

# Quantum error correction with the toric code

Atom Computing and Collaborators  
(Dated: June 2026)

Quantum computing platforms based on arrays of tweezer-confined neutral atoms have recently emerged as a competitive modality thanks to a direct path toward high qubit count, rapidly advancing operation fidelities, and their ability to execute circuits with arbitrary qubit connectivity. These features will enable the use of efficient error correction schemes with high encoding-rates, time-efficient decoding, and resource-efficient architectures based on transversal gates. With these goals in mind, recent state of the art neutral atom demonstrations focus on the transition from the use of physical qubits to error-corrected logical qubits, but to date there has been no demonstration of repeated error correction scalable to arbitrary depth. Here, we demonstrate many cycles of syndrome extraction in a toric quantum error correcting code, using mid-circuit measurement and replacement of lost qubits, including reloading of a qubit reservoir for indefinite coherent operation. We characterize the logical error rate after up to 90 cycles, showing that logical information can be preserved through multiple rounds of qubit reloading. Comparing two distances of the code up to 8 rounds of syndrome extraction shows a lower absolute logical error rate for the larger distance code.

## I. INTRODUCTION

Quantum computing platforms based on arrays of tweezer-confined neutral atoms have recently emerged as a competitive modality thanks to a direct path toward high qubit count [1–5], rapidly advancing operation fidelities [6–10], and their ability to execute circuits with arbitrary qubit connectivity [11–13]. These features will enable the use of efficient error correction schemes with high encoding-rates [14–18], time-efficient decoding [19, 20], and resource-efficient architectures based on transversal gates [21, 22]. With these goals in mind, recent state of the art neutral atom demonstrations focus on the transition from the use of physical qubits to error-corrected logical qubits [12, 23–26], but to date there has been no demonstration of repeated error correction scalable to arbitrary depth. Here, we demonstrate many cycles of syndrome extraction in a toric quantum error correcting code, using mid-circuit measurement and replacement of lost qubits, including reloading of a qubit reservoir for indefinite coherent operation. We characterize the logical error rate after up to 90 cycles, showing that logical information can be preserved through multiple rounds of qubit reloading. Comparing two distances of the code up to 8 rounds of syndrome extraction shows a lower absolute logical error rate for the larger distance code.

Syndrome extraction—performing measurements of a subset of qubits that reveal the presence of errors—is a key requirement for fault-tolerant quantum computation [27]. In order to be used in universal quantum computation, and to avoid a qubit number overhead that grows with circuit depth, syndrome extraction measurements must be performed in a true mid-circuit manner—the state of some qubits must be read out while coherence is preserved in others. Recently, the error rate associated with repeated rounds of syndrome extraction using mid-circuit measurement has been characterized in superconducting [28–32] and trapped ion [33] platforms, and correlations in logical observables under repeated syndrome extraction have been explored in neutral-atom [25] sys-

tems.

Performing many rounds of syndrome extraction requires maintaining qubits available in a usable state, which is a challenge for atomic platforms. Atoms are easily heated during computation, degrading the performance of future gates, and can also be lost from their traps, especially during readout. This leads to a hierarchy of error mechanisms that become relevant at different depths, as illustrated in Fig. 1(a). Overcoming these challenges requires some combination of qubit reuse and replacement after measurement. Nondestructive measurement techniques can reduce the rate at which atoms need to be replaced, but true continuous operation requires to execute arbitrarily deep computations requires a way to draw new atoms from an inexhaustible source such as a magneto-optical trap (MOT) without causing excessive errors in existing atomic qubits.

The ability to form a magneto-optical (MOT) trap while maintaining coherence in computational qubits has been demonstrated using transport between spatially separated zones [5, 34, 35], or dual species systems [36]. Separated zones or shelving techniques also enable iterative loading of large arrays, where new atoms can be repeatedly loaded without loss of those already present [1, 37]. The ability to refill tweezers either stochastically [35] or in ordered arrays [5] from a spatially-separated MOT while maintaining coherence in existing qubits has recently been demonstrated using simple characterization circuits involving only single-qubit operations. In this work, we integrate continuous loading, mid-circuit measurement, and qubit reuse into a fault-tolerant quantum memory circuit, showing that syndrome extraction can be performed indefinitely at constant error rate.

Repeated syndrome extraction forms the core of fault-tolerant schemes for quantum computation. A quantum memory experiment characterizes the performance of the logical identity operation, and is the standard test of a platform’s ability to suppress the logical error rate via error correction. While no platforms currently achieve logical error rates compatible with utility-scale compu-

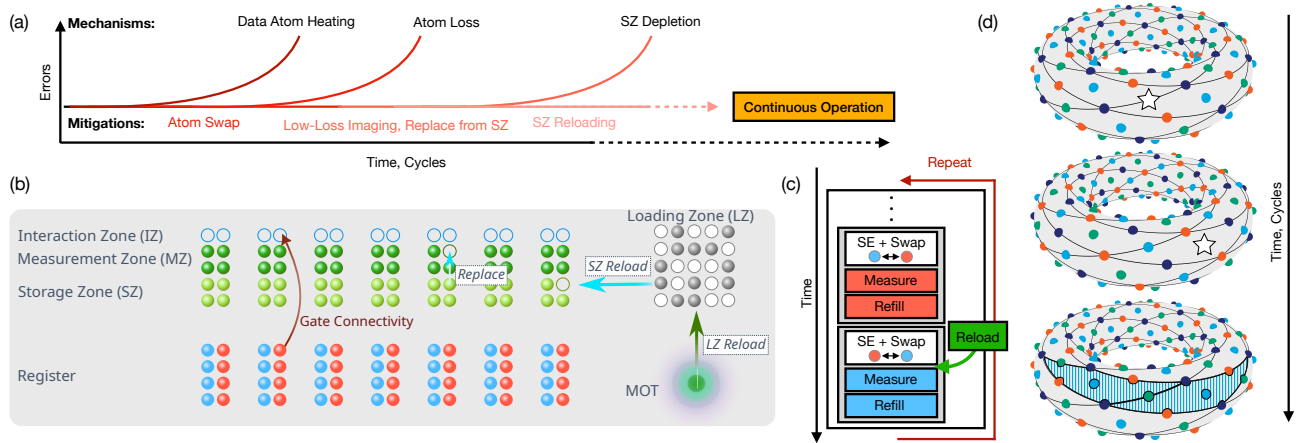


FIG. 1. Architecture for continuous operation of arbitrary logical computations. (a) Achieving continuous operation requires overcoming a series of mechanisms that limit the depth of computation that may be performed with neutral atoms. Swapping the role of data and ancilla qubits allows all atoms to be cooled during mid-circuit measurement, mitigating the effects of heating on gate performance. Low-loss imaging and refilling of the computational array from a local reservoir of atoms (the storage zone) extends loss-limited circuit execution, but the storage zone must be reloaded from an inexhaustible source to achieve true continuous operation. All operations indicated must be performed without causing decoherence of existing atoms. (b) Our zone-based neutral-atom architecture enables arbitrary quantum computations to be performed through the use of separate static register, measurement and interaction zones. Atoms are moved between sites using movable tweezers, enabling arbitrary connectivity. A storage zone provides a nearby source of atoms to replace those lost in computation. The storage zone is periodically refilled mid-circuit from a loading zone, which is in turn filled from a MOT. (c) Continuous operation is achieved by repeating blocks consisting of several rounds of syndrome extraction, followed by reloading of the SZ. The roles of data and ancilla qubits are swapped and missing atoms are replaced during each syndrome extraction round. (d) A depiction of a twisted toric code across multiple cycles of syndrome extraction.  $X$ -type ancillas are shown on vertices,  $Z$ -type ancillas are shown on the faces, and data qubits are on the edges. The four colors show how the qubits swap from one cycle to the next. The star indicates a lost data qubit, that becomes a lost ancilla, which gets replaced after measurement. After completing a round of syndrome extraction, the erased qubit collapses into the wrong frame, triggering stabilizer violations shown with shaded plaquettes.

tations, an observed decrease in the logical error rate as the code distance is increased is an encouraging indication that a platform may reach useful error rates upon scaling. Strictly, the threshold of a code family is an asymptotic property: the physical error rate below which the logical error rate of an infinite family of codes can be driven arbitrarily low by increasing the code distance. Recent experiments instead probe a finite-size analogue, colloquially referred to as sub-threshold behavior—a decrease in the logical error rate between two or more code sizes. Such behavior requires that the total error rate associated with gate, idle, state preparation and measurement are sufficiently low that the added complexity of higher distance codes is offset by the increased redundancy. Subthreshold behavior has been demonstrated in platforms using superconducting qubits [29, 32], and improvements in post-selection fraction with code distance have been shown in a trapped ion system [33]. These demonstrations benefit from the low error rate for physical operations currently available in those modalities, though face certain limitations: superconducting platforms are currently constrained to nearest-neighbor connectivity which restricts them to “local” error-correcting codes, while trapped ion platforms have yet to demonstrate syndrome extraction beyond 8 cycles.

Indications of sub-threshold behavior—evidenced by lower logical error rate for a distance-5 surface code relative to distance-3 at a fixed number of syndrome extraction cycles—have recently been reported in a neutral-atom system, though without the use of mid-circuit measurements (blocks of ancilla qubits used for syndrome extraction were retained and measured after all gate operations were performed) [25]. In this work, we characterize the logical error rates for two distances of toric codes—one using 16 data qubits and 16 ancilla qubits and the other using 32 data qubits and 32 ancilla qubits—and find lower error rates for the larger distance variant when performing 4, 6, or 8 syndrome extraction cycles. Furthermore, by implementing mid-circuit reloading, we extend syndrome extraction up to 90 cycles, and observe that the logical error rates do not increase with code distance.

Finally, we demonstrate that logical information can be maintained in our system far longer than the lifetime of any physical qubit. Using a repetition code of distance 3 or 7 to periodically correct for errors occurring in interleaved wait-times, we show that logical information decays on a timescale of over 3 minutes for the larger distance, even when underlying qubits have lifetimes in the system on the order of 10 seconds.

## II. ARCHITECTURE

We use a zone-based quantum processor [11], with qubits defined by the ground nuclear-spin states of  $^{171}\text{Yb}$  atoms confined within arrays of optical tweezers. The functional zones are shown in Fig. 1(b), and a detailed description of the system can be found in [12, 38] and references therein. This platform enables execution of arbitrary universal quantum computation, with single-qubit gates performed in a register zone that is currently configured to support 128 qubits, two-qubit controlled- $Z$  ( $CZ$ ) gates performed in an interaction zone (IZ), and arbitrary connectivity provided by reconfigurable parallel movement of atoms from the register to IZ.

Beyond gates, this platform enables measurement with qubit reset, reuse, and replacement (MCM, described in [6]), alongside mid-circuit qubit reloading. MCM is performed by transferring atoms in a measurement zone (MZ) into a co-located cavity-enhanced 3D optical lattice [1] and performing non-destructive, state-selective imaging (using techniques described in [34]), as well as state reset and cooling. Atoms identified as missing are replaced from a near-by storage zone (SZ).

In this work, we add the ability to reload the SZ mid-circuit from a MOT located 30 cm below the tweezer array. At pre-defined points in a computation, atoms are transported from the MOT using a moving optical lattice, and extracted into a loading zone (LZ) where light-assisted collisions (LACs) are performed. As in our previous work [1], extraction is performed without the application of cooling light. During a subsequent MCM, the locations of atoms are identified in the LZ, MZ and SZ, and atoms from the LZ are used to refill both the MZ and SZ before circuit execution resumes see Fig. 1(c). MOT loading and lattice transport are performed in parallel with other operations, so the only time overhead to reload comes from the LACs and the time required to move atoms from the LZ to SZ. Details of the reloading sequence can be found in the Supplementary Material.

## III. LOSS-TOLERANT TWISTED TORIC CODE MEMORY

The toric code is a foundational topological quantum error-correcting code that stores information on a two-dimensional periodic lattice (represented in Fig. 1(d)), making its encoded information intrinsically robust against local perturbations [39, 40]. The toric code and the related surface code are among the most well-studied quantum error correcting codes and exhibit many favorable properties for practical fault-tolerant quantum computation [41–43].

It has been theoretically shown that twisting the torus on which the toric code is defined modifies its geometry while preserving the local stabilizer structure, giving access to a family of code variants with reduced qubit overhead and tunable asymmetric distances [18, 44]. As in

the standard construction, the code is defined on the two-dimensional square lattice associating qubits with edges,  $X$ -checks with vertices, and  $Z$ -checks with faces. Following Ref. [18], we specify a twisted torus via a  $2 \times 2$  integer matrix  $M$  in Hermite normal form. A non-zero upper right entry of  $M$  describes a shear to the square lattice, introducing the twist. See Supplementary Material for a more detailed description.

