

Experimental demonstration of logical magic state distillation

<https://doi.org/10.1038/s41586-025-09367-3>

Received: 19 December 2024

Accepted: 7 July 2025

Published online: 14 July 2025

 Check for updates

Pedro Sales Rodriguez^{1,5}, John M. Robinson^{1,5}, Paul Niklas Jepsen^{1,5}, Zhiyang He^{1,2}, Casey Duckering¹, Chen Zhao¹, Kai-Hsin Wu¹, Joseph Campo¹, Kevin Bagnall¹, Minho Kwon¹, Thomas Karolyshyn¹, Phillip Weinberg¹, Madelyn Cain³, Simon J. Evered³, Alexandra A. Geim³, Marcin Kalinowski³, Sophie H. Li³, Tom Manovitz³, Jesse Amato-Grill¹, James I. Basham¹, Liane Bernstein¹, Boris Braverman¹, Alexei Bylinskii¹, Adam Choukri¹, Robert J. DeAngelo¹, Fang Fang¹, Connor Fieweger¹, Paige Frederick¹, David Haines¹, Majd Hamdan¹, Julian Hammett¹, Ning Hsu¹, Ming-Guang Hu¹, Florian Huber¹, Ningyuan Jia¹, Dhruv Kedar¹, Milan Kornjača¹, Fangli Liu¹, John Long¹, Jonathan Lopatin¹, Pedro L. S. Lopes¹, Xiu-Zhe Luo¹, Tommaso Macri¹, Ognjen Marković¹, Luis A. Martínez-Martínez¹, Xianmei Meng¹, Stefan Ostermann¹, Evgeny Ostroumov¹, David Paquette¹, Zexuan Qiang¹, Vadim Shofman¹, Anshuman Singh¹, Manuj Singh¹, Nandan Sinha¹, Henry Thoreen¹, Noel Wan¹, Yiping Wang¹, Daniel Waxman-Lenz¹, Tak Wong¹, Jonathan Wurtz¹, Andrii Zhdanov¹, Laurent Zheng¹, Markus Greiner³, Alexander Keesling¹, Nathan Gemelke¹, Vladan Vuletić⁴, Takuya Kitagawa¹, Sheng-Tao Wang¹, Dolev Bluvstein³, Mikhail D. Lukin³, Alexander Lukin¹, Hengyun Zhou¹ & Sergio H. Cantú¹

Realizing universal fault-tolerant quantum computation is a key goal in quantum information science^{1–4}. By encoding quantum information into logical qubits using quantum error correcting codes, physical errors can be detected and corrected, enabling a substantial reduction in logical error rates^{5–11}. However, the set of logical operations that can be easily implemented on these encoded qubits is often constrained^{1,12}, necessitating the use of special resource states known as ‘magic states’¹³ to implement universal, classically hard circuits¹⁴. A key method to prepare high-fidelity magic states is to perform ‘distillation’, creating them from multiple lower-fidelity inputs^{13,15}. Here we present the experimental realization of magic state distillation with logical qubits on a neutral-atom quantum computer. Our approach uses a dynamically reconfigurable architecture^{8,16} to encode and perform quantum operations on many logical qubits in parallel. We demonstrate the distillation of magic states encoded in $d = 3$ and $d = 5$ colour codes, observing improvements in the logical fidelity of the output magic states compared with the input logical magic states. These experiments demonstrate a key building block of universal fault-tolerant quantum computation and represent an important step towards large-scale logical quantum processors.

Quantum error correction (QEC) enables scalable quantum computation by exponentially suppressing logical error rates. However, the set of logical operations that can be efficiently implemented on these encoded qubits is constrained, making it challenging to perform universal quantum processing¹². For example, many QEC codes support only the realization of so-called Clifford gates^{17–19}. As Clifford gates can be efficiently simulated classically¹⁴, additional non-Clifford resources are required to achieve computational universality and quantum advantage. To overcome this difficulty, special quantum states, aptly named ‘magic states’, can be used to complete a universal set of logical operations by gate teleportation¹⁵. Owing to their relatively high cost of preparation, these magic states are often considered the key resource for scalable processing⁴.

High-fidelity magic states can be produced by refining multiple noisy copies through magic state distillation (MSD)¹³. The noisy states, encoded in data QEC codes, are concatenated into a distillation code and purified through protected operations on the data codes (Fig. 1). If the input fidelity exceeds the so-called distillation threshold, the fidelity of the output state is improved compared with the input. An attractive feature of MSD is that the logical-level circuit is independent of the data QEC code. Consequently, by increasing the data code distance and implementing the logical circuit with encoded operations, the logical fidelity of the output magic states can be improved, in principle, to any desired level through multiple rounds of distillation^{20–22}.

Important recent experiments have demonstrated MSD with physical qubits^{23,24}, but the direct physical encoding prevents suppression

¹QuEra Computing, Boston, MA, USA. ²Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA. ³Department of Physics, Harvard University, Cambridge, MA, USA. ⁴Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁵These authors contributed equally: Pedro Sales Rodriguez, John M. Robinson, Paul Niklas Jepsen. [✉]e-mail: hyzhou@quera.com; scantu@quera.com

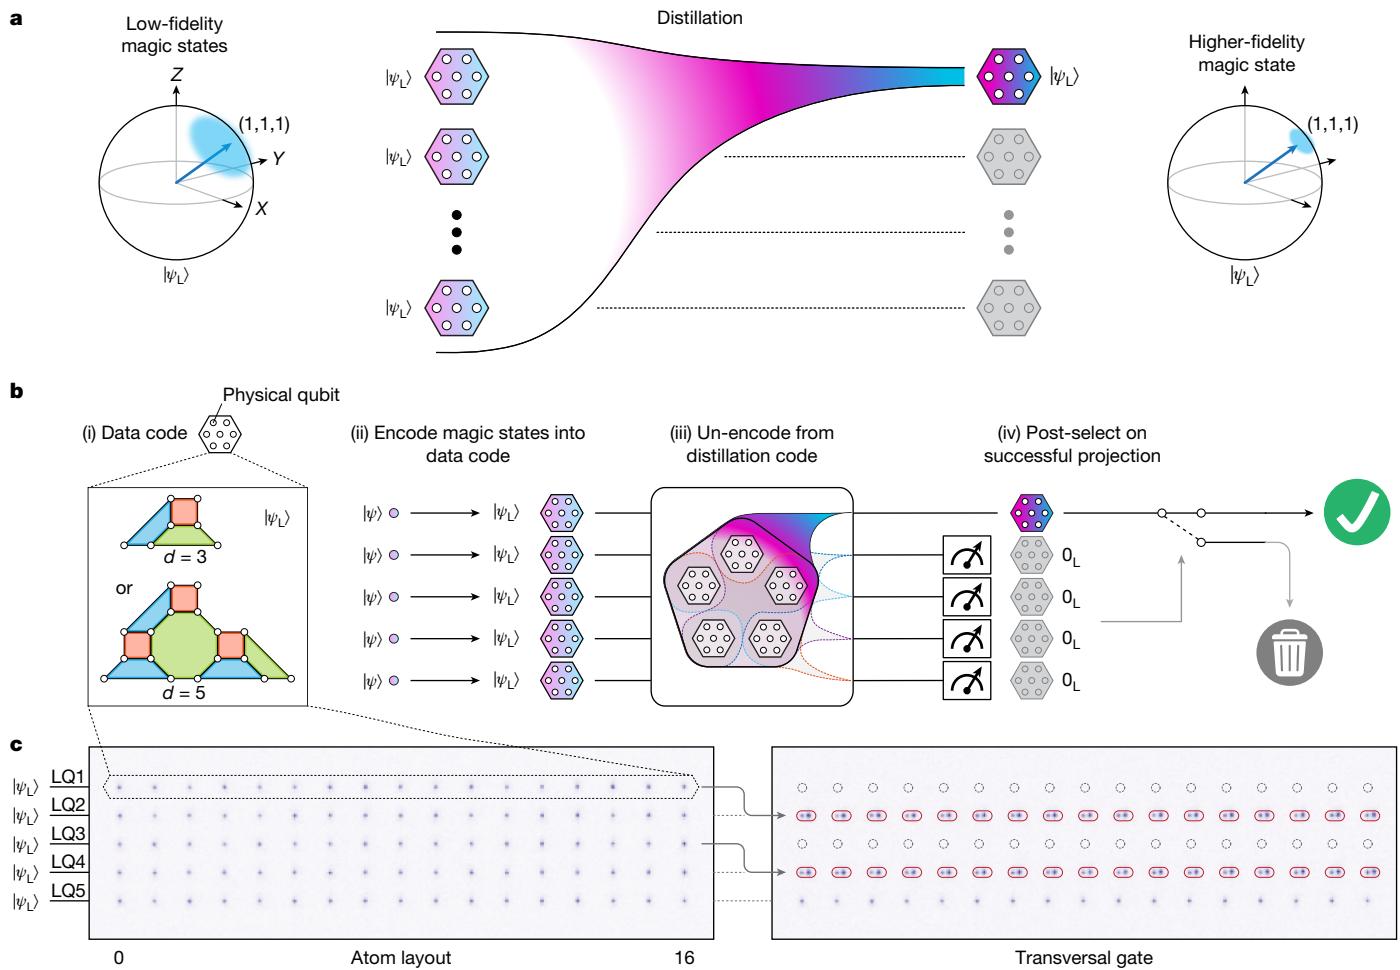


Fig. 1 | Logical MSD factory. **a**, Schematic of the Bloch sphere representation of magic state $|\psi_L\rangle$ (left) pointing in the $(1, 1, 1)$ direction with shaded region indicating noise. Distillation (centre) takes multiple noisy logical inputs and produces a higher fidelity magic state (right). **b**, A 5-to-1 distillation procedure (left to right). Non-fault-tolerant encoding of physical magic states into five data code logical qubits $|\psi_L\rangle$ protects logical operations (i and ii). In particular, we encode into distance 3 and 5 colour codes (i). Encoded states (ii) are purified using a distillation code. By running the un-encoding circuit of the distillation code (iii) and conditioning on distillation syndromes (iv), we have simultaneously

projected into the code state of the distillation code and un-encoded it into the output magic state. On measuring the correct distillation syndromes, the output qubit has been distilled to a higher fidelity along the $(1, 1, 1)$ direction. **c**, Averaged atom images from the $d = 5$ distillation experiment, showing 85 physical qubits encoded into five logical qubits (LQ1–LQ5) with 17 physical qubits each (left), shown here in spatial light modulator traps. Rows of logical qubits are coherently reconfigured for transversal CZ gates throughout the distillation circuit (right), shown here with LQ1 and LQ3 in acoustic-optic deflector traps.

of logical gate errors. Complementary experiments have also shown remarkable progress in error-suppressed encoding of magic states into logical qubits using flag protocols^{25–27}. However, without the protection provided by the data code, these approaches generally have higher complexity and low success probability when targeting very low logical error rates, although recent work has markedly improved their performance for low physical error rates^{28–30}.

We realize MSD at the logical level on a neutral-atom quantum computer. Magic states are encoded using two-dimensional (2D) colour codes¹⁹, and subsequently a 5-to-1 logical MSD is performed¹³. The factory outputs a single magic state and the remaining four qubits, which we call distillation syndromes, determine successful distillation. Central to our approach is the dynamic reconfigurability and high degree of parallel control of the neutral-atom processor^{8,16}. We realize gate and layout-efficient encoding circuits for arbitrary logical states in the $d = 3$ colour codes, executing 10 logical qubit encoding circuits and $d = 5$ colour codes, executing 5 logical qubit encoding circuits in parallel. MSD is carried out using transversal Clifford gates, efficiently implemented with parallel atom rearrangement across all code distances. Correlated decoding^{8,31} is applied to the distillation syndromes, and

their stabilizer values are further leveraged as flags^{28,32} to enhance output logical fidelity. The operation of the MSD factory is verified by distilling states with varying input fidelity and confirming the error suppression scaling. Conditioned on observing the correct logical outcome and suitable stabilizer patterns on the four distillation syndrome logical qubits, we obtain an enhancement of logical magic state fidelity from $95.1_{-0.1}^{+0.1}\%$ to $99.5_{-0.4}^{+0.4}\%$ for $d = 3$ and from $92.5_{-0.2}^{+0.1}\%$ to $99.1_{-1.3}^{+0.7}\%$ for $d = 5$. This corresponds to a factor of 12_{-6}^{+25} infidelity suppression for $d = 3$ and a factor of 10_{-6}^{+7} for $d = 5$.

Arbitrary logical state encoding

Our experiments use a newly built, Gemini-class quantum processor built and operated at QuEra (QuEra Computing and Collaborators, manuscript in preparation, 2025a). Inspired by earlier experiments from Harvard⁸, it involves control over a 2D array of neutral-atom qubits in a reconfigurable architecture.

We start by preparing magic states encoded in the data QEC code^{25–27,33}, with fidelity above the distillation threshold (for 5-to-1 distillation, the threshold fidelity for depolarizing errors is 83%; ref. 13).

