Universal Neutral-Atom Quantum Computer with Individual Optical Addressing and Nondestructive Readout

A.G. Radnaev,[1,*] W.C. Chung,[1,*] D.C. Cole,[1,*] D. Mason,[1,*] T.G. Ballance,[4] M.J. Bedalov,[2] D.A. Belknap,[2] M.R. Berman,[1] M. Blakely,[2] I.L. Bloomfield,[1] P.D. Buttler,[1] C. Campbell,[3] A. Chopinaud,[4] E. Copenhaver,[1] M.K. Dawes,[1] S.Y. Eubanks,[1] A.J. Friss,[1] D.M. Garcia,[1] J. Gilbert,[1] M. Gillette,[2] P. Goiporia,[3] P. Gokhale,[3] J. Goldwin,[2] D. Goodwin,[2] T.M. Graham,[5] C.J. Guttormsson,[1] G.T. Hickman,[1] L. Hurtley,[2] M. Iliev,[1] E.B. Jones,[1] R.A. Jones,[1] K.W. Kuper,[1] T.B. Lewis,[1] M.T. Lichtman,[2] F. Majdeteimouri,[1] J.J. Mason,[2] J.K. McMaster,[1] J.A. Miles,[2] P.T. Mitchell,[1] J.D. Murphree,[2] N.A. Neff-Mallon,[2] T. Oh,[1] V. Omole,[3] C. Parlo Simon,[1] N. Pederson,[2] M.A. Perlin,[3] A. Reiter,[1] R. Rines,[3] P. Romlow,[2] A.M. Scott,[1] D. Stiefvater,[1] J.R. Tanner,[1] A.K. Tucker,[1] I.V. Vinogradov,[1] M.L. Warter,[1] M. Yeo,[1] M. Saffman,[2,5] and T.W. Noel[1,†]

[1] *Infleqtion, Boulder, Colorado 80301, USA*
[2] *Infleqtion, Madison, Wisconsin 53703, USA*
[3] *Infleqtion, Chicago 60604, Illinois, USA*
[4] *Infleqtion UK, Oxford OX5 1GN, United Kingdom*
[5] *Department of Physics, University of Wisconsin-Madison, Madison, Wisconsin 53706, USA*

(Received 19 January 2025; accepted 11 July 2025; published 25 August 2025)

Quantum computers must achieve large-scale fault-tolerant operation to deliver on their promise of transformational processing power. This will require thousands or millions of high-fidelity quantum gates and similar numbers of qubits. Demonstrations using neutral-atom qubits trapped and manipulated by lasers have shown that this modality can provide high two-qubit gate (CZ) fidelities and scalable operation. However, the gates in these demonstrations are driven by lasers that do not resolve individual qubits, with universal computation enabled by physical midcircuit shuttling of the qubits. This relatively slow operation may greatly extend run times for useful large-scale computation. Here, we demonstrate a universal neutral-atom quantum computer with gate rates limited by optical switching times, rather than shuttling, by individually addressing tightly focused laser beams at an array of single atoms. We achieve CZ fidelity of 99.35(4)% and local single-qubit R_Z gate fidelity of 99.902(8)% in both cases accounting for leakage out of the computational basis. Moreover, we demonstrate nondestructive readout of alkali-atom qubits with 0.9(3)% loss, which boosts operational speed. This technique also enables us to measure a state-of-the-art CZ fidelity of 99.73(3)% when excluding atom-loss events, which may be mitigated through erasure conversion. Our results represent a critical step toward large-scale fault-tolerant neutral-atom quantum computers that can execute computations on practical timescales.

DOI: 10.1103/66s8-jj18

I. INTRODUCTION

Quantum computers promise to reduce computational energy consumption and solve classically expensive or intractable problems in a range of fields including material science, machine learning, security, chemistry, energy, finance, medicine, and pharmaceutical development [1–9]. Recent analyses suggest that realization of practical utility will require fault-tolerant operation [10]. This in turn requires quantum error correction (QEC), which comes with large overhead in the number of qubits and requires high one- and two-qubit quantum gate fidelities. The progress in the field along this path has entered a phase of experimental demonstrations [11–16]. Among the diverse modalities proposed for large-scale fault-tolerant quantum computation [17], the neutral-atom approach is one of the most promising candidates and has become a leader in

*These authors contributed equally to this work.
†Contact author: tom.noel@infleqtion.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

the combination of qubit count, universal gate-based circuit operation, connectivity, and quantum error correction [11–13,18–22]. The rapid progress of neutral atoms for quantum computing can be attributed to the combination of naturally identical qubits; long coherence times due to decoupling from the environment [23]; switchable qubit interaction strengths across 12 orders of magnitude [24]; submicrosecond high-fidelity scalable gates [25–27]; and the ability to trap and manipulate large dense multidimensional qubit arrays [18,28,29], as has recently been highlighted in demonstrations of error-corrected logical qubits [11–13].

Two architectural approaches have emerged for neutral-atom quantum processors. One [11–13] relies on simultaneously shuttling multiple qubits to and from a single dedicated zone for implementing gates. This allows flexible qubit connectivity and avoids the need for near-diffraction-limited addressing beams but increases compilation complexity [30] and limits gate rates (the number of gates per unit of time within a circuit) due to constraints on the speed of atom movement [31]. High gate rates are important for maintaining practical computation times, particularly in light of the physical-qubit count and gate-depth overheads required for fault-tolerant operation. An alternative approach [19] keeps the qubits stationary and addresses them individually with steerable laser beams, enabling gate rates limited only by optical switching, which is inherently faster than atom movement. This stationary-qubit approach poses challenges for achieving high gate fidelities and does not natively support all-to-all connectivity but it offers the promise of faster circuits, in direct support of long-term fault-tolerant operations [32]. Here, we present a full-stack neutral-atom quantum computer using gates based on individual optical addressing instead of atom motion. Overcoming the challenges specific to this architecture with dual-beamline optical design, optimized Rydberg-based gates, advanced tuning protocols, and other system upgrades described in this paper, we achieve an entangling gate fidelity of 99.35(4)%, already exceeding QEC thresholds for certain codes [33–36]. Our results open the door to faster approaches for fault-tolerant operation that exclusively use individual optical addressing of stationary qubits or that combine this with some atom shuttling to optimize execution of error-corrected circuits. We further advance the state of the art by incorporating nondestructive qubit readout, enabling faster computations through qubit reuse.

II. A NEUTRAL-ATOM ARCHITECTURE WITH INDIVIDUAL OPTICAL ADDRESSING

Building on prior work [19], we present a full-stack neutral-atom quantum computer using an individual optical addressing architecture. Arbitrary quantum circuits can be submitted via the cloud, where they are compiled [37] into a universal native gate set consisting of global microwave $R(\theta,\phi)$ rotations, arbitrary local $R_Z(\phi)$ rotations, and nearest-neighbor controlled-Z (CZ) gates. The qubits are optically trapped Cs atoms, with $|0\rangle$ and $|1\rangle$ encoded in the $|F=3,m_F=0\rangle$ and $|F=4,m_F=0\rangle$ clock states of the ground-state hyperfine manifold, respectively. The atoms are trapped out of a cold atomic beam [38], cooled, and rearranged by an optical tweezer to form a defect-free array. In Fig. 1, we illustrate a 24-qubit computational array alongside 12 reservoir atoms, formed in a 1064-nm tweezer array (see Sec. VI A). After state preparation via optical pumping, arbitrary circuits are implemented with microwave and laser pulses, followed by nondestructive state-selective readout (NDSSR) of the qubit state. During this readout, the qubits remain trapped in the optical tweezers with high probability, allowing reinitialization and reuse of the same atomic array for the next shot (see Sec. IV). After the requested number of measurements is collected, circuit results are returned to the user. Throughout this process, feedforward stabilization is applied to background electric and magnetic fields, and lasers are stabilized on various timescales using both optical-power measurements and atom-based signals.

Unique to this system is the dual-beamline optical design supporting laser-based local $R_Z(\phi)$ rotations and CZ gates. The former requires application of 459-nm light to a single atom and the latter requires simultaneous application of 459-nm and 1040-nm light to a pair of atoms to drive two-photon Rydberg excitation. To implement these gates, we send each laser through a sequence of three acousto-optic devices: one for pulse shaping and two for beam steering. The pulsed and steered beams are then focused to nearly diffraction-limited spots on the atom array (see Fig. 1). Using both direct optical measurements and atom-based beam profiling, we confirm that our 459-nm and 1040-nm beams are focused to spot sizes of 2.8 μm and 3.0 μm ($1/e^2$ intensity radii) in the atom plane. Critically, these spots are smaller than the atom spacing of 6 μm, allowing us to address all qubits individually. With this approach, we can alternate between CZ and local $R_Z(\phi)$ gates (on select target qubits) with a timescale limited only by the beam-steering devices. Our current acousto-optic deflectors (AODs) used as beam-steering devices can be re-aimed with a timescale of 2 μs. In future devices, driving AODs with multitone rf signals and/or combining AODs with spatial light modulators (SLMs) can enable fast individually addressed gates on multiple qubits simultaneously [39]. With this approach, our individual-addressing architecture supports parallel gate operations while maintaining high gate rates.

A practical consideration for this individual-addressing design is the crosstalk that can occur from the residual beam intensity reaching a neighboring atom. For instance, given our beam waists, one predicts 0.03% intensity reaching neighboring qubits. In practice, we find crosstalk

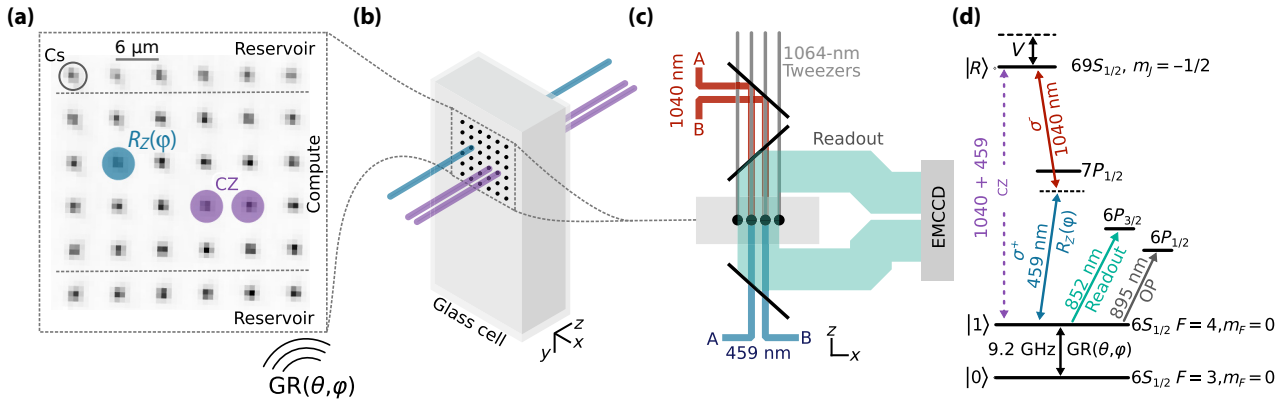


FIG. 1. A system overview. (a) An image of an array of Cs atoms, trapped in a glass cell. (b) Locally addressed $R_Z(\phi)$ (beam shown in blue) and CZ (beams show in purple) gates are applied via beams propagating perpendicular to the atom plane. Microwave-based global $R(\theta, \phi)$ gates (labeled GR) address the entire array. (c) The layout of the key lasers implementing trapping and gates. Rydberg excitation is implemented by counterpropagating 459-nm (blue) and 1040-nm (red) lasers. Rydberg illumination for the control and target qubits is sourced from separate pulse-shaping systems (labeled A and B). An optical-tweezer array (1064 nm) is generated via acousto-optic deflectors (AODs) and combined with the 1040-nm light. Readout fluorescence is collected and imaged on an electron-multiplying CCD (EMCCD). (d) The atomic levels and addressing fields, including lasers for two-photon Rydberg excitation, readout, and optical pumping (OP). Rydberg beams are circularly polarized (σ^+ and σ^- for 459 nm and 1040 nm, respectively). The Rydberg-blockade shift is denoted by V.

intensities up to approximately 10 times higher, likely due to optical aberrations. We provide a thorough discussion of the implications of this crosstalk in Sec. VI G and summarize the results here. One effect of this crosstalk is residual light reaching spectator atoms not participating in the gate. For our current parameters, we estimate that this produces an error contribution smaller than 10^{-4} per gate. Moreover, it is a predictable unitary error that can be compiled away [40]. The other important impact of crosstalk is on the entangling gate itself, where interference between the control and target illumination can significantly perturb the gate dynamics. However, we find that this error contribution can be made negligible by ensuring that the fields interfere at a frequency much higher than the Rabi frequency of the Rydberg excitation. We design our addressing system to satisfy this requirement. For each wavelength, we branch the single source laser into two fully independent beam-shaping and beam-steering paths (labeled A and B in Fig. 1). Thus, in a CZ gate, the Rydberg laser light reaching the control and target qubits has disparate acousto-optic frequency shifts, resulting in controllable crosstalk beat-note frequencies in the range 0–30 MHz (see Sec. VI G). This dual-beamline system enables our processor to tolerate greater crosstalk intensity at higher array density, achieving higher blockade strength, while still addressing qubits individually and minimizing the impact of crosstalk on neighboring atoms. Each qubit can be assigned uniquely to one of these beamlines only and this approach supports all nearest-neighbor connections in a square grid without incurring extra calibration overhead.