We performed a search over twisted toric codes for small instances compatible with our array geometry and movement constraints that support favorable circuit-level logical error rates. We focus on two codes specified by:

$$M_{\text{det } 8} = \begin{bmatrix} 1 & 2 \\ 0 & 8 \end{bmatrix}, \quad M_{\text{det } 16} = \begin{bmatrix} 1 & 3 \\ 0 & 16 \end{bmatrix}.$$

These codes support two logical qubits, but by treating one of the logical qubits as a gauge qubit we obtain subsystem codes with parameters

$$\begin{aligned} \text{det } 8 &: [[16, 1, (4, 3)]] \\ \text{det } 16 &: [[32, 1, (6, 4)]] \end{aligned}$$

where we indicate the asymmetric distances ( $d_X, d_Z$ ). Note that the gauge qubit interchanges the  $X$ - and  $Z$ -distances. We refer to the logical qubit operators as  $\bar{X}_1$  and  $\bar{Z}_1$  and the gauge qubit operators as  $\bar{X}_{\text{gauge}}$  and  $\bar{Z}_{\text{gauge}}$ . The lattice has a spatial self-duality which maps  $\bar{Z}_1 \leftrightarrow \bar{X}_{\text{gauge}}$  and  $\bar{Z}_{\text{gauge}} \leftrightarrow \bar{X}_1$ . In a single  $Z$ -type memory experiment, we measure both  $\bar{Z}_1$  and  $\bar{Z}_{\text{gauge}}$ . By the lattice self-duality, the latter is equivalent to an  $X$ -type memory experiment for the logical qubit. Thus the same experimental sequence gives access to both the  $Z$ - and  $X$ -type memory performance of the logical qubit.

The syndrome-extraction circuit (see Supplementary Material for a diagram) interleaves  $X$ - and  $Z$ -type stabilizers to avoid hook errors so that our circuits preserve code distance [40]. A distance  $d_{X/Z}$  guarantees the ability to correct  $X/Z$  errors of weight  $t_{X/Z} = \lfloor (d_{X/Z} - 1) / 2 \rfloor$ . We have  $t_X = 1$  for the det 8 code and  $t_X = 2$  for the det 16 code. Both codes have  $t_Z = 1$  and therefore offer similar protection to the  $X$ -type observable against Pauli noise, while det 16 offers stronger protection to the  $Z$ -type observable. On the other hand, since heralded erasures are correctable up to weight  $d - 1$ , the larger  $d_Z$  and  $d_X$  of the det 16 code affords both observables strictly better protection against atom loss.

In order to mitigate the effects of atom heating and loss during repeated cycles of syndrome extraction, we incorporate role-swapping of data and ancilla qubits, as well as reset and replenishment of ancillas after measurement. By compiling periodic swap operations into the syndrome extraction procedure [38, 45, 46], as depicted in Fig. 2(a), encoded information is preserved while allowing each atom to be reset after a limited number of operations—a typical atom participates in two syndrome extraction cycles before being measured, reset, and if necessary, replaced.

Detectors represent inconsistent correlations between syndromes of consecutive cycles [47], and the detection

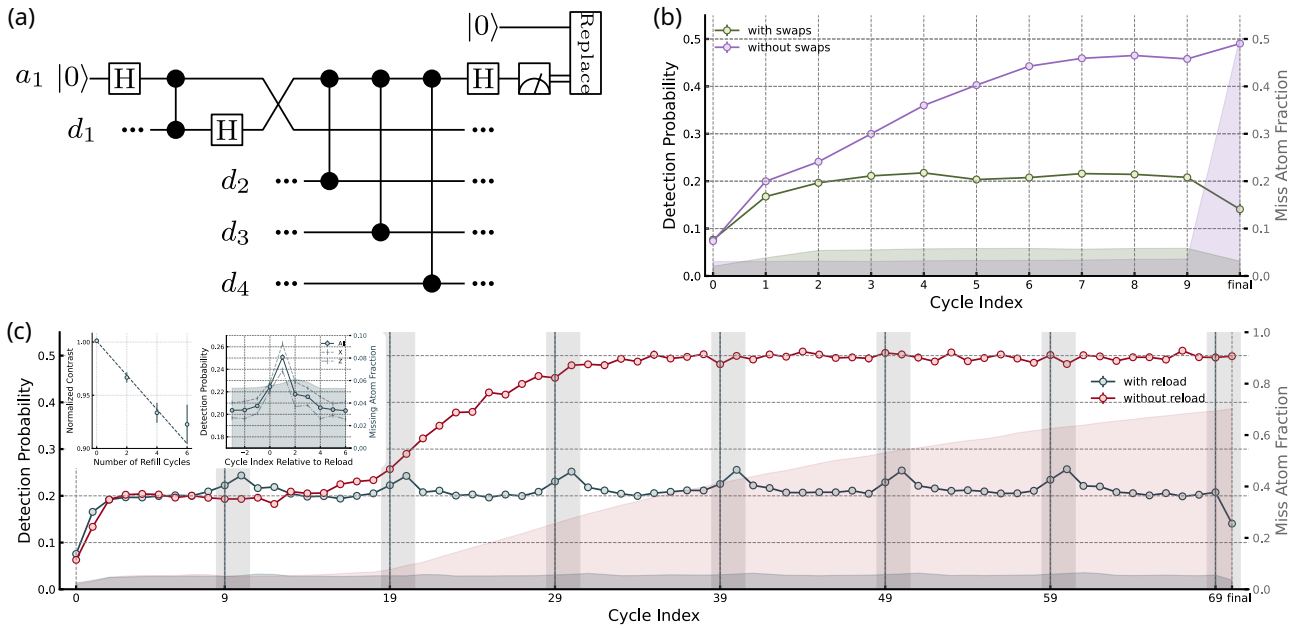


FIG. 2. Toric code memory using qubit reset, reuse, role-swap, and replenishment. (a) A single  $Z$ -type syndrome extraction, with a compiled swap operation. An incoming data qubit interacts with a fresh qubit via a CZ gate, after which their roles have been exchanged. After mid-circuit measurement, the circuit branches conditionally; a detected atom loss event triggers the replacement of the vacant ancilla site with a fresh qubit from the storage zone (SZ). (b) Stabilizer detection probability under ancilla-data swaps. Performances of a det 8 Toric code over 10 consecutive syndrome extraction cycles with and without role swapping are shown in green and purple respectively. Error bars represent the standard error of the mean (SEM) and are smaller than the marker size. Shaded regions show the corresponding missing atom fractions per cycle. The final loss value of near 50% without swaps indicates loss of nearly all data qubits. (c) Long-term toric code memory with and without SZ refilling. A det 16 toric code detection probability over 70 consecutive syndrome extraction cycles with SZ reloading every 10 cycles (olive) is compared to the no reloading baseline (red). Shaded regions track the missing atom fraction. Left inset shows the normalized MCM contrast over consecutive refilling cycles; right inset shows the cycle-averaged detection probability and missing atom fraction relative to the 10-cycle refill interval.

probability  $P_d$  can be used to characterize the temporal stability of physical error rates throughout the execution of the circuit [29, 48]. Our imaging method used for qubit measurement independently identifies qubit state and loss, enabling us to perform delayed erasure conversion [24, 49]. We handle loss events by randomly assigning a state prior to computing detection probabilities.

#### IV. ARBITRARILY REPEATABLE SYNDROME EXTRACTION

The combination of role-swapping, ancilla reuse and replacement, and reloading of atoms from the MOT combine to enable continuous syndrome extraction for arbitrary numbers of rounds. In Fig. 2(b), we compare  $P_d$  with or without role-swapping when operating the det 8 toric code over 10 cycles of syndrome extractions. In the static configuration with fixed data and ancilla roles,  $P_d$  rises steeply due to heating and loss of data qubits. In contrast, with the swap protocol,  $P_d$  remains near a constant baseline of  $\approx 0.2$  (the cycles indexed 0 and 1, as well as the final cycle are expected to show lower  $P_d$  due to edge-effects). The elevated ancilla qubit loss ob-

served in the role-swapping configuration (see green and purple shades in Fig. 2(b)) results from the fact that a measured atom experiences 8 two-qubit gates between measurements with role-swapping, compared to 4 two-qubit gates without. This demonstration shows the effectiveness of resetting and periodically replacing atoms as enabled by the role-swapping protocol.

While role-swapping and ancilla replacement prevents the buildup of heating and loss errors on specific atoms, this process will eventually deplete the occupancy of the SZ. To mitigate this, we employ mid-circuit reloading of the SZ once every 10 cycles of syndrome extraction. In Fig. 2(c), we track  $P_d$  for a det 16 toric code over 70 cycles of syndrome extraction (the total number of cycles is currently limited by compilation overhead). Without any atom replenishment to SZ (red points, shaded region), we observe a rapid increase in  $P_d$  as accumulated atom losses can no longer be addressed upon depletion of SZ, occurring here around cycle 15. In contrast, the ‘with-refill’ case (gray points, shaded region) maintains a steady-state  $P_d$  indefinitely, albeit with periodic increase of a few % during the refill cycles (indicated by vertical dashed lines in Fig. 2(c)).

The elevated detector frequency can be caused by de-

coherence induced by the reloading process. We estimate the magnitude of the decoherence by comparing the contrast of two Ramsey sequences: a baseline condition with MCM and a hold-time equal to the det 16 syndrome extraction cycle time containing a single echo pulse, and the same sequence with a full reloading procedure included in parallel with the hold time. The relative contrast versus the number of blocks included is shown in the inset of Fig. 2(c), and indicates an added 1.6(1)% contrast loss per refilling cycle. While this is a coherent error that will accumulate differently depending on what operations are performed in parallel with reloading, it is on the right scale to explain the bumps in detection probability.

The elevated detection probability typically persists for 3 cycles, with no significant increase ( $< 1\%$ ) in the missing atom fraction (see first inset of Fig. 2(c)). The first two cycles of elevated  $P_d$  are a direct consequence of the differential nature of the detector (indicated by vertical shades in Fig. 2(c)), which compares parity outcomes between successive cycles. The observation of two additional elevated cycles indicates a brief entropy dissipation period, during which repeated mid-circuit measurement (MCM) and reset operations remove the injected noise.

## V. LOGICAL ERRORS SUPPRESSION WITH CODE DISTANCE

We next characterize the logical error rate for the two code variants after different numbers of syndrome extraction cycles. As described above, lattice self-duality allows a single  $Z$ -type memory experiment to access logical error rates for both the  $Z$ - and  $X$ -type logical observables of one logical qubit: the  $Z$  logical observable is measured directly, while the complementary  $X$  logical observable is obtained through the dual  $Z$ -type observable of the gauge qubit. The final cycle uses a transversal measurement of the ancilla and data qubits to obtain the two final rounds of detectors as well as measurements of the logical observables.

The recorded bitstrings were decoded using PyMatching, a minimum-weight perfect matching decoder [50, 51]. We use Stim to construct an initial detector error model [52, 53]. On a per-shot basis, we use a heuristic method to re-weight the prior in the detector-error model around atom loss events (further details are provided in the Supplementary Material). Notably, this procedure requires a significantly reduced computational overhead compared to supercheck decoders [54, 55].

In Fig. 3, we present the logical error rate for both codes in regimes that do and do not include mid-circuit reloading from a MOT. All conditions involve role-swapping, MCM with reset, and replacing lost ancilla qubits from the SZ.