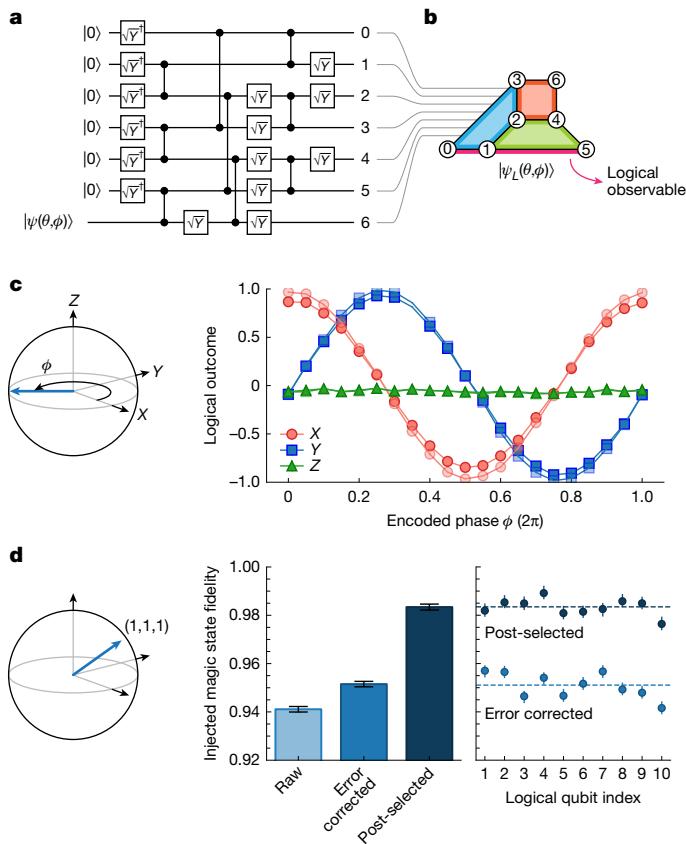


Fig. 2 | Parallel logical encoding of arbitrary states. **a**, Circuit for injecting an arbitrary state $|\psi(\theta, \phi)\rangle$ into the $[[7, 1, 3]]$ colour code. **b**, Schematic of $d = 3$ colour code stabilizers indicated by the three coloured regions, with a logical operator highlighted. **c**, Bloch sphere representation of the injected state with varying angle ϕ on the XY plane (left). Error-corrected logical outcomes for X, Y, Z measurement basis versus the injected phase. Faded markers indicate outcomes on post-selection on perfect stabilizers. **d**, Left, Bloch sphere representation of the $(1, 1, 1)$ magic state. Centre, injected $d = 3$ magic state fidelity corresponding to raw, error-corrected and post-selected on perfect stabilizers, averaged across all 10 logical qubits. Right, spatial distribution of injected magic state fidelities.

We choose the 2D colour code as our data QEC code, as the full Clifford group can be implemented transversally with it¹⁹. The $[[7, 1, 3]]$ colour code is illustrated in Fig. 2b, in which an X and Z stabilizer is associated with each coloured region, and the logical operators lie along the edge. Errors flip stabilizer values, so measuring stabilizers allows us to detect and correct physical errors in the circuit. The parameters $[[n, k, d]]$ denote a QEC code with n physical data qubits, k logical qubits, and code distance d , which can correct $\left\lfloor \frac{d-1}{2} \right\rfloor$ errors or detect $d-1$ errors. We use an arbitrary state encoding circuit that takes a physical qubit as input and encodes its state into a logical qubit (also known as state injection). In particular, to encode into the $d = 3$ 2D colour code, we use the circuit in Fig. 2a, optimized for atom movement and number of entangling gate layers (Methods).

We verify the encoding circuit by injecting a state lying on the XY plane and varying its angle (Fig. 2c). This results in a rotation of the encoded logical information, which can be read out as an oscillation in logical measurements in the X or Y basis. Logical measurements in the Pauli (X, Y, Z) basis are performed transversally, by measuring each physical qubit in the corresponding basis. To interpret data qubit measurement results, we calculate the stabilizer and logical outcome parities from physical measurement results. If errors are detected, we can perform error correction on the logical result or alternatively discard the measurement (error detection).

With the ability to encode arbitrary states, we shift our focus to encoding magic states for further use in distillation. In this work, we encode magic states that point in the $(1, 1, 1)$ direction on the Bloch sphere, for use in the MSD procedure based on the $[[5, 1, 3]]$ code¹³. We prepare this state by initializing in $|0\rangle$ followed by a local single-qubit rotation of angle $\arccos(1/\sqrt{3})$ about the $(-1, 1, 0)$ axis on the physical qubit to be injected. We perform logical quantum state tomography to estimate magic state fidelity (Methods). We find that the encoded logical magic states have raw logical fidelity $94.1_{-0.1}^{+0.1}\%$ (no error correction), error-corrected logical fidelity $95.1_{-0.1}^{+0.1}\%$ and error-detected logical fidelity $98.3_{-0.1}^{+0.1}\%$ (post-select on perfect stabilizers) (Fig. 2d). The error-detected state fidelity is close to the original physical magic state fidelity of $98.9_{-0.1}^{+0.1}\%$, indicating that most of the added errors during the encoding process will also trigger syndromes. To scalably use the resource states in a larger circuit, we cannot rely on post-selecting on stabilizers that only become available when performing transversal measurements of the logical qubit later on. Therefore, we focus on comparing the logical fidelity of magic states when only error correction (and no further post-selection) is applied on the target magic state.

5-to-1 MSD

The logical encoding circuit described above is not fault-tolerant, because physical errors on the injected physical qubit will lead to logical errors, resulting in a logical error rate that scales linearly with the physical error rate. To further suppress the logical error rate, we make use of MSD, which uses the properties of a distillation QEC code and the fault-tolerant gates of the data QEC code to improve the magic state quality (Fig. 1).

Our magic state factory is based on the $[[5, 1, 3]]$ perfect code^{13,34}. Schematically, the factory takes five noisy logical magic states as input and applies a unitary un-encoding circuit of the distillation $[[5, 1, 3]]$ code, which we optimize to have only three layers of entangling gates (Fig. 3a and Methods). Measuring four of these logical qubits effectively measures the stabilizers of the distillation code, whereas the remaining logical qubit contains the output magic state. By post-selecting on the appropriate logical outcome of the four logical qubits (factory post-selection), we achieve quadratic suppression of the logical error rate. In the absence of errors, the factory acceptance rate of the 5-to-1 distillation factory is expected to be 1/6 (ref. 13).

To decode the logical measurement results, we use a most-likely error (MLE) correlated-decoder based on mixed-integer programming³¹, with error weights obtained from a separate characterization of our system (Methods). For $d = 3$, we also explore a maximum likelihood decoder (MLD) that simulates most of the logical outcomes for a given stabilizer pattern and use it to determine the correction (Methods). These decoders can also be used to characterize the confidence of a given logical outcome assignment, allowing further sliding-scale post-selection based on observed stabilizer patterns^{8,35,36} similar to flag protocols^{25,27,28,37,38}. This post-selection is commonly used in theoretical analysis when preparing resource states, and in accordance with this, we use only the stabilizers of the four distillation syndrome logical qubits to perform decoding for post-selection, because the output logical qubit is meant to be used for subsequent operations. We refer to the use of this physical stabilizer information to flag bad executions of the distillation circuit as stabilizer post-selection.

Experimental results of our logical MSD factory are shown in Fig. 3b. Starting with the error-corrected input logical magic state with fidelity $95.1_{-0.1}^{+0.1}\%$, without any stabilizer post-selection, the output magic state fidelity is worse than the injected state, because of the added physical errors during the distillation process. However, we find that approximately 50% stabilizer post-selection is sufficient to improve the output magic state fidelity, and full post-selection on perfect stabilizers of distillation syndrome qubits results in a fidelity of $99.5_{-0.4}^{+0.4}\%$.

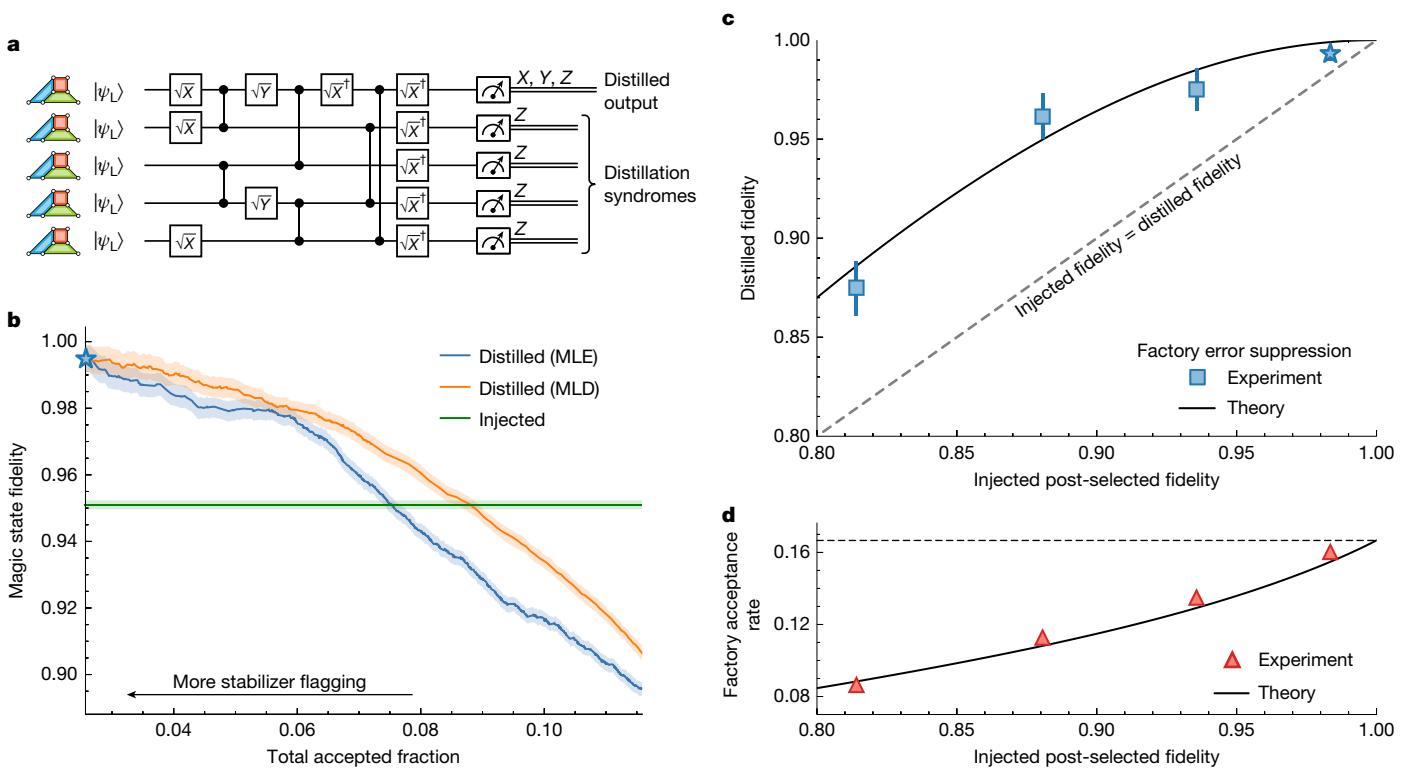


Fig. 3 | The 5-to-1 MSD. **a**, MSD circuit based on the $[[5, 1, 3]]$ code (distillation code). We measure distillation syndromes in the Z basis and perform tomography on the distilled output. The successful distillation syndrome for this circuit is 1011 (Methods). **b**, Fidelity of the output magic state for the $d = 3$ distillation (blue line for the MLE decoder and orange line for the MLD decoder; see main text) as a function of the total accepted fraction, which includes both sliding-scale post-selection on distillation syndrome stabilizers, and the factory acceptance (1/6 in the noiseless case). With sufficient stabilizer flagging, the output fidelity exceeds that of the input error-corrected magic state fidelity (green). The shaded regions indicate 68% confidence intervals, equivalent to 1σ . **c**, We examine the distilled fidelity with full stabilizer post-selection, after introducing coherent Z errors to the input magic states ($0.32\pi, 0.24\pi, 0.16\pi$ and 0 , left to right, blue points). The results are in good agreement with the theoretical expectation (grey line). The stars in **b** and **c** indicate the same data point. **d**, Factory acceptance rate of distillation syndromes after perfect stabilizer post-selection (acceptance rate between 13% and 16%) with the same coherent errors as in **c**. The dashed line indicates the 1/6 acceptance rate of the 5-to-1 magic state factory in the noiseless case.

Both decoders show similar performance, with the MLD decoder performing slightly better by accounting for the entropy of error configurations.

We further probe the physics of error suppression of the distillation code by artificially introducing coherent errors across the five input logical qubits (Fig. 3c,d), which we achieve by applying a Z rotation on the physical qubits before state injection. After encoding, this results in a magic state rotated around the X -axis, which we use as input to the factory. After full stabilizer post-selection on the four distillation syndrome qubits, we compare the error-corrected output fidelity against the post-selected input fidelity, to highlight the distillation behaviour on the logical information. We calculate the distillation gain for all added rotation angles.

