a$ state and perform a sequence of CZ gates, with data qubits in the appropriate bases, which realizes the $|0\rangle_a + Z_1 X_2 X_3 Z_4 |1\rangle_a$ state, where the numbers label the data qubits along an example ZXXZ operator. The subsequent measurement of the ancilla qubits in the X basis fixes the parity of the ZXXZ operators on the data qubits. The measurement outcomes and, thus, the projected operator parities are random (up to certain constraints on their products). A feedforward step, acting on the same sublattice, ensures that all parities end up being +1, which corresponds to the subspace with the ground state of the Kitaev model²⁷¹. For the fermion dynamics, this ensures that no magnetic fluxes are present.




Figure A.1: Honeycomb lattice layout and experiment sequence. **a**, Mapping from honeycomb geometry into experimental array geometry used, indicating link orientations, an example column, and an example plaquette orientation. We use this mapping in order to shrink the number of rows needed, as well as put the atoms into a rectangular grid which is convenient for atom motion and local Raman operations. **b**, Loss radius definition for error detection based on atom loss. The central (white) atom is the reference point, and for loss radius of 0, we postselect on this atom being present at the end of the circuit. For a larger loss radius, we postselect on atoms within a certain local region being present. For string observables, we perform this procedure for all atoms within the string. **c**, High-level overview of the experimental sequence used in these experiments, including feedforward topological state preparation, Floquet evolution, and measurement steps. The Fermi-Hubbard quantum simulations have an additional part of the Floquet circuits for engineering the onsite density-density interactions.




Figure A.2: Measurement-based topological state preparation. **a**, Experimental protocol used for measurement-based preparation of the topological state. First, a ZXXZ toric code state is prepared on one data qubit sublattice (black circles) using ancilla qubits (blue circles). After mid-circuit measurement and feedforward, this prepares the ZXXZ toric code state. The weight-four ZXXZ operators are then extended to weight-6 operators with parallel controlled-Y (CY) operations with a fresh data qubit sublattice originally in $|0\rangle$ (green circles). In the case that the ZZ link (where the CY gates were applied) are +1, this weight-6 operator is equal to the plaquette operators W_p . This sequence indeed ensures that these ZZ operators are all +1, owing to the propagation of the single-qubit Z stabilizer on the green qubit sublattice through the CY operation, which results in the weight-2 ZZ stabilizers. **b**, Implementation of the encoding steps in parallel across the full array. The top row of ancillas also performs gates with the bottom row of data qubits.

The order of entangling bases we use is Z, X, X, Z and since all qubits are initially in the $|0\rangle$ state, we can omit the first Z measurement; the resulting circuit for this part is depth-3. These weight-4 ZXXZ operators are grown to weight-6 ZYXZXY operators by performing parallel controlled-Y (CY) operations between the entangled sublattice (black) and the one remaining in the $|0\rangle$ state (green). The resulting weight-6 operators are equivalent to the plaquettes $W_p = ZYXZXY$ up to the ZZ-link operators. However, the structure of the circuit ensures that the ZZ-link terms are +1 and, thus, can be freely multiplied into other operators. As a result, this depth-4 circuit efficiently prepares the fermionic vacuum state, or equivalently the A_Z^I ground state of the Kitaev honeycomb model²⁴³. Such measurement-based methods can be used for preparing a wide range of topological states in finite depth^{281,416,368}.

Initially, the ancilla qubits are located in moveable AOD traps and the data qubits are in stationary SLM traps. The ancilla qubits are reconfigured and then entangled with data qubits in the correct basis. The first entanglement step includes the periodic boundary direction and has an ad-

ditional component where the top row of ancilla qubits is entangled with the bottom row of data qubits. We perform this step first so that all other qubits can be in state $|0\rangle$, avoiding additional errors due to Rydberg excitations.

To realize mid-circuit measurement, the ancilla qubits are transported to the readout zone and locally imaged⁴¹. The measurement outcomes are then used to perform real-time decoding and feedforward. The feedforward correction is applied before the parallel CY operations.

For the feedforward corrections, we use a field-programmable gate array, FPGA (Xilinx ZCU111), to gate on and off 32 local single-qubit Z gates applied across one sublattice of the array. An example of mid-circuit decoding and feedforward is shown in Fig. A.3a. The decoding algorithm uses single-site Z gates, which flip the two vertically adjacent plaquettes, to pair the -1 results in each column by pushing the -1 values until another one is encountered. The initial site for each column and the direction of the pushing procedure is randomized to avoid biasing any given row of plaquettes. The decoder can additionally be modified to prepare an initial state with a deterministic pattern of ± 1 plaquette values, as long as the configuration does not violate parity constraints. We use this flexibility to initialize different states with fewer local pulses than naively necessary: if a single-qubit gate pattern used to initialize the fermion sites flips an even number of plaquettes per column, we can pre-compensate for those flips in the decoding step.

In Extended Data Fig. A.3d-e, we numerically explore various constant-depth circuits realizing the long-range entangled state of interest. The method we employ in this work significantly outperforms other approaches, including the direct measurement of the hexagonal plaquettes.

A.3 DECODING ERROR DETECTION

The product of ancilla measurement outcomes in every column must be even since the plaquettes in any column multiply to strings enclosing the cylinder that are composed of Z operators only (Ex-

tended Data Fig. A.3b), which are fixed to be +1 due to our initial product state. This constraint can be used for error detection. In particular, whenever there are an odd number of -1 ancilla measurement results in a given column, we know that an error must have occurred during the state preparation circuit. Utilizing this decoding postselection method, we observe that all lengths of loops improve in value, as shown in Extended Data Fig. A.3c.

A.4 POSTSELECTION BASED ON ATOM LOSS

We utilize a state-selective qubit readout, described in our work in Ref.⁴², which distinguishes between the $\{|0\rangle, |1\rangle\}$ states and atom loss. In this method, a state-dependent circularly-polarized 1D optical lattice is used to pin one of the two qubit states^{287,446}, and an AOD tweezer moves the other qubit state a few microns away⁴², converting the atomic state to position which is then imaged via polarization gradient cooling (PGC). The information about lost atoms does not allow us to correct the state but we can use it for error detection and postselection^{381,282}. In particular, for each observable, we only use the experimental shots where all qubits constituting that observable are present. Moreover, we can employ a sliding-scale postselection scheme where we additionally postselect on qubits being present within a given distance on the lattice (Extended Data Fig. A.1b), which can mitigate the effects of error spreading throughout the circuit. To quantify this, we introduce a *loss radius*, which describes the distance on the honeycomb lattice within which the atoms must be present.

A.5 FLOQUET EVOLUTION CIRCUITS

After the measurement-based state preparation, we keep one data qubit sublattice in moveable AOD traps, and at each time step, transport them next to the atoms in the even data qubit sublattice to perform the tunable entangling gates. Global single-qubit rotations between entangling gates




Figure A.3: Midcircuit decoding and error detection method. **a**, Example of decoding and feedforward pairing based on ancilla results. White (blue) plaquettes have the value of -1 (+1). After the pairing procedure, all plaquettes are prepared with +1 value in the absence of errors. Note that the local Z gates flip plaquettes within the column to perform the 1D matching, but do not flip plaquettes in adjacent columns. **b**, Illustration of decoding postselection method based on parity of ancilla measurement results within a column, owing to the cylindrical geometry used in this work. The product of plaquettes within a column is equal to two Z loops enclosing the cylinder, which must be +1 in the absence of errors. **c**, Larger loops measured for the initial state, as a function of decoding postselection. We attribute the fact that the loop enclosing four hexagons is higher parity than the loop enclosing three hexagons due to certain differences in the measurement, for example the loops enclosing the cylinder are measured in the Z basis which doesn't require local Raman gates to rotate individual qubits. **d**, Comparison of different state preparation methods. The method we use in this work is based on the ZXXZ toric code prepared with 32 ancilla qubits. This method can also be adapted to require only 16 ancilla qubits. A third method ("Hexagons") consists of directly measuring the weight-6 plaquette operators with 32 ancilla qubits, requiring a slightly deeper state preparation circuit. **e**, Numerical simulation of plaquette parity for these different state preparation methods, for the phenomenological noise model which matches our experimental data (depolarizing noise of strength 0.027 on all qubits after every two-qubit gate layer). The method that we use in this work (ZXXZ) performs the best in numerics, both for no postselection as well as with decoding postselection.

are used to change the basis between X, Y, and Z. We perform dynamical decoupling throughout, including to cancel single-qubit dephasing from moves^{43,41}. During the Floquet evolution, we implement periodic cylindrical boundary conditions. In particular, due to our array geometry, only the XX links couple the top row to the bottom row. For these links, we apply it in two steps, first moving the data qubits in AOD traps up one lattice site, and then moving the top row to perform gates with the bottom row. During this second step, the three other rows are moved out of the Rydberg beam to avoid extra gate errors. Conveniently, the structure for all the Floquet circuits in this work is the same, and so we only need to change the CPHASE gates between the different Floquet circuits. This enables the implementation of a wide variety of fermionic evolution with minimal experimental changes.

Since the Floquet circuit commutes with closed loop operators, the topological order is preserved as we evolve into phase B, up to a small decay from gate errors. In principle, the final state could be more fully characterized by measuring all closed loop operators, including the additional larger loops measured in the initial state.

For the low-energy states of the Kitaev Hamiltonian, we optimize the state preparation circuits to maximize the overlap between bulk Majorana correlations (Extended Data Fig. A.4a-b) of the prepared state and that of the Floquet ground state with Floquet time $\tau=1/4J$. We confirm that the resulting state has the correct Chern number and a similar energy as a state optimized purely based on energetic considerations. We note that the circuit depth of 6 is sufficient for this system size, but will necessarily increase for larger system sizes due to the gap closing²³¹.

A.6 ENCODING FERMIONS IN QUBITS

The Hilbert space of N qubits and N fermions has the same dimension but due to non-local properties of fermions (anticommutation relations), mapping between them can be complicated. A direct

translation on the operator level is given by the Jordan-Wigner (JW) transformation²²⁹, where for a particular site ordering, $[1, \dots, N]$, we can identify complex-fermion creation and annihilation operators

$$a_j = \frac{1}{2}(X_j + iY_j) \prod_{k=1}^{j-1} Z_k,$$

$$a_j^\dagger = \frac{1}{2}(X_j - iY_j) \prod_{k=1}^{j-1} Z_k,$$

and the corresponding Majorana operators,

$$c_j = a_j^\dagger + a_j = X_j \prod_{k=1}^{j-1} Z_k, \quad \bar{c}_j = i(a_j^\dagger - a_j) = Y_j \prod_{k=1}^{j-1} Z_k,$$

which can be directly checked to satisfy the canonical anticommutation relations, $\{a_i, a_j^\dagger\} = \delta_{ij}$ and $\{c_i, c_j\} = 2\delta_{ij}$. Beyond one dimension, this approach leads to macroscopic operator weight on the qubit side, even for simple local fermion operations such as nearest-neighbor hopping,

$$c_i c_j = X_i X_j \prod_{k=i}^{j-1} Z_k, \tag{A.1}$$

for $i < j$.

The idea of local fermion-to-qubit encodings⁴³¹ relies on introducing a long-range entangled “background state” $|\emptyset\rangle$ that is stabilized by the non-local operator strings, effectively canceling them out. For example, if $\prod_k Z_k$ in Eq. (A.1) is a stabilizer of $|\emptyset\rangle$, i.e. $(\prod_k Z_k) |\emptyset\rangle = |\emptyset\rangle$, then the hopping term effectively becomes a simple weight-2 operator, $c_i c_j \sim X_i X_j$. However, introducing these additional constraints on the state reduces the available Hilbert space for fermion degrees of freedom; thus, the system needs to be expanded by introducing additional data qubits. In the honeycomb encoding studied here, this manifests as the long-range entangled ZXXZ state on one

sublattice being coupled to matter degrees of freedom through the parallel controlled-Y operation.

In this setting, we have twice as many qubit degrees of freedom compared to fermionic ones, which grants enough space to enforce the stabilizer constraints.

We now show that the two-qubit interactions correspond to hopping terms of Majorana fermions, c_j , localized at each site j . We begin by constructing JW operators that are similar to those in Eq. (A.1) but modified to better suit our honeycomb lattice. We choose a site ordering starting in the top-left corner of the lattice and creating a continuous path $\mathcal{L}$ through the entire system. We use a JW operator²³⁷,

$$c_j = Z_1 \prod_{l \in \mathcal{L}_j} \sigma_{l(1)}^{(l)} \sigma_{l(2)}^{(l)},$$

where $\mathcal{L}_j$ is the sequence of links along path $\mathcal{L}$ ending at site j , $l(1/2)$ denotes the vertices of link l , and $\sigma^{(l)} \in \{X, Y, Z\}$ is the Pauli operator along l . Since the consecutive link operators always anticommute, the c_j operators defined this way satisfy the correct anticommutation relations and, as products of Paulis, are Hermitian and square to the identity operator. Now, if the hopping terms is between two consecutive sites on path $\mathcal{L}$, then we trivially recover $\sigma_{l(1)}^{(l)} \sigma_{l(2)}^{(l)}$ as the rest of the string squares to identity. Finally, if the hopping term is between two nearest-neighbor sites that are not adjacent on $\mathcal{L}$, we end up with

$$c_i c_j = \prod_{l \in \mathcal{L}_j \setminus \mathcal{L}_i} \sigma_{l(1)}^{(l)} \sigma_{l(2)}^{(l)},$$

where $\mathcal{L}_j \setminus \mathcal{L}_i$ is the ordered set of links between sites i and j along $\mathcal{L}$. However, by applying the link (i, j) , that path can be completed to a closed loop (on a cylinder it also needs to be multiplied by a conserved non-trivial loop around the cylinder) which, in turn, is exactly the product of all enclosed plaquette operators. If the enclosed plaquettes are +1, the hopping term is effectively $\sigma_i^{(l)} \sigma_j^{(l)}$, where $l = (i, j)$. Thus, if all plaquettes are projected to be +1, such nearest-neighbor terms become the link

operators and more general Majorana correlations are mapped to a Pauli string constructed from products of link operators. The plaquettes with -1 values act as $\mathbb{Z}_2$ magnetic field fluxes, since they result in a π -phase for fermions hopping around them.

Complex fermions can always be formed by arbitrary pairing of Majoranas and here we choose to combine them along the ZZ links. The hopping operator for such complex fermions can be realized through a linear combination of length-2 and length-4 Pauli strings, which symmetrically couple the different Majorana constituents. The operators of this form can be realized through Floquet engineering. The two-qubit Pauli operators corresponding to the nearest-neighbor Majorana hopping, as derived here, constitute the exact interactions K_{ij}^X , K_{ij}^Y , and K_{ij}^Z , implemented in this work. Moreover, the Hamiltonian given by nearest-neighbor Majorana hopping terms results in the original Kitaev honeycomb model²⁴³.

A.7 FREE-FERMION STATES AND HAMILTONIANS

The Kitaev honeycomb model has an exact solution in terms of free fermions²⁴³, which enables efficient numerical simulation and benchmarking of most circuits in this work. Here we briefly summarize the main properties of such free-fermion states and Hamiltonians, focusing on Majorana operators, c_i , satisfying the canonical anticommutation relations $\{c_i, c_j\} = 2\delta_{ij}$.

Free-fermion states are captured by the two-point correlation matrix Γ ,

$$\Gamma_{ij} = \frac{i}{2} \langle [c_i, c_j] \rangle, \quad (\text{A.2})$$

with all higher-order terms decomposing into products of two-point functions via Wick's formula⁴⁴⁰. A general quadratic Majorana Hamiltonian,

$$H = \frac{i}{4} \sum c_i A_{ij} c_j, \quad (\text{A.3})$$

is defined through a real, skew-symmetric matrix $\mathbf{A}^\top = -\mathbf{A}$, where matrix elements A_{ij} encode Majorana hopping between sites i and j . The correlations of the ground state are related to those of the Hamiltonian²⁴³, and a unitary evolution of an arbitrary free-fermion state Γ is given by,

$$\Gamma(t) = \mathbf{U}^\dagger(t)\Gamma(0)\mathbf{U}(t), \quad (\text{A.4})$$

where $\mathbf{U}(t) = \exp(-\mathbf{A}t)$ is a matrix describing the time-evolution of the two-point correlation matrix²⁵¹.

We focus on a translationally invariant system, which can be described by the unit-cell position $\mathbf{R}$, and a label for the site within that unit cell, λ . For the honeycomb lattice, the unit cell has two sites, $\lambda = \{e, o\}$, corresponding to the even and odd sublattices, respectively, and Eq. (A.3) may be re-written as,

$$H = \frac{i}{4} \sum_{\mathbf{R}, \mathbf{r}} \begin{bmatrix} c_{\mathbf{R}}^{(e)} & c_{\mathbf{R}}^{(o)} \end{bmatrix} \begin{bmatrix} A_{\mathbf{r}}^{(e,e)} & A_{\mathbf{r}}^{(e,o)} \\ A_{\mathbf{r}}^{(o,e)} & A_{\mathbf{r}}^{(o,o)} \end{bmatrix} \begin{bmatrix} c_{\mathbf{R}+\mathbf{r}}^{(e)} \\ c_{\mathbf{R}+\mathbf{r}}^{(o)} \end{bmatrix}, \quad (\text{A.5})$$

where $\mathbf{r}$ is the relative position between the two relevant unit cells. In terms of momentum modes, $c_{\mathbf{k}} \propto \sum_{\mathbf{x}} e^{-i\mathbf{k} \cdot \mathbf{x}} c_{\mathbf{x}}$, where $\mathbf{x}$ is the position operator, the Hamiltonian takes the form $H = \sum_{\mathbf{k}} H_{\mathbf{k}}$,

$$H_{\mathbf{k}} = \frac{i}{4} \begin{bmatrix} c_{-\mathbf{k}}^{(e)} & c_{-\mathbf{k}}^{(o)} \end{bmatrix} \begin{bmatrix} \Delta_{\mathbf{k}} & \xi_{\mathbf{k}} \\ -\xi_{\mathbf{k}}^* & -\Delta_{\mathbf{k}} \end{bmatrix} \begin{bmatrix} c_{\mathbf{k}}^{(e)} \\ c_{\mathbf{k}}^{(o)} \end{bmatrix}, \quad (\text{A.6})$$

where Δ and ξ functions are the Fourier transforms of $A_{\mathbf{r}}^{(e,e)}$ and $A_{\mathbf{r}}^{(e,o)}$, respectively. The Hamiltonian H contains two energy bands with dispersions $\pm \varepsilon_{\mathbf{k}} \propto \sqrt{\xi_{\mathbf{k}}^2 + \Delta_{\mathbf{k}}^2}$. In the original Kitaev model, the gap closes in phase B²⁴³, $\Delta_{\mathbf{k}} = 0$, but our effective Floquet Hamiltonians have a spectral gap due to natural breaking of the time-reversal symmetry²³¹. Knowing the numerical values of the $\Delta_{\mathbf{k}}, \xi_{\mathbf{k}}$ functions on a grid of points in the Brillouin zone enables evaluation of various band properties. In particular, evaluation of the Chern number requires only a few points in momentum space¹⁴⁷.




Figure A.4: Chern number analysis. **a**, The central 12 columns are chosen as the bulk region of the system. **b**, The measured open Pauli strings. **c**, Majorana correlations in the unit-cell basis, obtained from the open strings. The correlation matrix is truncated at distance ± 1 (length-6 strings) to avoid system boundaries when starting from a reference site in the bulk. The (o, e) , (o, o) matrices are evaluated similarly, and related to (e, o) , (e, e) , respectively, through symmetry relations. **d**, The total variation distance of the measured string distribution (in phase B) from the ideal theory values with increasing error detection. The value is normalized by the point with least postselection and improves by $\sim 10\%$ for optimal error detection. **e**, Discretized Berry curvature F_k whose sum over the 1st Brillouin zone is the Chern number. The red arrows represent the discretized phase potential obtained from normalized eigenstate overlaps, with the arrow vector given by $\arg \tilde{U}_k = (\arg U_k^x, \arg U_k^y)$. **f**, The Chern number for different values of loss radius (LR) postselection, evaluated through bootstrapping with 300 trials as a function of a single batch size. The samples are drawn uniformly from the data set with replacement, and the samples from strings identified with each other are afterwards further averaged. **g**, Change in the Chern number between postselection at loss radius 3 and 0, as a function of batch size. **h**, Phase diagram of the Chern number as a function of per-layer error and initial plaquette parity values, obtained through noisy numerical simulations. The stars denote points corresponding to the phenomenological noise model inferred from experimental measurements (red with maximal postselection).

A.8 CHERN NUMBER

The gapped non-Abelian phase B of the Kitaev honeycomb model is characterized by a non-zero Chern number^{35,86} of the excitation band. An odd Chern number guarantees that a magnetic flux is accompanied by an unpaired Majorana zero-mode with non-Abelian (Ising anyon) statistics^{243,421,10}. We prepare a vortex-free ground state whose free-fermion parent Hamiltonian has a band with an odd Chern number. In principle, the Majorana zero modes could be prepared and probed directly, but the required circuit depths are larger than those used in this work²³¹.

The Chern number is a topological invariant of the energy band, which characterizes the geometry of the single-particle eigenstates of H_k ^{35,99}. Those eigenstates, satisfying $H_k |n_k\rangle = E_{n,k} |n_k\rangle$ for

energy E , are defined up to an overall choice of phase gauge,

$$|n_{\mathbf{k}}\rangle \rightarrow e^{-i\varphi_{\mathbf{k}}} |n_{\mathbf{k}}\rangle, \quad (\text{A.7})$$

where $\varphi_{\mathbf{k}}$ is the gauge parameter. To probe the local geometry of these eigenstates, as the momentum $\mathbf{k}$ is varied, we define the Berry potential (connection),

$$A_{n,\mathbf{k}} = \langle n_{\mathbf{k}} | \nabla_{\mathbf{k}} | n_{\mathbf{k}} \rangle, \quad (\text{A.8})$$

where $\nabla_{\mathbf{k}} = (\partial_{k_x}, \partial_{k_y})$ is the gradient in momentum space. The Berry potential is not gauge-invariant, since it transforms as $A_{n,\mathbf{k}} \rightarrow A_{n,\mathbf{k}} - i \nabla_{\mathbf{k}} \varphi_{\mathbf{k}}$ under Eq. (A.7), but the Berry curvature³⁵,

$$F_{n,\mathbf{k}} = \nabla_{\mathbf{k}} \times A_{n,\mathbf{k}}, \quad (\text{A.9})$$

is invariant under Eq. (A.7) because $\nabla \times (\nabla \varphi) = 0$.

The Chern number of the n th band is the integral of the Berry curvature over the 1st Brillouin zone,

$$C_n = \int_{\text{1st B.Z.}} d\mathbf{k} F_{n,\mathbf{k}} \quad (\text{A.10})$$

and is guaranteed to be an integer³⁶¹. Throughout this work we focus on the lowest energy band and omit the subscript n . For a given Bloch Hamiltonian in momentum space, $H_{\mathbf{k}}$, the Chern number can be evaluated either by direct numerical integration of Eq. (A.10) or with a specialized numerical algorithm¹⁴⁷. Therefore, the task at hand is reduced to learning the momentum-space parent Hamiltonians of the states prepared in our experiments^{341,210,318}.

The free-fermion states are defined by their two-point correlation functions and, similarly, the free-fermion ground states are related to the single-particle eigenstates of the parent Hamilto-

nian²⁴³. We measure the open Pauli strings corresponding to the two-point Majorana correlations, and average their values over the bulk region (Extended Data Fig. A.4a). We include all strings that span the bulk of the system, as depicted in Extended Data Fig. A.4b, and effectively recover all the $A_{\mathbf{r}}$ matrix elements of the parent Hamiltonian (up to the norm) with $r_x, r_y \in [-1, 0, 1]$ due to the finite-size restrictions (Extended Data Fig. A.4c). We then Fourier-transform these correlations onto a regular 5×5 grid in momentum space, resulting in an estimate of the parent Hamiltonian $H = \sum_{\mathbf{k}} H_{\mathbf{k}}$ up to an energy scale $\varepsilon_{\mathbf{k}}$.

Finally, we apply the algorithm of Ref.¹⁴⁷ to evaluate the Chern number. We diagonalize the $H_{\mathbf{k}}$ Hamiltonians at each $\mathbf{k}$ and calculate the phases of eigenstate overlaps between neighboring momentum points and collect them in a tensor $U_{\mathbf{k}}^{\mu}$, where $\mu \in \{\hat{x}, \hat{y}\}$ denotes the direction in momentum space, which depends only on the eigenstates of $H_{\mathbf{k}}$ and not $\varepsilon_{\mathbf{k}}$. Then, we calculate the discretized Berry curvature in Eq. (A.9) (Extended Data Fig. A.4d) and sum its matrix elements to obtain the Chern number, which is guaranteed to be an integer¹⁴⁷. An interesting future effort would be to compare this approach to other methods for extracting the Chern number, such as the real-space formula of Ref.²⁴³.

The Chern number is obtained from the learned Hamiltonian, and cannot be evaluated on individual snapshot data. When applying this procedure with the mean values plotted in the string distributions, we obtain $C=0$ for the Abelian phase and $C=1$ for phase B. To study the robustness of this result and the effect of postselection, we bootstrap the data¹⁰⁷ by evaluating the Chern number on Hamiltonians learned from randomized subsets (with replacement) of the entire dataset. The averaged Chern number approaches 1 as the batch size grows, and is reduced for small batches due to projection noise (Extended Data Fig. A.4f). For a batch size above 200, which is still a small fraction of our available data, the averaged Chern number is above 0.9 and quickly approaches 1 with increasing batch size. In that intermediate regime, postselecting on loss leads to a small but noticeable increase in the averaged Chern number (Extended Data Fig. A.4g). These observations provide

further evidence that the Chern number of our output distribution is consistent with having the value of 1.

We further study such Chern number evaluated on a noisy ensemble through numerical simulations, with the same approach as that used to calculate the string distribution. In Extended Data Fig. A.4h, we evaluate the Chern number on string distributions simulated for various initialization and per-layer errors. We find that our parameters are comfortably within the regime of the unit Chern number.

A.9 NUMERICAL SIMULATIONS WITH ERRORS

We perform circuit-level noisy numerical simulations for both the initial state-preparation step and the subsequent Floquet dynamics. The state preparation circuit consists of Clifford operations only, and we simulate it using the `stim` package¹⁵⁶. The following Floquet circuit has non-Clifford operations but the effective free-fermion dynamics enable efficient simulation by keeping track of the correlation matrix in Eq. (A.2) through unitary dynamics, as described in Eq. (A.4).

We incorporate the effect of errors to the numerical simulation of free-fermion systems in a stochastic fashion, with an error channel with strength p_l applied after each circuit layer and the result averaged over many noise realizations. The coherent errors simply modify the phase of applied gates. The lost qubits are kept track of in a separate data structure and all subsequent gates that include those qubits are removed (with the error model still applied). When evaluating the observables, we postselect the data such that all qubits of the target observable are present (as is done in the experiment). The stochastic single-site Paulis are also kept track of in a dedicated data structure, and they flip the sign of all the following XX, YY, and ZZ link operators that they anticommute with. At the end of the circuit, the final set of Pauli errors is propagated through the observables being evaluated and flips them accordingly. The dynamics of interacting fermions are simulated with the

approximate fermionic Gaussian state approach²⁵¹.

We use a simple ansatz for our noise per gate layer, which consists of single-qubit incoherent Pauli errors and atom loss, and tune the overall noise strength to match select observables and use those fixed values for the majority of simulations. This phenomenological approach is not directly connected to any particular fidelity, as it models (in a naive way) all experimental contributions. Since the state-preparation circuit cannot be simulated within the free-fermion framework, we initialize the noisy Floquet simulation with a layer of noise whose strength is chosen to match our experimental plaquette data, which gives the single-qubit initialization error of $p_{\text{ini}}=0.1$. Similarly, the effective noise per gate layer was calibrated by applying isotropic fermion evolution ($\theta=1$) and looking at the ZZ-link observables at depth 12, resulting in the fitted single-qubit error per circuit layer $p_l=0.01$. These error rates are divided between atom loss and (unbiased) single-qubit Pauli noise, with the loss constituting $\approx 6\%$ and 40% of p_{ini} and p_l , respectively. The simulations of circuits with interacting fermions assume perfect gate operations and include initialization errors calibrated to match the initial value. In general, for the fermion simulations, we found that the dominant effect of noise was an overall rescaling of quantities rather than a large change in qualitative trends.

A.10 EMERGENT PARTICLE NUMBER CONSERVATION

In the fermion quench experiments, increasing J_Z/J leads to emergent particle number conservation at certain time intervals (depths 6 and 12). Here we describe this process in more detail and provide basic derivations of the effective Floquet Hamiltonians. The Floquet unitary for a single cycle of the quench experiment is,

$$U_F = e^{i \sum_{\langle i,j \rangle_Z} K_{ij}^Z} e^{i \sum_{\langle i,j \rangle_Y} K_{ij}^Y} e^{i \sum_{\langle i,j \rangle_X} K_{ij}^X}, \quad (\text{A.11})$$

where $\langle i, j \rangle_{X/Y/Z}$ are the XX, YY, and ZZ links, respectively. The K_{ij}^Z term is, in terms of complex-fermion operators,

$$\sum_{\langle i, j \rangle_Z} K_{ij}^Z \sim J_Z \sum_i n_i = J_Z N_{\text{tot}}, \quad (\text{A.12})$$

where N_{tot} is the total particle number operator. Thus, for large J_Z values there is a strong term in the Hamiltonian proportional to the total particle number. Since the initial state is an eigenstate of N_{tot} , the subsequent evolution is projected into the subspace of N_{tot} with the same eigenvalue, which can be understood as orthogonal wavefunction components averaging out due to fast oscillations at the timescale of $1/J_Z$, effectively preserving the particle number and realizing complex-fermion dynamics. In Floquet evolution, we need at least two applications of U_F for the hopping along XX and YY to be affected by the N_{tot} operator and, thus, the shortest particle-conserving Floquet circuit is depth-6.

Such intuition can be further substantiated through analytical arguments on the operator level. As we show below, large-angle $ZZ(\theta)$ rotations effectively grow the nearest-neighbor link operators to length-4 strings in a way that can realize complex-fermion hopping. The effect of varying J_Z can be understood by looking at a composite two-site unitary,

$$\begin{aligned} ZZ(\theta) e^{i\varphi(X \otimes I)} ZZ(\theta) &= ZZ(2\theta) \\ &\times \exp[i\varphi(\cos(\theta\pi/2)(X \otimes I) + \sin(\theta\pi/2)(Y \otimes Z))], \end{aligned} \quad (\text{A.13})$$

where θ is proportional to J_Z and the remaining $ZZ(2\theta)$ rotation can, in principle, be removed by setting $\theta \rightarrow -\theta$ in one of the two initial $ZZ(\theta)$ s. This operation can be understood as a $ZZ(\theta)$ unitary growing the X operator into a $Y \otimes Z$ one, and similar relations hold for all basis combinations. Moreover, Eq. (A.13) governs the growth of general strings, as X can denote a particular site on a larger string operator. The special case of $\theta=1$, corresponding to the CZ gate, results in a complete

propagation.

Consider a hopping operator in a single direction, for example along an XX link. In particular, take four qubits, labeled 1, 2, 3, and 4, arranged such that the pairs (1, 2) and (3, 4) form ZZ links and qubits (2, 3) are connected by an XX link. After applying two cycles of U_F (with discarded YY terms), we can instead interpret it as the first hopping along XX followed by the second one that is propagated according to Eq. (A.13). In the average Hamiltonian, for large angle θ , this effectively realizes a $X_2X_3 + Z_1Y_2Y_3Z_4$ operator which corresponds to complex-fermion dynamics. Note that X_2X_3 and $Z_1Y_2Y_3Z_4$ commute so there are no Trotter errors, but in principle there can be additional terms from other sites. We numerically evaluate the particle conservation of the effective depth-6 Hamiltonian (without the 2θ term in Eq. (A.13)) and see that indeed the particle creation is suppressed most when $ZZ(\theta)$ realizes a CZ gate.

A.1.1 FERMION HOPPING AND DENSITY-DENSITY CORRELATIONS

In our quench experiments with fermionic Hamiltonians, we study time dynamics of two initialized particles. While the density of complex fermions at each step can reveal many spatial features, transport properties need to be inferred from other observables. For example, we could measure longer Pauli strings that include hopping terms. Alternatively, density-density correlations,

$$G_{ij} = \langle n_i n_j \rangle - \langle n_i \rangle \langle n_j \rangle, \quad (\text{A.14})$$

can also capture transport behavior. Intuitively, if a particle starts at site A and moves to site B, the density at sites A and B should be anticorrelated due to the particle leaving site A in order to appear at B.

Furthermore, for free-fermion states, we can show that the connected density-density correlations in Eq. (A.14) directly capture hopping strength. After expanding the density operators, $n_i = a_i^\dagger a_i$,

the correlation function becomes

$$G_{ij} = \langle a_i^\dagger a_i a_j^\dagger a_j \rangle - \langle a_i^\dagger a_i \rangle \langle a_j^\dagger a_j \rangle ,$$

in terms of complex-fermion operators. For free-fermion states, the four-body term can be further decomposed into two-body terms through Wick's theorem⁴⁴⁰,

$$\begin{aligned} \langle a_i^\dagger a_i a_j^\dagger a_j \rangle &= \langle a_i^\dagger a_i \rangle \langle a_j^\dagger a_j \rangle - \langle a_i^\dagger a_j \rangle \langle a_j^\dagger a_i \rangle \\ &\quad + \langle a_i^\dagger a_j^\dagger \rangle \langle a_j a_i \rangle , \end{aligned}$$

where the last term vanishes when the particle number is conserved. This gives the expression for connected density-density correlations,

$$G_{ij} = -|\langle a_j^\dagger a_i \rangle|^2 + |\langle a_i^\dagger a_j^\dagger \rangle|^2 ,$$

which contains the negative magnitude of the hopping current and a positive contribution from pair creation. Thus, for free-fermion states with well-defined particle number, such correlations should be negative, with a magnitude proportional to squared hopping strength.

A.12 FERMION EXCHANGE PROTOCOL

Here we describe elements of the fermion exchange experiment. These experiments focus on four complex fermion sites embedded in the full experimental array. In order to perform exact evolution without disturbing the rest of the system, we perform local gates by moving a subset of the atoms not involved in the exchange protocol to the storage zone. With this method, we can conveniently perform the required local gates without disturbing other atoms with either the Rydberg excitation




Figure A.5: Fermion exchange protocol. **a**, Four complex-fermion sites involved in the fermion exchange experiments, and where they are embedded in the full lattice. **b**, Full gate sequence for the exchange protocol (partial pair creation, first hop, second hop, and readout of the exchange phase by repeating the partial pair creation). Four-body interaction terms for the hopping sections are realized by a sequence of two-qubit gate applications. **c**, Depiction of qubits undergoing the two-qubit gates at each time step, denoted by red lines. The black qubits are the atoms that are picked up by AOD traps to perform the gates, and the layout is designed such that gates can be applied in parallel at each time step, as illustrated in the diagrams.

or moving AOD traps.

In the first step of the fermion exchange protocol, we use a sequence of three two-qubit unitaries to realize the three-body interaction term R which creates the superposition of zero and two fermions at sites A and D. Then, we apply two different hopping steps, along YY and XX links. Since the Majorana hopping terms along XX and YY links do not conserve the particle number, we realize particle-conserving hopping through a sequence of four parallel two-qubit gate operations. Finally, we apply R a second time to read out the exchange phase. In the absence of hopping, this final gate would complete the creation of two fermions at sites A and D, but due to the -1 exchange phase, we expect zero fermions in the final image (if no errors are present).

We benchmark our ability to realize particle-conserving complex fermion hopping terms by deterministically creating either 0 or 2 fermions. To create two fermions, here we apply R twice in a row. We observe that applying the hopping unitary preserves the vacuum state and hops the two

fermions at sites A and D to sites B and C, as expected. The results for both the full exchange and the control protocols match expectations and improve with postselection, providing evidence for observation of the exchange phase. A more robust protocol could involve measuring the full Ramsey fringe, for example by applying a variable phase to the two-fermion basis state before the final partial-creation operator.

A.13 FERMI-HUBBARD IMPLEMENTATION

Here we provide more information about the implementation of Fermi-Hubbard quantum simulations. To prepare the initial checkerboard configuration, we apply local gates on 16 of the sites after state preparation. For this circuit, we start with all operators on ZZ links with -1 value (by flipping the state of one of the data qubit sublattices after the CY operation). Then we use local Raman gates to flip the required ZZ operators to prepare the correct pattern of localized fermions. Here we use local single-qubit Z gates and convert them to local X and Y gates using global $\pi/2$ Raman pulses. These local gates also flip four columns of plaquettes, which we pre-compensate for by flipping ancilla measurement results in the decoder.

We decompose each Floquet step of our simulations into (a) fermion hopping within each half (b) coupling between the two halves to realize the spin interaction term. For independent hopping within each half of the system, we turn off the gates along XX and YY links connecting the two halves by modifying the atom motion pattern. The spin interaction gates commute with the final measurement Z basis, so we omit the final interaction step (e.g., after the first Floquet round both circuits are exactly the same data).

Concretely, through the entangling operations between the two halves, we realize contact fermion interactions with four-body terms $(Z_\uparrow Z_\uparrow) \times (Z_\downarrow Z_\downarrow)$ across the two halves. In terms of Majorana operators, these terms are proportional to $(ic_i c_{i'})_\uparrow (ic_i c_{i'})_\downarrow$; in terms of complex fermions, they are of

the form $4n_{\downarrow}n_{\uparrow}-2n_{\downarrow}-2n_{\uparrow}$. The first-order Floquet Hamiltonian is therefore given by,

$$H = \sum_{\sigma \in \uparrow, \downarrow} \sum_{\langle i, j \rangle} c_{\sigma, i} c_{\sigma, j} + U(c_{i\uparrow} c_{i\downarrow}), \quad (\text{A.15})$$

which reproduces the usual Fermi-Hubbard model when the complex fermion particle number is conserved.

B

Supplementary information about experiments in logical algorithms

Additional supplementary information can be found in Ref.^{[41](#)}.

B.1 CORRELATED DECODING

During transversal CNOT operations, physical CNOT gates are applied between the corresponding data qubits of two logical qubits. These physical CNOT gates propagate errors between the data qubits in a deterministic way: X errors on the control qubit are copied to the target qubit, and Z errors on the target qubit are copied to the control qubit (see Fig. B.1b). As a result, the syndrome of a particular logical qubit can contain information about the errors that have occurred on another logical qubit, at the point in time in which the pair underwent a transversal CNOT operation. We can leverage the information about these correlations and improve the circuit fidelity by jointly decoding the logical qubits involved in the algorithm. We note that this is closely related to other recent developments in decoding entire circuits, or so-called space-time decoding^{156,200,173,114}. It is also related to Steane error correction⁴¹¹, where errors are intentionally propagated from a data logical qubit onto an ancilla logical qubit, which is then projectively measured to extract the syndrome of the data logical qubit.

To perform correlated decoding, we solve the problem of finding the most likely error given the measured syndrome. We start by constructing a decoding hypergraph based on a description of the logical algorithm, which describes how each physical error mechanism (e.g. a Pauli-error channel after a two-qubit gate) propagates onto the measured stabilizers^{156,295}. The hypergraph vertices correspond to the stabilizer measurement results. Each edge or hyperedge corresponds to a physical error mechanism that affects the stabilizers it connects, with an edge weight related to the probability of that error. Each hyperedge can connect stabilizers both within and between logical qubit blocks (see Fig. 2b). We then run a decoding algorithm which uses this hypergraph, along with each experimental snapshot, to find the most likely physical error consistent with the measurements. This correction is then applied in software (with the exception of Fig. 4e, which is decoded real-time). Additional details for implementation of correlated decoding in these experiments can be found in




Figure B.1: Additional surface code data. **a**, Depiction of Bell state circuit and $d = 7$ surface codes. **b**, Diagram showing the transversal CNOT and physical error propagation rules. **c**, Covariance of the 48 measured stabilizers in both bases. The correlations near the diagonal corresponds to adjacent stabilizers within each block. Strong correlations are also observed with the stabilizers of the other block due to the error propagation in the transversal CNOT. **d**, Bell pair infidelity upper bound (as opposed to estimated Bell pair error in Fig. 2d of Main Text, see Methods text), showing improvement with increasing code distance. **e**, Probability of no detected error for each of the 96 measured stabilizers, showing agreement when compared to the theoretical values from empirically chosen error rates (experiment average = 77%, theory average = 82%). Note that X basis logical 1 and Z basis logical 2 have higher stabilizer error probability due to the error propagation in the transversal CNOT (reducing expectation values relative to if the transversal CNOT is not performed). **f**, Using the empirical error rates that correspond to data-theory agreement for the measured stabilizers in **e**, our simulations for improvement in Bell pair error, as a function of code distance, are in good agreement with experiments. The empirical error rates used are consistent with the 99.3% two-qubit gate fidelity, measured for this larger array, as well as the roughly 4% data qubit decoherence error (integrated over the entire circuit and measured by Ramsey method). These dephasing error rates are dominated by a complex moving sequence as we prepare the two surface codes in a serial fashion (see Supplementary Movie), and would be significantly smaller for a repetitive error correction experiment.

Ref.⁴¹.

B.2 DIRECT FIDELITY ESTIMATION AND TOMOGRAPHY

One challenge with logical qubit circuits is that convenient probes that are accessible with physical qubits may no longer be accessible. The GHZ state studied here provides such an example, since conventional parity oscillation measurements cannot be performed³⁰⁷. Instead, we use a technique known as direct fidelity estimation¹⁴⁰, which can be understood as follows. The target state ψ is the simultaneous eigenstate of the N stabilizer generators $\{S_i\}$, and so the projector onto the target state is $|\psi\rangle\langle\psi| = \prod_i^N (S_i + 1)/2$ (which is 1 if $S_i = 1 \forall i$, and 0 otherwise). One can thereby directly measure fidelity by measuring the expectation values of all terms in this product, which in other words refers to measuring the expectation values of all elements of the stabilizer group given by the exponentially many products of all the S_i . The logical GHZ fidelity is defined as the average expectation value of all measured elements of the stabilizer group. With our 4-qubit GHZ state, with 4 stabilizer generators $\{XXXX, ZZII, IZZI, IIZZ\}$, the 16-element stabilizer group is given by all possible products. We measure the expectation values of all 16 of these operators; for each element, we simply rotate each logical qubit into the appropriate logical basis and then calculate the average parity of the four logical qubits in this measurement configuration. We then directly average all 16 elements equally (with appropriate signs, as some of the stabilizer products should have -1 values), and in this way compute the logical GHZ state fidelity. This is an exact measurement of the logical state fidelity¹⁴⁰. Scaling to larger states can be achieved by measuring elements of the stabilizer group at random¹⁴⁰. To perform full tomography in Fig. 3e, we measure in all $3^4 = 81$ bases, thereby measuring the expectation values of all 256 logical Pauli strings, and reconstruct the density matrix by solving the system of equations with optimization methods.

B.3 SLIDING-SCALE ERROR DETECTION

Here we provide additional information about the sliding-scale error detection protocol applied for Figs. 3,5,6. Typically, error detection refers to discarding (or postselecting) measurements where any stabilizer errors occurred. In the context of an algorithm, however, discarding the result of an entire algorithm if just one physical qubit error occurred may be too wasteful, and one may want to only discard measurements where many physical qubits fail and the probability of algorithm success is greatly reduced. For this reason, for the algorithms here we explore error detection on a sliding-scale, where one can set a desired “confidence threshold”, where based on the syndrome outcomes one determines whether to accept a given measurement. Sliding this confidence threshold enables a continuous trade-off (in data analysis) between the fidelity of the algorithm and the acceptance probability. When sliding-scale error detection is applied, in all applicable cases we also apply error correction to return to the codespace.

We apply such a sliding-scale error detection for the color-code logical GHZ fidelity measurements in Fig. 3d. One possible method would be to discard measurements based on the number of detected stabilizer errors. However, this is suboptimal, both because on the color code a single physical qubit error can result from anywhere between 1 and 3 stabilizer errors, and also because errors deterministically propagate between codes during the transversal CNOT gates, such that a single physical error on one code can lead to detected errors on all codes, but which are still all correctable errors. As such, we perform the sliding-scale error detection utilizing the correlated decoding technique, and set the confidence threshold as a threshold weight of the overall correction weight on the decoding hypergraph. For example, in the color code GHZ experiment, a stabilizer error on all 4 logical qubits which is just consistent with a single physical qubit error that propagated to all 4 logical qubits, is in fact a low-weight (or high-probability) error as it corresponds to just a single physical qubit error. If the weight of hypergraph correction (inversely related to log of the probability that a

given error mechanism would have occurred leading to the observed syndrome outcome) is below the cut-off threshold weight, then the measurement is accepted; otherwise, it is rejected. For each threshold we then calculate the average algorithm result (y-axis of Fig. 3d), as well as the fraction of accepted data (x-axis of Fig. 3d).

In Fig. 5 with $[[8,3,2]]$ codes, for 3, 6, 24, and 48 logical qubits we apply our sliding-scale detection simply as given by the total number of stabilizer errors detected, although as illustrated above this can likely be improved by considering which stabilizer error patterns are more likely to cause an algorithmic failure. For the 12 logical qubits, in order to have a more fine-grained sliding-scale, for each of the $2^4 = 16$ possible stabilizer outcomes we calculate the XEB to rank the likelihood that each of the observed stabilizer outcomes leads to an algorithmic failure, and then use this ranking when deciding whether a given measurement is above/below the cut-off threshold. In Fig. 6b we set the threshold by the number of stabilizer errors, and in Fig. 6d, to have more fine-grained sliding-scale information we take different subsets of stabilizer outcome events that are all below the threshold of allowed number of stabilizer errors, and calculate the y-axis (Pauli expectation value) and x-axis (purity) for all of them. Broadly, there are many ways to perform this sliding-scale error detection, and this can be useful both as continuous trade-offs between fidelity and acceptance probability, as well as for use in techniques such as zero-noise extrapolation in data analysis.

B.4 OVERVIEW OF QEC METHODS

Here we provide a brief overview of key QEC methods used in our work.

Code distance, decoding, and thresholds. $[[n, k, d]]$ notation describes a code with a number of physical qubits n , a number of logical qubits k , and a code distance d . The code distance d sets how many errors a code can detect or correct. The code distance is the minimum Hamming distance between valid codewords (logical states), i.e. the weight of the smallest logical operator¹⁷². In the case of the

two-dimensional surface and color codes studied here, d is equivalent to the linear dimension of the system¹⁴¹.

Following this definition, quantum codes of distance d can detect any arbitrary error of weight up to $d - 1$. Such errors cause stabilizer violations, indicating that errors occurred. Postselecting on the results with no such stabilizer violations corresponds to performing error detection, which protects the quantum information up to $d - 1$ errors at the cost of postselection overhead. Conversely, codes can correct fewer errors than they detect (but without any postselection overhead). The correction procedure brings the system back to the closest logical state (codeword); thus, if more than $d/2$ errors occur, the resulting state may be closer to a codeword different from the initial one, resulting in a logical error¹⁷². For this reason, codes of distance d can correct any arbitrary error of weight up to $(d - 1)/2$ (rounded down if d is even¹⁴¹). The process of decoding refers to analyzing the observed pattern of errors and determining what correction to apply to return back to the original code state and undo the physical errors created. In many cases, such as with the 2D surface and color codes, one does not need to apply the correction in hardware (physically flipping the qubits); instead it is sufficient to undo an unintended X_L/Z_L operator that was applied by hardware errors by simply applying a “software” X_L/Z_L operator¹⁴¹, also described as Pauli frame tracking²⁴⁷.

As the size of an error correcting code and the corresponding code distance is increased, so are the opportunities for errors to occur as the number of physical qubits increases. This leads to a threshold behavior in quantum error correction: if the density of errors p is above a (possibly circuit-dependent) characteristic error rate p_{th} , then increasing code distance will worsen performance. However, if $p < p_{\text{th}}$, then increasing code distance will improve performance¹⁴¹. Theoretically, since one requires $(d + 1)/2$ errors to create a logical error, the logical error rate will be exponentially suppressed as $\propto (p/p_{\text{th}})^{(d+1)/2}$ at sufficiently low error rates¹⁴¹. The performance improvement with increasing code distance, observed for the preparation and entangling operation in Fig. 2, implies that we surpass the threshold of this circuit. We note that in this regime, improving fidelities by

e.g. a factor of $2x$ can then lead to an error reduction of $2^4 = 16x$ for the distance-7 code studied, and further exponential suppression with increasing code distance. This rapid suppression of errors with reduced error rate and increased code distance is the theoretical basis for realizing large-scale computation. We emphasize that thresholds can be circuit-dependent, as discussed in detail in the surface code section below.