Memory experiments out to 8 cycles do not require reloading the storage zone (SZ). To ensure consistent machine performance for the two codes, 35 shots of the det 8 codes were interleaved with 20 shots of the det 16 codes

across 6 separate datasets totaling approximately 180 hours of data collection ( $\approx 39,000$  shots per det 8 point and  $\approx 23,500$  shots per det 16 point). Fig. 3(a) shows the decoded logical error rates as a function of cycles with overlaid fits to a fixed logical error per cycle,  $\epsilon_{\text{cycle}}$  model [25, 29] with parameters  $\epsilon_{\text{cycle,det8}}^Z = 0.74 \pm .01\%$  and  $\epsilon_{\text{cycle,det16}}^Z = 0.56 \pm .01\%$ . The det 16 data do not appear to follow a fixed logical error rate per cycle in this range, likely due to extended boundary effects from the high initial state preparation fidelity and increased loss with cycles due to slow atom heating (see additional data and simulations in the Supplementary Material). We note that directly converting the 4-cycle point to a logical error rate per cycle as in [25] yields  $\epsilon_{\text{cycle,det16}}^Z = 0.38 \pm .02\%$  and an error suppression factor  $\Lambda_{\text{no reload}}^Z = 1.9 \pm .1$ . Similar behavior is observed in the  $X$ -observable at low cycles (Fig. 3, right), but with significantly higher logical error rates and a smaller error suppression factor (again at 4 cycles)  $\Lambda_{\text{no reload}}^X = 1.20 \pm .03$ . Averaging over the two observables, we observe  $\Lambda_{\text{no reload}}^{\text{Avg}} = 1.30 \pm .03$ .

While the observed decrease in logical error rate for the higher-distance code is consistent with sub-threshold behavior in this regime, several aspects of the result suggest care in extrapolating these results. First, our error probability exhibits trends versus numbers of cycles that are not well-captured by an error model that neglects a careful treatment of coherent errors and atom heating. This observation merits caution in extrapolating per-cycle error rates from measurements performed at a limited sampling of cycle numbers, even if backed by a stochastic error model. Second, we cannot rule out the possibility of per-operation error rates that grow with system size. While this is a somewhat generic statement about hardware, it may be more relevant for systems that take advantage of arbitrary connectivity.

Reloading the SZ every 10th cycle extends the memory experiment out as far as 70 cycles for the det 16 code and 90 cycles for the det 8 code (Fig. 3 (b); again, depth is currently limited by compilation overhead). The two codes were interleaved at 10 shots each within a single dataset spanning 67 hours ( $\approx 660$  shots per det 8 point and  $\approx 710$  shots per det 16 point). We find reasonable agreement between data with reloading and the fixed logical error rate per cycle model. The resulting fits yield  $\epsilon_{\text{cycle,det8}}^Z = 0.63 \pm .03\%$  and  $\epsilon_{\text{cycle,det16}}^Z = 0.64 \pm .04\%$ . This demonstrates that over an arbitrary number of cycles, the logical error rates do not increase with larger code distance. Importantly, atom reloading does not appear to cause a significant increase in logical error rates per cycle.

We simulate errors in our system using a model that accounts for independently calibrated stochastic errors as well as loss (see Supplementary Material for a description of the simulation and input parameters). As a reference case, we consider the  $Z$ -type logical error rates after six cycles here—a broader comparison between the model and observed logical error rates, detector frequencies, and loss is explored in the Supplementary Material. Our

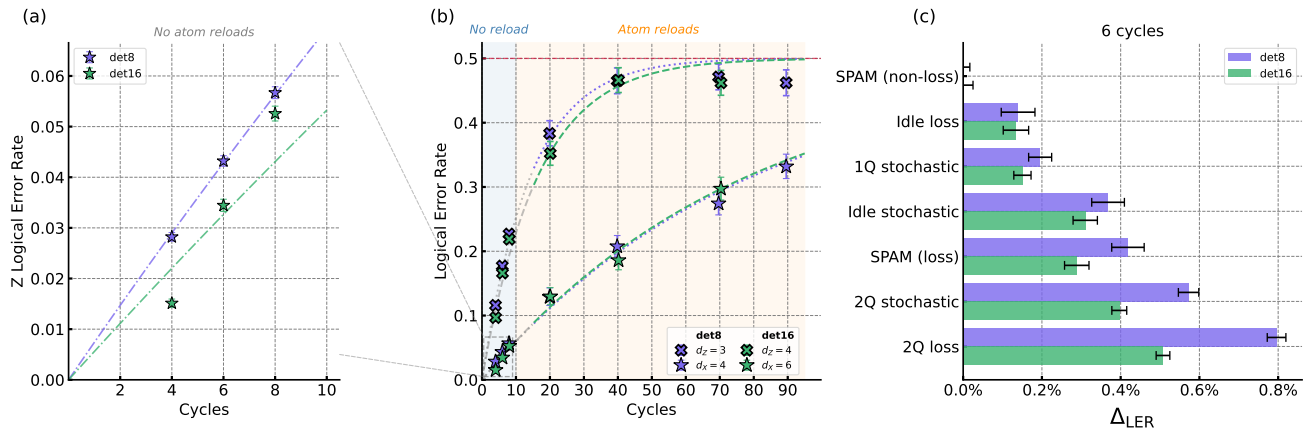


FIG. 3. (a) Z-type logical error rates without SZ reloads averaged over 6 datasets. The det 16 code shows logical error rates at or below those of the det 8 code, even through fluctuations over time within this dataset (see Supplementary Material). Lines are fits to a fixed logical error per cycle model to guide the eye,  $\text{LER}(c) = \frac{1}{2}(1 - (1 - 2\epsilon_{\text{cycle}})^c)$  [25, 29] with parameters  $\epsilon_{\text{cycle,det8}}^Z = 0.74 \pm .01\%$  and  $\epsilon_{\text{cycle,det16}}^Z = 0.56 \pm .01\%$ . The det 16 data shows significant deviations from this model likely due to boundary detector effects and slow atom heating (see Supplementary Material). (b) Logical error rate per cycle for both the Z- and X-type extended to deep cycles by reloading atoms to the SZ every 10th cycle. The X-type observable is measured through the Z-type observable of the gauge qubit. Dotted purple (det8) and dashed green (det16) lines are fits including only data with reloads to the fixed logical error per cycle model with parameters  $\epsilon_{\text{cycle,det8}}^Z = 0.63 \pm .03\%$ ,  $\epsilon_{\text{cycle,det8}}^X = 3.4 \pm .3\%$ ,  $\epsilon_{\text{cycle,det16}}^Z = 0.64 \pm .04\%$  and  $\epsilon_{\text{cycle,det16}}^X = 3.0 \pm .2\%$ . Atom reloading does not appear to have a significant effect on logical error rate per cycle. (c) Z-type error gradients for the det 8 ( $d_X = 4$ , purple) and det 16 ( $d_X = 6$ , green) codes at six cycles simulated from an independently calibrated error model (see Supplementary Material).

model predicts Z-type error rates of 1.1% for det 16 and 1.9% for det 8, somewhat lower than the average observed values of Fig. 3, but similar to the best-performing time block of the data (see Supplementary Material). This difference may be attributed to potential coherent errors which are not captured in our model (for example, uncalibrated qubit frequency differences or drifts that are imperfectly canceled by spin-echo techniques), or to difficulties in accurately calibrating the in-situ performance of operations, where effects like atom heating may be relevant. Incorporating additional data qubit loss in the model partially resolves this discrepancy and better captures certain features of detector frequencies and loss (see Supplementary Material).

In Fig. 3(c) we represent the simulated sensitivity of Z-type logical error rates after six cycles to the calibrated physical error sources by plotting the simulated reduction in the logical error rate when that error source is removed entirely. Because logical error rates are not linear versus physical error rates, we do not expect the component sensitivities to add up to the observed logical error rate, but do expect the relative values to be representative of the marginal contribution of the associated error source. Two features are apparent from this analysis: the 2Q gates contribute the largest single source of error, and the various channels for loss, taken together, exceed the contribution from stochastic errors.

## VI. A ROBUST REPETITION CODE MEMORY SURPASSING PHYSICAL QUBIT DECAY

One of the defining milestones of fault-tolerant quantum computing is the creation of a logical qubit with an operational lifetime that surpasses the decay of its physical components, which is essential for preserving quantum information indefinitely. In neutral-atom platforms, the operational upper bound is fundamentally governed by physical atom loss from the optical traps. To establish a rigorous operational milestone, we benchmark our encoded logical memory directly against this physical atom survival baseline. We characterize idle atom survival independently within both the science zones (IZ, MZ and SZ) and register using a gaussian decay fit, extracting baseline atom  $1/e$  lifetimes to be 10.5(4) s and 15.6(5) s, respectively (Fig. 4(a) top). We measure the logical memory of distance 3 ( $d = 3$ ) and distance 7 ( $d = 7$ ) phase-sensitive repetition codes [38] up to 95 QEC cycles with MCM and atom replacement every cycle, complemented by an SZ refill routine every 20 cycles (Fig. 4(a) bottom). We intentionally introduce extra wall-clock idling time ( $\tau = 0.5$ s) into each individual cycle (see Fig. 4(b)), artificially stretching the total circuit duration to expose the encoded logical states to maximum environmental error. The value  $\tau = 0.5$  s was chosen as a practical operating point, balancing increased idle error against logical error suppression per unit time; it is not intended to be an optimized delay time. Remarkably, the logical error rates for both  $d = 3$  and  $d = 7$  repetition

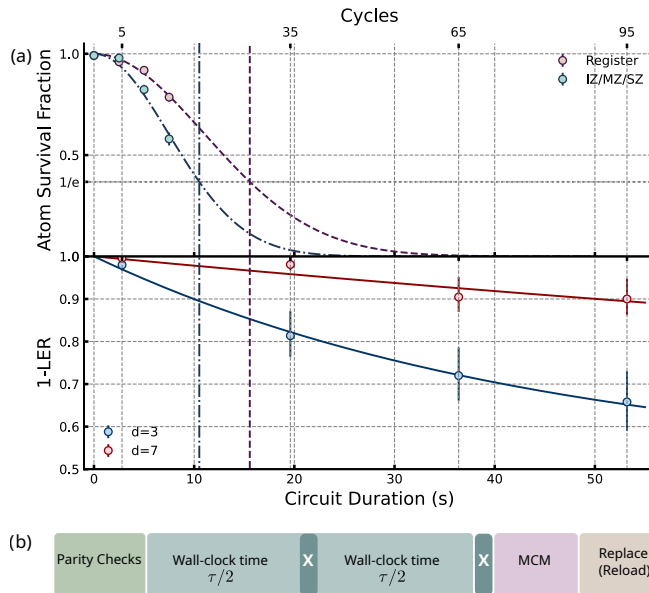


FIG. 4. Sustained Quantum Memory beyond physical qubit lifetime. (a) Repetition code performance over extended operational times. The top panel shows physical atom survival fractions over time in the register (purple) and science zones (IZ, MZ, and SZ; green). The bottom panel displays logical error rates for the  $d = 3$  (red) and  $d = 7$  (blue) repetition codes, plotted against dual axes mapping both the total number of QEC cycles and cumulative circuit durations. (b) Timing diagram of a QEC cycle extended by wall-clock idling time. Following the stabilizer parity checks, qubits idle in the register for a duration  $\tau/2$ , after which an echo  $\pi$ -pulse is applied. Following a second  $\tau/2$  idling window, a final  $\pi$ -pulse restores the qubits to their original states, immediately followed by MCM and atom replacement, and SZ reloading every 20th cycle.

codes remain strictly below the completely random 0.5 limit well beyond the physical lifetime of atoms held in tweezers. To quantify this behavior, we plot  $1 - \text{LER}$  and characterize the characteristic logical lifetime by fitting the decay profile to  $0.5 + 0.5 \exp(-t/\tau_L)$ , from which we extract logical lifetimes of  $\tau_L = 44(1)$  s for the  $d = 3$  and  $\tau_L = 225(33)$  s for  $d = 7$  repetition codes respectively, for the specific idle times chosen. This persistent error suppression over macroscopic operational duration demonstrates a sustained logical memory that outlasts the physical presence of individual atoms.

## VII. OUTLOOK

In this work we have demonstrated error suppression in a toric code quantum memory that integrates all of the physical components necessary for continuous operation of an error-corrected neutral atom quantum processor for the first time. Critically, we have demonstrated a syndrome extraction scheme that frequently swaps the identity of data and ancilla qubits such that lost qubits are quickly identified as an erasure errors and replaced by freshly initialized qubits. This scheme is generalizable to other QEC codes that will be used in utility-scale quantum computers. To further demonstrate the viability of this platform, we operated a repetition code logical memory maintaining encoded information up to one minute,

well beyond the lifetime of the atomic qubits. The reduction in logical error rate with increased code distance is consistent with sub-threshold performance of the quantum processor. Making use of the flexible connectivity of neutral atom qubits, this is the first demonstration of continuous syndrome extraction in a code that requires non-local connectivity inaccessible in a 2D planar geometry and enables practical utility-scale codes.