As the added rotation angle error increases, we observe that the output state infidelity is consistent with quadratic suppression of the added error. We also find that the factory acceptance rate decreases with added errors, with an initial decrease that scales linearly with the added error. This can be understood from the fact that a single input logical error will lead to an outcome different from the correct distillation syndrome, reducing the factory acceptance rate without contributing to the distilled fidelity. Two input logical errors are needed to affect the distilled fidelity, giving rise to quadratic error suppression¹³.

Extending to larger code distance

Larger data codes offer stronger protection against physical errors when operated below threshold and are crucial for scaling to low logical error rates. To this end, we investigate data codes with larger code distances

(green). The shaded regions indicate 68% confidence intervals, equivalent to 1σ . **c**, We examine the distilled fidelity with full stabilizer post-selection, after introducing coherent Z errors to the input magic states ($0.32\pi, 0.24\pi, 0.16\pi$ and 0 , left to right, blue points). The results are in good agreement with the theoretical expectation (grey line). The stars in **b** and **c** indicate the same data point. **d**, Factory acceptance rate of distillation syndromes after perfect stabilizer post-selection (acceptance rate between 13% and 16%) with the same coherent errors as in **c**. The dashed line indicates the 1/6 acceptance rate of the 5-to-1 magic state factory in the noiseless case.

by performing MSD on five copies of a $[[17, 1, 5]]$ $d = 5$ colour code. After optimizing for our native gate set, we obtain the logical encoding circuit with five entangling layers shown in Fig. 4a, with its corresponding stabilizers shown in Fig. 4b. Encoding all five $d = 5$ magic states involves 85 physical qubits, which are coherently manipulated in parallel within the entanglement zone of the processor (Methods). The transversal single- and two-qubit distillation gates and atom moves are the same as in $d = 3$. We apply the same correlated decoding procedure with the MLE decoder and post-selection criteria as for the $d = 3$ case.

The experimental results for the $d = 5$ colour code are shown in Fig. 4c. We first note that encoding magic states into larger distance codes results in lower injected fidelity $92.5^{+0.1}_{-0.2}\%$, as the encoding circuit involves more physical gate operations. We note that more stabilizer post-selection is required to achieve a comparably high fidelity as $d = 3$, because our gate fidelities are not yet past the circuit threshold for the data code; this can be improved with further reduction of the physical error rate. With full stabilizer and factory post-selection, we observe a distillation gain of $6.6^{+0.7}_{-1.3}\%$, from an encoded fidelity of $92.5^{+0.1}_{-0.2}\%$ to a distilled fidelity of $99.1^{+0.7}_{-1.3}\%$ (Fig. 4c).

We compare the MSD performance across code distances in Fig. 4d, including physical MSD ($d = 1$) and logical MSD with the $d = 3$ and $d = 5$ colour codes. For physical MSD, preparation of physical magic states is limited only by qubit initialization, measurement and single-qubit gate fidelity. Without the ability to perform error correction, physical distillation introduces additional errors, leading to a lower output fidelity. Shifting to logical qubits, we observe that the injected state fidelity drops as the distance increases because of the added errors during the non-fault-tolerant encoding circuit. However, the data code provides

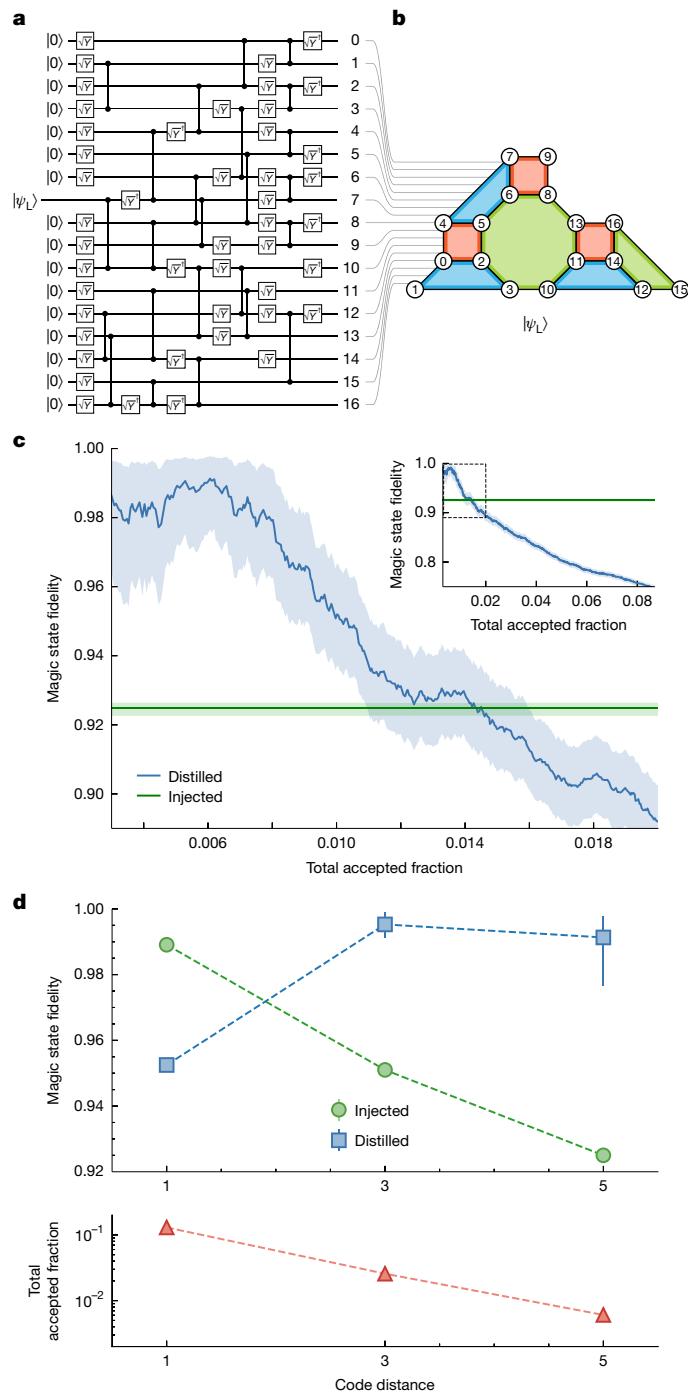


Fig. 4 | Scaling of distillation as a function of data code distance. **a**, The $d=5$ encoding circuit. **b**, The $d=5$ colour code stabilizers. **c**, Output magic state fidelity for $d=5$ distillation (blue) as a function of the total accepted fraction, again showing improvement over the input magic state fidelity (green). The inset shows the extended range. **d**, Injected (green circles) and distilled (blue squares) magic state fidelity with total acceptance fraction (red), when performing full stabilizer and factory post-selection, all as a function of code distance.

sufficient protection of distillation operations to achieve distillation gain for both $d=3$ and $d=5$.

Discussion and outlook

These experiments demonstrate the key ingredients of MSD for universal fault-tolerant quantum computation. Leveraging the dynamic

reconfigurability and transversal gate operations of the neutral-atom platform to realize a logical MSD factory, our approach allows us to probe key aspects of the distillation process. This factory can be combined with mid-circuit measurement and feedforward^{8,39–43}, to execute universal quantum algorithms by magic state teleportation. Although the present experiments demonstrate the performance of MSD past the distillation threshold, further improvements in both the fidelity and the rate of the MSD factory are required to enable the execution of deep logical circuits. Although at present, the use of higher distance codes results in lower acceptance fraction to achieve large fidelity gain, by improving gate fidelities to values well below the 2D colour code threshold, the accepted fraction can remain comparable as the code distance increases, and multiple distillation rounds can be executed for further error suppression. More specifically, we estimate (Methods) that a twofold reduction in physical error rates can result in distillation gain without stabilizer post-selection.

To enable efficient large-scale universal quantum computation, these fidelity improvements should also come hand-in-hand with further co-design of magic state preparation. Although MSD represents a foundational approach for implementing non-Clifford operations and has the advantage of being flexibly adaptable to many data codes, alternative methods with various trade-offs should also be explored. These include the use of QEC codes with transversal non-Clifford gates^{30,44,45}, as well as advanced flag protocols^{25,27,28,46–48} and the recently proposed magic state cultivation²⁹ schemes. Moreover, alternative MSD factories with improved input-to-output ratios or better error suppression^{13,20–22,49} can be co-designed and explored experimentally in the current framework. Paving the way towards reliable operation in large-scale quantum computers, our work, therefore, provides opportunities for exploration of hardware efficient generation of quantum magic.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-025-09367-3>.

1. Shor, P. W. Fault-tolerant quantum computation. In *Proc. 37th IEEE Symposium on the Foundations of Computer Science* 56–65 (IEEE, 1996).
2. Aharonov, D. & Ben-Or, M. Fault-tolerant quantum computation with constant error rate. *SIAM J. Comput.* **38**, 1207–1282 (2008).
3. Gottesman, D. An introduction to quantum error correction and fault-tolerant quantum computation. Preprint at <https://doi.org/10.48550/arXiv.0904.2557> (2010).
4. Campbell, E. T., Terhal, B. M., & Vuillot, C. Roads towards fault-tolerant universal quantum computation. *Nature* **549**, 172–179 (2017).
5. Google Quantum AI and Collaborators. Quantum error correction below the surface code threshold. *Nature* **638**, 920–926 (2025).
6. Google Quantum AI. Suppressing quantum errors by scaling a surface code logical qubit. *Nature* **614**, 676–681 (2023).
7. Ryan-Anderson, C. et al. Implementing fault-tolerant entangling gates on the five-qubit code and the color code. Preprint at <https://doi.org/10.48550/arXiv.2208.01863> (2022).
8. Bluvstein, D. et al. Logical quantum processor based on reconfigurable atom arrays. *Nature* **626**, 58–65 (2024).
9. Putterman, H. et al. Hardware-efficient quantum error correction using concatenated bosonic qubits. *Nature* **638**, 927–934 (2025).
10. Sivak, V. V. et al. Real-time quantum error correction beyond break-even. *Nature* **616**, 50–55 (2023).
11. Paetznick, A. et al. Demonstration of logical qubits and repeated error correction with better-than-physical error rates. Preprint at <https://doi.org/10.48550/arXiv.2404.02280> (2024).
12. Eastin, B. & Knill, E. Restrictions on transversal encoded quantum Gate Sets. *Phys. Rev. Lett.* **102**, 110502 (2009).
13. Bravyi, S. & Kitaev, A. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A* **71**, 022316 (2005).
14. Gottesman, D. The Heisenberg representation of quantum computers. In *Proc. XXII International Colloquium on Group Theoretical Methods in Physics* (eds Corney, S. P. et al.) 32–43 (International Press, 1999).
15. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2010).
16. Bluvstein, D. et al. A quantum processor based on coherent transport of entangled atom arrays. *Nature* **604**, 451–456 (2022).

17. Bravyi, S. & Koenig, R. Classification of topologically protected gates for local stabilizer codes. *Phys. Rev. Lett.* **110**, 170503 (2012).

18. Fowler, A. G., Mariantoni, M., Martinis, J. M. & Cleland, A. N. Surface codes: towards practical large-scale quantum computation. *Phys. Rev. A* **86**, 032324 (2012).

19. Bombin, H. & Martin-Delgado, M. A. Topological quantum distillation. *Phys. Rev. Lett.* **97**, 180501 (2006).

20. Litinski, D. Magic state distillation: not as costly as you think. *Quantum* **3**, 205 (2019).

21. Gidney, C. & Fowler, A. G. Efficient magic state factories with a catalyzed $|CCZ\rangle$ to $2|T\rangle$ transformation. *Quantum* **3**, 135 (2019).

22. O'Gorman, J. & Campbell, E. T. Quantum computation with realistic magic-state factories. *Phys. Rev. A* **95**, 032338 (2017).

23. Souza, A. M., Zhang, J., Ryan, C. A. & Laflamme, R. Experimental magic state distillation for fault-tolerant quantum computing. *Nat. Commun.* **2**, 169 (2011).

24. Brown, N. C. et al. Advances in compilation for quantum hardware – a demonstration of magic state distillation and repeat-until-success protocols. Preprint at <https://doi.org/10.48550/arXiv.2310.12106> (2023).

25. Postler, L. et al. Demonstration of fault-tolerant universal quantum gate operations. *Nature* **605**, 675–680 (2022).

26. Ye, Y. et al. Logical magic state preparation with fidelity beyond the distillation threshold on a superconducting quantum processor. *Phys. Rev. Lett.* **131**, 210603 (2023).

27. Gupta, R. S. et al. Encoding a magic state with beyond break-even fidelity. *Nature* **625**, 259–263 (2024).

28. Chamberland, C. & Cross, A. W. Fault-tolerant magic state preparation with flag qubits. *Quantum* **3**, 143 (2019).

29. Gidney, C., Shutt, N. & Jones, C. Magic state cultivation: growing T states as cheap as CNOT gates. Preprint at <https://doi.org/10.48550/arXiv.2409.17595> (2024).

30. Daguerre, L. & Kim, I. H. Code switching revisited: low-overhead magic state preparation using color codes. *Phys. Rev. Research* **7**, 023080 (2025).