Fault-tolerance and transversal gates. A common definition of fault-tolerance in quantum circuits¹⁷² (which we use in this work) is that a weight-1 error (i.e. an error affecting one physical qubit), cannot propagate into a weight-2 error (now affecting two physical qubits) within a logical block. This property implies that errors cannot spread within a logical block, and thereby prevents a single error from growing uncontrollably and causing a logical error.

Distance-3 codes, which are of significant historical importance^{398,410}, can correct any weight-1 error. Fault-tolerance is particularly important for these codes, since otherwise a weight-1 error can lead to a weight-2 error and thereby cause a logical fault. An important characteristic of a fault-tolerant circuit that uses distance-3 codes¹⁷² is that (in the low error rate regime), physical errors of probability p lead to logical errors with probability $\propto p^2$. We emphasize that the notion of fault-tolerance refers to circuit structuring to control propagation of errors, but a circuit can be fault-tolerant with low fidelity, or non-fault-tolerant with high fidelity. For example, even if a weight-1 error can lead to a weight-2 error, but the code has high distance, or if this error propagation sequence is possible but highly unlikely, then this property may not be of practical importance (for this reason definitions of fault-tolerance may vary). In practice, the goal of QEC is to execute specific algorithms with high fidelity, and fault-tolerant structuring of a circuit is one of many tools in the design and execution of high-fidelity logical algorithms.

Transversal gates, defined here as being composed of independent gates on the qubits within the code block (i.e., entangling gates are not performed between qubits within the same code block)¹²², constitute a direct approach to ensure fault-tolerant structuring of a logical algorithm. Since transver-

sal gates imply performing independent operations on the physical constituents of a code block, errors cannot spread within the block, and fault-tolerance is guaranteed. In the present work, all logical circuits we realize (following the logical state preparation) are fault-tolerant, as all logical operations we perform are transversal. Note, in particular, that even though the transversal CNOT allows errors to propagate between code blocks, this is still fault-tolerant as it does not lead to a higher weight error within the block, and thereby a single physical error can neither lead to a logical failure nor an algorithmic failure. Importantly, the large family of codes referred to as Calderbank-Shor-Steane (CSS) codes all have a transversal CNOT³⁹⁹, all of which can be implemented with the single-step, parallel-transport approach here.

Although all the logical circuits we implement are fault-tolerant, the logical qubit state preparation is fault-tolerant for our $d = 3$ color code (Figs. 3,4) and $d = 3$ surface code (part of Fig. 2), but is non-fault-tolerant for the state preparation of our $d = 5, 7$ surface codes and $[[8,3,2]]$ codes. Thus, all of our experiments with the $d = 3$ color codes are fault-tolerant from beginning to end, and so the entire algorithm is fault-tolerant and theoretically has a failure probability that scales as p^2 . However, we note that having a fault-tolerant algorithm also does not imply that errors do not build up during execution of the circuit. For this reason deep circuits require repetitive error correction^{170,253} to constantly remove errors and continuously benefit from the, e.g., p^2 suppression.

Our logical GHZ state theoretically has a failure probability scaling as p^2 . Nevertheless, the error build-up (increasing p) during the operations of the circuit and the spreading of errors through transversal gates, limits our logical GHZ fidelity to 72%. This is consistent with numerical modeling. Similar to the surface code modeling (Fig. B.1) we use empirical error rates consistent with 99.4% two-qubit gate fidelity as well as roughly 4% data qubit decoherence error (including SPAM) over the entire circuit. We simulate the experimental circuit (including the FT state preparation with the ancilla logical flag) and measurements of all 16 elements of the stabilizer group (see direct fidelity estimation section), and extract a simulated logical GHZ fidelity of 79%. This is slightly higher than

our measured 72% logical GHZ fidelity, possibly originating from imperfect experimental calibration. This modeling indicates that our logical GHZ fidelity is limited by residual physical errors, which will be reduced quadratically as p^2 with reduction in physical error rate p , in particular by reducing residual single-qubit errors which were larger during this measurement and are dominating the error budget here.

B.5 SURFACE CODE IMPLEMENTATION.

In 2D planar architectures, such as those associated with superconducting qubits^{170,253}, stabilizer measurement is the most important building block of error-corrected circuits¹⁴¹. In such systems, stabilizers need to be constantly measured in order to correct qubit dephasing and increase coherence time, as demonstrated recently¹⁷⁰. Logic operations are implemented by changing stabilizer measurement patterns, enabling realization of techniques such as braiding¹⁴¹ and lattice surgery²⁰⁵. Similar techniques can be used to move logical degrees of freedom in order to implement nonlocal logical gates¹⁵⁷. Due to this gate execution strategy, d rounds of stabilizer measurement are required for each entangling gate for ensuring fault-tolerance¹⁴¹.

Neutral atom quantum computers feature different challenges and opportunities. Specifically, they feature long qubit coherence times ($T_2 > 1\text{s}$), which can be further increased to the scale of 10 - 100s with well-established techniques²⁷. By using the storage zone, qubits can be idly, safely stored for long periods without repeated stabilizer measurements. Hence, from a practical perspective increasing qubit coherence by using a logical encoding does not provide immediate gains in improving quantum algorithms, and the gains will be from improving the fidelity of entangling operations. Moreover, logic gates and qubit movement do not have to be performed with stabilizer measurements. Instead, they can be executed with nonlocal atom transport and transversal gates. Since such transversal gates are intrinsically fault-tolerant, they do not necessarily require d rounds of correction

after each operation. Even syndrome measurement may be better executed in certain cases by techniques such as Steane error correction⁴¹¹ (similar to our ancilla logical flag with color codes as used in Fig. 3) as opposed to repeated stabilizer measurement. For these reasons, the transversal CNOT is among the most important building blocks in error-corrected circuits. Hence, we here focus on improving the transversal CNOT by scaling code distance.

Specifically, we use the so-called rotated surface code¹⁷⁰, which has code parameters $[[d^2, 1, d]]$. Our distance-7 surface codes (as drawn in Fig. 2d) are composed of 49 physical data qubits, with 24 X stabilizers (light blue squares) and 24 Z stabilizers (dark blue squares), and 1 encoded logical qubit described by anticommuting weight-7 operators, the horizontally oriented X_L and the vertically oriented Z_L . The X and Z stabilizers commute with the X_L and Z_L logical operators, allowing one to measure the stabilizers without disturbing the underlying logical degrees of freedom. In our experiments, we prepare one surface code in $|+_L\rangle$ and one surface code in $|0_L\rangle$. In the first code, this is realized by preparing all physical data qubits in $|+\rangle$, thereby preparing an eigenstate of X_L and the 24 X stabilizers, and then projectively measuring the 24 Z stabilizers with 24 ancilla qubits (Fig. 2d red dots) using four entangling gate pulses¹⁴¹. The second code is prepared similarly, but with all physical qubits initialized in $|0\rangle$, thus preparing an eigenstate of Z_L and the 24 Z stabilizers, and then projectively measuring the 24 X stabilizers with 24 ancillas. The CNOT is directly transversal because these two surface code blocks have the same orientation, and does not require rotation of the lattice to implement a H . The projective measurement of the ancillas defines the values of the stabilizers. During the transversal CNOT, the values of the stabilizers are copied onto the other code as well, and is tracked in software.

Since we only perform a single round of stabilizer measurement, our state preparation scheme is non-fault-tolerant (nFT) for the $d = 5, 7$ codes. Consider, for instance, the case when all stabilizers are defined as $+1$, and no errors are present in the system, but an ancilla measurement error in the middle of the surface code lattice yields a stabilizer measurement $= -1$. Correction then causes

a large-weight pairing of this apparent stabilizer violation to the boundary¹⁷. Hence this single ancilla measurement error can lead to several data qubit errors, resulting in nFT operation. The $d = 3$ code initialization is a special case which does not suffer from this issue¹⁷¹. Higher-order considerations about fault-tolerance given by gate ordering during stabilizer measurement can also be considered¹⁷⁰.

The effect of these nFT errors from noisy syndrome extraction is to cause X physical errors on the $|+_L\rangle$ state, and Z physical errors on the $|0_L\rangle$ state. Thus in performing just state preparation and measurement, the presence of these errors would not be directly apparent, as these errors commute with measuring the $|+_L\rangle$ in the X -basis and $|0_L\rangle$ in the Z basis. As such this circuit would not be a good benchmark of surface code state preparation. Conversely, the transversal CNOT experiment is sensitive to the various aspects of the circuit and a good probe of performance. Since we measure the Bell state in both the $X_L^1 X_L^2$ and $Z_L^1 Z_L^2$ bases, the nFT errors in both bases will propagate through the logical CNOT and cause errors on both logical qubits in both the X and Z bases. For these reasons, unlike a surface code SPAM measurement, this experiment is a good probe of logical performance. In fact, the effect of these nFT errors is such that if we just apply conventional decoding within each logical block, then we find that the Bell state degrades substantially with increased code distance (Fig. 2d).

The effects of this nFT preparation are suppressed (but are not entirely removed) by using the correlated decoding technique. For example, consider a nFT-induced apparent stabilizer violation to the left of the middle line in the lattice of the $d = 7$ $|+_L\rangle$ state, corresponding to a chain of 3 physical X errors to the boundary. These errors will propagate through the logical CNOT onto the second logical qubit, and affect the independent measurement of both logical qubits in the Z basis when probing the $Z_L^1 Z_L^2$ stabilizer. When decoded independently, if another single X error occurs on the first block after the CNOT moving the stabilizer violation to the right of the middle line, becoming a chain of 4 X physical errors, this will cause an incorrect pairing and lead to an indepen-

dent X_L^1 error on this code only and thereby corrupt the $Z_L^1 Z_L^2$ stabilizer, and would correspond to a total weight-6 correction between the two codes. However, when decoded jointly with correlated decoding, these errors can be effectively decoded since they will appear on the stabilizers of both logical qubits. In this example, the lowest weight pairing would remove this chain of 3 X errors from both codes, and leave only the 1 remaining X error on the first block, which can also successfully be decoded (total pairing weight here is only 2). Our correlated decoding technique is thus essential to our observation of improved Bell performance with code distance.

Finally, we elaborate on our evaluation of Bell pair error. Bell state fidelity is given by the average of the populations and the coherences, which for physical qubits can be measured as the ZZ populations and the amplitude of parity oscillations. In the language of stabilizers, the parity oscillation amplitude is given by the average of $\langle XX \rangle$ and $-\langle YY \rangle$ ^{42,4}. With the surface code we cannot conveniently measure the Y_L operators fault-tolerantly (and is why we use color codes for programmable Clifford algorithms and full tomography, see next section). For this reason, we estimate the logical coherences as $\langle X_L X_L \rangle$, which we then average with the populations for calculating the Bell pair error. To support the validity of this analysis, we can instead calculate a lower bound on the Bell state fidelity^{42,4}, which also shows the same improvement in performance as we increase code distance (Fig. B.1d).

Correlated decoding in the surface code. Following the above discussion, we provide additional insights related to the correlated decoding in the case of the surface code transversal CNOT. Consider a circuit where perfect (noiseless) surface codes are initialized, a transversal CNOT is executed, and then projective readout is performed. If errors occur before the transversal CNOT, then these errors can propagate; e.g., an X physical error on the control logical qubit will propagate onto the target logical qubit, and thereby double the density of errors on the target logical qubit. By multiplying the projectively measured Z stabilizers of the target logical qubit with those of the control logical qubit, the propagation is undone. Now the target logical qubit only has to decode its original den-

sity of X errors. The same considerations can be made for Z errors originating on the target logical qubit that propagate onto the control logical qubit. However, if there are errors after the transversal CNOT, then multiplying the stabilizers instead doubles the density of such errors. Thus, the optimal decoding strategy if errors are only after the transversal CNOT is to perform independent matching within both codes. The general case where errors are present both before and after the transversal CNOT corresponds to neither case, and is modeled by our decoding hypergraph that has edges and hyperedges connecting the two logical qubits, with edgeweights informed by our experimental error model. Fig. B.5 explores decoding performance with different values of the scaled weights of the edges and hyperedges that connect the stabilizers of the two logical qubits. These results illustrate that the correlated decoding is robust (but not completely insensitive) to the nFT errors associated with ancilla measurement errors. This feature would also be recovered by the simpler multiplication decoder, which would be entirely insensitive to errors from ancilla measurement, but however is more sensitive to errors after the CNOT. Specifically, Fig. 5c shows that our optimized decoder is not simply a “multiplication decoder” as the ancilla measurement values indeed contribute to the correction procedure and make the correlated decoding more robust to decoder parameters. For a given logical circuit, our correlated decoding procedure generates a decoding hypergraph which we then solve using most likely error methods, which is done here for both surface code and color code experiments, and can generically be applied to any stabilizer codes and Clifford circuits¹¹⁴. More theoretical details and discussion of correlated decoding will be presented in Ref.⁷⁰.

2D color codes. Two-dimensional color codes are topological codes, which are similar to surface codes²⁵⁵. Often portrayed in a triangular geometry, the color codes used here are a tiling of three colors of weight-4 and weight-6 stabilizers, with X_L and Z_L operator strings running along the boundary of the code²⁵⁵. In this work we study two-dimensional $d = 3$ color codes, as portrayed in Fig. 3a, which only contain weight-4 stabilizers given by the products of X and Z on the qubits of

each colored plaquette. This $d = 3$ color code is identical to the 7-qubit Steane code. However, we emphasize that the techniques used here directly apply to larger-distance color codes⁷⁹.

Although the color codes are similar to surface codes, an important difference is that in the color code, the X and Z stabilizers lie directly on top of the same qubits (as opposed to being on dual lattices with respect to each other), and similarly the X_L and Z_L operators lie on top of each other (as opposed to propagating in the orthogonal directions on the surface code). In other words, the operators here are symmetric and related by a global basis transformation. This has important consequences for the allowed transversal gate set^{254,336}. In particular, although the surface code technically has a transversal H which transforms $X_L \leftrightarrow Z_L$, it requires a physical 90-degree rotation of the code block. While such lattice rotations are possible using atom motion techniques, for many circuits it is inconvenient. Conversely, in the color code H is transversal: it directly exchanges $X_L \leftrightarrow Z_L$ as well as the X and Z stabilizers. This difference is even more important for the transversal S gate which is possible for the color code. Here, transversal S exchanges $X_L \leftrightarrow Y_L$ (where Y_L is given by the product of X_L and Z_L which lie on top of each other) as intended, and the X stabilizer of a given plaquette returns back to itself by multiplying the Z stabilizer of that same plaquette. (This is in contrast to the surface code, which does not have a transversal S , where the Y_L operator is a product of horizontally propagating X_L and vertically propagating Z_L ¹⁴¹.) Since the color code has the entire transversal gate set of $\{H, S, \text{CNOT}\}$ and also does not require tracking any lattice rotations, it is well-suited to exploration of programmable logical Clifford algorithms.

For fault-tolerant preparation of the $d = 3$ color code we use a modified version of the scheme summarized in Ref.¹⁷¹, where instead of the 8-gate encoding circuit, we use a 9-gate encoding circuit that is more conveniently mapped to specific atom movements in our system (corresponding to graph state preparation similar to Ref.⁴³), followed by a transversal CNOT with an ancilla logical flag. The logical SPAM fidelity is then calculated as the probability of observing $|0_L\rangle$ after decoding. We note that in Fig. 3 we could also have made a 5-qubit GHZ state but made a 4-qubit GHZ

state for simplicity of performing full tomography. In Fig. 4, when Bell state fidelities with feedforward are reported we estimate the logical coherences as the average of $\langle X_L X_L \rangle$ and $-\langle Y_L Y_L \rangle$, which we then average with the $Z_L Z_L$ populations (not plotted) for calculating the Bell pair fidelity. Finally, we note that the feedforward Bell state in Fig. 4e could also be performed with a software Z_L rotation on either one of the two qubits allowing one to correct to the appropriate Bell state, but here we do the feedforward S on both qubits to test our feedforward capabilities; this technique is directly compatible with performing magic state teleportation¹⁴¹.

Clifford and non-Clifford gates, and universality. 2D topological codes such as the surface and color codes have transversal implementation of Clifford gates (e.g., $\{H, S, \text{CNOT}\}$). This gate set is not universal, i.e. it cannot alone be used to realize an arbitrary quantum computation, and requires a non-Clifford gate such as $\{T, \text{CCZ}\}$ for achieving universal computation. Moreover, circuits composed solely of stabilizer states and Clifford gates can be simulated in polynomial time due to the Gottesman-Knill theorem¹. This can be understood as stabilizer tracking: for example, consider a 3-qubit system where a stabilizer of the state is $X \otimes I \otimes I$, where X stabilizes the $|+\rangle$ state and I is the identity. Applying two CZ entangling gates $\text{CZ}_{1,2} \otimes \text{CZ}_{1,3}$ transforms this stabilizer to $X \otimes Z \otimes Z$, since an X flip before the CZ simply changes whether a Z flip will be applied to the other qubits. Even though Clifford circuits create superposition and entanglement between qubits, the N initial stabilizers of the state can simply be tracked as they propagate through the circuit (so-called operator spreading²⁹⁸), and thereby simulation of the circuit can be easily accomplished.

The effect of non-Clifford gates, however, is significantly more complex. For example, passing the stabilizer $X \otimes I \otimes I$ through a CCZ maps into a superposition of Pauli strings, i.e. $X \otimes I \otimes I \rightarrow 1/4(X \otimes I \otimes I + X \otimes Z \otimes I + X \otimes I \otimes Z - X \otimes Z \otimes Z)$, as an X flip now changes whether a CZ operator will be applied on the other qubits, resulting in 4x more operators to track after the single CCZ. (The CZ operator matrix is simply equal to $1/4 [I \otimes I + Z \otimes I + I \otimes Z - Z \otimes Z]$). This causes not only operator spreading, but also so-called operator entanglement²⁹⁸. As we apply additional non-

Clifford gates, the number of operators to track will grow exponentially, and eventually will become computationally intractable. E.g. state-of-the-art Clifford + T simulators can handle roughly 16 CCZ gates⁵⁴. This is the basis behind our complex sampling circuits, where the 48 CCZs on the 48 logical qubits create a high degree of scrambling and magic (defined below), rendering Clifford + T simulation impractical.

B.6 $[[8,3,2]]$ CIRCUIT IMPLEMENTATION

Here we provide additional detail about our $[[8,3,2]]$ circuit implementations. The $[[8,3,2]]$ code blocks are initialized in the $| -_L, +_L, -_L \rangle$ state with the circuit in Fig. B.2, which can be understood as preparing two 4-qubit GHZ states (corresponding to $[[4,2,2]]$ codes²⁷⁶), i.e. $\text{GHZ}_Z^{1,3,5,7} \otimes \text{GHZ}_X^{2,4,6,8}$, and subsequently entangling them as illustrated in Fig. B.2a (as well as applying Z gates). In our circuit implementations, for system sizes of 3 to 24 logical qubits both for sampling and two-copy measurements, we prepare 8 blocks encoded over 64 physical qubits. For the 48 logical qubit circuit (128 physical qubits total) we encode 8 blocks and entangle them, and then drop them into storage; then, we pick up 64 new physical qubits from storage, encode them into 8 blocks in the entangling zone and entangle them. Finally, we bring the original 8 blocks from storage and entangle them with the second group of 8 blocks in the entangling zone (Fig. B.2) (see Supplementary Movie).

The transversal gate set of the $[[8,3,2]]$ code is enabled as follows (see also Refs. ^{428,204,435,71}). The transversal CNOT between blocks immediately follows from the fact that the $[[8,3,2]]$ code is a CSS code. In-block CZ gates between two logical qubits L_i and L_j (CZ_{L_i,L_j}) can be realized by $S, S^\dagger$ gates on the face corresponding to logical qubit L_k . For example, consider applying the pattern of $S, S^\dagger$ gates to the top face in Fig. 5, i.e., $S_1 S_3^\dagger S_5^\dagger S_7$, which transforms $X_{L1} = X_1 X_2 X_3 X_4$ to $X'_{L1} = -Y_1 X_2 Y_3 X_4$, which is equal to $X'_{L1} = X_{L1} Z_{L2}$, and the same applies to give $X'_{L2} = X_{L2} Z_{L1}$;

i.e., a CZ is realized between logical qubits 1 and 2. This procedure can also be used to understand why the pattern of $T, T^\dagger$ realizes a CCZ between the three encoded qubits. CCZ gates should map X_{L3} to $X_{L3} \otimes \text{CZ}_{L1,L2}$. By applying the pattern of $T, T^\dagger$ in Fig. 5a, each X face maps to itself multiplied by a pattern of $S, S^\dagger$, e.g. $X_{L3} = X_1 X_3 X_5 X_7$ maps to $X'_{L3} = X_{L3} S_1 S_3^\dagger S_5^\dagger S_7$, or then $X'_{L3} = X_{L3} \otimes \text{CZ}_{L2,L3}$. This happens for all three X_L faces, thereby realizing a CCZ gate. Finally we detail the permutation CNOT, which was also developed in Ref. ⁴³⁵. Physically permuting atoms to swap qubits $4 \leftrightarrow 8$ and $3 \leftrightarrow 7$ takes $X_{L1} = X_1 X_2 X_3 X_4$ to $X'_{L1} = X_1 X_2 X_7 X_8$ or instead $X'_{L1} = X_{L1} X_{L2}$ (also by multiplying the global X stabilizer), and similarly it can be seen by tracking the qubit permutations that $Z'_{L2} = Z_{L2} Z_{L1}$, i.e. realizing a CNOT. Finally, although these 3D codes do not have a transversal H , as they are CSS codes they can be initialized and measured in either the X or Z basis, effectively allowing H gates at the beginning or end of the circuit.

In-block logical entangling gates are applied block-by-block, and any in-block gate combination can be realized. For conceptual simplicity we apply only two particular local Raman patterns in layers. The first is the gate combination $\text{CCZ}_{L1,L2,L3} \cdot \text{CZ}_{L1,L2} \cdot \text{CZ}_{L1,L3} \cdot \text{CZ}_{L2,L3} \cdot Z_{L1} \cdot Z_{L2} \cdot Z_{L3}$, given by applying $T^\dagger$ on the entire physical qubit block, and the second gate combination we apply is $\text{CCZ}_{L1,L2,L3} \cdot \text{CZ}_{L2,L3} \cdot \text{CZ}_{L1,L3} \cdot Z_{L3}$, given by applying T on the top row and $T^\dagger$ on the bottom row. In our circuits we alternate layers of in-block transversal entangling gates and out-block transversal CNOTs, entangling logical blocks on up to 4D hypercubes (see Fig. B.2)^{195,256,228}. We keep the control and target qubits the same throughout the circuit for conceptual simplicity, allowing the local physical H gates on the target qubits to be compiled with the in-block gate layers, but the control-target direction can also be chosen arbitrarily. We ensure that in-block logical entangling gates are applied such that they do not trivially commute through and cancel with earlier entangling gate applications. As an experimental note, we note that for the Clifford states realized in the other parts of this work, stabilizers take on values of either +1 or -1 (due to, e.g., use of physical $\pi/2$ rotations instead of H), which is then simply re-defined in software. Since for our $[[8,3,2]]$ circuits we




Figure B.2: $[[8,3,2]]$ and hypercube encoding. **a**, State preparation circuit for the $[[8,3,2]]$ code, in which two 4-qubit GHZ states are simultaneously prepared and subsequently entangled. This initializes an $[[8,3,2]]$ code with logical states $|-L_1, +L_2, -L_3\rangle$. **b**, 4D hypercube circuit performed on 48 logical qubits (128 physical qubits). The circuit is drawn on the block-level, where each block consists of 3 logical qubits and 8 physical qubits. The first in-block gate layer is performed with a global $T^\dagger$. The local gate patterns, and the corresponding logical gates they execute within each code block, are illustrated in the inset. **c**, Diagram illustrating the code block movements and use of the processor's zoned architecture throughout the circuit. Initially, eight $[[8,3,2]]$ code blocks are prepared in the entangling zone and atoms for later state preparation of eight additional code blocks are loaded in the storage zone. The code blocks in the entangling zone are then picked up and interlaced with adjacent blocks to perform three transversal CNOT layers. The two groups of eight code blocks are then swapped and the same procedure is repeated with the second group of code blocks. The first group of code blocks are then moved back into the entangling zone and interleaved with the atoms of the first group to perform a final parallel transversal CNOT. The layers of CNOT gates connect the code blocks such that a 4D hypercube on 16 blocks of $[[8,3,2]]$ codes is constructed.

implement non-Clifford gates on the physical level, it is important to ensure all stabilizers are initialized and maintained as $+1$; e.g., if a Z stabilizer is -1 , then the logical CCZ implementation sends the X stabilizer expectation value to 0 . This can be understood as a physical X on a single site transforming to a superposition $(X + Y)/\sqrt{2}$ by physical T 's, going into an equal superposition of X stabilizer being $+1$ and -1 .

B.7 CLASSICALLY HARD CIRCUITS WITH $[[8,3,2]]$ CODES

Our implemented circuits are equivalent to Instantaneous Quantum Polynomial (IQP) circuits⁵⁸, which gives a theoretical basis for understanding why our circuits could in principle be classically hard to simulate, for which we also provide numerical evidence of so-called anticoncentration¹⁸⁹. IQP circuits are defined as initializing $|+\rangle^{\otimes n}$ on n qubits, applying a diagonal entangling unitary such as those composed by $\{CCZ, CZ, Z\}$, and then measuring in the X -basis^{58,59}. A uniform superposition of 2^n bitstrings is created, the diagonal gates apply -1 signs in a complicated fashion to the exponentially many bit strings, and then “undoing the superposition” with the final H before measurement now results in an intricate “speckle” interference pattern¹⁷. Sampling from the output distribution of this speckle pattern can be done efficiently on a quantum device that implements the circuit, but is exponentially costly on a classical device for certain choices of IQP circuits⁵⁹. The transversal gate-set of the $[[8, 3, 2]]$ code, as described above, contains diagonal gates $\{CCZ, CZ, Z\}$ that apply -1 signs to the bitstrings, but is made non-diagonal by the application of CNOTs, which permute bitstrings. Since this bitstring permutation does not break the IQP framework, these circuits are equivalent to an effective IQP circuit, but which is significantly more complex: e.g., circuits with 48 CCZs and 96 CNOTs map to effective IQP circuits with roughly 1000 CCZ gates. Nevertheless, since IQP circuits are a well-understood framework we can discuss our circuit properties with this toolset.




Figure B.3: Additional $[[8,3,2]]$ circuit sampling data. **a**, Overlap of error-detected 12-qubit sampling data with the theoretical distribution (same data as fully error-detected case in Fig. 5b). Progressive zoom-in's show the agreement between theory and experiment, down to the level of 10^{-4} probability per bit string. This error-detected data set is composed of 23,545 shots (raw data set is 138,626 shots). Note that we simultaneously measure on two groups of 12 logical qubits; plotted here is only one of the two 12-logical groups with an XEB of 0.69(1), while in plots Figs. 5e,f and Fig. B.3b we average the two logical groups, which gives a measured XEB of 0.616(7). **b**, Same data as Fig. 5f of Main Text but with purity (orange), as measured by two-copy measurement, additionally plotted. The measured XEB is slightly below the measured purity, providing evidence that XEB is a faithful fidelity proxy. We further note that under error detection, the logical XEB for these IQP circuits should be a good fidelity proxy. Interestingly, the behavior can be different for the raw, uncorrected data, as the circuit we apply on the physical level is not IQP. Without applying error detection, not all errors are logical errors and therefore the circuit differs from IQP behavior and can lend itself to a different scaling. For systems of 3, 6, and 12 logical qubits, multiple systems are measured in parallel and their results are averaged together. We note that although our preparation of $[[8,3,2]]$ code states makes these states on a cube, it does not have CNOTs between two pairs of qubits in the first step and, therefore, does not have the full gate connectivity of a cube. Instead, one can interpret these CNOTs as having been included but then compiled away as they commute with the state. We neglect this in plotting our physical qubit connectivity, which is derived from entangling 3D cubes on a 4D hypercube connectivity, realizing a 7D hypercube. **c**, 48-qubit XEB sliding-scale error-detection data. The point with full postselection on all stabilizers being perfect returned only 8 samples, so we omit this point from the Main Text plot for clarity.

We experimentally probe these circuits with the XEB¹⁷, defined as $\text{XEB} = 2^{N_L} \sum_i p(x_L^i) q(x_L^i) - 1$, where N_L is the number of logical qubits, $q(x_L^i)$ is the measured probability distribution for our logical qubits and $p(x_L^i)$ is the calculated probability distribution; here, we normalize the XEB by its ideal value such that the XEB for the noiseless circuit is 1. In typical cases, if noise overwhelms the circuit, the measured distribution will be uniform¹⁷, and the measured XEB will be 0.

The IQP circuits are a good setting for quantum-advantage-type experiments, as the bitstring distribution of IQP circuits with randomly applied $\{\text{CCZ}, \text{CZ}, \text{Z}\}$ gates (random degree-3 polynomials) is known to be classically hard to simulate^{59,60}. In Ref.¹⁹¹, we show that the ensemble of random hypercube IQP circuits (hIQP), whose instances are experimentally explored here, converges to the uniform IQP ensemble as the depth and size of the hypercube is increased. In Fig. B.6a, we show that hIQP circuits with random in-block operations and randomized control-targets on the out-block CNOT layers (realizing the hypercube) anti-concentrate quickly as the dimension of the hypercube is increased, with XEB eventually reaching the uniform-IQP value of 2. We also find that the presence of non-Clifford CCZ gates, which are critical for the computational hardness here, further improves anti-concentration properties, as we observe the ideal XEB of experimental circuits approach 2 as well, even without much randomization.

Moreover, XEB turns out to be a better benchmark for IQP circuits than for generic random circuit sampling settings (such as Haar-random circuits)^{150,309,436}. For IQP, XEB is close to the many-body fidelity and the difference can be theoretically bounded under reasonable noise assumptions¹⁹¹. Intuitively, this fact is related to the diagonal structure of the IQP circuits, which allows XEB to capture errors in a manner closer to fidelity, despite being defined only in the computational basis. In other words, a Z error will always corrupt the X basis measurement, and an X error (except one immediately before measurement) will create new Z errors that also corrupt the X basis measurement. Thus, in the fully postselected regime, where errors at the end of the circuit are well-described by logical errors, we expect the XEB to be a good measure of fidelity. We further note that,

in addition to the efficient generation of complex IQP circuits here, the $[[8,3,2]]$ gate set presented here can realize arbitrary IQP circuits composed of $\{\text{CCZ}, \text{CZ}, Z\}$ gates³⁹⁴. The in-block $\{\text{CCZ}, \text{CZ}, Z\}$ operations can be applied to any groupings of qubits by noting that combining the in-block and out-block CNOTs allow us to compose arbitrary transversal SWAP operations of targeted individual logical qubits between different blocks.

B.8 SIMULATION OF BITSTRING PROBABILITIES

To calculate the logical bitstring probabilities necessary for evaluating XEB and benchmarking our circuits, we employ a hybrid simulation approach combining wavefunction and tensor-network³²³ methods. It works best only when performing all of the entangling gates of the hypercube a single time, and relies on the fact that the final round of CNOTs is immediately followed by a measurement, simplifying network contraction. Concretely, for a D -dimensional logical hypercube, the two subsystems consisting of 2^{D-1} blocks are simulated independently and then the final layer of CNOTs and in-block operations is combined with the measurement outcomes (the bitstring of interest), which results in a contraction of two $8^{2^{D-1}}$ tensors (see Fig. B.6b). This is a square-root reduction in the memory requirement compared to the full wavefunction simulation which uses $O(8^{2^D})$ space. The ideal XEB value is calculated by sampling bitstrings from the ideal output distribution and then averaging the corresponding probabilities. The bitstrings are sampled using a marginal sampling algorithm, which utilizes the same contraction scheme described above.

We next consider if the finite XEB scores in this problem can be easily “spoofed” by foregoing exact simulation of the implemented circuit and using a classical algorithm with fewer resources, similar in spirit to the algorithm introduced in Ref. ¹⁵⁰. For the circuits studied in Fig. 5 of the main text, containing only a single layer of gates on the hypercube, there is only a single round of CNOTs connecting the two 2^{D-1} -block partitions; thus, removing them from the circuit and sampling from

the two independent halves might not decrease the XEB substantially while reducing the memory requirement to $8^{2^{D-2}}$. In Fig. B.6c, we study the performance of this spoofing attack and find that the obtained XEB rapidly decreases, once additional gate layers are introduced, for a particular extension of our circuit.

The contraction scheme above, utilized both for the ideal simulation and the XEB spoofing, scales exponentially with the number of qubits. However, the exponent is significantly reduced, by utilizing the fact that the hypercube circuits can be naturally partitioned into smaller blocks, with only a single inter-partition layer of CNOT's at the end of the circuit. This simulation method therefore becomes less efficient if we introduce additional CNOT layers (within a single partition) after the inter-partition layer, as we estimate in Fig. 8.5d. Applying $l = \{0, \dots, D - 1\}$ additional intra-partition CNOT layers forces the CNOT tensors in Fig. B.6b to be blocked into groups of 2^l , which results in the execution time to scale roughly as $O(8^{2^l}/2^l)$, where the numerator comes from the tensor contraction complexity and the denominator accounts for the reduced number of contractions due to blocking. The explicit times quoted in Fig. 8.5d as a function of additional CNOT layers are based on the above matrix-multiplication estimate and fitted such that the depth-1 hypercube time matches 1.44 s, which corresponds to our implementation. In practice, the actual run-times might differ due to hardware and software optimization, and additional factors such as the cost of tensor permutations; however, we expect the general trend to hold. Finally, if the 2^l blocked tensors were to be stored directly, the memory requirement of this approach would grow as $8^{2^{l+1}}$, recovering the full 8^{2^D} memory complexity for $l = D - 1$.

In this work, we use these circuits and XEB results for benchmarking our logical encoding, which requires the ability to simulate these circuits. Future logical algorithm experimentation can explore quantum-advantage^{17,45,465,283,447} tests with encoded qubits, as will be detailed in Ref.²³¹.

B.9 PHYSICAL QUBIT CIRCUIT IMPLEMENTATIONS

To compare our logical qubit algorithms with analogous circuits on physical qubits, we work out a concrete implementation of our sampling / scrambling circuits on physical qubits using the same physical gate set, Clifford + T , as used in the logical circuit, which we also then attempt to realize experimentally. We replace each $[[8,3,2]]$ block with a 3-physical-qubit block, decomposing the “in-block” CCZ gates into 6 CNOTs and 7 $\{T, T^\dagger\}$ gates, and implement “transversal” CNOTs directly between the 3-qubit blocks. We note that the CZ can be compiled into the CCZ implementation, but this has minor effect on our analysis and estimates. These physical circuits are complex: 48 qubits with 48 CCZs and 228 two-qubit gates (as realized with our logical qubits) decomposes into an effective 516 two-qubit gates (384 if the CZ gates are compiled into the CCZs). In trying to implement these circuits in practice, the build-up of coherent errors resulted in a vanishing XEB for our physical circuits. These experiments made it clear the logical circuit equivalent was greatly outperforming the physical circuit, thereby providing direct evidence that our logical algorithm outperforms our physical algorithm for this specific sampling circuit.

More quantitatively, with a concrete physical implementation, we calculate an upper-bound by assuming optimistic performance. We assume our best-measured fidelities: SPAM of 99.4%¹²⁹, local single-qubit gate fidelity of 99.91% (Fig. ??), two-qubit gate fidelity of 99.55%¹²⁹, and $T_2 = 2$ s. We then count the total number of entangling gate pulses for the CZ gates, the total number of compiled local single-qubit gates, and the estimated circuit duration, and use this to calculate the estimate presented in Fig. 5f. We further confirm this analysis for small-scale circuit implementations. For a short 3-qubit circuit, we benchmark the XEB for the physical circuit as ≈ 0.87 , below the estimated 3-qubit upper-bound of ≈ 0.92 . We note that in Fig. 5f we plot estimates of physical qubit fidelity and not XEB, but we expect XEB and fidelity to be closely related as discussed previously.

We note several observations made in comparing physical and logical implementations of these

complex circuits. First, empirically it appears that the logical circuit is significantly more tolerant to coherent errors^{56,221,222}, and understanding the manifestations of this is a subject of ongoing investigation. Specifically, it appears that the logical circuit realizes inherently digital operation, where the small coherent errors do not significantly shift / distort the bitstring distribution, but just reduce the overall fidelity^{56,221} (see e.g. the agreement in Fig. B.3a). This is in contrast to the physical implementation, where coherent errors are seen to dramatically alter the shape of the bitstring distribution, e.g. changing relative amplitudes. Second, we note that we optimize our $[[8,3,2]]$ circuits only by optimizing the stabilizer expectation values and not by optimizing the XEB or two-copy result directly. When running complex circuits, the stabilizers serve as useful intermediate fidelity benchmarks, both for optimizing circuit design and ensuring proper execution, especially in regimes where output distributions or other observables cannot be calculated. Overall, we find that these complex circuits appear to perform significantly better with logical qubits than physical qubits.

B.10 TWO-COPY MEASUREMENTS

A powerful method to extract various quantities of interest are Bell basis measurements between two copies of the same state^{103,208,190}. First, we use these measurements to calculate the purity or entanglement entropy of the resulting state^{103,235,43,190}. Measuring the occurrences of the singlet state $\frac{|01\rangle - |10\rangle}{\sqrt{2}}$ ($|11\rangle$ outcome for our measurements after applying the final pair-wise entangling operations) probes the eigenvalue of the SWAP operator $\hat{s}_i$ at a given pair of sites i . This is in turn related to the purity of the state by observing that $\text{Tr}[\rho_A^2] = \text{Tr}[\Pi_{i \in A} \hat{s}_i \rho_A \otimes \rho_A]$ for any subsystem A . Thus, the average purity can be estimated by the average parity $\text{Tr}[\rho_A^2] = \langle (-1)^{\#\text{observed singlets}} \rangle$ within A , and thus also the second-order Renyi entanglement entropy $S_2(A) = -\log_2 \text{Tr}[\rho_A^2]$.

The entanglement entropy calculation only involves the singlet outcomes. By making use of the full outcome distribution, we can also evaluate the absolute value of all 4^N Pauli strings, from a sin-

gle data set, where N is the number of qubits involved in each copy of the state²⁰⁸. More concretely, consider a given Pauli string $O = \prod_i P_i$, where $P_i \in \{X_i, Y_i, Z_i, I_i\}$ are individual Pauli operators on site i (and the identity), and a given observed bit string $\{\vec{a}, \vec{b}\}$, where $\vec{a}, \vec{b}$ label the outcomes in the control and target copy. The rules of reconstructing the Pauli strings through these Bell basis bit strings can be worked out through considering the computational states that the Bell states are mapped to and considering which operators of XX, YY, ZZ have $+1$ or -1 eigenvalue for the various Bell states. We explicitly list the analysis procedure: for Pauli term X_i , we assign parity $+1$ if $a_i = 0$ and -1 otherwise; for Pauli term Y_i , we assign parity $+1$ if $a_i \neq b_i$ and -1 otherwise; for Pauli term Z_i , we assign parity $+1$ if $b_i = 0$ and -1 otherwise; for I_i we assign parity $+1$ always. The contribution of the bit string $\{\vec{a}, \vec{b}\}$ to $|\text{tr}(O\rho)|^2$ is then given by the product of the individual parities.

We can perform the same analysis as a function of the amount of error detection applied. As shown in Fig. B.4a, as more error detection is applied, the distribution of Pauli expectation values that are expected to be zero and nonzero separate apart further. This also provides a natural method to perform error mitigation via zero noise extrapolation: by performing sliding-scale error detection, we can extract the Pauli expectation value squared for groups of Pauli strings with the same expected value, as a function of the logical purity. We perform a linear fit of the Pauli expectation value squared vs. the logical purity, and extrapolate to purity $\text{tr}(\rho^2) = 1$, corresponding to the case of zero noise, to estimate the error-mitigated values. The choice of a linear fit is motivated by the fact that both $|\text{tr}(O\rho)|^2$ and $|\text{tr}(\rho^2)|$ scale with power 2 of the density matrix. We expect that more detailed considerations of the noise model, using knowledge about the weight of each operator, as well as whether detected errors in each shot overlap with a given Pauli operator, can further improve the error mitigation results.

We can also compute measures of distance from stabilizer states, also known as “magic”, using the additive Bell magic measure in Ref. ¹⁹⁶, which only requires $O(1)$ number of samples and $O(N)$ classical post-processing time. To do so, we randomly sample subsets of 4 measured Bell basis bit strings




Figure B.4: Additional Bell basis measurement results. **a**, Histogram of $|\text{tr}(P\rho)|^2$ for all 4^6 Pauli strings P in the 6 logical qubit circuit, as a function of stabilizer postselection threshold (i.e. number of correct stabilizers across the 6x2 logical qubits). Blue (red) indicate Pauli strings that are expected to have $|\text{tr}(P\rho)|^2 = 0.0625$ (0). The separation between the histograms improves as more postselection is applied. **b**, Signal-to-noise (purity divided by statistical uncertainty of purity) as a function of sliding-scale error detection (converted into accepted fraction), for the 12 logical qubit two-copy measurements, where subsystem size 1 indicates a single logical qubit in one copy, and subsystem size 12 indicates all logical qubits. For subsystem size 1, the signal-to-noise ratio gets worse as data is discarded, since the signal does not change (maximally mixed) but the number of repetitions decreases. In contrast, for the global purity, the signal-to-noise increases as near-unity purities are faster to measure⁶⁷. **c,d**, Entanglement entropy when analyzing the circuit as a physical Bell basis measurement as opposed to a logical Bell basis measurement. For logical entanglement entropy calculations we average over all possible subsystems of that given subsystem size, which we find behaves very similarly to e.g. contiguous subsystems due to the high-dimensional hypercube connectivity. In the physical qubit entanglement entropy calculations, we randomly choose from the possible subsystems as there are many. **c**, 6 logical (16 physical) qubits per copy; **d**, 12 logical (32 physical) qubits per copy. The finite sampling imposes a noise floor for very high entanglement entropy values. **e**, Entanglement entropy measurements as a function of logical subsystem size. **f**, Logical circuits used for benchmarking magic. For 1 CCZ, we include U_1 and omit U_0 ; for 2 CCZ, we include U_0 and omit U_1 ; for the 3 CCZ, we include both U_0 and U_1 .

$\mathbf{r}, \mathbf{r}', \mathbf{q}, \mathbf{q}'$ and calculate their contribution to the Bell magic using the check-commute method of Ref. ¹⁹⁶: $\mathcal{B} = \sum_{\substack{\mathbf{r}, \mathbf{r}', \mathbf{q}, \mathbf{q}' \\ \in \{0,1\}^{2N}}} P(\mathbf{r})P(\mathbf{r}')P(\mathbf{q})P(\mathbf{q}') \|\sigma_{\mathbf{r} \oplus \mathbf{r}'}, \sigma_{\mathbf{q} \oplus \mathbf{q}'}\|_{\infty}$, with $\|\sigma_{\mathbf{r} \oplus \mathbf{r}'}, \sigma_{\mathbf{q} \oplus \mathbf{q}'}\|_{\infty}$ being 0 when the two Pauli strings commute and 2 otherwise. $\mathbf{r} \oplus \mathbf{r}'$ denotes bitwise XOR between the two bitstrings. $P(\mathbf{r})$ is the probability of observing bitstring $\mathbf{r}$. The Pauli string $\sigma_{\mathbf{r}}$ is of length N , and the i th element is I, X, Z, Y when the target and control qubit at site i read 00, 01, 10, 11, respectively. We convert this result to additive Bell magic via the formula $\mathcal{B}_a = -\log_2(1 - \mathcal{B})$. We use approximately 10^7 samples to estimate the additive Bell magic for each dataset. The results for the estimated additive Bell magic as a function of the number of non-Clifford gates applied (circuits shown in Fig. B.4f) are shown in Fig. 6c. These results additionally leverage the purity estimates in the same dataset, which is used for error mitigation as described in Eq. (13-15) of Ref. ¹⁹⁶. All additive Bell magic data shown is with full error detection applied.

The same experiments we perform here can also be interpreted as a physical Bell basis measurement. Using this insight, in Fig. B.4c,d, we show the entanglement entropy for different subsystem sizes, when analyzing the data as physical Bell pair measurements, and applying different levels of stabilizer-based postselection. Interestingly, the full-system parity when postselecting on all stabilizers being correct is identical when analyzing the outcomes as either a physical or logical circuit. This is because in this limit, the results of the physical circuit analysis can be viewed as taking the (imperfect) logical state and running a perfect encoding circuit, hence giving identical results.




Figure B.5: Surface code preparation and decoding data. **a**, Surface code stabilizers for the two independent $d = 7$ codes following state preparation. The entire movement circuit corresponding to the transversal CNOT is implemented, and the transversal entangling gate pulse is simply turned off. The mean stabilizer probability of success across the 96 total stabilizers is 83%. The high probability of stabilizer success of the two independent codes in both the X and Z bases shows that topological surface codes were prepared (and Fig. B.1 shows that they were preserved during the transversal CNOT). We note that physical fidelities were slightly lower during this measurement due to calibration drift and so these results slightly underestimate performance relative to the data in Fig. 2 and Fig. B.1. **b**, Logical Bell pair error while optimizing the decoder by (inversely) scaling the weights of the inter-logical edges and hyperedges that connect the stabilizers of the two logical qubits (higher values correspond to lower pairing weights). More concretely, the probability p of the error mechanism corresponding to the inter-logical edges / hyperedges is scaled and the weights are calculated as $\log((1 - p)/p)$. Qualitatively, optimizing this scaling value optimizes with respect to the probability that errors are before or after the transversal CNOT, as errors before the CNOT will lead to correlations between the two logical qubits, corresponding to the inter-logical edges. As the decoder is optimized by tuning the inter-logical scaling factor, the performance for all three code distances improves, and the larger code distances improve faster when approaching the optimal decoding configuration, as expected. This data is consistent with the decoder being properly optimized for all three code distances, consistent with the fact that our improvement with code size does not originate from suboptimal decoder performance for low distance. Note the y-axis is log-scale. **c**, Logical Bell pair error when using (black) and not using (gray) the ancilla stabilizer measurement values, as a function of the scaling of the inter-logical edges and hyperedges that connect the stabilizers of the two logical qubits. The ancilla measurements contribute to the correction procedure, and contribute more for smaller values of the inter-logical scaling as they correspond to errors that happen before the transversal CNOT. 0x inter-logical scaling corresponds to conventional decoding within the two independent surface codes. For the 1x inter-logical scaling plotted here, the $d = 7$ inter-logical scaling parameter is chosen slightly different than in Fig. 2d in order to have consistency across the three code distances (which produces measured values within errorbars).




Figure B.6: Theoretical exploration of hypercube IQP circuits. **a**, Anti-concentration property of our circuits. The circuit is said to be anti-concentrated if its output distribution is spread almost uniformly amongst all outcomes, without the probability being concentrated on a subset of bitstrings. This property is crucial for many proofs of classical hardness^{59,52} and, thus, it is desired for our sampling circuits to anti-concentrate. The plot shows that the output distribution of random hypercube circuits (randomized in-block operations and randomized control/target in out-block CNOT layers) anti-concentrates as the dimension of the hypercube is increased and the XEB (which captures the output collision probability) converges to the uniform-IQP value of 2 (here using Clifford circuits; i.e., circuits composed only of random CZ and Z gates)⁵⁹. This suggests that sampling from the ideal output distribution can be classically hard. In general, the hypercube IQP circuit ensemble converges to the uniform IQP ensemble in total variation distance as the depth and hypercube dimension are increased²³¹. The specific circuit instances implemented in the experiment also anticoncentrate quickly with increasing hypercube dimension. **b**, A single layer of the hypercube circuit admits an efficient tensor-network contraction scheme, which allows us to evaluate the ideal and experimental XEB values. The final out-block CNOT layer is immediately followed by the measurement, which can be incorporated into a non-unitary tensor that is contracted between the two halves of the system (controls and targets of the final CNOT layer). This contraction scheme reduces the memory requirements to half the system size, which enables bitstring amplitude evaluation for the 48-qubit experiment. This simulation approach can be made significantly more expensive by applying additional out-block operations within the two subsystems, forcing the blocking of the intra-partition tensors, which increases the memory and runtime requirements (Fig. 8.5d). **c**, To understand the effects of finite XEB on required classical simulation time, we explore if our circuit families can be “spoofed” with cheaper, approximate simulation that achieves moderately high XEB scores¹⁵⁰, studied here for a 24-qubit system with full state-vector simulation. The spoofing algorithm works by independently sampling from the two halves of the system (two groups of 12 qubits), effectively removing the final layer of CNOTs. This further reduces the simulation complexity, as each of the halves can, in principle, be independently simulated with the efficient approach from **b**. The plot shows that the spoofed XEB for the 24-qubit non-Clifford circuit can be exponentially reduced by extending the circuit with additional gate layers (similarly to the approach used to decrease the performance of the efficient hypercube contraction), for a particular extension of our circuit. This result shows that future work can consider adding additional CNOT layers into these circuits to demonstrate quantum advantage (in the presence of finite experimental noise).