Moving from these demonstrations to useful, utility-scale quantum computing requires advancement on multiple fronts. We anticipate that improved two-qubit gate performance (particularly loss) can be achieved through the use of faster gate schemes that are less sensitive to atomic temperature, which in turn will enable faster atomic movement at acceptable heating levels and lower associated idle errors as well as cycle times. The system size may be increased to utilize more advanced error correction codes with larger distance and logical operations (Clifford and non-Clifford) with efficient space-time cost. This will require both increased laser power and the ability to maintain flexible parallelization of atom movement—for example with multiple independent sets of mobile tweezers. Further, while the moderate system size and low-loss operations demonstrated in this work mean that steady-state operation can be achieved with a modest flux of atoms, increased system size and decreased logical cycle times may benefit from performing either multiple extractions of atoms from the transport lattice per loading cycle [5, 35].

## Atom Computing and Collaborators

David Aasen<sup>2</sup>, Alexander Aeppli<sup>1</sup>, Stephen Armstrong<sup>1</sup>, Sambit Banerjee<sup>1</sup>, Katrina Barnes<sup>1</sup>, Ethan Becker<sup>1</sup>, Juan M. Bello-Rivas<sup>2</sup>, Trent Bjorkman<sup>1</sup>, Benjamin J. Bloom<sup>1</sup>, Ian Bloomfield<sup>1</sup>, Thomas Bohdanowicz<sup>1,\*</sup>, Graham Booth<sup>1</sup>, Mahdi Bornadel<sup>1</sup>, Andrew Brown<sup>1</sup>, Mark Brown<sup>1</sup>, Jonathan Bruce<sup>1</sup>, William Cairncross<sup>1</sup>, Samuel Carman<sup>1</sup>, Cheng An Chen<sup>1</sup>, Jonathan Chen<sup>1</sup>, Grace Cowan<sup>1</sup>, Daniel Crow<sup>1</sup>, Markus DeMartini<sup>1</sup>, Jeffrey Epstein<sup>1</sup>, Benjamin Familetto<sup>1</sup>, Wentao Fan<sup>1</sup>, Max Feldkamp<sup>1</sup>, Adam Friss<sup>1</sup>, Christopher Griger<sup>1</sup>, Evan Hanson<sup>1</sup>, Youssef Hassan<sup>1</sup>, Ben Heberlein<sup>1</sup>, Andre Heinz<sup>1</sup>, Elayne Heynen<sup>1</sup>, Theodore Hoff<sup>1</sup>, Thomas Hofler<sup>1</sup>, Samuel Hotaling<sup>1</sup>, Frederic Hummel<sup>1</sup>, Andrei Iskra<sup>1</sup>, Matthew Jaffe<sup>1, 4</sup>, Yungpeng Ji<sup>1,\*</sup>, Gwendolyn Johnson<sup>1</sup>, Antonia Jones<sup>1</sup>, Eliot Kapit<sup>1</sup>, Robert Keus<sup>1</sup>, Hyosub Kim<sup>1</sup>, Jonathan Kindem<sup>1</sup>, Jonathan King<sup>1</sup>, Samuel Korn<sup>1</sup>, Zachary Lane<sup>1</sup>, Shaoheng Li<sup>1</sup>, Justin Lim<sup>1</sup>, De Luo<sup>1</sup>, Jan Marjanovic<sup>1</sup>, Julia McMaster<sup>1</sup>, Eli Megidish<sup>1,\*</sup>, Matthew Meredith<sup>1</sup>, Tanya Mittal<sup>1</sup>, Ryan Morshead<sup>1</sup>, Ciro Nishiguchi<sup>1</sup>, Eyal Niv<sup>1</sup>, Matthew Norcia<sup>1</sup>, Botond Oreg<sup>1</sup>, Adway Patra<sup>1,\*</sup>, David Rodriguez Perez<sup>1</sup>, Matthew Peters<sup>1</sup>, Fischer Pettner<sup>1</sup>, Sebastian Pucher<sup>1</sup>, Kristen Pudenz<sup>1</sup>, Brenden Roberts<sup>1,\*</sup>, Nicholas Rogers<sup>1</sup>, Melody Rouault<sup>1</sup>, Albert Ryou<sup>1</sup>, Parth Sabharwal<sup>1</sup>, Sven Schmidt<sup>1</sup>, Calvin Schwadron<sup>1</sup>, Ian Scott<sup>1</sup>, Juan Andres Muniz Silva<sup>1</sup>, Elaine Silverman<sup>1</sup>, Jon Simon<sup>1, 3</sup>, Michael Sorensen<sup>1</sup>, Daniel Stack<sup>1</sup>, Mark Stone<sup>1</sup>, Peter Stromberger<sup>1</sup>, Ioan Sturzu<sup>1</sup>, Lavanya Taneja<sup>1</sup>, Raanan Tobey<sup>1</sup>, Adam Turflinger<sup>1,\*</sup>, Miroslav Urbanek<sup>1</sup>, Gerard Valenti-Rojas<sup>1</sup>, Eric Victorson<sup>1</sup>, Laura Wadleigh<sup>1</sup>, Yiping Wang<sup>1</sup>, Jonathan Ward<sup>1</sup>, Robert Weverka<sup>1</sup>, Kennard White<sup>1</sup>, Anthony Wiese<sup>1</sup>, Ayesha Wilson<sup>1</sup>, Tsung-Yao Wu<sup>1,\*</sup>, Evan Zalys-Geller<sup>1</sup>, Xiaogang Zhang<sup>1</sup>,

\* These authors contributed to data taking, analysis and simulation for this work.

<sup>1</sup> Atom Computing, Inc.

<sup>2</sup> Microsoft Quantum

<sup>3</sup> Department of Physics and Department of Applied Physics, Stanford University

<sup>4</sup> Department of Physics, Montana State University

- 
- [1] M. Norcia, H. Kim, W. Cairncross, M. Stone, A. Ryou, M. Jaffe, M. Brown, *et al.*, Iterative assembly of 171 yb atom arrays with cavity-enhanced optical lattices, *PRX Quantum* **5**, 030316 (2024).
- [2] R. Tao, M. Ammenwerth, F. Gyger, I. Bloch, and J. Zeiher, High-fidelity detection of large-scale atom arrays in an optical lattice, *Physical Review Letters* **133**, 013401 (2024).
- [3] H. J. Manetsch, G. Nomura, E. Bataille, X. Lv, K. H. Leung, and M. Endres, A tweezer array with 6,100 highly coherent atomic qubits, *Nature* **647**, 60 (2025).
- [4] R. Lin, H.-S. Zhong, Y. Li, Z.-R. Zhao, L.-T. Zheng, T.-R. Hu, H.-M. Wu, Z. Wu, W.-J. Ma, Y. Gao, *et al.*, Ai-enabled parallel assembly of thousands of defect-free neutral atom arrays, *Physical Review Letters* **135**, 060602 (2025).
- [5] N.-C. Chiu, E. C. Trapp, J. Guo, M. H. Abobeih, L. M. Stewart, S. Hollerith, P. L. Stroganov, M. Kalinowski, A. A. Geim, S. J. Evered, *et al.*, Continuous operation of a coherent 3,000-qubit system, *Nature* **646**, 1075 (2025).
- [6] J. A. Muniz, M. Stone, D. T. Stack, M. Jaffe, J. M. Kindem, *et al.*, High-fidelity universal gates in the  $^{171}\text{Yb}$  ground-state nuclear-spin qubit, *PRX Quantum* **6**, 020334 (2025).
- [7] R. B.-S. Tsai, X. Sun, A. L. Shaw, R. Finkelstein, and M. Endres, Benchmarking and linear response modeling of high-fidelity Rydberg gates (2024), arXiv:2407.20184.
- [8] 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  $^{171}\text{Yb}$  rydberg states for high-fidelity two-qubit gates, *Phys. Rev. X* **15**, 011009 (2025).
- [9] A. Senoo, A. Baumgartner, J. W. Lis, G. M. Vaidya, Z. Zeng, G. Giudici, H. Pichler, and A. M. Kaufman, High-fidelity entanglement and coherent multi-qubit mapping in an atom array, arXiv preprint arXiv:2506.13632 (2025).
- [10] S. J. Evered, M. Xu, S. H. Li, A. A. Geim, J. Ataides, M. Kalinowski, D. Bluvstein, N. Maskara, C. Kokail, M. Greiner, *et al.*, High-fidelity entangling gates and nonlocal circuits with neutral atoms, arXiv preprint arXiv:2604.25987 (2026).
- [11] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, *et al.*, A quantum processor based on coherent transport of entangled atom arrays, *Nature* **604**, 451 (2022).
- [12] B. W. Reichardt, A. Paetznick, D. Aasen, I. Basov, J. M. Bello-Rivas, *et al.*, Logical computation demonstrated with a neutral atom quantum processor, arXiv preprint arXiv:2411.11822 (2024).
- [13] R. Chinnarasu, C. Poole, L. Phuttitarn, A. Noori, T. Graham, S. Coppersmith, A. Balantekin, and M. Saffman, Variational simulation of the lipkin-meshkov-glick model on a neutral atom quantum computer, *PRX Quantum* **6**, 020350 (2025).
- [14] C. Zhao, C. Duckering, A. Gu, N. Maskara, and H. Zhou, Towards ultra-high-rate quantum error correction with reconfigurable atom arrays, arXiv preprint arXiv:2604.16209 (2026).
- [15] Q. Xu, H. Zhou, G. Zheng, D. Bluvstein, J. P. B. Ataides, M. D. Lukin, and L. Jiang, Fast and parallelizable logical computation with homological product codes, *Physical Review X* **15**, 021065 (2025).
- [16] W. Yang, J. Chadwick, M. H. Teo, J. Vizslai, and F. Chong, Spacetime-efficient and hardware-compatible complex quantum logic units in qldpc codes, arXiv preprint arXiv:2602.14273 (2026).
- [17] P. Webster, L. Berent, O. Chandra, E. T. Hockings, N. Baspin, F. Thomsen, S. C. Smith, and L. Z. Cohen, The pinnacle architecture: Reducing the cost of breaking rsa-2048 to 100 000 physical qubits using quantum ldpc codes, arXiv preprint arXiv:2602.11457 (2026).
- [18] D. Aasen, J. Haah, M. B. Hastings, and Z. Wang, Geometrically enhanced topological quantum codes (2025), arXiv:2505.10403 [quant-ph].
- [19] H. Bombín, Single-shot fault-tolerant quantum error correction, *Physical Review X* **5**, 031043 (2015).
- [20] M. Cain, D. Bluvstein, C. Zhao, S. Gu, N. Maskara, M. Kalinowski, A. A. Geim, A. Kubica, M. D. Lukin, and H. Zhou, Fast correlated decoding of transversal logical algorithms, arXiv preprint arXiv:2505.13587 (2025).
- [21] S. Sunami, A. Goban, and H. Yamasaki, Transversal surface-code game powered by neutral atoms, arXiv preprint arXiv:2506.18979 (2025).
- [22] H. Zhou, C. Zhao, M. Cain, D. Bluvstein, N. Maskara, C. Duckering, H.-Y. Hu, S.-T. Wang, A. Kubica, and M. D. Lukin, Low-overhead transversal fault tolerance for universal quantum computation, *Nature* **646**, 303 (2025).
- [23] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, *et al.*, Logical quantum processor based on reconfigurable atom arrays, *Nature* **626**, 58 (2024).
- [24] B. Zhang, G. Liu, G. Bornet, S. P. Horvath, P. Peng, S. Ma, S. Huang, S. Puri, and J. D. Thompson, Leveraging erasure errors in logical qubits with metastable 171yb atoms, arXiv preprint arXiv:2506.13724 (2025).
- [25] D. Bluvstein, A. A. Geim, S. H. Li, S. J. Evered, J. P. Bonilla Ataides, G. Baranes, A. Gu, T. Manovitz, M. Xu, M. Kalinowski, *et al.*, A fault-tolerant neutral-atom architecture for universal quantum computation, *Nature* **649**, 39 (2026).
- [26] O. Lib, H. Timme, M. Ammenwerth, F. Gyger, R. Tao, S. Sun, I. Bloch, and J. Zeiher, Velocity-enabled quantum computing with neutral atoms, arXiv preprint arXiv:2603.15561 (2026).
- [27] D. Gottesman, Stabilizer codes and quantum error correction (1997), arXiv:quant-ph/9705052 [quant-ph].
- [28] G. Q. AI and collaborators, Suppressing quantum errors by scaling a surface code logical qubit, *Nature* **614**, 676 (2023).
- [29] G. Q. AI and collaborators, Quantum error correction below the surface code threshold, *Nature* **638**, 920 (2025).
- [30] A. Eickbusch, M. McEwen, V. Sivak, A. Bourassa, J. Atalaya, J. Claes, D. Kafri, C. Gidney, C. W. Warren, J. Gross, *et al.*, Demonstration of dynamic surface codes, *Nature Physics* , 1 (2025).
- [31] N. Lacroix, A. Bourassa, F. J. Heras, L. M. Zhang, J. Bausch, A. W. Senior, T. Edlich, N. Shutty, V. Sivak, A. Bengtsson, *et al.*, Scaling and logic in the colour code on a superconducting quantum processor, *Nature* **645**, 614 (2025).