31. Cain, M. et al. Correlated decoding of logical algorithms with transversal gates. *Phys. Rev. Lett.* **133**, 240602 (2024).

32. Chao, R. & Reichardt, B. W. Fault-tolerant quantum computation with few qubits. *npj Quantum Inf.* **4**, 42 (2018).

33. Kim, Y., Sevior, M. & Usman, M. Magic state injection on IBM quantum processors above the distillation threshold. Preprint at <https://doi.org/10.48550/arXiv.2412.01446> (2024).

34. Laflamme, R., Miquel, C., Paz, J. P. & Zurek, W. H. Perfect quantum error correcting code. *Phys. Rev. Lett.* **77**, 198–201 (1996).

35. Smith, S. C., Brown, B. J. & Bartlett, S. D. Mitigating errors in logical qubits. *Commun. Phys.* **7**, 386 (2024).

36. Bombin, H., Pant, M., Roberts, S. & Seetharam, K. I. Fault-tolerant post-selection for low overhead magic state preparation. *PRX Quantum* **5**, 010302 (2024).

37. Chao, R. & Reichardt, B. W. Quantum error correction with only two extra qubits. *Phys. Rev. Lett.* **121**, 050502 (2018).

38. Ryan-Anderson, C. et al. Realization of real-time fault-tolerant quantum error correction. *Phys. Rev. X* **11**, 041058 (2021).

39. Graham, T. M. et al. Midcircuit measurements on a single-species neutral alkali atom quantum processor. *Phys. Rev. X* **13**, 041051 (2023).

40. Deist, E. et al. Mid-circuit cavity measurement in a neutral atom array. *Phys. Rev. Lett.* **129**, 203602 (2022).

41. Singh, K. et al. Mid-circuit correction of correlated phase errors using an array of spectator qubits. *Science* **380**, 1265–1269 (2023).

42. Lis, J. W. et al. Midcircuit operations using the *omg* architecture in neutral atom arrays. *Phys. Rev. X* **13**, 041035 (2023).

43. Norcia, M. A. et al. Midcircuit qubit measurement and rearrangement in a ^{171}Yb atomic array. *Phys. Rev. X* **13**, 041034 (2023).

44. Bombin, H. Gauge color codes: optimal transversal gates and gauge fixing in topological stabilizer codes. *New J. Phys.* **17**, 083002 (2015).

45. Kubica, A., Yoshida, B. & Pastawski, F. Unfolding the color code. *New J. Phys.* **17**, 083026 (2015).

46. Chamberland, C. & Noh, K. Very low overhead fault-tolerant magic state preparation using redundant ancilla encoding and flag qubits. *npj Quantum Inf.* **6**, 91 (2020).

47. Hirano, Y., Itogawa, T. & Fujii, K. Leveraging zero-level distillation to generate high-fidelity magic states. In *IEEE International Conference on Quantum Computing and Engineering* 843–853 (IEEE, 2024).

48. Itogawa, T., Takada, Y., Hirano, Y. & Fujii, K. Even more efficient magic state distillation by zero-level distillation. *PRX Quantum* **6**, 020356 (2025).

49. Jones, C. Multilevel distillation of magic states for quantum computing. *Phys. Rev. A* **87**, 042305 (2013).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2025

Methods

System overview

All experiments described in this work were performed on the Gemini-class neutral-atom quantum computer of QuEra. More detailed information and characterization will be described in QuEra Computing and Collaborators (manuscript in preparation, 2025a). The system is based on neutral ^{87}Rb atoms trapped in reconfigurable optical tweezers^{50–53}, following the dynamically reconfigurable architecture described in refs. 8,16. We use laser cooling to trap atoms in a magneto-optical trap, and subsequently load them into a fixed set of optical tweezers generated by a spatial light modulator (SLM). Atoms are then coherently rearranged by dynamic tweezers at 852 nm generated by a crossed pair of acoustic-optic deflectors (AODs).

Qubits are encoded in the $m_f = 0$ hyperfine ground states, $|0\rangle \equiv |F = 1, m_f = 0\rangle$ and $|1\rangle \equiv |F = 3, m_f = 0\rangle$, with a T_2 coherence time of approximately 2 s (QuEra Computing and Collaborators, manuscript in preparation, 2025a). Single-qubit gates are performed by Raman transitions⁵⁴, with a laser red-detuned from $5P_{1/2}$ by 350 GHz. We drive global single-qubit rotations at 650 kHz by illuminating the entire array along the quantization axis and locally at 250 kHz by addressing atoms using another pair of AODs⁸. Two-qubit gates are mediated by Rydberg interactions, which we achieve by driving atoms in $|1\rangle$ to $53S_{1/2}$ in a two-photon process by $6P_{3/2}$ with intermediate state detuning of 6 GHz (refs. 55–59). We perform readout globally, by heating and ejecting atoms in $|1\rangle$ with resonant light, followed by fluorescence imaging of the remaining atoms.

Circuit details and calibration

We deterministically load and prepare the atoms into a rectangular grid of 17×5 SLM traps. The same regular grid of SLM sites is used to run both the $d = 3$ and $d = 5$ distillation experiments (Extended Data Figs. 1 and 2). During circuit execution, atoms are rearranged entirely within the 5-row-wide entanglement zone and illuminated with 1,013 nm and 420 nm light that couples qubit state $|1\rangle$ to the Rydberg state. To execute CZ gates, we coherently move atoms such that gate pairs are 2 μm away, within the Rydberg blockade radius, while keeping 8 μm separation between the independent gate pairs. We perform parallel horizontal moves during data code encoding to prepare a logical qubit in each row. Once encoded, we move the rows of logical qubits using parallel vertical moves. To preserve qubit coherence, all moves are accompanied by dynamical decoupling implemented with global single-qubit pulses. Local single-qubit gates around the XY plane are executed in between the CZ layers, and we echo the induced local light shifts using a global single-qubit gate pulse.

We now detail the changes in the logical operator and stabilizer conventions due to the circuit optimizations that we apply. In both $d = 3$ and $d = 5$ encoding, we physically implement the first layer of \sqrt{Y} gates with a global pulse. We also substitute local \sqrt{Y} with \sqrt{Y} to improve parallelism. This changes the basis of the physical input, so to inject a logical $(1, 1, 1)$, we physically prepare the $(-1, -1, -1)$ state. We further optimize by pre-applying the first set of \sqrt{X} gates required for distillation on the physical qubits before encoding: qubits 1, 2 and 5 are prepared into $(-1, 1, -1)$, which after encoding becomes $(1, -1, 1)$. Finally, single-qubit echoes during the encoding and distillation circuit redefine the colour code stabilizer basis. We classically track this through encoding and distillation to recover the original stabilizer basis. Echoes during distillation also flip the distillation syndrome outcomes. This means that our actual acceptance case is 0011, different from the acceptance case of 1011 from the circuit in Fig. 3a.

We use quantum state tomography to evaluate the fidelity of logical magic states. Measurement of the injected state fidelity is done by applying a global tomography pulse to all qubits and subsequently measuring all three bases (see Extended Data Fig. 3 for examples). To measure the fidelity after distillation, we apply transversal single-qubit

gates to the output logical qubit to sample the three bases. In this way, the other four logical qubits are always measured in the Z basis. Loss during the circuit can lead to a biased error in the magic state fidelity. To mitigate this, we interleave measurements in all basis states: $\pm X, \pm Y, \pm Z$ and average the results. We track the injected fidelity by interleaving one shot of magic state injection with no distillation for every seven shots of full factory execution. This protects against bias due to potential systematic drift during data taking. Throughout the run, we monitor the perfect stabilizer rate of the injection circuit as a proxy for gate calibrations (Extended Data Fig. 3b,d).

A single factory instance using $d = 3$ and $d = 5$ data codes requires 35 and 85 atoms, respectively. To improve data rates, we run two independent parallel instances of the $d = 3$ factory, requiring a total of 70 atoms. We ran 658,562 shots of the $d = 3$ experiment, split into a total of $2 \times 576,131 = 1,152,262$ factory runs and $10 \times 82,431 = 824,310$ encoding tomography runs. The added error datasets consist of 143,000 shots for each added error, corresponding to 251,000 factory runs and 175,000 encoding tomography runs. For $d = 5$, we ran a total of 259,261 factory shots and $5 \times 37,108 = 185,540$ encoding tomography shots. Shots are evenly split into X, Y, Z bases for all experiments.

Error model

We use randomized benchmarking⁶⁰ to calibrate and benchmark single- and two-qubit gates. We measure a global amplitude robust single-qubit gate fidelity of 99.978(1)% and a local amplitude robust single-qubit gate fidelity of 99.978(1)%. We benchmark controlled-phase (CZ) two-qubit gates by driving pairs of blockaded atoms with alternating single-qubit and two-qubit gates⁵⁶. We measure the return probability to $|00\rangle$ as a function of the number of entangling gates, resulting in a fidelity of 99.42(1)% per entangling gate. We estimate the state preparation and measurement errors to be a total of 1%.

To simulate the impact of various error sources on the circuits, we model the error sources as depolarizing Pauli channels. Errors due to global and local single-qubit gates are incorporated as single-qubit channels. Two-qubit gate error is modelled by a two-qubit depolarizing Pauli channel, biased towards Z and ZZ phase flip channels. Movement of atoms by AOD also induces errors, but in two distinct ways. On the moving atoms, tweezers light induces a qubit frequency shift, resulting in Z errors. During the move duration we also account for the idling errors on all qubits. For each error type, we assume uniform error across the atom array and the magnitude of errors is derived from independent benchmarking of each operation. Overall, our error model shows good agreement with the experimentally observed stabilizer and logical outcomes (Extended Data Fig. 4b).

Probing distillation code error suppression

The 5-to-1 MSD achieves a quadratic suppression in infidelity of the input magic states. We probe this phenomenon in experiment and (noiseless) simulation by applying coherent Z errors to input magic states and recording output magic state fidelity (Fig. 3c) as well as factory acceptance rate (Fig. 3d).

In experiments, we apply coherent Z errors of $0.32\pi, 0.24\pi, 0.16\pi$ and 0 to the five physical magic states, which are then injected into $d = 3$ colour codes and distilled. The output fidelity is plotted against the injected post-selected fidelity in Fig. 3c. The distillation output fidelity performs only error correction on the output logical qubit, but post-selects on perfect stabilizers on the distillation syndrome logical qubits. By contrast, the injected post-selected fidelity post-selects on perfect stabilizers on the target logical qubit itself. We choose this comparison to highlight more clearly the distillation behaviour of the logical information. We see that all four data points show distillation gain—namely, the error-corrected output fidelity is higher than the post-selected injected fidelity. In Extended Data Fig. 4c, we further show the results for different distillation stabilizer post-selection thresholds.

We also numerically simulate the performance of the ideal distillation circuit subject to coherent input errors. The output fidelity is plotted against input physical magic state fidelity, calculated based on the applied error (Fig. 3c, grey curve). We observe the expected quadratic suppression in input infidelity, which is in good agreement with our experimental data (blue points). Note that when the input fidelity is at 0.80, which is lower than the frequently quoted 5-to-1 MSD distillation threshold of 83% (ref. 13), we still observe an improved output fidelity. This is because the usual distillation threshold is computed for incoherent errors, whereas the threshold for our applied coherent errors is lower. We also observe good agreement with experimental data for the factory acceptance rate (Fig. 3d). Overall, our experimental data closely align with the theoretical predictions of 5-to-1 MSD.

Comparison with alternative methods for magic state preparation

In this section, we compare different methods for magic state preparation, including alternative injection or projection-based schemes, and other MSD factories.

There are a few natural approaches to preparing logical magic states. We can use an (often non-fault-tolerant) encoding circuit, measure the data code stabilizers to project into the target logical state, or measure certain operators for which the target logical magic state is an eigenstate. Some of these operations can further serve to detect errors to boost the fidelity of the resulting magic state. These protocols can be further expanded with flag qubits or extended with error-correction cycles for improved fidelity.

Previous experiments on trapped ions²⁵ and superconducting qubits^{26,27} have demonstrated different combinations of these techniques. This includes magic state preparation based on unitary encoding²⁵ and stabilizer measurement projection^{26,33}, as well as the further use of flagged verification schemes, error detection, and/or correction to achieve fault tolerance against any single physical error^{25,27}.

The above techniques of magic state injection and verification can produce magic states with fairly high fidelity, which could serve as the input into MSD factories in the future⁴⁷. However, these techniques have some noteworthy limitations because of their direct use of physical operations, in contrast to protected logical operations of the data code as in MSD. First, direct injection of physical magic states without further verification will have a performance limited by the physical magic state fidelity, which is insufficient for large-scale quantum computing. Second, operating the verification protocols at higher distances or higher physical error rates generally increases the complexity of ancilla preparation and/or post-selection overhead significantly^{28,29,46}.

For these reasons, protocols that make use of an inner data code to protect operations, such as MSD, are a crucial primitive as we scale to lower logical error rates. Existing implementations of MSD achieve this with physical qubits^{23,24}, which do not provide protection of Clifford operations within the MSD factory. Thus, our demonstration of logical MSD with an inner data code is a crucial step towards further improvements of magic state preparation.