III. QUANTUM PROCESSOR IMPLEMENTATION AND CHARACTERIZATION

Optically trapped cesium atoms offer a ground-state hyperfine clock qubit with excellent coherence properties. We find a trap vacuum lifetime of 26(2) s, limited by background gas collisions. Raman sideband cooling (RSC) is used to produce an atom temperature of 2.6(3) μK, as indicated by both release-recapture measurements [41] and a Doppler-limited $T_2^* = 15(1)$ μs for the ground-Rydberg qubit. This temperature alone would permit a clock-qubit coherence time of about 42 ms [42]. In practice, we measure $T_2^* = 12.7(5)$ ms and we attribute this additional decoherence to magnetic field shot-to-shot fluctuation of approximately 1 mG at our bias of 11.15 G, as supported by measurements of $T_2^* < 1$ ms on magnetically sensitive hyperfine transitions. By concatenating repeated XY8 dynamical decoupling sequences, we can extend T_2 beyond the capability of our hardware to accurately measure it—our best estimate is 2.8 s, with a 97% confidence interval lower bound of 128 ms. The fidelity of initializing atoms into the qubit state $|1\rangle$ is limited by two error sources, namely, the infidelity of optical pumping into the stretched state during RSC and the infidelity of global microwave transfer pulses from the stretched state to the $|1\rangle$ state. We estimate the stretched-state optical-pumping fidelity to be 99.8(2)% based on the ratio of the mean optical-pumping time (53(6) μs) to the mean depumping time (23(18) ms) by the optical-pumping beam without repump. For microwave transfer, we use a Knill composite-pulse sequence for robustness against

detuning and Rabi frequency error. The total state preparation and measurement (SPAM) fidelity characterization is discussed in Secs. IV and VI H.

Our single-qubit gate set consists of global $R(\theta,\phi)$ (GR) and local $R_Z(\phi)$ rotations. The former are implemented via a uniform pulsed microwave field resonant with the magnetically insensitive qubit $|0\rangle \rightarrow |1\rangle$ transition. The duration and phase of the pulse determine θ and ϕ, respectively. We illustrate Rabi oscillations driven by this microwave field in Fig. 2(a). With this system, we have reached Rabi rates as high as $\Omega_{\mathrm{GR}} = 2\pi \times 120$ kHz, enabling an $R(\pi,\phi)$ gate time of 4.1 μs. We characterize the full set of GR Cliffords across a 24-qubit system using randomized benchmarking [Figs. 2(c) and 2(e)] and find an average fidelity of 99.980%, with an ensemble standard deviation of 0.003% and a typical fit uncertainty of 0.004%.

Local $R_Z(\phi)$ rotations are implemented via the differential light shift arising from a locally addressed 459-nm laser beam that off-resonantly addresses the transition from $|1\rangle$ to the $7P_{1/2}$ excited state with a detuning Δ from the $7P_{1/2}$ center of mass. In Fig. 2(b), we present Ramsey measurements of a differential light shift of $2\pi \times 2$ MHz, enabling an $RZ(\pi)$ gate time of 250 ns. We characterize $R_Z(\phi)$ fidelity for two values of Δ using interleaved randomized benchmarking (IRB), with rotation angles sampled uniformly over $[0, 2\pi)$ (see Sec. VI C). We first characterize the full 24-qubit array with $\Delta/2\pi = 885$ MHz (and, incidentally, a reduced $R_Z(\phi)$ rate of $2\pi \times 720$ kHz). This characterization uses a maximum IRB depth of 200 $R_Z(\phi)$ gates per qubit, which entails redirecting the beam-steering AODs 4800 times in a 15-ms circuit. We find an average $R_Z(\phi)$ fidelity of 99.76%, with an ensemble standard deviation of 0.03% [Fig. 2(e)]. To properly account for the difference between leakage and true depolarization, the fidelity calculation includes an estimate of the scattering error due to nonzero population of the $7P_{1/2}$ intermediate state (Sec. VI C). Our observed $R_Z(\phi)$ error is dominated by this scattering, which can be suppressed by using larger detuning. Indeed, changing $\Delta/2\pi$ to -2.1 GHz (and operating with an $R_Z(\phi)$ rate of $2\pi \times 2$ MHz), we measure, for a single qubit, a higher $R_Z(\phi)$ fidelity of 99.902(8)% [see Fig. 2(d)]. The residual error of 0.098(1)% is largely explained by the remaining leakage probability of 0.066% and laser-pulse energy fluctuation.

Our quantum processor implements the two-qubit controlled-Z (CZ) gate, which, combined with arbitrary single-qubit rotations, is sufficient for universal quantum computation. The gate is based on simultaneous excitation of the control and target qubits to a Rydberg state, where blockade interactions provide interqubit coupling. In particular, we use two-photon excitation (459 nm + 1040 nm) to couple the $|1\rangle$ state to the $|69S_{1/2}, m_j = -1/2\rangle$ Rydberg level. The single-photon intensities are defined by the

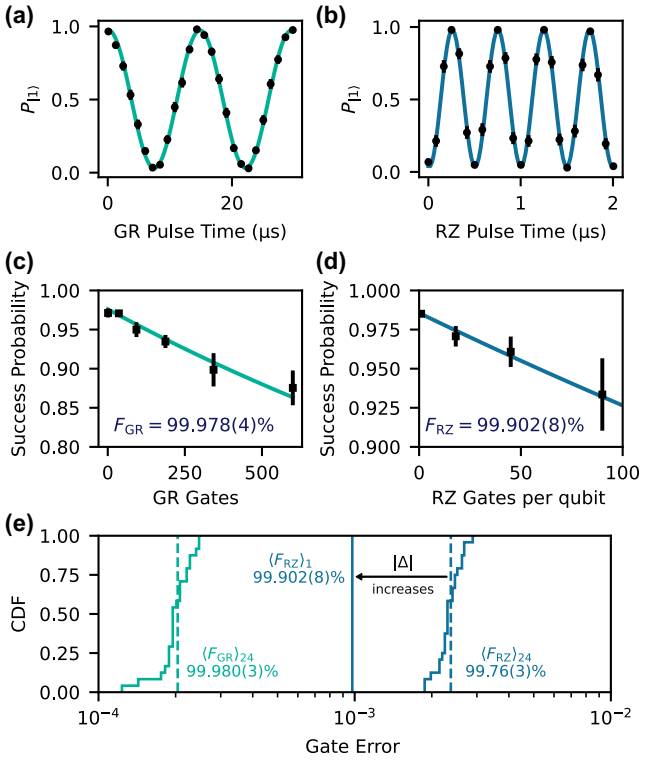


FIG. 2. Single-qubit gates. (a) Rabi oscillations driven by a global microwave field. (b) Ramsey oscillations from local qubit $R_Z(\phi)$ control. The error bars in (a) and (b) represent projection noise. (c),(d) Randomized benchmarking results for global $R(\theta,\phi)$ and local $R_Z(\phi)$ gates, respectively. Each point is an average over multiple random circuits [ten per depth for $R(\theta,\phi)$ and five per depth for $R_Z(\phi)$], with the error bar given by standard error of the mean. The solid lines are fits used to extract the fidelity. (e) Randomized benchmarking results, plotted as the cumulative distribution function (CDF) of 24-qubit (or one-qubit) results. For the 24-qubit results, gates are applied to each of the 24 qubits in each circuit. For the single-qubit results, the reported uncertainty is the benchmarking fit uncertainty. For the 24-qubit results, the reported uncertainty is the ensemble standard deviation and the dashed line indicates the ensemble mean. The fidelity improvement between the 24-qubit results and the one-qubit result comes primarily from an increase in the magnitude of the intermediate-state detuning, Δ.

$\Delta/2\pi = -2.1$ GHz detuning from the $7P_{1/2}$ intermediate state, the $2\pi \times 2$ MHz light shift applied to the qubit by the 459-nm laser, and the $\Omega = 2\pi \times 3$ MHz Rabi frequency of Rydberg excitation [Fig. 3(a)]. We measure a blockade interaction energy of $2\pi\hbar \times 10.3(1)$ MHz [Fig. 3(b)]. We implement an approximation of the time-optimal Rydberg CZ gate [27] using a pulse of duration $7.85/\Omega = 416$ ns, with a sinusoidal phase profile [implemented via electro-optic modulation of the 1040-nm beam; see the measured phase profile in Fig. 3(c)]. The full CZ gate consists of this phase-modulated Rydberg pulse followed by single-qubit phase corrections to produce the canonical gate. We first

calibrate the uncorrected gate, using a cost function that is insensitive to the single-qubit phases [25] [Fig. 3(d)], and then calibrate the phase corrections (see Sec. VI D). Note that this gate modulation is implemented upstream of the beam-steering AODs, so our individual-addressing architecture does not impose any extra calibration overhead for this protocol.

We characterize our CZ gates using randomized benchmarking circuits in which the CZ gates alternate with global single-qubit gates [see Fig. 3(e)]. Symmetric benchmarking is discussed in detail in Ref. [43] and approaches similar to ours have been used in Refs. [20,25,44]. This approach has limitations; e.g., swap errors go completely undetected and it strictly only characterizes the performance of the gate on a subset of two-qubit input states. Nevertheless, it is known to perform well given the error profile of neutral-atom platforms. To validate application of the approach to our gate, we perform density-matrix simulations of our benchmarking procedure using a realistic microscopic error model for the gate (see Sec. VI F) and find agreement between the fidelity extracted from benchmarking and the true underlying gate fidelity, reproducing findings from prior applications of the technique [25,44].

In Fig. 3, we present CZ benchmarking results for a single pair of atoms. In these circuits, single-qubit phase corrections are implemented via a combination of local $R_Z(\phi)$ and virtual phase shifts to the local oscillator, as in Ref. [45]. Because we characterize a single qubit pair, the beam-steering AODs are statically driven throughout the control sequence. By following each NDSSR measurement immediately with an occupancy measurement, we are able to separately characterize the two-atom survival probability ($R = 99.62(3)\%$ per gate) and the cycle polarization ($P = 99.67(4)\%$). Combining these with a model-based estimate of the leakage probability (see Sec. VI E), we extract a CZ fidelity of 99.35(4)%. The leading physical mechanisms contributing to the observed 0.65% CZ-gate error are leakage due to the finite lifetime of the $7P_{1/2}$ intermediate state, leakage and loss due to the finite lifetime of the Rydberg level, and shot-to-shot intensity fluctuations of the Rydberg laser beams. We estimate the contribution to the CZ error from optical crosstalk to be $\lesssim$ 0.1% in our current configuration, with a straightforward path to further reduction (see Sec. VI G).

NDSSR allows us to filter the data based on atom retention and calculate a postselected CZ fidelity of 99.73(3)% (see Sec. VI E). This approximately 0.38% atom-loss error arises from population that remains in a Rydberg state at the end of the gate, as indicated by separate measurements that show that loss from the temporary extinction of the tweezers during the gate is negligible. Atom-loss error may be converted to erasure through a leakage-detection measurement in which an ancilla becomes entangled with the presence or absence of an atom in a qubit state (rather than its logical state) [46,47], enabling the use of erasure-specific techniques to achieve improved correction performance [48].

IV. NONDESTRUCTIVE STATE-SELECTIVE READOUT

While neutral-atom platforms have demonstrated significant improvements in gate fidelity and qubit-count scaling that make neutral atoms a leading modality for quantum computing, there has been relatively less emphasis on their shot rate. Noncircuit operations such as atom loading and arrangement are expected to be a non-negligible part of the total operation time for measurement schemes based on destructive readout; a nondestructive-readout scheme with repetitive measurements can significantly increase the shot rate by amortizing the cost of these noncircuit operations and can enable midcircuit measurements for fault-tolerant operation. We implement the nondestructive state-detection method demonstrated in Ref. [49] with a few critical improvements (see Sec. VI H). During the state-detection sequence, the qubit state $|1\rangle$ is optically pumped to the stretched state $|F = 4, m_F = 4\rangle$, where we can detect fluorescence from the closed $|F = 4, m_F = 4\rangle \rightarrow |F' = 5, m'_F = 5\rangle$ transition, while $|0\rangle$ remains dark. By strategically alternating between illumination with trap light and readout light, we mitigate the deleterious effects of the excited-state light shift and the dipole force fluctuations, and we reach a state-averaged loss rate of 0.7(3)% (0.8(4)% for the bright state and 0.7(3)% for the dark state) while acquiring sufficient fluorescence from the atoms with 6 ms of camera exposure time. We estimate a base-discrimination fidelity (i.e., our raw ability to discriminate bright- or dark-count distributions, ignoring any readout-induced depumping; this is the "discrimination fidelity" of Ref. [50]) of 99.96(2)% (see Sec. VI H 2)

To quantify the combined performance of our state preparation and state-selective imaging in a single metric, we define the SPAM fidelity as

$$1 - \frac{P(D|\,|1\rangle_{\text{prep}}) + P(B|\,|0\rangle_{\text{prep}})}{2}, \tag{1}$$

where $P(D|\,|1\rangle_{\text{prep}})$ is the probability of detecting a dark signal after attempting to prepare $|1\rangle$ and, similarly, $P(B|\,|0\rangle_{\text{prep}})$ is the probability of detecting a bright state after attempting to prepare $|0\rangle$. We estimate $P(B|\,|0\rangle_{\text{prep}})$ from the results of a circuit with a single $R(\theta = \pi, \phi = 0)$ gate and $P(D|\,|1\rangle_{\text{prep}})$ from the results of a circuit without any gate; the two circuits comprise an iteration of our SPAM characterization experiment. In Fig. 4, we present SPAM characterization data aggregated from 43 repetitions of the experiment performed across a 5×5 array characterizing 25 qubits. The SPAM fidelity after postselecting for atom loss is 98.2(5)%, limited by