- [32] A. Vezvaei, C. Benito, M. Morford-Oberst, A. Bermudez, and D. A. Lidar, Surface code scaling on heavy-hex superconducting quantum processors, arXiv preprint arXiv:2510.18847 (2025).
- [33] S. Dasu, M. DeCross, A. Y. Guo, A. Lavasani, J. Behrends, A. Benhemou, Y.-H. Chen, K. Mayer, C. N. Self, S. Simsek, *et al.*, Computing with many encoded logical qubits beyond break-even, arXiv preprint arXiv:2602.22211 (2026).
- [34] M. A. Norcia, W. B. Cairncross, H. Kim, *et al.*, Mid-circuit qubit measurement and rearrangement in a  $^{171}\text{Yb}$  atomic array, *Phys. Rev. X* **13**, 041034 (2023).
- [35] Y. Li, Y. Bao, M. Peper, C. Li, and J. Thompson, Fast, continuous and coherent atom replacement in a neutral atom qubit array (2025), arXiv preprint arXiv:2506.15633.
- [36] K. Singh, C. 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).
- [37] F. Gyger, M. Ammenwerth, R. Tao, H. Timme, S. Snigirev, I. Bloch, and J. Zeiher, Continuous operation of large-scale atom arrays in optical lattices, *Physical Review Research* **6**, 033104 (2024).
- [38] J. Muniz, D. Crow, H. Kim, J. Kindem, W. Cairncross, A. Ryou, T. Bohdanowicz, C.-A. Chen, Y. Ji, A. Jones, *et al.*, Repeated ancilla reuse for logical computation on a neutral atom quantum computer, *Physical Review X* **15**, 041040 (2025).
- [39] A. Kitaev, Fault-tolerant quantum computation by anyons, *Annals of Physics* **303**, 2 (2003).
- [40] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *Journal of Mathematical Physics* **43**, 4452 (2002), [https://pubs.aip.org/aip/jmp/article-pdf/43/9/4452/19183135/4452.1\\_online.pdf](https://pubs.aip.org/aip/jmp/article-pdf/43/9/4452/19183135/4452.1_online.pdf).
- [41] G. Duclos-Cianci and D. Poulin, Fast decoders for topological quantum codes, *Phys. Rev. Lett.* **104**, 050504 (2010).
- [42] D. Horsman, A. G. Fowler, S. Devitt, and R. V. Meter, Surface code quantum computing by lattice surgery, *New Journal of Physics* **14**, 123011 (2012).
- [43] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
- [44] M. B. Hastings, J. Haah, and R. O'Donnell, Fiber bundle codes: breaking the  $n^{1/2}$  polylog( $n$ ) barrier for quantum ldpc codes, in *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021 (Association for Computing Machinery, New York, NY, USA, 2021) p. 1276–1288.
- [45] M. N. H. Chow, V. Buchemavari, 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).
- [46] P. Liu, S. J. S. Tan, E. Huang, U. A. Acar, H. Zhou, and C. Zhao, Achieving optimal-distance atom-loss correction via pauli envelope (2026), arXiv:2603.04156 [quant-ph].
- [47] A. G. Fowler, D. Sank, J. Kelly, R. Barends, and J. M. Martinis, Scalable extraction of error models from the output of error detection circuits, arXiv preprint arXiv:1405.1454 (2014).
- [48] I. Hesner, B. Hetényi, and J. R. Wootton, Using detector likelihood for benchmarking quantum error correction (2024), arXiv:2408.02082 [quant-ph].
- [49] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, Erasure conversion for fault-tolerant quantum computing in alkaline earth rydberg atom arrays, *Nature Communications* **13**, 4657 (2022).
- [50] A. G. Fowler, Optimal complexity correction of correlated errors in the surface code (2013), arXiv:1310.0863 [quant-ph].
- [51] O. Higgott and C. Gidney, Sparse Blossom: correcting a million errors per core second with minimum-weight matching, *Quantum* **9**, 1600 (2025).
- [52] C. Gidney, Stim: a fast stabilizer circuit simulator, *Quantum* **5**, 497 (2021).
- [53] P.-J. H. Derks, A. Townsend-Teague, A. G. Burchards, and J. Eisert, Designing fault-tolerant circuits using detector error models, *Quantum* **9**, 1905 (2025).
- [54] T. M. Stace, S. D. Barrett, and A. C. Doherty, Thresholds for topological codes in the presence of loss, *Physical review letters* **102**, 200501 (2009).
- [55] S. D. Barrett and T. M. Stace, Fault tolerant quantum computation with very high threshold for loss errors, *Phys. Rev. Lett.* **105**, 200502 (2010).
- [56] K. Barnes, P. Battaglini, B. J. Bloom, K. Cassella, R. Coxe, N. Crisosto, J. P. King, S. S. Kondov, K. Kotru, S. C. Larsen, *et al.*, Assembly and coherent control of a register of nuclear spin qubits, *Nature Communications* **13**, 2779 (2022).
- [57] J. A. Jones, Designing short robust not gates for quantum computation, *Phys. Rev. A* **87**, 052317 (2013).
- [58] G. Baranes, M. Cain, J. Ataides, D. Bluvstein, J. Sinclair, V. Vuletic, H. Zhou, and M. D. Lukin, Leveraging atom loss errors in fault tolerant quantum algorithms, arXiv preprint arXiv:2502.20558 (2025).
- [59] T. Müller, T. Alexander, M. E. Beverland, M. Bühler, B. R. Johnson, T. Maurer, and D. Vandeth, Improved belief propagation is sufficient for real-time decoding of quantum memory (2025), arXiv:2506.01779 [quant-ph].
- [60] A. Hashim, L. B. Nguyen, N. Goss, B. Marinelli, R. K. Naik, H. Mitchell, D. I. Santiago, T. Tomesh, M. Ippoliti, M. Kliesch, I. Roth, R. Harper, S. T. Flammia, J. J. Wallman, J. Emerson, I. Hincks, C. Granade, J. Combes, and I. Siddiqi, A practical introduction to benchmarking and characterization of quantum computers, *PRX Quantum* **6**, 030202 (2025).
- [61] M. A. Nielsen, A simple formula for the average gate fidelity of a quantum dynamical operation, *Phys. Lett. A* **303**, 249 (2002).
- [62] B. Clader, C. J. Trout, J. P. Barnes, K. Schultz, G. Quiroz, and P. Titum, Impact of correlations and heavy tails on quantum error correction, *Physical Review A* **103**, 052428 (2021).

## Supplemental Material

Atom Computing and Collaborators

Dated: June 2, 2026

### SI. SZ RELOADING SEQUENCE

Extended Data Fig. S1 illustrates the sequence for atom reload operations. Detailed physical implementations of each operation were described in previous work [S38]. The order of operations is adapted for parallel circuit execution, while their effects on the circuits are discussed in the main text.

The MOT and transport stage happens in parallel with the syndrome extraction (SE) circuit. Background loading of the MOT contributes an additional  $0.014(6)$  1/s to the overall decoherence rate. Before rearrangement moves atoms to the MZ for MCM, but after the ancilla atoms are brought to the measurement basis, a light-assisted collision (LAC) block prepares the LZ with a 50% average filling. The LAC beams are local to the LZ, IZ, MZ, and SZ. To echo out the residual light shift from the LAC beams to finite extinction in the RZ, a composite global  $\pi$  pulse is added between the two LAC blocks.

Inside the MCM block, a similar echo strategy is deployed between two MCM images. The first image is state-selective and the second is state-insensitive (optical pumping is applied) for atom-loss detection. Following the MCM block, all atoms in the LZ, MZ, and SZ are cooled and reinitialized. Subsequently, our software service uses the MCM images to determine occupancy and calculate the rearrangement moves required to refill both the SZ and MZ from the LZ, prioritizing the MZ. The overall duration of the refilling process depends on the number of sites to fill and their specific locations; on average, it takes approximately 90 ms, which includes the software processing time and a sequence buffer time of  $\sim 20$  ms.

To mitigate the effects of static qubit frequency differences or drifts, we deploy a dynamical decoupling echo at the midpoint of each SE circuit, when all active data and ancilla qubits occupy the static register sites.

### SII. 1Q GATES

Our system uses two forms of single-qubit gates: global and local. The local gates have been described in previous work [S12, S56], and use two pairs of crossed-acousto-optic deflectors to apply fully arbitrary gates to atoms in the register in a row-parallel manner. For this work, we have added the ability to perform global 1Q gates, where the same gate is applied to all atoms in parallel. This reduces the time overhead associated with parallelizing over rows, and is less sensitive to alignment-induced gate

quality fluctuations associated with the tightly focused local addressing beams. We use the global 1Q gates for spin-echo pulses associated with MCM, and the local 1Q gates for computational operations.

The global gates are driven with a single beam with two RF tones separated in frequency by the qubit transition frequency, with the polarization carefully tuned to both provide the two components needed to drive the Raman transition between qubit states and to balance the light shifts on both qubit states. Each component pulse features a rounded, flat-top profile with a single  $\pi$ -pulse duration of  $\sim 65 \mu\text{s}$ . For spin echo pulses, we use a five-pulse Knill-type configuration to mitigate amplitude and second-order off-resonance errors. For the specific phase sequence, see Table I (corresponding to Fig. 5b) of Ref. [S57]. We bound the error associated with the full five-pulse composite  $\pi$  pulse to  $6 \times 10^{-4}$  by fitting the decay in return-probability when we perform many rounds of interleaved composite and random operations, followed by a pre-computed return operation. The bound includes the contribution from the two non-composite  $\pi/2$  pulses used to achieve randomization in each round.

Because we do not perform a site-by-site reference-frame calibration between the global and local gates, mixing the two kinds of operation in a circuit requires that even numbers of global echo pulses are applied between any pair of local gates.

### SIII. TWISTED TORIC CODES

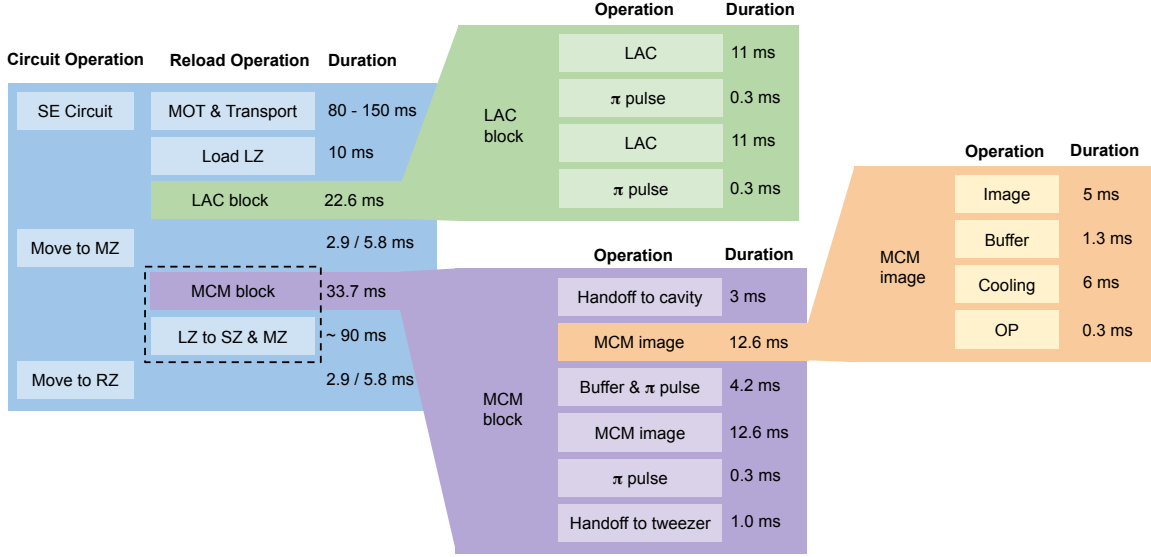
In this section we describe the codes used in our experiment, relating their structure to the standard toric code. The construction is summarized here for convenience; the general framework and higher-dimensional generalizations are in Ref. [S18].