Our experiment focuses on the implementation of 5-to-1 MSD¹³, as it exemplifies the principles of MSD with relatively low resource requirements. It has the downside that with perfect distillation code operations, the factory only has a 1/6 acceptance rate, and it only achieves quadratic suppression despite the code being distance 3. For future, large-scale operation, it may instead be desirable to use MSD factories with higher distillation rate, better error suppression scaling, and which have an acceptance rate of unity in the absence of input errors^{13,20–22,49}. However, some of the key ingredients we demonstrated, such as the use of parallel operations by transversal gates and sliding-scale post-selection based on stabilizer readouts, are probably broadly useful for future experiments.

Design and optimization of state injection circuits

To achieve high fidelities, we optimize the implementation of several key quantum circuits. In this section, we focus on arbitrary state injection circuits for the $d = 3$ and $d = 5$ colour codes, whereas in the next section we discuss optimizations of logical MSD circuits. We primarily focus on reducing the number of entangling gate layers, because the gate infidelity and associated move errors are an important contributor to our error budget. To the best of our knowledge, previous unitary injection circuits for the $d = 3$ colour code require four entangling gate layers^{25,61,62}. Although extensive search over all possible 7-qubit, 3-layer injection circuits is sufficient for $d = 3$, we develop more efficient methods for $d = 5$ to find low-depth circuits with good atom layouts.

We present an algorithm based on matrix row reduction^{63,64} to find an injection circuit for 2D colour codes. Simple extensions to this algorithm may work well for any CSS code in general. Executing the algorithm gives some injection circuit, which is unlikely to be optimal.

The algorithm operates on a matrix representation of the checks and the logical operators of the code. Each row corresponds to a data qubit of the code and each column is either a check or a logical operator. The matrix entry M_{qc} is 1 if the check or logical c contains qubit q , and is 0 otherwise. Performing row operations (adding one row to another) on this matrix corresponds to the application of CNOTs, whereas performing column operations (adding one column to another) corresponds to redefining stabilizers and logical operators. We find the circuits shown in Figs. 2a and 4a using the following simple heuristics and search methods:

1. Choose row operations in layers, in which we pick the best $n/2$ disjoint pairs of rows for row operations before reusing rows in the next layer. This maximizes circuit parallelism because each row operation will become a CNOT.
2. Every row operation ideally reduces (or sometimes maintains) the total number of 1 entries.
3. A row operation is preferred if it leaves the updated row more similar (by Hamming distance) to another row. This enables a future row operation to be more effective.
4. Prioritize row operations to remove 1 entries from high-weight columns (high relative numbers of 1 entries). If certain columns are very high-weight near row-reduction completion, backtrack and prioritize them sooner.
5. While backtracking to try new choices, prioritize minimizing the number of operation layers over the number of operations.
6. Column operations do not need to be optimal because they do not affect the circuit; they only redefine the stabilizer basis.

For example, the $d = 3$ colour code shown in Fig. 2b has checks $S_0 = Z_0Z_1Z_2Z_3$, $S_1 = Z_1Z_2Z_4Z_5$, $S_2 = Z_2Z_3Z_4Z_6$, $S_3 = X_0X_1X_2X_3$, $S_4 = X_1X_2X_4X_5$ and $S_5 = X_2X_3X_4X_6$ and logical operators $L_{z0} = Z_0Z_1Z_5$ and $L_{x0} = X_0X_1X_5$. Owing to the self-dual structure where X and Z checks match, this can be represented with the matrix

$$M^0 = \begin{matrix} & S_0 & S_1 & S_2 & L_0 \\ q_0 & 1 & - & - & 1 \\ q_1 & 1 & 1 & - & 1 \\ q_2 & 1 & 1 & 1 & - \\ q_3 & 1 & - & 1 & - \\ q_4 & - & 1 & 1 & - \\ q_5 & - & 1 & - & 1 \\ q_6 & - & - & 1 & - \end{matrix}$$

where zero is shown as ‘-’ for visual clarity.

Our goal is to find a sequence of row and column operations by matrix row reduction under addition modulo 2. For example, the row operation $0 \rightarrow 2$ results in

Article

	S_0	S_1	S_2	L_0
q_0	1	—	—	1
q_1	1	1	—	1
q_2	—	1	1	1
q_3	1	—	1	—
q_4	—	1	1	—
q_5	—	1	—	1
q_6	—	—	1	—

Note that columns corresponding to the logical operator(s) must not be source columns (for example, $L_i \rightarrow$ any column is not allowed), but may be target columns.

We find the best row operations $R_{\text{ops}} = [0 \rightarrow 1, 3 \rightarrow 2, 5 \rightarrow 4, 0 \rightarrow 3, 2 \rightarrow 5, 4 \rightarrow 6, 2 \rightarrow 1, 4 \rightarrow 3, 6 \rightarrow 5]$ and column operations $C_{\text{ops}} = [S_0 \rightarrow L_{z0}, S_2 \rightarrow L_{z0}]$ result in our final matrix M^{final}

	S_0	S_1	S_2	L_0
q_0	1	—	—	—
q_1	—	—	—	—
q_2	—	1	—	—
M^{final}	q_3	—	—	—
	q_4	—	—	1
	q_5	—	—	—
	q_6	—	—	1

This solution explicitly defines our encoding circuit:

1. On the qubit, where $M_{q_i, L_{zj}}^{\text{final}} = 1$, prepare the injected state $|\psi\rangle$.
2. For other q_i where $M_{q_i, S_j}^{\text{final}} = 1$ for some S_j , prepare q_i in the $|+\rangle$ state.
3. Prepare all other qubits in the $|0\rangle$ state.
4. For each entry $s \rightarrow t$ of R_{ops} , in reverse order, add a CNOT gate with control q_s and target q_t . C_{ops} does not affect the circuit.
5. Use circuit identities to convert to hardware-supported gates:
 - (a) Preparing a qubit in the $|+\rangle$ state becomes preparing in the $|0\rangle$ state followed by \sqrt{Y} .
 - (b) CNOT _{ij} becomes $\sqrt{Y_j}^\dagger, \text{CZ}_{ij}$ and $\sqrt{Y_j}$.
 - (c) Adjacent $\sqrt{Y_j}$ and $\sqrt{Y_j}^\dagger$ cancel.

The resulting circuit non-fault-tolerantly prepares the logical code state $|\bar{\psi}\rangle$. For the $d=3$ colour code, the solution above gives the 9-gate, 3-layer encoding circuit. For the $[[17, 1, 5]]$ colour code, the best solution we find has five layers and 24 CNOT/CZ gates (Fig. 4a) with $R_{\text{ops}} = [1 \rightarrow 0, 3 \rightarrow 2, 4 \rightarrow 5, 7 \rightarrow 6, 9 \rightarrow 8, 15 \rightarrow 12, 2 \rightarrow 0, 6 \rightarrow 3, 8 \rightarrow 5, 12 \rightarrow 10, 13 \rightarrow 11, 2 \rightarrow 4, 8 \rightarrow 6, 9 \rightarrow 7, 10 \rightarrow 13, 16 \rightarrow 14, 4 \rightarrow 7, 8 \rightarrow 10, 14 \rightarrow 11, 15 \rightarrow 16, 3 \rightarrow 1, 7 \rightarrow 10, 14 \rightarrow 12, 16 \rightarrow 13]$.

Design and optimization of distillation circuit

The 5-to-1 distillation protocol consists of running the un-encoding circuit of the $[[5, 1, 3]]$ perfect code, followed by measurements of the four logical qubits that correspond to stabilizers of the distillation code. To implement this protocol, we start with an un-encoding circuit with low entangling gate count (Gidney, C., tweet), previously optimized from ref. 34, and further optimize it for our hardware. Our optimizations aim to reduce the number of local single-qubit gates as well as the number of entangling gate layers, because these have larger contributions to the infidelity. We use a variety of techniques to achieve this:

1. Reordering of qubits and commuting gates. The final circuit includes three rounds of CZ gates separated by local single-qubit gates.
2. Use of circuit identities, such as $\frac{1+i}{\sqrt{2}}H = X^{1/2}SX^{1/2}$.
3. Absorbing certain operations into the initial state or measurement, without changing the ideal initial state, the post-selection basis, or affecting the quadratic error suppression of the distillation circuit.

Our optimized circuit is shown in Fig. 3a. Note that in standard 5-to-1 MSD, the successful distillation syndrome is 0000, as shown in Fig. 1. Our optimizations flipped it to 1011. As we used identities related to the initial state inputs and final post-selection, this circuit is an un-encoding circuit of a 5-qubit code, which is equivalent to the perfect code up to Clifford operations. For MSD, it achieves quadratic suppression in infidelity.

Design and optimization of atom layout

The optimized circuits described in the previous two sections need efficient implementations of atom movement. Here, we describe our design process for finding circuit-specialized atom move sequences.

We design logical circuits with transversal operations to have a 2D product structure, in which transversal operations are horizontally parallel and logical state injection is vertically parallel (Fig. 1). Thus, we lay out each logical qubit linearly in the same row. All atoms have a home position in a static SLM trap, and for each layer of gates, we pick fewer than half the atoms, move them horizontally or vertically near their gate partners, and move them back. To minimize atom transfer, we optimize for an atom ordering and circuit layers in which none of the moves reorder atoms and the move distances are minimized. An atom order is valid for given circuit layers if

$$\begin{aligned} \max(i, j) &< \max(k, l) \\ \forall \text{CZ}_{ij}, \text{CZ}_{kl} \in \text{Layer} \quad &|i| < k \\ \forall \text{Layer} \in \text{Circuit} \end{aligned}$$

where a layer is a set of CZ gates that may be executed in parallel without changing the meaning of the circuit.

We use a combination of hand-optimization over choice of circuit layers and brute-force search over atom orders. The index numbers labelled in the encoding circuit show these optimal qubit orders. The order of the five (logical) qubits in Fig. 3a has the first two qubits swapped.

Approach to simulation of MSD circuit performance

Our full circuit, which injects five physical magic states into five logical magic states in the colour code, and then performs logical MSD, is supported on 35 qubits in the $d=3$ case and 85 qubits in the $d=5$ case. The injection and distillation circuits are entirely Clifford, with the non-Cliffordness coming only from the input states. This poses a challenge towards using standard simulation methods. As the input states are magic states, standard Clifford circuit simulation tools such as Stim⁶⁵ cannot be applied directly. The circuit size of 85 means state-vector simulation is intractable, and approximation methods such as matrix product states simulations become technically and computationally consuming. Although methods such as extended Clifford simulation could be used⁶⁶, existing open-source implementations only support up to 64 physical qubits⁶⁷. For these reasons, we developed a simulation technique, which we refer to as Input-Decoupled Noise Learning (QuEra Computing and Collaborators, manuscript in preparation, 2025b), in which learning of the noise channel is separated from simulating the actual state of the logical circuit.

The key idea of our approach is that the analysis of noise can largely be separated from the analysis of the ideal logical action itself. The ideal logical circuit can be viewed as a channel that maps some input quantum state to classical bit strings, $\mathcal{C} : (\mathbb{C}^2)^{\otimes 5} \rightarrow \mathbb{F}_2^5$, where the bit strings correspond to logical measurement outcomes. As this is an ideal logical circuit involving only five qubits, it can be readily simulated. Under a Pauli noise model σ and a Clifford circuit C , the combined effect of noise and error correction is to apply additional logical Pauli operations, which further map the logical outcomes $\mathcal{E}_\sigma : \mathbb{F}_2^5 \rightarrow \mathbb{F}_2^5$. As this combined effect involves only Pauli operators and Clifford circuits, we can efficiently simulate it by error sampling and decoding.

More generally, the same approach can be applied whenever error sampling and decoding can be done efficiently. The full simulation result can then be obtained by composing the two channels $\mathcal{E}_o \circ \mathcal{C}$.

Learning the channel \mathcal{E}_o is implemented as follows. We use Stim⁶⁵ to simulate the noisy logical circuit, but we replace the input physical magic states by a special 5-qubit entangled state chosen to make the logical measurement results deterministic. This special 5-qubit state is generated by running the noise-free inverse of the logical circuit. When measuring the output logical in the X , Y , or Z basis, the state-prep for the special state begins by preparing $|+_L\rangle$, $|+_L\rangle$ or $|0_L\rangle$ of the 5-qubit distillation code, respectively. This ensures that the measurement results on all five logical qubits are $+1$ in the absence of errors and that the simulation is fully Clifford and therefore efficient. We perform decoding based on the simulated syndromes \tilde{S} (see the following sections for details of our decoder), resulting in the final logical measurement result $l \in \mathbb{F}_2^5$, characterizing logical flips caused by circuit noise. The decoding is done either using the syndrome information of the four logical qubits of factory post-selection or using the syndrome information of all five logical qubits during tomography, resulting in the appropriate channel in each case. With a large amount of samples, which can be efficiently generated, we can learn the classical logical error channel \mathcal{E}_o to high accuracy.