Supplementary information about
ingredients for universal fault-tolerant
quantum computing

Additional information related to this Appendix can be found in Ref.⁴².

C.1 ANALYSIS OF ERROR CORRELATIONS

Correlations in errors, in either space or time, can have important implications on quantum error correction. Here we explain various correlation analyses in our system.

De-correlation of global coherent errors by projective measurement. Parallel control enables us to, for example, realize a transversal entangling gate with a single global pulse of our entangling laser⁴¹.

One may be concerned that such a global control can lead to globally correlated errors that can affect error correction performance. However, error correction natively de-correlates such errors.

Consider a code block of qubits with X and Z stabilizers. Applying a global θ will map each of the X operators to $\rightarrow (X + i\theta Y) = X \cdot (1 - \theta Z)$. Consequently, measuring the X-basis component of this qubit will probabilistically lead to a Pauli Z error on this site with probability θ^2 . Note that, for global rotation θ , the logical operator $X_L = XXXX\dots$ maps to $\rightarrow (X + i\theta Y)(X + i\theta Y)(X + i\theta Y)(X + i\theta Y)\dots = XXXX\dots + i\theta YXXX\dots + (i\theta)^d YYYYY\dots$. As such, for small θ , logical rotations are exponentially suppressed with the code distance d . As such, even though all the physical qubits receive a global rotation θ , the logical qubit state does not receive that same rotation, and after syndrome measurements these errors are converted into incoherent-type errors and can be corrected. This is the basis behind the observed suppression in Figure 1. In ED Figure C.3b, we further show that the error correction prevents an unintended logical rotation. The logical rotation here is even further suppressed by the random stabilizer signs (below).

Role of stabilizer signs under coherent errors. Stabilizer signs affect the logical qubit's response to global coherent rotations. For transversal non-Clifford gates, deterministic stabilizer eigenvalues (e.g. = +1) are necessary to correctly implement the logical gate (Fig. 4). In contrast, Clifford circuits allow the eigenvalues to be either +1 or -1, as the signs can be simply tracked through the circuit, giving freedom to engineer how coherent errors interfere. For example, choosing negative signs can generate decoherence-free subspaces¹¹¹ and give greater robustness of the logical operator against

coherent errors. The same principle also suppresses logical coherent errors during computation⁵⁶.

In particular, stabilizer measurement projects the logical state onto a specific stabilizer configuration with random ± 1 values, which corresponds to a random configuration of physical X and Z flips which do not commute with the coherent rotation. On top of the exponential suppression of logical coherent errors, this further results such that the specific rotation angle of scale $\sim \theta^d$ is random on each shot, effectively turning these again into incoherent errors on the logical level.

Surface code measurements. In our repeated quantum error correction on the surface code, we search for unexpected error correlations by plotting the distribution of detector errors in a shot. Crucially, we find these are closely consistent with the expected distribution as seen by Clifford simulations that assume uncorrelated one- and two-qubit errors. Although this data is only composed of approximately 10^5 detector rounds (14855 shots, 96 detectors per shot across 5 rounds), it is nevertheless indicative of the absence of such events. We have yet to observe error burst events such as those observed in solid-state systems⁴.

Logical teleportations and deep-circuit measurements. In a physical system, diverse physical errors and imperfections can cause complex correlations. For instance, a leakage event can lead to complex correlations that - without its knowledge - can greatly affect QEC performance. As discussed in the next section, incorporating logical teleportations in an architecture can ensure such errors are removed. In Figure 6 of the main text we verify that such teleportations indeed rapidly remove errors and ensure that errors are not correlated in either time or space. It is important that the atomic qubits are re-initialized properly for this to work. In ED Figure C.1c,d we plot the correlations when turning off cooling and sorting and find that correlations in such a case do not rapidly decay.




Figure C.1: Entropy in deep circuits. **a**, General hypercube encoding. The same entangling circuit structure is combined with programmable input physical states to prepare $[[7,1,3]]$, $[[15,1,3]]$ and $[[16,6,4]]$ codes (members of the family of quantum Reed-Muller codes¹⁶⁶). For the punctured codes, either the top or bottom qubit is removed. **b**, Turning off entropy removal mechanisms in the 2D $[[7,1,3]]$ cluster state circuit. In the absence of re-cooling or refilling of loss, physical stabilizer correlations persist in time. Lost atoms are assigned as qubit state 0 for this plot, and only correlations between codes constructed from the same atomic qubits in every other layer are shown. **c,d** Circuit structure and physical error correlations. By replacing CZ gates by CNOT gates in the 1D $[[7,1,3]]$ cluster state circuit, physical errors can propagate beyond a single layer, extending the stabilizer correlations. The product of adjacent stabilizers commutes with these propagated errors, recovering the rapid decay in physical correlations.

C.2 DECODING

MLE and error-model tuning. To decode the surface code experiments in Figs. 1-3, we use the delayed-erasure MLE decoder described in Ref.²³, augmenting the MLE decoder developed in Ref.⁷⁰ to leverage loss information. In particular, the MLE decoder takes as input the stabilizer measurements and the probabilities of the physical error sources in the circuit, and outputs the most likely combination of errors consistent with the syndrome. We construct the circuit error model using Stim¹⁵⁶ to initially contain information about the Pauli error sources in the circuit, then update it for each shot to reflect the detected atom losses. In particular, after an atom is lost, all subsequent gates are canceled, generating different potential errors depending on when the loss occurred. We therefore consider all potential locations a qubit loss could have originated (e.g., initialization, gates, movement, or idling prior to measurement), then add each of the resulting error patterns and their probabilities to the error model for that shot. Errors producing the same syndrome are combined into a composite error mechanism and their probabilities are correspondingly reweighted, as in

Ref.¹⁵⁶. Note that this process explicitly accounts for both propagated Pauli errors from the gate cancellations and the invalidation of stabilizers, which are handled using superchecks.

To optimize the performance of the MLE decoder, we fine-tune the probabilities of different error sources in the circuit error model. In particular, we associate each physical operation with both a Pauli and loss error rate. The error probabilities in these channels are then treated as variables which we optimize using the covariance matrix adaptation evolution strategy¹⁹² in order to minimize the logical error rate on a dataset of approximately 10000 shots (different shots from the final dataset used for evaluating the fidelity).

To quantify the benefit of using loss information, the ‘bare MLE’ decoder does not update the circuit error model based on the loss information, and assigns each loss event to a $|0\rangle$ measurement. We find the loss information improves the measured $d = 3/d = 5$ error ratio from 1.24(5) to 1.69(8).

Finally, we quantify the confidence of the MLE correction for each shot by comparing the probability of the most likely error p_0 and the probability of the most likely error which gives the correction to the logical Pauli observable p_1 ^{214,49,405,158}. The more similar these two error probabilities are, the less confident the decoder is in its correction. We can therefore postselect on increasing $p_0/(p_0 + p_1)$ to improve the accuracy of the results, which we use in Fig. 3d when studying repeated logical gates.

ML decoder - surface code. We employ a fully connected neural network to decode measurement outcomes from the Fig. 2 surface code experiment using machine learning (ML)^{250,19,155,30}. The decoding task is formulated as a supervised binary classification problem: the input features are measurement outcomes from the experiment, and the output is a label indicating whether the initial state was $|0_L\rangle$ or $|1_L\rangle$. The ML architecture is a fully connected feedforward network comprising four linear layers, each followed by batch normalization and a Gaussian Error Linear Unit (GELU) activation, as illustrated below:

```

decoder = nn.Sequential(
    nn.Linear(input_size, 1024),
    nn.BatchNorm1d(1024),
    nn.GELU(),
    nn.Linear(1024, 512),
    nn.BatchNorm1d(512),
    nn.GELU(),
    nn.Linear(512, 256),
    nn.BatchNorm1d(256),
    nn.GELU(),
    nn.Linear(256, 1),
    nn.Sigmoid()
)

```

Training proceeds in three stages: raw training, ensembling, and fine-tuning.

Raw training — We begin by training the decoder on simulated data generated through circuit-level simulations that incorporate both Pauli and loss errors. Measurement outcomes take values of 0, 1, or 2, corresponding to the qubit being in the $|0\rangle$ state, the $|1\rangle$ state, or being lost, respectively. These are one-hot encoded, so the feature vector of a given shot is $3 \times (\# \text{ of measurements})$. To create balanced training data, random software flips are applied with probability $1/2$ along the relevant logical operator, yielding ensembles of $|0_L\rangle, |1_L\rangle$ for the Z memory and $|+_L\rangle, |-_L\rangle$ for the X memory. In addition to the raw $\{0, 1, 2\}$ measurement values, we provide the neural network with calculated detector outcomes and logical operator values. These additional features help the model learn from structured correlations in the data. Detector values are computed as binary parities (0 or 1) over specified stabilizer regions; if a measurement gives a loss (2), it is assigned a value of

o when computing detector parities. Logical operator values are calculated along each row or column, depending on the basis. We use a hidden layer size of 1024, the Adam optimizer with an initial learning rate of 10^{-3} , and a weight decay of 10^{-2} . The learning rate is decreased by a factor of 0.3 if the validation loss does not improve for 10 epochs. Training is performed independently for 10 total experimental configurations: two with code distance $d = 5$ (in the Z and X bases) and eight with $d = 3$ (covering four spatial quadrants in both bases). In the pre-training phase, each model is trained on 200 million simulated shots and validated on 20 million simulated shots. We find that decoder performance is largely robust to small perturbations in the error model, and thus precise tuning of simulation parameters is not necessary. For a batch size of $\sim 10^4$ shots, the inference time per shot is $0.33 \mu\text{s}$ on a GPU (NVIDIA-A100).

Ensembling — To account for training variability and enhance robustness, we repeat the full training procedure with 10 different random seeds, resulting in 10 independently trained models per experiment. These are ensembled together by computing the geometric mean of their output probabilities. The resulting ensembled ML decoder achieves a logical error per round (LEPR) of 0.78(4)% for $d = 5$ and 1.37(3)% for $d = 3$.

Fine-tuning — To improve decoding performance, we fine-tune each pre-trained decoder on experimental data taken from designated training sets (independent of the final dataset). For the $d = 5$ decoders, we fine-tune on approximately 37000 shots per basis. For the $d = 3$ decoders, we use approximately 2500 shots per basis, per quadrant. The neural network architecture remains unchanged, and fine-tuning is performed using the Adam optimizer with a learning rate of 10^{-3} and a weight decay of 8×10^{-2} . The resulting ensemble of fine-tuned ML decoder achieves a LEPR of 0.71(4)% for $d = 5$ and 1.33(4)% for $d = 3$.

Hybrid — When comparing the MLE and ML decoders, we find they do not predict the same logical state on all shots and, in particular, differ on shots where one of the decoders has low confidence in its prediction. To further enhance performance, we therefore construct a hybrid decoder that

combines the output confidences of the ensembled ML decoder with those from the delayed-erasure MLE decoder, where the MLE confidence is derived from comparing the probabilities of the most-likely error and the most-likely error which gives the opposite logical outcome. The final prediction is given by the weighted geometric mean of the two confidence values with weights of 0.4 and 1 for the MLE and ML, respectively. This results in a final value for the reported LEPR of 0.62(3)% for $d = 5$ and 1.33(4)% for $d = 3$, which corresponds to the ML with loss decoder reported in Fig. 2. *MLE decoder - lattice surgery.* In the lattice surgery experiment (Figs. 3c,d), we perform a joint ZZ measurement using additional stabilizer checks along the common vertical edge between the two surface codes (“seam”). We start with both codes prepared in $|+_L\rangle$ and perform two rounds of stabilizer checks on the effective $d \times 2d$ surface code lattice, measuring both the codes’ stabilizer checks and the new seam checks in each round. To measure the ZZ parity of the resulting logical Bell state, with the delayed-erasure MLE decoder^{23,70}, we use two decoding procedures. First, we use only the ancilla measurements to obtain the result of the lattice surgery $Z_1^L Z_2^L$ measurement given by the product of the seam Z stabilizers. Second, we obtain $Z_1^L Z_2^L$ directly from the data qubit measurements, using the prior ancilla measurements in decoding. This measures the ZZ parity of the logical Bell state obtained using lattice surgery. Note that the seam checks from the final data qubit measurement are not included. A shot is counted as an error, $Z_L^1 Z_L^2 = -1$, if these two decoding procedures disagree.

To obtain the XX Bell state parity, we measure all data qubits in the X basis and decode the joint $X_L^1 X_L^2$ operator spanning both codes. The final logical error probability is given by the mean of the XX and ZZ parities.

ML decoder - deep circuits. To decode the 1D and 2D cluster states of logical $[[7,1,3]]$ and $[[16,6,4]]$ codes in Fig. 6, we employ a convolutional neural network (CNN). Since error correlations in the cluster state do not propagate beyond two CZ gates (see Fig. 6e), a convolutional window of size 3 is sufficient to capture the relevant correlations. The decoder architecture comprises three compo-

nents: an encoder, a convolutional block, and a readout module. Both the encoder and readout are constructed from linear layers interleaved with GELU activations, with `hidden_size = 128`.

```
encode = nn.Sequential(  
    nn.Linear(input_size, 1024),  
    nn.GELU(),  
    nn.Linear(1024, 512),  
    nn.GELU(),  
    nn.Linear(512, 256),  
    nn.GELU(),  
    nn.Linear(256, hidden_size)  
)  
  
readout = nn.Sequential(  
    nn.Linear(hidden_size*8, 512),  
    nn.GELU(),  
    nn.Linear(512, 256),  
    nn.GELU(),  
    nn.Linear(256, 128),  
    nn.GELU(),  
    nn.Linear(128, out_size)  
)
```

The convolutional block applied between the encoder and readout modules is defined as:

```
conv = nn.Sequential(  
    nn.Conv2d(hidden_size, hidden_size*2,  
              kernel_size=3, padding='same'),
```

```

nn.GELU() ,
nn.BatchNorm2d( hidden_size * 2 ),
nn.Conv2d( hidden_size * 2, hidden_size * 4,
           kernel_size = 3, padding = 'same' ),
nn.GELU() ,
nn.BatchNorm2d( hidden_size * 4 ),
nn.Conv2d( hidden_size * 4, hidden_size * 8,
           kernel_size = 3, padding = 'same' ),
nn.GELU() ,
nn.BatchNorm2d( hidden_size * 8 ),
)

```

For 1D cluster state decoders, we replace the 2D convolutions with 1D convolutions and omit the batch normalization layers.

Training is performed using circuit-level simulations. The decoder is tasked with inferring the signs of the logical cluster state stabilizers, which are of the form X on a given qubit and Z on its neighbors. By performing measurements in alternating X and Z bases, half of the stabilizers can be reconstructed. The remaining stabilizers are recovered by repeating the experiment with the measurement bases swapped. The decoder predicts the stabilizer signs by inferring the initial state of the qubits measured in the X basis. All logical qubits are initialized in the $|+_L\rangle$ state, and software logical Z flips are applied with probability $1/2$ to those measured in the X basis, to generate a balanced training set.

The decoder input includes the raw measurement outcomes (0, 1, or loss), detector values computed from the measurements, and the raw logical operator values, similar to the input format used in the surface code decoder. We train four distinct decoders: one for each combination of code type

($[[7,1,3]]$, $[[16,6,4]]$) and cluster state geometry (1D, 2D). Each model is trained on over 100 million simulated shots.

For further details on this decoder architecture, and on ML-based decoders for general quantum algorithms, see Ref. ⁵⁰.

C.3 BENCHMARKING SURFACE CODE PERFORMANCE

NZNZ stabilizer gate pattern. Here we describe the effective distance, defined as the minimum number of physical errors required to create a logical error, of d rounds of repeated syndrome extraction using alternating “N” or “Z” movement patterns (Fig. C.2d). By alternating gate orderings, the effective distance is close to the optimal value. To see this, note that without alternating orderings, the effective code distance in the rotated surface code is reduced by a factor of two due to hook errors¹¹⁷ (from a theoretical perspective - see discussion at end). A hook error is a physical error on the ancilla qubit halfway through the stabilizer measurement cycle that propagates onto two data qubits oriented parallel to the corresponding logical operator (e.g., X_L for a physical X error). One of these propagated data qubit errors is immediately detected, while the other is detected in the following round by the next-nearest stabilizer along the direction of error propagation. As a result, if the same gate ordering is used for each round of stabilizer measurements, a sequence of $\lceil \frac{d+1}{4} \rceil$ hook errors, one occurring in each round along the direction of error propagation, can generate a logical error upon correction. This issue is circumvented by alternating gate orderings between rounds, as only every other round has the unfavorable propagation. In this case, $\frac{d-1}{2}$ physical errors on consecutive rounds are needed to generate a logical error.

In our experiments in particular, we choose an ordering of $NZ Z_r N_r$, where r represents performing the reverse ordering (see Fig. C.2d). In addition to such structuring helping preserve fault-tolerance against hook errors, we also note that the dominance of Z -type errors means that most er-




Figure C.2: Additional data for repeated QEC characterization circuit. **a**, Processor layout for repeated QEC on a $d = 5$ surface code. The data qubits and one ancilla block are located in the entangling zone (top) and four additional ancilla blocks are in storage (bottom); one ancilla block is unused in the four-round circuit. **b**, Processor layout for repeated QEC on a $d = 3$ code in one of four possible quadrants. **c**, Circuit for four rounds of QEC. For the XZZX rotated surface code⁵¹, $Y(\pi/2)$ gates are applied to one data qubit sublattice (A or B) for preparing and measuring in the X or Z basis. The equivalent circuit is obtained for stabilizer measurement of the CSS rotated surface code upon compiling $Y(\pi/2)$ gates. **d**, Stabilizer gate ordering. The same pattern is used globally in each round. **e**, Detector error probability for $d = 5$. Faint dashed curves correspond to the 24 individual detectors (12 for the first and last rounds), and solid curves are the mean of all deterministic detectors. In this subfigure only, to illustrate the supercheck error distribution, we assign the error of each supercheck to all detectors from which it is composed, and plot the resulting effective detectors; the overall mean is the average of these individual detector values. The first round has lower error due to the neighboring transversal state preparation. The best detector over the central three rounds after loss postselection has 27% lower error than the overall mean; one atom has anomalously high error in the Z basis. **f**, Comparison to simulation. Left: detector error probability, converting loss to qubit state 0. Right: detection correlation matrix p_{ij} ¹⁷⁰. Both metrics show good agreement between simulation and experiment in the structure and magnitude of the errors. **g**, CZ gate fidelity measured via randomized benchmarking¹²⁹. Infidelity due to Rydberg P states is removed by leaving sufficient time or distance between gates. Sparse corresponds to 2x larger separation between gate sites.

rors do not lead to propagated errors between the middle two CZ gate layers. Due to these reasons, in simulations we do not observe having spatially alternating “N” and “Z” patterns helps performance (not plotted).

Simulations. We perform simulations using the Stim simulation package¹⁵⁶. We sample both Pauli errors and qubit losses. Pauli errors are generated using Stim’s sampling routines, based on circuit-level noise models. Qubit losses are sampled according to the loss probabilities associated with each instruction, and when a loss occurs, subsequent gates acting on the lost qubit are removed to reflect the absence of the qubit. The simulations detect $\{0, 1, \text{loss}\}$ during qubit readout, like in our experiments. For each set of physical parameters, we estimate the logical error rate via Monte Carlo sampling. Logical errors are declared when the decoder’s prediction for the logical observable differs from the true value. See Supplementary Information for details of the noise model, as well as the discussion below.

Analysis of below-threshold performance for deep circuits. In Figure 2 we perform four rounds of repeated QEC as a benchmark. However, increasing the circuit depth can affect the threshold in various ways, depending on the particular circuit. Fig. C.4a shows how the LEPR ratio r changes for a single logical qubit as we increase the number of QEC rounds using a theory error model, showing a roughly 17% decrease in r from 4 rounds to 20. Similarly, Fig. C.4b plots the same quantity for a single logical qubit with an approximate experimental error model, showing an analogous 9% decrease in r from 4 rounds to 50. Further, by interspersing 1 transversal gate every 1 QEC round under an approximate experimental error model, we find the ratio r changes by 2% at 25 QEC rounds. Similarly, prior work with a theory error model has shown that the threshold can change by $\approx 10\%$ with 1 gate per QEC round⁷⁰. These simulations indicate that the benchmark studied in Fig. 2 is representative, but depending on context, can be different on the scale of $\approx 15\%$ for deep circuits. We note, however, that in transversal architectures, the prevalence of logical gate teleportations (e.g., in magic state distillation and angle synthesis) makes it such that there are typically only several sta-

bilizer measurement rounds before transversal measurement.

Our benchmark results are comparable to those in Ref. ⁴. For instance, although a one-to-one comparison is not direct due to the presence of loss information, using the supercheck metric reveals a 9.04% mean detector error, comparable to the 8.5-8.7% mean detector error in Ref. ⁴.

C.4 ADDITIONAL EXPERIMENT AND DATA ANALYSIS DETAILS

Figure 1. In Fig. 1d, we prepare either $|+_L\rangle = |+\rangle^{\otimes 25}$ or $|0_L\rangle = |0\rangle^{\otimes 25}$ and apply up to five rounds of stabilizer measurement followed by measurement in the X or Z basis, respectively. A global $Z(\theta)$ rotation is applied to the data qubits at every gate layer (20 time-steps in total). For fewer than five stabilizer measurement rounds, the relevant CZ gates are removed but single qubit rotations still applied. 1 QEC round has gates in the first round only, 2 QEC rounds has gates in the first and fourth rounds, and 5 QEC rounds has gates in all five rounds. The data is averaged over both initial states and uses MLE decoding with a 50% acceptance fraction for visual clarity, as well as pre-selection on perfect initial qubit filling. The right plot uses an injected error of $\theta/2\pi = 0.016$ and additional error rates are shown in Fig. C.3b.

Figure 2. No postselection is used in the analysis of the surface code. Pre-selection of initial qubit rearrangement (standard in the literature) is used. Data in Figs. 2b-d are averaged over $|+_L\rangle$ and $|0_L\rangle$, and the distributions in Figs. 2e,g aggregate the two bases. Figs. 2b,c plot the detector error probability averaged over all rounds. Fig. 2b uses the same data set with shots binned according to data qubit loss. The four metrics are (i) ‘bare’, where loss is converted to qubit state 0 (ii) ‘detect loss’, where projective measurements whose value is ‘loss’ are not counted erroneously (iii) ‘supercheck’, where stabilizers with a lost data qubit are formed into superchecks for all prior rounds, and (iv) ‘postselected’, where detectors involving any lost atoms are ignored. The plots show the mean error of all deterministic detectors (96 per basis). The supercheck error is calculated over all samples

per round per basis, and the mean of these 8 values plotted. Superchecks paired to the boundary are removed from the averaging as these return no error by construction; if included, the mean error decreases from 9.0% to 8.8%. The contribution of each supercheck is normalised by the supercheck weight, e.g., a weight-6 supercheck contributes an error of $4/6$ (to account for the greater amount of information in the check - e.g., multiplying checks even in absence of loss raises the detector error without reducing the amount of information). Without reweighting, the error probability increases to 9.6%.

In Figs. 2d,e, the logical error per round is calculated as $\text{LEPR} = 0.5(1 - (1 - 2p_L)^{1/r})$ where p_L is the final logical error after $r = 4$ rounds, same as the definition in Ref.⁴. The $d = 5$ dataset contains 9021 shots in the X basis and 5834 in the Z basis. The $d = 3$ dataset contains 2523 shots for X and 2534 for Z (on average per quadrant). To make a $d = 3$ surface code in each of four possible quadrants, we only remove atoms and do not modify the circuit. The specific circuit for the repeated stabilizer measurement is shown in ED C.2c. Not shown are local $Y(\pi)$ and $Y(\pi/2)$ gates on the boundary ancillas (see Supplementary Information and Stim circuit). Additionally, local detunings²⁸⁹ are applied to the lowest row of gate sites (where there are only isolated ancilla qubits) to mitigate inhomogeneity in the 1013-nm lightshift during entangling gates.

In Fig. 2f, the error budget shows the contributions to the detector error (with loss detection) and is obtained by removing error sources individually from the simulation error model. We obtain a similar error model breakdown by simulating the relative contribution to the logical error. Fig. 2g plots the detector distribution with loss detection. Fig. C.2f compares the bare detector error to simulation.

Figure 3. All data uses MLE decoding and is pre-selected on perfect initial qubit filling. In Figs. 3c,d, we prepare logical Bell states using either transversal gates or lattice surgery, and measure the mean error in the resulting XX and ZZ parities. The error per logical operation is defined as $\varepsilon = 0.5(1 - (1 - 2p_L)^{1/3N})$ for N transversal gates per round and Bell state infidelity p_L , and $\varepsilon = p_L$ for the lat-

tice surgery logical product measurement. In Fig. 3c, the transversal CNOT is shown for 3 CNOTs per QEC round, and the injected measurement error is applied in-software with probability p independently to every ancilla qubit. An acceptance fraction of 1 is used for the transversal CNOT plots unless otherwise stated. In Fig. 3d, the lattice surgery point uses error detection on the middle three ancillas each having the same value in both rounds of stabilizer measurement, to compensate for having fewer than d rounds of repeated syndrome measurement for this result. We find in numerical simulations using our experimental error model that the optimal number of QEC rounds for this circuit is approximately 3 (as opposed to 5), and that by using error detection with 2 rounds we recover a similar performance to this optimal value found in numerics.

Figure 4. To modify the stabilizer signs in Fig. 4a, local π pulses are applied at the end of the encoding circuit. “Negative” stabilizers corresponds to flipped qubits on the four corners of the Reed-Muller tetrahedron. We use a lookup table for decoding and plot all three 3D color code curves with an acceptance fraction of 46% and the 2D color code with 74%, corresponding to a rescaling by the number of physical qubits in the code. For postselection, the shots are ordered by the weight of the detected error. The curves are further normalized to highlight key features, with unnormalized data shown in ED Figure ??a. Fig. 4c uses error detection and plots the angles for $\leq NT$ gates. All plots are postselected on no loss and perfect initial qubit filling.

Figure 5. In Fig. 5c, we study the atom temperature and loss as a function of cycle using the circuit for state preparation of Steane codes (Fig. 6b) with only entangling gates removed. In the fifth cycle, we turn off all imaging and cooling light. For comparison, the same measurements are repeated with conventional 3D PGC imaging and cooling in place of the local techniques. To extract the atom temperature shown in the upper panel, we use a drop-recapture measurement after N cycles and fit the resulting loss to a Monte-Carlo simulation⁴²⁶. Shaded regions indicate the range of fitted temperatures due to uncertainty in trap parameters.

Figure 6. The $[[7,1,3]]$ and $[[16,6,4]]$ codes in Fig. 6, as well as the $[[15,1,3]]$ in Fig. 4, are members

of the family of quantum Reed-Muller codes based on the hypercube encoding circuit illustrated in Fig. 10a¹⁶⁶. For each code, a different pattern of local $Y(\pi/2)$ pulses is applied while the entangling gate structure is the same; for the 2D $[[7,1,3]]$ code, the fourth layer of gates is turned off. In Fig. 6b, groups of sixteen independent $[[7,1,3]]$ codes are prepared in parallel in each time layer, repeated for 27 layers. The stabilizer error probability as a function of layer is plotted for no loss in the code block.

To characterize the propagation of physical and logical information in deep circuits, we further entangle the codes into 1D and 2D cluster states. Starting with two groups of logical qubits, group A and group B in Fig. 6a, these are entangled to form the first two time layers of the cluster state. Due to the local entanglement structure of a cluster state, group A undergoes no more entangling gates and is idle up until its measurement (in the appropriate basis), and thereby the measurement can be performed and the same physical qubits re-used to form the third layer of the cluster state (in typical MBQC fashion). Group B can then be measured and re-used to form the fourth layer of the cluster state, and so on. This alternating structure is typical in MBQC using cluster states³⁴⁵.

The physical correlations in Figs. 6c,d,g are calculated as the covariance between errors (stabilizer = -1) on the same stabilizer between codes at different coordinates in the cluster state. The covariance is then averaged across all co-propagating cluster states and the different stabilizers (three for $[[7,1,3]]$ and five for $[[16,6,4]]$). The logical correlations in Figs. 6c,g are calculated as the appropriate product of cluster state stabilizers between the two target coordinates (cluster states have stabilizers corresponding to $X_i \prod_j Z_j$ where i is a specific site and j is its neighbors). For example, we define $\langle Z_0 Z_4 \rangle \equiv \langle (Z_0 X_1 Z_2) \cdot (Z_2 X_3 Z_4) \rangle$. The single-qubit expectation values $\langle Z_i \rangle$ are calculated using a lookup table decoder for $[[7,1,3]]$ and raw values for $[[16,6,4]]$. Due to the underlying assumption of time- and space- invariance, i.e., that correlations depend only on relative coordinates, we truncate time layers where the reservoir begins to be depleted and this assumption breaks down. This corresponds to 16 layers for Fig. 6c, 13 layers for Fig. 6d (correlations plot only), and 12 layers for Fig. 6g.

This has only a small effect on the measured logical correlations but otherwise leads to a longer-tail of physical correlations because atoms are not properly refilled once the reservoir begins depleting.

All logical operators in Figure 6 are decoded with machine learning which directly predicts the cluster state stabilizers. The 2D cluster state stabilizers in Fig. 6d use a global acceptance fraction of 0.24%. In Figs. 6c,g we instead use a global confidence threshold for each curve, which is then converted to a mean acceptance fraction. In this way, each curve corresponds to a constant effective error rate for the logical operator independent of its weight, resulting in a reduced acceptance fraction for higher weight operators. The confidence for products of the weight-3 logical stabilizers is given as the geometric mean of the constituent confidences. Fig. 6g uses a mean acceptance fraction of 3.4% (same data for both curves). The 2D $[[16,6,4]]$ cluster state in Fig. 6i also uses the confidence-based postselection, where the confidence per cluster state stabilizer is the geometric mean of the six decoded co-propagating 2D cluster states. On top of this decoding postselection, the logical stabilizer expectation value is shown as a function of the minimum number of co-propagating operators, N , with the same measurement outcome. We take the mean of all combinations of choosing N out of 6 such operators.

The permutation CNOT in Fig. 6g is applied in software and its effect here is to increase the weight of the operator connecting coordinates t_i and t_j , labeled as an effective separation $i - j$. Following the definitions in Ref. ³⁴⁷, two permutation CNOTs (swapping a pair of rows and a pair of columns) convert the cluster state stabilizers supported on logical qubits 3 to 6 from four weight-3 to one weight-3, two weight-6 and one weight-12 operator.




Figure C.3: Exploring entropy removal during single- and two-logical-qubit operations. **a**, Additional data for entropy removal via stabilizer measurement. As studied in Fig. 1d, the final logical error depends on the balance between the injected error rate (corresponding to the injected $\theta/2\pi$ per time-step) and the entropy removal rate (number of QEC rounds in the fixed total time window). For small injected error, there is an optimal number of QEC rounds (here, 3-4) since stabilizer measurement is imperfect and introduces entropy of its own. The plot uses MLE decoding and an acceptance fraction of 66% to enhance salient features. **b**, Absence of coherent logical error. Using the same error-injection protocol as in (a), the final logical state is measured in both the X and Z basis. With no QEC, the global coherent error results in a coherent logical rotation. With one round of QEC, or more, this coherent rotation vanishes. The stabilizer measurement is performed immediately after transversal state preparation, before the majority of the error is injected, such that the lack of logical rotation compared to no QEC can be attributed to the non-deterministic X(Z) stabilizer signs randomizing the response of the Z(X) logical operator to coherent error. Here we ignore the ancillas in the MLE decoding and use a 50% acceptance fraction. All curves are the average of both bases; for left plot, the measurement is the same basis as preparation, and for the right plot it is the orthogonal basis. **c-e** Analysis of logical gate performance of two logical qubits undergoing repeated transversal CNOTs and QEC, with the circuit studied in Fig. 3 of main text. **c**, The measured detector error probability increases linearly as a function of number of CNOTs applied in each round. **d**, Logical error probability as a function of number of CNOTs per round. Fits are to a functional form of $\mathcal{A} \cdot (p_{\text{qec}} + N\Delta p_{\text{det}})^3$, where $\mathcal{A}$, p_{qec} , and Δp_{det} are fitted parameters (blue). Using the fitted values of p_{qec} and Δp_{det} in **c** produces the predicted logical error probability $P_L(p_{\text{det}})$ (red). **e**, Results of **d** divided by total number of CNOT gates. Logical gate fidelity $F_L(p_{\text{det}})$ is $1 - P_L(p_{\text{det}})/(3N)$, where 3 is the number of rounds and N the number of gates per round.




Figure C.4: Theoretical characterization of logical-error-per-round ratio. **a**, Numerical simulations of rotated surface code initialized in $|0_L\rangle$ undergoing repeated QEC rounds with stabilizer measurement gate ordering in Ref.¹⁷⁰. As a simple model, we apply a theory error channel with error probability $p = 0.6\%$, where qubit resets and measurements experience uniform single-qubit depolarizing noise, CNOT gates are followed by uniform two-qubit depolarizing noise, and data qubits experience an idling single-qubit uniform depolarizing channel during ancilla qubit resets. We plot different LEPR ratios ($r_{d/(d+2)} = \text{LEPR}(d)/\text{LEPR}(d+2)$) using the PyMatching decoder²⁰¹, observing a change of at most 17% as the number of rounds is increased from four. The LEPR for a circuit with N logical qubits and n QEC rounds is defined as $\text{LEPR} = P_{L,\max} \left(1 - (1 - P_L/P_{L,\max})^{1/n} \right)$, where $P_{L,\max} = 1 - \frac{1}{2^N}$ is the logical error rate of a fully mixed state. **b**, Numerical simulations of repeated QEC on a single $|+_{\text{L}}\rangle$ surface code using the same circuit as the $d = 5$ surface code experiment in Fig. 3 of the main text. We use a simplified error model where we take the experimental error model and then turn all loss rates to 0 and double the Pauli error rates for idle errors, reset, measurement, and gate errors (single and two-qubit gates) on both data and ancilla qubits. We plot LEPR ratios for varying numbers of QEC rounds using an MLE decoder⁷⁰, observing changes of at most 9%, indicating that performance remains stable even for deeper circuits. **c**, Numerical simulations for circuits with 25 QEC rounds on four logical qubits, with interleaved transversal gate layers using the approximate experimental error model described above. Half of the qubits, selected uniformly at random, are initialized $|+_{\text{L}}\rangle$, while the other half are initialized in $|0_L\rangle$. Each gate layer consists of random pairing of transversal CNOT gates followed by logical single-qubit Pauli gates selected uniformly at random. After applying the random gate sequence U , the inverse $U^\dagger$ is applied, followed by transversal measurement in the same basis as initialization. By varying the number of gate layers interspersed between rounds, the LEPR ratio averaged over ≈ 100 randomly sampled circuits varies by at most 2%.

D

Supplementary information about analog quantum simulation of quantum coarsening and collective dynamics

D.1 EXPERIMENTAL PLATFORM

A detailed description of our experimental platform is given in ^{123,41}. All measurements are realized using a two-dimensional programmable quantum simulator based on Rydberg atom arrays. Single ^{87}Rb atoms are stochastically loaded into optical tweezers shaped by a spatial light modulator (SLM), and then rearranged into defect-free patterns using a pair of crossed acousto-optic deflectors (AODs). Both sets of tweezers use 852 nm-wavelength light. The atoms are then laser-cooled and optically pumped to the $|5S_{\frac{1}{2}}, F=2, m_F=-2\rangle$ state, which we denote as $|g\rangle$ in the main text. A pair of counterpropagating lasers at 420 nm and 1013 nm wavelengths couple $|g\rangle$ to the highly excited Rydberg state $|r\rangle = |70S_{\frac{1}{2}}, J=\frac{1}{2}, m_J=-\frac{1}{2}\rangle$ via a two-photon transition through the $6P_{\frac{3}{2}}$ orbital, blue-detuned by approximately $2\pi \times 2.4$ GHz with respect to the $5S_{\frac{1}{2}} \rightarrow 6P_{\frac{3}{2}}$ transition.

We use 16×16 lattices of atoms, and maintain an R_b/a ratio of $1.12 - 1.15$, such that only nearest neighbors lie within the blockade radius. The data shown in the main text is taken with two-photon Rabi frequencies of either $\Omega/2\pi = 3.8, 6.0$, or 3.1 MHz, with corresponding lattice spacings $a = 6.8, 6.45$, or $7.15 \mu\text{m}$ and $R_b/a = 1.15, 1.12$, or 1.13 . The 1013 nm and 420 nm single-photon Rabi frequencies are approximately balanced ($\Omega_{1013} \approx \Omega_{420}$). The experiment time T , sweep

parameters, and lattice spacings are all rescaled for different Ω , such that Ω/Δ , R_b/a , and the total phase accumulated ΩT are constant when comparing different experimental configurations.

In order to observe coarsening, the sweep rate through the critical point must fall within a certain range. A sweep rate that is too slow would create a low-energy state, and consequently, the coarsening dynamics may be too slow to measure. In contrast, a very fast sweep could inject too much energy and bring the system out of the ordered phase (which persists up to a finite energy density). Here, all measurements use linear sweeps with sweep rates of $\frac{\Delta/\Omega}{\Omega T} \approx \frac{1}{2\pi} \times 3.0$, which, we find, falls within the desired range.

For the global sweeps, we apply post-selection based on the success of the rearrangement protocol, selecting shots with ≤ 3 defects. Additionally, we post-select on measurement results in order to discard runs in which we suspect large-scale errors have occurred. Due to the high energy of the Rydberg blockade, a large number of blockade violations are extremely unlikely to be naturally generated by the coherent dynamics. Furthermore, since our projective measurement cannot differentiate $|r\rangle$ occupation from loss induced by other mechanisms, we attribute the presence of a large number of apparent blockade violations to manifestations of unwanted noise processes, such as blackbody-induced avalanche decays¹³⁶. We therefore discard runs where the longest chain of consecutive atoms in a single row or column detected in state $|r\rangle$ has a length of more than four sites.

D.2 LOCAL CONTROL

To enable individual single-site addressing of atoms with a local light shift, we use an SLM (Hamamatsu LCOS-SLM X15213-02) to generate optical tweezers in arbitrary spatial patterns with beam waist $1\ \mu\text{m}$, ensuring robustness to atomic position fluctuations (Extended Data Fig. D.1). The wavelength we choose to operate at, $784\ \text{nm}$, achieves a measured differential AC Stark shift between the $5S_{1/2}$ and $5P_{3/2}$ states of $12.2(3)\ \text{MHz}$ with $\approx 160\ \mu\text{W}$ per spot, but a negligible scatter-

ing rate (≈ 35 Hz) (the scaling of the light shift with laser amplitude is shown in Extended Data Fig. D.1c). The light is linearly polarized to minimize vector light shifts on the ground state hyperfine manifold. We further measure the shift on the $|g\rangle \rightarrow |r\rangle$ transition and find that the light shift is well approximated by the differential $5S_{1/2} \rightarrow 5P_{3/2}$ light shift as $\delta_0 = -2\pi \times 12(2)$ MHz at the same power per spot.

The phase holograms for the SLM are generated using the phase-fixed weighted Gerchberg–Saxton (WGS) algorithm, taking into account the desired position and relative intensity of the local light-shift pattern²⁴⁰. We first generate a local addressing pattern that closely matches the positioning of the atomic tweezer array; however, perfect matching of the two arrays is computationally expensive as it requires an extremely high sampling rate of the image plane of the local addressing pattern. In order to overcome this computational barrier, after creating an initial local addressing pattern, we align it to the atom positions by transforming the phase hologram. By stretching, rotating, and applying tilts and defocus, we can match the two patterns with feedback on the atom signal. The latter three can be easily controlled using Zernike polynomials, while the stretching and rotation require more care to preserve the intensity homogeneity of the desired pattern. We find that naïve rescaling or rotating of the hologram results in unwanted distortion of the intensity pattern, attributed to software interpolation when working with a pixelated hologram. This is mitigated by applying the computational corrections in the image plane. We take the Fourier transform of the hologram, convolve the intensity profile with a 2D Gaussian to broaden each spot over several pixels (in order to minimize effects of interpolation), and then apply the rotation and stretching. Lastly, we apply an inverse Fourier transform back to the Fourier plane and use the resultant phase hologram for the SLM. Using this procedure, we first coarsely align the individual addressing pattern to the tweezers on a camera, and then precisely align the two using a spin-echo measurement of the light shift (Extended Data Fig. D.1b) to optimize the alignment parameters such that the intensity is maximized at the atom sites. Good alignment is also crucial to prevent atom loss from turning on

a misaligned potential. Finally, we correct the tweezer intensities as required using the fitted light shifts to feedback on the target weights in the hologram generation.

Examples of states prepared using such tweezer profiles, where the detuning is used to strongly pin atoms to the ground state, are shown in Extended Data Fig. D.1d. At the boundary between different AF orders and the edges of the array, the mean-field repulsive interaction strength decreases for sites with fewer Rydberg neighbors; we therefore weight the local detuning strength inverse-proportionally to the number of neighbors. Note that when arbitrary weighting is used, the total power remains constant (number of addressed sites $\times 2\pi \times 12$ MHz), but the power is redistributed in the tweezers accordingly. Nevertheless, particularly at large Δ/Ω , neighboring Rydberg excitations start to be energetically favored (antiblockaded) along the domain boundaries. Excluding such edge effects, the preparation probability of preparing the single-atom ground state on the pinned sites is 93-95% (Extended Data Fig. D.1e).

For other realizations of single-site addressing using light shifts in atom arrays, see e.g. ^{84,110}.

D.3 THEORETICAL BACKGROUND OF COARSENING DYNAMICS

In this section, we summarize the theoretical details of the different kinds of coarsening processes that govern the dynamics of the system as long-range order is formed. Although our focus will be on the Rydberg atom array, to begin, let us consider the generic situation of a system driven through a continuous quantum phase transition (QPT) by tuning some parameter of the Hamiltonian, g , linearly with time. Without loss of generality, we assume that the quantum critical point (QCP) is located at $g = 0$ and the zero of time is set such that $g(t) = t/\tau$; hence, the system crosses the QCP at $t = 0$. For the specific case of the neutral atom array considered in this work, the time-dependent parameter g can be defined as $g(t) = (\Delta(t) - \Delta_c)/\Omega$.

As the system approaches the quantum critical point, its relaxation time diverges and it necessar-

ily falls out of equilibrium. However, *when* it does so depends on the velocity of the linear ramp, $\dot{g}(t) = 1/\tau$. The quantum Kibble-Zurek mechanism posits that the time at which the system's evolution ceases to be adiabatic is $t = -t_{\text{KZ}}$ with $t_{\text{KZ}} \sim t_0(\tau/t_0)^{\nu z/(\nu z+1)}$, where ν is the correlation length exponent, z is the dynamical critical exponent, and t_0 is some microscopic time scale. Thereafter, since the system cannot dynamically respond fast enough to the changing parameter of the Hamiltonian, it remains “frozen” through a so-called impulse regime until a later time $t = +t_{\text{KZ}}$, when it unfreezes on the other side of the quantum phase transition. During this impulse regime, the KZM presumes that the system's correlation length remains the same as when it initially froze: $\xi_{\text{KZ}} \sim l_0(\tau/t_0)^{\nu/(\nu z+1)}$, where l_0 is some microscopic length scale. As a consequence, in this picture, the correlation length in the ordered phase is also set by ξ_{KZ} with no subsequent dynamics.

However, the nonequilibrium correlation length of the system, $\xi(t)$, can and does grow both in the impulse regime as well as in the ordered phase since the long-range correlations take time to develop. In the experiments described in the main text, this occurs via a two-step process. First, as the system passes through the quantum critical regime, it undergoes *quantum critical coarsening*, which is governed by the dynamical critical exponent z of the particular QCP; for the $(2 + 1)\text{D}$ Ising transition, $z = 1$. Then, as time progresses and the ramp continues, the system eventually enters the ordered phase. Here, once the growing nonequilibrium correlation length $\xi(t)$ exceeds the equilibrium correlation length of the quantum ground state (which, recall, scales as $\xi_q \sim |g|^{-\nu}$), the dynamics cross over to a regime of *noncritical coarsening*, for which

$$\frac{d\xi(t)}{dt} \sim \frac{\xi_q^{z_d} \varepsilon}{(\xi(t))^{z_d-1}}, \quad (\text{D.1})$$

where ε is the many-body gap between the ground and first-excited states. The dynamical exponent z_d is dependent on the dimensionality and conservation laws of the system. For curvature-driven coarsening dynamics with a nonconserved scalar order parameter—as is indeed the case

experimentally— $z_d = 2 > z$. A particular feature of noncritical coarsening worth emphasizing is the dependence of the dynamics on the distance to the QCP encoded in Eq. (D.1). Specifically, the ground-state equilibrium correlation length scales as $\xi_q \sim |g|^{-\nu}$ and the gap $\varepsilon \sim |g|^{\nu z}$. Plugging in the exponents of the (2 + 1)D Ising QPT, $\nu = 0.629$ and $z = 1$, along with $z_d = 2$, in Eq. (D.1), we find the growth law

$$\frac{d\xi(t)}{dt} \sim \frac{(\Delta - \Delta_c)^{-0.629}}{\xi(t)}, \quad (\text{D.2})$$

This relation can be observed in Fig. 6.4, which studies the rate at which a locally introduced domain in the center of the array shrinks. The area of such a domain decreases at a rate $dr^2/dt \sim -\xi d\xi/dt$, which scales as $(\Delta - \Delta_c)^{-0.629}$.

For a ramp that continues indefinitely without stopping, the entire dynamical evolution of the correlation length can be described by a single universal scaling function encompassing the adiabatic, quantum critical coarsening, and noncritical coarsening regimes^{373,40,80}:

$$\xi(t) \approx \xi_{\text{KZ}} f\left(\frac{t}{t_{\text{KZ}}}\right) \equiv \xi_{\text{KZ}} f(x), \quad (\text{D.3})$$

where $f(x)$ is some universal function, and ξ_{KZ} and t_{KZ} depend on the ramp rate τ as specified earlier. The scaling variable x delineates the three regimes discussed above as

$$\begin{aligned} x \ll -1 : & \quad \text{adiabatic,} \\ |x| \lesssim \mathcal{O}(1) : & \quad \text{quantum critical coarsening,} \\ x \gg 1 : & \quad \text{noncritical coarsening.} \end{aligned} \quad (\text{D.4})$$

More generally, if the ramp is stopped at a time t_s , the dynamical scaling form is altered to

$$\xi(t) \approx \xi_{\text{KZ}} \mathcal{F}\left(\frac{t}{t_{\text{KZ}}}, \frac{t_s}{t_{\text{KZ}}}\right) \equiv \xi_{\text{KZ}} \mathcal{F}(x, x_s); \quad (\text{D.5})$$

for $x \leq x_s$, one recovers the earlier scaling as $\mathcal{F}(x, x_s) = f(x)$. If the ramp is stopped at $x_s = t_s/t_{\text{KZ}} \gg 1$ in the noncritical coarsening regime, the behavior of the universal scaling function for $x > x_s$ describes the physics during the hold time and is given by

$$\mathcal{F}(x, x_s) \approx x_s^{-\nu + (\nu z/z_d)} (\mathcal{C}x - \mathcal{C}_s x_s)^{1/z_d}, \quad (\text{D.6})$$

for some $\mathcal{O}(1)$ constants $\mathcal{C} > \mathcal{C}_s$. Note that because $z < z_d$, the coarsening speeds up as we stop earlier in the ordered phase, closer to the QCP (which results in a lower x_s). This is indeed what we observe experimentally during the hold time following global sweeps across the phase transition. Intuitively, this is because a smaller Δ/Ω corresponds to a greater relative influence of critical coarsening, which is faster than noncritical coarsening.

In contrast, near the thermal phase boundary, the system can undergo an interval of *classical critical coarsening*, which is described by a growth law $\xi(t) \sim t^{1/\bar{z}}$ with a distinct dynamical exponent. For the 2D classical Ising phase transition, $\bar{z} \approx 2.16 > z_d^{213}$, so the growth of correlations via classical critical coarsening is slower than for noncritical coarsening. Correspondingly, the dynamics should decelerate as one approaches the classical critical point, in sharp contrast to the speedup outlined above in the vicinity of a QCP.

