


# High Rate Magic State Cultivation on the Surface Code

Yotam Vaknin<sup>1,2,\*</sup>, Shoham Jacoby<sup>1,2</sup>, Arne Grimsmo<sup>1</sup>, and Alex Retzker<sup>1,2</sup>

<sup>1</sup>*AWS Center for Quantum Computing, Pasadena, California 91125, USA*

<sup>2</sup>*Racah Institute of Physics, The Hebrew University of Jerusalem, Jerusalem 91904, Givat Ram, Israel*

 (Received 10 February 2025; revised 6 October 2025; accepted 23 January 2026; published 16 March 2026)

Magic state cultivation is a leading approach for generating the resource states required for fault-tolerant quantum computation. Here we present a cultivation protocol that increases the success probability of magic state generation in platforms with flexible and nonlocal connectivity. Our method implements cultivation directly on the surface code, avoiding the detour through alternative, less efficient error-correcting codes used in prior approaches. This both improves the acceptance rate and preserves compatibility with the geometry of the code and the hardware. Numerical simulations show that our protocols improve success probabilities and reduce output error rates compared with protocols tailored to locally connected platforms. Under realistic noise models for cold-atom and trapped-ion systems, we improve the rate of magic state generation by more than a factor of 20. Finally, we study qubit loss and erasure and show that very low error rates can be achieved with minimal overhead, reaching below  $10^{-6}$  infidelity using only nine physical erasure qubits.

DOI: [10.1103/prxquantum-7-010353](https://doi.org/10.1103/prxquantum-7-010353)

## I. INTRODUCTION

Quantum error correction offers the potential to perform highly accurate quantum computations even when using noisy physical qubits [1,2]. Executing arbitrary quantum computations this way requires the generation of a state in a class of states called magic states. Generating a magic state requires unique machinery, different from standard error correction [3–7]. Early estimates predicted that roughly 95% of the physical qubits of a quantum computer would be designated for magic state generation [8], but many subsequent studies reduced these numbers considerably [9–15], up to the recent magic state cultivation (MSC) paper [16], which achieved magic state generation with resources comparable to a  $CX$  gate. Resource estimates based on MSC suggest that between 10% and 50% of the qubits are required for magic state generation [17,18], depending on the computer’s architecture and the specifics of the implementation.

The MSC result assumes the limitations that often occur for superconducting qubits: qubits are arranged in a square grid with local connectivity, and experience similar levels of noise while participating in gates and while idling. This

local connectivity motivates the use of an error-correcting code called the color code [19–21], which has a simple implementation of the  $H$  gate using only single-qubit gates, but generally performs worse than the surface code [8,22]. Because one typically wants to use the surface code for a significant part of the computation, Ref. [16] expands the color code into a surface code in a complex procedure that requires many additional qubits and substantial postselection.

Our work investigates how platforms with effective all-to-all connectivity, such as cold-atom arrays and trapped ions, can improve magic state cultivation. We also note that superconducting architectures could benefit from adding a limited number of long-range connections [23–27]. Recent work [28] showed that such systems can reduce the cost of magic state cultivation, but at the price of using codes with nonplanar connectivity, which are more difficult to implement in practice, mainly when qubit movement is required. In contrast, we use the surface code together with three-qubit gates in a way that respects the geometry of both the code and the physical qubits. Three-qubit gates have been demonstrated in both cold-atom arrays [29,30] and trapped-ion systems [31]. However, even without direct physical implementation, the added overhead of compiling three-qubit gates into sequences of single- and two-qubit gates is small, because they appear sparsely throughout the protocol. We further observe that, when we adopt a realistic noise model for such devices, our scheme yields an additional order-of-magnitude reduction in the required resources.

\*Contact author: [yotam.vaknin@mail.huji.ac.il](mailto:yotam.vaknin@mail.huji.ac.il)

*Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.*

We therefore summarize below the specific contributions that enable this improvement.

We make three main contributions. First, we construct and simulate two magic state cultivation protocols for the surface code based on transversal implementations of the gates  $H$  and  $H_{XY}$  using three-qubit gates. These schemes realize MSC-style magic state generation directly on the surface code, avoiding the color-code detour of Ref. [16] and yielding substantial savings in the magic state generation rate on architectures with effective all-to-all connectivity. Second, we show that, when combined with a realistic noise model for cold-atom and trapped-ion devices, the  $H$ - and  $H_{XY}$ -based protocols achieve an additional order-of-magnitude reduction in magic state resources compared with superconducting-style assumptions. Third, we revisit the logical  $CX$ -eigenstate protocol of Dennis *et al.* [9,15], recast it as a  $CX$  cultivation scheme on the surface code. Our simulations show that although this protocol is not competitive with our  $H$  and  $H_{XY}$  constructions in our setting, it is code-agnostic and may be attractive for error-correcting codes where transversal  $CX$  is particularly natural.