For our logical circuit involving five logical qubits, we can easily obtain the ideal logical circuit channel \mathcal{C} . Using Qiskit⁶⁷, we implement the ideal distillation circuit with magic state inputs. Note that this simulation also supports noise applied to the magic state input. We directly calculate the 5-bit logical output for this circuit, producing the channel \mathcal{C} . The final output magic state fidelity can then be computed by composing the channels $\mathcal{E}_o \circ \mathcal{C}$.

With this approach, for a Clifford physical circuit with non-Pauli inputs, our method decouples the Pauli noise in the physical circuit from the input states and learns the noise-induced logical errors efficiently. Beyond this example, we expect our techniques to have further applications as we scale to larger quantum codes and more complex logical circuits.

Estimation of confidence intervals

When performing quantum state tomography to estimate the logical fidelity, it is possible that the reconstructed density matrix is not positive semi-definite, causing the calculated fidelity confidence interval to exceed 1 (ref. 68). To address this and obtain meaningful confidence intervals, we use Bayesian analysis to calculate posterior probabilities⁶⁹.

Consider quantum state tomography, with $\mathbf{n} = (n_x, n_y, n_z)$ measurements in the X , Y , Z basis, respectively. Denote the number of $|0\rangle$ outcomes decoded as $\mathbf{m} = (m_x, m_y, m_z)$. We would like to extract the probability distribution of true fidelity values \mathcal{F} that could produce these measurement results. To this end, we apply Bayes's rule:

$$P(\mathcal{F}=F|\mathbf{m}, \mathbf{n}) = \frac{P(\mathbf{m}, \mathbf{n}|\mathcal{F}=F)P_{\text{prior}}(\mathcal{F}=F)}{P(\mathbf{m}, \mathbf{n})}, \quad (1)$$

where \mathbf{m} and \mathbf{n} denote the observed measurement outcomes and P_{prior} is the prior distribution. For Figs. 3b and 4c, we use the Bures prior, as recommended in ref. 70. For the remaining data, we assume a simpler prior distribution of density matrices that has a uniform random distribution within the Bloch sphere, because the large number of samples causes the difference between different priors to have a negligible effect (<0.1%) on them. See Extended Data Fig. 5 for a comparison of estimation methods.

We (numerically) compute this distribution over the Bloch sphere using

$$P(\mathbf{m}, \mathbf{n}|\mathcal{F}=F) = \int_{\mathbf{v}|F_{|\text{SH}}(\mathbf{v})=F} P(\mathbf{m}, \mathbf{n}|\mathbf{v})f(\mathbf{v})d\mathbf{v}$$

$$P_{\text{prior}}(\mathcal{F}=F) = \int_{\mathbf{v}|F_{|\text{SH}}(\mathbf{v})=F} f(\mathbf{v})d\mathbf{v}$$

where $\mathbf{v} = (x, y, z)$ is the Bloch sphere vector representation of the density matrix,

$$F_{|\text{SH}}(\mathbf{v}) = \frac{1}{2} + \frac{x+y+z}{2\sqrt{3}} \quad (2)$$

is the fidelity of the mixed state \mathbf{v} relative to our desired $|\text{SH}\rangle$ magic state $\mathbf{v}_{|\text{SH}} = \frac{1}{\sqrt{3}}(1, 1, 1)$ and $f(\mathbf{v})$ denotes the probability density function of the prior. We label our magic state as $|\text{SH}\rangle$, as it is the eigenstate of the single-qubit logical gate SH (a product of the S and H gates). Intuitively, we integrate over all mixed states with the same fidelity.

Decoding and post-selection methods

At the end of our MSD protocol, we transversally measure all physical qubits of the four distillation syndrome logical qubits in the Z basis, and all physical qubits of the output magic state in one of the X , Y or Z basis for logical tomography. We use two decoding methods for our data: an MLD constructed by direct sampling of a lookup table, and an MLE decoder based on mixed-integer programming⁷¹.

Given either decoder, we first perform decoding using only the syndromes of the four ancillary logical qubits to infer their logical outcomes. The syndromes of the output logical qubit are not used at this stage, because the factory post-selection should be done without measuring the output logical qubit, so that the output can continue to be used in subsequent logical operations. We perform factory post-selection on the distillation logical outcome being 0011, which is the desired outcome for our distillation circuit. We then optionally perform further stabilizer post-selection, which can further boost the fidelity of the output magic state by flagging bad executions of the distillation circuit. After post-selection, we decode with the full syndrome of all five logical qubits to infer the logical outcome of the output magic state (with no further post-selection), which we use to compute the output fidelity. We note that agreement between the results of the two rounds of decoding (four compared with five logical qubits) could be further used to herald logical errors in the execution of the full circuit.

The MLD decoder is only tractable for $d = 3$, in which the number of syndrome combinations is limited. To construct the MLD decoder, we sample 10^9 measurement samples for our full 35-qubit circuit, under the noise model described above. Our lookup table T will have 2^{15} keys corresponding to all possible syndromes, each key storing 2^5 entries corresponding to the number of occurrences of each logical observable pattern among our samples with the given syndrome. Sampling can be done efficiently in Stim, by replacing the input magic states with stabilizer states (see our noise learning method described above). With each sample, we store the 15-bit syndrome information and the 5-bit logical error string into T . After all samples are collected, each syndrome s will have a most-likely logical error ℓ_s , which will be our decoder output for s . To perform post-selection, we can sort the stabilizer patterns based on the logical fidelity of the output they lead to and perform sliding-scale post-selection based on this.

The MLD decoding method described in the previous paragraph works only for small code distances, because the space complexity for the table is exponential. Therefore, it is not realistic to use it for decoding at $d = 5$. We, therefore, use an MLE decoder, adapted from refs. 31,72, to decode the logical measurement results and evaluate their confidence for post-selection.

We construct an MLE decoder based on mixed-integer programming (MIP) formalized as follows. We denote all stabilizers as $\Sigma = \{\sigma_1, \dots, \sigma_k\}$, and all logical observable as $O = \{O_1, \dots, O_d\}$. We enumerate all possible elementary Pauli errors $\mathcal{E} = \{\epsilon_1, \dots, \epsilon_m\}$ in the injection and distillation circuits, and each error ϵ_j can flip a subset of stabilizers $\Sigma_j \subset \Sigma$ as well as a subset of logical observables $O_j \subset O$ with probability p_j . If we define

$$\delta_{i,j} = 1 \text{ if } \sigma_i \in \Sigma_j, L_{i,j} = 1 \text{ if } O_i \in \Omega_j, \quad (3)$$

then given an error configuration $\mathbf{e} \in \mathbb{F}_2^m$, the resulting stabilizer and observable configuration will be $\partial\mathbf{e}$ and $L\mathbf{e}$, respectively. The input of

the MLE decoder is a stabilizer configuration $\mathbf{s} = \begin{pmatrix} s_1 \\ \vdots \\ s_k \end{pmatrix} \in \mathbb{F}_2^k$, and it will

return the MLE configuration that results in the same stabilizer configuration. More precisely, the MLE is defined

$$\mathbf{e}_{\text{MLE}} = \text{argmax } P(\mathbf{e}), \text{ such that } \partial\mathbf{e} = \mathbf{s}, \quad (4)$$

where e_j, s_j are binary variables. Equivalently, the MLE can be determined by the following MIP problem by regarding all variables as integers and introducing new slack variables λ_j :

$$\mathbf{e} = \text{argmax } \sum_j \log \frac{1-p_j}{p_j} e_j, \text{ s.t. } \sum \delta_{i,j} e_j = s_j + 2\lambda_j, \quad (5)$$

where e_j, s_j, λ_j are integers.

To post-select a shot based on the stabilizer configuration, we analyse the logical gap^{35,36,73,74}, which characterizes the confidence in the chosen correction. We seek to characterize the confidence by analysing the likelihood of this error compared with those resulting in other logical corrections. We define the second most-likely error (SMLE) to be

$$\mathbf{e}_{\text{SMLE}} = \text{argmax } P(\mathbf{e}), \text{ such that } \partial\mathbf{e} = \sigma \text{ and } L\mathbf{e} \neq L\mathbf{e}_{\text{MLE}}, \quad (6)$$

then the logical gap of a given stabilizer configuration is defined as

$$g = \log \frac{P(\mathbf{e}_{\text{MLE}})}{P(\mathbf{e}_{\text{SMLE}})}. \quad (7)$$

The logical gap provides a confidence measure for decoding—the gap approximates the likelihood difference between the most-likely logical outcome and the second most-likely logical outcome.

In the case of 5-to-1 distillation, there are four logical qubits that are measured, and we will use the stabilizer information from those four to post-select the shots. We enumerate all 2^4 logical representatives over these four measured qubits and add the corresponding logical observable as a new constraint into the MIP solver to obtain the MLE and the SMLE. To determine whether we accept a shot, we compute the logical gap based on the detector information on the measured four logical qubits, and see if it is greater than a logical gap threshold we set ahead of time.

We observe that for $d = 3$, the logical error performance for the MLD and MLE is comparable (Fig. 3b). This suggests that the additional entropic contribution from considering all error cosets is smaller than that coming from analysing the MLE itself.

Physical error rate to achieve distillation gain without stabilizer post-selection

We now perform numerical simulations of our $d = 3$ distillation process at a variety of different physical error rates and evaluate the injected and distilled magic state fidelities in the absence of stabilizer post-selection. This provides an estimation of how much the physical error rate should be improved to see distillation gain without extra post-selection penalties, and future work can extend this to a comparison between different code distances.

The results are shown in Extended Data Fig. 4a. As we globally rescale the physical error rate, both the injected and distilled fidelity improve, with the distilled fidelity improving faster due to its quadratic scaling. To match the experimentally observed fidelities (crosses), the physical error rates are rescaled by 1.25 times compared with the error model used for decoding. We find that an approximately twofold improvement

in physical error rate suffices to achieve distillation gain without stabilizer post-selection.

Data availability

Peer reviewer reports are available. All data supporting the findings of this study are available from Zenodo at <https://doi.org/10.1038/s41586-025-09367-3> (ref. 75).

50. Barredo, D., de Léséleuc, S., Lahaye, T. & Browaeys, A. An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science* **354**, 1021–1023 (2016).
51. Endres, M. et al. Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. *Science* **354**, 1024–1027 (2016).
52. Kim, H. et al. *In situ* single-atom array synthesis using dynamic holographic optical tweezers. *Nat. Commun.* **7**, 13317 (2016).
53. Ebadi, S. et al. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature* **595**, 227–232 (2021).
54. Levine, H. et al. Dispersive optical systems for scalable Raman driving of hyperfine qubits. *Phys. Rev. A* **105**, 032618 (2022).
55. Levine, H. et al. Parallel implementation of high-fidelity multiqubit gates with neutral atoms. *Phys. Rev. Lett.* **123**, 170503 (2019).
56. Evered, S. J. et al. High-fidelity parallel entangling gates on a neutral-atom quantum computer. *Nature* **622**, 268–272 (2023).
57. Ma, S. et al. High-fidelity gates and mid-circuit erasure conversion in an atomic qubit. *Nature* **622**, 279–284 (2023).
58. Jandura, S. & Pupillo, G. Time-optimal two- and three-qubit gates for Rydberg atoms. *Quantum* **6**, 712 (2022).
59. Scholl, P. et al. Erasure conversion in a high-fidelity Rydberg quantum simulator. *Nature* **622**, 273–278 (2023).
60. Emerson, J. et al. Scalable noise estimation with random unitary operators. *J. Opt. B Quantum Semiclassical Opt.* **7**, S347 (2005).
61. Goto, H. Minimizing resource overheads for fault-tolerant preparation of encoded states of the Steane code. *Sci. Rep.* **6**, 19578 (2016).
62. Mayer, K. et al. Benchmarking logical three-qubit quantum Fourier transform encoded in the Steane code on a trapped-ion quantum computer. Preprint at <https://doi.org/10.48550/arXiv.2404.08616> (2024).
63. Gottesman, D. E. *Stabilizer Codes and Quantum Error Correction*. PhD thesis, California Institute of Technology (1997).
64. Cleve, R. & Gottesman, D. Efficient computations of encodings for quantum error correction. *Physical Review A* **56**, 76–82 (1997).
65. Gidney, C. Stim: a fast stabilizer circuit simulator. *Quantum* **5**, 497 (2021).
66. Bravyi, S. et al. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum* **3**, 181 (2019).
67. Javadi-Abhari, A. et al. Quantum computing with Qiskit. Preprint at <https://doi.org/10.48550/arXiv.2405.08810> (2024).
68. Granade, C., Ferrie, C., & Flammia, S. T. Practical adaptive quantum tomography. *New J. Phys.* **19**, 113017 (2017).
69. Rice, J. A. *Mathematical Statistics and Data Analysis*, 3rd edn, Vol. 371 (Thomson Brooks/Cole, 2007).
70. Schmid, R. Quantum state tomography of a single qubit: comparison of methods. *J. Mod. Optics* **63**, 1744–1758 (2016).
71. Gurobi Optimization. *Gurobi Optimizer Reference Manual* (Gurobi Optimization, 2024).
72. Landahl, A. J., Anderson, J. T. & Rice, P. R. Fault-tolerant quantum computing with color codes. Preprint at <https://doi.org/10.48550/arXiv.1108.5738> (2011).
73. Meister, N., Patterson, C. A. & Preskill, J. Efficient soft-output decoders for the surface code. Preprint at <https://doi.org/10.48550/arXiv.2405.07433> (2024).
74. Gidney, C., Newman, M., Brooks, P. & Jones, C. Yoked surface codes. *Nat. Commun.* **16**, 4498 (2025).
75. QuEra Computing and collaborators. Data for “Experimental demonstration of logical magic state distillation”. Zenodo <https://doi.org/10.1038/s41586-025-09367-3> (2025).