D.4 STRUCTURE FACTOR AND CORRELATION LENGTH

In order to extract the structure factor and the correlation length³⁷⁴, we first calculate the two-point connected correlation function $G(\mathbf{r}_1, \mathbf{r}_2) = \langle \tilde{Z}_{\mathbf{r}_1} \tilde{Z}_{\mathbf{r}_2} \rangle - \langle \tilde{Z}_{\mathbf{r}_1} \rangle \langle \tilde{Z}_{\mathbf{r}_2} \rangle$ and then average over all pairs of points with identical displacements $\mathbf{r}$:

$$G(\mathbf{r}) = \frac{\sum_{\mathbf{r}_1, \mathbf{r}_2} G(\mathbf{r}_1, \mathbf{r}_2) \delta_{\mathbf{r}_1 - \mathbf{r}_2, \mathbf{r}}}{\sum_{\mathbf{r}_1, \mathbf{r}_2} \delta_{\mathbf{r}_1 - \mathbf{r}_2, \mathbf{r}}}. \quad (\text{D.7})$$

We first derive the standard structure factor by computing the Fourier transform of $G(\mathbf{r})$,

$$S(\mathbf{k}) \equiv \mathcal{F}_k[G(\mathbf{r})] = \sum_{\mathbf{r}} e^{-i\mathbf{k}\cdot\mathbf{r}} G(\mathbf{r}), \quad (\text{D.8})$$

and then calculate the radially averaged structure factor,

$$S(k) = \frac{\sum_{\mathbf{k}} S(\mathbf{k}) \delta_{|\mathbf{k}|,k}}{\sum_{\mathbf{k}} \delta_{|\mathbf{k}|,k}}. \quad (\text{D.9})$$

To extract a correlation length, we fit $S(k)$ to:

$$S(k) \approx \frac{S_0}{(1 + \xi^2 k^2)^{\frac{3}{2}}}, \quad (\text{D.10})$$

and we factorize S_0 as $S_0 = b\xi^2/\pi$. This form of Eq. (D.10) is equivalent to assuming that the position-space correlations follow an exponential decay, $G(r) \approx A \exp(-\frac{r}{\xi})$ ¹²³, up to finite-size corrections. While equilibrium considerations for an infinite-size system suggest that for small k , $S(k)$ should obey the Ornstein-Zernike form $S(k) \approx \frac{S_0}{1+\xi^2 k^2}$ ²³⁸, we empirically find that Eq. (D.10) better captures our observed nonequilibrium distributions.

For universal coarsening dynamics, we theoretically expect b to be constant. While our data indeed exhibit scaling collapse as in Eq. (D.10), we find that b varies during the dynamics, indicating the presence of an additional length scale(s). We observe that b is correlated with ξ and depends on Δ/Ω , as shown in Fig. D.2(d). Additional length scales that may affect the dynamics include the finite system size, the finite width of the domain walls (which depends on Δ/Ω), spatial inhomogeneity in Δ and/or Ω , and length scales introduced by decoherence effects such as decay due to the finite lifetime of the Rydberg state.

Specifically, the expected universal scaling regime is expected to hold for distances r and correlation lengths ξ such that $l \ll r, \xi \ll L$, where l is the width of a domain wall and L is the sys-

tem size⁵⁷, indicating that finite size effects are likely playing an important role in the present experiments. Hence, observing the theoretically expected universal coarsening behaviour in global quenches would likely require access to larger system sizes and correspondingly longer experiment times. While recent experimental advances in neutral atom array platforms suggest that lattices more than an order of magnitude larger than the one presented in this paper are within reach²⁸⁸, elsewhere we describe the use of local control to deterministically nucleate and study domain dynamics away from the system’s boundaries, allowing us to study universal properties of coarsening under present experimental conditions.

D.5 ANALYSIS OF DOMAINS IN GLOBAL SWEEPS

Using single-site-resolved detection, we can map out the domains in each snapshot. First, we calculate the local staggered magnetization. Each domain is then identified as a region of the array where the same ordering AF_1 or AF_2 is connected by nearest neighbors. We do not consider single spins of opposite order as a separate domain. These are identified as single atoms which are of the opposite order compared to all of their nearest and next-nearest neighbors. Only after we have identified single spin flips and corrected them to match their surrounding bulk order do we identify the domain boundaries. A domain’s area is defined as the total number of atoms comprising the domain. For the probability distribution of domain occupations, the frequency of each domain area is weighted by the area of that domain. We normalize the distribution by the sum of all area-weighted frequencies at each time step.

D.6 CLASSICAL ENERGY ANALYSIS

To calculate the classical energy per single shot of the experiment, we first perform the single-spin-flip correction as described above. We then identify regions of the array which do not belong to

either AF ordering by calculating a coarse-grained local staggered magnetization with a similar approach to previous works¹²³. In this work, specifically, we calculate the convolution $C_{x,y}$ of the Rydberg occupation $n_{x,y}$ with the kernel $W = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$ for each snapshot. The output values of $C_{x,y}$ range from 0 to 4, where the extremal values correspond to atoms surrounded by nearest and next-nearest neighbors that all belong to the same AF orderings, as shown in Extended Data Fig. D.3a. We consider an atom to be at a boundary if $n(x,y) = 1(0)$ and $C_{x,y} \neq 0(4)$ (see Extended Data Fig. D.3b). In the raw array (not single-spin-flip corrected), we then compute the classical energy for each snapshot. The value of the interaction energy is calculated from the lattice spacing (a) and V_0 as $V_{nn} = V_0/a^6$. By using the spin-flip-corrected lattice for domain identification and the uncorrected one for the subsequent energy calculation, single spin flips that occur contribute to the classical energy in identified domain walls and bulk orderings but not as separate domains. We exclude the layer of atoms closest to the edge of the array for all contributions to the classical energy. Note that for the classical energy calculation, we post-select such that the maximum number of directly adjacent Rydberg atoms can be no more than three (compared to four used throughout this work). Due to the sensitivity of the boundary identification procedure used here to correlated decays, this post-selection is slightly stricter.

D.7 ANALYSIS OF LOCALLY PREPARED DOMAINS

The radius, r , of the central domain is defined as the Manhattan distance, d_m , at which the radially averaged local staggered magnetization, m_r , crosses zero (see Extended Data D.3c). We consider Manhattan distances instead of Euclidian distances from the center of the injected domains as the former is more representative of the nearest-neighbor interactions dominating the dynamics for very short times (on the order of one Rabi cycle). For long times, both measurements of distances in our lattice reveal a collapse to the linear form. For this analysis, we only consider the unpinned

sublattice. The local state-preparation protocol prepares states where atoms that are locally detuned are prepared in $|g\rangle$ with high probability for all values of Δ/Ω (Extended Data Fig. D.1d). We observe that for Δ/Ω close to the QPT, there is a larger discrepancy of the local order parameter m_s between the pinned and unpinned sublattices. Therefore, when considering both sublattices, the radius is less clearly defined by a single point at which the order parameter crosses zero. We fit a linear relationship to extract dr^2/dt at each Δ/Ω .

A similar procedure is followed for the analysis of the coarsening dynamics in the zigzag domain wall. Here too, the mean value of the local order parameter $m_{x,y}$ is calculated at each lattice site per time step. The domain wall's horizontal position is calculated as the point at which the linearly interpolated line between points crosses $m_s = 0$. In Extended Data Fig. D.3d, we show the variation of the domain-wall position with hold time for two additional values of Δ/Ω . These data points reinforce the strong Δ/Ω -dependence of the domain-wall velocity. Note that for this analysis, we include atoms which were initially locally detuned as we are considering each row separately. We therefore see larger uncertainty in the domain walls' positions for low Δ/Ω (see Extended Data Fig. D.3d).

D.8 NUMERICAL SIMULATIONS OF LOCAL DOMAINS

We simulate the dynamics of locally prepared domains using the time-dependent variational principle (TDVP)^{186,187}. We use a two-site variant of this algorithm, which allows the bond dimension to grow with the evolution time at the expense of forgoing strict energy conservation due to the truncation step involved. In our calculations, we find that the energy is conserved to within 0.004% of that of the initial state up to the longest times simulated.

The initial states for the numerics are chosen to be a mean-field approximation of the experimentally prepared state. Specifically, we pin certain lattice sites to $|g\rangle$, as specified by the target configura-

tion, while the remaining ones are set to the vector on the Bloch sphere that minimizes the system's mean-field energy for a given Δ/Ω (instead of the fully polarized $|r\rangle$ state). The many-body evolution is simulated using a maximum bond dimension of $\chi = 1200$ with a time step Δt of $0.2\Omega^{-1}$ (the dynamics are also consistent with those for a smaller $\Delta t = 0.1\Omega^{-1}$). The results thus obtained are showcased in Extended Data Fig. D.4 and are found to be in good agreement with the experiments.

D.9 AMPLITUDE/HIGGS MODE

Background and theory

To describe the observed amplitude mode, we consider the low-energy effective action that describes the transition between the disordered and antiferromagnetic phases. Its Lagrangian is the φ^4 -theory

$$\mathcal{L}[\varphi] = \frac{1}{2} \left[(\partial_t \varphi)^2 + (\nabla \varphi)^2 - q \varphi^2 \right] - \frac{\lambda}{4} \varphi^4, \quad (\text{D.11})$$

where φ corresponds to the coarse-grained order parameter of the antiferromagnetic phase. While the phase transition is described via the Wilson-Fisher fixed point, here we perform a simple mean-field treatment of Eq. (D.11) in order to capture the physics away from the immediate vicinity of the transition. In particular, the classical mean-field equation of motion for the order parameter's expectation value is given by

$$\partial_t^2 \varphi = -(q + \lambda \varphi^2) \varphi, \quad (\text{D.12})$$

which corresponds to a classical anharmonic oscillator. The stationary value of the order parameter is given by $\varphi_0 = 0$ in the disordered phase ($q > 0$) and $\varphi_0 = \pm \sqrt{-q/\lambda}$ on the ordered side of the transition ($q < 0$). Expanding Eq. (D.12) for small amplitudes, $\varphi = \varphi_0 + \delta\varphi$, around the potential minima leads to harmonic oscillations of the order parameter with frequencies $\omega(q > 0) = \sqrt{q}$ and

$$\omega(q < 0) = \sqrt{2|q|}.$$

Numerical simulations

We first investigate the low-energy spectrum of the 2D Rydberg Hamiltonian at different values of Δ/Ω using the excited-state density matrix renormalization group (DMRG) method, which iteratively finds the lowest-energy eigenstate that is orthogonal to previous lower-energy eigenstates. Our simulations take into account van der Waals interactions up to third-nearest neighbors on the square lattice. We identify the gapped paramagnetic and spontaneous-symmetry-breaking (SSB) phases, and obtain the ground-state energy as well as the energy gaps ΔE_1 , ΔE_2 of the first two excited states above the ground state (in the SSB phase, the first excited state that we identify is the symmetry-related ground state). In Extended Data Fig. D.5a, we perform bond-dimension scaling for ΔE_1 , ΔE_2 on a 10×10 lattice with open (OBC) and periodic (PBC) boundary conditions up to bond dimensions $\chi = 512$ and $\chi = 256$, respectively. We note that while the energy gap in the paramagnetic (disordered) phase is robust to boundary conditions, in the SSB (ordered) phase, we identify a distinct boundary energy gap (panel i) smaller than the bulk energy gap (panel ii). In the dynamics, the coupling of the order parameter to the boundary mode vanishes with increasing system size, and we thus consider the bulk gap extracted from periodic boundary conditions as the relevant frequency of the amplitude (Higgs) mode.

Furthermore, we perform DMRG calculations to obtain the initial state of the quench dynamics, which is the ground state of the Hamiltonian with additional local detunings $|\partial_l|/\Omega = 0.7$ that pin one sublattice of the checkerboard to the ground state. We evaluate the energy expectation value of this initial state with respect to the unpinned 2D Rydberg Hamiltonian, and compare this with the energy of a $\mathbb{Z}_2$ product state, confirming that the pinned state is indeed a low-energy state.

Dynamics from the low-energy pinned initial state lead to amplitude oscillations at frequencies matching the ground-state gap, which we simulate using TDVP^{186,187} on a 10×10 lattice at a bond dimension $\chi = 256$. We use time steps of $0.25\Omega^{-1}$ and have verified that the resulting dynamics

agrees with smaller time steps of $0.1\Omega^{-1}$. As seen in Extended Data Fig. D.5b, for small systems with OBC, the dynamics in the ordered phase (here, $\Delta/\Omega = 2$) are dominated by a slow mode with frequency matching the boundary gap shown in Extended Data Fig. D.5a.i. On top of this slow mode, we see the presence of a mode with higher frequency matching the bulk gap shown in Extended Data Fig. D.5a.ii. Therefore, in the following, we consider dynamics with PBC, which allow us to isolate the bulk mode.

In Extended Data Fig. D.5c, we observe clear oscillations of the order parameter deep in either phase, which we fit to a damped harmonic oscillator $\varphi(t) \approx \varphi_0 + A \cos(\omega t + \theta_0)e^{-\gamma t}$ with frequency ω , damping γ , offset φ_0 , and amplitude A . We show the dependence of these parameters on the ratio Δ/Ω in Extended Data Fig. D.5d. Away from the phase transition, the oscillation frequencies overlap with the previously obtained ground-state energy gaps and are robust to system size. We further see that the damping and amplitude become larger towards the transition, where the offset acquires a nonzero value. Close to the transition, the TDVP dynamics fails to converge in bond dimension and fit to the damped harmonic oscillator's functional form, as apparent in Extended Data Fig. D.5c.ii. Moreover, even away from the critical point, the limited bond dimension does not capture a finite damping rate of the oscillations.

Dependence on local detuning strength

We experimentally investigate the dependence of the amplitude mode on the strength of the applied local detunings at two points, $\Delta/\Omega = 0$ and $\Delta/\Omega = 1.1$. This data is obtained by performing the state-preparation sequence for varying strengths of the applied local Stark shift δ . For all other uses of the local detunings in this work, the pinning is applied at a constant magnitude, $\delta_0 = -2\pi \times 12(2)\text{MHz}$, and the corresponding state preparation for various Δ/Ω is documented in Extended Data Fig. D.1c, d. Here, this strength is varied to much lower values than for the saturated pinning of the ground state ($0.04\delta_0 - 0.1\delta_0$) in the rest of the work (as indicated in Extended Data Fig. D.8a,b) before significant differences in the resultant oscillations are observed.

We find that at both values of Δ/Ω considered, the amplitude of the oscillations is progressively reduced as $|\delta|$ is decreased. The frequency of the oscillations also decreases with $|\delta|$. The change in the oscillation frequency is more pronounced for oscillations near the critical point than in the disordered phase. The δ dependence of the oscillation frequency, as well as the increased sensitivity near the phase transition, are qualitatively consistent with the behavior of an anharmonic oscillator, as predicted by the mean-field Eq. (D.12).

Amplitude mode in global sweeps

The Higgs-mode dynamics are also apparent in parallel to coarsening, as manifested in the oscillations of the magnetization n_i (Extended Data Fig. D.7a). Additionally, they can be clearly discerned by observing the dynamics of the total magnetization of the two-point correlation function in position space $C(r)$, as shown in Extended Data Fig. D.7b–d. The frequency of the oscillations of the correlation length closely support those extracted from the quench protocol described below (Extended Data Fig. D.8) and the calculated ground-state energy gap. The interplay of the amplitude mode with coarsening dynamics is generally unexplored theoretically. Therefore, we note further exploration of the two processes occurring in parallel, as a possible future extension of this work.

Quenches in the ordered phase

Although Higgs oscillations in the ordered phase can be extracted through the state-preparation sequence through deterministic preparation with local detunings, the amplitude of the oscillations is substantially reduced when compared to that inside the disordered phase. We therefore perform an alternate state-preparation sequence to extract the amplitude-mode frequencies deep in the ordered phase, as shown in Extended Data Fig. D.8a. First, local detunings are applied in a checkerboard pattern as the global detuning Δ is swept from negative values to $\Delta/\Omega = 3.3$, a point far inside the ordered phase. We hold the global detuning constant while quenching off the site-dependent δ . At this point, we quench Δ to its final detuning value in the ordered phase, yet closer to the phase transition. By way of this protocol, we observe, as with all other sweep protocols presented thus far,

long-lived oscillations of the order parameter and correlation lengths as shown in Extended Data Fig. D.8 b,c. The extracted oscillation frequencies ω are in close agreement with the ground-state energy gap in the ordered phase. Note that in Fig. 8.5d, the points in red at $\Delta/\Omega = 2.0, 2.5$ are measured via this quench protocol.

Frequency doubling

From the above-mentioned quenches to the ordered phase as well as in data from the local protocol, we extract the oscillations in both the order parameter and the correlation lengths. We find, as shown in Extended Data Fig. D.8d, that in the ordered phase, these two frequencies are approximately equal, while in the disordered phase they vary by $\omega_\xi/\omega_{m_s} \approx 2$. The changing relationship between the two observables can be understood by the following symmetry argument.

We begin by taking into account the dynamics of order-parameter fluctuations within a Gaussian approximation. Neglecting corrections to the effective mass due to fluctuations, the relevant equations of motion are^{81,119}:

$$\partial_t \mathcal{D}_{\mathbf{k},t}^{\varphi\varphi} = 2\mathcal{D}_{\mathbf{k},t}^{\varphi\pi}, \quad (\text{D.13})$$

$$\partial_t \mathcal{D}_{\mathbf{k},t}^{\varphi\pi} = \mathcal{D}_{\mathbf{k},t}^{\pi\pi} - (k^2 + q + 3\lambda\varphi^2)\mathcal{D}_{\mathbf{k},t}^{\varphi\varphi}, \quad (\text{D.14})$$

$$\partial_t \mathcal{D}_{\mathbf{k},t}^{\pi\pi} = -2(k^2 + q + 3\lambda\varphi^2)\mathcal{D}_{\mathbf{k},t}^{\varphi\pi}, \quad (\text{D.15})$$

where $\pi_{\mathbf{k}}(t) \equiv \partial_t \varphi_{\mathbf{k}}(t)$ and $\mathcal{D}_{\mathbf{k},t}^{\varphi\varphi} \equiv \langle \varphi(-\mathbf{k}, t) \varphi(\mathbf{k}, t) \rangle_c$. From the correlation function $\mathcal{D}_{\mathbf{k},t}^{\varphi\varphi}$, which corresponds to the structure factor discussed in the main text, one can extract the evolution of the correlation length.

Expanding $\varphi = \varphi_0 + \delta\varphi$ and $\mathcal{D}_{\mathbf{k},t}^{\varphi\varphi} = \mathcal{D}_{\mathbf{k}}^{\varphi\varphi} + \delta\mathcal{D}_{\mathbf{k},t}^{\varphi\varphi}$, in the disordered phase, an eigenmode analysis of Eqs. (D.13)–(D.15) yields a frequency spectrum $2\sqrt{k^2 + q}$. The smallest frequency, which is expected to set the correlation-length oscillations¹¹⁹, is thus $2\sqrt{q}$, i.e., twice that of the order parameter. In contrast, in the ordered phase, Eq. (D.14) contains a term $6\lambda\varphi_0\mathcal{D}_{\mathbf{k}}^{\varphi\varphi}\delta\varphi(t)$, and the

oscillation of the order parameter thus acts as a linear drive on the dynamics of two-point correlation functions. As such, the correlation length will oscillate at the corresponding frequency $\sqrt{2|q|}$ of the order parameter.

Frequency ratio of oscillations

In Extended Data Fig. D.9 we present the full dataset of amplitude-mode oscillations as a function of the distance from the phase transition. As described above, Landau mean-field theory predicts the relationship between oscillation frequencies to be $\omega(-|q|)/\omega(|q|) = \sqrt{2}$ ^{356,359}. However, beyond mean-field theory, the phase transition is described by the Wilson-Fisher fixed point, and the universal frequency ratio is shifted accordingly. Theoretical estimates based on both analytical and numerical methods yield a frequency ratio around $\omega(-|q|)/\omega(|q|) \approx 1.9$ ^{61,75,72}. We find that both the experimental data and MPS simulations deviate from the mean-field prediction and are suggestive of a similarly higher ratio. However, very close to the critical point, $|(\Delta - \Delta_c)/\Omega| \lesssim 0.3$, we observe deviations from this theoretical ratio. Possible explanations (besides the limited bond dimension for the MPS data) include finite-size effects (resulting e.g. in non-vanishing gap at QPT point), possible errors in QPT location, and the possibility that sufficiently close to the transition, the overdamped oscillations may no longer track the ground-state excitation gap. In future work, a detailed exploration of the region near the critical point via the amplitude mode could allow for higher precision tests of this universal ratio.




Figure D.1: Application of local detunings. **a**, In addition to the SLM used for atom trapping, a second SLM generates a set of superimposed tweezers that can apply arbitrary detuning profiles to the atoms. The pair of crossed AODs for local hyperfine control⁴¹ are not used in this work. **b**, A spin-echo sequence is used to measure and calibrate the differential AC Stark shift of the local detuning beams on the ground-state clock transition. Darker curves indicate a higher target detuning, weighted between 1 and 4, and the grey curves for unaddressed sites demonstrate negligible crosstalk between sites. The pattern here corresponds to a rotated version of the domains in the main text. **c**, Measured differential AC Stark shift of the $5S_{1/2} \rightarrow 5P_{3/2}$ transition. Unless stated otherwise (e.g., in Extended Data Fig. D.6), the laser amplitude used is 1, corresponding to $\approx 160 \mu\text{W}$ per site and imparting an approximate light shift on the differential $|g\rangle \rightarrow |r\rangle$ transition of $\delta_0 = -12(2)\text{MHz}$. The amplitude is the square root of power relative to $160 \mu\text{W}$. When sites are weighted, the total power imparted on the array remains constant but is proportionally redistributed. **d**, The local detunings are used to prepare deterministic checkerboard orderings, both for a single global order (left) or for domains with different orders (right). A site polarization of unity, measured at the start of the hold time for varying Δ/Ω , corresponds to $|g\rangle$ ($|r\rangle$) for addressed (unaddressed) sites, while -1 corresponds to the flipped state. Note that at finite Δ/Ω , the polarization of $|r\rangle$ sites $\neq 1$; at large Δ/Ω , mean-field shifts reduce the spin polarization of both $|g\rangle$ and $|r\rangle$ along the domain wall. **e**, Dependence of the site polarization on Δ/Ω in the bulk of the domain. As the local light shift becomes weaker relative to the global detuning, the probability of preparing $|g\rangle$ sites decreases slightly.




Figure D.2: Correlation lengths and structure factor. **a**, Structure factor $S(\mathbf{k})$ calculated at early (left) and late (right) times. **b**, Radially averaged structure factor $S(k)$ at early (left) and late (right) times for $\Delta/\Omega = 2.0$. The correlation length ξ is extracted by fitting $S(k)$ to $b\xi^2/\pi(k^2\xi^2 + 1)^{3/2}$. **c**, At various values of Δ/Ω , the structure factors collapse onto a single curve. In this case, we show data for $\Delta/\Omega = 3.0$. **d**, We find a clear dependence of amplitude b on the nonequilibrium correlation length ξ , and that their relation is Δ -dependent. **e**, For early times, the squared correlation length ξ^2 evolves linearly with time for all values of Δ/Ω tested. **f**, Dependence of $\partial_t \xi^2$ on the distance from the quantum critical point. We find that the dynamics accelerate as the system is held closer to the phase transition. The departure from the theoretically expected scaling (dashed line) stands in contrast to the case with locally seeded domains and points to the possible role of finite-size effects due to boundaries for these global sweeps.




Figure D.3: Analysis details on domain wall dynamics. **a**, Convolution to calculate the coarse-grained-local staggered magnetization for each single shot. This is used to distinguish the bulk checkerboard orderings from the domain walls. Extremal values of 0 and 4 indicate the ideal coarse-grained local staggered magnetization of either order while values of 1–3 indicate deviations. **b**, Demarcation of domain wall boundaries and the bulk for the same shot as shown in **a**. Red (white) indicates the domain walls (bulk). **c**, Staggered magnetization radially averaged at each Manhattan distance from the center of the prepared square. The radius is extracted as the d_m at which the linearly interpolated m_s crosses 0 (red lines). **d**, Sweeps to various end detunings for the motion of the zigzag domain wall. The acceleration of the dynamics near the QPT is also clearly supported for this shape of the domain wall. For both **c** and **d**, the left and right plots show data for $\Delta/\Omega = 2.0$ and $\Delta/\Omega = 3.5$, respectively. All data shown is for $\Omega/2\pi = 6$ MHz.




Figure D.4: Numerical simulations of locally prepared domains. **a**, TDVP simulations of coarsening dynamics observed when the system is initialized with a single square domain in the center (performed on a 15×15 lattice). The initial state approximates the experimental state preparation. **b, c**, The change in the radially averaged local staggered magnetization at Manhattan distances $d_m = 0, 2, 8$ from the center of the injected square. The layer at $d_m = 2$, closer to the domain boundary, changes order before $d_m = 0$. We also find that similarly, the atoms in the bulk ($d_m = 8$) remain in their initially prepared order. In these simulations, a slight acceleration in the dynamics is also observed on going closer to the phase transition. **d**, Simulation of the zigzag initial state on a 16×15 lattice. **e**, As in the experimental results, the locally curved points of the domain wall move towards the center of the domain while points with no domain-wall curvature remain stable.




Figure D.5: Numerical simulations of the amplitude mode. **a**, Bond-dimension scaling of the energy gaps ΔE_1 , ΔE_2 of the first two excited states above the ground state as a function of Δ/Ω , obtained via DMRG on a 10×10 lattice with (i) open boundary conditions (OBC), and (ii) periodic boundary conditions (PBC). **b**, Order-parameter dynamics at $\Delta/\Omega = 2$, of an initial state prepared with pinning fields $\delta_I/\Omega = -0.7$, on a 10×10 lattice with OBC, simulated via TDVP with $\chi = 64, 128$. **c**, Order-parameter dynamics at (i) $\Delta/\Omega = 0$, (ii) $\Delta/\Omega = 1.15$, (iii) $\Delta/\Omega = 2$, of an initial state prepared with pinning fields $\delta_I/\Omega = -0.7$, on a 10×10 lattice with PBC, simulated via TDVP with $\chi = 64, 128, 256$. The order-parameter dynamics are well modeled as a damped harmonic oscillator, as shown by the fitting. **d**, Oscillation parameters (blue): (i) frequencies, (ii) damping rates, (iii) offsets, and (iv) amplitudes as a function of Δ/Ω , obtained from functional fits to the numerical simulations on a 10×10 PBC lattice with $\chi = 128$ shown in **c**. The frequencies agree with the bulk (PBC) energy gaps for a 10×10 PBC lattice obtained in panel **a(ii)**, and with the oscillation frequencies for a 8×8 PBC lattice, for comparison.




Figure D.6: Effect of applied local detunings on the Amplitude mode **a**, Amplitude mode oscillations in the disordered phase at $\Delta/\Omega = 0$. The magnitude of the applied local detunings is varied as a fraction of $\delta_0 = -2\pi \times 12(2)\text{MHz} \approx -4\Omega$ (see Extended Data Fig. D.1c). **b**, Higgs-mode oscillations for varying magnitudes of the local detunings at $\Delta/\Omega = 1.1$. Note that for **a** and **b**, the frequency extracted here for $\delta = 0.06\delta_0 \approx -\Omega/4$ is plotted in yellow for comparison. **c**, With decreasing magnitude of δ , we observe a decrease in the oscillation frequency. The dependence is more pronounced closer to the phase transition. This behaviour can be qualitatively captured by considering the anharmonicity of the Landau mean-field potential in Eq. (D.11) as the phase transition is approached. **e**, The damping is stronger near the phase transition for both detunings. **f**, The oscillations at $\Delta/\Omega = 1.1$ are centered around finite ϕ_0 suggesting a transition into the ordered phase at this value. The data shown here is taken for the same Ω as all other Higgs-mode oscillations, $\Omega/2\pi = 3.1\text{MHz}$.




Figure D.7: Amplitude mode after a sweep through the QPT. Higgs-mode oscillations are apparent in parallel to coarsening in multiple observables. Note that we fit the the oscillation after an initial decay on the time-scales less than a Rabi-cycle. Extracted frequencies are consistent across observables. **a**, Oscillations of the magnetization $\langle n_i \rangle$. **b, c, d**, Dynamics of the two-point correlation function in position space for various end detunings.




Figure D.8: Ordered-phase quenches and frequency doubling. **a**, Sweep profile for quenches into the ordered phase. Under a constant Rabi drive at $\Omega/2\pi = 6$ MHz, the global detuning Δ is ramped from large negative values to large positive values deep into the ordered phase ($\Delta/\Omega = 3.3$). The local detuning δ is quenched off quickly over 50 ns. The global Δ is then also quenched down to varying final detuning values closer to the QPT. **b**, Oscillations of the staggered magnetization in the ordered phase following this preparation sequence. In agreement with the Landau mean-field picture, these oscillations occur around a nonzero value of the order parameter m_s . Note that we fit the oscillation after an initial decay on the time-scale less than a Rabi-cycle. **c**, Oscillations in the correlation length on top of a growing background. **d**, Ratio between oscillation frequencies extracted from the staggered magnetization (ω_{m_s}) and the correlation length (ω_ξ). In the disordered phase, $\omega_\xi/\omega_{m_s} \approx 2$ while in the ordered phase, this changes to $\omega_\xi/\omega_{m_s} \approx 1$.




Figure D.9: Ratio between oscillation frequencies. **a**, Experimentally extracted oscillation frequency of the amplitude mode. The frequency is plotted with respect to the distance from the critical point, $\Delta_c = 1.12$ measured in ¹²³. We see an approximate agreement with the theoretically predicted factor, $\omega(-|q|)/\omega(|q|) = 1.9$ ^{61,75,72}. Deviations close to the phase transition are apparent. **b**, Numerical simulations using MPS methods on 10x10 sites with periodic boundary conditions described in Extended Data Fig. D.5. We again see approximate agreement with the expected ratio, as well as deviations near the numerically determined critical point $\Delta_c = 1.0$.

E

Supplementary information about architecture for analog-digital quantum simulations

E.1 EXPERIMENTAL SYSTEM

Our experimental apparatus is described in detail in Refs. ^{289,130,42}, with several key modifications to enable these analog-digital experiments (ED Fig. 1). Individual ^{87}Rb atoms are stochastically loaded into a static array of 852-nm optical tweezers generated by a spatial light modulator (SLM), with a loading efficiency of 75% using Λ -enhanced gray molasses cooling on the D1 line ⁶⁶. A second set of 852-nm optical tweezers is generated with a pair of crossed acousto-optic deflectors (AODs; DTSX-400, AA Opto-Electronic) and enable atoms to be rearranged into defect-free arrays ^{28,380,123}. We use fluorescence imaging with a 0.65-NA objective (Special Optics) onto a CMOS camera (Hamamatsu ORCA-Quest C15550-20UP) for readout.

The hyperfine qubit is encoded in the clock states $|0\rangle = |F=1, m_F=0\rangle$ and $|1\rangle = |F=2, m_F=0\rangle$ and single-qubit operations on this qubit are implemented via two-photon Raman excitation ²⁶². The Raman light is directed either through a global path for global single-qubit gates or through a pair of crossed AODs for fully-programmable local gates ⁴². In this work, we operate the Raman laser at a decreased intermediate-state detuning of +53 GHz in order to reach a large two-photon Rabi frequency of ~ 8 MHz, which is necessary for coherent mapping (see “Coherent mapping”).

Owing to the large differential light shift at this detuning (~ 1 MHz), we choose to implement local X rotations indirectly via local Z rotations combined with global $Y^{A/2}$ rotations^{41,130}.

For excitation of the Rydberg qubit, composed of $|1\rangle = |5S_{1/2}, F=2, m_F=0\rangle$ and $|r\rangle = |nS_{1/2}, m_J = -1/2\rangle$, we use a two-photon excitation scheme with 420-nm and 1013-nm Rydberg beams. Here n denotes the Rydberg state principal quantum number and we use either $n = 53$ or $n = 70$, as discussed below. In contrast to our previous analog quantum simulation experiments^{289,123,386}, which used excitation from the stretched states $|5S_{1/2}, F=2, m_F = \pm 2\rangle$, here we drive transitions from the clock state. This choice comes with two downsides: the two-photon Rabi frequency is reduced by $\sqrt{2}$, for the same laser power and scattering rate, and the clock state can couple to both Rydberg magnetic sublevels, $m_J = \pm 1/2$. To suppress the latter effect, we first increase our magnetic field to 11.3 G such that the Zeeman splitting of 32 MHz is more than $2\times$ larger than the maximum detuning Δ used in this work. Second, we choose to excite to the lower-energy $m_J = -1/2$ state to ensure that all blockade-violating states—including those involving atoms in both m_J sublevels—lie at higher energy than the desired blockaded subspace. In the inverse configuration, cancellation of the van der Waals interaction and Zeeman splitting, which are of comparable magnitude, can spuriously bring such blockade-violating states near to resonance. Finally, where possible, we avoid large positive Δ . This is relevant for the Floquet engineering drives where we have the freedom to choose negative detuning π -pulses.

Our analog-digital approach to quantum simulation is based on coherently transferring the quantum state of each atom between the digital, long-lived hyperfine qubit and the interacting Rydberg qubit. Similar methods have been demonstrated previously in Refs.^{43,387}, and here we extend this approach to larger two-dimensional arrays. An important addition in this work is the use of non-destructive readout of the hyperfine qubit via spin-to-position conversion⁴². This enables loss detection and, since most atoms are retained during the readout, the atoms can subsequently be reinitialized and reused. The full experimental sequence for coherent mapping and qubit reuse is




Figure E.1: Analog-digital quantum simulations. **a**, Level diagram showing the relevant atomic transitions in ^{87}Rb . Raman transitions are used for single-qubit gates, the spin-dependent lattice is used for spin-to-position conversion, and the local SLM is used for local light shifts. The clock qubit $|1\rangle = |F=2, m_F=0\rangle$ is excited to the lower-energy $m_J = -1/2$ Rydberg state via a two-photon transition. All Rydberg pair states lie higher in energy than the blocked subspace. We suppress coupling to the nearby $m_J = +1/2$ Rydberg state, which has $3\times$ smaller matrix element, by increasing our DC magnetic field to 11.3 G. The intermediate state detuning of +1 GHz and $n = 70$ Rydberg state is used for all 2D experiments. For the 1D chain in Figs. 2,3 we instead use -6.3 GHz and $n = 53$. **b**, Schematic experiment timeline. The total loop takes 16.4 ms for 1D experiments and 24 ms for 2D experiments; in 2D, spin-to-position conversion is performed in two groups due to limited laser power for the AOD tweezers, and image processing and rearrangement are slower for the larger lattice. **c**, Example pulse sequence used for analog-digital quantum simulation. Not all steps are used in all experiments. For the QSL experiments, the first Rydberg π -pulse is omitted when starting from the all- $|0\rangle$ initial state before mapping up. Gaussian Rydberg π -pulses are used to minimize coupling to the other m_J state. The Rydberg Rabi frequency is reduced for the analog dynamics via the 420-nm intensity. For 1D experiments at $n = 53$ and performed in AOD tweezers, the atom spacing is reconfigured from $a = 6.0\ \mu\text{m}$ for local Raman rotations to $3.6\ \mu\text{m}$ for Rydberg dynamics. In 2D, $a = 6.0\ \mu\text{m}$ is used throughout. **d**, Characterization of coherent mapping. As in Fig. 1d, single qubits are arranged at next-nearest neighbor sites in 1D and 2D to emulate the vdW tails in a many-body state. Microstate-dependent phase accumulation from vdW tails occurs during the gap for the Raman π -pulse after the analog dynamics (here, a single-qubit $\pi/2$ -pulse). We observe the decay in Ramsey contrast for increasing vdW tail energy, E_{vdW} , by varying the lattice spacing in 1D (180 ns gap time) and also for increasing gap time in 2D ($2\pi \times 670$ kHz tail energy at $\sqrt{3}a$). Analytical formula includes only phase from vdW tails during the gap time t , given as $\cos^N(\phi/2)$ for N neighboring qubits and $\phi = E_{\text{vdW}}t$. Numerics simulate the entire measurement where the effect of vdW tails during the Rydberg pulses leads to a further reduction in contrast. Both curves are rescaled by the measured contrast of the hyperfine qubit without mapping and no other sources of error are included in the numerics.

illustrated in ED Fig. 1b,c and Supplementary Video 1, and involves the following steps:

1. Global state initialization of $|0\rangle$ using Raman-assisted optical pumping.
2. Local and global single-qubit gates to prepare a product state of arbitrary local superpositions of $|0\rangle$ and $|1\rangle$.
3. A fast Rydberg π -pulse to transfer population from $|1\rangle$ to $|r\rangle$.
4. A fast Raman π -pulse to transfer population from $|0\rangle$ to $|1\rangle$.
5. Analog Rydberg evolution by driving the $|1\rangle \leftrightarrow |r\rangle$ transition with time-dependent Rabi frequency Ω and detuning Δ .
6. A fast Raman π -pulse to transfer population from $|1\rangle$ to $|0\rangle$.
7. A fast Rydberg π -pulse to transfer population from $|r\rangle$ to $|1\rangle$.
8. Local and global single-qubit gates to measure in arbitrary local bases.
9. Loss-resolved non-destructive qubit readout.
10. Refilling of missing sites with atoms from the reservoir.

This cycle is repeated until the reservoir is depleted, typically allowing between 20 and 200 cycles (in 2D and 1D arrays, respectively) before reloading a new magneto-optical trap to refill the array. Details of mid-circuit rearrangement from the reservoir are given in Ref. ⁴², with the difference that here we use global imaging, allowing the reservoir occupations to be updated after each cycle rather than keeping track of used atoms in software. This is especially important in this work as loss of reservoir atoms is non-negligible when the SLM tweezers are turned off for up to $4\mu\text{s}$ during the analog evolution; this is required since the Rydberg state is anti-trapped by 852-nm light. Future work would benefit from independent control of the reservoir traps and the rest of the array.

E.2 DETAILS OF EXPERIMENT CONFIGURATION

To enable these experiments, we employ a range of configurations, including different array geometries and Rydberg excitation parameters. We use a zoned architecture in which the static SLM traps are arranged into a simulation zone and a reservoir zone. The reservoir, fixed throughout this work to be 11 rows tall and 32 columns wide, serves as a storage for atoms that are used to refill missing sites at the end of each experimental cycle (see above). The geometry of the simulation zone is programmed individually for each experiment, and it is where all operations—including single-qubit gates, Rydberg evolution and qubit readout—are performed. In addition, the experiments in this work are carried out using two different choices of Rydberg state, n , and intermediate state detuning. These are chosen to balance trade-offs associated with scattering, Rydberg state lifetime, interaction energies, and excitation matrix elements.

For the 1D experiments in Figs. 1–3, we operate at $n = 53$, enabling us to use an intermediate-state detuning of -6.3 GHz to minimize scattering error, as well as suppressing decoherence from electric field fluctuations, at the cost of reduced Rydberg lifetime. For efficient use of laser power, we use narrow Gaussian Rydberg beams (waist $\approx 35 \mu\text{m}$). These parameters are also compatible with high-fidelity two-qubit gates¹²⁹. For the analog Rydberg dynamics, the atoms are spaced by $a = 3.6 \mu\text{m}$ ($V_0 = 14.2$ MHz) and we use a Rydberg Rabi frequency of 2.3 MHz, except for in Fig. 1c. Since the local Raman tweezer waist of $2.5 \mu\text{m}$ is too large to address individual sites at this close spacing, we use moveable AOD tweezers and increase the atom spacing to $a = 6.0 \mu\text{m}$ for the state initialization and measurement (ED Fig. 3a). Dynamical decoupling of the hyperfine qubit is used to preserve coherence during this reconfiguration. Two additional rows of static SLM traps in the simulation zone are used for spin-to-position conversion.

For all remaining experiments performed in a 2D region, we instead operate at $n = 70$, motivated by the larger spacing used for Rydberg dynamics, here $a = 6.0 \mu\text{m}$ ($V_0 = 18.1$ MHz). This spacing




Figure E.2: Loss detection in analog quantum simulations. **a**, Distribution of losses per shot in QSL experiments. Baseline losses exclude the many-body dynamics and follow a Poisson distribution (solid line). The loss increases when performing analog dynamics and a non-Poissonian tail from correlated losses appears. There are 271 atoms in total. **b**, Spatial correlations of loss events. The two-point connected correlator follows a dipolar $1/r^3$ decay, suggesting the presence of Rydberg S - P interactions resulting from an atom, originally in the S state, undergoing blackbody decay to a nearby P state. These long-range interactions facilitate so-called avalanche events^{136,165,53} in which nearby atoms come into resonance with unwanted transitions and are also lost, explaining the tail in **a**. A background offset of $\approx 2.5 \times 10^{-5}$ is subtracted. Data in **a,b** are postselected on the atom being present in the prior image to remove correlations from mid-circuit rearrangement which are created when the reservoir starts to deplete. **c,d** Comparison of local and global loss postselection methods. Data are binned according to the number of missing atoms in each snapshot, including both losses and imperfect initial lattice filling. Postselection on no lost qubits within the local operator approximates the expectation value obtained from shots with a vanishing number of missing atoms across the entire 271-site lattice (dark curves). Randomly converting loss to qubit state $|0\rangle$ or $|1\rangle$ in post-processing strongly affects the calculated expectation value (light curves). We illustrate this for the **c**, three-body dimer populations and **d**, six-body potential energy from the QSL experiments. Dashed lines show the means of the entire dataset with local postselection. There are 285,948 shots in total and 6 shots have zero missing atoms.

is more comfortable for both mid-circuit rearrangement and for resolving sites separated by only $a/2$ during imaging, a consequence of the spin-to-position conversion which effectively doubles the density of the array. The simulation zone is thus $62 \mu\text{m}$ tall and $185 \mu\text{m}$ wide; the exact geometries are shown Fig. 1a (kagome lattice) and ED Fig. 4a (15 hexagonal plaquettes). To uniformly address this zone, we use tophat Rydberg intensity profiles which we homogenize to 2-3% peak-to-peak variation using the automated calibration described in Ref. ¹³⁰. This is crucial to ensure homogeneous Rydberg laser intensities for coherent mapping, which enforces a more stringent requirement on homogeneity than our previous analog simulation experiments (see below). Given limited laser power, here we operate at an intermediate-state detuning of +1 GHz.


E.3 COHERENT MAPPING

There are several sources of infidelity that must be considered when performing coherent mapping between the hyperfine and Rydberg manifolds, which we group into population-transfer errors and phase errors.

Population-transfer errors.

A principal limitation is the long-range vdW interaction tails which detune the Rydberg π -pulse away from resonance, reducing its fidelity. While we compensate for the mean-field interaction shift, the exact local resonance frequency differs between microstates and cannot be compensated. The dominant contribution comes from next-nearest neighbor sites which have a tail energy of 220 kHz in the 1D chain and 670 kHz in the 2D Kagome lattice. To mitigate this effect, we increase the Rydberg Rabi frequency for the mapping pulses to 9.7 MHz (1D) or 12.6 MHz (2D) by increasing the 420-nm intensity by $15 - 20\times$ relative to the analog dynamics. For these parameters, numerical simulations estimate an upper bound on the Rydberg π -pulse infidelity of $\approx 1.3\%$ in 2D when all possible microstates are equally populated (for an otherwise perfect pulse without vdW tails). Importantly, this infidelity leaves the population in $|r\rangle$, which is subsequently lost. These errors are therefore detected with loss-resolved readout and do not lead to a measurement error.

This phenomenon of loss-conversion of errors by the mapping protocol occurs for two other error sources. First, blockade-violating states (i.e., states with nearest-neighbor atoms both in $|r\rangle$) are strongly shifted off-resonance and are thus similarly left in the $|r\rangle$ state and lost; we estimate a $\approx 40\%$ loss probability per qubit in these cases. Second, any population erroneously in $|0\rangle$ prior to the mapping down is transferred to $|r\rangle$ and also lost; this can originate from imperfect initial state preparation as well as scattering errors and Rydberg decay during the analog dynamics. Finally, when mapping up, we always initialize all qubits in $|0\rangle$ and apply single-qubit gates to prepare $|1\rangle$, ensuring that gate errors do not prepare blockade-violating states.



Phase errors.

In addition to imperfect population transfer, both spurious coherent phase accumulation and incoherent dephasing reduce the fidelity of off-diagonal (i.e., X - and Y -basis) measurements. In particular, population in $|r\rangle$ accumulates a phase relative to the $\{|0\rangle, |1\rangle\}$ ground-state manifold during the gap time for the Raman π -pulse (100 ns unless stated otherwise). The origin of this phase includes:

- *Long-range Rydberg interaction tails.* A microstate-dependent phase is accrued due to the vdW interaction between atoms idling in $|r\rangle$.
- *420-nm, 1013-nm and Raman light shifts.* During the Raman π -pulse, the Raman and 1013-nm beams both impart phase shifts on the ground-Rydberg transition of $\sim \pi$. During the Rydberg π -pulse the 420-nm beam additionally induces a phase of $\sim 2\pi$ on the hyperfine qubit; this is particularly large when blue-detuned by +1 GHz as the light shift has opposite sign for $|0\rangle$ and $|1\rangle$. We calibrate and compensate these global phases, but residual inhomogeneity of the Rydberg beam intensities results in a spatially-varying phase. We find this to be negligible here, but if necessary this could also be calibrated and locally corrected.
- *Dephasing.* Doppler shifts and fluctuations of the 1013-nm light shift on the ground-Rydberg transition result in a random phase of $\sim 2\pi \times 0.02$ radians for 2D experiments ($T_2^* \approx 1-2 \mu\text{s}$ at $n = 70$, including 1013-nm spatial inhomogeneity). The 420-nm light is not on during the 100 ns gap and so does not contribute.

Mitigation and future improvements.

To mitigate these effects, we minimize the time between the analog evolution and the coherent mapping steps. This is accomplished using fast 70 ns π -pulses for both the Raman and the Rydberg (square and Gaussian pulse shapes, respectively). The 100 ns gap ensures timing jitter between the

Raman and Rydberg pulses (typically < 10 ns) does not lead to pulse overlap. We further synchronize all the arbitrary waveform generators (Spectrum Instrumentation) used to generate the Raman and Rydberg pulses to the same 10 MHz frequency standard (SRS FS725) to prevent drift in the pulse alignment over the repeated cycles.

In Figure 1, we benchmark the coherent mapping for both populations and coherent superpositions and find that it works well in the 1D chain, consistent with the above error sources and known state preparation and readout fidelities^{42,129}. In 2D, the larger vdW tail energy is the main limitation for mapping coherent superpositions. In ED Fig. 1d we measure the coherence after mapping down as a function of interaction energy and gap time, confirming that the reduction in contrast in experiment is in agreement with unitary numerical simulations including vdW tails.

Despite these limitations, we find coherent mapping to be an invaluable tool for probing complex many-body quantum states in one and two dimensions. In particular, directly measuring coherences of a many-body state in the $\{|1\rangle, |r\rangle\}$ manifold is otherwise impractical: to apply a single-qubit $\pi/2$ -pulse would require increasing the Rydberg Rabi frequency by an additional one to two orders of magnitude. Two possible approaches to further improve this technique in 2D are (1) reducing the magnitude of the vdW interaction strength relative to other energy scales or (2) reconstructing the desired off-diagonal observable using additional local measurements. For (1), a comparable fidelity to the 1D chain is reached at $\approx 4\times$ reduced tail energy for the Kagome lattice. This would require a 25% increase in the lattice spacing, with a corresponding increase in Rydberg laser power to maintain the same intensity, and an $\approx 4\times$ slower analog evolution to maintain the same R_b/a . For (2), one can infer an expectation value $\langle O \rangle$ without any phase error by measuring the evolved operator $U_{\text{vdW}} O U_{\text{vdW}}^\dagger$, where U_{vdW} is the known unitary evolution during the gap time. In general this will be a sum of Pauli strings requiring local measurements in mixed bases, although symmetries may reduce the number of measurement settings. Moreover, rather than assuming a known vdW interaction strength, in certain cases one can use parity oscillations to directly extract and compensate

for the accumulated phase by conditioning on the Rydberg populations of neighboring qubits in post-processing. The remaining limitation of this approach is that the Rydberg π -pulse still experiences a microstate-dependent detuning shift.

E.4 LOSS IN ANALOG SIMULATIONS

Atom loss is one of the main error sources in analog Rydberg simulation, arising predominantly from blackbody-induced decay of the Rydberg state, loss from turning off the optical tweezers during the Rydberg evolution, and scattering from the intermediate state in the two-photon excitation scheme. Furthermore, other error channels can be converted to loss errors via the coherent mapping protocol (see above) or other erasure-conversion schemes^{381,448,282}. To mitigate the impact of loss, we use loss-resolved readout to directly detect which atoms are missing. This identifies true losses as well as vacancies from imperfect rearrangement. In these experiments, we find that postselection using this loss information is an efficient approach that improves effective fidelities^{387,381}: only the qubits within the operator being measured need to be postselected on atom presence, such that the sample overhead scales with operator weight rather than system size (ED Fig. 2c,d).

We further use the ability to detect loss to characterize correlations between loss events in Rydberg simulations. In particular, blackbody decay of Rydberg S states to opposite-parity Rydberg P states is known to trigger the loss of many atoms, known as avalanche events^{136,165,74}. Indeed, we find that the distribution of the number of lost atoms per snapshot deviates from a Poisson distribution for large number of losses (ED Fig. 2a), indicative of such large-scale correlated loss events. Furthermore, the two-point loss correlations decay as $1/r^3$ for radial separation r (ED Fig. 2b), consistent with dipolar S - P interactions that can tune otherwise unpopulated blockade-violating states into resonance.