Last, we consider the effect of erasure and loss on magic state cultivation. Similar to Ref. [32], we show that the final logical fidelity is limited only by the nonidentifiable noise, i.e., the residual Pauli rate. We describe a minimal protocol that can utilize significant erasure bias, and achieve below  $10^{-6}$  infidelity using only nine physical erasure qubits.

Our paper is organized as follows: We begin by outlining the theory of magic state cultivation in detail. Next, we define our concrete protocol based on three-qubit gate implementations. We then present numerical results under two noise models tailored to different quantum computing architectures. Finally, we show numerical simulations of erasure-based implementations of our scheme.

## II. MAGIC STATE CULTIVATION

Magic state cultivation builds on the observation that measurements can achieve extremely high fidelity by being repeated and by coupling the data qubit to multiple ancilla qubits [33]. Magic state preparation is inherently noisy: each attempt typically produces a state with a component outside the target magic state. The basic idea is to generate such a noisy magic state and repeatedly measure an observable  $A$  for which the ideal magic state is a  $+1$  eigenstate. Conditioning on the  $+1$  outcome progressively increases our confidence that the state is correct. Ultimately, the fidelity will be constrained by the probability of flipping the state during measurement, which scales as the gate fidelity to the power of the code distance.

Let  $A$  be a unitary and Hermitian single-qubit operator with eigenvalues  $\pm 1$ . The  $+1$  eigenstate is the target magic state  $|A\rangle$ , and  $|\bar{A}\rangle$  denotes the  $-1$  eigenstate.




FIG. 1. Schematics of magic state cultivation—(Top) Partial Markov chain description of the cultivation protocol. Arrows are labeled by the leading-order scaling of their transition probabilities. We write  $p_{\bar{A} \rightarrow A} \propto p^l$ , where  $l = 1$  for phase-kickback and  $l = 2$  for *double-checking*. The final fidelity is limited by  $p^d$ , the same scaling as the error detecting distance of the underlying code. (Bottom) Schematic circuit for a single phase-kickback projection using a GHZ ancilla and a transversal implementation of  $A$  on a quantum error-correcting code. The  $X$ -basis measurement of the top ancilla qubit measures the eigenvalue of  $A$ , while the  $Z$ -basis measurements check the stabilizers of the GHZ state. We only keep runs in which all of these measurements, together with a subsequent round of surface code stabilizer measurements, return the trivial outcomes.

The process begins with a superposition or mixture of states  $|A\rangle$  and  $|\bar{A}\rangle$  and aims to increase the population of  $|A\rangle$ . This is achieved by repeatedly measuring the observable  $A$ . The process remains effective even when the measurement error rate is high, provided that measurement faults rarely change the underlying state without also producing an observable syndrome. Cultivation achieves this by supplementing the measurement with ancillary checks and a stabilizer measurement of the surface code. We abort and restart the protocol whenever any one of these checks reports a syndrome.

Even if we observe the correct measurement result with no syndromes, a noisy measurement of  $A$  can still fail in two ways: (1) with probability  $p_{A \rightarrow \bar{A}}$ , we can flip a correct  $|A\rangle$  while measuring it or (2) with probability  $p_{\bar{A} \rightarrow \bar{A}}$  a wrong state  $|\bar{A}\rangle$  is missed due to measurement errors. This idea is illustrated in Fig. 1.

The core of cultivation is that if there is a separation of scale between the two probabilities,

$$p_{\bar{A} \rightarrow \bar{A}} \gg p_{A \rightarrow \bar{A}},$$

repeated measurements improve the accuracy, up to a floor of the flipping probability  $p_{A \rightarrow \bar{A}}$ .

The more precise result depends on the input probabilities and obeys

$$p_{\bar{A}}^{\text{out}} \leq p_A^{\text{in}} \frac{p_{A \rightarrow \bar{A}}}{1 - p_{\bar{A} \rightarrow \bar{A}}} + p_{\bar{A}}^{\text{in}} (p_{\bar{A} \rightarrow \bar{A}})^k + \mathcal{O}(p_{A \rightarrow \bar{A}}^2).$$

Here  $p_{\bar{A}}^{\text{out}}$  is the probability of incorrectly accepting a  $|\bar{A}\rangle$  state,  $p_{A(\bar{A})}^{\text{in}}$  is the probability of correctly starting the procedure in  $|A(\bar{A})\rangle$  (i.e., successful injection), and  $k$  is the number of repetitions.