Acknowledgements We acknowledge the discussions with L. Jiang, M. Kang, G. Masella, C. Patterson, A. Piñeiro Orioli, Q. Xu and M. Yuan and technical contributions from I. Paus, C. Skinker and Q. Yu. This work, including the design, assembly and operation of the Gemini-class neutral-atom quantum computer, was supported by QuEra Computing. Pathfinding work at Harvard and MIT was supported by IARPA and the Army Research Office, under the Entangled Logical Qubits program (cooperative agreement no. W911NF-23-2-0219), the DARPA ONISQ program (W911NF2010021), the DARPA MeasQuIT program (HR0011-24-9-0359), the Center for Ultracold Atoms (an NSF Physics Frontiers Center, PHY-1734011), the National Science Foundation (grant no. PHY-2012023 and grant no. CCF-2313084), the Army Research Office MURI (grant no. W911NF-20-1-0082) and QuEra Computing. Z.H. acknowledges support from the NSF Graduate Research Fellowship Program (grant no. 2141064). D.B. acknowledges support from the NSF Graduate Research Fellowship Program (grant no. DGE1745303) and The Fannie and John Hertz Foundation. S.J.E. acknowledges support from the National Defense Science and Engineering Graduate (NDSEG) fellowship. T. Manovitz acknowledges support from the Harvard Quantum Initiative Postdoctoral Fellowship in Science and Engineering. M.C. acknowledges support from the Department of Energy Computational Science Graduate Fellowship (award no. DE-SC0020347).

Author contributions The QuEra Computing staff designed, built and ran the experiment and performed data analysis. P.S.R., J.M.R., P.N.J., Z.H., C.D., C.Z., K.-H.W., J.C., K.B., M. Kwon, T. Karolishyn, P.W., M.C., S.J.E., A.A.G., M. Kalinowski, S.H.L., T. Manovitz, J.A.-G., J.I.B., L.B., B.B., A.B., A.C., R.J.D., F.F., C.F., P.F., D.H., M. Hamdan, J.H., N.H., M.-G.H., F.H., N.J., D.K., M. Kornjač, A.

F.L., J. Long, J. Lopatin, P.L.S.L., X.-Z.L., T. Macri, O.M., L.A.M.-M., X.M., S.O., E.O., D.P., Z.Q., V.S., A.S., M.S., N.S., H.T., N.W., Y.W., D.W.-L., T.W., J.W., A.Z., L.Z., M.G., A.K., N.G., V.V., T. Kitagawa, S.-T.W., D.B., M.D.L., A.L., H.Z. and S.H.C. discussed the results, revised and contributed to the writing of this paper.

Competing interests M.G., V.V. and M.D.L. are co-founders and shareholders of QuEra Computing. Authors affiliated with QuEra Computing are employees or interns at QuEra Computing at the time of their contributions.

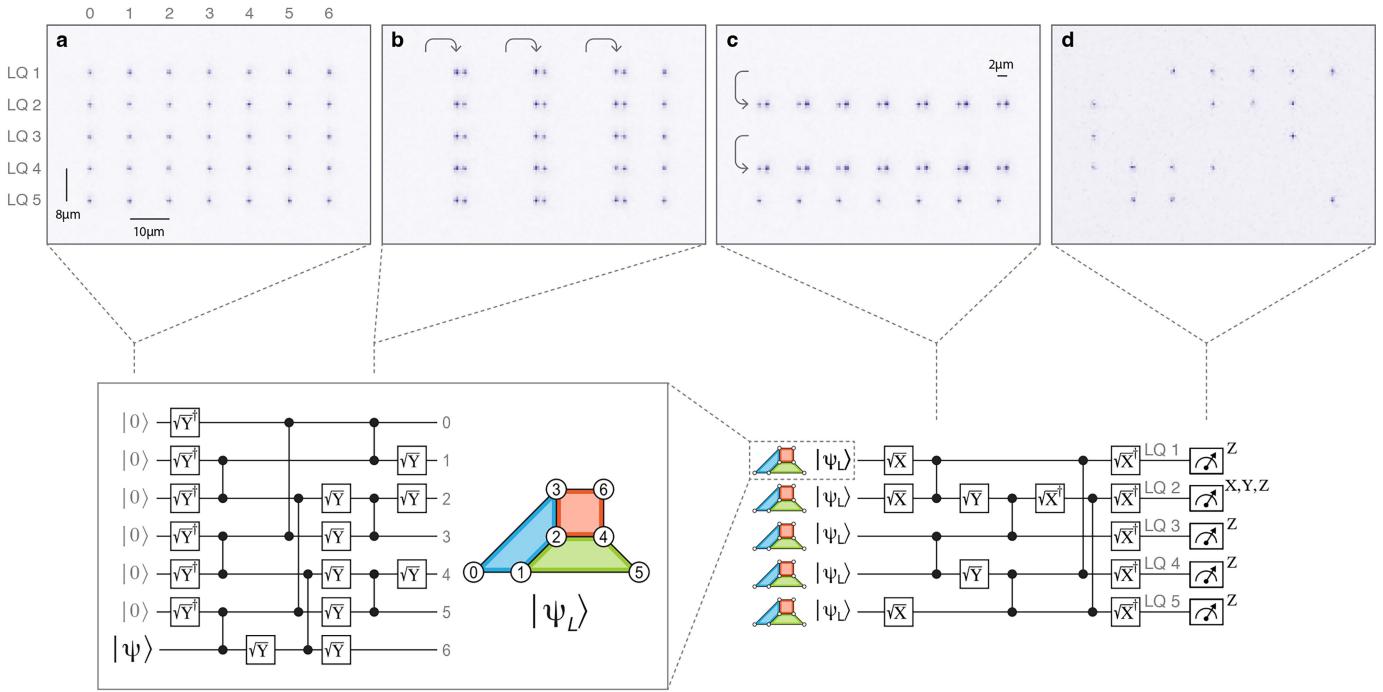
Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-025-09367-3>.

Correspondence and requests for materials should be addressed to Hengyun Zhou or Sergio H. Cantú.

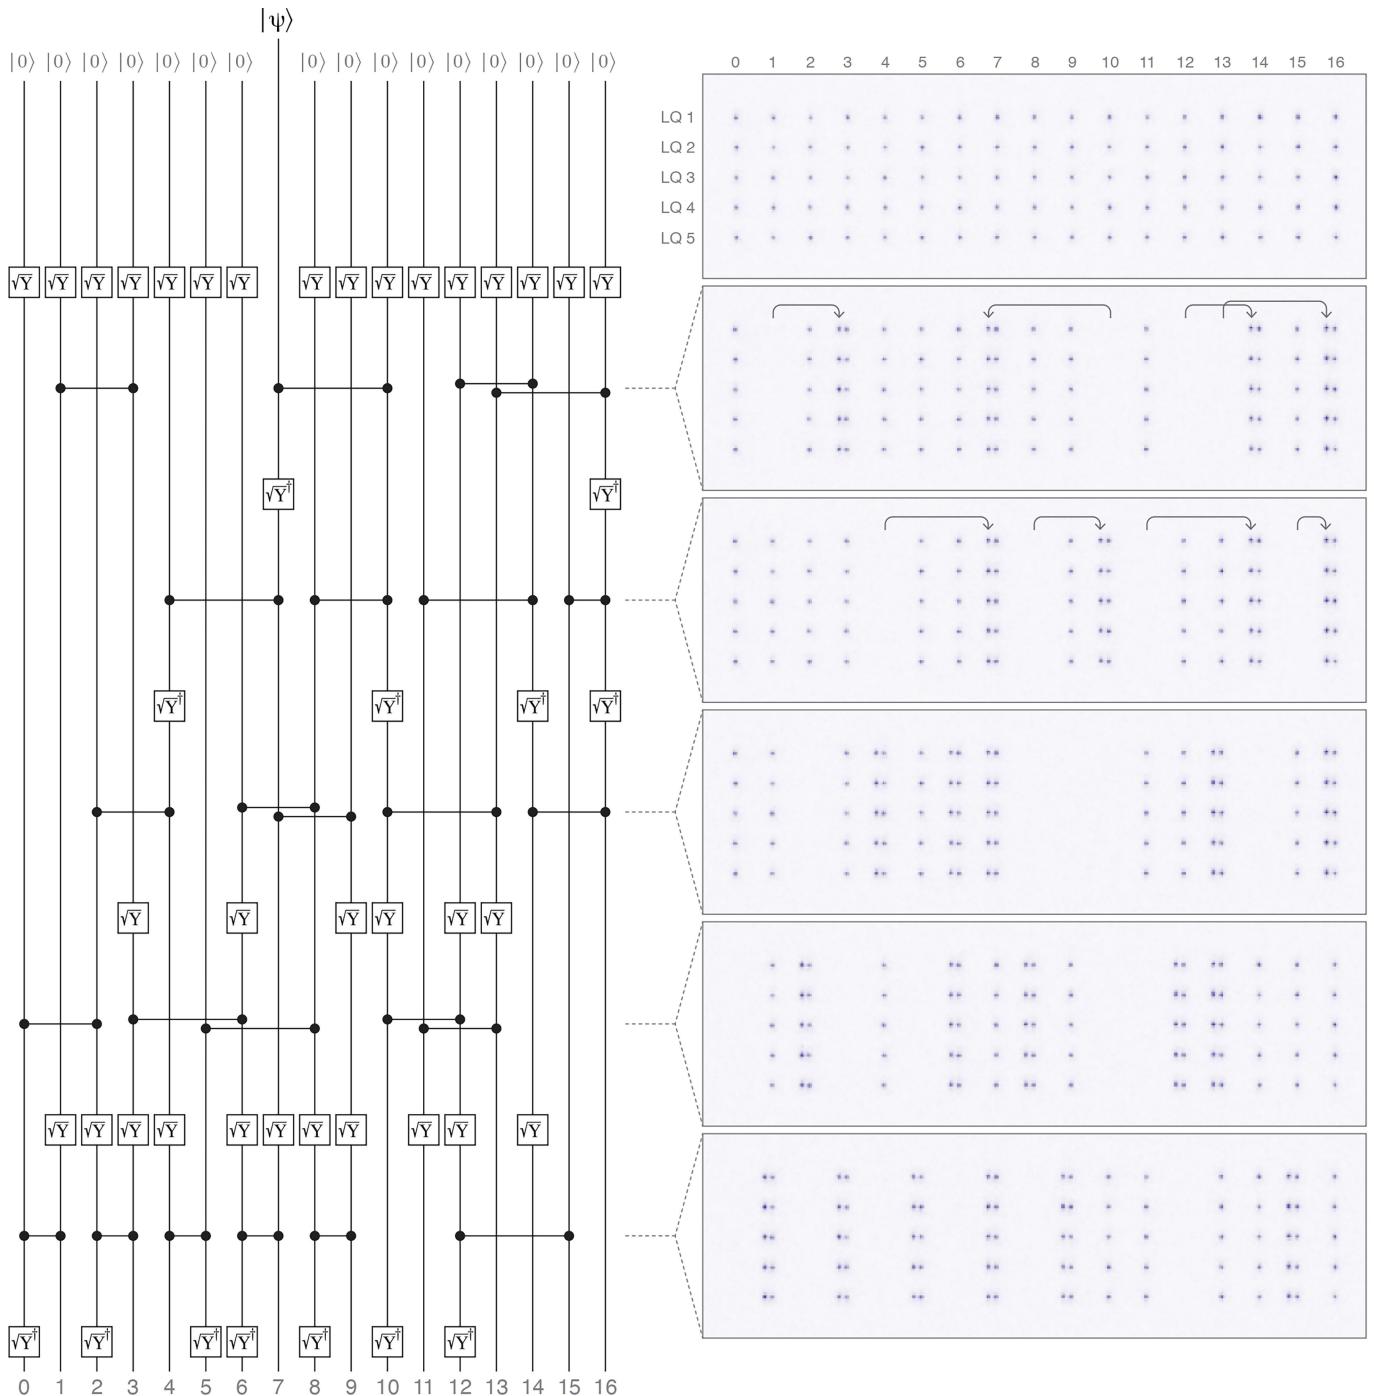
Peer review information *Nature* thanks Michael Vasmer and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at <http://www.nature.com/reprints>.