For Floquet experiments in the 1D chain (Fig. 2a–e and Fig. 3), we retain $\approx 95\%$ of the data after

loss postselection. In Figs. 2c,d we postselect on all 19 sites and keep 55%. For the hexagonal plaquettes, we postselect on all six sites, keeping 77% of the data. For the quantum spin liquid experiments in Figures 4–6, as representative examples, we retain 96% for the mean density in Fig. 4d, 90% for dimer populations in Fig. 4e, and 81% for the six-body RK energies in Fig. 4f.

E.5 APPROACH TO FLOQUET HAMILTONIAN ENGINEERING

Periodic driving can transform a system’s native Hamiltonian into an effective Hamiltonian with qualitatively different dynamics^{164,3}, commonly achieved through modulation of the single-qubit Pauli frame^{382,152} or the interaction strengths³³⁷. Here we use the Floquet engineering protocol developed in Refs.^{249,134}, summarized below, to realize effective Hamiltonians beyond the native Rydberg Hamiltonian (H , Eq. 7.1) using a global detuning drive with period τ , $\Delta(t + n\tau) = \Delta(t)$ and a constant Rabi frequency, Ω . This approach relies on constrained dynamics generated by Rabi oscillations within the nearest-neighbor blockade approximation, where the many-body unitary evolution is governed by the ‘PXP’ model¹³⁵

$$H/\hbar \approx \underbrace{\frac{\Omega}{2} \sum_i \mathcal{P} \sigma_i^x \mathcal{P}}_{H_{\text{PXP}}} - \Delta(t)N \quad (\text{E.1})$$

with $N = \sum_i n_i$. The projector $\mathcal{P} = \prod_{\langle i,j \rangle} (1 - n_i n_j)$ enforces the blockade constraint by removing configurations in which nearest-neighbor sites $\langle i, j \rangle$ are simultaneously excited to the Rydberg state. Within this constrained manifold, π -pulses of the global detuning effectively reverse the sign of the interacting PXP evolution²⁹³, $e^{i\pi N} H_{\text{PXP}} e^{i\pi N} = -H_{\text{PXP}}$. This property enables a ‘many-body echo’ by means of equally-spaced π -pulses located at $t_e + n\tau$ and $t_e + \tau/2 + n\tau$ within each Floquet cycle

n ,

$$\Delta_e(t|t_e, w_e) = -\pi \sum_{n \in \mathbb{N}_0} f\left(t|t_e + n\frac{\tau}{2}, w_e\right)$$

where $t_e \leq \tau/2$ and $f(t|\bar{t}, w) = \frac{1}{w\sqrt{\pi}} e^{-(t-\bar{t})^2/w^2}$ describes narrow Gaussian pulses centered at $\bar{t}$ with width w (Fig. 2a). In the instantaneous pulse limit $w_e \rightarrow 0$, the PXP evolution is reversed during the interval between the two π -pulses: within each Floquet period, the system evolves forward for time t_e , backward for $\tau/2$, and forward again for $\tau/2 - t_e$. The resulting net PXP evolution up to time t can be expressed as $U(t) = e^{-iH_{\text{PXP}} t_{\text{eff}}(t)}$, with an effective time $t_{\text{eff}}(t)$,

$$t_{\text{eff}}(0 \leq t \leq \tau) = \begin{cases} t & 0 \leq t < t_e \\ -t + 2t_e & t_e \leq t < t_e + \tau/2 \\ t - \tau & t_e + \tau/2 \leq t < \tau \end{cases},$$

and $t_{\text{eff}}(t + \tau) = t_{\text{eff}}(t)$. In particular, the system exhibits periodic revivals as there is no net evolution at stroboscopic times, $t_{\text{eff}}(n\tau) = 0$, and thus the Floquet unitary $U_F = U(\tau) = 1$.

By perturbing the many-body echo drive, $\Delta(t) = \Delta_e(t) + \Delta_p(t)$, we controllably generate a variety of higher-body interactions in the Floquet effective Hamiltonian. The effect of the global detuning perturbation $\Delta_p(t)$ can be captured in the interaction picture with respect to the many-body echo evolution^{249,134}, resulting in the Floquet unitary $U_{\text{int}} = \mathcal{T} \exp \left\{ i\hbar \int_0^\tau dt \Delta_p(t) N(t_{\text{eff}}(t)) \right\}$. Here, the total Rydberg occupation operator N is conjugated by the effective interacting PXP evolution,

$$N(t_{\text{eff}}) \equiv e^{it_{\text{eff}} H_{\text{PXP}}} N e^{-it_{\text{eff}} H_{\text{PXP}}}.$$

The key insight is that $N(t_{\text{eff}})$ acquires overlap with multi-body interactions due to operator spreading. The choice of perturbation $\Delta_p(t)$ then determines the weights of these terms in the engineered effective Floquet Hamiltonian $U_{\text{int}} \equiv \exp \{-i\tau H_{\text{eff}}\}$. To leading order,

$$\frac{H_{\text{eff}}^{(0)}}{\hbar} = - \int_0^\tau dt \Delta_p(t) N(t_{\text{eff}}(t)) \quad (\text{E.2})$$

with higher-order corrections obtained systematically from a high-frequency Floquet-Magnus expansion^{127,2}.

Deviations of $\Delta_e(t)$ from ideal instantaneous π -pulses ($w_e \rightarrow 0$) can also be absorbed into the global detuning perturbation: $\tilde{\Delta}_p(t) = \Delta_p(t) + [\Delta_e(t|t_e, w_e) - \Delta_e(t|t_e, 0)]$, and therefore enter the effective Hamiltonian on the same footing. Furthermore, van der Waals interaction tails beyond the blockade radius, which are neglected in the PXP approximation, are likewise dressed by the effective interacting PXP evolution, thereby generating additional contributions to the effective Hamiltonian.

E.6 ENGINEERING HOPPING DYNAMICS

In Figures 2 and 3, we engineer an effective blockaded XX Hamiltonian using a global detuning drive, parameterized by $\{\varepsilon, \gamma, \theta, \Delta_0\}$ with the form

$$\begin{aligned} \Delta(t)/2\pi = & (\varepsilon - 0.5) \sum_{n \in N} f\left(t \Big| \frac{\tau}{4} + n \frac{\tau}{2}, w_e\right) \\ & + (\gamma - \theta) \sum_{n \in N} f(t|n\tau, w_p) + \theta \sum_{n \in N} f\left(t \Big| \frac{\tau}{2} + n\tau, w_p\right) + \Delta_0/2\pi \end{aligned}$$

where $w_e = 16$ ns and $w_p = 8$ ns are chosen to be as narrow as possible whilst ensuring the magnitude of $\Delta(t)$ does not exceed ≈ 20 MHz. This minimizes accidental modulation of the 420-nm intensity, and hence Ω , arising from the frequency-dependent deflection efficiency of the acousto-

optic modulators (AOMs) used to control the detuning profile; we correct for this effect within the ± 20 MHz range. We use $\tau = 0.36 \mu\text{s}$ and $\Omega = 2\pi \times 2.3 \text{ MHz}$ ($R_b/a = 1.35$), ensuring strong blockade such that the PXP approximation is valid. Since the effective PXP evolution is bounded by $\Omega t_{\text{eff}}/2\pi \leq \Omega\tau/8\pi \approx 0.2$, we can exploit a short-time expansion of the time-evolved global number operator $N(t_{\text{eff}})$ in the blockaded subspace²⁴⁹,

$$N(t_{\text{eff}}) = N + \frac{\Omega t_{\text{eff}}}{2} H_{\text{PYP}} + \frac{(\Omega t_{\text{eff}})^2}{4} \left(-H_{\text{PXP}} + \frac{1}{4} H_{\text{ZIZ}} - 3N \right) + \mathcal{O}((\Omega\tau)^3).$$

Here, $H_{\text{PYP}} \equiv \sum_i \mathcal{P} \sigma_i^y \mathcal{P}$ enacts blockaded local spin-flips with a relative phase, $H_{\text{PXP}} \equiv \frac{1}{2} \sum_i \mathcal{P} (\sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y) \mathcal{P}$ generates blockaded nearest-neighbor spin exchange, and $H_{\text{ZIZ}} = \sum_i \sigma_i^z \sigma_{i+2}^z$ is a diagonal next-nearest-neighbor interaction. We define $\sigma_i^x = |1_i\rangle \langle r_i| + |r_i\rangle \langle 1_i|$, $\sigma_i^y = i|1_i\rangle \langle r_i| - i|r_i\rangle \langle 1_i|$, and $\sigma_i^z = 2n_i - 1 = |r_i\rangle \langle r_i| - |1_i\rangle \langle 1_i|$. These operators form the basis for the following ansatz for the effective Floquet Hamiltonian:

$$H_F/\hbar \approx -\mu N + U H_{\text{ZIZ}} - J H_{\text{PXP}} - g_X H_{\text{PXP}} - g_Y H_{\text{PYP}} \quad (\text{E.3})$$

In Fig. 2a, we illustrate the detuning profile for $\varepsilon = 0.1$, $\theta = 0.05$, and $\gamma = -0.25$. For all data in the main text, $\theta = 0$ and $\Delta_0 \approx 2\pi \times -0.5 \text{ MHz}$ (up to a systematic offset in Δ) are fixed. For the many-body echo in Fig. 2c, we use $\varepsilon = -0.06$ and $\gamma = 0$. For Fig. 2e we use, from left to right, (ε, γ) to $(-0.05, -0.05)$, $(0.0, -0.15)$, and $(0.1, -0.35)$, where γ is chosen for good number conservation. In Fig. 2c (right) and Figure 3, we target fast hopping dynamics and use $\varepsilon = 0.2$ and $\gamma = -0.4$. Even though many of these parameters are not small, we find that the engineered dynamics are well-described by Eq. (E.3), where the coefficients are renormalised by the finite pulse

widths and van der Waals tails. For example, we find that the hopping energy scales linearly with ε (ED Fig. 3b).

Corrections from higher-order terms in the short-time and Floquet-Magnus expansions may also introduce additional operators into the effective Hamiltonian, with coefficients beyond quadratic order in $\Omega\tau, \varepsilon, \gamma, \theta$ ²⁴⁹. In particular, for drives with $\gamma > \pi/2$, we can redefine $\gamma' = \gamma + \pi$, such that the drive is recast as a perturbation around a different many-body echo containing additional π -pulses at $t = n\tau$. This redefinition effectively symmetrizes the drive, $t_{\text{eff}}(\tau \leq t \leq 2\tau) = -t_{\text{eff}}(0 \leq t \leq \tau)$, and thereby removes terms linear in t_{eff} from the leading-order effective Hamiltonian. Because $N(t_{\text{eff}}) = N(0)$ at $t = n\tau$, another effect of this transformation is a modification of the prefactor of the global detuning associated with γ in the leading-order effective Hamiltonian.

E.7 VERIFICATION OF THE EFFECTIVE HAMILTONIAN

To independently verify the effective Hamiltonian governing the Floquet-engineered dynamics, we fit experimentally measured stroboscopic observables to exact classical simulations of the ansatz model^{324,184,441}. Specifically, we numerically optimize the Hamiltonian coefficients to minimize the discrepancy between simulated and experimentally measured bitstring probability distributions on a contiguous 8-site subsystem, which fully determine all observables diagonal in the chosen measurement basis and supported on the block (i.e., all local $k \leq 8$ -body correlators in that basis). To mitigate boundary effects, we analyze the central eight sites of the experimentally realized open chain, and compare to eight consecutive sites of a 12-site periodic chain in simulation.

We consider several initial states $|\psi_0\rangle$, Floquet periods n , and measurement unitaries $\mathcal{M}$ to constrain the Hamiltonian terms that govern the hopping, detuning, and creation of excitations. First, we separately prepare and Floquet evolve the vacuum (all-ground) state $|\psi_0\rangle = |\Psi_\downarrow\rangle$, a single central excitation $|\psi_0\rangle = \sigma_0^x |\Psi_\downarrow\rangle$, and two next-nearest-neighbor excitations $|\psi_0\rangle = \sigma_{-1}^x \sigma_{+1}^x |\Psi_\downarrow\rangle$,

followed by measurement in the Z -basis ($\mathcal{M} = \mathbb{I}$). To probe phase-sensitive dynamics, we additionally prepare four different superpositions of the vacuum and single-excitation states, $|\psi_0\rangle = \frac{1}{\sqrt{2}} (|\Psi_\downarrow\rangle \pm \sigma_0^{x,y} |\Psi_\downarrow\rangle)$, and measure in the X -basis ($\mathcal{M} = e^{-i\frac{\pi}{4}\sum_j \sigma_j^y}$). Aggregating these datasets, we perform a multi-parameter minimization of the cost function

$$C(\mathcal{S}) = \sum_{\substack{\{(|\psi_0\rangle, \mathcal{M})\} \\ \mathbf{s} \in \{0,1\}^{\otimes 8}}} \sum_{n=0}^{n_{\max}-1} \frac{1}{2^8 \cdot n_{\max}} \left[p_{\text{exp}}(\mathbf{s} | |\psi_0\rangle, n, \mathcal{M}) - \left| \langle \mathbf{s} | \mathcal{M} e^{-i(n\tau)H|_{J,b,g}} | \psi_0 \rangle \right|^2 \right]^2$$

where $\mathcal{S} = (J, \mu, U, g_X, g_Y)$ is the set of input parameters and $p_{\text{exp}}(\mathbf{s} | |\psi_0\rangle, n, \mathcal{M})$ is the experimentally measured probability of observing bitstring $\mathbf{s}$ at time $n\tau$ in each of the above configurations.

We apply this procedure to verify the effective Hamiltonian engineered in Figure 3, yielding a cost function $C_{\min} = 1.34 \cdot 10^{-4}$ with optimal coefficients $J/2\pi = -0.143$ MHz, $\mu/2\pi = 0.411$ MHz, $U/2\pi = 0.023$ MHz, $g_X/2\pi = 0.000$ MHz, and $g_Y/2\pi = 0.021$ MHz (ED Fig. 3c,d). We find the fitted J and $(\mu + 4U) = 2\pi \times 0.502$ MHz are close to those extracted from the single-particle dispersion measurement (Fig. 3d), corroborating the blockaded XX Hamiltonian as a good description of the Floquet-engineered dynamics.

E.8 MANY-BODY SPECTROSCOPY

Here we outline the procedure for measuring the single-particle Green's function, $\mathcal{G}(r, t) = \langle \Psi_\downarrow | \sigma_r^x(t) \sigma_0^x(0) | \Psi_\downarrow \rangle$, in Fig. 3d^{294,246}.

Measurement protocol.

We prepare one of four different superposition states $|\psi_{\pm x, y}\rangle = (|\Psi_\downarrow\rangle \pm \sigma_0^{x,y} |\Psi_\downarrow\rangle) / \sqrt{2}$ where $\sigma_0^{x,y}$ acts on the central site, $r = 0$, and all other spins are in the ground state. After evolving under the Floquet-engineered dynamics for a time $t = n\tau$, we measure all sites in the X -basis in parallel to

obtain the site-resolved expectation values

$$\begin{aligned}
\langle \sigma_r^x(t) \rangle_{\pm x, y} &\equiv \langle \psi_{\pm x, y} | \sigma_r^x(t) | \psi_{\pm x, y} \rangle \\
&= \langle \Psi_{\downarrow} | \left[\sigma_r^x(t) + \sigma_0^{x, y}(0) \sigma_r^x(t) \sigma_0^{x, y}(0) \right. \\
&\quad \left. \pm \sigma_r^x(t) \sigma_0^{x, y}(0) \pm \sigma_0^{x, y}(0) \sigma_r^x(t) \right] | \Psi_{\downarrow} \rangle.
\end{aligned}$$

Taking differences cancels the first two terms and isolates the real and imaginary parts of the target Green's function:

$$\begin{aligned}
\Re \mathcal{G}(r, t) &= \frac{1}{2} \langle \Psi_{\downarrow} | [\sigma_r^x(t) \sigma_0^x(0) + \sigma_0^x(0) \sigma_r^x(t)] | \Psi_{\downarrow} \rangle \\
&= \frac{\langle \sigma_r^x(t) \rangle_{+x} - \langle \sigma_r^x(t) \rangle_{-x}}{4} \\
\Im \mathcal{G}(r, t) &= \frac{1}{2i} \langle \Psi_{\downarrow} | [\sigma_r^x(t) \sigma_0^x(0) - \sigma_0^x(0) \sigma_r^x(t)] | \Psi_{\downarrow} \rangle \\
&= \frac{1}{2} \langle \Psi_{\downarrow} | [\sigma_r^x(t) \sigma_0^y(0) + \sigma_0^y(0) \sigma_r^x(t)] | \Psi_{\downarrow} \rangle \\
&= \frac{\langle \sigma_r^x(t) \rangle_{+y} - \langle \sigma_r^x(t) \rangle_{-y}}{4}
\end{aligned}$$

where $\sigma_0^y | \Psi_{\downarrow} \rangle = -i \sigma_0^x | \Psi_{\downarrow} \rangle$ is used in the second line. Finally, by performing a discrete Fourier transform in time and space, we obtain the density of states (Fig. 3d) as a function of momentum, k , and frequency, ω , across the 1D Brillouin zone:

$$\mathcal{G}(k, \omega) = \sum_{n, r} e^{-in\omega\tau} e^{-ikr} \mathcal{G}(r, n\tau).$$

For the hopping Hamiltonian, Eq (E.3), with chemical potential μ , blockaded hopping energy J , and next-nearest neighbor interaction U , the peak of the imaginary part follows the single-particle dispersion $\omega(k) = \mu + 4U + 2J \cos ka$ ²⁴⁹.

Calibration of phases.

Any relative phase accumulation between the two qubit manifolds must be compensated to ensure the prepared state, analog evolution, and measurement are all in the desired basis. Using single qubits, we first map up $|+\rangle_{0,1} = (|0\rangle + |1\rangle)/\sqrt{2}$ and calibrate the phase of a readout Rydberg $\pi/2$ -pulse, corresponding to faithfully preparing $|+\rangle_{1,r} = (|1\rangle + |r\rangle)/\sqrt{2}$. We then remove the Rydberg $\pi/2$ -pulse, map down, and calibrate the phase of a readout Raman $\pi/2$ -pulse, corresponding to returning to $|+\rangle_{0,1}$. Finally, we calibrate the 420-nm differential light shift on the hyperfine qubit imparted during the analog evolution. These phases are then used in the four measurements described above. We further apply an in-software compensation of the mean-field phase from vdW tails during the mapping down ($\approx 0.04\pi$), estimated from the measured populations with the $\pi/2$ -pulse before readout removed. Miscalibration of these phases mixes the real and imaginary parts of the Green's function.

E.9 DOMAIN WALL DYNAMICS UNDER RYDBERG BLOCKADE

The constrained XX model in the absence of spin flips, Eq. (E.3) with $g_{X,Y} = 0$, can be mapped onto the integrable XXZ Heisenberg chain for which spin-wave excitations are known to be interacting⁴³⁰, as demonstrated by the spin scattering in Fig. 3b. In ED Fig. 3e-g, we further investigate the effect of these interactions on the dynamics of domain walls on top of the maximally-packed Néel state, $|\mathbb{Z}\rangle_2$. By allowing for a weak spin-flip term $g_{X,Y}$, domain wall pairs can be created and subsequently propagate under the hopping dynamics, washing out the initial Néel order. Importantly, the band of domain wall excitations must be resonant with the Néel state for this process to occur. By sweeping the energy of the Néel state across the band, tuned via the chemical potential, we observe this generation of gapless domain walls (ED Fig. 3g). This can be interpreted as an out-of-equilibrium probe of the underlying Luttinger liquid ground state found when $g_{X,Y} \rightarrow 0$ ²⁴⁹. This Luttinger liquid phase is a consequence of the emergent $U(1)$ symmetry^{430,8,249} and, interest-

ingly, is the one-dimensional counterpart of the $U(1)$ QSL¹⁴².

E.10 ENGINEERING RING-EXCHANGE DYNAMICS

Ring-exchange dynamics can be realized in our Floquet approach either by directly engineering a six-body interaction ($g_6 \times \mathcal{P} (\prod_{i=1}^6 \sigma_i^x) \mathcal{P}$) or indirectly via a second-order process generated from three-body interactions ($g_3 \times \sum_i \mathcal{P} \sigma_i^y \sigma_{i+1}^y \sigma_{i+2}^y \mathcal{P}$) together with a large static detuning. Here we employ the latter approach. Since accessing these higher-weight operators requires leaving the short-time regime, we follow the numerical optimization of the detuning profile of Ref.¹³⁴, which generates profiles that produce large overlap of the effective Hamiltonian with the desired three-body interactions. Concretely, we approximate the profile by the closed-form expression

$$\begin{aligned} \Delta(t)/2\pi = & (\varepsilon_0 - 0.5) \sum_{n \in N} f(t|n\tau, w_e) + \\ & (\varepsilon_1 - 0.5) \sum_{n \in N} f\left(t \middle| \frac{\tau}{2} + n\tau, w_e\right) + \mathcal{A} \cos\left(\frac{2\pi}{\tau}t\right) + \Delta_0/2\pi \end{aligned}$$

parameterized by $\{\varepsilon_0, \varepsilon_1, \mathcal{A}, \Delta_0\}$. We use $\tau = 0.30 \mu\text{s}$ and $\Omega = 2\pi \times 2.3 \text{ MHz}$ ($R_b/a = 1.41$), such that the effective PXP evolution satisfies $\Omega t_{\text{eff}}/2\pi \leq \Omega\tau/4\pi \approx 0.35$, providing sufficient evolution to build appreciable overlap with the target three-body operators¹³⁴. For the data in Fig. 2f-h, we optimize the detuning profile by hand, targeting fast dynamics and good number conservation, and obtain $w_e = 20 \text{ ns}$, $\varepsilon_0 = 0.025$, $\varepsilon_1 = 0.13$, $\mathcal{A} = -1.75$ and $\Delta_0 \approx 2\pi \times 0.65 \text{ MHz}$, where we observe that $\mathcal{A}$ is the primary parameter that controls the speed of the dynamics (ED Fig. 4b). For these parameter values, the leading-order effective Hamiltonian of Eq. (E.2) exhibits three-body interactions of strength $g_3 \approx -2\pi \times 0.3 \text{ MHz}$ and an effective detuning $\mu = -2\pi \times 0.8 \text{ MHz}$. Neglecting all terms other than g_3 and μ , this induces an effective ring-exchange of strength $g_6^{(\text{eff})} \approx 2\pi \times 0.22 \text{ MHz}$, which predicts plaquette oscillations of period $\pi/g_6^{(\text{eff})} \approx 2.3 \mu\text{s}$. This time scale




Figure E.4: Engineering ring exchange via Floquet driving. **a**, Atom image of the array of 15 isolated hexagonal plaquettes. **b**, Continuously tuneable six-body interactions. The speed of the ring-exchange dynamics is varied via the parameter $\mathcal{A}$ of the Floquet drive, shown here via the populations on the two sublattices. $\mathcal{A} = -1.75$ uses the same data as in Fig. 2h. **c**, Homogeneity of the engineered Hamiltonian. Sublattice dynamics are shown for individual plaquettes, where the several-percent-scale inhomogeneity of the 420-nm and 1013-nm light shifts result in the spatial variations of the static detuning offset. **d**, Numerical simulations of the engineered Floquet dynamics. Good agreement is found when including the measured $T_2^* \approx 1 \mu\text{s}$ in numerics and further multiplying by an empirical exponential decay, $\exp[-\langle n \rangle t / T_1]$, where $\langle n \rangle \approx 2.5$ is the mean Rydberg density and the Rydberg lifetime is $T_1 \approx 30 \mu\text{s}$, dominated by 1013-nm scattering. The T_2^* is dominated by inhomogeneity between plaquettes and shot-to-shot fluctuations of the 1013-nm light shift and therefore is modelled as a Gaussian-distributed detuning shift with zero mean and $2\pi \times 200 \text{ kHz}$ standard deviation. **e**, Stabilization of entangled states. The coherence time of a quasi-adiabatically prepared GHZ state, $\approx (|010101\rangle + |101010\rangle) / \sqrt{2}$, can be extended by Floquet driving, as compared to holding the state at large static detuning, $\Delta / \Omega = 5.9$. We attribute this to the combined effect of an energy gap generated by the ring-exchange interaction and the reduced incoherent error owing to the lower average Rydberg population between stroboscopic times. The engineered Hamiltonian contains non-number-conserving three-body interactions which admix the state $|100100\rangle$ and its cyclic permutations, resulting in oscillations in both the populations and coherences. Decay of populations at static detuning is consistent with Rydberg T_1 .

is consistent with the oscillations observed in the experiment (Fig. 2h, ED Fig. 4d). Moreover, $H_{\text{eff}}^{(0)}$ contains blockaded hopping of strength $J \approx -2\pi \times 0.23$ MHz, comparable to the strength of the effective ring-exchange and consistent with the experimental result of Fig. 2g. We note however that in this protocol, perturbations around the many-body echo are strong, both due to the detuning drive as well as the stronger tails of the vdW interaction on the chosen ring geometry. The first-order expression of Eq. (E.2) should therefore be viewed as a qualitative guide for the direct optimization of the dynamics on the simulator.

Finally, we find that inhomogeneity between different plaquettes contributes to the decaying contrast of the oscillations (ED Fig. 4c). This decay is reproduced in numerical simulations that include Rydberg dephasing, $T_2^* \approx 1 \mu\text{s}$, incorporating both shot-to-shot fluctuations and spatial inhomogeneity of the 420-nm and 1013-nm light shifts. We similarly find that the engineered dynamics are sensitive to proper optimization of Δ_0 . This sensitivity is the main limitation to extending this Floquet approach to two dimensions, where we find the large micro-state dependent vdW tails reduce the fidelity of the engineered number conservation; the symmetry of the two Néel states on the single plaquette provides a special case where this effect is small. This therefore motivates exploring protocols that are robust to such detuning shifts, as well as architectures such as dual-species processors with individually tunable inter- and intra-species interactions^{439,77}.

E.11 OVERVIEW OF THE $U(1)$ QUANTUM SPIN LIQUID

Here we provide a theoretical overview of the Rokhsar-Kivelson Hamiltonian and the field theory description of the $U(1)$ quantum spin liquid.

Rokhsar-Kivelson state and Hamiltonian.

The Rokhsar-Kivelson (RK) state is the equal-amplitude and equal-phase superposition of dimer

configurations on the honeycomb lattice:

$$|\Psi_{\text{RK}}\rangle \propto \sum_{\mathcal{D}} |\mathcal{D}\rangle \quad (\text{E.4})$$

where $\mathcal{D}$ is any configuration of dimers with exactly one dimer touching each vertex, i.e., satisfying the local dimer constraint. In literature, such a wavefunction first appeared as the ground state of the RK dimer model³⁵³, $H_{\text{RK}} = -tT + vV$, at its RK point ($t = v$). Here, T is the kinetic energy and V is the potential energy, defined in Fig. 4a. At this RK point, the Hamiltonian is positive semi-definite and can be written as a sum of (non-commuting) projectors. The wavefunction of Eq. (E.4) is a zero energy state of the Hamiltonian and is hence one of its ground states.

Predictions for diagonal correlations.

The diagonal observables of the RK state can be rewritten as observables in a classical gas of dimers about which much is known^{142,302}. In particular, for any diagonal observable $\hat{\mathcal{Z}}$ in the dimer basis, $\langle \mathcal{D} | \mathcal{Z} | \mathcal{D}' \rangle = \mathcal{Z}(\mathcal{D}) \delta_{\mathcal{D}, \mathcal{D}'}$, hence

$$\langle \Psi_{\text{RK}} | \hat{\mathcal{Z}} | \Psi_{\text{RK}} \rangle = \frac{1}{\sum_{\mathcal{D}} 1} \sum_{\mathcal{D}} \mathcal{Z}(\mathcal{D}). \quad (\text{E.5})$$

where $\sum_{\mathcal{D}} 1$ is the infinite-temperature partition function of the classical dimer gas. The correlation functions of the classical dimer gas are known exactly^{232,233,419,270} and can be reproduced by the $(2 + 0)\text{D}$ compact-boson conformal field theory (CFT)¹⁴² whose action is

$$S = \frac{1}{8\pi K} \int d^2x (\partial_\mu \phi)^2 \quad (\text{E.6})$$

where for the RK state $K = 1$, and ϕ is a compact bosonic field. This can be seen through a number of methods, including height-field representations of dimer configurations¹⁴². Most explicitly, the

classical dimer gas is free-fermion solvable via a transfer matrix method²⁷⁰ and its correlations on the lattice can be explicitly related to those of Eq. (E.6) in the continuum⁴⁴².

The relevant correlation functions of CFT are of $\partial\phi = \partial_x\phi - i\partial_y\phi$, which obey $\langle\partial\phi(x)\partial\phi(0)\rangle \sim 1/x^2$, and are further associated with vertex operators $e^{im\phi}$, which are local operators for $m \in \mathbb{Z}$. These vertex operators have scaling dimensions given by $\Delta_m = m^2K$, implying that correlations of these vertex operators decay as:

$$\langle e^{im[\phi(x)-\phi(0)]} \rangle \sim \frac{1}{x^{2m^2K}} \quad (\text{E.7})$$

For $m = 1$, $e^{i\phi}$ is related to the nematic order parameter. In particular, let x label the $\mathcal{A}$ sublattices of the honeycomb lattice, which forms the tripartite triangular lattice. If $\sigma(x) = 0, 1, 2$ labels the sublattices of the triangular lattice, then:

$$\omega^{\sigma(x)} \sum_{\alpha} n_{\alpha}(x) \omega^{\alpha} \sim e^{i\phi(x)} + \partial\phi(x) \quad (\text{E.8})$$

where $n_{\alpha}(x)$ is onsite density at the three edges $\alpha = 0, 1, 2$ surrounding a vertex x and $\omega = e^{2\pi i/3}$ is the third root of unity. As such we expect the correlations of the nematic order parameter to decay as $\sim 1/x^2$.

Moreover, when $m = 1/2$, $e^{im\phi}$ is not a local operator and appears as a string on the lattice. Indeed, we find that⁴⁶⁷:

$$\omega^{\sigma(x)} \sum_{\alpha, \beta} T_{\gamma_{\alpha, \beta}} \omega^{\alpha - \beta} \sim e^{i[\phi(x) - \phi(0)]/2} + \dots \quad (\text{E.9})$$

where the $\dots$ are fields whose decay is faster than the leading $e^{i[\phi(x) - \phi(0)]/2}$ term, $\gamma_{\alpha\beta}$ is a string that starts at edge α around vertex 0 and ends at edge β around vertex x and $T_{\gamma_{\alpha\beta}}$ is a product of σ^z operators at these edges. Consequently, we expect that the string correlation function decays as $\sim 1/x^{1/2}$.

These vertex operators of the CFT correspond to the rectified dimer-dimer and dimer-string

correlation functions studied in experiment (see “Calculation of correlation functions” and ED Fig. 9).

Connection to $U(1)$ gauge theory.

To establish the wavefunction of Eq. (E.4) as a microscopic realization of the $U(1)$ quantum spin liquid, we recast its properties in the language of the $U(1)$ gauge theory. The local dimer constraint can be defined on any lattice and, for example, on the kagome lattice it leads to an emergent $\mathbb{Z}_2$ gauge theory description of significant interest in the study of topological quantum matter^{302,142,429,386}. However, a crucial property of the honeycomb lattice is that it is *bipartite*. Consequently, it is possible to orient each dimer on the lattice to point from the A to B sublattice.

This directionality is captured by the ‘electric field’ degree of freedom, defined as $E_\alpha(x) = e^{iQ_0 \cdot x} n_\alpha(x)$, where x is a vertex, $\alpha = 0, 1, 2$ labels the three edges surrounding the vertex, and $Q_0 = \frac{2\pi}{3a} \cdot (1, \frac{1}{\sqrt{3}})$ is chosen such that the electric field is positive (goes out of) the A sublattice and negative (goes into) the B sublattice. The dimer constraint can thus be expressed as a Gauss’s law, $\nabla \cdot E(x) |\Psi_{\text{RK}}\rangle = e^{iQ_0 \cdot x} |\Psi_{\text{RK}}\rangle$ where $q(x) = \nabla \cdot E(x) \equiv \sum_{\alpha \in \text{dx}} E_\alpha(x)$ is the electric charge operator given by lattice divergence of the electric field. The integral (sum) of $q(x)$ over an enclosed region counts the net number of dimers crossing the enclosing loop, with opposite dimer orientations contributing with opposite sign. The spectrum of this operator therefore lives in the integers, identifying this as a $U(1)$ gauge constraint—i.e. the sort of gauge constraint that appears in electromagnetism, where charges are quantized to the integers.

The physics of the RK point can be more firmly connected to the deconfined phase of a $U(1)$ gauge theory through the quantum Liftshitz theory^{142,16}. While the diagonal static correlators of the RK state are described by the compact-boson CFT of Eq. (E.6), the dynamical correlations in

(2 + 1)D are described by the action

$$S = \int d^3x \left[\frac{1}{2}(\partial_t \phi)^2 + \frac{1}{2}g(\nabla^2 \phi)^2 \right] \quad (\text{E.10})$$

$$= \int d^3x \left[\frac{1}{2}\mathbf{B}^2 + \frac{g}{2}(\nabla \times \mathbf{E})^2 \right]. \quad (\text{E.11})$$

where $g = 1/4\pi K$, $\mathbf{E} = \nabla \times (\phi \hat{z})$ and $\mathbf{B} = \partial_t(\phi \hat{z})$. This action is to be contrasted with the usual action of the $U(1)$ gauge theory—equivalently, the action of electromagnetism—whose Lagrangian is $\propto (\mathbf{E}^2 - \mathbf{B}^2)$. As it will turn out, the absence of the $\mathbf{E}^2$ term in the action is essential for the physics of the RK point. In particular, it is known that standard electromagnetism in (2 + 1)D is generically confining (in the absence of additional fermionic fields)³³⁴. Consequently, standard electromagnetism does not have algebraic correlations, whereas the electromagnetic action of the RK point does.

Instabilities in equilibrium.

The missing leading $\mathbf{E}^2$ term in the ‘RK electromagnetic action’ of Eq. (E.10) shows that the RK point is fine-tuned. In particular, to realize this action each term with a lower or equal number of derivatives as the $(\nabla \times \mathbf{E})^2$ term must be fine-tuned to zero. The only such terms that are consistent with time-reversal and spatial rotation symmetry and are not total derivatives are $\mathbf{E}^2$ and $(\nabla \cdot \mathbf{E})^2$. In an exact dimer model, the Gauss law $(\nabla \cdot \mathbf{E})$ is fixed in the Hilbert space and can be ignored, implying that the Eq. (E.10) requires one parameter of fine-tuning—tuning the coefficient on the $\mathbf{E}^2$ term to zero. In an emergent gauge theory, such as those that appear in spin systems, the coefficient on the $(\nabla \cdot \mathbf{E})^2$ term also needs to be fine-tuned.

Quantum spin lakes.

Despite this instability in equilibrium, the RK state can be prepared robustly from non-equilibrium quantum dynamics^{369,161}. The dynamics of interest keeps the system at low-energy (in contrast to,




Figure E.5: Experimental optimization of QSL state preparation. **a**, Parameterization of $\Omega(t)$ and $\Delta(t)$ profiles. The detuning sweep is a cubic interpolation between five points (red markers) defined by seven parameters. The Rabi frequency is constant with a 20 ns turn-off time. **b**, Summary table of parameters used in the main text, found using a Bayesian optimizer. Detuning sweep is later clipped between ± 20 MHz. Final row is a by-hand modification of optimized profile in row 3. **c**, Results of one round of closed-loop optimization. Metrics used in the multi-objective optimization are shown as a function of trial number. Each trial takes roughly 7 minutes to obtain 2000 Z -basis and 4000 X -basis shots. At the end of the round, we repeat the best 2-4 trials to select the final sweep. For some rounds (not shown), we also included minimization of the magnitude of open length-5 X strings but found this had little effect and was within error bars for all reasonable sweeps. **d**, The measured RK potential energy is positively correlated with both the mean density and RK kinetic energy across different preparation sweeps. **e**, Sensitivity to sweep parameters. By rescaling the time of the 0.8μ s sweep (row 3 in b), we find the RK energies can change significantly in this out-of-equilibrium regime. The local maxima of $\langle T \rangle$ and $\langle V \rangle$ coincide. For slower sweeps, $\langle T \rangle$ decreases faster than $\langle V \rangle$ which we explore numerically in ED Figure 6.

say, quench dynamics) and thus the physics is governed by near-ground-state energetics. Here we consider the PXP approximation for simplicity and discuss the effect of vdW tails in “Closed-loop optimization of state preparation”.

At large detuning Δ , the low-energy states of our system includes all dimer coverings on the honeycomb lattice, which are split at sixth order in perturbation theory $\sim \Omega^6/\Delta^5$. This leads to a nematic ground state wavefunction that spontaneously breaks rotation and translation symmetry. As Δ is reduced, monomers—vertices with no adjacent dimers—will start to fluctuate in the system. These are interpreted as electric charges in the language of an emergent gauge theory. Eventually, monomers in the system will proliferate leading to a trivial ‘Higgs’ phase with a low density of dimers; indeed, when $\Delta \rightarrow -\infty$, the ground state is $|0\rangle^{\otimes N}$.

Ground state behavior in hand, consider the following simplified dynamical protocol. Starting in the ground state at $\Delta = -\Delta_0$ where $\Delta_0/\Omega \gg 1$, the wavefunction is (remarkably) well approximated by:

$$|\psi(0)\rangle \approx \mathcal{P} \bigotimes_{i=1}^N (|0\rangle + \varepsilon |1\rangle)_i \quad (\text{E.12})$$

where $\mathcal{P}$ denotes the blockade constraint, encoding the small quantum fluctuations induced by Ω at large negative Δ . If we adiabatically evolve the system from $-\Delta_0$ to $+\Delta_0$, i.e., at a rate that is much smaller than the Ω^6/Δ_0^5 splitting between dimer configurations, then the final state $|\psi(T)\rangle$ would well approximate the nematic ground state at $+\Delta_0$. Conversely, for an instantaneous sweep, $|\psi(T)\rangle \approx |\psi(0)\rangle$.

Suppose we sweep at a rate that is much slower than Δ_0 , corresponding to adiabatic evolution relative to dimer/monomer excitations (i.e., tracking the ground-state filling), but much faster than Ω^6/Δ_0^5 such that the dynamics is diabatic with respect to this small energy splittings between dimer coverings. A natural conjecture for the state of the wavefunction after performing this sweep is the

wavefunction of the system is unchanged except that monomers are projected out.

$$|\psi(T)\rangle \propto \mathcal{P}_{\text{Gauss}} |\psi(0)\rangle = |\Psi_{\text{RK}}\rangle \quad (\text{E.13})$$

where $\mathcal{P}_{\text{Gauss}}$ projects the wavefunction into the space of dimer configurations. This was shown to faithfully reproduce the result of the dynamics in this intermediate-rate ‘hemidiabatic’ regime in Ref. ³⁶⁹. Crucially, because the rate of the sweep is determined by rough energy scales, it does not require the same fine-tuning as required for the equilibrium realization of the RK point’s correlations.

Since as Δ is swept from $-\Delta_0$ to $+\Delta_0$ the gap to the monomer/dimer excitations closes, it is not possible to truly be adiabatic with respect to these excitations. As a result, the Gauss law is not perfectly projected out and instead we are left with a dilute density of monomer excitations atop $|\Psi_{\text{RK}}\rangle$, per an analog of the Kibble-Zurek mechanism ⁴⁶⁹. The density of these dilute excitations sets an effective lake size L_{lake} over which one expects the system to have correlations similar to the RK wavefunction. Crucially, L_{lake} is an intensive quantity in the experiment and is not expected to depend on system size.

Finite size effects.

Considering the RK wavefunction on an infinite cylinder of width L_y , the diagonal correlations are expected to follow those of the $(2 + 0)\text{D}$ compact-boson CFT, Eq. (E.6), which reduces to the $(1 + 1)\text{D}$ compact-boson CFT placed on a ring of size L_y ^{199,142}. The finite size effects of these correlations are well known to follow

$$\langle e^{im[\phi(x) - \phi(0)]} \rangle \sim \exp\left(-\frac{2\pi\Delta_m x}{L_y}\right) \quad (\text{E.14})$$

For comparing to the experiment, we expect that L_y is replaced with L_{lake} , which sets the effective

size of the RK wavefunction. This leads to two interesting predictions:

- I. The observed monomer density in the experiment sets the scale over which the system matches the RK wavefunction. Consequently, postselecting on no monomers should improve consistency with the RK wavefunction (Fig. 5a).
- II. The length scale of the decay of the experiment—predicted to be set by L_{lake} —should be largely independent of the height of the lattice used in the experiment (ED Fig. 10e).

We find that both of these predictions are borne out in the data. Furthermore, from Eq. (E.14) we expect that the ratio of the logarithms of the rectified dimer-dimer ($m = 1$) and dimer-string ($m = 1/2$) correlators is equal to $\Delta_{m=1}/\Delta_{m=1/2} = 4$ for any system size. We find good agreement with this prediction in both classical Monte Carlo simulations and experiment (Fig. 6e and ED Fig. 10).

E.12 CLOSED-LOOP OPTIMIZATION OF STATE PREPARATION

In Figure 4, we optimize detuning sweeps to target the $U(1)$ quantum spin liquid using a Bayesian optimizer, implemented with Ax³¹⁹. Our initial ansatz consisted of a sum of Gaussians for both $\Omega(t)$ and $\Delta(t)$ (not shown), inspired by pulsed counterdiabatic driving schemes¹⁶². However, the large positive detunings required in this approach led to increased atom loss and blockade violations, and the optimization converged to smooth sweeps, thus motivating the parameterization shown in ED Fig. 5a,b. There are eight free parameters, $\{\Omega, T, \Delta_{\min}, \Delta_{\max}, \Delta_{\inf}, t_{\inf}, s, t_{\text{slope}}\}$ and we impose loose bounds such that $\Delta(t)$ is physically motivated. For example, we enforce $\Delta_{\min} < 0$ and $\Delta_{\max} > 0$, but do not enforce monotonicity (i.e., s can be negative). We employ a multi-objective optimization over four metrics ($\langle T \rangle, \langle V \rangle$, mean Rydberg density, and atom loss) rather than a scalar cost function (ED Fig. 5c). In particular, we do not minimize the RK energy directly, since the different measurement error for $\langle T \rangle$ and $\langle V \rangle$ would bias the results. We find that the observables are




Figure E.6: Preparation of quantum spin lakes in numerical simulations. **a**, Schematic of the energy spectrum of the Rydberg Hamiltonian on the kagome lattice as a function of detuning. For $\Delta \gg \Omega$, monomer excitations are strongly gapped whereas vdW tails and perturbative couplings lead to a small splitting between dimer coverings. Hemidiabatic preparation^{369,162} corresponds to sweeping through the phase transition at a rate (dashed line) adiabatic relative to monomer excitations but sudden relative to the low-energy splitting. **b**, To confirm that the final sweep rate in Fig. 8.4d corresponds to the hemidiabatic regime, we modify the $0.8\mu\text{s}$ detuning profile by accelerating the sweep through the phase transition (gray dashed line) by a multiplicative rate factor after the inflection point (dotted line). **c**, Evidence for hemidiabatic preparation of quantum spin lakes. We use exact diagonalization to simulate a 3×3 unit cell lattice with periodic boundary conditions along each lattice vector. Starting from the ground state at $t = 0$, we evolve under the modified sweep from **b** and plot the final potential and kinetic energies for both the PXP approximation and the full simulation of the long-range vdW interaction. For large rate factors, we see that accelerating the sweep decreases these RK energies due to the large number of monomers in the final state. For small rate factors, although the PXP model approaches the RK state ($\langle T \rangle = \langle V \rangle \approx 0.19$ for this system size), the inclusion of vdW tails changes the relative phases and amplitudes between dimer configurations which reduces both T and V . Rather, the optimal rate with tails corresponds to the hemidiabatic regime where the sweep is slow enough to remove monomers but fast enough that vdW tails dephasing is limited. As expected, the experimentally optimized sweep corresponds to rates near the maxima of T and V . **d**, Effect of disorder on unprojected energies. To account for sources of disorder in the experiment, we simulate the $0.8\mu\text{s}$ sweep starting from the empty product state $|1\rangle^{\otimes N}$ 100 times, randomly sampling on-site detuning shifts, global detuning shifts, in-plane positional shifts, and out-of-plane positional shifts from Gaussians with widths of $2\pi \times 0.043\text{ MHz}$, $2\pi \times 0.2\text{ MHz}$, $0.02a$, and $0.08a$, respectively. We plot V , T , and T after 100 ns of idle evolution under the vdW interaction to simulate the gap for the Raman pulse during mapping down. We attribute the dip in these energies at the end of the sweep to finite-size effects which differ from the experiment. **e**, Effect of same disorder on projected energies, where bonds around the loop are projected onto no dimer (see Fig. 5c).




Figure E.7: Dimer configurations in the Rokhsar-Kivelson state. **a**, Comparison between the standard and weakly modified detuning sweeps used in the main text. Zoom-in highlights the scale of the modification. **b**, Probability distribution of dimer configurations in isolated subsystems using the modified sweep. The experimental distribution follows a thermal distribution with a fitted inverse temperature $\beta = 9.0(2)/V_0$, where V_0 is the nearest-neighbor interaction strength (light gray bars). The same data is used for the inset of Fig. 5a. The dimer configurations are explicitly shown for two subsystem sizes. The vdW tail energy is larger for configurations with more flippable plaquettes within the subsystem, shaded in red, which explains the reduced probability of these configurations. **c**, Probability distribution within a 2×1 subsystem for varying number of dimers surrounding the subsystem, using the standard $0.8 \mu\text{s}$ sweep. The uniform distribution in Fig. 5a corresponds to 0 outer dimers and the single configuration with 8 outer dimers is omitted. The outer dimers restrict the state of the lattice surrounding the subsystem to be in one of a finite number of allowed dimer configuration which satisfy the boundary conditions, such that even the perfect RK state has a non-uniform probability distribution within subsystems in this analysis. The RK distribution is shown for a 12×6 lattice with periodic boundary conditions obtained from Monte Carlo simulations; we find a qualitatively similar distribution for other lattice dimensions although the precise amplitudes differ slightly (not shown). We find experiment is in good agreement with this qualitative distribution.

largely correlated across trials (ED Fig. 5d), supporting this approach. For each optimization we run 50-80 trials, at the end of which we repeat the best 2-4 trials to select the final detuning sweep (Fig. 4d).

Using exact diagonalization to simulate the experimental drive on a 3×3 unit cell kagome lattice, we find that the optimized detuning sweeps naturally target the hemidiabatic regime^{369,162} necessary to prepare quantum spin lakes. In particular, the final leg of the 0.8 μ s drive sweeps through the ground state phase transition at a rate which is adiabatic relative to monomer excitations but sudden relative to the low-energy splitting between different dimer coverings induced by perturbative resonances and vdW tails (ED Fig. 6a,b). Indeed, as this final sweep rate is increased, $\langle T \rangle$ and $\langle V \rangle$ decrease due to monomer proliferation; as the rate is decreased, $\langle T \rangle$ and $\langle V \rangle$ decrease due to relative phases accumulated between different dimer coverings, primarily due to the vdW interaction (ED Fig. 6c).

E.13 BOUNDARY ENGINEERING WITH LOCAL DETUNINGS

To prevent seeding of Rydberg excitations from the array edges, we apply local light shifts to the boundary atoms using a programmable array of optical tweezers (1.5 μ m waist) generated with a SLM²⁸⁹. The light is sourced from a titanium:sapphire laser (M Squared) operated 1 THz red-detuned of the D2 line. For the Floquet engineering experiments, we do not use any boundary detuning. To generate the appropriate boundary detunings for the quantum spin liquid, we start with an empirical weighting of roughly three times larger intensity on sites near the lattice corners and use the weighted Gerchberg-Saxton (WGS) algorithm to compute the phase hologram displayed on the SLM²⁴⁰. Next, we measure the prepared Rydberg density distribution using the optimized detuning sweep and update the weights of each tweezer to target a uniform population over the boundary, repeating this roughly five times to converge to the final weights. When updating the

weights, we re-run the WGS algorithm with the previous phase profile as a starting guess and, in the simulated fourier plane, update only the amplitudes of each wavevector component and not the phases, improving convergence⁸⁸. Finally, we optimize the overall magnitude of the light shift by minimizing the nematic order parameter Φ (Fig. 4h). Vertices that are connected to boundary atoms are excluded from Φ , since perfect dimer coverings on a lattice with edges and corners are expected to have higher dimer populations on the boundary¹³⁹.

E.14 PROBABILITY DISTRIBUTION OF DIMER CONFIGURATIONS

The number of allowed dimer configurations on the lattice grows exponentially with system size³⁰² such that, to study the population distribution in practice, we restrict to subsystems with a sufficiently small number of possible states.

Isolating small subsystems.