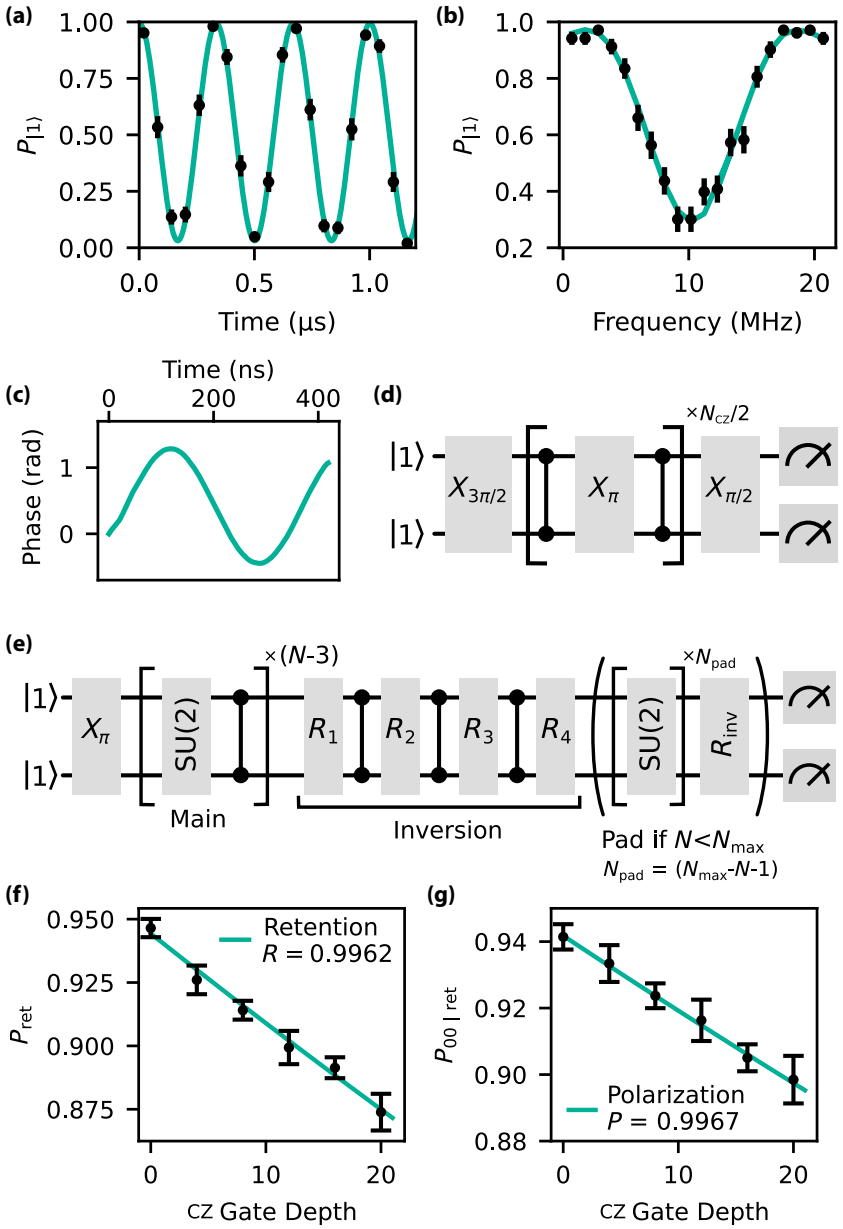


FIG. 3. Characterization of the CZ gates. (a) Rabi oscillations between $|1\rangle$ and $|R\rangle$, at a rate of $2\pi \times 3$ MHz. (b) Rydberg spectroscopy, measured while a neighboring atom is already in the Rydberg state. The decreased qubit $|1\rangle$ population indicates simultaneous excitation to the Rydberg state, which is only possible when the excitation is detuned (x axis) by the Rydberg blockade. The error bars in (a) and (b) represent projection noise. (c) The phase profile of the 1040-nm optical pulse used in the CZ gate, measured using optical heterodyne interferometry. (d) The circuit used as the cost function for CZ-gate optimization. Global $R(\theta, \phi = 0)$ gates are denoted as X_θ. (e) The circuits used for characterization of CZ gates include a main block of $(N-3)$ pairs of alternating Haar-random global SU(2) gates and CZ gates, as well as an inversion operation (consisting of three CZ gates and four global rotation gates R_i) and a padding circuit that reduces the spurious impact of $R(\theta, \phi)$ errors on the extracted CZ fidelity. For circuits with $N = 0$, the "Main" and "Inversion" blocks are omitted entirely. (f) The observed probability of retaining both atoms as a function of the number of CZ gates (N) in the circuit, along with the fit used to extract the retention probability. The decay is fitted to an asymptote of zero. (g) The observed probability of obtaining the expected state-selective measurement outcome, given that both atoms have been retained, as a function of the number of CZ gates in the circuit, along with the extracted cycle polarization. The decay is fitted to an asymptote of 25%. The data in (f) and (g) are obtained from sequential (g) Nondestructive state-selective readout (NDSSR) and (h) occupancy-readout measurements on the same quantum state. Each point is obtained as an average over between 18 and 48 circuits (183 circuits in the data set in total), with an error bar corresponding to the standard error of the mean. Combining the fitted depolarization and retention per cycle yields a fidelity of 99.35(4)%.

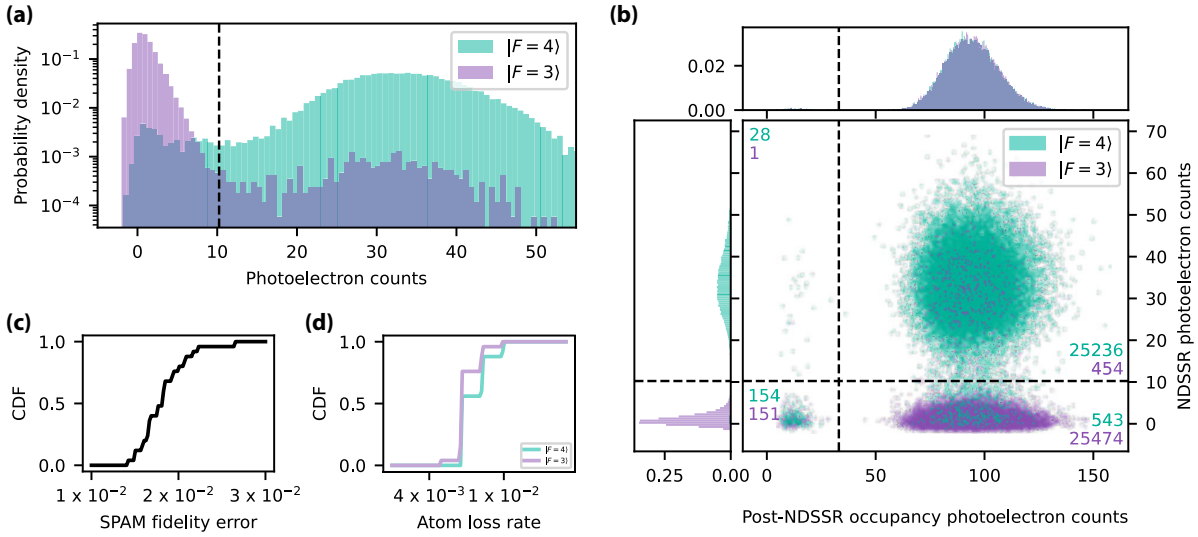


FIG. 4. Nondestructive state-selective readout. The NDSSR fluorescence signal distribution and SPAM statistics from 43 repetitions of SPAM characterization experiment performed on a 25-qubit array. Each experiment executes two circuits (see the text) with approximately 600 shots each, resulting in approximately 25 800 camera frames for each circuit. The turquoise-colored (purple-colored) distribution describes the signals where atoms are expected to be in $|F=4\rangle$, which is bright ($|F=3\rangle$, which is dark). Photoelectron counts for each site have been calculated from the raw EMCCD camera analog-digital units using nominal EM gain and preamplifier gain. (a) The sampled probability distribution of the photoelectron counts for bright and dark states for a representative site in the array. The dashed black line denotes the discrimination threshold. (b) The NDSSR and post-NDSSR occupancy-readout photoelectron counts for the representative site. A dark post-NDSSR occupancy signal indicates atom loss. Each quadrant of the plane defined by NDSSR and post-NDSSR photoelectron counts contains two text labels showing total counts from the $|F=4\rangle$ SPAM circuit (turquoise) and the $|F=3\rangle$ SPAM circuit (purple), respectively. (c) The cumulative distribution function (CDF) of the postselected SPAM fidelity without atom loss across 25 qubits. (d) The CDF of the atom loss rate for each initial state across the array.

$P(D|\,|1\rangle_{\text{prep}})$ of 2.1(7)% and $P(B|\,|0\rangle_{\text{prep}})$ of 1.6(6)%. Values are reported as the mean (standard deviation) calculated from measurements of individual qubits. We note that state-preparation error, which includes failure to optically pump into $|F=4,m_F=4\rangle$ during RSC and failure to coherently transfer $|F=4,m_F=4\rangle$ to $|F=4,m_F=0\rangle$ using microwave pulses (see Sec. VI B), is distributed across both types of state-detection errors. The shot-rate enhancement from repetitive nondestructive measurements depends on the quantum circuit being run, especially on the number of CZ gates, as they can induce small atom loss. For an example circuit, we have chosen the preparation of a 15-qubit Greenberger-Horne-Zeilinger (GHZ) within a 24-qubit array, which contains 14 CZ gates. We have recorded an average of an 8.2(2)-Hz shot rate for the example circuit compared to an estimated single-measurement shot rate of approximately 3 Hz (see Sec. VI H 3). These experiments have been performed before the integration of RSC, which adds additional time to each measurement that cannot be amortized.

The nondestructive-readout method for alkali atomic species demonstrated in this work provides the starting point for repetitive midcircuit readout of ancilla qubits in dual-alkali-species-based logical qubit architectures. The atom loss rate can be further reduced by using shorter exposure times and advanced camera-image binarization algorithms that work at low photon counts [51]. The depumping error during imaging can be improved with technical improvements in polarization purity and with microwave-assisted state-selective repumping during imaging [52].

V. CONCLUSIONS AND OUTLOOK

This work advances the state of the art for individually addressed entangling gates in a neutral-atom platform and we foresee a clear path toward the error levels required for fault tolerance. Control errors may be reduced by reaching lower atom temperatures, better stabilizing laser power and frequency, and using robust gate protocols [53,54]. Straightforward changes such as increasing the intermediate-state detuning will immediately reduce scattering error. Moving forward, scattering error can be eliminated by transitioning to single-photon Rydberg excitation [44,45], which also reduces the complexity of light shifts on atomic resonances. Parallel gate operation using multiple beams is already possible through multitone AOD drives and/or SLMs [38] and novel technologies are under

development, e.g., photonic integrated circuits [55]. The path for scaling our qubit array is also well established, with demonstrations [18] suggesting the feasibility of trapping $> 10^4$ qubits in a single array, with the option to improve trapping lifetimes by operating at cryogenic temperatures [56]. The nondestructive readout that we present here already enables state-of-the-art shot rate and paths toward even faster readout are being developed [57,58]. Finally, we note that our ability to achieve high ground-Rydberg coherence ($T_2^* = 15$ μs) at $n = 69$ is a unique advantage in the field, enabled by our electric field control. Since the Rydberg interaction scales with n^{11}, operating at higher n unlocks greater connectivity, which is valuable for QEC protocols and applications [32].

For quantum computers to provide meaningful scientific and commercial value, they will likely need to support fault-tolerant operation on at least approximately 100 logical qubits [10,59,60] and neutral-atom architectures are evolving rapidly toward that goal. This work demonstrates important new capabilities on that path, namely, the ability to implement individually addressed high-fidelity entangling gates and NDSSR, which collectively enable faster qubit operations and a broader set of QEC protocols. Ultimately, a compelling architecture for fast fault-tolerant quantum computing is likely to integrate these techniques alongside other demonstrated neutral-atom capabilities, including dual-species platforms [61,62], midcircuit measurement [11,52,61–64], and alternative addressing architectures [11,19,65,66]. In particular, future work should investigate the optimal combination of individual addressing, global beams, and midcircuit rearrangement techniques to minimize circuit run times. For example, individual addressing can enable fast parallel short-range CZ gates, while motion unlocks long-range connectivity. And while a shared global beam can support parallel identical CZ gates, individual addressing is preferable in the case of $R_Z(\phi)$, to support parallel gates with distinct rotation angles. Additionally, individual addressing with tighter beams provides greater flexibility in achieving the desired laser intensity at significantly lower laser power, which is an important consideration for scaling and achieving higher-fidelity gates. The advanced optical control techniques that we demonstrate here also extend directly to alternative atomic species [11–13,19,44,67], as well as multiqubit ($N > 2$) gates [68], which can further optimize logical operations [69].

VI. METHODS

A. QPU details

At the core of our platform, the qubit array is housed in a custom Infleqtion-manufactured dual-chamber ultra-high-vacuum glass cell, which lives at the center of the quantum processing unit (QPU, Fig. 5). From the lower chamber, a cold atomic beam is generated with a two-dimensional (2D) magneto-optical-trap (MOT), which loads a three-dimensional (3D) MOT in the upper chamber. Surrounding the upper chamber are components to stabilize the electric and magnetic field environment, as well as a 405-nm light source for inducing atom desorption from the cell walls and a microwave horn for implementing GR gates (Sec. III). Large beams address the whole atom array for implementing the 3D MOT, polarization-gradient cooling, Raman sideband cooling [70], readout, and optical pumping. A 2D array of 1064-nm traps is generated by a pair of acousto-optic deflectors (AODs). Rydberg-excitation light is provided by 1040-nm and 459-nm lasers, both locked to a stable high-finesse reference cavity. The two-photon detuning from the atomic transition is actively stabilized using a ground-Rydberg Ramsey signal. For each Rydberg wavelength (1040 nm and 459 nm), separate pulse shapers (one each for control or target qubit) feed separate beam-steering modules. The pulse shapers consist of double-pass acousto-optic modulators (AOMs), optimized for fast rise times. The 1040-nm beamline additionally contains an electro-optic modulator (EOM) for implementing the phase modulation used for CZ gates. The beam-steering modules each consist of a pair of AODs with orthogonal axes, for rastering to any location within the array. For each Rydberg wavelength, the control and target beams are combined on a nonpolarizing beam splitter. The 1040-nm light is combined with 1064-nm light and focused onto the atoms from one side of the cell, via high-numerical-aperture (NA = 0.7) objective lenses, while the 459-nm light is focused from the opposing side, using a matching high-NA lens. These lenses also collect atom fluorescence during readout and direct it to an electron-multiplying CCD camera (EMCCD). Atom-based beam profiling of the 1064-nm tweezer has measured a Gaussian-beam-waist radius at $1/e^2$ of peak intensity of 1.4(2) μm; beam-waist estimates of 2.8 and 3 μm for 459-nm and 1040-nm beams, respectively, are discussed in Sec. II. The trap depth is calibrated by measuring the intensity dependence of the qubit microwave transition; for the benchmarking data presented in this work, the estimated trap depth is 1 mK. The QPU is enclosed for electromagnetic, thermal, and vibrational isolation and stabilization.