The codes used in this work are two-dimensional toric codes on twisted tori, with the twist specified by a  $2 \times 2$  integer matrix  $M$  in Hermite normal form (HNF):

$$M = \begin{pmatrix} a & b \\ 0 & c \end{pmatrix}, \quad a, c > 0, \quad 0 \leq b < c. \quad (\text{S1})$$

The lattice  $\Gamma \subset \mathbb{Z}^2$ , defined as the points in the integer span of the rows of  $M$ , uniquely specifies the torus: two points of  $\mathbb{R}^2$  represent the same point on the torus if and only if their difference is a lattice vector (see Extended Data Fig. S2).

Each integer-coordinate vertex in the fundamental domain contributes two outgoing edges—horizontal and vertical—carrying physical qubits, along with one plaquette. The domain contains  $ac = \det M$  such vertices,



Extended Data Fig. S1. Timing for a SE cycle with SZ reloading. The duration for MOT and transport,  $\approx 60$  ms, is adjusted to be shorter than the length of the SE circuit, which is roughly 80 ms for det 8 and 150 ms for det 16. The illustrated timing diagram applies to a full SZ reloading cycle (occurring, e.g., every 10th cycle). For standard cycles involving only a replace from SZ, the reload operation column under the blue block consists solely of the black dashed box, and the “LZ to SZ & MZ” operation is replaced by “SZ to MZ”, reducing the overall duration to  $\sim 30$  ms. The rearrangement duration to shuttle atoms from the MZ to the RZ (or vice versa) is 2.9 ms when running the det 8 code which uses half of the MZ sites, and 5.8 ms when running the det 16 code which uses the entire MZ.

giving  $n = 2ac$  physical qubits. The standard local toric code prescription applies: an  $X^{\otimes 4}$  stabilizer at each vertex and a  $Z^{\otimes 4}$  stabilizer at each plaquette.

For the standard toric code ( $b = 0$ , diagonal  $M$ ), the minimum-weight logical operators are coordinate-aligned strings of weights  $a$  and  $c$ , giving code parameters  $[[2ac, 2, \min(a, c)]]$ . Twisting distorts these minimum-weight representatives and the diagonal entries of  $M$  no longer correspond to code  $s$ . In general the code distance is the  $\ell_1$  systole of  $\Gamma$ ,

$$d = \min_{v \in \Gamma \setminus \{0\}} \|v\|_1. \quad (\text{S2})$$

We use Gurobi to find minimum-weight representatives of logical operators of these codes.

Using the qubit numberings shown in Extended Data Fig. S2 (b) and (c), one can verify the following minimum weight representatives of the  $X$  and  $Z$  logical operators of the logical qubits probed by our experiments:

det 8 code,  $(d_X, d_Z) = (4, 3)$ :

$$\bar{X}_1 = X_0 X_2 X_4 X_6, \quad (\text{S3})$$

$$\bar{Z}_1 = Z_6 Z_7 Z_{15}, \quad (\text{S4})$$

det 16 code,  $(d_X, d_Z) = (6, 4)$ :

$$\bar{X}_1 = X_2 X_5 X_8 X_{11} X_{14} X_{16}, \quad (\text{S5})$$

$$\bar{Z}_1 = Z_{10} Z_{11} Z_{12} Z_{27}. \quad (\text{S6})$$


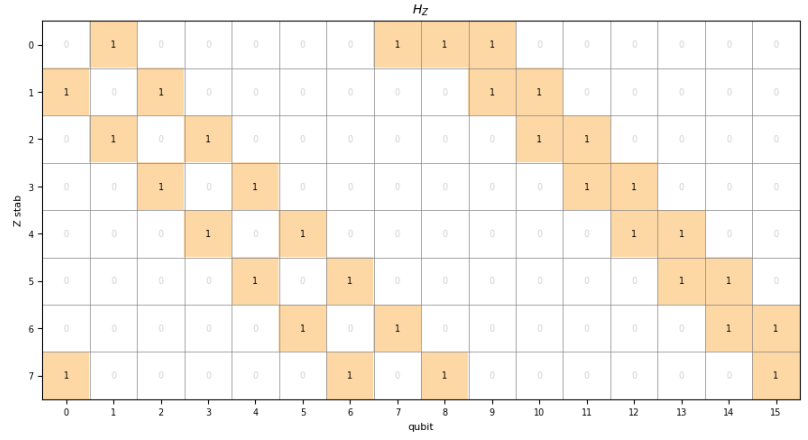

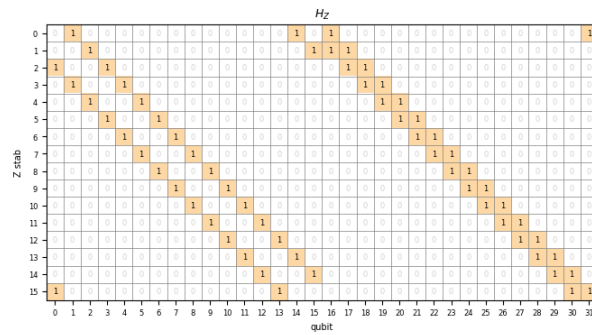

For explicit calculations yielding the convention chosen for Extended Data Fig. S2, we enumerate the  $ac$  vertices of the fundamental domain on a rectangular integer grid as  $v_{i,j} = (i, j)$  for  $i = 0, 1, \dots, a - 1$  and  $j = 0, 1, \dots, c - 1$ . Each vertex carries two qubits, on its outgoing horizontal and vertical edges:  $q_h(i, j)$  is the horizontal edge from  $(i, j)$  to  $(i + 1, j)$ , and  $q_v(i, j)$  is the vertical edge from  $(i, j)$  to  $(i, j + 1)$ . The boundary identifications follow from the lattice generators:  $(a, j) \equiv (0, j - b \bmod c)$  via the first generator  $(a, b)$ , and  $(i, c) \equiv (i, 0)$  via the second.

#### A. Swapped Syndrome Extraction Schedule

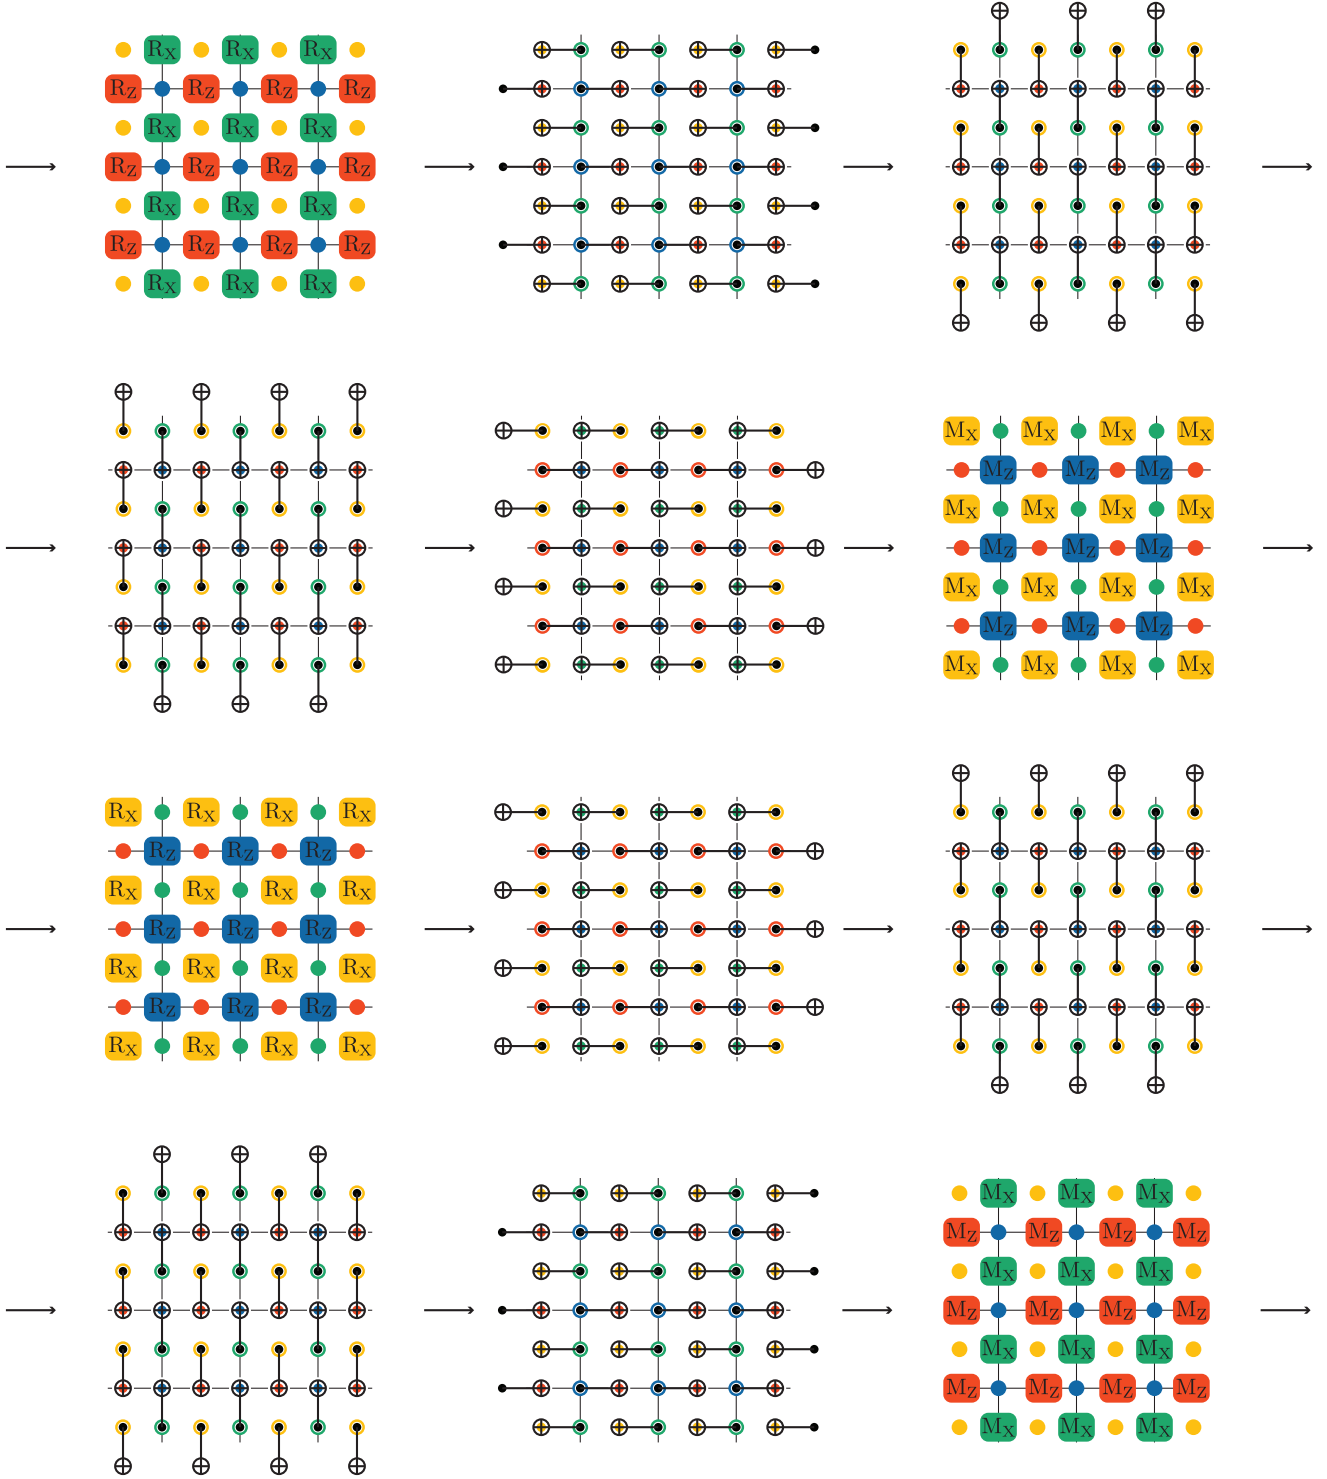
For both  $X$ - and  $Z$ -type stabilizers, data qubits were scheduled in a E-S-N-W pattern as shown in Fig. S3.