We postselect for snapshots (per subsystem) where the atoms on the outer bonds connecting the subsystem to the surrounding lattice are in the desired state. From these instances, we compute the population distribution over subsystem configurations that satisfy this imposed boundary condition, aggregating over all possible subsystems related by lattice translations. This boundary-fixing procedure requires an exponential number of samples but, crucially, the scaling is governed by the subsystem size rather than the lattice size. In experiment, the overhead is increased by postselection to exclude monomers, double-dimers, and atom loss, as required for comparison to the RK state and to reduce the number of configurations. In Fig. 5a, for no dimers around the subsystem, the percentage of kept data per subsystem is 7.2%, 1.8%, 0.4%, 0.3% and 0.04% for the 1×1 , 2×1 , 3×1 , 2×2 and 3×2 subsystem sizes, respectively. These correspond to 6, 11, 16, 18, and 27 sites inside the subsystem, plus 6, 8, 10, 10, and 12 sites on the boundary. The increased data acquisition rate, enabled by qubit reuse, makes these overheads sufficiently practical.

Thermal distribution.

The probability distribution of the state prepared using the modified sweep (ED Fig. 7a) is well-described by a thermal distribution $p_i \propto e^{-\beta E_i}$ where p_i is the observed probability of configuration i , β is the inverse temperature, and E_i is the energy of the configuration. Probabilities are defined within a given subsystem size, S , so that $\sum_{i \in S} p_i = 1$, and β is the same across all subsystems. The energy of each configuration is given by the sum of the vdW tail energies, excluding small corrections from Rydberg atoms outside the subsystem and perturbative ring exchange $\propto \Omega^6 / \Delta^5$. We fit the observed probabilities to these thermal probabilities with β as the only free parameter, obtaining the fitted thermal distribution shown in Fig. 5a (inset) and ED Fig. 7b.

E.15 CALCULATION OF CORRELATION FUNCTIONS

Dimer-string correlator.

The expectation value of Z loops in the RK state is ± 1 , with the sign given by the emergent Gauss's law³⁶⁹. For the dimer-string correlator, C_s , we restrict to loops that are (ideally) $+1$ in order to remove sign ambiguities, otherwise the two open strings whose product forms the closed loop will have the opposite sign. We use symmetric loop geometries in which the two open strings have the same length and take the (signed) geometric mean of their expectation values for the numerator of C_s . Given this, we exclude strings that terminate on boundary plaquettes where a symmetric closed loop cannot always be formed.

Rectified correlators.

The rectified dimer-dimer and dimer-string correlators for the honeycomb lattice are shown schematically in ED Fig. 9a,d. We follow the convention of Ref.⁴⁶⁷ where the three plaquette sublattices (a,b,c) determine the phases in the correlation functions. Concretely, the rectified dimer-dimer correlator, corresponding in field theory to Eq. (E.8), is $\tilde{C}_d = \langle \varphi^*(\mathbf{r}_i) \varphi(\mathbf{r}_j) \rangle$ where $\varphi(\mathbf{r}_i) =$




Figure E.8: Loops and strings in the quantum spin liquid. **a**, Time-dependent matrix product state (MPS) simulations of $\langle V \rangle$ and $\langle T \rangle$, compared to experiment. For the final state, MPS predicts $\langle T \rangle / \langle V \rangle = 0.92$. The lattice used for the numerics is three plaquettes tall and infinite in the horizontal direction, with the van der Waals interaction truncated at $\sqrt{7}a$ and a bond dimension of $\chi = 256$. **b**, Including the phase accumulation from vdW tails during the 100 ns mapping down gap reduces the expectation value in MPS numerics to $\langle T \rangle / \langle V \rangle = 0.76$, compared to $\langle T \rangle / \langle V \rangle = 0.47(1)$ in experiment; numerics cannot account for the Rydberg π -pulse in the three-level system. In experiment, plaquettes on the boundary (bdy) have higher $\langle T \rangle$ than in the bulk, where van der Waals tails are larger, but similar $\langle V \rangle$. For this plot only, the four corner plaquettes are excluded. **c**, Spatial plot shows the energy ratio in experiment. **d**, Evolution of experimental RK energies after state preparation, where the Rabi drive is turned off during the hold time. Decay of $\langle V \rangle$ is consistent with Rydberg T_1 error, while $\langle T \rangle$ decays faster as the different dimer configurations accumulate relative phases but the populations remain approximately unchanged. **e**, Perimeter-law decay of closed X and Z loops. Data for X loops are the same as in Fig. 5f. **f**, Normalization of open X strings. See Fig. 6c for the corresponding definition for Z strings. **g**, Open X strings are small and decay to zero with increasing string length, consistent with the state being described by a deconfined field theory^{386,429}. **h**, Dynamics of open strings. X strings, which create monomers at their endpoints, are finite when there is a high density of monomers. Z strings are trivially +1 for the initial product state. We show open strings with length 3 which are normalized by the length-6 closed loops that enclose a single plaquette. **i**, Effect of defects on the dimer-string correlator. Postselection on no monomers or double-dimers inside the closed loop (such that it is +1) slightly improves agreement between experiment and the RK prediction from Monte Carlo simulations.

$\sigma_{i,bc}^z + \omega \sigma_{i,ac}^z + \omega^2 \sigma_{i,ab}^z$ is the sum over the three bonds touching vertex i , $\omega = e^{2\pi i/3}$, and the bonds are labeled by the sublattices of the adjacent plaquette pair. The rectified dimer-string correlator, corresponding in field theory to Eq. (E.9), is $\tilde{C}_s = \langle \psi^*(\mathbf{r}_i) \psi(\mathbf{r}_j) \rangle$ where $\psi(\mathbf{r}_i) = S_i^a + \omega S_i^b + \omega^2 S_i^c$ is the sum over the three adjacent plaquettes and $\langle S_i S_j \rangle \equiv C_s(|\mathbf{r}_i - \mathbf{r}_j|)$ is the unrectified string correlator studied in Fig. 6d. In a translationally invariant system $\tilde{C}_{d,s}$ are real quantities; here we find the imaginary part is negligible for the $\tilde{C}_d$ but finite for $\tilde{C}_s$. We compute the correlators between all pairs of vertices, again excluding the boundary layer. These are then averaged over endpoints with the same radial separation, resulting in the rectified correlators shown in Fig. 6e and ED Fig. 9, 10.

For all correlation functions, error bars are calculated using bootstrap resampling with replacement over 20 resamples of the full dataset; increasing the number of resamples does not appreciably change the estimated uncertainties.

E.16 MEASURING MANY-BODY COHERENCES

To understand the effect of off-diagonal measurement error, we use MPS simulations of the preparation sweep, finding $\langle T \rangle / \langle V \rangle = 0.92$ for the final state. We then simulate the mapping down by evolving this state under the vdW tails interaction for 100 ns, which reduces the ratio to $\langle T \rangle / \langle V \rangle = 0.76$ (ED Fig. 8b). This is compared to $\langle T \rangle / \langle V \rangle = 0.47(1)$ in experiment, where an additional reduction is expected as the MPS numerics cannot account for the infidelity from the Rydberg π -pulse when mapping down in a three-level system (see also ED Fig. 1d). The decay of $\langle T \rangle$ as the mapping down gap time increases is several times faster than expected for single-qubit dephasing, consistent with the phase from vdW tail being the dominant error (ED Fig. 8d).

In addition, we use exact diagonalization on a smaller 3×3 unit cell lattice to simulate the effect of ground-Rydberg dephasing on the dynamics, including on-site Doppler dephasing and global detuning fluctuations, as well as radial and axial spread of the wavefunction corresponding to $\approx 10 \mu\text{K}$

atom temperature. These are modelled as local and global detuning shifts (randomly sampled from Gaussians of width $2\pi \times 0.043$ MHz and $2\pi \times 0.2$ MHz, respectively), and in-plane and out-of-plane positional disorder (randomly sampled from Gaussians of width $0.02a$ and $0.08a$, respectively). In ED Fig. 6d,e, we find that the RK energies are largely insensitive to these error sources and attribute this insensitivity to two factors. First, hemidiabatic preparation is largely insensitive to the exact start and end of the parameter sweep, since, by definition, it is too rapid to resolve such low-energy details of the spectrum³⁶⁹. As such, small global detuning shifts have little effect on the final fidelity of the prepared state. Second, the kinetic energy operators live in a decoherence-free subspace²⁶⁹ within the blockaded manifold, such that a global Z rotation (i.e., a detuning shift) does not change the measured kinetic energy, as we observe in experiment.

E.17 NUMERICAL METHODS

Here we provide details of the two numerical methods used to understand our results in Fig. 6 and ED Figs. 8–10.

Tensor Network Simulations.

Matrix product state (MPS) simulations of unitary dynamics are performed using the “W-II” matrix product operator method of Ref. ⁴⁶⁰ which is implemented in the Tensor Network Python (TeNPy) package¹⁹⁷. Specifically, we simulate the quantum dynamics under the Rydberg Hamiltonian, Eq. (7.1), for the standard $0.8 \mu\text{s}$ preparation protocol used in the main text. While the true vdW potential of the experiment is long-range, this is challenging to exactly simulate in MPS-based simulations as the bond dimension of the Hamiltonian’s MPO gets prohibitively large. Consequently, we study two approximations:

1. *PXP*. This enforces the nearest-neighbor blockade constraint but excludes long-range vdW tails, Eq. (E.1). In practice, the blockade is enforced exactly for the three atoms in the upper

triangles of the kagome lattice, thereby replacing these three atoms with an effective 4-state qudit, and is enforced energetically between these triangles with a nearest neighbor interaction $V_{ij} = 100 \Omega$.

2. *PXP plus tails.* In addition to enforcing nearest-neighbor blockade as above, we include long-range vdW tails up to a truncation distance of $R_{\text{trunc}} = \sqrt{7}a$ (i.e. up to the next-next-next-nearest neighbor contributions). The truncation distance is chosen such that observables change negligibly upon slightly decreasing this distance.

In both cases, our numerics are performed on strip geometries that are infinite in the x -direction and finite in the y -direction. Following the experiment, we use a boundary detuning of $\delta_{\text{bdy}} = 0.52 \Omega$ for sites along the top and bottom edges of the strip. We use a Trotter time step of $dt = 0.01 \Omega^{-1}$, which we find is sufficient to achieve convergence. To simulate the kinetic and potential energies (ED Fig. 8), we use the ‘PXP plus tails’ approximation with a three-plaquette-tall lattice and a bond dimension of $\chi = 256$. To simulate the Z -basis correlation functions (ED Fig. 9), we instead use the ‘PXP’ approximation to allow us to use a four-plaquette-tall lattice to compute these long-range observables. The bond dimension is $\chi = 350$. For such Z -basis observables, we find that neglecting vdW tails is a good approximation and the correlations closely agree with those on the smaller lattice that include vdW tails.

Classical Monte Carlo.

Following the mapping to a classical dimer gas, Eq. (E.5), observables for the RK state can be evaluated as classical averages over $\mathcal{N}$ dimer coverings $\mathcal{D}$ sampled uniformly by Monte Carlo. In addition to Z -basis observables, expectation values of off-diagonal observables $\hat{\mathcal{X}}$ can be obtained from averaging over transition matrix elements, with contributions only from finitely many $\mathcal{D}'$ connected to




Figure E.9: Diagonal correlators in the emergent U(1) gauge theory. **a**, Schematic of the rectified dimer-dimer correlator between two vertices. The sum involves nine (unconnected) dimer-dimer correlation functions, as highlighted, and the colors indicate the relevant phases, determined by the sublattices (a,b,c) of the neighboring plaquette pair. Here, $\omega = e^{2\pi i/3}$. **b**, Rectified dimer-dimer correlator in experiment. **c**, Comparison to MPS simulations of the experimental preparation protocol and to Monte Carlo simulations of the RK state on a 36×6 lattice with periodic boundary conditions. Inset shows real and imaginary components in MPS. MPS numerics use the PXP approximation, a bond dimension of $\chi = 350$, and a four-plaquette-tall lattice that is infinite in the horizontal direction. **d**, Schematic of rectified dimer-string correlator between two vertices. The sum involves nine unrectified dimer-string correlators between plaquette centers and the colors indicate the plaquette sublattices that define the phases. **e**, Rectified dimer-string correlator in experiment. While $\tilde{C}_s$ is a real quantity for a translationally invariant system, here we find an appreciable imaginary component (inset). **f**, Comparison of the dimer-string correlator to MPS and Monte Carlo simulations. Inset shows real and imaginary parts in MPS numerics, with the experimental data shown as faint markers. In **c** and **f** the experimental correlations are shown along the horizontal direction only to compare to numerics, while in **b** and **e** they are radially averaged over the lattice, excluding vertices on the boundary.

$\mathcal{D}$ via allowed moves along the support of $\mathcal{X}$:

$$\langle \hat{\mathcal{X}} \rangle = \frac{1}{\sum_{\mathcal{D}} 1} \sum_{\mathcal{D}, \mathcal{D}'} \langle \mathcal{D} | \mathcal{X} | \mathcal{D}' \rangle. \quad (\text{E.15})$$

For instance, the kinetic energy of a single plaquette is equivalent to its potential energy. We evaluate both diagonal and off-diagonal observables via classical Monte-Carlo sampling of $\mathcal{D}$ with uniform weight using a worm-update Markov Chain, which performs nonlocal loop updates and yields short autocorrelation times^{9,375}. Each worm update temporarily introduces a pair of monomers, propagates one of the monomers through the lattice via local dimer pivots, and terminates when the monomers reunite to form a dimer. As a result, each worm update flips the dimer configuration along a closed loop.

The worm update starts from a valid close-packed dimer covering $\mathcal{D}$ on the honeycomb lattice, where each vertex v is incident to exactly one dimer oriented along one of three lattice directions $l \in \{1, 2, 3\}$. The update then proceeds as follows.

- *Initialize.* Choose a vertex u_0 uniformly at random to be the *tail* of the worm. Let (u_0, v_1) be the (unique) dimer incident on u_0 , connecting it to the initial *head* of the worm at neighboring vertex v_1 . Remove (u_0, v_1) , creating two monomers at u_0 and v_1 .
- *Propagate head.* At step $i = 1, 2, \dots$, the *head* is at v_i , and (u_{i-1}, v_i) denotes the recently removed dimer entering v_i with orientation l_i .
 1. Reinsert a dimer emanating from v_i in a new orientation k_i , chosen with transition probability

$$P_{l_i \rightarrow k_i} = \begin{cases} 1/2 & l_i \neq k_i \\ 0 & l_i = k_i \end{cases} \quad \text{so that } \sum_{k_i} P_{l_i \rightarrow k_i} = 1, \quad (\text{E.16})$$

which satisfies local detailed balance at infinite temperature and forbids immediate backtracking. The inserted dimer (v_i, u_i) annihilates the monomer at v_i and connects it to a neighboring vertex u_i .

2. If $u_i = u_0$, terminate (see below). Otherwise, u_i is also incident to a (unique) pre-existing dimer (u_i, v_{i+1}) . Remove (u_i, v_{i+1}) to restore the dimer constraint at u_i and advance the *head* to the connecting vertex v_{i+1} .
- *Terminate.* The worm closes the loop at the first step $i = \ell > 1$ for which $u_\ell = u_0$. The final dimer insertion (v_ℓ, u_0) connects the head v_ℓ back to the tail u_0 and annihilates the monomer pair. This restores close packing and yields a new dimer covering $\mathcal{D}'$, which differs from $\mathcal{D}$ by dimers flipped along a closed loop of length 2ℓ .

Monte Carlo samples are obtained by iterating this worm update to generate a Markov chain $\mathcal{D} \rightarrow \mathcal{D}' \rightarrow \dots$.

In this work, we sample Monte Carlo trajectories initialized from one of the three columnar dimer coverings, which maximize the number of flippable plaquettes³⁰³. We consider periodic honeycomb lattices of size $L_x \times L_y$ in units of plaquette separation $(1/2a)$, containing $3 \cdot L_x \cdot L_y$ dimers. For each lattice, we average not only along a single trajectory but also across independent runs with identical initialization and different random seeds. For the 36×6 lattice, we generate 10 trajectories of 19,600,001 to 30,100,001 samples each, for a total of 229,300,010 samples. For the 36×9 lattice, we generate 10 trajectories of 11,600,001 to 19,700,001 samples each, for a total of 142,600,010 samples. For the 36×12 lattice, we generate 10 trajectories of 8,500,001 to 9,400,001 samples each, for a total of 90,300,010 samples. For the largest lattice (36×36), we average over 10 trajectories of 2,800,001 to 3,400,001 samples each, for a total of 32,600,010 samples. These data are used to compute RK state observables as a function of system size.

For the RK state on a torus, there are four topological sectors $(\pm 1, \pm 1)$, specified by the Z string

eigenvalues of the two non-contractible loops wrapping around the horizontal and vertical directions of the lattice. For both the dimer-dimer and dimer-string correlators, we find that conditioning on a single topological sector enhances the estimates correlations relative to sector-averaged estimates, and reduces finite-size effects. We show this in ED Fig. 10d where the data shown in the $(+1, +1)$ topological sector were obtained by retaining 1,550,589 to 1,753,196 samples per trajectory across 10 trajectories, for a total of 16,805,690 samples. Within a fixed sector, the correlators are symmetric about $L_x/2$ due to periodic boundary conditions. When averaging over sectors, the dimer-string correlator continues to decay beyond half-length as the non-contractible loops average to zero. For comparison to experiment, we average over all four sectors, as the sector is not well-defined for open boundary conditions.




Figure E.10: Finite-size scaling in numerics and experiment. **a,b** Finite-size scaling of rectified correlators in Monte Carlo simulations of the RK state. The correlators converge towards the predicted $1/r^2$ and $1/r^{1/2}$ decays at large system sizes. Dashed lines are guides to the eye showing exponential decays with correlation length $\xi(L_y) \propto L_y$. Experiment is well-captured by the 36×6 data (ED Fig. 9). **c**, Relation between the magnitudes of the rectified dimer-dimer and dimer-string correlators in Monte Carlo simulations. The ratio of 4 from $U(1)$ gauge theory holds for all system sizes until the dimer-dimer correlator plateaus. **d**, Effect of topological sector in Monte Carlo with periodic boundary conditions. Within the $(+1, +1)$ topological sector, defined as $+1$ eigenvalues of the two non-contractible Z -strings, both correlators are symmetric about half-length (dashed line), as enforced by translation and reflection symmetries. Averaging over all four topological sectors suppresses the correlators overall, and leads to a distance-dependent decay of the dimer-string correlator beyond half-length. In all other plots, correlators are averaged over sectors. Data are for a 36×36 lattice. **e**, Finite-size scaling in experiment. Upon varying the lattice height from $L_y = 3$ to $L_y = 6$, at fixed width of $L_x = 13$, we find the rectified dimer-string correlator is largely independent of system size while the rectified dimer-dimer decays until the correlations plateau. This plateau arises from explicit breaking of the lattice's $\mathbb{Z}_3$ symmetry and decreases in magnitude as the lattice height increases. For all other data in this work, $L_y = 6$. **f**, The field theory prediction (solid line) is in good agreement with the data at short distances for all system sizes until the plateau in the dimer-dimer correlations is reached.

F

Supplementary information about further improvements to entangling CZ gate fidelity

F.1 EXPERIMENTAL SYSTEM

Our experiments begin by stochastically loading rubidium-87 atoms from a magneto-optic trap (MOT) into an array of static 852-nm optical tweezers generated with a spatial light modulator (SLM, Hamamatsu X13138-02) using 795-nm lambda-enhanced gray molasses for enhanced loading of $\sim 75\%$ ⁶⁶. We rearrange the atoms using a second set of 852-nm optical tweezers generated from two-dimensional acousto-optic deflectors (AODs, DTSX-400, AA Opto-Electronic). For benchmarking gates and performing high-fidelity quantum circuits, we use a reservoir of 374 atoms, and bring up to 20 atoms at a time into the entangling zone (Fig. F.1a). We image the atoms with a 0.65-NA objective (Special Optics) onto a CMOS camera (Hamamatsu ORCA-Quest C15550-20UP). Compared to our previous work benchmarking high-fidelity gates in Ref.¹²⁹, here we only use 3D polarization gradient cooling (PGC), as we find that errors from atomic temperature form a small part of our overall error budget.

The qubit states are encoded in the $m_F = 0$ states in the ground hyperfine manifolds, $|0\rangle = |F=1, m_F=0\rangle$ and $|1\rangle = |F=2, m_F=0\rangle$, with coherence time $T_2 > 1\text{s}$ ⁴². At the start of the quantum circuits, the qubits are prepared in $|0\rangle$ using Raman-assisted optical pumping. A global




Figure F.1: Architecture for high-fidelity gates and circuits. (A) Atom array layout utilized in this work, with a large reservoir zone for fast cycle rates, an entangling zone with gate sites demarcated (dark sites are the subset used for randomized benchmarking), and additional sites in the entangling zone used for state-selective readout⁴². The top five rows of sites are used for initial loading and rearrangement of atoms into the reservoir. (B) Level diagram for Rydberg excitation to $n = 53$ Rydberg state $|r\rangle$. The lasers used for other aspects of the architecture, including for Raman single-qubit gates and state-selective readout, are described in our previous work⁴².

Raman beam illuminating the entire array, including atoms in all zones, is used to implement single-qubit operations²⁶², with an intermediate state detuning of ~ 550 GHz to reduce scattering error and achieve fidelities for BB1 composite pulses of $\sim 99.99\%$ ⁴².

We implement loss-resolved qubit readout, as described in Ref.⁴². In this approach, state-selective qubit readout is achieved using a one-dimensional state-selective optical lattice potential, which pins one of the two qubit states and allows us to move the other with a moving AOD trap. We split the two qubit states into two different vertical positions in the entangling zone, and then we move them in parallel to the bottom of the reservoir zone to image them. This approach retains the majority of atoms, enabling us to reuse the qubits in subsequent repetitions, in addition to refilling from the large reservoir. Throughout this work, we perform 40-80 repeated experiments before loading an atom array from the MOT.

The entangling zone used for benchmarking CZ gates and realizing high-fidelity circuits consists of 17 total gate sites. For benchmarking, we typically use a sparse array of 8 gate sites, in order to reduce the effect of decay of Rydberg states to nearby P states on the measured fidelity¹²⁹. Since

this effect does not impact quantum circuits, in which CZ gates are spaced further apart in time due to additional time overhead from movement⁴², we use the denser arrangement of gate sites for the circuit implementations.

F.1.1 RYDBERG EXCITATION

For realizing CZ gates, we excite atoms from the qubit state $|1\rangle$ to the Rydberg state $53S_{1/2}$, using a smooth-amplitude gate (see section below for details)¹²⁹. We use circularly polarized 420-nm and 1015-nm light in a two-photon transition, with an intermediate state detuning of 7.8 GHz (Fig. F.1b). We operate with the pair of atoms within each gate site separated by about $1.73\mu\text{m}$, where the blockade energy is around 950 MHz. We pulse the traps off for a few hundred nanoseconds during the gate.

The laser systems for Rydberg excitation are the same as reported in Ref.¹²⁹. By operating red-detuned from the intermediate state, we ensure that the σ^- polarized 420-nm light is further detuned for the other qubit state $|0\rangle$. The σ^+ polarized 1015-nm laser primarily couples to $|r\rangle$ with $m_J = -1/2$ in this magnetic field configuration, but also couples to the other $m_J = +1/2$ Rydberg state $|r'\rangle$ with a two-photon Rabi frequency suppressed by a factor of three. To mitigate the effect of this spurious coupling, we increase our bias magnetic field strength to $B = 13.7$ G, with a larger 38.5 MHz splitting between the two Rydberg levels than in our previous work. As in previous work, staying close to the dark state during the two-photon Rydberg excitation is crucial for reducing scattering error¹²⁹. This suppression is achieved by choosing the detuning of the Rydberg gate to start with a positive sign. As a result, the peak Rabi frequency occurs at negative detuning. The combination of magnetic field direction and laser polarization ensures that at peak Rabi frequency, we remain far detuned from $|r'\rangle$, to reduce unwanted excitation to this Rydberg state.

We use high Rabi frequencies reaching a peak of $\Omega \approx 2\pi \times 17$ MHz in this work. The effective single-photon Rabi frequencies of the two legs of the two-photon transition are $\Omega_{420,\text{peak}} =$

$2\pi \times 670$ MHz and $\Omega_{1015} = 2\pi \times 460$ MHz for the peak 420-nm and 1015-nm intensities, respectively. The two-photon Rabi frequency is given by $\Omega = \sqrt{3/4} \frac{\Omega_{420,\text{peak}} \Omega_{1015}}{2\Delta}$, where the prefactor of $\sqrt{3/4}$ comes from the presence of multiple intermediate states¹²⁹. To achieve these high Rabi frequencies, we generate circularly symmetric beams with Gaussian waists of roughly 30 microns and 35 microns for the 420-nm and 1015-nm beams, respectively, to cover a one-dimensional entangling zone. The 420-nm laser beam profile is Gaussian and is created by displaying a phase grating on the SLM (Hamamatsu). The 1015-nm laser beam profile is a circular tophat profile created by the SLM (Hamamatsu). These Rydberg beams are stabilized on cameras between each sequence, using active feedback to piezo-controlled mirrors.

F.1.2 CZ GATE PROFILE AND CALIBRATION

We use the following parametrization for the smooth-amplitude gate in this work, with phase profile $\varphi(t)$ and Rabi frequency profile $\Omega_{420}(t)$:

$$\begin{aligned}\Omega_{420}(t)/\Omega_{420,\text{peak}} &= \Omega_0 + A \text{sech}[\omega_\Omega \tau]^\alpha, \\ \Omega_0 &= \frac{A_{\min} \cosh[\omega_\Omega T/2]^\alpha - A_{\max}}{\cosh[\omega_\Omega T/2]^\alpha - 1}, \\ A &= \frac{(A_{\max} - A_{\min}) \cosh[\omega_\Omega T/2]^\alpha}{\cosh[\omega_\Omega T/2]^\alpha - 1}, \\ \varphi(t) &= \varphi_0 \tau + B \tanh(\lambda \tau),\end{aligned}$$

where $\tau = t - T/2$.

Theoretical parameters for ideal smooth amplitude gates are given in our earlier work¹²⁹. The following parameters are the experimental values used in Fig. 2b:

$$\begin{aligned}
A_{\min} &= -0.1565, & B &= -2\pi \times 0.3569, \\
A_{\max} &= 0.9622, & \lambda &= 1.1195 \, \Omega, \\
\omega_{\Omega} &= -0.2962 \, \Omega, & \delta_0 &= 0.775 \, \Omega, \\
\alpha &= 0.3,
\end{aligned}$$

For the other experiments, the parameters are similar, with small adjustments from performing recalibration. To compensate for pulse imperfections arising from our experimental setup, we additionally apply amplitude pulse corrections in the form of Chebyshev polynomials $T_n(t) = \cos(n \arccos(t/t_{\max}))$ ²⁸², which we sum together with coefficients c_n and multiply by the amplitude profile. We find that most of these corrections have optimal coefficients near zero (< 0.01). There is, however, a finite value of $c_1 = -0.214$ for the linear correction $T_1(t) = t/t_{\max}$. Adding this term improves the fidelity at the $\sim 0.02 - 0.03\%$ level. We attribute its nonzero optimal value to imperfections in our experimental setup which make the pulse asymmetric.

We calibrate the CZ gate experimentally by starting from close to a theoretically optimal gate, and then scanning individual gate parameters to maximize the return probability of the echo RB sequence. Compared to our previous work¹²⁹, here we are able to detect loss errors using state-selective readout, and we observe that this can be helpful in gaining intuition for the calibration process. First, the loss postselected fidelity is very sensitive to certain parameters such as detuning, but relatively insensitive to other parameters such as pulse duration. This observation is consistent with the role of different gate parameters: certain parameters more strongly affect the gate trajectory, which determines the population returning to the ground hyperfine manifold and loss error, whereas other parameters primarily determine the accumulated phase during the gate and the resulting phase error. Relatedly, we observe that at the beginning of calibration, the optimal param-




Figure F.2: Readout methods and spatial homogeneity for echo RB. Echo RB curves for the data plotted in Fig. 8.2c. Measured return probability as a function of number of CZ gates (A) with and without a single-qubit π pulse before state-selective readout (B) with pushout versus state-selective readout (SSR). (C) CZ fidelity across the gate sites used randomized benchmarking.

ters for raw and loss-postselected fidelity can differ, but as the gate becomes calibrated, the optimal parameter values converge to be the same. We find that this convergence can be a useful proxy for determining whether the gate is optimally calibrated. While in this work we find that our simple calibration method works robustly, more optimal methods could improve the speed of calibration, such as using a more ideal Hessian basis for scanning³¹¹, and automated calibration via Nelder-Mead optimization or machine learning.

F.2 BENCHMARKING PROTOCOLS

We use global randomized benchmarking (RB) sequences in this work, mostly focusing on RB sequences which include global single-qubit X gates between pairs of CZ gates. We find that this benchmarking scheme is a robust way to calibrate and track the fidelity of CZ gates over time. The results are averaged over 8 CZ gate sites, for which the gate performance is relatively homogeneous (Fig. F.2c). To assess the effects of our readout protocol on the measured fidelity, we perform two variants of this RB sequence: (a) with an additional single-qubit Raman π pulse before readout and (b) using resonant blowout of $F = 2$ instead of state-selective readout. We observe that all three

sequences result in consistent measured fidelities (Fig. F.2a-B).

To check the sensitivity of the measured gate fidelities to the exact benchmarking sequence, we additionally measure the fidelity in the symmetric stabilizer benchmarking (SSB) sequence developed in Ref. ⁴²⁵. This benchmarking sequence was designed to be insensitive to single-qubit gate errors compared to the echo RB sequence, for which sufficiently large single-qubit error rates can lead to an underestimation of the CZ fidelity ⁴²⁵. For our experiment, we expect that the single-qubit fidelity is sufficiently high, and that the two benchmarking schemes should offer consistent results ⁴²⁵. In contrast to the echo RB sequence, in which the global X gate between a pair of consecutive CZ gates eliminates the effect of accumulated single-qubit phase during the gate pulse, the SSB sequence is sensitive to the single-qubit phase of the CZ gate, which arises both from the gate trajectory and from the 420-nm differential lightshift on the qubit states. We first calibrate the CZ gate using echo RB, and then optimize the single-qubit phase $Z(\varphi)$ in the SSB sequence (Fig. F.3b). We verify that the optimal single-qubit phase is independent of the number of CZ gates, which we find useful for verifying that there is no bias from other possible spurious single-qubit phases in the sequence. We find that the CZ gate fidelity measured with the SSB sequence agrees with that obtained from our echo RB sequence, which shows that the fidelity can be maintained in applications where an echo pulse cannot be applied between adjacent CZ gates.

Finally, we test two different trap configurations for benchmarking the gate. For the CZ gate randomized benchmarking data in Fig. 2, we place two SLM traps close to each other to form the gate site. In quantum circuits with reconfigurable atom positions and connectivity, we typically perform gates between static and moving arrays of traps. In Fig. 4a, we check that the fidelity is comparable when the gate sites are composed of an AOD and an SLM trap to when both sites are SLM traps. The AOD-SLM trap configuration is used when performing CZ gates for the cluster state and non-local circuits presented in Figs. 4-5.




Figure F.3: Symmetric stabilizer benchmarking. (A) Circuit used for symmetric stabilizer benchmarking (SSB); figure adapted from Ref. ⁴²⁵. (B) Optimization of single-particle phase ϕ after the CZ gate, here looking at the return probability after 12 CZ gates. (C) Comparison of SSB and echo RB approaches, taken interspersed with each other to avoid any potential bias from experimental drift. Measured fidelities are within error bars of each other.

F.3 CZ GATE ERROR SOURCES

The experimental gate errors can be largely explained by our numerical error model. We compute the errors in the system through a combination of exact diagonalization and simulation of incoherent error sources ^{129,425}, with relevant parameters such as T_1 and T_2^* informed by experimental measurements. We numerically model the Hamiltonian of our two-photon Rydberg excitation of ^{87}Rb atoms, which includes the qubit states $|0\rangle$, $|1\rangle$, three excited states in the $6P_{3/2}$ hyperfine manifold with $m_F = 1$ and $F = 1, 2, 3$, and two Rydberg states with $m_J = +1/2$ and $m_J = -1/2$ (see diagram in Fig. F.1b). We apply the exact smooth-amplitude gate pulse in the numerical model, ignoring the effect of the finite rise time of the AOM. A finite rise time effectively applies a Gaussian window function in the frequency domain to the amplitude profile. The measured rise time of the 420-nm AOMs is around 10ns, therefore, the smooth pulse profile does not contain high-frequency components which would fall outside the window function.

To model the incoherent error sources, intermediate-state scattering, and Rydberg lifetime, we compare the full density-matrix simulation, including an additional leakage state ¹²⁹, with calcula-

Parameter	Value
Peak two-photon Rabi frequency	$\Omega = 2\pi \times 17 \text{ MHz}$
Single-photon 420-nm Rabi frequency	$\Omega_{420} = 2\pi \times 670 \text{ MHz}$
Single-photon 1015-nm Rabi frequency	$\Omega_{1015} = 2\pi \times 460 \text{ MHz}$
Intermediate state detuning	$\Delta = 2\pi \times 7.8 \text{ GHz}$
T_2^* for $ 1\rangle$ - $ r\rangle$ with no Rydberg light	$4.7 \mu\text{s}$
T_2^* for $ 1\rangle$ - $ r\rangle$ with 420-nm light on	$3.8 \mu\text{s}$
T_2^* for $ 1\rangle$ - $ r\rangle$ with 1015-nm light on	$2.8 \mu\text{s}$
T_1 of Rydberg state (53S)	$72 \mu\text{s}$

Table F.1: Experiment parameters. Parameters which are used for our numerical simulations of the experiment, informed by our measurements. The value for T_1 is given by the Alkali Rydberg Calculator (ARC)⁴⁷⁰.

Error source	Smooth-amplitude gate in this work
Scattering $ 1\rangle$	0.028%
Scattering $ 0\rangle$	0.018%
Rydberg T_1	0.046%
1015-nm T_2^*	0.008%
420-nm T_2^*	0.001%
Doppler T_2^*	0.007%
Rydberg blockade	0.001%
Rydberg $m_J = \frac{1}{2}$	0.002–0.030%
Total fidelity	99.861–99.889%

Table F.2: Error budget for the smooth-amplitude gate used in this work. Scattering errors for the qubit state $|0\rangle$ originate from off-resonant scattering from the 420-nm light. The range for the contribution from coupling to the other Rydberg m_J level $|r'\rangle$ relates to optimization of the gate profile (see text).

tions using the decay rate equation with the corresponding population in the intermediate states and Rydberg state. We find that the difference between these two methods is negligible. We measure the Rydberg lifetime T_1 to be roughly $80 \mu\text{s}$ in our experiment (with Rydberg lasers off), similar to the theoretical prediction of $72 \mu\text{s}$ given by the Alkali Rydberg Calculator⁴⁷⁰. We assume the lifetime of the intermediate states to be 110 ns ³⁶⁵. To model the effects of intensity, frequency, and position noise of the Rydberg beams, we adapt the Fidelity Response Theory (FRT) developed in Ref.⁴²⁵ to the two-photon excitation configuration. One key difference for two-photon excitation

is that intensity noise can be converted into frequency noise because of the large ground-Rydberg lightshift. We numerically find that the relative intensity noise (RIN) and frequency noise of both the 1015-nm laser, which is filtered by a ULE cavity, and the 420-nm laser, which is frequency-doubled from a Ti-Sapphire laser¹¹⁵, lead to negligible errors. Instead, frequency noise of the gate pulse is dominated by position fluctuations of the laser beams, which leads to time-dependent variations in the lightshifts. This is validated by direct measurements of the T_2^* coherence between $|1\rangle$ and $|r\rangle$ (Table F.1). In particular, we measure independent decoherence contributions from finite atomic temperature, compared to frequency noise induced by the 420-nm and 1015-nm lasers through three different $|1\rangle$ - $|r\rangle$ Ramsey T_2^* measurements. During the wait time between the two $\pi/2$ Rydberg pulses we either (a) keep both Rydberg lasers off, (b) keep only the 1015-nm laser on, or (c) keep only the 420-nm laser on. From the T_2^* measured in the absence of any Rydberg lasers, we estimate the atom temperature to be around 12 μ K from a Doppler decoherence model, similar to estimates from drop-recapture measurements⁴²⁶.

We additionally model the position errors of the atoms, which are composed of the motional spread of atoms at finite temperatures and variation in tweezer trap positions. The latter may originate from misalignment or relative focus between the two tweezers, in particular for the AOD-SLM trap configuration, which leads to a variation of the interatomic spacing. As a result, the blockade energy, which scales approximately as r^{-6} , will vary across the array, for example causing different detunings due to the light shift from the doubly-occupied Rydberg pair state $|rr\rangle$. At 12 μ K with a radial trapping frequency of around 80kHz and axial trapping frequency of around 20kHz, the atom motional spread has a standard deviation of approximately 70nm in the radial direction and 270nm in the axial direction. To estimate the infidelity due to position error in the SLM-SLM trap configuration, where the position spread mainly comes from finite temperature effects, we sample from a Gaussian distribution with a standard deviation of 100nm (400nm) to generate radial (axial) position errors. We obtain the blockade energy $V_{\text{blockade}}(r)$ as a function of gate site spacing r us-




Figure F.4: Understanding leakage in CZ gates. Measurement of the buildup of population in $F = 1$ during the echo RB sequence. At the end of the RB sequence, population in $F = 2$ is expelled from the trap with near-resonant pushout light. Then, a single-qubit π pulse transfers $|0\rangle$ to $F = 2$, and we perform state-selective readout. Plotted is the residual population in $F = 1$ as a function of the number of CZ gates, consisting of buildup of leakage population. The green curve is for the raw data, and the tan curve is for postselecting on the other atom being in the correct state at the end of the sequence.

ing the Alkali Rydberg Calculator⁴⁷⁰. We generate 10,000 sampled blockade energy values and use them in the gate pulse Hamiltonian optimized at $V_{\text{blockade}} = 950$ MHz to obtain a mean infidelity of around 0.001% resulting from position errors. In the AOD-SLM configuration, we scan the focal plane of the SLM tweezers and ensure the traps are aligned in the axial direction. We further fine tune the relative in-plane positions of the AOD-SLM tweezers by optimizing on site-resolved fidelities.

We also estimate the error due to off-resonant excitation to the $m_J = +1/2$ Rydberg state $|r'\rangle$. This coupling grows rapidly as we increase the Rabi frequency (Fig. F.6a,c). This can be alleviated by increasing the Zeeman splitting between the Rydberg states and fine tuning the phase profile of the gate. With 38.5 MHz Zeeman splitting and $2\pi \times 17$ MHz peak Rabi frequency, we numerically find that the error due to coupling to $|r'\rangle$ can be reduced to below 0.005% with further optimization of the phase profile. In particular, the off-resonant excitation fraction is not simply given by $\Omega^2/4\Delta_{\text{Zeeman}}^2$ but can form a nearly closed trajectory with appropriately chosen phase profiles. We specify a range for the $m_J = +1/2$ error, corresponding to how well-optimized the pulse profile is.

F.3.1 UNDERSTANDING LOSS AND LEAKAGE ERRORS

Loss errors in Rydberg gates can arise from several different sources, including Rydberg decay, intermediate state scattering, miscalibration, T_2^* , and coupling to the other Rydberg level $|r'\rangle$. Additionally, decay from the Rydberg state to lower-lying excited states and off-resonant scattering from the intermediate state can potentially result in leakage errors into other hyperfine m_F states in the ground-state hyperfine manifold. To estimate the breakdown of different error types, we consider a simple numerical model described in Ref. ¹²⁹. We assume that any remaining population in the Rydberg state, which can arise from error in the closure of the gate, blackbody-mediated decay to a nearby Rydberg state, or coupling to the other Rydberg state $|r'\rangle$, is converted to loss. This loss conversion occurs because the tweezer traps are turned on within 100ns after the end of the gate

pulse, and the Rydberg state is anti-trapped. With this numerical model, for the parameters used in our experiments, we find that the total numerical error of 0.127% can be decomposed into 0.061% loss error, 0.035% potential leakage error, 0.024% Pauli-Z error, and 0.003% Pauli-X(Y) error. We note that this error budget may overestimate the leakage error and more precise simulations would need to account for several additional effects. For example, excited states are anti-trapped which can lead to loss when the traps are turned back on depending on the exact timing, excited state lifetime, and branching ratios. Additionally, leakage population can be converted to loss by re-excitation in subsequent CZ gates in the benchmarking circuit. Below, we present experimental characterizations which can guide these more detailed numerical simulations.

To understand the effect of leakage into other m_F sublevels, we performed several different measurements. First, as described previously, we compare our randomized benchmarking results with and without the addition of a single-qubit π pulse on both qubits before readout (Fig. F.2a). If the leakage errors are strongly asymmetric between $F = 2$ and $F = 1$ sublevels, we would measure a different CZ gate fidelity for these two protocols. We instead find that the outcomes are consistent with each other.

Second, in Fig. F.4, we measure leakage into $F = 1$ sublevels from the CZ gates^{42,5}. For this, we perform the echo RB sequence with up to 40 CZ gates, and afterwards push out $F = 2$ and perform a single-qubit π pulse to transfer $F = 1, m_F = 0$ to $F = 2, m_F = 0$. Then we use our state-selective readout to separately read out $F = 1$ and $F = 2$ populations, in which $F = 2$ population is optically pumped to the dark state of the state-selective lattice and $F = 1$ population remains in a bright state⁴². We plot the mean population measured in $F = 1$, which consists of population left in the leakage states $m_F = +1$ and $m_F = -1$. We find that the population increases as 0.008(1)% per qubit per CZ. However, often a leakage error on one atom results in a different type of error on the other atom and therefore is detected in our randomized benchmarking. To understand this, we postselect on the other atom in the pair having correctly returned to the $|0\rangle$ state at the end of

the sequence, and find that the leakage population is $0.004(1)\%$ per qubit per CZ. This suggests that our measured fidelity might be overestimated by $\sim 2 \times 0.004\%$. Future work can employ similar techniques combined with population transfer between m_F levels to map out the full leakage populations, in particular for the $F = 2$ manifold. Additionally, studying how leakage is converted to loss through re-excitation to Rydberg states in subsequent CZ gates is another important avenue for understanding leakage and loss in this high-fidelity regime and their impact in fault-tolerant quantum circuits. Practically, we observe no signature of large-scale correlated loss events down to the $\sim 10^{-5} - 10^{-6}$ level in the cluster state and nonlocal circuits (Fig. F.8), suggesting that leakage does not lead to detrimental correlated loss events in our complex quantum circuits.

We additionally consider possible leakage from the 420-nm light driving m_F -changing Raman transitions. Polarization impurity alone is insufficient to cause substantial leakage, since a mixture of σ^+ and σ^- light cannot drive m_F transitions in the limit of being far-detuned from $6P_{3/2}$ ²⁶². By contrast, if the external magnetic field is at an angle with respect to the beam's propagation axis, then the 420-nm light has a π polarization component which can drive transitions with $\Delta m_F = \pm 1$. Our choice of σ^- polarization results in a vector lightshift that increases the Zeeman splitting between the m_F levels, which is favorable for reducing this effect. We experimentally observe coherent oscillations between m_F levels if we intentionally misalign the magnetic field, consistent with numerical simulations. However, when we carefully align the magnetic field, this effect becomes negligible and does not factor into our error budget.

F.3.2 UNDERSTANDING SENSITIVITY TO EXPERIMENTAL DRIFT

In Fig. 2, we demonstrate that we can robustly achieve roughly 99.85% fidelity gates for 10 hours without recalibration of the gate. To understand the requirements needed to achieve such stability, we perturb the system and measure how much the fidelity decreases. In particular, for our CZ gates, the main sources of drift are the positions and powers of the Rydberg beams; we find that per-

formance can generally be recovered by tuning the position back and adjusting the powers slightly. For the 1015-nm Rydberg beam, due to the large (roughly 50 MHz) differential light shift on the $|1\rangle - |r\rangle$ transition, the stability requirements are at roughly 0.1% for intensity and a few hundred nanometers for beam position for our choice of Rydberg parameters and beam size (Fig. F.5). Since the 420-nm beam is lower in intensity, the gate is less sensitive to perturbations for that beam.

The experimentally measured change in fidelity can be well captured by our numerical modeling with no free parameters other than offsets. We observe that the fidelity change as a function of laser power can be described by a quadratic fit. The sensitivity to the 420-nm laser intensity is approximately a factor of four weaker than the sensitivity to the 1015-nm laser intensity, which is consistent with the fact that the lightshift from the 420-nm beam is about four times smaller than the lightshift from the 1015-nm beam. This suggests that the primary mechanism for such fidelity reduction is the detuning change from the lightshift. Next, we consider the effect of changing the Rydberg beam positions (Fig. F.5c-D). For a Gaussian laser beam profile, at the center of the beam where the atoms are located, the intensity changes quadratically as a function of beam displacement, which would result in a quartic response in CZ fidelity to beam position when the detuning is suboptimal. However, the fluctuations of lightshift due to position noise also get worse once atoms are not at the center of the beams and, therefore, the resulting fidelity change can deviate from a quartic fit. Our numerical simulation takes into account the deteriorating T_2^* due to increased sensitivity to pointing noise when the beam is misaligned, and agrees with the experimental data.

F.3.3 CHOICE OF n

We choose the same principal quantum number $n = 53$ for the Rydberg state as in previous work¹²⁹. Many properties of the Rydberg state depend on n , including the Rydberg state lifetime ($\propto n^3$ for radiative decay, $\propto n^2$ for black-body decay), the dipole matrix element of the 1015-nm transition ($\propto n^{-3/2}$), the blockade strength given by the van der Waals interaction ($\propto n^{11}$), and




Figure F.5: Understanding sensitivity of CZ fidelity to drifts. Sensitivity of the measured CZ gate fidelity to Rydberg beam (A-B) powers and (C-D) positions. Red points are loss-postselected, and blue points are raw CZ fidelities. Faint colored curves are quadratic fits for beam powers and quartic fits for beam positions, and the black curves are from numerical simulations with offsets added as the only free parameter (see text for more details). (E) Positions of Rydberg beams on camera used for tracking and active feedback, averaged between each experimental repetition.

sensitivity to electric fields ($\propto n^7$). It may seem beneficial to increase n until sensitivity to electric fields starts to limit coherence time T_2^* , which reduces errors due to increased Rydberg lifetime and blockade strength. However, for two-photon Rydberg excitation, our error model indicates a more stringent tradeoff between coherence time, which is limited by Rydberg beam lightshift fluctuations, and suppression of Rydberg decay. The polarizabilities at 1015nm of the ground state, and of the Rydberg state at our large intermediate state detuning of 7.8 GHz, remain nearly constant as the transition dipole matrix element changes with n . For constant laser intensity, the error contribution from the 1015-nm lightshift therefore scales as $1/\Omega^2 \propto n^3$. In contrast, the error contribution from Rydberg decay scales as $n^{-1/2}$ for black-body decay and as $n^{-3/2}$ for radiative decay. We find that our choice of n is nearly optimal in balancing these different error sources. In particular, this suggests that the position and intensity stability of the 1015-nm laser need to be further improved for higher n to be favorable. This may provide an alternative route to increasing Rabi frequency, in




Figure F.6: Path to higher CZ fidelity. Numerical projection of the CZ fidelity for improved parameters. **(A)** CZ error as a function of Rabi frequency (relative to current Rabi frequency Ω_0), with and without contributions of the other Rydberg m_J state $|r'\rangle$. **(B)** CZ error as a function of detuning, relative to the current detuning Δ_0 , where the optimum balances contributions from Rydberg T_2^* and intermediate state scattering. The other m_J level is not included for this plot. **(C)** CZ error as a function of Δ and Ω , compared to our current Rabi frequency Ω_0 and intermediate state detuning Δ_0 (colorbar is log-scale). Dashed lines denote equal beam intensity, for $1\times$, $2\times$, and $3\times$ current intensity I_0 used in this work. The effect of coupling to $|r'\rangle$ prevents going to much higher Rabi frequencies. **(D)** Projection of CZ fidelity improvements achievable with higher laser intensities, as a function of Δ and Ω . For this plot, we consider a configuration with $2\times$ improved 1015-nm T_2^* and no coupling to $|r'\rangle$. The white stars mark the optimal configuration that minimizes CZ error for each intensity, with the associated CZ error labeled.

order to suppress Rydberg decay error. However, we additionally note that in circuit contexts, an extra benefit of using lower n is the ability to pack gate sites more densely in the entangling zone, especially within a two-dimensional region, owing to the weaker interaction strength.