B. Cooling and state preparation

Here, we give a brief description of how we cool atoms inside the optical trap and initialize them in the $|0\rangle = |F = 3, m_F = 0\rangle$ state. Once the atoms are loaded into the optical trap, we turn on a pair of counterpropagating beams on all three axes to provide polarization-gradient cooling (PGC). The same beams are used for occupancy readout but at smaller detuning and different intensity balance. After PGC, the magnetic bias field is stepped from zero to approximately 6 G along the x axis and a

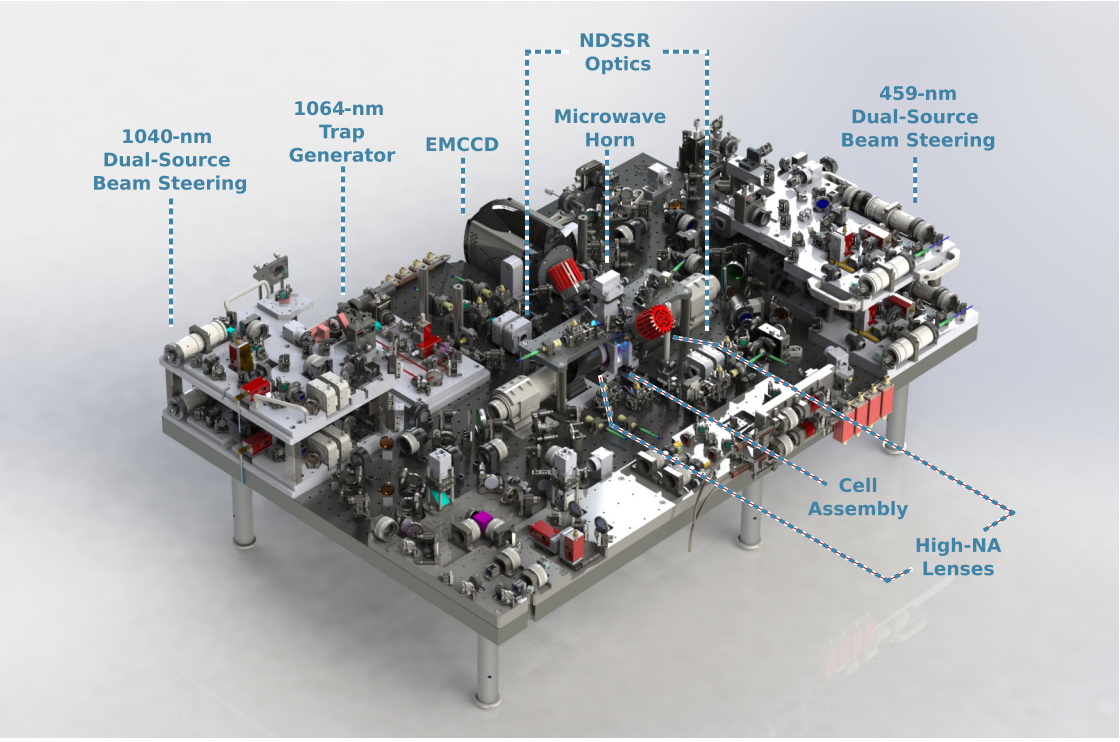


FIG. 5. The three-dimensional (3D) model of the quantum processing unit (QPU), highlighting the dual-source beam-steering modules for both Rydberg wavelengths, a microwave horn for global $R(\theta,\phi)$ gates, AODs for generating a two-dimensional (2D) array of 1064-nm optical tweezers, the electron-multiplying CCD (EMCCD) for imaging, the optics supporting NDSSR, the ultra-high-vacuum glass cell, and the high-NA lenses for focusing addressing beams and collecting readout light.

circularly polarized 895-nm laser optically pumps atoms to the stretched state $|F=4,m_F=4\rangle$. After 10 ms of optical pumping, we typically measure a temperature of approximately 4–5 μK using the release-and-recapture method. This is followed by an RSC step [70], in which optical-pumping pulses and Raman pulses alternate. Using sideband thermometry, we measure mean trap phonon numbers of $\bar{n}_{\mathrm{ax}}=6.1(65)$ for an axial trap frequency of 11.5 kHz and $\bar{n}_{\mathrm{rad}}=0.49(22)$ for a radial trap frequency of 60.9 kHz. We believe that these mean phonon numbers are currently limited by nonlinear tone modulation in the AODs producing our trap light, as changing the tone separation significantly changes the achievable temperature. Future devices that use spatial light modulators (SLMs) should be able to obtain lower temperatures. At the end of RSC, a sequence of composite microwave pulses transfers atoms from the stretched state to qubit state $|1\rangle$. After the microwave-pulse sequence is complete, the magnetic bias field is rotated from 6 G along the x axis to the final bias field of 11.15 G along the z axis. Once the rotation is complete, a microwave π pulse transfers the atoms from $|1\rangle$ to $|0\rangle$ and this marks the beginning of a quantum circuit.

The quantum circuit shot-rate data discussed in Sec. VI H 3 use a different optical-pumping method, in which a linearly polarized 895-nm beam pumps atoms toward $|1\rangle=|F=4,m_F=0\rangle$ at 11 G.

C. Benchmarking of $R_Z(\phi)$ gates

We characterize our $R_Z(\phi)$ gates using interleaved randomized benchmarking (IRB), randomly sampling each gate over $\phi\in[0,2\pi)$. Given the implementation of our $R_Z(\phi)$ gate, we expect a significant scattering error contribution due to off-resonant excitation to the $7P_{1/2}$ state. This can produce leakage error that populates ground-state Zeeman sublevels outside the qubit manifold. In these states, the fidelity with the target state is $F_{\mathrm{leak}}=0$, in contrast with true depolarizing error with fidelity $F_\sigma=1/d$, where $d=2$ is the dimension of the Hilbert space for a single-qubit gate. Denoting this leakage probability as P_{leak} and the depolarizing probability (conditioned on not experiencing a leakage error) as σ, we can express the probability per gate of implementing the desired unitary as $C=(1-P_{\mathrm{leak}})(1-\sigma)$. With this definition, the data are then fitted to $\frac{1}{2}+\frac{\mathrm{SPAM}}{2}C^N$, where SPAM is a fitting parameter that accounts for reduced contrast arising from SPAM error.

We can express the fidelity as F_{RZ}:

$$F_{\mathrm{RZ}}=P_{\mathrm{leak}}F_{\mathrm{leak}}+(1-P_{\mathrm{leak}})\sigma F_\sigma+(1-P_{\mathrm{leak}})(1-\sigma). \tag{2}$$

With $F_{\text{leak}} = 0$ and $d = 2$ as defined above, this yields

$$F_{\text{RZ}} = (1 - P_{\text{leak}})(1 - \sigma/2). \tag{3}$$

The fidelity extracted from a given cycle polarization C will thus depend on the ratio of depolarization to leakage. Since scattering error is well understood, we choose to model its contribution and use this assumption in calculating the fidelity. We can predict the scattering error based only on the intermediate-state lifetime and detuning (Δ). Table I gives the relevant parameter values for the data sets presented in Fig. 2. We attribute the improved depolarizing error at $\Delta/2\pi = -2.1$ GHz to improved laser intensity stability.

We expect to improve the fidelity of $R_Z(\phi)$ operations with larger intermediate-state detuning. For detunings significantly greater than the ground-state hyperfine splitting, the differential light shift and photon scattering rate both scale as $1/\Delta^2$, resulting in a fundamental upper bound on the figure of merit—the $R_Z(\phi)$ rotation angle achieved per scattered photon—of approximately $2\pi \times 700$ radians. By comparison, at $\Delta/2\pi = -2.1$ GHz, this figure of merit is approximately $2\pi \times 250$ radians. It is worth noting that this upper bound can also be approached at $\Delta/2\pi = +4.6$ GHz, which is approximately half the ground-state hyperfine splitting. However, at this detuning, both the differential light shift and scattering rate are relatively large. This increased differential light shift is undesirable if the same laser is used for $R_Z(\phi)$ and CZ operations, as it can exacerbate qubit-pair decoherence after a CZ gate. Therefore, to minimize such coherence loss, we aim to operate at a larger red detuning to improve $R_Z(\phi)$ fidelity.

D. Calibration of CZ gates

The CZ-gate calibration proceeds in two phases. The first phase establishes AOM pulse parameters for resonant Rydberg excitation with the desired two-photon Rabi rate Ω and component single-photon Rabi rates Ω_{459} and Ω_{1040}. Operationally, these single-photon rates are established by first setting the 459-nm laser power to give the desired differential light shift on the qubit and then varying the 1040-nm laser power to achieve the desired two-photon Rabi frequency. This is accomplished by iterated time- or frequency-domain characterizations of Rydberg Rabi excitation. This calibration is executed sequentially on all qubits in the array in a single circuit, allowing parallelized calibration of all qubits.

In the second phase, we calibrate the CZ-gate pulse sequence, which consists of two components: a Rydberg pulse with sinusoidally modulated phase, followed by single-qubit phase corrections required to produce the canonical CZ gate. The uncorrected gate is calibrated first, by maximizing a cost function [Fig. 3(d)] that is insensitive to the single-qubit phases. The circuit begins and ends with $R(\theta = \pi/2, \phi = 0) \equiv RX(\theta = \pi/2)$ gates, surrounding N pairs of CZ gates, with an $RX(\pi)$ gate interleaved within each pair. These interleaved gates provide the desired insensitivity to the single-qubit phases. A properly tuned gate results in all population returning to $|11\rangle$ and detecting as bright-bright, while any dark signal indicates either a state $\neq |11\rangle$ or atom loss, both of which correspond to a gate imperfection.

Using this cost function, we optimize a representation of the gate using six parameters: the amplitude (A_0), frequency (Ω_{mod}), and phase offset (ϕ_0) of the 1040-nm phase modulation, as well as the overall pulse duration and two-photon detuning. To compensate for calibration errors, we include separate two-photon detuning corrections for each qubit, though the protocol itself should be symmetric. To optimize these six parameters, we utilize a Bayesian optimizer (built on the BoTorch library) with a Gaussian-process (GP) surrogate model. This allows us to optimize without making any assumptions about the underlying parameter landscape. Indeed, simulations suggest that our phase-modulation waveform is somewhat over-parametrized, leading to the possibility of multiple local optima with similar performance. The optimizer run that preceded the benchmarking data from Fig. 3 is shown in Fig. 6. In Fig. 6(b), we see the overall trajectory of the optimizer, over the course of 85 cost-function evaluations, divided into an initial random-sampling phase, a main-exploration phase, and a final-polish phase. In Fig. 6(a), we show the final GP model, evaluated at the sampled parameter values, projected onto each parameter axis. The same data are shown in Fig. 6(c), projected onto 2D slices of the space. This latter view can be useful for identifying correlations between parameters and finding improved parametrizations. Once an optimized waveform is identified [Fig. 3(c)], the single-qubit phases are measured via a qubit Ramsey sequence and appropriate corrections are determined.

E. Benchmarking of CZ gates

We compute the CZ fidelity from the results of our benchmarking circuits by measuring the probabilities of

TABLE I. Calculated and fitted parameters related to the fits and fidelities presented in Fig. 2.

$\Delta/2\pi$	Fitted C	Modeled P_{leak}	Extracted $\sigma/2$	$1 - F_{\text{RZ}}$
+885 MHz	99.65(5)%	0.125%	0.11(3)%	0.24(3)%
−2.1 GHz	99.87(2)%	0.066%	0.03(1)%	0.098(8)%