### SIV. LOSS-AWARE DECODING

Although our circuits are constructed to identify and correct for atom loss, dealing with such losses during decoding presents a separate challenge. Although atom loss resembles an ideal erasure channel, we do not identify the exact temporal location of atom loss. Furthermore, we must deal with this delayed erasure in the presence of other errors that leave atoms within the qubit subspace

(a)  $M = [[1, 2], [0, 8]]$ ,  $n = 16$ ,  $k = 2$ (b)  $M = [[1, 2], [0, 8]]$ (c)  $M = [[1, 3], [0, 16]]$ 

Extended Data Fig. S2. (a) Obtaining the twisted torus of the det 8 code from the HNF. Starting with the square cellulation of  $\mathbb{R}^2$  at integer coordinates, the rows of  $M$  span a parallelogram: the fundamental domain of the torus, tiling  $\mathbb{R}^2$  under translation by lattice vectors. Opposing edges of the fundamental domain are identified to yield the cellulated twisted torus. We can understand the geometry of the twisted stabilizer generators by drawing those of the standard toric code on the plane and translating any qubit outside of the fundamental domain into a qubit inside the fundamental domain via a unique lattice vector. (b) The stabilizer generators of the det 8 code. (c) The stabilizer generators of the det 16 code.



Extended Data Fig. S3. Two rounds of loss-tolerant syndrome extraction. In each syndrome cycle, fresh qubits are initialized, four layers of gates perform syndrome extraction, and the former data qubits are measured. As a result, data and ancilla qubits exchange roles once per syndrome cycle so that every physical qubit is measured once every two syndrome cycles. These measurements enable qubit loss to be detected and corrected while simultaneously performing the usual stabilizer measurements.

[S46, S58].

We follow the terminology of Stim, and formulate the decoding problem in terms of detector error models [S52, S53]. For each possible error channel, a detector error model stores the prior likelihood that an error occurs, which detectors are flipped, and which logical observables are flipped. Using Stim it is straightforward to compute the set of detectors flipped by a given Pauli error at a specific circuit location. Our heuristic approach to decoding with atom loss is based on a simple model for Rydberg  $CZ$  gates. If a qubit is lost prior to a  $CZ$  gate, there is no Rydberg interaction between the lost qubit and its  $CZ$  partner and we treat this absent interaction as a possible phase error on the partner qubit.

We observe that for individual atoms participating in a  $CZ$  gate we expect  $\langle Z \rangle = 0$ . Ancilla qubits are initialized in  $|+\rangle$  states (or other states in the  $X - Y$  plane of the Bloch sphere) so that they are sensitive to the conditional phase of a  $CZ$  gate. Similarly, single-qubit Pauli operators acting on data qubits have no expectation value. If we imagine such a qubit being lost before a  $CZ$  gate and then reinserted after the  $CZ$  gate, there is a relative phase “missing” between the  $|0\rangle$  and  $|1\rangle$  components of the lost-and-reinserted qubit’s wavefunction.

We now outline the procedure for updating a detector error model in response to a lost qubit. For each lost qubit, we compute the likelihood  $L_t$  that the qubit was lost when participating in a gate at timestep  $t$ . For each  $CZ$  gate between the qubit’s initialization and measurement, we introduce a new Pauli  $Z$  error mechanism on the partner qubit at the circuit location immediately following the  $CZ$  gate. We give this error mechanism a likelihood of  $\epsilon = L_t/2$ . The factor of  $1/2$  reflects that the lost qubit had no expectation value for its single-qubit  $Z$  operator prior to loss. Note that although an atom lost at timestep  $t$  is also lost at timestep  $t+1$ , we ignore possible correlations between induced error mechanisms. In addition to the error mechanisms introduced on the partner qubits, we introduce a readout error with a likelihood of 50% on the lost qubit immediately before readout while assigning the lost measurement a random bit-value.

In practice, we use Stim to pre-compute the set of updates associated with each possible observed loss. Thus, obtaining an updated detector error model can be efficiently computed for each bitstring. We note that this approach is compatible with any algorithmic decoder that accepts a detector error model (or some equivalent representation) as input. Examples of such decoders are Relay-BP and PyMatching [S51, S59].

Priors for the detector-error model are generated from the independently calibrated noise model detailed in Section SV for data with atom reloads. For data without atom reloads, we use a noise model which reduces  $CZ$  loss and increases loss during ancilla reset after MCM.

## SV. ERROR PARAMETERS

The simulated error sensitivity presented in Fig. 3(c) of the main text adopts a parameter set derived from independent measurements outlined below and summarized in Table S1. The uncertainties quoted below represent the statistical fluctuations over time for most physical measurements; for parameters specifically derived from Gaussian decay channels, the error bars reflect the uncertainty in the fitting procedure.

The single-site single-qubit (1Q) gate performance is benchmarked via standard Clifford randomized benchmarking (CRB) [S60], yielding an average error per-Clifford decay of  $4.6(1.6) \times 10^{-4}$ . This error is converted to a depolarizing Pauli stochastic error per native  $SX$  gate of  $p_{X,Y,Z} = 1.15 \times 10^{-4}$  [S61] (note we assume our virtual  $Z$  rotations to be of unit fidelity).

The two-qubit (2Q) errors are determined via interleaving  $CZ$  gates within a reference sequence of two different random 1Q Cliffords applied to pairs of qubits (interleaved randomized benchmarking henceforth IRB) [S60]. The reference sequence is generated such that the inverting 2Q Clifford for the interleaved sequence gives the total sequence  $depth$   $CZ$  gates. The reference sequence includes the atom movement necessary to switch between 1Q (register) and 2Q gates (interaction zone), even though no 2Q gates are performed. The 2Q loss (here referring to at least one lost atom) is extracted from the additional loss in the interleaved sequence relative to the reference sequence, yielding  $9.6(2.0) \times 10^{-3}$  per gate. This is mapped to a model where each atom is treated independently and equally likely to be lost. The error per  $CZ$  gate post-selected on the survival of both atoms is extracted to be  $4.3(1.9) \times 10^{-3}$ . According to a theoretical model of our  $CZ$  gates, this error is mapped to a 2Q Pauli stochastic error channel with probabilities  $p_{IZ} = p_{ZI} = 10p_{ZZ}$ , where  $p_{ZZ} = 1.9 \times 10^{-4}$ .

We benchmark errors during movement and idling on a per-unit-time basis. Many of these are Gaussian processes, but our simulation pipeline is not currently able to use time-varying error probabilities to reproduce this behavior. Instead, we approximate Gaussian decay by matching the zeroth moment of a Gaussian function with time constant  $\tau_{\text{decay}}$  over a maximum experiment timescale  $t_{\text{total}}$  to an exponential decay with an inverse time constant  $\frac{2t_{\text{total}}}{3\tau_{\text{decay}}^2}$ . Here,  $t_{\text{total}} \sim 400$  ms is set corresponding to the duration of two full cycles of our det 16 code, after which the qubits are reset via MCM. Note that this acts as an upper bound on errors in experiments with timescales shorter than  $t_{\text{total}}$  due to the higher initial slope of the exponential function. These errors are marked with inverse time units.

With the grouping “Idle stochastic”, we account for three categories of single-qubit stochastic Pauli errors: qubits idling at RZ sites during circuits (*idle Pauli*), qubits undergoing physical rearrangement (*move Pauli*), and data qubits idling at RZ sites during mid-circuit measurement (*MCM data qubit Pauli*). The *idle Pauli* con-

tribution is benchmarked using a Hahn echo sequence designed to mimic the circuit’s built-in echo configuration. A Gaussian decay fit yields a coherence time of  $T_2 = 7.2(0.6)$  s from which we extract an exponential time constant of  $5.1(0.9) \times 10^{-3} \text{ s}^{-1}$ . The *move Pauli* error rate is assumed to be identical, as independent measurements reveal no discernible performance degradation during rearrangement. The *MCM data qubit Pauli* error is parametrized per MCM cycle. Its Pauli  $X$  component is characterized by averaging the spin-flip probabilities of the two qubit states over repeated MCM, yielding a stochastic Pauli  $X$  error of  $1.9(0.5) \times 10^{-3}$  per MCM cycle. The Pauli  $Z$  component is benchmarked via a Ramsey sequence spanning an increasing number of MCM cycles; the resulting Ramsey contrast decay of  $1.08(0.24) \times 10^{-2}$  is converted to a stochastic Pauli  $Z$  error of  $5.4(1.2) \times 10^{-3}$  per MCM cycle.

“Idle loss” mechanisms include qubit loss during rearrangement (*move loss*), loss accrued while idling within static tweezers traps (*idle loss*), and loss accrued for qubits within the static RZ traps during MCM (*MCM data qubit loss*). The *move loss* is measured by tracking atom survival probability with moves of increasing duration. This follows a non-exponential Gaussian decay with time constant  $0.9(.2)$ s yielding an exponential decay rate  $2.7(1.1) \times 10^{-2} \text{ s}^{-1}$  over the time of an average number of physical rearrangement moves per atom over two complete cycles. For the *idle loss*, we apply our Gaussian approximation directly to the data from Fig. 4, yielding a loss rate of  $2.3(0.2) \times 10^{-3} \text{ s}^{-1}$ . The *MCM data qubit loss* is benchmarked with repeated MCM cycles and averaged over both qubit states to be  $3.5(0.7) \times 10^{-3}$  per MCM cycle. Note that the data qubits incur an extra two-photon loss channel compared to the imaged ancilla qubits from the trapping light and residual MCM imaging light.

Finally, state preparation and measurement errors (SPAM) include contributions from readout errors and loss. The readout error profiles differ slightly between regular and MCM images; however, they remain consistently  $< 2.0(0.4) \times 10^{-3}$  per image and constitute only a minor fraction of the overall error budget. SPAM loss is divided into two distinct factors: initial array filling imperfections (*initial load loss*, evaluated at  $6.8(6.8) \times 10^{-3}$ ) and atom loss during reset after MCM (*MCM reset loss*, evaluated at  $9.5(4.6) \times 10^{-3}$  per MCM cycle).

Note that while some errors like Pauli errors and readout infidelities are universal to all detectors, the manifestations of different loss mechanisms on detector probability are dependent on the chronological order and our qubit role-swap schedule (see Extended Data Fig. S4). This role-swapping of data and ancilla qubits effectively symmetrizes the error channels in the bulk sequence; however, this symmetry is broken at the boundaries. Specifically, the 0th, 1st, and final detectors capture additional or incomplete errors.

## SVI. SIMULATIONS

Simulations were performed using a Clifford simulator including loss. Coherent errors and atom heating were not included due to the additional computational overhead required. Additionally, errors introduced by atom reloading were not considered in this simulation (we note that atom reloading causes an observable spike in detector frequencies around the reloading rounds, but it does not appear to have a substantial effect on logical error rate in our data).

We first compare simulations using a set of independently benchmarked parameters (Table S1, “Sim A”) to several aspects of the experimental data. This model captures many features of the detector frequencies (Extended data Fig. S5(a)) including the slight elevation of  $X$  detectors over  $Z$  detectors as well as the firing frequency of early rounds; however, it underestimates “bulk cycle” detector frequencies. Turning to atom loss resolved by cycle (Extended Data Fig. S5(b)), we see a similar trend. Looking at the distribution of the number of detectors fired in a bulk cycle (Extended Data Fig. S5 (c)), we see that the model captures the mean of the distribution reasonably, but underestimates the long tail which could be indicative of unaccounted for correlated errors [S62]. This model underestimates logical error rates (Extended Data Fig. S5 (e)), and further predicts an error suppression factor  $\Lambda_Z > 1$  at arbitrary cycles whereas our data appears to plateau to  $\Lambda_Z \approx 1$  at and above 8 cycles.