F.3.4 PATH TO FURTHER IMPROVEMENTS IN CZ FIDELITY

The error budget and our experimental characterization point to a few areas where Rydberg gates could be improved in the future. The error budget is dominated by decay from the Rydberg state due to finite lifetime and by the intermediate excited state in the form of off-resonant scattering. At the current laser intensities, the finite coherence time T_2^* between the qubit state $|1\rangle$ and the Ryd-

berg state $|r\rangle$, which is limited by finite atom temperature and intensity fluctuations of the two Rydberg lasers, also contributes to gate errors. Residual errors mainly come from off-resonant coupling to the other Rydberg level $|r'\rangle$ with $m_J = +1/2$.

Using our numerical error budget, we project that a raw CZ gate fidelity at the level of 99.9-99.95% should be achievable with this gate scheme. It may seem straightforward to reduce the dominant errors with higher laser intensity. In particular, a higher Rabi frequency reduces the duration of the gate pulse, thereby decreasing error from Rydberg state lifetime. Additionally, a larger intermediate state detuning suppresses off-resonant scattering error during the gate, for which higher laser intensity is necessary to maintain the same Rabi frequency. However, in practice, there are additional experimental challenges, as illustrated in Fig. F.6c. As the Rabi frequency is increased, we begin to couple strongly to the other Rydberg state $|r'\rangle$ which is shifted up in energy by the bias magnetic field (Fig. F.6a,C). This off-resonant coupling constitutes a hard limit on our ability to increase Rabi frequency. However, this effect can be mitigated in the future, for example through the use of a higher magnetic field to increase the Zeeman splitting or by coupling through the $6P_{1/2}$ intermediate state where selection rules only allow coupling to one m_J level.

As the intermediate state detuning is increased, the T_2^* error eventually increases considerably due to the high 1015-nm laser intensities causing large lightshifts on the $|1\rangle - |r\rangle$ transition. The contribution from finite coherence time scales as $(\Omega T_2^*)^{-2}$ or $(\partial f/\Omega)^2$, where Ω is the two-photon Rabi frequency and ∂f is the integrated frequency noise. In the case where ∂f is dominated by the 420-nm lightshift fluctuation, it is proportional to the two-photon Rabi frequency $\partial f = \Omega_{420}^2/4\Delta = \Omega/2$ and therefore the infidelity contribution is also a constant, assuming the single-photon Rabi frequencies are balanced between the 420-nm and 1015-nm lasers. In contrast, the polarizability from the 1015-nm laser is nearly a constant when intermediate state detuning is increased away from 7.8 GHz, for both the ground state and the Rydberg state. As a result, the infidelity due to the 1015-nm lightshift scales as $(\Omega_{1015}^2/\Omega)^2 = \Delta^2$, growing quadratically with increasing intermediate state de-




Figure F.7: Numerical simulation of fidelity in cluster state circuit. Comparison between fitted fidelity of the decay in global return probability and the true CZ fidelity put into the numerical simulation, for the cluster state circuit in Fig. 8.4. Black lines denote $y = x$.

tuning. As a result, by improving the $1015\text{-nm } T_2^*$ by a factor of 2 from our current value, we can increase Δ further (Fig. F.6b).

With these two technical improvements of suppressed coupling to $|r'\rangle$ and improved $1015\text{-nm } T_2^*$, the CZ gate fidelity can then be improved further through scaling the laser intensities (Fig. F.6d). In particular, for a given intensity available for the two Rydberg lasers, there exists an optimal configuration of Rabi frequency and intermediate-state detuning that minimizes the error by making the most efficient use of the available laser power. As shown in Fig. F.6d, this gives a path toward reaching raw CZ gate fidelities approaching the $\sim 99.95\%$ level through an increase in available laser power.

F.4 HIGH-FIDELITY QUANTUM CIRCUITS

To benchmark our ability to perform quantum circuits based on atom movement, we perform the sequence in which one-dimensional cluster states are created and then unmade repeatedly over time. We then calculate the global return probability, which is defined as the probability that all 20 atoms are measured as $|0\rangle$ after the final $\pi/2$ single-qubit gate used to measure in the X basis. We fit an exponential decay of this global return as a function of the total number of CZ gates applied across the array. For all numbers of CZ gates, we include the same full single-qubit sequence, in-

cluding all atom moves, in order to isolate the effect of the CZ gates themselves (similar to the RB sequences). In addition to the CZ gates performed between pairs of atoms during the circuit, the unpaired atoms at the edges of the array are also excited to the Rydberg state and can experience gate errors. To be more conservative in our analysis, we do not count these unpaired atoms in the sum for the total number of CZ gates, which leads to a small decrease in the extracted CZ fidelity. Finally, we note that while such a sequence is not as rigorous of a CZ gate benchmarking protocol as the RB sequences, we find in numerical simulations that the extracted fidelity is similar to the CZ fidelity (Fig. F.7). Therefore, our measurements suggest that maintaining low-error regimes during quantum circuits is achievable.

Fig. F.8a shows images of atoms in the locations of the CZ gates used in the nonlocal circuits; for the cluster state circuit, only nearest-neighbor moves are used. The circuit structure is designed such that one half of the atoms are in static SLM traps and the other half are in moving AOD traps, and there are no transfers between different traps during the circuit.

For the nonlocal circuit implementation in Fig. 8.5, we find that the measured cross-entropy benchmarking (XEB) results match well to an error budget consisting of a CZ gate Pauli error of 0.06% with 0.09% loss error, initialization Pauli error of 0.36% with 0.37% loss, measurement Pauli error of 0.5% with 0.69% loss, and total error from idling, single-qubit gates, and atom moves during each move layer of 0.091% Pauli error and 0.033% loss error. We further empirically measure the loss during the circuit and upper bound the loss per qubit per move at $< 0.03\%$. For the loss per atom per CZ gate measurement shown in Fig. 8.5c, we only consider atoms which perform CZ gates in all of the circuit layers, excluding the atoms on the edge of the array. Finally, we find that for both the cluster state and nonlocal circuits, there is no evidence for large-scale correlated loss events, as shown in Fig. F.8c.




Figure F.8: Nonlocal circuit setup and supplementary data. (A) Atom images for the CZ gate locations for the nonlocal circuit (cluster state circuit repeats the first two nearest-neighbor moves). (B) Sampling from 10-qubit circuit, showing first 30 unordered outcomes. (C) Distribution of atoms lost per shot (out of 20 atoms total), for the 670,890 shots for the cluster state circuit and the 26,140 shots for the nonlocal circuit. The solid curves are a sum of Poissonian fits of the loss distributions for each different number of CZ gates.

F.5 THEORETICAL ANALYSIS OF NONLOCAL CIRCUITS

Here we discuss the theoretical properties of our nonlocal circuits. We first numerically study the circuits by evaluating the entanglement entropy, which we define throughout this work as the half-system-size von Neumann entropy (first Rényi entropy), $S_1(A) = -\text{Tr}(\rho_A \log \rho_A)$, where we consider the half-system A with a cut in the middle of the chain. Specifically, we choose nonlocal quantum circuits that generate states with entanglement entropy close to the Page entropy³²¹, which has a value of ~ 9.28 for 20 qubits. We find that the circuit structure gives rise to super-ballistic scrambling, in which the time to saturate the Page entropy scales as $\sqrt{N}$ for system size N . The intuition underlying this scaling is that the range of CZ gate connectivity expands linearly with circuit depth, as compared to the case of only nearest-neighbor CZ gates where the connectivity is constant with depth. Further, performing the same circuit with only nearest-neighbor gates causes the entanglement entropy not to saturate at the Page entropy (Fig. F.9a). In this case, the failure to reach the Page entropy is attributed to the integrability of the circuit dynamics, which constrains the system's evolution and leads to revivals in the block entanglement entropy³⁰¹. Note that a modification of the nearest-neighbor circuit, such as additional single-qubit z -axis rotations in the bulk $e^{-i\varphi \sum_j Z_j}$, could break integrability and lead to scrambling, but in a time that scales linearly with system size N . An additional component of our circuits is the choice of a 45-degree single-qubit rotation $R = X(\pi/4)$ between CZ layers. We explore this rotation angle in theory, finding that larger or smaller angles do not saturate to the Page theory value (Fig. F.9b).

When sampling the output of these circuits in the Pauli Z basis, we find that the resulting ideal distribution of outcome probabilities approaches the Porter-Thomas distribution $\text{PT}(x)$ as a function of depth (Fig. F.9c). We can compare between our circuit output distribution $q(x)$ and $\text{PT}(x)$ more quantitatively by calculating the KL divergence, $D_{\text{KL}} = \sum_x \text{PT}(x) \log \frac{\text{PT}(x)}{q(x)}$. In Fig. F.9d, we find that D_{KL} approaches zero for our circuits with sufficient depth. Finally, we characterize the




Figure F.9: Theoretical analysis of nonlocal circuits. (A) Entanglement entropy as a function of circuit depth, comparing our nonlocal circuit to applying only nearest-neighbor gates at all depths. (B) Entanglement entropy as a function of depth, for different rotation angles θ for the single-qubit gates in the circuit, where we use $\theta = 0.25$ corresponding to a 45-degree X rotation. (C) Distribution of ordered outcome probabilities as a function of nonlocal connectivity, compared to the Porter-Thomas distribution (PT). (D) KL divergence from $PT(x)$ as a function of nonlocal connectivity. (E) Distribution of eigenvalues of the reduced density matrices for subsystems on the boundary, compared to the Marchenko-Pastur distribution (black curve) expected for Haar-random pure states, for nonlocal connectivity of 11.

scrambling in these nonlocal circuits more rigorously by examining the distribution of eigenvalues of the reduced density matrices for subsystems on the boundary. In particular, for sufficiently deep circuits (nonlocal connectivity of 11), we find that this eigenvalue distribution is well-described by a Marchenko–Pastur distribution⁹⁷, consistent with the entanglement spectrum expected for Haar-random pure states (Fig. F.9e).



Example experiment command strings

Here we present example experiment command strings, for the full surface code repeated QEC experiments in Ref.⁴².

G.1 MOVING AWG

The command string used to program the RF waveforms for the Moving AWG, which control the positions and powers of moveable AOD tweezers, is given below. Waveforms are uploaded to dif-

ferent segments of the AWG using the SEGMENT[] command, which can then be individually looped or triggered with an external trigger.

Each move is defined using a MOVEAOD command. The initial position and power is specified for the first command when tweezers are turned on. For subsequent commands, the arguments specify where to move to (and optionally the power to ramp to), and take the form:

`MOVEAOD[ [x_positions,y_positions] , [xpowers,ypowers] , time_for_move]`

Additional arguments such as SHIFT apply a global (x,y) coordinate shift to the coordinates in that move. Coordinates are given in 'SLM units' where 1 SLM unit here is approximately 4.6um in the x direction and 4.9um in the y direction (corresponding to 7 and 7.5 units in the corresponding Supplementary Video, respectively)

```
### Define coordinates and powers
COMB[xd_ent, [13,3,6]]
COMB[yd_ent, [2,2,6]]
COMB[ykill_ent, [0, 2, 7]]
REPEAT_VAL[xkill_pow, [0.4, 6]]
REPEAT_VAL[ykill_pow, [1.65, 7]]
COMB[xanc_ent, [14.86,3,6]]
COMB[yanc_ent, [1,2,6]]
REPEAT_VAL[xno6_pow, [0.14, 6]]
REPEAT_VAL[yno6_pow, [0.0001, 6]]
REPEAT_VAL[yno7_pow, [0.0001, 7]]
REPEAT_VAL[xpow_imaging, [0.25, 6]]
REPEAT_VAL[ypow_imaging, [1.01,6]]
REPEAT_VAL[xpow_gates, [0.18, 6]]
REPEAT_VAL[ypow_gates, [0.91, 6]]
COMB[x_st, [-0.98, 3, 6]]
COMB[y_st, [21, 1.98, 6]]
REPEAT_VAL[xpow_pickup, [0.36, 6]]
REPEAT_VAL[ypow_pickup, [1.83, 6]]
REPEAT_VAL[x2pow_pickup, [0.36, 6]]
```

```

REPEAT_VAL[y2pow_pickup, [1.33, 6]]
COMB[x_ssr, [-0.98, 2, 18]]
VAR[xpow1_ssr, [0.26, 0.26, 0.24, 0.24, 0.22, 0.21, 0.21, 0.21, 0.21, 0.21, 0.21, 0.23,
0.23, 0.23, 0.24, 0.24, 0.24, 0.20]]
REPEAT_VAL[xno18_pow, [0.0001, 18]]
VAR[y3pow_pickup, [1.19, 1.20, 1.23, 1.28, 1.29, 1.47]]

### First segment, used to clear unwanted atoms from entangling zone traps
SEGMENT[1, LOOP_UNTIL_TRIG]
MoveAOD[Initial[xd_ent, ykill_ent], InitialPows[xkill_pow, yno7_pow],
FinalPows[xkill_pow, ykill_pow], 0.1 ms]
MoveAOD[[xd_ent, ykill_ent], SHIFT[[-1.5], [-1]], 0.1 ms]
WAIT[0.02 ms]

### Second segment, move atoms to their positions for the start of the circuit
SEGMENT[2]
MoveAOD[Initial[x_st, y_st], InitialPows[xno6_pow, xno6_pow],
FinalPows[xpow_pickup, ypow_pickup], 0.2 ms]

### Move data qubits to the entangling zone
MoveAOD[[x_st, y_st], SHIFT[[0.5], [0]], 0.1 ms]
MoveAOD[[x_st, yd_ent], SHIFT[[0.5], [-0.5]], 0.4 ms]
MoveAOD[[xd_ent, yd_ent], SHIFT[[0], [-0.5]], 0.2 ms]
MoveAOD[[xd_ent, yd_ent], SHIFT[[0], [0]], 0.1 ms]
MoveAOD[[xd_ent, yd_ent], FinalPows[xno6_pow, xno6_pow], 0.3 ms]

### Move first ancilla set to entangling zone
MoveAOD[Initial[x_st, y_st], SHIFT[[1], [0]], InitialPows[xno6_pow, xno6_pow],
FinalPows[xpow_pickup, ypow_pickup], 0.2 ms]
MoveAOD[[x_st, y_st], SHIFT[[1.5], [0]], 0.1 ms]
MoveAOD[[x_st, yanc_ent], SHIFT[[1.5], [0]], 0.3 ms]
MoveAOD[[xanc_ent, yanc_ent], 0.3 ms]
MoveAOD[[xanc_ent, yanc_ent], FinalPows[xpow_imaging, ypow_imaging], 0.05ms]

### Start of circuit
SEGMENT[4]

```

```

Wait[1 us]
MoveAOD[[xanc_ent,yanc_ent], FinalPows[xpow_gates, ypow_gates], 0.1 ms]
### Data pi/2 pulses for first round
Wait[403 us]
### Fake ancilla pi/2 pulses, first round
Wait[70 us]

### Data qubit pi pulse, followed by Ancilla pi/2 pulses
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[-1.5],[1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[-1.5],[-1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0.75],[1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0.78],[-1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]

### Move back to original spot
MoveAOD[[xanc_ent,yanc_ent], FinalPows[x2pow_pickup, y2pow_pickup], 0.1 ms]
MoveAOD[[x_st,yanc_ent],SHIFT[[1.53],[0]],0.4 ms]
MoveAOD[[x_st,y_st],SHIFT[[1.53],[0]],0.4 ms]
MoveAOD[[x_st,y_st], SHIFT[[1.03],[0]],0.1 ms]
MoveAOD[[x_st,y_st], FinalPows[xno6_pow, yno6_pow], SHIFT[[1.03],[0]], 0.250 ms]

### Pick up next group of atoms, and do again.
### Exchanging the ancilla arrays is currently given 1 ms total time.
MoveAOD[Initial[x_st,y_st], InitialPows[xno6_pow, yno6_pow],

```

```

FinalPows[x2pow_pickup, y2pow_pickup],SHIFT[[1.938],[0]], 0.250 ms]
MoveAOD[[x_st,y_st],SHIFT[[2.438],[0]],0.1 ms]
MoveAOD[[x_st,yanc_ent],SHIFT[[2.438],[0]],0.4 ms]
MoveAOD[[xanc_ent,yanc_ent], 0.4 ms]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], FinalPows[xpow_gates, ypow_gates], 0.1 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[-1.5],[1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0.78],[1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[-1.5],[-1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0.78],[-1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], FinalPows[x2pow_pickup, y2pow_pickup], 0.1 ms]
MoveAOD[[x_st,yanc_ent],SHIFT[[2.438],[0]],0.4 ms]
MoveAOD[[x_st,y_st],SHIFT[[2.438],[0]],0.4 ms]
MoveAOD[[x_st,y_st], SHIFT[[1.938],[0]],0.1 ms]
MoveAOD[[x_st,y_st], FinalPows[xno6_pow, yno6_pow], SHIFT[[1.938], [0]], 0.250ms]

```

Third ancilla block

```

MoveAOD[Initial[x_st,y_st], InitialPows[xno6_pow, yno6_pow],
    FinalPows[x2pow_pickup, y2pow_pickup],SHIFT[[17.959],[0]], 0.250 ms]
MoveAOD[[x_st,y_st],SHIFT[[18.459],[0]],0.1 ms]
MoveAOD[[x_st,yanc_ent],SHIFT[[18.459],[0]],0.4 ms]
MoveAOD[[xanc_ent,yanc_ent], 0.4 ms]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], FinalPows[xpow_gates, ypow_gates], 0.1 ms]
Wait[70 us]

```

```

MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0.78],[-1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[-1.5],[-1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0.78],[1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[-1.5],[1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], FinalPows[x2pow_pickup, y2pow_pickup], 0.1 ms]
MoveAOD[[x_st,yanc_ent],SHIFT[[18.459],[0]],0.4 ms]
MoveAOD[[x_st,y_st],SHIFT[[18.459],[0]],0.4 ms]
MoveAOD[[x_st,y_st], SHIFT[[17.959],[0]],0.1 ms]
MoveAOD[[x_st,y_st], FinalPows[xno6_pow, yno6_pow], SHIFT[[17.959],[0]], 0.250 ms]

```

Fourth ancilla block

```

MoveAOD[Initial[x_st,y_st], InitialPows[xno6_pow, yno6_pow],
    FinalPows[x2pow_pickup, y2pow_pickup],SHIFT[[18.972],[0]], 0.250 ms]
MoveAOD[[x_st,y_st],SHIFT[[19.472],[0]],0.1 ms]
MoveAOD[[x_st,yanc_ent],SHIFT[[19.472],[0]],0.4 ms]
MoveAOD[[xanc_ent,yanc_ent], 0.4 ms]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], FinalPows[xpow_gates, ypow_gates],
0.1 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0.78],[-1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0.78],[1]], 0.195 ms]

```

```

Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[-1.5],[-1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[-1.5],[1]], 0.195 ms]
Wait[10 us]
MoveAOD[[xanc_ent,yanc_ent], SHIFT[[0],[0]], 0.195 ms]
Wait[70 us]
MoveAOD[[xanc_ent,yanc_ent], FinalPows[x2pow_pickup, y2pow_pickup], 0.1 ms]
MoveAOD[[x_st,yanc_ent],SHIFT[[19.472],[0]],0.4 ms]
MoveAOD[[x_st,y_st],SHIFT[[19.472],[0]],0.4 ms]
MoveAOD[[x_st,y_st], SHIFT[[18.972],[0]],0.1 ms]
MoveAOD[[x_st,y_st], FinalPows[xno6_pow, yno6_pow], SHIFT[[18.972],[0]], 0.250ms]

```

Move data back to storage zone, 1.2 ms

```

MoveAOD[Initial[xd_ent, yd_ent], InitialPows[xno6_pow, yno6_pow],
FinalPows[x2pow_pickup,y2pow_pickup], 0.35 ms]
MoveAOD[[xd_ent,yd_ent],SHIFT[[0],[-0.5]],0.2 ms]
MoveAOD[[x_st,yd_ent], SHIFT[[0.53],[-0.5]],0.3 ms]
MoveAOD[[x_st,y_st], SHIFT[[0.53],[0]],0.5 ms]
MoveAOD[[x_st,y_st], SHIFT[[0.03],[0]],0.2 ms]
MoveAOD[[x_st,y_st], FinalPows[xno6_pow, yno6_pow], SHIFT[[0.03],[0]], 0.3 ms]

```

SEGMENT[6]

SSR on half, 1.2 ms

```

WAIT[0.2 ms]
MoveAOD[Initial[x_ssr, y_st],
InitialPows[xno18_pow, yno6_pow],
FinalPows[xpow1_ssr,y3pow_pickup],0.5 ms]
WAIT[100 us]
MoveAOD[[x_ssr, y_st], SHIFT[[0.5],[0]], 0.7 ms]
MoveAOD[[x_ssr, y_st], SHIFT[[0.5],[-1]],0.2 ms]
MoveAOD[[x_ssr, y_st], SHIFT[[0],[-1]],0.1 ms]

```



```
MoveAOD[[x_ssr,y_st], FinalPows[xno18_pow, yno6_pow], SHIFT[[0],[-1]], 0.3 ms]
```

```
### SSR on half, 1.23 ms
```

```
MoveAOD[Initial[x_ssr, y_st],
InitialPows[xno18_pow, yno6_pow],
FinalPows[xpow1_ssr,y3pow_pickup], SHIFT[[1],[0]],0.5 ms]
WAIT[100 us]
MoveAOD[[x_ssr, y_st], SHIFT[[1.5],[0]], 0.7 ms]
MoveAOD[[x_ssr, y_st], SHIFT[[1.5],[-1]],0.2 ms]
MoveAOD[[x_ssr, y_st], SHIFT[[1],[-1]],0.1 ms]
MoveAOD[[x_ssr,y_st], FinalPows[xno18_pow, yno6_pow], SHIFT[[1],[-1]], 0.3 ms]
END_COMMAND[]
```

G.2 RAMAN AWG

The command string used to program the I/Q waveforms and AOM amplitudes for local and global Raman control is given below. The command specifies the parameters for each single-qubit gate, which is then converted into the appropriate waveforms to modulate the light intensity and microwave phase and frequency (defined relative to the resonance frequency). For local gates, the pulses are optionally gated on and off by the Raman AOD command.

```
### Examples of typical commands are
# Z : global Z rotation, in units of 2*pi
# E : wait time, idling at the specified frequency
# SC1_ECHO : global SCROFULOUS Y(pi) pulse for dynamical decoupling,
# idling at the specified frequency
# SC1_H_FIXEDLOOPED_SMOOTH_PCOMP : SCROFULOUS Y(pi/2) local gaussian pulses
# SC1PI_FIXEDLOOPED_SMOOTH_PCOMP : SCROFULOUS Y(pi) local gaussian pulses
# BB1_FIXED: global BB1 pulse

### State preparation before circuit begins
E[50.5 us, 12 MHz]
```

```

CPROBUST180_OFFSETFREQSMOOTH_COMPENSATED[1.73 us, 0.035 us, 12.1056 MHz, 0.75,
4]
E[0.2 us, 8.3 MHz]
CPROBUST180_OFFSETFREQSMOOTH[2.455 us, 0.035 us, 4.063 MHz, 1, 4]
E[2.65 us, 0.1 MHz]
SEGMENTBREAK[1896]
E[16.6 us, 0.1 MHz]
BB1_FIXED[0.5,0.12 us,5.65 us]
E[4 us, 5 MHz]
SMOOTHPUMPINGLOOPEDTTL[0, 2.18 us, -11.7 MHz, 11.7 MHz, 1, 4, 24, 8.53 us,
0.1us, 3us]
E[494.052 us, 5 MHz]

```

Start of circuit. Prepare data qubits

```

SC1_ECHO[6 us, 0 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
Z[0.45710]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]

```

Start first QEC round

```

Z[0]
E[397.968 us, 0.07 MHz]
Z[0.000]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
SC1_ECHO[76 us, 0.07 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
Z[0.287]
E[397.968 us, 0.07 MHz]
Z[0.000]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.46744]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.46744]
Z[0.000]

```

```

Z[0.287]
E[397.968 us, 0.07 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
SC1_ECHO[4 us, 0.07 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
E[397.968 us, 0.07 MHz]
Z[0.287]
Z[0.000]
Z[0.46744]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
Z[0.46744]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.287]
E[397.968 us, 0.07 MHz]
Z[0.000]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
Z[0.000]
SC1_ECHO[2574.028 us, 0.07 MHz]
Z[0.000]

# Start second QEC round
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
Z[0.287]
E[397.968 us, 0.07 MHz]
Z[0.000]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.46744]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.46744]
Z[0.000]
Z[0.287]
E[397.968 us, 0.07 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
SC1_ECHO[4 us, 0.07 MHz]

```

```

SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
E[397.968 us, 0.07 MHz]
Z[0.287]
Z[0.000]
Z[0.46744]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
Z[0.46744]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.287]
E[397.968 us, 0.07 MHz]
Z[0.000]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
Z[0.000]
SC1_ECHO[2574.028 us, 0.07 MHz]
Z[0.000]

```

Start third QEC round

```

SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
Z[0.287]
E[397.968 us, 0.07 MHz]
Z[0.000]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.46744]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.46744]
Z[0.000]
Z[0.287]
E[397.968 us, 0.07 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
SC1_ECHO[4 us, 0.07 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
E[397.968 us, 0.07 MHz]
Z[0.287]
Z[0.000]

```

```

Z[0.46744]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
Z[0.46744]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.287]
E[397.968 us, 0.07 MHz]
Z[0.000]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
Z[0.000]
SC1_ECHO[2574.028 us, 0.07 MHz]
Z[0.000]

# Start fourth QEC round
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
Z[0.287]
E[397.968 us, 0.07 MHz]
Z[0.000]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.46744]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.46744]
Z[0.000]
Z[0.287]
E[397.968 us, 0.07 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
SC1_ECHO[4 us, 0.07 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
E[397.968 us, 0.07 MHz]
Z[0.287]
Z[0.000]
Z[0.46744]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
Z[0.46744]

```

```

SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
Z[0.287]
E[397.968 us, 0.07 MHz]
Z[0.000]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
SC1_ECHO[76 us, 0.07 MHz]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 6, local]
Z[0]
E[397.968 us, 0.07 MHz]
Z[0.000]

# Readout data qubits
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]
SC1PI_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.45710, 6, local]
Z[0.46744]
SC1_H_FIXEDLOOPED_SMOOTH_PCOMP[0.01us, 0.88, 4, 5.8180 us, 0.46744, 3, local]

# Bright state transfer for spin-to-position conversion
E[2.872 ms, 5 MHz]
BB1_FIXED[0.5,0.1us,6us]
E[1.819ms, 5MHz]
CPROBUST180_OFFSETFREQSMOOTH[4.55 us, 0.035 us, 4.083 MHz, 1, 4]
E[0.2 us, 8.3 MHz]
CPROBUST180_OFFSETFREQSMOOTH_COMPENSATED[3.1 us, 0.035 us, 12.3256 MHz, 0.75, 4]
END_COMMAND[]

```

G.3 LOCAL RAMAN AODs AWG

The command string used to program the RF waveforms for the local Raman AODs AWG, which controls the positions and powers of 2D grids of optical tweezers used for local single-qubit gates, can be found in the Supplementary Information of Ref.⁴²

G.4 RYDBERG AWG

The command string, used to program the amplitude, frequency, and phase for the 420-nm AOM for Rydberg entangling gates and local SLM AOM for local detunings, is given below. Gate profiles (amplitude, detuning, and phase) are loaded from saved calibration files.

```
### Start first QEC round
Wait[500.91 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Wait[2570.02 us]
```

Start second QEC round

```
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Wait[2570.02 us]
```

Start third QEC round

```
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
```



```

420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Wait[2570.02 us]

```

Start fourth QEC round

```

Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]

```

```
Wait[467.800 us]
Pulse[420Amp[amplitudeFromFile,idx=1001,loops=1,delay=0.2],
420Freq[detuningFromFile,idx=1001],
420Phase[phaseFromFile,idx=1001, loops=1,delay=0.2],
localShiftAmp[RabiPulseLoop,time=0.40,loops=1,A=1,delay = 0.18],
localShiftFreq[fixedFrequency,detuning = 0.0]]
Wait[467.800 us]
```

References

- [1] Aaronson, S. et al. (2004). Improved simulation of stabilizer circuits. *Physical Review A*, 70, 052328.
- [2] Abanin, D. et al. (2017a). A Rigorous Theory of Many-Body Prethermalization for Periodically Driven and Closed Quantum Systems. *Communications in Mathematical Physics*, 354(3), 809–827.
- [3] Abanin, D. A. et al. (2017b). Effective Hamiltonians, prethermalization, and slow energy absorption in periodically driven many-body systems. *Phys. Rev. B*, 95(1), 014112.
- [4] Acharya, R. et al. (2025). Quantum error correction below the surface code threshold. *Nature*, 638, 920–926.
- [5] Acharya, R. et al. (2023). Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7949), 676–681.
- [6] Aharonov, D. et al. (1999). Fault-tolerant quantum computation with constant error rate.
- [7] Aidelsburger, M. et al. (2015). Measuring the Chern number of Hofstadter bands with ultracold bosonic atoms. *Nat. Phys.*, 11(2), 162–166.
- [8] Alcaraz, F. C. et al. (1999). An Exactly Solvable Constrained XXZ Chain. [arXiv:9904042](#).
- [9] Alet, F. et al. (2006). Classical dimers with aligning interactions on the square lattice. *Physical Review E*, 74(4), 041124.
- [10] Alicea, J. (2012). New directions in the pursuit of Majorana fermions in solid state systems. *Rep. Prog. Phys.*, 75(7), 076501.
- [11] Allen, R. R. et al. (2025). Quantum computing enhanced sensing. *arXiv preprint arXiv:2501.07625*.
- [12] Altman, E. et al. (2021). Quantum Simulators: Architectures and Opportunities. *PRX Quantum*, 2(1), 017003.

- [13] Andersen, T. I. et al. (2025). Thermalization and criticality on an analogue–digital quantum simulator. *Nature*, 638(8049), 79–85.
- [14] Andersen, T. I. et al. (2023). Non-Abelian braiding of graph vertices in a superconducting processor. *Nature*, 618(7964), 264–269.
- [15] Anderson, P. W. (1972). More Is Different. *Science*, 177(4047), 393–396.
- [16] Ardonne, E. et al. (2004). Topological order and conformal quantum critical points. *Annals of Physics*, 310(2), 493–551.
- [17] Arute, F. et al. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 505–510.
- [18] Arute, F. et al. (2020). Hartree-Fock on a superconducting qubit quantum computer. *Science*, 369(6507), 1084–1089.
- [19] Baireuther, P. et al. (2018). Machine-learning-assisted correction of correlated qubit errors in a topological code. *Quantum*, 2, 48.
- [20] Bakr, W. S. et al. (2010). Probing the superfluid-to-Mott insulator transition at the single-atom level. *Science*, 329(5991), 547–550.
- [21] Baldwin, C. H. et al. (2020). Subspace benchmarking high-fidelity entangling operations with trapped ions. *Phys. Rev. Res.*, 2, 013317.
- [22] Ballance, C. J. et al. (2016). High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.*, 117, 060504.
- [23] Baranes, G. et al. (2026). Leveraging qubit loss detection in fault-tolerant quantum algorithms. *Physical Review X*, 16(1), 011002.
- [24] Barenblatt, G. I. (1996). *Scaling, Self-Similarity, and Intermediate Asymptotics*. Cambridge: Cambridge University Press.
- [25] Barends, R. et al. (2014). Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497), 500–503.
- [26] Barends, R. et al. (2015). Digital quantum simulation of fermionic models with a superconducting circuit. *Nat. Commun.*, 6(1), 7654.
- [27] Barnes, K. et al. (2022). Assembly and coherent control of a register of nuclear spin qubits. *Nature Communications*, 13(1), 2779.
- [28] Barredo, D. et al. (2016). An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science*, 354(6315), 1021–1023.

- [29] Barrett, S. D. et al. (2010). Fault tolerant quantum computation with very high threshold for loss errors. *Physical Review Letters*, 105(20), 200502.
- [30] Bausch, J. et al. (2024). Learning high-accuracy error decoding for quantum processors. *Nature*, 635(8040), 834–840.
- [31] Bell, J. S. (1964). On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3), 195.
- [32] Benhelm, J. et al. (2008). Towards fault-tolerant quantum computing with trapped ions. *Nat. Phys.*, 4(6), 463–466.
- [33] Benioff, P. (1980). 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.
- [34] Bernien, H. et al. (2017). Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682), 579–584. Number: 7682.
- [35] Berry, M. V. (1984). Quantal phase factors accompanying adiabatic changes. *Proc. R. Soc. Lond.*, 392(1802), 45–57.
- [36] Beterov, I. I. et al. (2016). Two-qubit gates using adiabatic passage of the Stark-tuned Förster resonances in Rydberg atoms. *Phys. Rev. A*, 94, 062307.
- [37] Beugnon, J. et al. (2007). Two-dimensional transport and transfer of a single atomic qubit in optical tweezers. *Nat. Phys.*, 3(10), 696–699.
- [38] Beverland, M. E. et al. (2021). Cost of Universality: A Comparative Study of the Overhead of State Distillation and Code Switching with Color Codes. *PRX Quantum*, 2(2), 020341.
- [39] Beverland, M. E. et al. (2022). Assessing requirements to scale to practical quantum advantage. *arXiv:2211.07629*.
- [40] Biroli, G. et al. (2010). Kibble-Zurek mechanism and infinitely slow annealing through critical points. *Physical Review E*, 81(5), 050101.
- [41] Bluvstein, D. et al. (2024). Logical quantum processor based on reconfigurable atom arrays. *Nature*, 626, 58–65.
- [42] Bluvstein, D. et al. (2026). A fault-tolerant neutral-atom architecture for universal quantum computation. *Nature*, 649(8095), 39–46.
- [43] Bluvstein, D. et al. (2022). A quantum processor based on coherent transport of entangled atom arrays. *Nature*, 604(7906), 451–456.

- [44] Bluvstein, D. et al. (2021). Controlling quantum many-body dynamics in driven Rydberg atom arrays. *Science*, 371(6536), 1355–1359.
- [45] Boixo, S. et al. (2018). Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6), 595–600.
- [46] Bollor, K.-J. et al. (1991). Observation of electromagnetically induced transparency. *Phys. Rev. Lett.*, 66, 2593–2596.
- [47] Bombín, H. (2015a). Gauge color codes: optimal transversal gates and gauge fixing in topological stabilizer codes. *New Journal of Physics*, 17(8), 083002.
- [48] Bombín, H. (2015b). Single-Shot Fault-Tolerant Quantum Error Correction. *Physical Review X*, 5(3), 031043.
- [49] Bombín, H. et al. (2024). Fault-Tolerant Postselection for Low-Overhead Magic State Preparation. *PRX Quantum*, 5(1), 010302.
- [50] Bonilla Ataides, J. P. et al. (2025). Neural decoders for universal quantum algorithms. *arXiv:2509.11370*.
- [51] Bonilla Ataides, J. P. et al. (2021). The xzzx surface code. *Nature Communications*, 12(1), 2172.
- [52] Bouland, A. et al. (2019). On the complexity and verification of quantum random circuit sampling. *Nature Physics*, 15(2), 159–163.
- [53] Boulier, T. et al. (2017). Spontaneous avalanche dephasing in large Rydberg ensembles. *Physical Review A*, 96(5), 053409.
- [54] Bravyi, S. et al. (2019). Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum*, 3, 181.
- [55] Bravyi, S. et al. (2024). High-threshold and low-overhead fault-tolerant quantum memory. *Nature*, 627(8005), 778–782.
- [56] Bravyi, S. et al. (2018). Correcting coherent errors with surface codes. *npj Quantum Information*, 4(1), 55.
- [57] Bray, A. J. (1994). Theory of phase-ordering kinetics. *Adv. Phys.*, 43(3), 357–459.
- [58] Bremner, M. J. et al. (2011). Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126), 459–472.
- [59] Bremner, M. J. et al. (2016). Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations. *Physical Review Letters*, 117(8), 080501.

- [60] Bremner, M. J. et al. (2017). Achieving quantum supremacy with sparse and noisy commuting quantum computations. *Quantum*, 1, 8.
- [61] Brézin, E. et al. (1974). Universal ratios of critical amplitudes near four dimensions. *Phys. Lett. A*, 47(4), 285–287.
- [62] Brion, E. et al. (2007). Implementing a neutral atom Rydberg gate without populating the Rydberg state. *J. Phys. B: At. Mol. Opt. Phys.*, 40(9), S159.
- [63] Browaeys, A. et al. (2020). Many-body physics with individually controlled rydberg atoms. *Nature Physics*, 16, 132–142.
- [64] Brown, B. J. (2020). A fault-tolerant non-clifford gate for the surface code in two dimensions. *Science Advances*, 6(21).
- [65] Brown, M. O. et al. (2023). Time-of-flight quantum tomography of single atom motion. *Nature Physics*, 19, 569–573.
- [66] Brown, M. O. et al. (2019). Gray-molasses optical-tweezer loading: Controlling collisions for scaling atom-array assembly. *Phys. Rev. X*, 9, 011057.
- [67] Brydges, T. et al. (2019). Probing Rényi entanglement entropy via randomized measurements. *Science*, 364(6437), 260–263.
- [68] Bytyqi, E. et al. (2026). Device for mhz-rate rastering of arbitrary 2d optical potentials. arXiv preprint arXiv:2602.16025.
- [69] Cain, M. et al. (2025). Fast correlated decoding of transversal logical algorithms. arXiv preprint arXiv:2505.13587.
- [70] Cain, M. et al. (2024). Correlated decoding of logical algorithms with transversal gates. *Physical Review Letters*, 133(24), 240602.
- [71] Campbell, E. T. (2016). The smallest interesting colour code.
- [72] Campostrini, M. et al. (2002). 25th-order high-temperature expansion results for three-dimensional ising-like systems on the simple-cubic lattice. *Phys. Rev. E*, 65(6), 066127.
- [73] Cao, A. et al. (2024). Multi-qubit gates and schrödinger cat states in an optical clock. *Nature*, 634, 315–320.
- [74] Cao, A. et al. (2025). Autoionization-enhanced Rydberg dressing by fast contaminant removal. *Physical Review Letters*, 134(13), 133201.
- [75] Caselle, M. et al. (1997). Universal amplitude ratios in the three-dimensional ising model. *J. Phys. A: Math. Gen.*, 30(14), 4963–4982.

- [76] Cerezo, M. et al. (2021). Variational quantum algorithms. *Nature Reviews Physics*, 3(9), 625–644.
- [77] Cesa, F. et al. (2026). Engineering discrete local dynamics in globally driven dual-species atom arrays. arXiv:2601.16961.
- [78] Chalopin, T. et al. (2026). Observation of emergent scaling of spin–charge correlations at the onset of the pseudogap. *Proceedings of the National Academy of Sciences*, 123(4), e2525539123.
- [79] Chamberland, C. et al. (2020). Triangular color codes on trivalent graphs with flag qubits. *New Journal of Physics*, 22(2), 023019.
- [80] Chandran, A. et al. (2012). Kibble-Zurek problem: Universality and the scaling limit. *Phys. Rev. B*, 86(6), 064304.
- [81] Chandran, A. et al. (2013). Equilibration and coarsening in the quantum $O(n)$ model at infinite n . *Phys. Rev. B*, 88(2), 024306.
- [82] Chang, K. et al. (2025). Surface Code with Imperfect Erasure Checks. *PRX Quantum*, 6, 040355.
- [83] Chang, T. H. et al. (2023). High-fidelity rydberg controlled-z gates with time-optimal pulses. *New Journal of Physics*, 25(12), 123007.
- [84] Chen, C. et al. (2023a). Continuous symmetry breaking in a two-dimensional Rydberg array. *Nature*, 616(7958), 691–695. Number: 7958.
- [85] Chen, Y.-A. et al. (2023b). Equivalence between Fermion-to-Qubit Mappings in two Spatial Dimensions. *PRX Quantum*, 4(1), 010326.
- [86] Chern, S.-S. (1946). Characteristic Classes of Hermitian Manifolds. *Ann. Math.*, 47(1), 85–121.
- [87] Chew, Y. et al. (2022). Ultrafast energy exchange between two single rydberg atoms on a nanosecond timescale. *Nature Photonics*, 16, 724–729.
- [88] Chew, Y. T. et al. (2024). Ultraprecise holographic optical tweezer array. *Physical Review A*, 110(5), 053318.
- [89] Childs, A. M. et al. (2025). Streaming quantum state purification. *Quantum*, 9, 1603.
- [90] Chiu, N.-C. et al. (2025). Continuous operation of a coherent 3,000-qubit system. *Nature*, 646, 1075–1080.
- [91] Choi, J. et al. (2023). Preparing random states and benchmarking with many-body quantum chaos. *Nature*, 613, 468–473.

- [92] Chow, C. H. et al. (2024). Fano resonance in excitation spectroscopy and cooling of an optically trapped single atom. *Physical Review Research*, 6(2), 023154.
- [93] Claes, J. et al. (2021). Character Randomized Benchmarking for Non-Multiplicity-Free Groups With Applications to Subspace, Leakage, and Matchgate Randomized Benchmarking. *PRX Quantum*, 2(1).
- [94] Clinton, L. et al. (2024). Towards near-term quantum simulation of materials. *Nat. Commun.*, 15(1), 211.
- [95] Cochran, T. A. et al. (2025). Visualizing dynamics of charges and strings in $(2 + 1)$ D lattice gauge theories. *Nature*, 642, 315–320.
- [96] Coleman, S. (1977). Fate of the false vacuum: Semiclassical theory. *Physical Review D*, 15(10), 2929.
- [97] Collins, B. et al. (2010). Random graph states, maximal flow and fuss–catalan distributions. *Journal of Physics A: Mathematical and Theoretical*, 43(27), 275303.
- [98] Cong, I. et al. (2022). Hardware-Efficient, Fault-Tolerant Quantum Computation with Rydberg Atoms. *Physical Review X*, 12(2), 021049.
- [99] Cooper, N. R. et al. (2019). Topological bands for ultracold atoms. *Rev. Mod. Phys.*, 91(1), 015005.
- [100] Covey, J. P. et al. (2023). Quantum networks with neutral atom processing nodes. arXiv preprint arXiv:2304.02088.
- [101] Cross, A. W. et al. (2019). Validating quantum computers using randomized model circuits. *Physical Review A*, 100, 032328.
- [102] Cummins, H. K. et al. (2003). Tackling systematic errors in quantum logic gates with composite rotations. *Physical Review A*, 67(4), 042308.
- [103] Daley, A. J. et al. (2012). Measuring Entanglement Growth in Quench Dynamics of Bosons in an Optical Lattice. *Physical Review Letters*, 109(2), 020505.
- [104] Dalibard, J. et al. (1989). Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *JOSA B*, 6(11), 2023–2045.
- [105] Darbha, S. et al. (2024a). False vacuum decay and nucleation dynamics in neutral atom systems. *arXiv preprint arXiv:2404.12360*.
- [106] Darbha, S. et al. (2024b). Long-lived oscillations of false and true vacuum states in neutral atom systems. *arXiv preprint arXiv:2404.12371*.

- [107] Davison, A. C. et al. (1997). *Bootstrap Methods and their Application*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge: Cambridge University Press.
- [108] Dawson, C. M. et al. (2005). The Solovay-Kitaev algorithm. *Quantum Information and Computation*, 6(1), 081–095.
- [109] de Léséleuc, S. et al. (2019). Observation of a symmetry-protected topological phase of interacting bosons with Rydberg atoms. *Science*, 365(6455), 775–780.
- [110] de Oliveira, A. et al. (2025). Demonstration of weighted graph optimization on a rydberg atom array using local light-shifts. *PRX Quantum*, 6, 010301.
- [111] Debroy, D. M. et al. (2021). Optimizing Stabilizer Parities for Improved Logical Qubit Memories. *Physical Review Letters*, 127(24), 240501.
- [112] Decross, M. et al. (2025). Computational power of random quantum circuits in arbitrary geometries. *Physical Review X*, 15, 021052.
- [113] Deist, E. et al. (2022). Mid-circuit cavity measurement in a neutral atom array. *Phys. Rev. Lett.*, 129, 203602.
- [114] Delfosse, N. et al. (2023). Spacetime codes of Clifford circuits. arXiv preprint arXiv:2304.05943.
- [115] Denecker, T. et al. (2025). Measurement and feedforward correction of the fast phase noise of lasers. *Phys. Rev. A*, 111, 042614.
- [116] Deng, X.-L. et al. (2005). Effective spin quantum phases in systems of trapped ions. *Physical Review A*, 72(6), 063407.
- [117] Dennis, E. et al. (2002). Topological quantum memory. *Journal of Mathematical Physics*, 43(9), 4452–4505.
- [118] Derby, C. et al. (2021). Compact fermion to qubit mappings. *Phys. Rev. B*, 104(3), 035118.
- [119] Dolgirev, P. E. et al. (2022). Periodic dynamics in superconductors induced by an impulsive optical quench. *Commun. Phys.*, 5(1), 234.
- [120] Dordević, T. et al. (2021). Entanglement transport and a nanophotonic interface for atoms in optical tweezers. *Science*, 373(6562), 1511–1514.
- [121] Duan, L. M. et al. (2003). Controlling spin exchange interactions of ultracold atoms in optical lattices. *Phys. Rev. Lett.*, 91.
- [122] Eastin, B. et al. (2009). Restrictions on transversal encoded quantum gate sets. *Physical Review Letters*, 102(11), 110502.

- [123] Ebadi, S. et al. (2021). Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature*, 595(7866), 227–232.
- [124] Eckner, W. J. et al. (2023). Realizing spin squeezing with rydberg interactions in an optical clock. *Nature*, 621(7980), 734–739.
- [125] Egan, L. et al. (2021). Fault-tolerant control of an error-corrected qubit. *Nature*, 598, 281–286.
- [126] Eisert, J. et al. (2015). Quantum many-body systems out of equilibrium. *Nature Physics*, 11(2), 124–130.
- [127] Else, D. V. et al. (2017). Prethermal Phases of Matter Protected by Time-Translation Symmetry. *Phys. Rev. X*, 7(1), 011026.
- [128] Endres, M. et al. (2012). The Higgs amplitude mode at the two-dimensional superfluid/Mott insulator transition. *Nature*, 487, 454–458.
- [129] Evered, S. J. et al. (2023). High-fidelity parallel entangling gates on a neutral-atom quantum computer. *Nature*, 622, 268–272.
- [130] Evered, S. J. et al. (2025). Probing the kitaev honeycomb model on a neutral-atom quantum computer. *Nature*, 645, 341–347.
- [131] Evered, S. J. et al. (2026a). Notes from the international school of physics “enrico fermi” course 214 on “quantum computers and simulators with atoms”. (Varenna, Italy, 5th–14th July 2024).
- [132] Evered, S. J. et al. (2026b). High-fidelity entangling gates and nonlocal circuits with neutral atoms. arXiv preprint arXiv:2604.25987.
- [133] Fazekas, P. et al. (1974). On the ground state properties of the anisotropic triangular antiferromagnet. *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics*, 30(2), 423–440.
- [134] Feldmeier, J. et al. (2024). Quantum simulation of dynamical gauge theories in periodically driven Rydberg atom arrays. arXiv:2408.02733.
- [135] Fendley, P. et al. (2004). Competing density-wave orders in a one-dimensional hard-boson model. *Physical Review B*, 69(7), 075106.
- [136] Festa, L. et al. (2022). Blackbody-radiation-induced facilitated excitation of Rydberg atoms in optical tweezers. *Physical Review A*, 105(1), 013109.
- [137] Feynman, R. P. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7).

- [138] Finkelstein, R. et al. (2024). Universal quantum operations and ancilla-based readout for tweezer clocks. *Nature*, 634, 321–327.
- [139] Fisher, M. E. et al. (1963). Statistical Mechanics of Dimers on a Plane Lattice. II. Dimer Correlations and Monomers. *Physical Review*, 132(4), 1411–1431.
- [140] Flammia, S. T. et al. (2011). Direct fidelity estimation from few Pauli measurements. *Physical Review Letters*, 106(23), 230501.
- [141] Fowler, A. G. et al. (2012). Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86, 032324.
- [142] Fradkin, E. (2013). *Field Theories of Condensed Matter Physics*. Cambridge: Cambridge University Press.
- [143] Fradkin, E. et al. (1990). Short range resonating valence bond theories and superconductivity. *Modern Physics Letters B*, 4(03), 225–232.
- [144] Fredenhagen, K. et al. (1983). Charged states in z_2 gauge theories. *Communications in Mathematical Physics*, 92(1), 81–119.
- [145] Fromenteil, C. et al. (2023). Protocols for rydberg entangling gates featuring robustness against quasi-static errors. *PRX Quantum*, 4, 020335.
- [146] Fu, Z. et al. (2022). High-fidelity entanglement of neutral atoms via a rydberg-mediated single-modulated-pulse controlled-phase gate. *Phys. Rev. A*, 105, 042430.
- [147] Fukui, T. et al. (2005). Chern Numbers in Discretized Brillouin Zone: Efficient Method of Computing (Spin) Hall Conductances. *J. Phys. Soc. Jpn.*, 74(6), 1674–1677.
- [148] Gaebler, J. P. et al. (2012). Randomized benchmarking of multiqubit gates. *Phys. Rev. Lett.*, 108, 260503.
- [149] Gagel, P. et al. (2015). Universal postquench coarsening and aging at a quantum critical point. *Phys. Rev. B*, 92(11), 115121.
- [150] Gao, X. et al. (2024). Limitations of Linear Cross-Entropy as a Measure for Quantum Advantage. *PRX Quantum*, 5, 010334.
- [151] Gazo, M. et al. (2025). Universal coarsening in a homogeneous two-dimensional Bose gas. *Science*, 389(6762), 802–805.
- [152] Geier, S. et al. (2021). Floquet Hamiltonian Engineering of an Isolated Many-Body Spin System. *Science*, 374(6571), 1149–1152.
- [153] Geim, A. A. et al. (2026). Engineering quantum criticality and dynamics on an analog-digital simulator. arXiv:2602.18555.