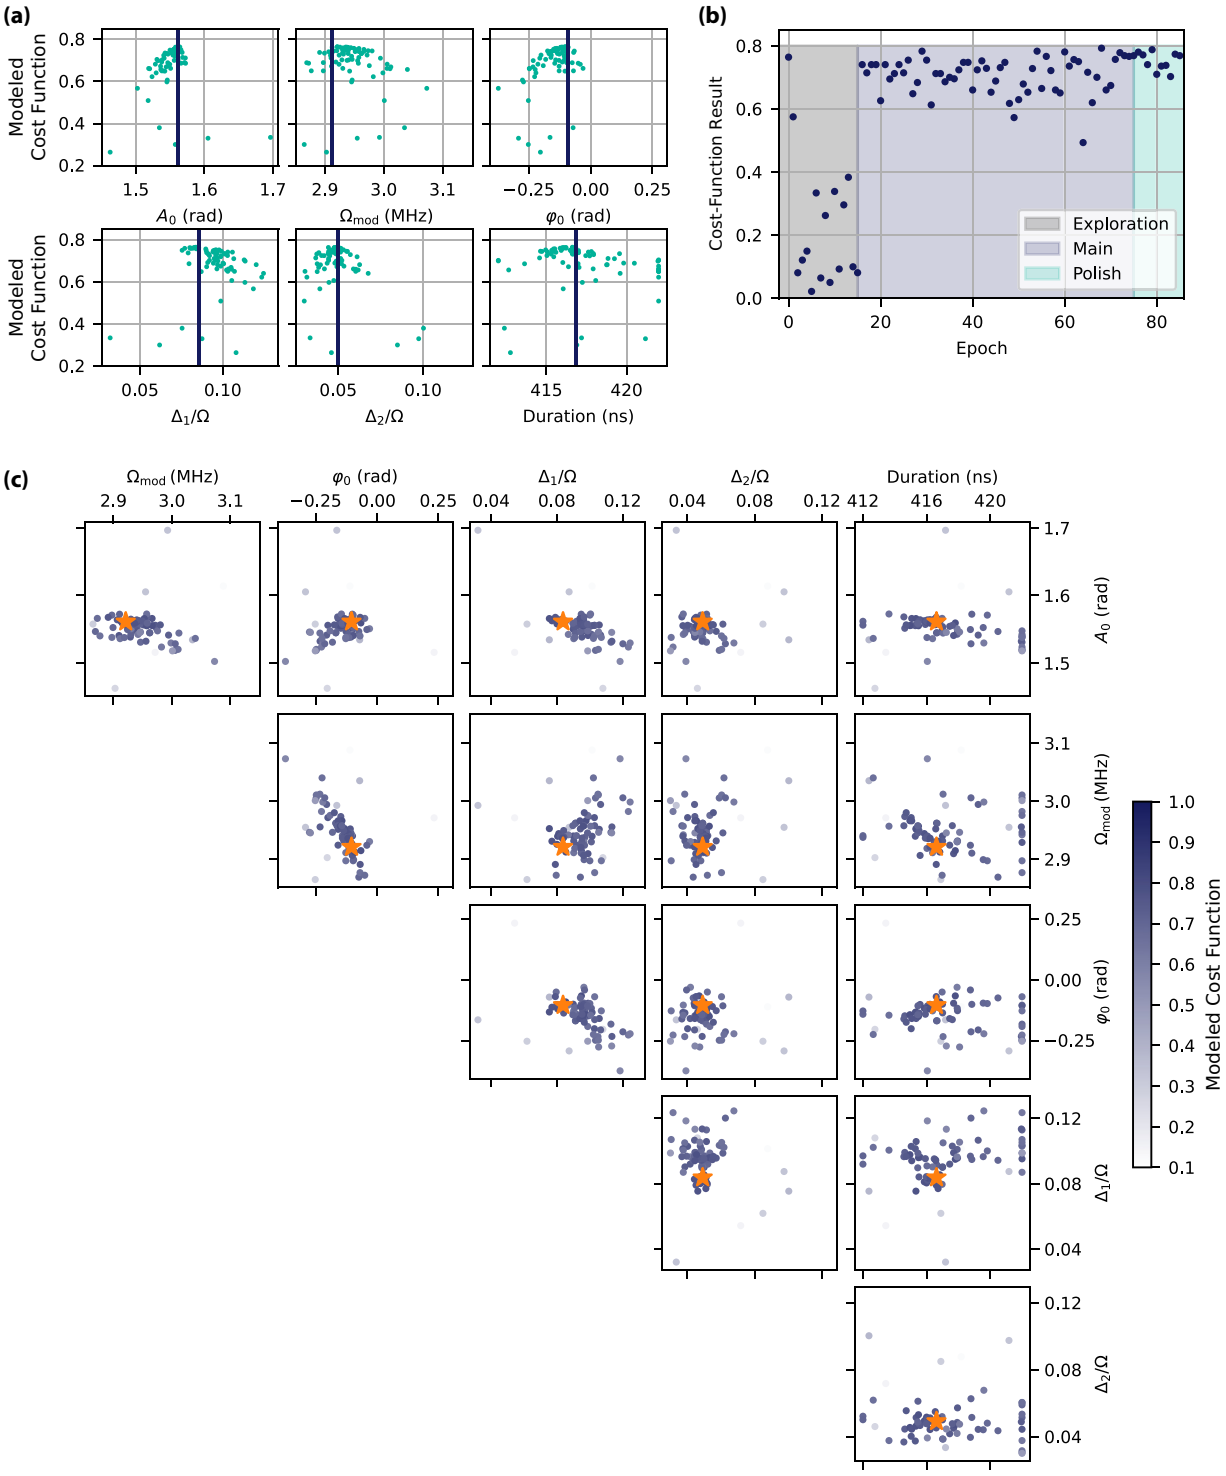


FIG. 6. Calibration of CZ gates. (a) One-dimensional projections of the final Gaussian-process (GP) model evaluated at the parameter values sampled by the optimizer. (b) The trajectory of the measured cost function, across three different phases of the optimizer. (c) 2D projections of the final model, evaluated at the parameter values sampled by the optimizer.

the circuits to produce the expected outcomes, which in this case are uniformly $|00\rangle$. In order to account for the possibility for each CZ gate to produce atom loss, leakage, or errors within the qubit subspace, we make use of the information afforded by our nondestructive state-selective readout capability (NDSSR; see Sec. IV). By sequentially performing NDSSR and occupancy measurements, we can determine whether a given atom was retained, and

if retained, whether it was in the $F = 4$ (bright signal) or $F = 3$ (dark signal) ground-state manifold. We therefore separately analyze, to a gate depth of 20, the probability of retaining both atoms (P_{ret}) and the probability of observing a dark-dark measurement outcome in the state-selective measurement, given that both atoms have been retained ($P_{00|\text{ret}}$) (see Fig. 3). We fit these data to decaying exponentials with asymptotes of zero and one quarter, respectively, and find a two-atom survival probability of $R = 99.62(3)\%$ per gate and a cycle polarization of $P = 99.67(4)\%$. Various effects may change the true values of these asymptotes but we expect the systematics introduced to be small in general. For example, we predict from our state-measurement errors that we would measure the dark-dark outcome 25.25% of the time for the asymptotic fully depolarized state; reanalysis of the data with any asymptote in the range [23.9%, 25.8%] yields the same result for the fidelity to four digits. The observed cycle depolarization has contributions from true depolarizing error and from leakage to states $|6S_{1/2}, m_F \neq 0\rangle$. The former asymptotically produces the fully mixed state (with 25% fidelity with the target state) and the latter produces a state that has zero overlap with the desired state. A detailed model for the dynamics of the intermediate-state and Rydberg populations during the gate provides an estimated leakage probability of $P_{\text{leak}} = 0.10\%$ per gate. Because the CZ gate can produce loss as well as leakage, we introduce Q, the state-averaged probability for the gate to not produce leakage or loss, and R_i and q_i, state-dependent retention probabilities and probabilities to remain within the qubit space, conditional on retention. Here, i indexes over input states to the gate: $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. The expression analogous to Eq. (3), with $d = 4$, is

$$F = Q \times \left(1 - \frac{3}{4}\sigma\right), \tag{4}$$

$$Q = \frac{1}{n}\sum_{i}^{n} Q_i \tag{5}$$

$$= \frac{1}{n}\sum_{i}^{n} R_i q_i. \tag{6}$$

We approximate the average of the product $R_i q_i$ with the product of the averages $Rq = R \times (1 - P_{\text{leak}})$. Whether this approximation produces a value that is above or below the true fidelity depends on the covariance between the sets R_i and q_i. We expect positive covariance between these outcomes because both are more likely for $|0\rangle$, so this approximation likely provides a small under-estimate of the true fidelity. We use this approximation, together with the calculated $P_{\text{leak}} = 0.10\%$ and the measured R, to calculate the CZ fidelity as $F = R \times (1 - P_{\text{leak}}) \times (1 - \frac{3}{4}\sigma)$, where $\sigma = (1 - P_{\text{leak}} - P)/(1 - P_{\text{leak}})$ is the calculated probability of depolarizing error; we find a CZ fidelity of 99.35(4)%. Using the same leakage rate P_{leak}, we can calculate a CZ fidelity postselected for atom loss of $F = (1 - P_{\text{leak}}) \times (1 - \frac{3}{4}\sigma) = 99.73(3)\%$.

The circuits that we use to benchmark CZ gates consist of four steps: (a) preparation of the $|00\rangle$ qubit state; (b) application of a variable number of gate pairs consisting of the same Haar-random SU(2) gate applied to both qubits, followed by a CZ gate; (c) an inversion operation consisting of alternated application of four global $R(\theta, \phi)$ rotations and three CZ gates; and (d) another layer of N_{pad} Haar-random global SU(2) gates, followed by a final inversion operation. Steps (b)–(d) in each circuit are constructed so that, in the absence of error, the circuit performs the identity operation on the input $|00\rangle$ state. The final layer of random SU(2) gates is included to hold the total number of single-qubit gates constant at a value N_{max} as the depth of layer (b) is varied. The data in Fig. 3 are a composite of two data sets, acquired concurrently, with $N_{\text{max}} = 20$ and 24; we verify that this composition does not introduce significant systematic error in the extracted fidelity by simulating the full benchmarking procedure (see Sec. VI F). We calculate the inversion operation in step (c) as an alternating sequence of four CZ gates and four SU(2) gates, and then we omit the final CZ gate from the circuit because it does not change measurement outcomes. The structure of the benchmarking circuits is depicted in Fig. 3(e). The global SU(2) operations are implemented using a single global $R(\theta, \phi)$ gate and an update to the qubit phase frame, which allows us to use only the global microwave field for the single-qubit gates in the CZ benchmarking circuits.

F. Simulation of CZ benchmarking

We conduct circuit-level simulations of our benchmarking procedure using a model for the CZ-gate process. The error model includes transitions from the $|1\rangle$ state to a loss state $|L\rangle$ (no analogous term for the $|0\rangle$ state is included, because only $|1\rangle$ is near-resonantly coupled to the Rydberg state), as well as transitions between $|0\rangle$, $|1\rangle$, $|0_L\rangle$ and $|1_L\rangle$; the latter two states are leakage states that represent the $|F = 3, m_F \neq 0\rangle$ and $|F = 4, m_F \neq 0\rangle$ ground-state Zeeman sublevels and that are detected as $|0\rangle$ and $|1\rangle$, respectively. This model approximates spontaneous transition dynamics between the ground-state Zeeman sublevels, where a more complete model would separately treat all 16 ground states.

We calculate spontaneous transition probabilities, accounting for (a) interference between coupling of population from $|1\rangle$ and from the Rydberg level $|R\rangle$ to the intermediate state that produces "dark-intermediate-state" ground-Rydberg coupling [24], (b) interference in coupling paths through $|7P_{1/2}, F = 4, m_F = 1\rangle$ and $|7P_{1/2}, F = 3, m_F = 1\rangle$, (c) the full decay chain from higher-lying $7P_{1/2}$ decay products $5D_{3/2}$ and $7S_{1/2}$, and

(d) Rydberg decay-to-ground events. The result is a matrix describing the probability of a spontaneous transition per gate, expressed in the $\{|0\rangle, |1\rangle, |0_L\rangle, |1_L\rangle\}$ ordered basis:

$$\begin{pmatrix} 17.4 & 185.0 & 4.9 & 165.4 \\ 18.5 & 197.5 & 4.6 & 177.7 \\ 31.1 & 420.3 & 45.8 & 1210.0 \\ 42.1 & 590.2 & 52.9 & 1901.0 \end{pmatrix} \times 10^{-6}. \quad (7)$$

Our error model includes these transition probabilities, as well as a probability per gate to transition from $|1\rangle$ to $|L\rangle$ of 0.39% and an additional single-qubit Z error probability of 0.065%. These latter two probabilities are chosen such that a simulation of the circuits used to obtain the benchmarking data shown in Fig. 3 exactly reproduces the observed retention and polarization probabilities per gate. The results of this simulation are shown in Fig. 7. The choice of Z as the only Pauli error is well motivated by the underlying physics, as the dominant source of bit-flip errors is already embodied in the scattering matrix of Eq. (7).

By employing both first-principles and phenomenological inputs, this model provides valuable validation of the benchmarking routine and its analysis. For example, the true asymptote of the probability to observe the dark-dark outcome in the state-selective measurement is certainly different from the 25% that is used for the fit to the data. This difference arises from the depumping from the $F = 4$ hyperfine level driven by the 459-nm laser. Nevertheless, we calculate gate fidelities averaged over Haar-random input states and over 12 symmetric stabilizer states that form a 2-design on the two-qubit symmetric subspace [45] to be 99.38% and 99.37%, respectively, for the error model described above, which reproduces the experimental observations. These fidelities are greater than and consistent with the 99.35(4)% fidelity extracted from the benchmarking data for both the experimental and simulated data sets. These comparisons suggest that our benchmarking routine adequately characterizes the gate for our gate physics.

To further validate our benchmarking approach, we compare the gate fidelity averaged over Haar-random input states to the fidelity extracted from simulated benchmarking for a variety of error models. These models are obtained from the one thus far described by independently scaling the leakage matrix in Eq. (7), the $|1\rangle \to |L\rangle$ probability, and the single-qubit dephasing probability. Benchmarking runs are simulated with the same number of circuits at each depth and with the same $N_{\max}$ as used to obtain the data shown in Fig. 3, and the simulated analysis routine includes projection noise. Across these runs, the extracted error ε_E is larger than the true error ε_T in 86%

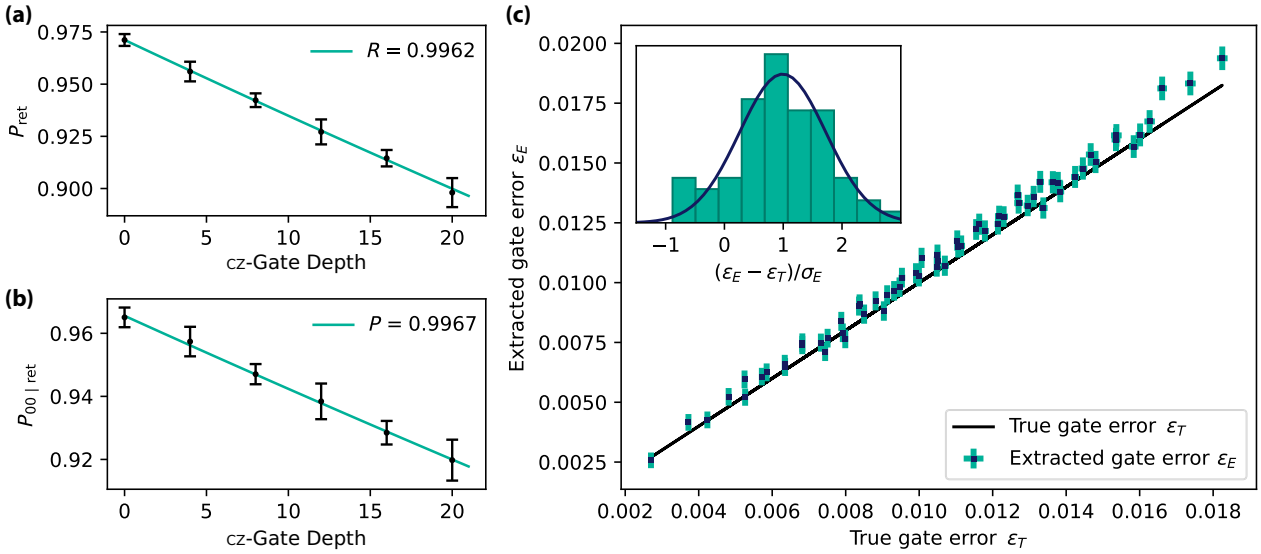


FIG. 7. Simulated CZ benchmarking. (a) The predicted probability of retaining both atoms in a simulated benchmarking run, using the error model described in Sec. VI F and the circuits used to collect the data presented in Fig. 3 [cf. Fig. 3(f)]. (b) The predicted probability of obtaining the expected state-selective measurement outcome, given that both atoms have been retained, in the same simulated benchmarking run [cf. Fig. 3(g)]. This simulation neglects projection noise. (c) The results of a simulation of 64 complete benchmarking runs, as described in Sec. VI F. This simulation includes projection noise. The figure depicts the extracted gate error and confidence interval versus the calculated gate error, obtained by averaging over Haar-random two-qubit input states. The inset histogram shows the observed distribution of the difference between these two quantities, normalized to the uncertainty in the extracted gate error in each case. The depicted Gaussian fit to this data has a mean of 0.99 and a standard deviation of 0.74, indicating that (i) our benchmarking procedure systematically underreports fidelity (by roughly 0.04%) and (ii) the extracted uncertainty somewhat overstates the statistical uncertainty.

of runs, with an average difference $\varepsilon_E - \varepsilon_T = 0.04\%$. We present further quantitative details in Fig. 7(c).