We define the probability of accepting a  $|A\rangle$  state as  $p_A^{\text{out}}$ . The final state infidelity is defined as

$$\text{IF} = \frac{p_A^{\text{out}}}{p_A^{\text{out}} + p_{\bar{A}}^{\text{out}}},$$

and the acceptance rate as

$$\text{Rate} = (p_A^{\text{out}} + p_{\bar{A}}^{\text{out}}).$$

In our plots, we show the expected number of attempts per accepted shot, which is the inverse of the rate. In the next section, we show that for code with distance  $d$ , cultivation can achieve  $\text{IF} = \mathcal{O}(p^d)$  and  $\text{Rate} = 1 - \mathcal{O}(p)$ .

### A. Implementation on an error-correcting code

Consider a logical state  $|\psi_L\rangle$  encoded in an error correcting code, and suppose the logical operator  $A$  admits a decomposition into single- and two-qubit gates  $\{A_i\}_i$ , so that

$$A = \prod_i A_i.$$

Note that  $A$  acts as an automorphism of the stabilizer group. The existence of such a decomposition means that we can implement  $A$  logically by acting only on individual physical qubits or pairs of qubits. We refer to such an  $A$  as a transversal logical operator.

Initializing an error-correcting code in the magic state  $|A\rangle$  requires a nontransversal unitary, which by the Eastin-Knill theorem [3] cannot be made fully fault tolerant. We therefore use a circuit that is not fault-tolerant, called a state-injection circuit [4,5,32]. State injection can produce an incorrect state with a probability that is linear in the physical error rate  $p$ .

To measure the  $A$  operator, we couple the encoded state to a GHZ ancilla state through controlled- $A_i$  gates and then

measure the GHZ state in the  $X$  basis, as illustrated in Fig. 1. We call this circuit a phase-kickback measurement circuit. For phase-kickback, the logical error probabilities scale as

$$p_{A \rightarrow \bar{A}} \propto p^d, \quad p_{\bar{A} \rightarrow \bar{A}} \propto p, \quad (1)$$

where  $d$  is the distance of the underlying code. Any pattern of faults that does not produce a syndrome and maps the logical state from  $|A\rangle$  to  $|\bar{A}\rangle$  must involve at least  $d$  physical errors, while incorrectly accepting an incoming  $|\bar{A}\rangle$  state can be caused by a single fault in the measurement protocol. Note that Eq. (1) omits the probabilities of identifiable error events, which are discarded rather than accepted. Consequently, for input  $|\bar{A}\rangle$  the most likely outcome is postselection, while for input  $|A\rangle$  it is acceptance; see Fig. 1. In practice, our best implementation uses a different circuit called *double-checking*, defined in Sec. II C. It preserves the favorable scaling  $p_{A \rightarrow \bar{A}} \propto p^d$  while improving the acceptance of incorrect states to

$$p_{\bar{A} \rightarrow \bar{A}} \propto p^2. \quad (2)$$

We establish these scalings by performing a complete search in error space, or when this becomes too computationally expensive, we run a full state-vector simulations and observe the same asymptotic behavior, see Appendix D.

These scalings imply that injecting a magic state into a distance-3 code achieves  $p^d$  logical error with a single *double-checking* round, and a distance-5 code achieves the same  $p^d$  scaling with only two iterations of *double-checking*.

### B. Transversal implementation on the surface code

We propose three cultivation protocols, each based on a distinct transversal operation (Fig. 2) and all relying on controlled-two-qubit gates (three-qubit gates), as shown in Figs. 3(b)–3(d). Two of these protocols use transversal implementations of the Hadamard gate  $H = (X + Z)/\sqrt{2}$  and the  $H_{XY} = (X + Y)/\sqrt{2}$  gate to produce the resource states  $|T_H\rangle$  and  $|T\rangle$ , respectively (defined below).  $|T_H\rangle$  and  $|T\rangle$  allow the realization of the non-Clifford gates  $Y^{1/4}$  and  $Z^{1/4}$ , which can be interleaved to synthesize arbitrary single-qubit rotations. The two states differ only by a Clifford transformation, so either one is sufficient for universal quantum computation, but having protocols for both based on different physical gates offers additional flexibility in practical implementations.