- [154] Ghosh, S. et al. (2025). Random regular graph states are complex at almost any depth. *PRX Quantum*, 6(4), 040328.
- [155] Gicev, S. et al. (2023). A scalable and fast artificial neural network syndrome decoder for surface codes. *Quantum*, 7, 1058.
- [156] Gidney, C. (2021). Stim: a fast stabilizer circuit simulator. *Quantum*, 5, 497.
- [157] Gidney, C. et al. (2021). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5, 433.
- [158] Gidney, C. et al. (2025). Yoked surface codes. *Nature Communications*, 16(1), 1–12.
- [159] Gidney, C. et al. (2024). Magic state cultivation: growing T states as cheap as CNOT gates. arXiv preprint arXiv:2409.17595.
- [160] Girvin, S. M. et al. (2019). *Modern Condensed Matter Physics*. Cambridge University Press.
- [161] Giudici, G. et al. (2022). Dynamical Preparation of Quantum Spin Liquids in Rydberg Atom Arrays. *Physical Review Letters*, 129(9), 090401.
- [162] Gjonbalaj, N. O. et al. (2025). Shortcuts to Analog Preparation of Non-Equilibrium Quantum Lakes. arXiv:2502.03518.
- [163] Glidden, J. A. P. et al. (2021). Bidirectional dynamic scaling in an isolated Bose gas far from equilibrium. *Nat. Phys.*, 17(4), 457–461.
- [164] Goldman, N. et al. (2014). Periodically Driven Quantum Systems: Effective Hamiltonians and Engineered Gauge Fields. *Physical Review X*, 4(3), 031027.
- [165] Goldschmidt, E. et al. (2016). Anomalous Broadening in Driven Dissipative Rydberg Systems. *Physical Review Letters*, 116(11), 113001.
- [166] Gong, A. et al. (2024). Computation with quantum Reed-Muller codes and their mapping onto 2D atom arrays. arXiv preprint arXiv:2410.23263.
- [167] Gonzalez-Cuadra, D. et al. (2025). Observation of string breaking on a $(2 + 1)$ D Rydberg quantum simulator. *Nature*, 642(8067), 321–326.
- [168] González-Cuadra, D. et al. (2023). Fermionic quantum processing with programmable neutral atom arrays. *Proceedings of the National Academy of Sciences of the United States of America*, 120(35), e2304294120.
- [169] Goo, J. et al. (2022). Universal Early Coarsening of Quenched Bose Gases. *Phys. Rev. Lett.*, 128(13), 135701.

- [170] Google Quantum AI (2023). Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7949), 676–681.
- [171] Goto, H. (2016). Minimizing resource overheads for fault-tolerant preparation of encoded states of the Steane code. *Scientific Reports*, 6, 19578.
- [172] Gottesman, D. (1997). Stabilizer Codes and Quantum Error Correction. arXiv preprint quant-ph/9705052.
- [173] Gottesman, D. (2022). Opportunities and Challenges in Fault-Tolerant Quantum Computation. arXiv:2210.15844.
- [174] Gottesman, D. et al. (1999). Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature*, 402(6760), 390–393.
- [175] Graham, T. M. et al. (2019). Rydberg-mediated entanglement in a two-dimensional neutral atom qubit array. *Physical Review Letters*, 123, 230501.
- [176] Graham, T. M. et al. (2023). Mid-circuit measurements on a single-species neutral alkali atom quantum processor. *Physical Review X*, 13(4), 041051.
- [177] Graham, T. M. et al. (2022). Multi-qubit entanglement and algorithms on a neutral-atom quantum computer. *Nature*, 604(7906), 457–462.
- [178] Gregor, K. et al. (2010). Diagnosing Deconfinement and Topological Order. arXiv:1011.4187.
- [179] Greiner, M. et al. (2002). Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, 415(6867), 39–44.
- [180] Grier, A. T. et al. (2013). Λ -enhanced sub-Doppler cooling of lithium atoms in DI gray molasses. *Physical Review A*, 87(6), 063411.
- [181] Grinkemeyer, B. et al. (2025). Error-detected quantum operations with neutral atoms mediated by an optical cavity. *Science*, 387(6740), 1301–1307.
- [182] Gross, C. et al. (2017). Quantum simulations with ultracold atoms in optical lattices. *Science*, 357(6355), 995–1001.
- [183] Gu, A. et al. (2026). Scalable neural decoders for practical fault-tolerant quantum computation. arXiv:2604.08358.
- [184] Gu, A. et al. (2024). Practical Hamiltonian learning with unitary dynamics and Gibbs states. *Nature Communications*, 15(1), 312.
- [185] Gyger, F. et al. (2024). Continuous operation of large-scale atom arrays in optical lattices. *Physical Review Research*, 6(3), 033104.

- [186] Haegeman, J. et al. (2011). Time-dependent variational principle for quantum lattices. *Phys. Rev. Lett.*, 107(7), 070601.
- [187] Haegeman, J. et al. (2016). Unifying time evolution and optimization with matrix product states. *Phys. Rev. B*, 94(16), 165116.
- [188] Halimeh, J. C. et al. (2025). Quantum simulation of out-of-equilibrium dynamics in gauge theories. arXiv:2509.03586.
- [189] Hangleiter, D. et al. (2018). Anticoncentration theorems for schemes showing a quantum speedup. *Quantum*, 2, 65.
- [190] Hangleiter, D. et al. (2023). Bell sampling from quantum circuits. arXiv:2306.00083.
- [191] Hangleiter, D. et al. (2025). Fault-tolerant compiling of classically hard iqp circuits on hypercubes. *PRX Quantum*, 6(2), 020338.
- [192] Hansen, N. (2019). A global surrogate assisted CMA-ES. *GECCO 2019 - Proceedings of the 2019 Genetic and Evolutionary Computation Conference*, (pp. 664–672).
- [193] Hartke, T. et al. (2022). Quantum register of fermion pairs. *Nature*, 601(7894), 537–541.
- [194] Hartung, L. et al. (2024). A quantum-network register assembled with optical tweezers in an optical cavity. *Science*, 385, 179–183.
- [195] Hashizume, T. et al. (2021). Deterministic Fast Scrambling with Neutral Atom Arrays. *Physical Review Letters*, 126(20), 200603.
- [196] Haug, T. et al. (2023). Scalable Measures of Magic Resource for Quantum Computers. *PRX Quantum*, 4(1), 010301.
- [197] Hauschild, J. et al. (2018). Efficient numerical simulations with Tensor Networks: Tensor Network Python (TeNPy). *SciPost Physics Lecture notes*, (pp. 005).
- [198] Hayden, P. et al. (2007). Black holes as mirrors: Quantum information in random subsystems. *Journal of High Energy Physics*, 2007(9), 120.
- [199] Henley, C. L. (2010). The “Coulomb phase” in frustrated systems. *Annual Review of Condensed Matter Physics*, 1(1), 179–210.
- [200] Higgott, O. et al. (2023). Improved Decoding of Circuit Noise and Fragile Boundaries of Tailored Surface Codes. *Physical Review X*, 13(3), 031007.
- [201] Higgott, O. et al. (2025). Sparse Blossom: correcting a million errors per core second with minimum-weight matching. *Quantum*, 9, 1600.

- [202] Hohenberg, P. C. et al. (1977). Theory of dynamic critical phenomena. *Rev. Mod. Phys.*, 49(3), 435–479.
- [203] Holman, A. et al. (2026). Trapping of single atoms in metasurface optical tweezer arrays. *Nature*, 649, 859–865.
- [204] Honciuc Menendez, D. et al. (2024). Implementing fault-tolerant non-clifford gates using the $[[8,3,2]]$ color code. *Physical Review A*, 109, 062438.
- [205] Horsman, C. et al. (2012). Surface code quantum computing by lattice surgery. *New Journal of Physics*, 14(12), 123011.
- [206] Howard, M. et al. (2014). Contextuality supplies the 'magic' for quantum computation. *Nature*, 510(7505), 351–355.
- [207] Hu, B. et al. (2025). Site-selective cavity readout and classical error correction of a 5-bit atomic register. *Physical Review Letters*, 134(12), 120801.
- [208] Huang, H. Y. et al. (2022). Quantum advantage in learning from experiments. *Science*, 376(6598), 1182–1186.
- [209] Huang, H. Y. et al. (2020). Predicting many properties of a quantum system from very few measurements. *Nat. Phys.*, 16(10), 1050–1057.
- [210] Huang, H.-Y. et al. (2023). Learning Many-Body Hamiltonians with Heisenberg-Limited Scaling. *Phys. Rev. Lett.*, 130(20), 200403.
- [211] Huang, X. (2018). Robust simplex algorithm for online optimization. *Phys. Rev. Accel. Beams*, 21, 104601.
- [212] Huh, S. et al. (2024). Universality class of a spinor Bose–Einstein condensate far from equilibrium. *Nat. Phys.*, 20, 402–408.
- [213] Humayun, K. et al. (1991). Non-equilibrium dynamics of the ising model for t less-than/equal-to t_c . *J. Phys. A: Math. Gen.*, 24(8), 1915–1932.
- [214] Hutter, A. et al. (2014). Efficient Markov chain Monte Carlo algorithm for the surface code. *Physical Review A*, 89(2), 022326.
- [215] Hutzler, N. R. et al. (2017). Eliminating light shifts for single atom trapping. *New Journal of Physics*, 19(2), 023007.
- [216] Hémery, K. et al. (2024). Measuring the Loschmidt Amplitude for Finite-Energy Properties of the Fermi-Hubbard Model on an Ion-Trap Quantum Computer. *PRX Quantum*, 5(3), 030323.

- [217] IonQ (2023). Introducing IonQ Forte.
- [218] Iqbal, M. et al. (2025). Non-Abelian topological order and anyons on a trapped-ion processor. *Nature*, 626(7999), 505–511.
- [219] Isenhowe, L. et al. (2011). Multibit C_k NOT quantum gates via Rydberg blockade. *Quantum Inf. Process.*, 10(6), 755.
- [220] Itano, W. M. et al. (1990). Quantum Zeno effect. *Physical Review A*, 41(5), 2295.
- [221] Iverson, J. K. et al. (2020). Coherence in logical quantum channels. *New Journal of Physics*, 22(7), 073066.
- [222] Iyer, P. et al. (2017). A Small Quantum Computer is Needed to Optimize Fault-Tolerant Protocols. *Quantum Science and Technology*, 3(3), 030504.
- [223] Jain, A. et al. (2017). Higgs mode and its decay in a two-dimensional antiferromagnet. *Nat. Phys.*, 13(7), 633–637.
- [224] Jaksch, D. et al. (2000). Fast quantum gates for neutral atoms. *Phys. Rev. Lett.*, 85, 2208–2211.
- [225] Jandura, S. et al. (2022). Time-Optimal Two- And Three-Qubit Gates for Rydberg Atoms. *Quantum*, 6, 712.
- [226] Jandura, S. et al. (2023). Optimizing rydberg gates for logical-qubit performance. *PRX Quantum*, 4, 020336.
- [227] Jenkins, A. et al. (2022). Ytterbium nuclear-spin qubits in an optical tweezer array. *Phys. Rev. X*, 12, 021027.
- [228] Jia, Y. et al. (2020). Chaos on the hypercube. *Journal of High Energy Physics*, 2020(11), 1–46.
- [229] Jordan, P. et al. (1928). über das Paulische äquivalenzverbot. *Z. Phys.*, 47(9-10), 631–651.
- [230] Jurcevic, P. et al. (2021). Demonstration of quantum volume 64 on a superconducting quantum computing system. *Quantum Sci. Technol.*, 6(2), 025020.
- [231] Kalinowski, M. et al. (2023). Non-abelian floquet spin liquids in a digital rydberg simulator. *Physical Review X*, 13(3), 031008.
- [232] Kasteleyn, P. W. (1961). The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice. *Physica*, 27(12), 1209–1225.
- [233] Kasteleyn, P. W. (1963). Dimer Statistics and Phase Transitions. *Journal of Mathematical Physics*, 4(2), 287–293.

- [234] Katz, O. et al. (2025). Hybrid digital-analog protocols for simulating quantum multi-body interactions. *arXiv:2512.21385*.
- [235] Kaufman, A. M. et al. (2016). Quantum thermalization through entanglement in an isolated many-body system. *Science*, 353(6301), 794–800.
- [236] Keesling, A. et al. (2019). Quantum Kibble-Zurek mechanism and critical dynamics on a programmable Rydberg simulator. *Nature*, 568(7751), 207–211.
- [237] Kells, G. et al. (2009). Description of Kitaev’s honeycomb model with toric-code stabilizers. *Phys. Rev. B*, 80(12), 125415.
- [238] Kennedy, T. (1991). Ornstein-Zernike decay in the ground state of the quantum Ising model in a strong transverse field. *Commun. Math. Phys.*, 137(3), 599–615.
- [239] Khazali, M. et al. (2020). Fast Multiqubit Gates by Adiabatic Evolution in Interacting Excited-State Manifolds of Rydberg Atoms and Superconducting Circuits. *Phys. Rev. X*, 10, 021054.
- [240] Kim, D. et al. (2019). Large-scale uniform optical focus array generation with a phase spatial light modulator. *Opt. Lett.*, 44(12), 3178–3181.
- [241] Kim, Y. et al. (2023). Evidence for the utility of quantum computing before fault tolerance. *Nature*, 618(7965), 500–505.
- [242] Kim, Y. et al. (2022). High-fidelity three-qubit iToffoli gate for fixed-frequency superconducting qubits. *Nat. Phys.*, 18(7), 783–788.
- [243] Kitaev, A. (2006). Anyons in an exactly solved model and beyond. *Ann. Phys.*, 321(1), 2–111.
- [244] Kitaev, A. Y. (1997). Quantum computations: Algorithms and error correction. *Russian Mathematical Surveys*, 52(6), 1191–1249.
- [245] Kitaev, A. Y. (2003). Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1), 2–30.
- [246] Knap, M. et al. (2013). Probing Real-Space and Time-Resolved Correlation Functions with Many-Body Ramsey Interferometry. *Physical Review Letters*, 111(14), 147205.
- [247] Knill, E. (2005). Quantum computing with realistically noisy devices. *Nature*, 434(7029), 39–44.
- [248] Kokail, C. et al. (2026). Inverse Quantum Simulation for Quantum Material Design. *arXiv:2601.12239*.
- [249] Koyluoglu, N. U. et al. (2025). Floquet Engineering of Interactions and Entanglement in Periodically Driven Rydberg Chains. *Physical Review Letters*, 135(11), 113603.

- [250] Krastanov, S. et al. (2017). Deep Neural Network Probabilistic Decoder for Stabilizer Codes. *Scientific Reports*, 7(1), 1–7.
- [251] Kraus, C. V. et al. (2010). Generalized Hartree–Fock theory for interacting fermions in lattices: numerical methods. *New J. Phys.*, 12(11), 113004.
- [252] Kribs, D. et al. (2005). Unified and Generalized Approach to Quantum Error Correction. *Physical Review Letters*, 94(18), 180501.
- [253] Krinner, S. et al. (2022). Realizing repeated quantum error correction in a distance-three surface code. *Nature*, 605(7911), 669–674.
- [254] Kubica, A. et al. (2015a). Universal transversal gates with color codes: A simplified approach. *Physical Review A*, 91(3), 032330.
- [255] Kubica, A. et al. (2015b). Unfolding the color code. *New Journal of Physics*, 17(8), 083026.
- [256] Kuriyattil, S. et al. (2023). Onset of Scrambling as a Dynamical Transition in Tunable-Range Quantum Circuits. *PRX Quantum*, 4(3), 030325.
- [257] Lampen, J. et al. (2018). Long-lived coherence between ground and rydberg levels in a magic-wavelength lattice. *Phys. Rev. A*, 98, 33411.
- [258] Landauer, R. (1961). Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development*, 5(3), 183–191.
- [259] Landauer, R. (1995). Is quantum mechanics useful? *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 353(1703), 367–376.
- [260] Le Kien, F. et al. (2013). Dynamical polarizability of atoms in arbitrary light fields: General theory and application to cesium. *European Physical Journal D*, 67(5), 92.
- [261] Lee, P. A. et al. (2006). Doping a Mott insulator: Physics of high-temperature superconductivity. *Rev. Mod. Phys.*, 78(1), 17–85.
- [262] Levine, H. et al. (2022). Dispersive optical systems for scalable raman driving of hyperfine qubits. *Physical Review A*, 105(3), 032618.
- [263] Levine, H. et al. (2018). High-fidelity control and entanglement of rydberg-atom qubits. *Phys. Rev. Lett.*, 121, 123603.
- [264] Levine, H. et al. (2019). Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms. *Phys. Rev. Lett.*, 123(17), 170503.
- [265] Li, R. et al. (2022). Optimal Model for Fewer-Qubit CNOT Gates With Rydberg Atoms. *Phys. Rev. Appl.*, 17, 024014.

- [266] Li, R. et al. (2023). Proposal for practical rydberg quantum gates using a native two-photon excitation. *Quantum Science and Technology*, 8(3), 035032.
- [267] Li, Y. et al. (2025). Fast, continuous and coherent atom replacement in a neutral atom qubit array. arXiv:2506.15633.
- [268] Lib, O. et al. (2026). Velocity-enabled quantum computing with neutral atoms. arXiv:2603.15561.
- [269] Lidar, D. A. (2014). Review of Decoherence-Free Subspaces, Noiseless Subsystems, and Dynamical Decoupling. In *Advances in Chemical Physics* (pp. 295–354). Wiley.
- [270] Lieb, E. H. (1967). Solution of the Dimer Problem by the Transfer Matrix Method. *Journal of Mathematical Physics*, 8(12), 2339–2341.
- [271] Lieb, E. H. (1994). Flux Phase of the Half-Filled Band. *Phys. Rev. Lett.*, 73(16), 2158–2161.
- [272] Lienhard, V. et al. (2018). Observing the Space- and Time-Dependent Growth of Correlations in Dynamically Tuned Synthetic Ising Models with Antiferromagnetic Interactions. *Phys. Rev. X*, 8, 021070.
- [273] Lifshitz, I. M. (1962). Kinetics of ordering during second-order phase transitions. *Sov. Phys. JETP*, 15(5), 939–942.
- [274] Lin, R. et al. (2026). Sustaining high-fidelity quantum logic in neutral-atom circuits via mid-circuit operations. arXiv:2603.01612.
- [275] Lin, R. et al. (2025). Ai-enabled parallel assembly of thousands of defect-free neutral atom arrays. *Phys. Rev. Lett.*, 135, 060602.
- [276] Linke, N. M. et al. (2017). Fault-tolerant quantum error detection. *Science Advances*, 3(10).
- [277] Lis, J. W. et al. (2023). Mid-circuit operations using the omg-architecture in neutral atom arrays. *Physical Review X*, 13(4), 041035. Published 22 November 2023.
- [278] Litinski, D. et al. (2022). Active volume: An architecture for efficient fault-tolerant quantum computers with limited non-local connections. arXiv:2211.15465.
- [279] Lloyd, S. (1996). Universal Quantum Simulators. *Science*, 273(5278), 1073–1078.
- [280] Lu, S. et al. (2021). Algorithms for Quantum Simulation at Finite Energies. *PRX Quantum*, 2(2), 020321.
- [281] Lu, T.-C. et al. (2022). Measurement as a Shortcut to Long-Range Entangled Quantum Matter. *PRX Quantum*, 3(4), 040337.

- [282] Ma, S. et al. (2023). High-fidelity gates and mid-circuit erasure conversion in a metastable neutral-atom qubit. *Nature*, 622(7982), 279–284.
- [283] Madsen, L. S. et al. (2022). Quantum computational advantage with a programmable photonic processor. *Nature*, 606(7912), 75–81.
- [284] Magesan, E. et al. (2012). Efficient measurement of quantum gate error by interleaved randomized benchmarking. *Phys. Rev. Lett.*, 109, 080505.
- [285] Mahesh, T. P. et al. (2024). Generation of 480 nm picosecond pulses for ultrafast excitation of rydberg atoms. arXiv preprint arXiv:2408.02324.
- [286] Maldacena, J. (2023). A simple quantum system that describes a black hole. arXiv:2303.11534.
- [287] Mandel, O. et al. (2003). Coherent Transport of Neutral Atoms in Spin-Dependent Optical Lattice Potentials. *Physical Review Letters*, 91(1), 010407.
- [288] Manetsch, H. J. et al. (2025). A tweezer array with 6,100 highly coherent atomic qubits. *Nature*, 647, 60–67.
- [289] Manovitz, T. et al. (2025). Quantum coarsening and collective dynamics on a programmable quantum simulator. *Nature*, 638(8049), 86–92.
- [290] Maraga, A. et al. (2015). Aging and coarsening in isolated quantum systems after a quench: Exact results for the quantum $O(N)$ model with $N \rightarrow \infty$. *Phys. Rev. E*, 92(4), 042151.
- [291] Martin, M. J. et al. (2021). A mølmer-sørensen gate with rydberg-dressed atoms. arXiv:2111.14677.
- [292] Maskara, N. et al. (2025a). Fast simulation of fermions with reconfigurable qubits. arXiv:2509.08898.
- [293] Maskara, N. et al. (2021). Discrete Time-Crystalline Order Enabled by Quantum Many-Body Scars: Entanglement Steering via Periodic Driving. *Physical Review Letters*, 127(9), 090602.
- [294] Maskara, N. et al. (2025b). Programmable simulations of molecules and materials with reconfigurable quantum processors. *Nature Physics*, 21(2), 289–297.
- [295] McEwen, M. et al. (2023). Relaxing Hardware Requirements for Surface Code Circuits using Time-dynamics. *Quantum*, 7, 1172.
- [296] Mezher, R. et al. (2020). Fault-tolerant quantum speedup from constant depth quantum circuits. *Physical Review Research*, 2(3), 033444.

- [297] Mi, X. et al. (2022a). Time-crystalline eigenstate order on a quantum processor. *Nature*, 601, 531–536.
- [298] Mi, X. et al. (2021). Information scrambling in quantum circuits. *Science*, 374(6574), 1479–1483.
- [299] Mi, X. et al. (2022b). Noise-resilient edge modes on a chain of superconducting qubits. *Science*, 378(6621), 785–790.
- [300] Miller, C. et al. (2024). Two-axis twisting using Floquet-engineered XYZ spin models with polar molecules. *Nature*, 633(8029), 332–337.
- [301] Modak, R. et al. (2020). Entanglement revivals as a probe of scrambling in finite quantum systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(8), 083110.
- [302] Moessner, R. et al. (2008). Quantum dimer models. arXiv:0809.3051.
- [303] Moessner, R. et al. (2001). Phase diagram of the hexagonal lattice quantum dimer model. *Physical Review B*, 64(14), 144416.
- [304] Mohan, M. et al. (2023). Robust control and optimal rydberg states for neutral-atom two-qubit gates. *Physical Review Research*, 5(3), 033052.
- [305] Møller, D. et al. (2008). Quantum gates and multiparticle entanglement by rydberg excitation blockade and adiabatic passage. *Phys. Rev. Lett.*, 100, 170504.
- [306] Monz, T. et al. (2009). Realization of the quantum toffoli gate with trapped ions. *Phys. Rev. Lett.*, 102, 040501.
- [307] Monz, T. et al. (2011). 14-qubit entanglement: Creation and coherence. *Physical Review Letters*, 106(13), 130506.
- [308] Morong, W. et al. (2023). Engineering Dynamically Decoupled Quantum Simulations with Trapped Ions. *PRX Quantum*, 4(1), 010334.
- [309] Morvan, A. et al. (2024). Phase transitions in random circuit sampling. *Nature*, 634(8035), 328–333.
- [310] Müller, M. M. et al. (2011). Optimizing entangling quantum gates for physical systems. *Phys. Rev. A*, 84, 042315.
- [311] Muniz, J. A. et al. (2025). Repeated ancilla reuse for logical computation on a neutral atom quantum computer. *Physical Review X*, 15(4), 041040.
- [312] Nelder, J. A. et al. (1965). A Simplex Method for Function Minimization. *Comput. J.*, 7(4), 308–313.

- [313] Nigmatullin, R. et al. (2025). Experimental demonstration of breakeven for a compact fermionic encoding. *Nat. Phys.*, 21(8), 1319–1325.
- [314] Noh, H.-R. et al. (2012). Transmittance signal in real ladder-type atoms. *Phys. Rev. A*, 85(3), 033817.
- [315] Noiri, A. et al. (2022). Fast universal quantum gate above the fault-tolerance threshold in silicon. *Nature*, 601(7893).
- [316] Norcia, M. A. et al. (2023). Mid-circuit qubit measurement and rearrangement in a ^{171}Yb atomic array. *Physical Review X*, 13(4), 041034.
- [317] Norcia, M. A. et al. (2024). Iterative Assembly of ^{171}Yb Atom Arrays with Cavity-Enhanced Optical Lattices. *PRX Quantum*, 5(3), 030316.
- [318] Olsacher, T. et al. (2025). Hamiltonian and Liouvillian learning in weakly-dissipative quantum many-body systems. *Quantum Sci. Technol.*, 10(1), 015065.
- [319] Olson, M. et al. (2025). Ax: A Platform for Adaptive Experimentation. In *Proceedings of the Fourth International Conference on Automated Machine Learning*, volume 293 of *Proceedings of Machine Learning Research* (pp. 21:1–21:25).: PMLR.
- [320] Pagano, A. et al. (2022). Error budgeting for a controlled-phase gate with strontium-88 Rydberg atoms. *Phys. Rev. Res.*, 4, 033019.
- [321] Page, D. N. (1993). Average entropy of a subsystem. *Phys. Rev. Lett.*, 71, 1291–1294.
- [322] Paletta, L. et al. (2024). Robust sparse IQP sampling in constant depth. *Quantum*, 8, 1337.
- [323] Pan, F. et al. (2022). Simulation of Quantum Circuits Using the Big-Batch Tensor Network Method. *Physical Review Letters*, 128(3), 030501.
- [324] Pastori, L. et al. (2022). Characterization and Verification of Trotterized Digital Quantum Simulation via Hamiltonian and Liouvillian Learning. *PRX Quantum*, 3(3), 030324.
- [325] Patterson, D. A. et al. (2018). *Computer Organization and Design: The Hardware Software Interface - RISC-V Edition*.
- [326] Pavešić, L. et al. (2024). Constrained dynamics and confinement in the two-dimensional quantum ising model. *arXiv preprint arXiv:2406.11979*.
- [327] Pekker, D. et al. (2015). Amplitude/Higgs modes in condensed matter physics. *Annu. Rev. Condens. Matter Phys.*, 6(1), 269–297.
- [328] Pelegrí, G. et al. (2022). High-fidelity multiqubit Rydberg gates via two-photon adiabatic rapid passage. *Quantum Sci. Technol.*, 7(4), 045020.

- [329] Peper, M. et al. (2025). Spectroscopy and modeling of ^{171}Yb rydberg states for high-fidelity entangling gates. *Physical Review X*, 15(1), 011009.
- [330] Periwal, A. et al. (2021). Programmable interactions and emergent geometry in an array of atom clouds. *Nature*, 600(7890), 630–635.
- [331] Perrin, H. et al. (2026). Correlated atom loss as a resource for quantum error correction. *arXiv preprint arXiv:2603.24237*.
- [332] Petrosyan, D. et al. (2017). High-fidelity Rydberg quantum gate via a two-atom dark state. *Phys. Rev. A*, 96, 042306.
- [333] Pogorelov, I. et al. (2021). Compact ion-trap quantum computing demonstrator. *PRX Quantum*, 2, 020343.
- [334] Polyakov, A. M. (1987). *Gauge Fields and Strings*. Taylor & Francis.
- [335] Porter, C. E. et al. (1956). Fluctuations of nuclear reaction widths. *Physical Review*, 104, 483.
- [336] Postler, L. et al. (2022). Demonstration of fault-tolerant universal quantum gate operations. *Nature*, 605, 675–680.
- [337] Potirniche, I.-D. et al. (2017). Floquet symmetry-protected topological phases in cold atomic systems. *Physical Review Letters*, 119(12), 123601.
- [338] Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum*, 2, 79.
- [339] Putterman, H. et al. (2025). Hardware-efficient quantum error correction using concatenated bosonic qubits. *Nature*, 638(8052), 927–934.
- [340] Pyka, K. et al. (2013). Topological defect formation and spontaneous symmetry breaking in ion Coulomb crystals. *Nat. Commun.*, 4, 2291.
- [341] Qi, X.-L. et al. (2019). Determining a local Hamiltonian from a single eigenstate. *Quantum*, 3, 159.
- [342] Radnaev, A. G. et al. (2025). A universal neutral-atom quantum computer with individual optical addressing and non-destructive readout. *PRX Quantum*, 6(3), 030334.
- [343] Rao, D. D. B. et al. (2014). Robust Rydberg-interaction gates with adiabatic passage. *Phys. Rev. A*, 89, 030301.
- [344] Raussendorf, R. (2012). Key ideas in quantum error correction. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1975), 4541–4565.

- [345] Raussendorf, R. et al. (2001). A One-Way Quantum Computer. *Physical Review Letters*, 86(22), 5188.
- [346] Read, N. et al. (1989). Valence-bond and spin-peierls ground states of low-dimensional quantum antiferromagnets. *Phys. Rev. Lett.*, 62, 1694–1697.
- [347] Reichardt, B. W. et al. (2024a). Demonstration of quantum computation and error correction with a tesseract code. *arXiv preprint arXiv:2409.04628*.
- [348] Reichardt, B. W. et al. (2024b). Logical computation demonstrated with a neutral atom quantum processor. *arXiv preprint arXiv:2411.11822*.
- [349] Rines, R. et al. (2025). Demonstration of a logical architecture uniting motion and in-place entanglement: Shor’s algorithm, constant-depth cnot ladder, and many-hypercube code. *arXiv preprint arXiv:2509.13247*.
- [350] Robens, C. et al. (2017). Atomic “bomb testing”: the Elitzur–Vaidman experiment violates the Leggett–Garg inequality. *Appl. Phys. B*, 123(1), 12.
- [351] Robicheaux, F. et al. (2021). Photon-recoil and laser-focusing limits to rydberg gate fidelity. *Phys. Rev. A*, 103, 022424.
- [352] Rodriguez, P. S. et al. (2025). Experimental demonstration of logical magic state distillation. *Nature*, 645, 620–625.
- [353] Rokhsar, D. S. et al. (1988). Superconductivity and the Quantum Hard-Core Dimer Gas. *Physical Review Letters*, 61(20), 2376–2379.
- [354] Rosi, S. et al. (2013). Fast closed-loop optimal control of ultracold atoms in an optical lattice. *Physical Review A*, 88(2), 021601.
- [355] Roychowdhury, K. et al. (2021). Dynamics and correlations at a quantum phase transition beyond Kibble-Zurek. *Phys. Rev. B*, 104(1), 014406.
- [356] Rüegg, C. et al. (2008). Quantum magnets under pressure: controlling elementary excitations in TiCuCl_3 . *Phys. Rev. Lett.*, 100(20), 205701.
- [357] Ryan-Anderson, C. et al. (2022). Implementing Fault-tolerant Entangling Gates on the Five-qubit Code and the Color Code. *arXiv preprint arXiv:2208.01863*.
- [358] Sachdev, S. (1997). Theory of finite-temperature crossovers near quantum critical points close to, or above, their upper-critical dimension. *Phys. Rev. B*, 55(1), 142–161.
- [359] Sachdev, S. (2009). Exotic phases and quantum phase transitions: model systems and experiments. In *Quantum Theory of Condensed Matter*, 24th Solvay Conference on Physics.

- [360] Sachdev, S. (2011). *Quantum Phase Transitions*. Cambridge: Cambridge University Press, 2 edition.
- [361] Sachdev, S. (2023). *Quantum Phases of Matter*. Cambridge: Cambridge University Press.
- [362] Sadler, L. E. et al. (2006). Spontaneous symmetry breaking in a quenched ferromagnetic spinor Bose-Einstein condensate. *Nature*, 443(7109), 312–315.
- [363] Saffman, M. et al. (2020). Symmetric Rydberg controlled-Z gates with adiabatic pulses. *Phys. Rev. A*, 101, 062309.
- [364] Saffman, M. et al. (2010). Quantum information with Rydberg atoms. *Rev. Mod. Phys.*, 82(3), 2313–2363.
- [365] Safronova, M. S. et al. (2011). Critically evaluated theoretical energies, lifetimes, hyperfine constants, and multipole polarizabilities in rb. *Physical Review A*, 83(5), 052508.
- [366] Sahay, K. et al. (2023a). High threshold codes for neutral atom qubits with biased erasure errors. *Physical Review X*, 13(4), 041013.
- [367] Sahay, R. et al. (2025). Emergent holographic forces from tensor networks and criticality. *Physical Review X*, 15(2), 021078.
- [368] Sahay, R. et al. (2024). Finite-Depth Preparation of Tensor Network States from Measurement. *arXiv:2404.17087*.
- [369] Sahay, R. et al. (2023b). Quantum Spin Puddles and Lakes: NISQ-Era Spin Liquids from Non-Equilibrium Dynamics. *arXiv:2211.01381*.
- [370] Sales Rodriguez, P. et al. (2025). Experimental Demonstration of Logical Magic State Distillation. *Nature*, 645(8081), 620–625.
- [371] Samajdar, R. et al. (2020). Complex density wave orders and quantum phase transitions in a model of square-lattice Rydberg atom arrays. *Phys. Rev. Lett.*, 124(10), 103601.
- [372] Samajdar, R. et al. (2021). Quantum phases of Rydberg atoms on a kagome lattice. *PNAS*, 118(4), e2015785118.
- [373] Samajdar, R. et al. (2024). Quantum and classical coarsening and their interplay with the kibble-zurek mechanism. *arXiv preprint arXiv:2401.15144*.
- [374] Sandvik, A. W. (2010). Computational studies of quantum spin systems. In *AIP Conference Proceedings*, volume 1297 (pp. 135–338).
- [375] Sandvik, A. W. et al. (2006). Correlations and confinement in nonplanar two-dimensional dimer models. *Physical Review B*, 73(14), 144504.

- [376] Satzinger, K. J. et al. (2021). Realizing topologically ordered states on a quantum processor. *Science*, 374(6572), 1237–1241.
- [377] Schlosser, M. et al. (2011). Scalable architecture for quantum information processing with atoms in optical micro-structures. *Quantum Information Processing*, 10(6), 907–924.
- [378] Schlosser, N. et al. (2001). Sub-poissonian loading of single atoms in a microscopic dipole trap. *Nature*, 411, 1024–1027.
- [379] Schmitt, M. et al. (2022). Quantum phase transition dynamics in the two-dimensional transverse-field Ising model. *Sci. Adv.*, 8(37), eabl6850.
- [380] Scholl, P. et al. (2021). Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms. *Nature*, 595(7866), 233–238.
- [381] Scholl, P. et al. (2023). Erasure conversion in a high-fidelity rydberg quantum simulator. *Nature*, 622(7982), 273–278.
- [382] Scholl, P. et al. (2022). Microwave Engineering of Programmable XXZ Hamiltonians in Arrays of Rydberg Atoms. *PRX Quantum*, 3(2), 020303.
- [383] Schreiber, M. et al. (2015). Observation of many-body localization of interacting fermions in a quasirandom optical lattice. *Science*, 349(6250), 842–845.
- [384] Sekino, Y. et al. (2008). Fast scramblers. *Journal of High Energy Physics*, 2008(10).
- [385] Self, C. N. et al. (2022). Protecting Expressive Circuits with a Quantum Error Detection Code. *arXiv:2211.06703*.
- [386] Semeghini, G. et al. (2021). Probing Topological Spin Liquids on a Programmable Quantum Simulator. *Science*, 374(6572), 1242–1247.
- [387] Senoo, A. et al. (2025). High-fidelity entanglement and coherent multi-qubit mapping in an atom array. *arXiv:2506.13632*.
- [388] Shao, H. J. et al. (2024). Antiferromagnetic phase transition in a 3d fermionic hubbard model. *Nature*, 632, 267–272.
- [389] Shao, X.-Q. et al. (2014). One-step achievement of robust multipartite Greenberger-Horne-Zeilinger state and controlled-phase gate via Rydberg interaction. *J. Opt. Soc. Am. B*, 31(4), 827–832.
- [390] Shaw, A. L. et al. (2025). A cavity array microscope for parallel single-atom interfacing. *arXiv preprint arXiv:2506.10919*.
- [391] Shea, M. E. et al. (2020). Submillisecond, nondestructive, time-resolved quantum-state readout of a single, trapped neutral atom. *Physical Review A*, 102(5), 053101.

- [392] Shen, C.-P. et al. (2019). Construction of robust Rydberg controlled-phase gates. *Opt. Lett.*, 44(8), 2036–2039.
- [393] Shende, V. V. et al. (2009). On the cnot-cost of toffoli gates. *Quantum Inf. Comput.*, 9(5-6).
- [394] Shepherd, D. et al. (2009). Temporally unstructured quantum computation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2105).
- [395] Shi, X.-F. (2017). Rydberg Quantum Gates Free from Blockade Error. *Phys. Rev. Appl.*, 7, 064017.
- [396] Shi, X.-F. et al. (2014). Entanglement of neutral-atom chains by spin-exchange Rydberg interaction. *Phys. Rev. A*, 90, 062327.
- [397] Shimano, R. et al. (2020). Higgs mode in superconductors. *Annu. Rev. Condens. Matter Phys.*, 11, 103–124.
- [398] Shor, P. W. (1995). Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4), R2493.
- [399] Shor, P. W. (1996). Fault-tolerant quantum computation. In *Annual Symposium on Foundations of Computer Science - Proceedings* (pp. 56–65).: IEEE.
- [400] Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5).
- [401] Simon, R. et al. (1987). Cross polarization in laser beams. *Applied Optics*, 26(9), 1589–1593.
- [402] Singh, K. et al. (2023). Mid-circuit correction of correlated phase errors using an array of spectator qubits. *Science*, 380(6651), 1265–1269.
- [403] Sivak, V. et al. (2025). Reinforcement learning control of quantum error correction. *arXiv:2511.08493*.
- [404] Sivak, V. V. et al. (2023). Real-time quantum error correction beyond break-even. *Nature*, 616(7955), 50–55.
- [405] Smith, S. C. et al. (2024). Mitigating errors in logical qubits. *Communications Physics*, 7(1), 1–10.
- [406] Spar, B. M. et al. (2022). Realization of a Fermi-Hubbard Optical Tweezer Array. *Phys. Rev. Lett.*, 128(22), 223202.
- [407] Srakaew, K. et al. (2023). A subwavelength atomic array switched by a single Rydberg atom. *Nat. Phys.*, 19, 714–719.

- [408] Stace, T. M. et al. (2009). Thresholds for topological codes in the presence of loss. *Physical Review Letters*, 102(20), 200501.
- [409] Stas, P.-J. et al. (2026). Entanglement-assisted non-local optical interferometry in a quantum network. *Nature*, 651, 326–332.
- [410] Steane, A. (1996). Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 452(1954), 2551–2577.
- [411] Steane, A. M. (1997). Active stabilization, quantum computation, and quantum state synthesis. *Physical Review Letters*, 78(11), 2252.
- [412] Su, S.-L. et al. (2023). Rabi- and Blockade-Error-Resilient All-Geometric Rydberg Quantum Gates. *Phys. Rev. Appl.*, 19, 044007.
- [413] Sun, Y. et al. (2020). Controlled Phase Gate Protocol for Neutral Atoms via Off-Resonant Modulated Driving. *Phys. Rev. Appl.*, 13, 024059.
- [414] Takagi, H. et al. (2019). Concept and realization of Kitaev quantum spin liquids. *Nat. Rev. Phys.*, 1(4), 264–280.
- [415] Tan, D. B. et al. (2024). Compiling quantum circuits for dynamically field-programmable neutral atom array processors. *Quantum*, 8, 1281.
- [416] Tantivasadakarn, N. et al. (2023). Hierarchy of topological order from finite-depth unitaries, measurement, and feedforward. *PRX Quantum*, 4.
- [417] Tao, R. et al. (2024). High-fidelity detection of large-scale atom arrays in an optical lattice. *Physical Review Letters*, 133, 013401.
- [418] Tazhigulov, R. N. et al. (2022). Simulating Models of Challenging Correlated Molecules and Materials on the Sycamore Quantum Processor. *PRX Quantum*, 3, 040318.
- [419] Temperley, H. N. V. et al. (1961). Dimer problem in statistical mechanics-an exact result. *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics*, 6(68), 1061–1063.
- [420] Teng, Y. et al. (2025). Learning topological states from randomized measurements using variational tensor-network tomography. *PRX Quantum*, 6(4), 040303.
- [421] Teo, J. C. Y. et al. (2010). Topological defects and gapless modes in insulators and superconductors. *Phys. Rev. B*, 82(11), 115120.
- [422] Theis, L. S. et al. (2016). High-fidelity rydberg-blockade entangling gate using shaped, analytic pulses. *Phys. Rev. A*, 94, 032306.

- [423] Thompson, J. D. et al. (2013). Coherence and Raman Sideband Cooling of a Single Atom in an Optical Tweezer. *Physical Review Letters*, 110(13), 133001.
- [424] Tóth, G. et al. (2005). Entanglement detection in the stabilizer formalism. *Physical Review A*, 72(2), 022340.
- [425] Tsai, R. B. S. et al. (2025). Benchmarking and Fidelity Response Theory of High-Fidelity Rydberg Entangling Gates. *PRX Quantum*, 6(1), 010331.
- [426] Tuchendler, C. et al. (2008). Energy distribution and cooling of a single atom in an optical tweezer. *Physical Review A*, 78(3), 033425.
- [427] Vandersypen, L. M. K. et al. (2005). NMR techniques for quantum control and computation. *Reviews of Modern Physics*, 76(4), 1037.
- [428] Vasmer, M. et al. (2022). Morphing Quantum Codes. *PRX Quantum*, 3(3), 030319.
- [429] Verresen, R. et al. (2021). Prediction of Toric Code Topological Order from Rydberg Blockade. *Physical Review X*, 11(3), 031005.
- [430] Verresen, R. et al. (2019). Stable Luttinger liquids and emergent $U(1)$ symmetry in constrained quantum chains. *arXiv:1903.09179*.
- [431] Verstraete, F. et al. (2005). Mapping local Hamiltonians of fermions to local Hamiltonians of spins. *J. Stat. Mech.: Theory Exp.*, 2005(09), P09012.
- [432] Vidal, G. et al. (2002). Computable measure of entanglement. *Physical Review A*, 65(3), 032314.
- [433] Vu, D. et al. (2025). Optimizing the dynamical preparation of quantum spin lakes on the ruby lattice. *arXiv:2512.09040*.
- [434] Walhout, M. et al. (1992). Sigma-plus–Sigma-minus Optical molasses in a longitudinal magnetic field. *JOSA B*, 9(11), 1997–2007.
- [435] Wang, Y. et al. (2024). Fault-tolerant one-bit addition with the smallest interesting colour code. *Science Advances*, 10(44), ead09024.
- [436] Ware, B. et al. (2023). A sharp phase transition in linear cross-entropy benchmarking. *arXiv:2305.04954*.
- [437] Wei, X. et al. (2026). A 10 megahertz spatial light modulator. *arXiv:2601.08906*.
- [438] Weiss, D. S. et al. (2004). Another way to approach zero entropy for a finite system of atoms. *Physical Review A*, 70(4), 040302.

- [439] White, R. et al. (2026). Quantum Cellular Automata on a Dual-Species Rydberg Processor. *arXiv:2601.16257*.
- [440] Wick, G. C. (1950). The Evaluation of the Collision Matrix. *Phys. Rev.*, 80(2), 268–272.
- [441] Wilde, F. et al. (2022). Scalably learning quantum many-body Hamiltonians from dynamical data. *arXiv:2209.14328*.
- [442] Wilkins, N. et al. (2023). Derivation of field theory for the classical dimer model using bosonization. *Physical Review E*, 107(5), 054126.
- [443] Will, M. et al. (2025). Probing non-equilibrium topological order on a quantum processor. *Nature*, 645(8080), 348–353.
- [444] Wimperis, S. (1994). Broadband, Narrowband, and Passband Composite Pulses for Use in Advanced NMR Experiments. *Journal of Magnetic Resonance, Series A*, 109(2), 221–231.
- [445] Wu, H.-Z. et al. (2010). Implementation of a multiqubit quantum phase gate in a neutral atomic ensemble via the asymmetric Rydberg blockade. *Phys. Rev. A*, 82, 034307.
- [446] Wu, T.-Y. et al. (2019). Stern–Gerlach detection of neutral-atom qubits in a state-dependent optical lattice. *Nature Physics*, 15(6), 538–542.
- [447] Wu, Y. et al. (2021). Strong Quantum Computational Advantage Using a Superconducting Quantum Processor. *Physical Review Letters*, 127(18), 180501.
- [448] Wu, Y. et al. (2022). Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays. *Nat. Commun.*, 13(1), 4657.
- [449] Xia, T. et al. (2015). Randomized benchmarking of single-qubit gates in a 2d array of neutral-atom qubits. *Phys. Rev. Lett.*, 114, 100503.
- [450] Xu, M. et al. (2025a). A neutral-atom Hubbard quantum simulator in the cryogenic regime. *Nature*, 642(8069), 909–915.
- [451] Xu, M. et al. (2023). Frustration- and doping-induced magnetism in a fermi–hubbard simulator. *Nature*, 620, 971–976.
- [452] Xu, Q. et al. (2024). Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays. *Nature Physics*, 20(7), 1084–1090.
- [453] Xu, Q. et al. (2025b). Fast and Parallelizable Logical Computation with Homological Product Codes. *Phys. Rev. X*, 15, 021065.
- [454] Xu, W. et al. (2021). Fast Preparation and Detection of a Rydberg Qubit Using Atomic Ensembles. *Phys. Rev. Lett.*, 127, 050501.

- [455] Young, A. W. et al. (2020). Half-minute-scale atomic coherence and high relative stability in a tweezer clock. *Nature*, 588, 408–413.
- [456] Young, A. W. et al. (2022). Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice. *Science*, 377(6608), 885–889.
- [457] Young, D. J. et al. (2024). Observing dynamical phases of BCS superconductors in a cavity QED simulator. *Nature*, 625(7996), 679–684.
- [458] Young, J. T. et al. (2021). Asymmetric Blockade and Multiqubit Gates via Dipole-Dipole Interactions. *Phys. Rev. Lett.*, 127, 120501.
- [459] Yu, D. et al. (2019). Adiabatic and high-fidelity quantum gates with hybrid Rydberg-Rydberg interactions. *Opt. Express*, 27(16), 23080–23094.
- [460] Zaletel, M. P. et al. (2015). Time-evolving a matrix product state with long-ranged interactions. *Physical Review B*, 91(16), 165112.
- [461] Zenesini, A. et al. (2024). False vacuum decay via bubble formation in ferromagnetic superfluids. *Nature Physics*, 20(4), 558–563.
- [462] Zeng, H.-B. et al. (2023). Universal breakdown of Kibble-Zurek scaling in fast quenches across a phase transition. *Phys. Rev. Lett.*, 130(6), 060402.
- [463] Zeng, Z. et al. (2025). Adiabatic echo protocols for robust quantum many-body state preparation. *arXiv:2506.12138*.
- [464] Zhang, B. et al. (2025). Leveraging erasure errors in logical qubits with metastable ^{171}Yb atoms. *arXiv:2506.13724*.
- [465] Zhong, H.-S. et al. (2020). Quantum computational advantage using photons. *Science*, 370(6523), 1460–1463.
- [466] Zhou, H. et al. (2025). Low-overhead transversal fault tolerance for universal quantum computation. *Nature*, 646(8046), 303–308.
- [467] Zhou, Z. et al. (2021). Emergent Rokhsar-Kivelson point in realistic quantum Ising models. *arXiv:2106.05518*.
- [468] Zupancic, P. et al. (2016). Ultra-precise holographic beam shaping for microscopic quantum control. *Opt. Express*, 24(13), 13881.
- [469] Zurek, W. H. et al. (2005). Dynamics of a Quantum Phase Transition. *Physical Review Letters*, 95(10), 105701.
- [470] Šibalić, N. et al. (2017). Arc: An open-source library for calculating properties of alkali rydberg atoms. *Computer Physics Communications*, 220, 319–331.