G. Impact and mitigation of optical crosstalk in the scanned-beam addressing architecture

In our addressing geometry, the lasers that drive Rydberg excitation propagate perpendicular to and are focused in the plane of the atomic qubit array, across which their foci are scanned using acousto-optic deflectors. System-level performance considerations (e.g., increasing qubit count and connectivity) exert pressure to reduce spacing in atom arrays. While in principle addressing the spatial profiles of beams can be shaped to effectively eliminate intensity on nontarget atoms (which we refer to as "crosstalk"), in practice crosstalk elimination is challenging over large optical field of view, for optical systems consisting of many optical surfaces, and for atom separations that approach the optical wavelength. For example, given 6-μm qubit spacing and 3-μm spot radii for Gaussian Rydberg beams, one predicts an intensity crosstalk η, the ratio of intensity on a neighboring spectator site to peak intensity, of 0.03%.

In practice, atomic measurements suggest η up to 10 times larger in the worst case, likely due to the impact of wave-front distortions introduced by optical surface imperfections. The practical inevitability of crosstalk motivates a discussion of two of its primary effects and their relevance to high-fidelity and fault-tolerant operation of our neutral-atom quantum processor.

1. Effect of optical crosstalk on entangling gate error

In an entangling gate, multiple atoms are addressed by different fields and crosstalk error contributions fundamentally arise due to the interference between the direct addressing field and the crosstalk field on each atom. If this interference can be compensated—e.g., by appropriately modulating the intensities of the fields—then the error contribution vanishes. However, such compensation presents a challenging control problem, suggesting that a different solution is more promising: we take an approach in which the frequency of this interference ω_X can be made much larger than the Rabi frequency Ω describing the gate dynamics. Simulations confirm that as the ratio ω_X/Ω increases, the Rabi dynamics average over this increasingly fast perturbation, allowing the desired gate dynamics to be realized. They also reveal that achieving a ratio of ω_X/Ω of about 50 is sufficient, for the particular crosstalk intensities that we investigate, to make the CZ error contribution from optical crosstalk negligible compared to error correction thresholds. Our system currently operates at $\omega_X/\Omega \approx 5$ but with modifications to the frequency or diffraction order of the pulse-shaping AOMs, it is straightforward to reach $\omega_X/\Omega > 50$. This modeling suggests that maintaining a sufficiently large beat note between neighboring addressing beams is a straightforward and effective way to mitigate crosstalk in an individual-addressing architecture. Future work could explore alternative mitigation strategies such as operation at zero beat-note frequency with active or passive stabilization of the relative phase between the two beams.

We depict simulations of the impact of optical crosstalk on single-atom Rabi oscillations to the Rydberg state and on CZ error in Fig. 8.

2. Effect of optical crosstalk on spectator qubits

During implementation of an $R_Z(\phi)$ or CZ gate, the addressing lasers irradiate spectator qubits with some small crosstalk intensity. In the case of an $R_Z(\phi)$ gate, crosstalk from the 459-nm laser produces an unintended $RZ(\phi_X)$ gate on spectator qubits, where $\phi_X = \eta\phi$ for relative crosstalk intensity η. The realized $RZ(\phi_X)$ rotation implements the desired identity operation on the spectator qubit with an expected error $\epsilon_X = \frac{2}{3}\sin^2(\phi_X/2)$ for a Haar-random input spectator state. This error is quadratic in the crosstalk intensity η. As an example, if we take $\eta = 0.003$ as an approximate description of our current system, for $\phi = \pi$ we have $\epsilon_X = 1.5 \times 10^{-5}$.

In the case of a CZ gate, the 459-nm laser applies a similar undesired $RZ(\phi_X)$ rotation to spectator qubits. For our current parameters, the rotation applied to a spectator qubit by a CZ gate is equivalent to that from an $(\phi = 1.7\pi)$ gate on the target qubit, for which we calculate $\epsilon_X = 4.3 \times 10^{-5}$ for $\eta = 0.003$.

The impact of these spectator-qubit rotations on the overall circuit fidelity is circuit dependent. We estimate the impact using simple assumptions: that every qubit is targeted by the same number of $R_Z(\phi)$ gates and that every nearest-neighbor pair is targeted by the same number of CZ gates. We calculate an effective adjustment to the fidelities of the $R_Z(\phi)$ and CZ gates on a given atom by including error accumulated during other $R_Z(\phi)$ and CZ gates in which that atom is a nearest-neighbor spectator. In a square lattice under this assumption of uniform gate distribution, the ratio of nearest-neighbor-spectator or target $R_Z(\phi)$ rotations a given qubit experiences is 4, while for CZ gates this ratio is 3. We can therefore add an effective heuristic error contribution to the $R_Z(\phi)$ gate error budget of $4 \times 1.5 \times 10^{-5} = 6 \times 10^{-5}$. To the CZ-gate error budget, we add an analogous $2 \times 3 \times 4.3 \times 10^{-5} = 2.6 \times 10^{-4}$, where the first factor of 2 arises from the number of qubits participating in the gate. This suggests that reducing the crosstalk intensity may be important for future fault-tolerant operation but a more complete analysis should calculate the circuit-specific impact. Another possibility that should be explored is crosstalk-aware compilation [39], in which the native gate set contains gates not of the form $R_Z(\phi) \otimes I \otimes \cdots \otimes I$ but, instead, $R_Z(\phi) \otimes R_Z(\eta_1\phi) \otimes \cdots \otimes R_Z(\eta_n\phi)$ and similar for the CZ gate.

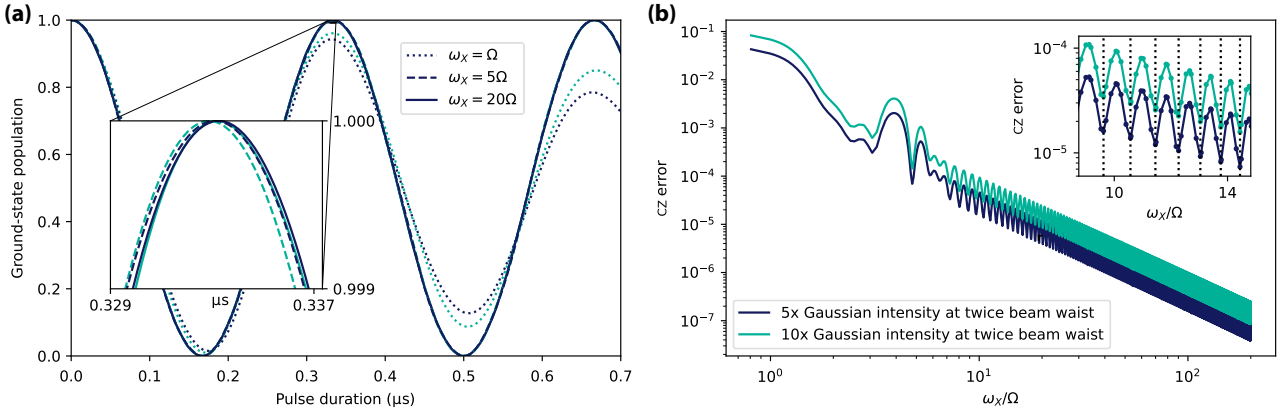


FIG. 8. The optical crosstalk effect on CZ gates. (a) Representative simulated trajectories for Rabi oscillations between the $6S_{1/2}$ ground state and the $69S_{1/2}$ Rydberg level, including the effect of optical crosstalk from Rydberg-excitation light addressing a neighboring site. The crosstalk intensity used in the simulation is 5× higher than the intensity 6 μm from the center of a Gaussian beam with 3 μm waist radius at $1/e^2$ of peak intensity (see the text). An important parameter is the crosstalk frequency ω_X relative to the frequency $\Omega = 2\pi \times 3$ MHz of Rabi oscillations and three values for this parameter are shown: $\omega_X/\Omega = 1$ (dotted trajectories), 5 (dashed trajectories), and 20 (solid trajectories). For each crosstalk frequency, we display two trajectories (distinguishable by color), with different phase relationships between the crosstalk fields for each trajectory, highlighting the sensitivity to this phase relationship for $\omega_X/\Omega = 1$. Simulations include both Rabi frequency and light-shift oscillations that arise due to optical crosstalk and include no other error sources. (b) The calculated CZ error as a function of the ratio ω_X/Ω. The error can be made smaller by increasing this ratio. At ratios greater than about 6, regular structure appears, with error minima when the gate duration is an integer multiple of the crosstalk period. The inset shows this structure over a smaller range, with vertical lines indicating the frequencies at which this integral relationship is satisfied; it also emphasizes our discrete sampling of the ratio ω_X/Ω in the simulation. We find that the gate error is approximately proportional to the crosstalk intensity, for which we depict two values here.

There are two more effects arising from the CZ gate that we consider: compared to the 459-nm beam, the differential qubit shift from the 1040-nm Rydberg-excitation beam is smaller by a factor larger than 10^3, so we neglect the very small rotation applied by the 1040-nm crosstalk. Finally, coupling to the Rydberg level may arise from illumination by 1040-nm crosstalk and 459-nm crosstalk. If we assume the same relative crosstalk intensity η for both colors, the Rabi frequency of this coupling is $\eta\Omega$, where Ω is the Rabi frequency for the target qubit. Additionally, the ground-to-Rydberg transition frequency of the spectator atom differs from that of the target atom by (approximately) the value of the light shifts applied to that transition by the Rydberg excitation beams. This produces a ground-to-Rydberg coupling on the spectator atom that is (a) weaker by a factor of η and (b) very far detuned. Concretely, for our current system parameters and $\eta = 0.003$, we estimate that this coupling is relatively detuned by an amount $\Delta_X/\Omega_X = 10^3$. The time-averaged spectator-qubit population in the Rydberg state arising from this coupling is 5×10^{-7}; we therefore fully neglect this effect as well.

H. NDSSR implementation

1. NDSSR sequence and optimization

Our implementation of nondestructive state-selective readout improves upon the setup demonstrated in Ref. [49]. At the end of circuit execution, the bias field is rotated from the z axis to the x axis [for orientation, see Fig. 1(c)] along which two counterpropagating circularly polarized near-resonant 852-nm beams optically pump atoms in the upper hyperfine manifold $F = 4$ to the stretched state $|F = 4, m_F = 4\rangle$ and scatter photons primarily along the closed transition $|F = 4, m_F = 4\rangle \rightarrow |F = 5', m_F = 5'\rangle$. The beams are red detuned from zero-field resonance by $2\pi \times 1.6$ MHz, with one beam further detuned by $2\pi \times 3$ kHz to wash out the standing-wave intensity pattern; the peak beam intensity per beam is estimated to be approximately 1 mW/cm^2, which is close to the saturation intensity of the transition (1.1 mW/cm^2). Atoms that are in the lower hyperfine manifold $F = 3$ stay dark during this process. The intensities of the near-resonant beam-pair and the trap laser are modulated out-of-phase, so that the trap laser intensity is minimized while the atoms have a high probability of being excited by the imaging beams. The trap intensity modulation has a nominal duty cycle of 75% and the imaging beam pair has a nominal duty cycle of 15%. The duty cycle and the phase offset of the intensity modulation waveforms are controlled by the direct digital synthesis (DDS) cards driving the acousto-optic modulators for the lasers. Because there is no 3D cooling, with only Doppler cooling expected on the imaging beam axis, the atom temperature rises quickly during this readout and hence the exposure time has to be limited such that atom loss from the trap is minimized

while achieving good discrimination between bright and dark distributions of the camera signal. An exposure time of 6 ms has been found to be sufficient to provide high base-discrimination fidelity (see Sec. VI H 2). By carefully optimizing the laser-power modulation parameters, the power balance between counterpropagating beams, and the bias field, we have been able to achieve a state-average atom-loss probability of $<1\%$ [Fig. 4(b)] in a time-averaged trap depth of 1.97 mK (peak trap depth 2.6 mK), which is significantly lower than the >10 mK used in Refs. [49] and [71] while achieving a better atom loss rate. Preliminary results of further optimization and technical improvements not discussed in this work suggest that the peak trap depth can be lowered to 2 mK while maintaining the state discrimination fidelity and atom loss rate mentioned in Sec. IV. In particular, we have observed an interesting correlation between the optimal intensity modulation frequency for minimal atom loss and the Zeeman energy difference between m_F levels of the upper-hyperfine manifold, which warrants further investigation; for the readout sequence used in this work, we have chosen an x-bias field of 3.59 G and an intensity modulation frequency of 1.32 MHz. Leakage of off-resonant repump laser light due to the finite extinction ratio of acousto-optic switches can result in increased dark-to-bright state leakage at deep trap depths, which accounts for some of the 1.6(5)% dark-state discrimination error mentioned in Sec. IV.