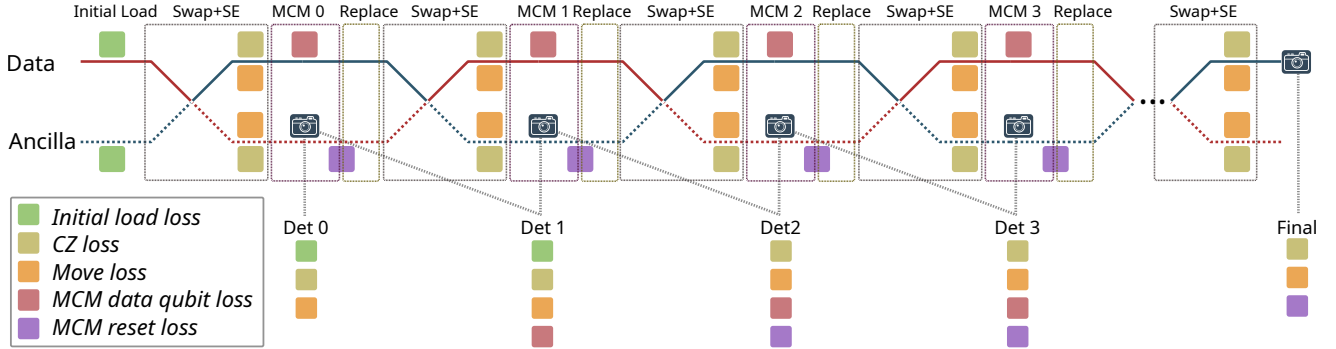
By phenomenologically adding an additional 1% data qubit loss during MCM (“Sim B”), we better capture the detector frequencies and loss during bulk cycles (Extended Data Fig. S5 (a) and (b)); however, we now overestimate these quantities in early cycles. This could be explained by slow atom heating during the circuit leading to increased loss with rounds. This model better predicts logical error rates in addition to the observed decrease in  $\Lambda_Z$  with cycles. It still underestimates the tail of the bulk detector firing distribution (Extended Data Fig. S5 (c)) as well as detector firing frequencies at low loss (Extended Data Fig. S5 (d)), indicating that unaccounted for coherent errors likely also play a role in the observed logical error rates.

We note that the individual datasets shown in Extended Data Fig. S5(a) and (b) show significant variation in detector frequency and loss. The origins are currently unclear; however, possible explanations include drifting commensurability between the imaging cavity and tweezer arrays causing atom heating and loss as well as calibration timing.

We explore the contributions to our observed LERs by simulating these codes using the parameters in Table S1, and scaling each parameter individually to get partial gradients to the full LER. To do this, we simulated the LER for the full noise model, except one of the parameters is scaled down to zero ( $\text{LER}_0$ ), and then we simulate the LER again with that parameter scaled by

TABLE S1. **Summary of error parameters.** Summary of the benchmarked physical error model parameters - see text for more details on benchmarking experiments and mapping to simulation parameters. These parameters are used in "Sim A".

Category	Parameter	Units	Benchmarked	Calibration Experiment
1Q stochastic	SX total Pauli	per gate	$3.45(1.) \times 10^{-4}$	1Q CRB
2Q stochastic	CZ total Pauli	per gate	$4.03(1.8) \times 10^{-3}$	IRB
2Q loss	CZ loss	per gate	$9.6(2.0) \times 10^{-3}$	IRB
Idle stochastic	idle Pauli $Z$	$s^{-1}$	$5.1(0.9) \times 10^{-3}$	Hahn echo (Gaussian)
	move Pauli $Z$	$s^{-1}$	$5.1(0.9) \times 10^{-3}$	Hahn echo (Gaussian)
	MCM data qubit Pauli $X$	per MCM	$1.9(0.5) \times 10^{-3}$	MCM state decay
	MCM data qubit Pauli $Z$	per MCM	$5.4(1.2) \times 10^{-3}$	MCM Ramsey contrast decay
Idle loss	move loss	$s^{-1}$	$2.7(1.1) \times 10^{-2}$	Repeated moves (Gaussian)
	idle loss	$s^{-1}$	$2.3(0.2) \times 10^{-3}$	Trap lifetimes (Gaussian)
	MCM data qubit loss	per MCM	$3.5(0.7) \times 10^{-3}$	MCM survival decay
SPAM (non-loss)	readout	per image	$< 2.0(0.4) \times 10^{-3}$	Prepare and measure
SPAM (loss)	initial load loss	per load	$6.8(6.8) \times 10^{-3}$	Prepare and measure
	MCM reset loss	per MCM	$9.5(4.6) \times 10^{-3}$	MCM survival decay



Extended Data Fig. S4. Impact of loss mechanisms on detector statistics. Data (solid line) and ancilla qubits (dashed line) are susceptible to distinct loss mechanisms throughout the circuit (see legend). Apart from the first and final cycles, each detector is sensitive to the losses within intervals over two consecutive MCMs.

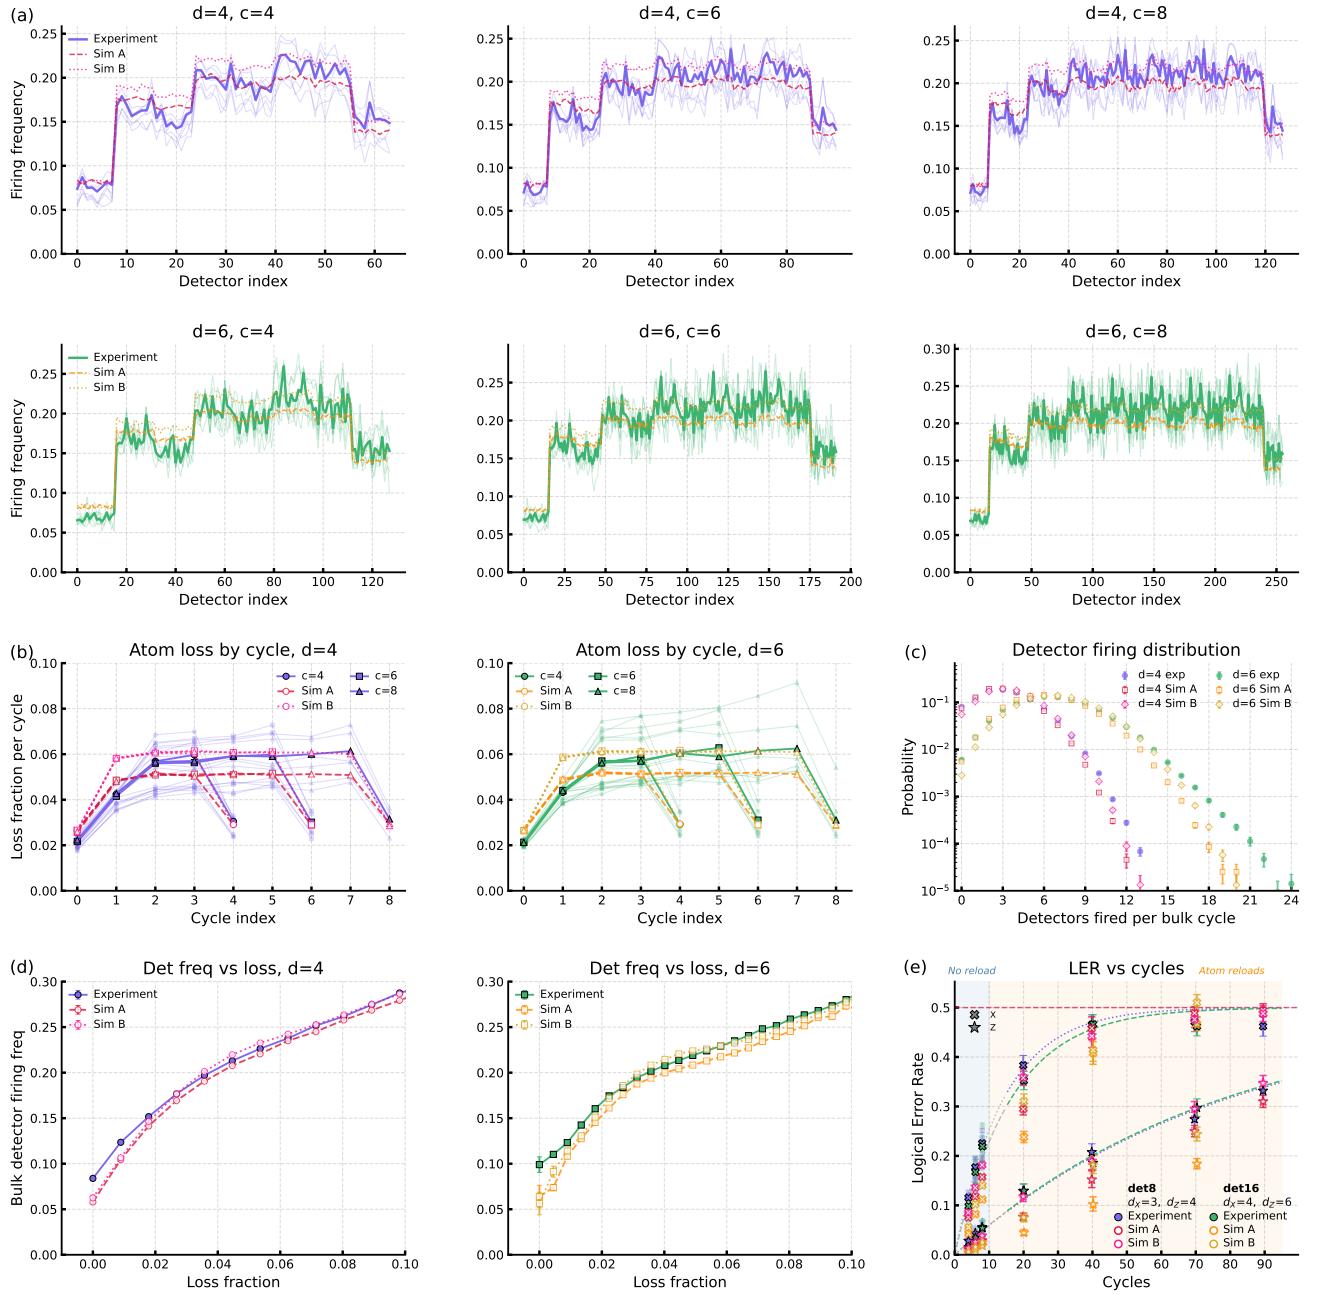
a factor  $f$  ( $LER_f$ ). Then, to first order, we can approximate the gradient to the LER for each parameter with  $\Delta_{LER} = (LER_f - LER_0)/f$ . The results of the partial gradients are categorized as in Table S1, and are depicted in Fig. 3(c) of the main text for the gradient range  $(0, 1)$  (i.e.  $f = 1$ ), with simulations using 100,000 shots.

## SVII. SUPPLEMENTAL DATA


For completeness, we include here the data for the detector probability of the det 8 toric code up to 90 cycles

with SZ reloading (Extended Data Fig. S6).


In Extended Data Fig. S7, we break down the  $Z$  logical error rates for the six individual datasets comprising Fig. 3 of the main text up to 8 cycles without reload. These itemized runs confirm that the observed logical error suppression of the det 16 code over the det 8 code remains persistent against experimental drift.



Extended Data Fig. S5. Comparison of experiment data to simulations. “Sim A” is using an independently benchmarked error model while “Sim B” phenomenologically adds 1% to MCM data qubit loss. Semi-transparent lines are individual datasets, and solid lines are averaged over all of these datasets. (a) Average detector frequency by index for all data without reloads. Note det 8 (det 16) contains 16 (32) detectors per cycle, and the first and last cycles only contain Z detectors. For all other cycles, all X detectors - we attribute this to slow atom heating during the circuit. Note Sim A appears to underestimate late cycles, while SimB overestimates early cycles - we attribute this to slow atom heating during the circuit. (b) Atom loss by syndrome extraction cycle - a slow increase in loss across bulk rounds again appears which is not captured well by simulation, and we attribute this to slow atom heating. There is significant variation across datasets in both loss and detector frequency indicating slow drifts with time. (c) Bulk detector (detectors in rounds 2 to  $c - 1$  indexing from 0) firing distributions. The experimental data has a slightly longer tail than either simulation, indicating correlated errors which we are not currently accounting for (atom heating, coherent errors). (d) Bulk detector firing frequency as a function of measured ancilla loss fraction in that cycle. (e) Comparison of experimental and simulated logical error rates. Sim A underestimates the experimental logical error rates and predicts error suppression with distance with increasing cycles. The addition of extra MCM data qubit loss in Sim B recovers decreased error suppression with increasing cycles, although it still seems to underestimate errors at increased cycles, particularly for the det 16 code.



Extended Data Fig. S6. A det 8 toric code detection probability over 90 consecutive syndrome extraction cycles with SZ reloading every 10th cycle. The missing atom fraction is shown in shades.



Extended Data Fig. S7.  $Z$  Logical error rates for the six individual datasets used in the demonstration of logical error suppression without atom reload. The number of shots per point and duration of each dataset is indicated in the sub-figures. While slow fluctuations of the experiment affect the logical error rates, we consistently see logical error rates for the det 16 code ( $d_X = 6$ ) at or below that of the det 8 code ( $d_X = 4$ ).