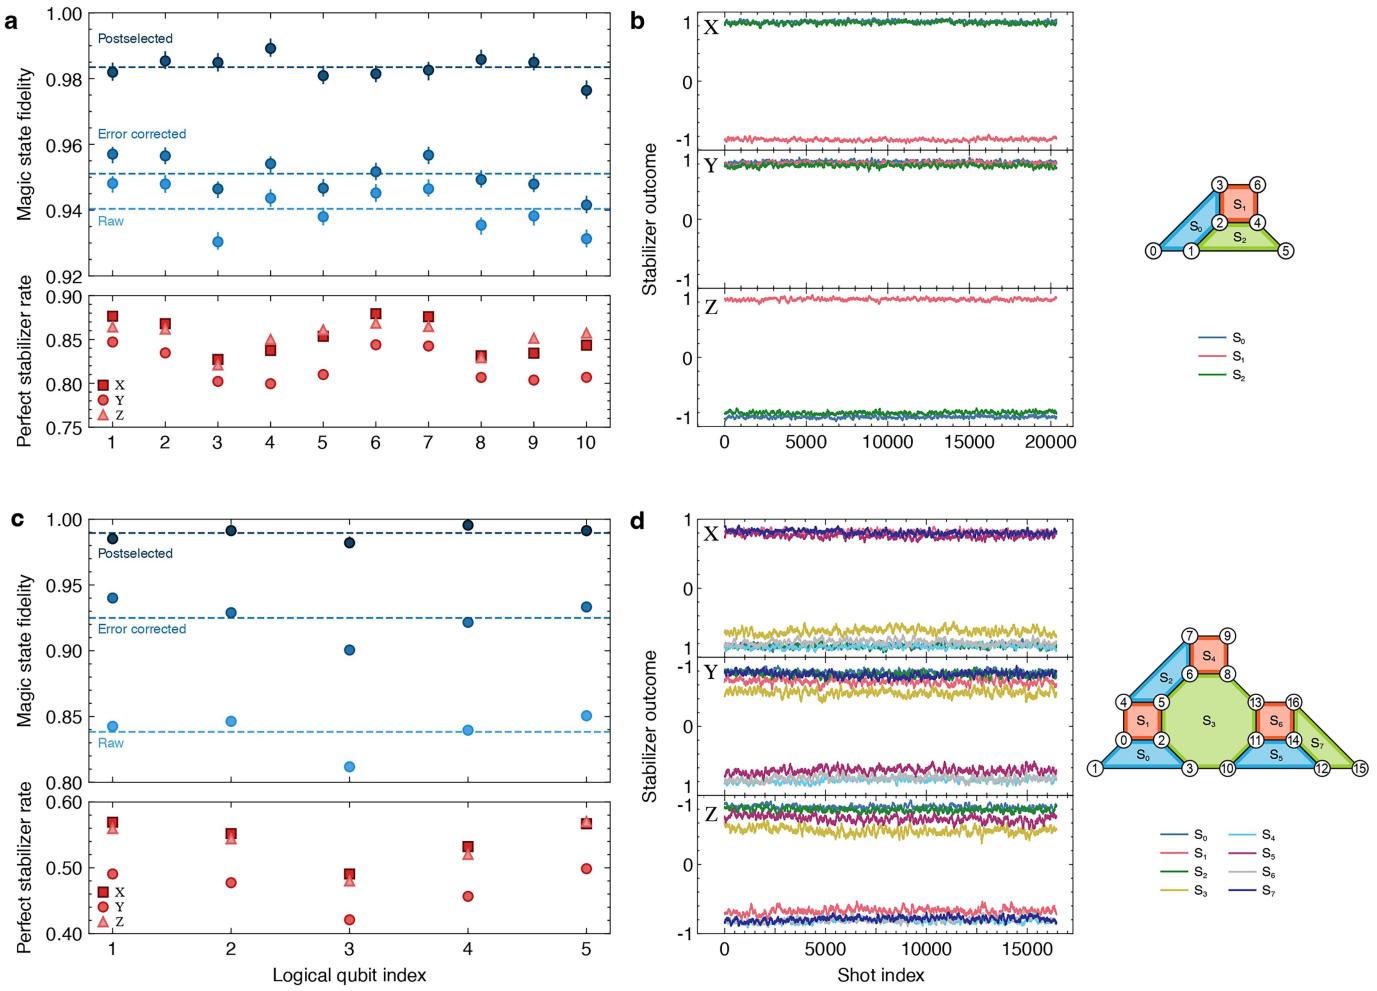
Extended Data Fig. 1 | Experimental layout of magic state distillation factory. **a**, We arrange 7 to 17 ^{87}Rb atoms, each corresponding to a physical qubit, into a row. This horizontal register represents a logical qubit, tiled into 5 rows for a total of five logical qubits (LQ1 to LQ5). **b**, Encoding. Once the register of physical qubits is prepared, we coherently rearrange atoms to perform two-qubit entangling gates using the Rydberg blockade mechanism. We break

up the circuit into “layers” each containing one set of local rotations, transport, and CZ gates. **c**, Coherent movement of logical qubits to perform transversal CZ gates. In the case of 5-to-1 distillation, this is achieved in three layers. The circuit as drawn here corresponds 1 to 1 to the atom layout, whereas in Fig. 3 logical qubits LQ1 and LQ2 are swapped for clarity. **d**, Global measurement of qubits after circuit execution.



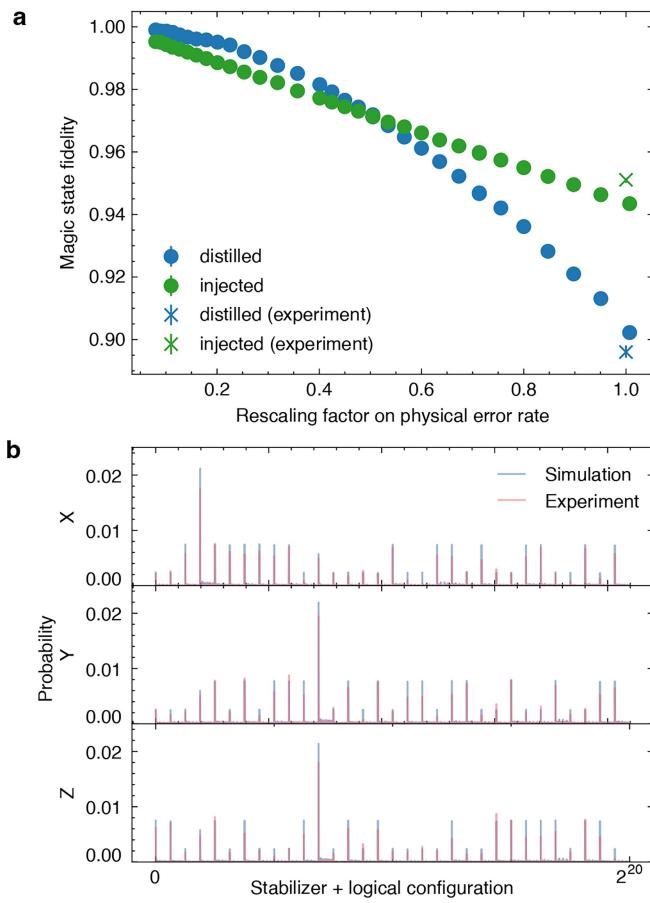
Extended Data Fig. 2 | Experimental layout of $d=5$ encoding. The arbitrary-state encoding circuit for the $d=5$ color code (left) is comprised of five entangling gate layers, illustrated by averaged images of the corresponding atom configurations (right), and local gates between the layers. We execute encoding with 5x parallelism, one instance per row (LQ1 to LQ5). The horizontal

AOD trap array is tiled vertically by the second AOD. For each layer, atoms start in SLM sites, we apply local rotations, pick up and move atoms to their gate location, execute parallel CZ gates, echo (omitted for clarity), and finally move back to SLM sites.

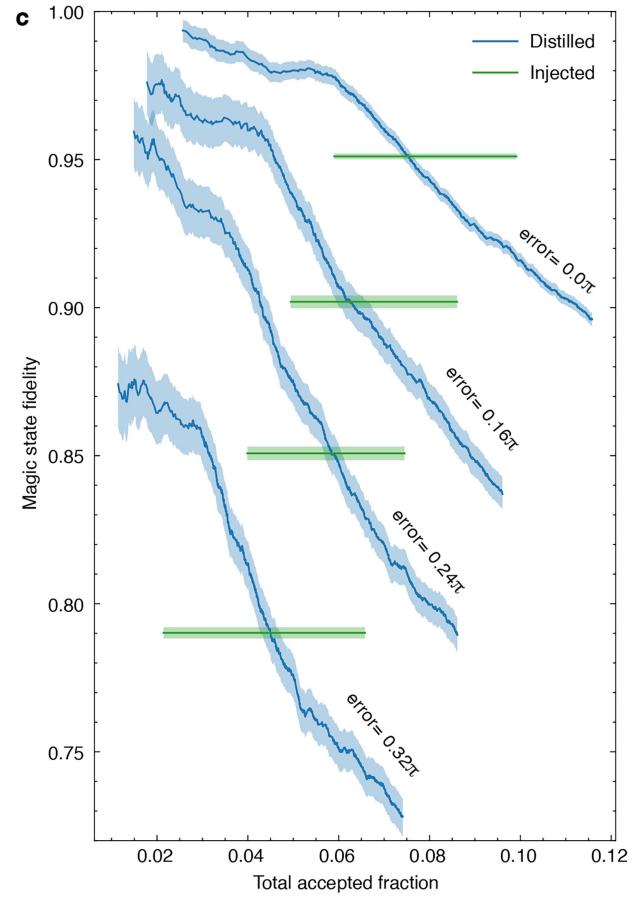


Extended Data Fig. 3 | Encoded magic state fidelity and stabilizers. **a**, Spatial dependence of distance-3 magic state encoding fidelity, for the experimental run with no added coherent error. Logical qubits numbered 1-5 and 6-10 are the input qubits to two parallel distillation circuits. We observe some spatial dependence on both the fidelity and perfect stabilizer rate, which we attribute

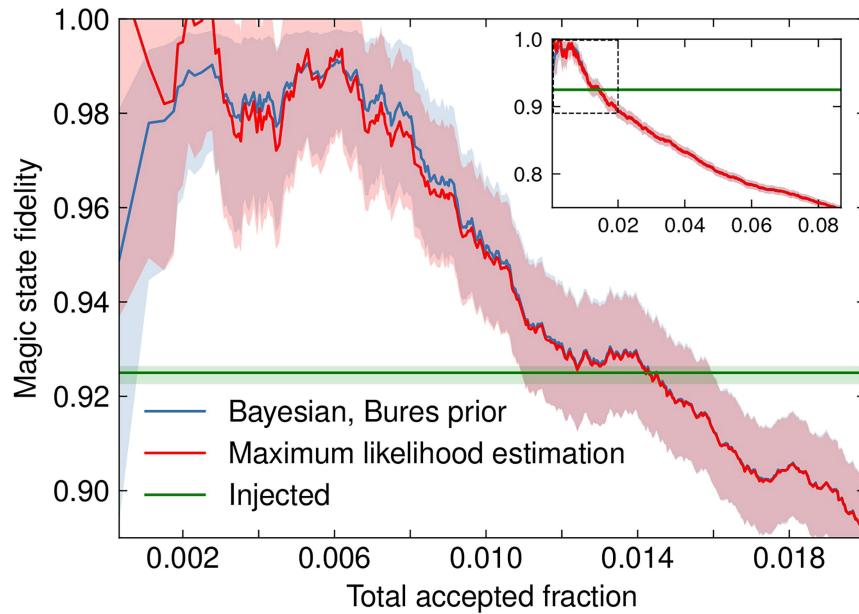
to local single-qubit gate inhomogeneity and two-qubit gate inhomogeneity. **b**, Time dependence of distance-3 color code stabilizers, for the experimental run with no added coherent error. Time traces are averaged with window size of 100. **c,d**, Same as **a** and **b** for distance-5.



Extended Data Fig. 4 | Additional decoding results. **a**, Simulated injected and distilled magic state fidelities as a function of global rescaling of physical error rates, when no stabilizer postselection is applied. Relative to our error model for decoding, the physical error rates have been further increased by $1.25 \times$ to match the experimental injected and distilled fidelities. **b**, Simulation and experimental data in table format for $d=3$, sorted into bins corresponding to $3 \times 5 = 15$ stabilizers and 5 logical observables, for a total of 2^{20} bins. We see



good agreement between simulation and experiment. **c**, Sliding-scale postselection of experimental distillation fidelity with added input errors. Fidelity of the output magic state (blue line) as a function of the total accepted fraction. The accepted fraction range decreases with added errors due to the factory acceptance rate decreasing. Horizontal line segments indicate the error corrected fidelity of the factory input states (green). Shaded regions indicate 68% confidence intervals.



Extended Data Fig. 5 | Comparison of different fidelity estimation methods for $d=5$ distillation. For low accepted fractions, the small number of samples causes the maximum likelihood estimate (red, allowed to exceed 1 here) and Bayesian estimates (blue) to differ noticeably, since the latter will be influenced

by the prior. In our figures in the main text, we therefore focus on the region in which the two methods give consistent estimates. Horizontal line indicates the error corrected fidelity of the factory input states (green). Shaded regions indicate 68% confidence intervals.