2. Base-discrimination-fidelity modeling

Following Ref. [50], we define the base-discrimination fidelity as the probability of correctly classifying the photoelectron counts as coming from atomic fluorescence. For each qubit-trapping site, we define a mask for selecting relevant camera pixels and then compute a weighted sum of raw camera analog-digital units (ADUs) for those pixels (note each pixel has a fixed offset, which must be subtracted); the sum is then converted into photoelectron counts by dividing the sum by the EM gain (300) and then multiplying by the preamplifier gain (5.16 electrons per ADU count) provided by the EMCCD vendor data sheet. We model the distribution of photoelectron counts $P(x)$ as a mixture of two distributions: exponentially modified Gaussian for the dark distribution $P_D(x)$ weighted by the dark-count fraction A_D, and skew-normal for the bright distribution $P_B(x)$ weighted by the bright-count fraction $1-A_D$:

$$P(x) = A_D P_D(x) + (1-A_D) P_B(x), \tag{8}$$

$$P_D(x) = \frac{1}{2K\sigma_D} \exp\left(\frac{1}{2K^2} - \left(\frac{x-\mu_D}{K\sigma_D}\right)\right) \times \operatorname{erfc}\left(-\frac{(x-\mu_D)/\sigma_D - 1/K}{\sqrt{2}}\right), \tag{9}$$

$$P_B(x) = \frac{1}{\sqrt{2\pi}\sigma_B} \exp\left(-\frac{(x-\mu_B)^2}{2\sigma_B^2}\right) \times \left(1 + \operatorname{erf}\left(a\frac{x-\mu_B}{\sqrt{2}\sigma_B}\right)\right). \tag{10}$$

The dark distribution is parametrized by μ_D, σ_D, and K, which describe the location, scale, and shape of the dark distribution, respectively. The bright distribution is parametrized by μ_B, σ_B, and a, which describe the location, scale, and shape of the bright distribution, respectively. $\operatorname{erf}(x)$ and $\operatorname{erfc}(x)$ are the error function and the complementary error function. We have obtained the following parameter values after fitting the proposed mixture distribution to the representative SPAM site data used in Fig. 4(a): $A_D = 0.506$, $\mu_D = -0.100$, $\sigma_D = 0.673$, $K = 1.836$, $\mu_B = 27.471$, $\sigma_B = 9.597$, and $a = 1.231$ [for the fit plot, see Fig. 9(a)]. Regarding the choice of using exponentially modified Gaussian distribution for the dark distribution, we believe that a convolution of normal and Gamma distributions is more appropriate, based on the observation that the EMCCD ADU count per pixel distribution in the dark can be modeled as a mixture of normal (from readout noise) and Erlang distributions (from parallel clock-induced charge noise) [72], but we have found that an exponentially modified Gaussian distribution fits the data just as well and is easier to compute; the skew-normal distribution is chosen for the bright distribution for a similar reason of convenience.

Once a mixture distribution is fitted to the data, the optimal discrimination threshold θ^* is then determined, as the photoelectron count minimizes the sum of classification errors $E_C(\theta)$, which is

$$E_C(\theta) = A_D \int_{\theta}^{\infty} P_D(x)\,dx + (1-A_D) \int_{-\infty}^{\theta} P_B(x)\,dx. \tag{11}$$

To estimate the discrimination fidelity $1 - E_C(\theta^*)$ using this model, we have applied bootstrapping to the SPAM image data set used in Fig. 4: we have created 1000 SPAM image subsets, with each subset containing 2000 images randomly sampled from the main data set with replacement, fitted the mixture distribution for the photoelectron counts in each subset, and then computed $E_C(\theta)$. The bootstrapping result [Fig. 9(b)] shows that the discrimination fidelity saturates at 99.96(2)%, around the discrimination threshold of 10.235 photoelectron counts. We note that the mean and the standard deviation of the optimal discrimination fidelities from the 25 qubits used in the SPAM data set also return a similar estimate. We interpret this high discrimination-fidelity estimate as the asymptotic performance limit of the imaging system if there were no depumping error during the imaging sequence, which manifests as a spread of photoelectron counts between the

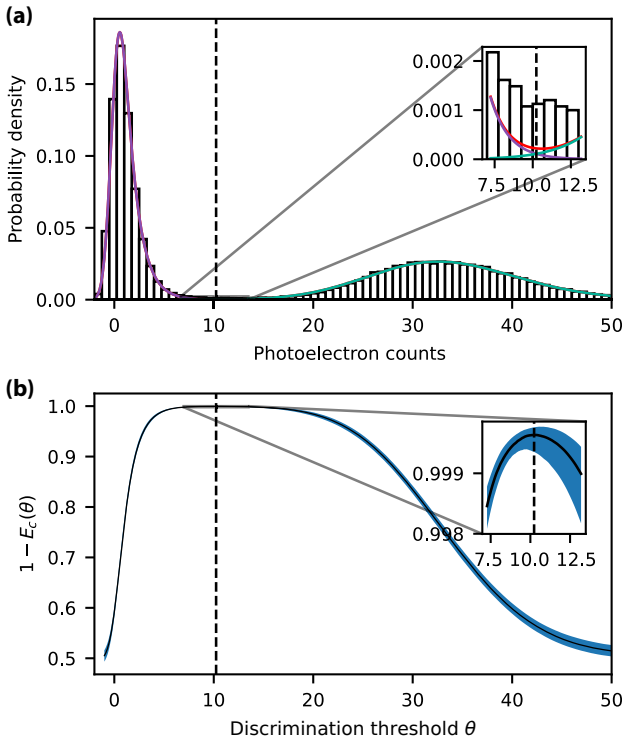


FIG. 9. NDSSR discrimination-fidelity modeling. The dashed black line denotes the optimal discrimination threshold of 10.235 photoelectron counts. (a) The photoelectron count distribution from the data set used in Fig. 4. The red line is a fit to a model that is a mixture of an exponentially modified Gaussian distribution (purple) and a skew-normal distribution (turquoise). The inset shows an enlargement of the sum of the two component distributions (red) around the optimal discrimination threshold, where there is a small overlap. (b) The discrimination-fidelity estimate as a function of the chosen threshold θ (dark line) with a 1σ confidence interval (blue band) obtained from bootstrapping. The inset shows an enlargement of the discrimination-fidelity estimate around the optimal discrimination threshold.

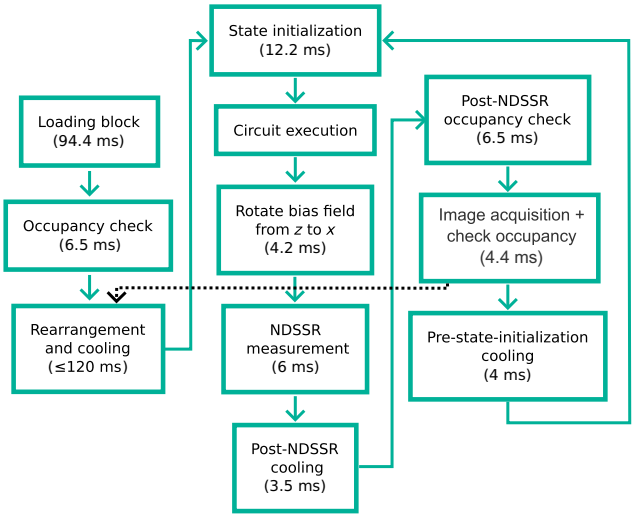


FIG. 10. The NDSSR control flow. The control flow of an NDSSR experiment cycle, with an approximate duration for each step. At the step with the solid green arrow and dotted black arrow, if an atom loss is detected from the qubit array or if the atom reservoir does not have enough resources for rearrangement, the control sequence follows the dotted black arrow. Otherwise, the green arrow is followed. When the required number of measurements has been collected or the atom reservoir runs out of atoms, the control sequence segues into the monitor block (not shown, duration approximately 10 ms) to check the laser intensity and lock status. Once the monitoring data are gathered and processed over the network (approximately 20 ms), feedback corrections are applied to actuators and another cycle is launched. If the Raman sideband cooling (RSC) step is enabled, an extra step is inserted between the state-initialization step and the circuit-execution step, which effectively adds an extra 52-ms duration with our existing parameters.

dark and the bright distributions not captured by the mixture model [see the inset in Fig. 9(a)]. If state-preparation error is negligible, one could attempt to calculate both the discrimination fidelity and the depumping probability by analyzing a data set of consecutive imaging measurements and relating the observed probabilities of dark-dark, dark-bright, and bright-bright events to the quantities of interest, similar to the approach taken in Ref. [73, Sec. 2.6.7].

3. Shot-rate enhancement from repetitive measurements

In our architecture, a measurement cycle is defined by a sequence of repetitive measurements that continues until the maximum measurements allowed per cycle are reached or the atom reservoir is depleted (see Fig. 10). The shot-rate enhancement effect of repetitive measurement has been measured by averaging the total execution times of a quantum circuit that prepares a 15-qubit GHZ state with CZ gates for entangling, a local $R_Z(\phi)$ gate for qubit phase correction, and global $R(\theta,\phi)$ gates for applying Hadamard and dynamical decoupling gates. The test circuit has consisted of 14 CZ gates, 14 $R_Z(\phi)$ gates, and 29 $R(\theta,\phi)$ gates. The circuit has been repeated 3 times, with each instance requiring at least 400 valid measurements, maximum measurements per measurement cycle limited to seven, and recorded total execution times of 47.7 s, 50.3 s, and 48.1 s, which results in an average shot rate of 8.2(2) Hz. In this context, a raw measurement is invalidated if the post-NDSSR occupancy measurement has detected an atom loss in the qubit array and multiple measurements in a measurement cycle are invalidated if the analog (laser intensity) and digital (laser lock status) validators have detected samples out of validation bounds at the end of a cycle; the total execution time is defined by the time interval between the beginning of the pulse-sequence compilation and control-device initialization and the end of the job-data (camera image set) upload to the database. It should be noted that these data have been taken when the CZ-gate fidelity has been significantly

lower (the average SPAM-corrected Bell-state preparation fidelity was 93.1%) than the best CZ-gate fidelity demonstrated in this work and the NDSSR bright-state atom loss rate has been slightly higher (approximately 1.5% instead of 1%), resulting in higher atom loss and incurring more measurement invalidation and atom loading and arrangement iterations. For comparison, concurrent measurements indicate that without NDSSR, even an empty circuit (i.e. with zero gate-induced loss) would have been limited to 3.2 Hz.

Note added. We have recently become aware of related work that addresses the impact of leakage errors on randomized benchmarking in detail [74]. The approach taken in that work is consistent with ours.

ACKNOWLEDGMENTS

We gratefully acknowledge contributions from Alan Doak, Alex Olivas, Brenden Villars, Bryna Guriel, Cody Mart, Chris Wood, David Owusu-Antwi, Denny Dahl, Emma Brann, Evan Salim, Jan Preusser, Jean-Philippe Feve, John Sivak, Jonathan Cohen, Jorge Gomez, Josh Behne, Josh Cherek, Karen Rosenthall, Marie Grubb, Marissa McMaster, Matt Ebert, Rob Williamson, Robin Keus, Ryan Prior, Salahedeen Issa, Stephanie Lee, Steven Kitchen, and Steven Sedig.

A.G.R., D.M., T.W.N., W.C.C., D.M.G., A.M.S., D.C.C., M.C.G., M.T.L., G.T.H., I.L.B., T.O., E.C., M.S., M.Y., S.D.S., A.J.F., R.A.J., K.W.K., F.M., I.V.V., M.L.W., M.I., P.T.M., J.R.T., M.K.D., J.R.G., M.R.B., J.K.M., A.R., and T.M.G. contributed to the design and construction of the apparatus hardware, A.G.R., D.M., T.W.N., W.C.C., A.M.S., N.G.P., D.C.C., T.G.B., M.T.L., G.T.H., D.G., I.L.B., T.B.L., C.C., P.G., V.O., P.G., E.C., M.J.B., M.E.B., A.K.T., E.B.J., R.A.J., K.W.K., S.Y.E., D.A.B., C.P.S., I.V.V., J.J.M., P.T.M., J.R.T., M.K.D., J.R.G., N.A.N.M., C.J.G.J., P.J.R., L.L.H., P.D.B., R.R., and M.A.P. contributed to the control system and software stack. A.G.R., D.M., T.W.N., W.C.C., D.C.C., M.T.L., G.T.H., C.C., P.G., V.O., E.C., E.B.J., R.A.J., K.W.K., I.V.V., P.T.M., J.R.T., M.K.D., and J.R.G. contributed to operating the apparatus and collecting the data. A.G.R., D.M., T.W.N., W.C.C., D.M.G., D.C.C., M.T.L., G.T.H., I.L.B., C.C., P.G., V.O., T.O., P.G., J.M.G., E.C., M.S., E.B.J., J.D.M., R.A.J., K.W.K., I.V.V., J.A.M., P.T.M., J.R.T., M.K.D., J.R.G., N.A.N.M., A.C., and T.M.G. contributed to analyzing data, A.G.R., D.M., T.W.N., W.C.C., D.C.C., M.T.L., G.T.H., C.C., E.C., M.S., M.Y., E.B.J., J.D.M., K.W.K., I.V.V., J.R.T., J.R.G., R.R., T.M.G., and M.A.P. contributed to theoretical analysis and simulation. The manuscript was written by A.G.R., D.M., T.W.N., W.C.C., D.C.C., and M.S. The project was supervised by A.G.R., T.W.N., W.C.C., D.C.C., D.G., P.G., T.O., M.E.B., M.S., J.J.M., and M.I.

A.G.R., D.M., T.W.N., W.C.C., D.M.G., A.M.S., N.G.P., D.C.C., M.C.G., T.G.B., M.T.L., G.T.H., D.G., I.L.B., T.B.L., C.C., P.G., V.O., T.O., P.G., J.M.G., E.C., M.J.B., M.E.B., M.S., A.K.T., M.Y., S.D.S., E.B.J., A.J.F., J.D.M., R.A.J., K.W.K., S.Y.E., D.A.B., F.M., C.P.S., I.V.V., J.J.M., M.L.W., J.A.M., M.I., P.T.M., J.R.T., M.K.D., J.R.G., N.A.N.M., M.R.B., C.J.G.J., P.J.R., J.K.M., A.R., L.L.H., P.D.B., R.R., A.C., and M.A.P. are employees and/or shareholders of ColdQuanta, Inc. d.b.a. Infleqtion.

DATA AVAILABILITY

The data that support the findings of this paper are not publicly available because they contain commercially sensitive information. The data are available from the authors upon reasonable request.

[1] P. W. Shor, in *Proceedings—Annual IEEE Symposium on Foundations of Computer Science, FOCS* (IEEE, Santa Fe, NM, USA, 1994), p. 124.

[2] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, Rev. Mod. Phys. **92**, 015003 (2020).

[3] X. Gao, E. R. Anschuetz, S. T. Wang, J. I. Cirac, and M. D. Lukin, Enhancing generative models via quantum correlations, Phys. Rev. X **12**, 021037 (2022).

[4] R. Santagati, A. Aspuru-Guzik, R. Babbush, M. Degroote, L. González, E. Kyoseva, N. Moll, M. Oppel, R. M. Parrish, N. C. Rubin, M. Streif, C. S. Tautermann, H. Weiss, N. Wiebe, and C. Utschig-Utschig, Drug design on quantum computers, Nat. Phys. **20**, 549 (2024).

[5] P. Date, D. Arthur, and L. Pusey-Nazzaro, QUBO formulations for training machine learning models, Sci. Rep. **11**, 1 (2021).

[6] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, npj Comput. Mater. **6**, 1 (2020).

[7] S. Koretsky *et al.*, in *Proceedings—2021 IEEE International Conference on Quantum Computing and Engineering, QCE 2021* (IEEE, Broomfield, CO, USA, 2021), p. 181.

[8] M. A. Perlin *et al.*, Q-CHOP: Quantum constrained Hamiltonian optimization, Arxiv:2403.05653v1.

[9] S. Ramesh, T. Tomesh, S. J. Riesenfeld, F. T. Chong, and A. T. Pearson, Quantum computing for oncology, Nat. Cancer **5**, 811 (2024).

[10] T. Hoefler, T. Häner, and M. Troyer, Disentangling hype from practicality: On realistically achieving quantum advantage, Commun. ACM **66**, 82 (2023).

[11] D. Bluvstein *et al.*, Logical quantum processor based on reconfigurable atom arrays, Nature **626**, 58 (2024).

[12] P. S. Rodriguez, J. M. Robinson, P. N. Jepsen, Z. He, C. Duckering, C. Zhao, K.-H. Wu, J. Campo, K. Bagnall, M. Kwon *et al.*, Experimental demonstration of logical magic state distillation, Nature (2025).

[13] B. W. Reichardt *et al.*, Logical computation demonstrated with a neutral atom quantum processor, ArXiv:2411.11822.

[14] G. Q. AI, Suppressing quantum errors by scaling a surface code logical qubit, Nature **614**, 676 (2023).

[15] M. P. da Silva *et al.*, Demonstration of logical qubits and repeated error correction with better-than-physical error rates, Arxiv:2404.02280v2.

[16] M. Bedalov *et al.*, Fault-tolerant operation and materials science with neutral atom logical qubits, ArXiv:2412.07670.

[17] J. Bergou, M. Hillery, and M. Saffman, *Quantum Information Processing: Theory and Implementation* (Springer, Cham, Switzerland, 2021), 2nd ed.

[18] H. J. Manetsch, G. Nomura, E. Bataille, K. H. Leung, X. Lv, and M. Endres, A tweezer array with 6100 highly coherent atomic qubits, arxiv:2403.12021v2.

[19] T. M. Graham *et al.*, Multi-qubit entanglement and algorithms on a neutral-atom quantum computer, Nature **604**, 457 (2022).

[20] M. Peper, Y. Li, D. Y. Knapp, M. Bileska, S. Ma, G. Liu, P. Peng, B. Zhang, S. P. Horvath, A. P. Burgers, and J. D. Thompson, Spectroscopy and modeling of ^{171}Yb Rydberg states for high-fidelity two-qubit gates, Phys. Rev. X **15**, 011009 (2025).

[21] R. Finkelstein, R. B.-S. Tsai, X. Sun, P. Scholl, S. Direkci, T. Gefen, J. Choi, A. L. Shaw, and M. Endres, Universal quantum operations and ancilla-based read-out for tweezer clocks, Nature **634**, 321 (2024), https://www.nature.com/articles/s41586-024-08005-8.

[22] T. Y. Wu, Neutral-atom quantum computer with ^{87}Sr nuclear-spin qubits (2024). https://meetings.aps.org/Meeting/DAMOP24/Session/J06.3.

[23] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with Rydberg atoms, Rev. Mod. Phys. **82**, 2313 (2010).

[24] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin, High-fidelity parallel entangling gates on a neutral-atom quantum computer, Nature **622**, 268 (2023).

[25] J. Muniz, M. Stone, D. T. Stack, M. Jaffe, J. M. Kindem, L. Wadleigh, E. Zalys-Geller, X. Zhang, C.-A. Chen, M. A. Norcia *et al.*, High-fidelity universal gates in the ^{171}Yb ground-state nuclear-spin qubit, PRX Quantum **6**, 020334 (2025).

[26] S. Jandura and G. Pupillo, Time-optimal two- and three-qubit gates for Rydberg atoms, Quantum **6**, 712 (2022), https://quantum-journal.org/papers/q-2022-05-13-712/.

[27] M. A. Norcia *et al.*, Iterative assembly of ^{171}Yb atom arrays with cavity-enhanced optical lattices, PRX Quantum **5**, 030316 (2024).

[28] P. Huft, Y. Song, T. M. Graham, K. Jooya, S. Deshpande, C. Fang, M. Kats, and M. Saffman, Simple, passive design for large optical trap arrays for single atoms, Phys. Rev. A **105**, 063111 (2022).

[29] D. B. Tan, D. Bluvstein, M. D. Lukin, and J. Cong, Compiling quantum circuits for dynamically field-programmable neutral atoms array processors, Quantum **8**, 1281 (2024).

[30] M. R. Lam, N. Peter, T. Groh, W. Alt, C. Robens, D. Meschede, A. Negretti, S. Montangero, T. Calarco, and A. Alberti, Demonstration of quantum brachistochrones between distant states of an atom, Phys. Rev. X **11**, 011035 (2021).

[31] C. Poole, T. M. Graham, M. A. Perlin, M. Otten, and M. Saffman, Architecture for fast implementation of qLDPC codes with optimized Rydberg gates, arxiv:2404.18809v1.

[32] E. Knill, Quantum computing with realistically noisy devices, Nature **434**, 39 (2005).

[33] K. Sahay, J. Jin, J. Claes, J. D. Thompson, and S. Puri, High-threshold codes for neutral-atom qubits with biased erasure errors, Phys. Rev. X **13**, 041013 (2023).

[34] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, Nature **627**, 778 (2024).

[35] Y. Hong, M. Marinelli, A. M. Kaufman, and A. Lucas, Long-range-enhanced surface codes, Phys. Rev. A **110**, 022607 (2024).

[36] via Infleqtion Superstaq [39].

[37] Infleqtion PICAS.

[38] T. M. Graham, E. Oh, and M. Saffman, Multiscale architecture for fast optical addressing and control of large-scale qubit arrays, Appl. Opt. **62**, 3242 (2023).

[39] C. Campbell *et al.*, in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Bellevue, WA, USA, 2023), p. 1020.

[40] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais, and P. Grangier, Energy distribution and cooling of a single atom in an optical tweezer, Phys. Rev. A **78**, 033425 (2008).

[41] S. Kuhr, W. Alt, D. Schrader, I. Dotsenko, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede, Analysis of dephasing mechanisms in a standing-wave dipole trap, Phys. Rev. A **72**, 023406 (2005).

[42] C. H. Baldwin, B. J. Bjork, J. P. Gaebler, D. Hayes, and D. Stack, Subspace benchmarking high-fidelity entangling operations with trapped ions, Phys. Rev. Res. **2**, 013317 (2020).

[43] K. Barnes *et al.*, Assembly and coherent control of a register of nuclear spin qubits, Nat. Commun. **13**, 1 (2022).

[44] S. Ma, G. Liu, P. Peng, B. Zhang, S. Jandura, J. Claes, A. P. Burgers, G. Pupillo, S. Puri, and J. D. Thompson, High-fidelity gates and mid-circuit erasure conversion in an atomic qubit, Nature **622**, 279 (2023).

[45] R. B.-S. Tsai, X. Sun, A. L. Shaw, R. Finkelstein, and M. Endres, Benchmarking and fidelity response theory of high-fidelity Rydberg entangling gates, PRX Quantum **6**, 010331 (2025).

[46] J. Preskill, Reliable quantum computers, Proc.: Math. Phys. Eng. Sci. **454**, 385 (1998).

[47] M. N. H. Chow, V. Buchemmavari, S. Omanakuttan, B. J. Little, S. Pandey, I. H. Deutsch, and Y.-Y. Jau, Circuit-based leakage-to-erasure conversion in a neutral-atom quantum processor, PRX Quantum **5**, 040343 (2024).

[48] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays, Nat. Commun. **13**, 4657 (2022).

[49] M. Kwon, M. F. Ebert, T. G. Walker, and M. Saffman, Parallel low-loss measurement of multiple atomic qubits, Phys. Rev. Lett. **119**, 180504 (2017).

[50] J. P. Covey, H. Weinfurter, and H. Bernien, Quantum networks with neutral atom processing nodes, npj Quantum Inf. **9**, 90 (2023).

[51] L. Su, A. Douglas, M. Szurek, A. H. Hebert, A. Krahn, R. Groth, G. A. Phelps, O. Marković, and M. Greiner, Fast

single atom imaging for optical lattice arrays, Nature **16** (2025).
[52] T. M. Graham, L. Phuttitarn, R. Chinnarasu, Y. Song, C. Poole, K. Jooya, J. Scott, A. Scott, P. Eichler, and M. Saffman, Midcircuit measurements on a single-species neutral alkali atom quantum processor, Phys. Rev. X **13**, 041051 (2023).
[53] S. Jandura, J. D. Thompson, and G. Pupillo, Optimizing Rydberg gates for logical-qubit performance, PRX Quantum **4**, 020336 (2023).
[54] M. Mohan, R. de Keijzer, and S. Kokkelmans, Robust control and optimal rydberg states for neutral atom two-qubit gates, Phys. Rev. Res. **5**, 033052 (2023).
[55] I. Christen, T. Propson, M. Sutula, H. Sattari, G. Choong, C. Panuski, A. Melville, J. Mallek, C. Brabec, S. Hamilton, P. Benjamin Dixon, A. J. Menssen, D. Braje, A. H. Ghadimi, and D. Englund, An integrated photonic engine for programmable atomic control, Nat. Commun. **16**, 82 (2025).
[56] K.-N. Schymik, B. Ximenez, E. Bloch, D. Dreon, A. Signoles, F. Nogrette, D. Barredo, A. Browaeys, and Th. Lahaye, *In situ* equalization of single-atom loading in large-scale optical tweezer arrays, Phys. Rev. A **106**, 022611 (2022).
[57] M. E. Shea, P. M. Baker, J. A. Joseph, J. Kim, and D. J. Gauthier, Submillisecond, nondestructive, time-resolved quantum-state readout of a single, trapped neutral atom, Phys. Rev. A **102**, 053101 (2020).
[58] M. N. Chow, B. J. Little, and Y. Y. Jau, High-fidelity low-loss state detection of alkali-metal atoms in optical tweezer traps, Phys. Rev. A **108**, 032407 (2023).
[59] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik, H. Neven, and R. Babbush, Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via Trotterization, Quantum **4**, 296 (2020).
[60] E. T. Campbell, Early fault-tolerant simulations of the Hubbard model, Quantum Sci. Technol. **7**, 015007 (2021).
[61] S. Anand, C. E. Bradley, R. White, V. Ramesh, K. Singh, and H. Bernien, A dual-species Rydberg array, Nat. Phys. **20**, 1744 (2024), https://www.nature.com/articles/s41567-024-02638-2.
[62] K. Singh, C. E. Bradley, S. Anand, V. Ramesh, R. White, and H. Bernien, Mid-circuit correction of correlated phase errors using an array of spectator qubits, Science **380**, 1265 (2023).
[63] N. Chen, L. Li, W. Huie, M. Zhao, I. Vetter, C. H. Greene, and J. P. Covey, Analyzing the Rydberg-based optical-metastable-ground architecture for ^{171}Yb nuclear spins, Phys. Rev. A **105**, 052438 (2022).
[64] J. W. Lis, A. Senoo, W. F. McGrew, F. Rönchen, A. Jenkins, and A. M. Kaufman, Midcircuit operations using the *omg* architecture in neutral atom arrays, Phys. Rev. X **13**, 041035 (2023).
[65] B. Zhang, P. Peng, A. Paul, and J. D. Thompson, Scaled local gate controller for optically addressed qubits, Optica **11**, 227 (2024).
[66] S. Ma, A. P. Burgers, G. Liu, J. Wilson, B. Zhang, and J. D. Thompson, Universal gate operations on nuclear spin qubits in an optical tweezer array of ^{171}Yb atoms, Phys. Rev. X **12**, 021028 (2022).
[67] P. Scholl, A. L. Shaw, R. B.-S. Tsai, R. Finkelstein, J. Choi, and M. Endres, Erasure conversion in a high-fidelity Rydberg quantum simulator, Nature **622**, 273 (2023).
[68] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, Parallel implementation of high-fidelity multiqubit gates with neutral atoms, Phys. Rev. Lett. **123**, 170503 (2019).
[69] M. A. Perlin, V. N. Premakumar, J. Wang, M. Saffman, and R. Joynt, Fault-tolerant measurement-free quantum error correction with multiqubit gates, Phys. Rev. A **108**, 062426 (2023).
[70] Z. Tian, H. Chang, X. Lv, M. Yang, Z. Wang, P. Yang, P. Zhang, G. Li, and T. Zhang, Resolved Raman sideband cooling of a single optically trapped cesium atom, Opt. Lett. **49**, 542 (2024).
[71] B. Nikolov, E. Diamond-Hitchcock, J. Bass, N. L. Spong, and J. D. Pritchard, Randomized benchmarking using non-destructive readout in a two-dimensional atom array, Phys. Rev. Lett. **131**, 030602 (2023).
[72] A. Bergschneider, V. M. Klinkhamer, J. H. Becher, R. Klemt, G. Zürn, P. M. Preiss, and S. Jochim, Spin-resolved single-atom imaging of ^{6}Li in free space, Phys. Rev. A **97**, 063613 (2018).
[73] I. S. Madjarov, Ph.D. thesis, California Institute of Technology, Pasadena, California, 2021. https://resolver.caltech.edu/CaltechTHESIS:01292021-001639979.
[74] Y.-H. Chen and C. H. Baldwin, Randomized benchmarking with leakage errors, Arxiv:2502.00154.