The third protocol is based on the  $CX$  gate and prepares the resource state  $|CX\rangle$  (defined below). This state is attractive because only two copies suffice to implement a Toffoli gate with probability  $4/9$  via gate teleportation [9,15], compared with seven copies of  $|T\rangle$  in standard constructions. Since the Toffoli gate is often the dominant




FIG. 2. Transversal implementation of  $H$ ,  $H_{XY}$ , and  $CX$  on the surface code—On the unrotated surface code,  $H$  is implemented via a pattern of SWAP and single-qubit  $H$  operations, while  $H_{XY}$  is implemented using  $H_{XY}$ ,  $H_{XYZ}$ , and  $CZ$  gates. Logical  $CX$  is transversal for both rotated and unrotated surface codes by applying physical  $CX$  gates between corresponding data qubits of two code patches. The figure depicts the construction on rotated surface-code patches, but the transversal pairing applies to either geometry.

non-Clifford operation in fault-tolerant implementations of arithmetic, an efficient source of  $|CX\rangle$  states can significantly reduce the overall magic state cost of many important algorithms. In practice, we find that the  $CX$  cultivation protocol is not efficient enough to outperform the two constructions described above. Nevertheless, we include it here because its definition does not rely on any particular features of the surface code, and thus it may be of independent interest as a code-agnostic protocol applicable to other CSS codes that encodes a single qubits.

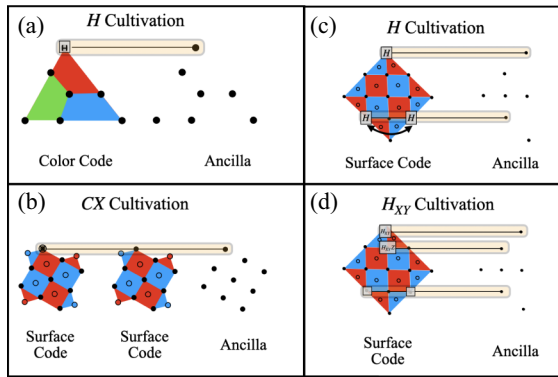


FIG. 3.  $H$ ,  $H_{XY}$ , and  $CX$  cultivation—Schematic diagrams of  $H$ ,  $H_{XY}$ , and  $CX$  cultivation. Black dots label the data qubits, and orange rectangles label representatives of controlled transversal gates (shown in full in Fig. 2). (a)  $H$  cultivation on 2d triangular color code. A GHZ ancilla measures the  $H$  operator by transversal application of the controlled-Hadamard gate. (b)  $CX$  cultivation on a pair of  $d = 3$  rotated surface codes. Blue and red dots label ancillas. The  $CX$  operator is measured on a pair of surface codes by transversal application of the  $CCX$  (Toffoli) gate with using an ancilla as control. (c)  $H$  cultivation on the unrotated surface code, implemented using a controlled transversal  $H$  on all the qubits, followed by controlled SWAP along the diagonal. (d)  $H_{XY}$  cultivation on the unrotated surface code, applied using alternating controlled  $H_{XY}$  and  $H_{XYZ}$  gates along the diagonal and controlled  $CZ$  gates across the diagonal.

We will now describe the transversal implementations on the unrotated surface code, as shown in Fig. 2. The unrotated surface code has fold-duality as defined in Ref. [28]; By swapping pairs of qubits across the code’s diagonal, the resulting code is a dual of the original code, in the sense that  $X$  stabilizers of the original code become  $Z$  stabilizers of the dual code and vice versa. The same goes for the  $X/Z$  logical operators.

(Fold) Transversal Hadamard can be implemented by following the SWAP operation with  $H$  gate on every data qubit, which swaps the  $X$  and  $Z$  operators back. An unrotated surface code has  $2d - 1$  qubits on its diagonal and  $n$  data qubits in total. We label the  $2d - 1$  diagonal qubits as  $a_i$ , and the  $(n - 2d + 1)/2$  pairs of off diagonal data qubits as  $(b_i, c_i)$  and write the transversal decomposition as

$$A_i = \begin{cases} H_{a_i} & i \leq (2d - 1), \\ \text{SWAP}(b_i, c_i) H_{b_i} H_{c_i} & i > (2d - 1), \end{cases}$$

where  $\text{SWAP}(b_i, c_i)$  swaps between the qubits  $b_i, c_i$ , and  $H_x$  is the  $H$  operator on qubit labeled by  $x$ . The resulting  $+1$  eigenstate of the logical  $H$  is the following magic state:

$$|T_H\rangle = \cos(\pi/8) |0\rangle + \sin(\pi/8) |1\rangle.$$

Similarly, the following transversal operator behaves as  $H_{XY}$  in the logical basis of the surface code,

$$A_i = \begin{cases} (H_{XY})_{a_i} & i \leq (2d - 1) \wedge i \bmod 2 = 1, \\ (H_{XYZ})_{a_i} & i \leq (2d - 1) \wedge i \bmod 2 = 0, \\ CZ(b_i, c_i) & i > (2d - 1), \end{cases}$$

where the subscript again labels the single-qubit gate’s support, and  $CZ(b_i, c_i)$  is the controlled- $Z$  gate between qubits  $b_i, c_i$ . The details of this transformation are described in Appendix A. The  $+1$  eigenstate of  $H_{XY}$  is

$$|T\rangle = [ |0\rangle + e^{i\pi/4} |1\rangle ] / \sqrt{2}.$$

The same ideas can be used on the rotated surface code, although it is not fold-dual. The rotated surface code requires fewer physical qubits than the unrotated layout to reach similar logical error rates, so it is generally the more resource-efficient choice. As described in Ref. [34], halfway through the stabilizer extraction (SE) cycle of the rotated surface code, the stabilizer group closely matches that of an unrotated surface code with the same distance. This mid-cycle equivalence allows us to implement both fold-transversal operators defined above for the *unrotated* surface code directly on a *rotated* surface code by applying them at this intermediate point in the stabilizer circuit, thereby combining the transversal structure of the unrotated geometry with the qubit savings of the rotated

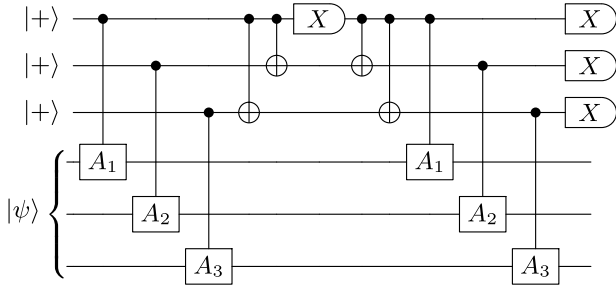


FIG. 4. Schematic description of double-checking—The top three qubits represent the ancilla qubits, while the bottom three qubits represent the data qubits of the surface code (or codes). The circuit projects  $|\psi\rangle$  to the  $+1$  eigenvector of  $A$  when all measurement outcomes are  $+1$ . The state is discarded if any measurement outcome is  $-1$  using *double-checking*.

one. After the projection, the measurement cycle is completed and we measure the stabilizers and treat them as an additional ancillary checks. Figure 9 shows the *double-checking* circuit implemented this way on the unrotated surface code.

Last,  $CX$  cultivation is based on the observation that logical  $CX$  is transversal for any CSS code: applying physical  $CX$  gates between corresponding data qubits of two code blocks implements a logical  $CX$  between the encoded qubits. Explicitly, if we label each pair of corresponding physical qubits from the two surface-code patches by  $(a_i, b_i)$  for  $1 \leq i \leq n$ , and we can define the transversal operator  $A_i = CX(a_i, b_i)$ . The  $+1$  eigenspace of logical  $CX$  contains multiple logical states, but by initializing the two surface codes in the logical state  $|+, 0\rangle$  and projecting onto this eigenspace, the joint logical state is prepared in

$$|CX\rangle = |0, +\rangle + |0, -\rangle + |1, +\rangle,$$

which we use as the entangled resource state for Toffoli implementation [9,15].

### C. Phase Kickback and double checking

For completeness, we will define in detail two different methods for state projection obeying Eqs. (1) and (2): phase kickback and *double-checking*. For phase kickback we initialize a GHZ ancilla by expanding a  $|+\rangle$  state from one qubit using  $CX$  gates. Then, a controlled transversal gate is applied with the GHZ state as control, as seen in Fig. 1.

If the logical qubit is in some superposition of the two eigenstates  $|\psi\rangle = \alpha |A\rangle + \beta |\bar{A}\rangle$ , the application of the transversal operation conditioned on the GHZ ancilla

results in the following state:

$$\left( \prod_i CA_i \right) \left[ \frac{|0\rangle^{\otimes n} + |1\rangle^{\otimes n}}{\sqrt{2}} \right] |\psi\rangle = \alpha [|0\rangle^{\otimes n} + |1\rangle^{\otimes n}] |A\rangle + \beta [|0\rangle^{\otimes n} - |1\rangle^{\otimes n}] |\bar{A}\rangle,$$

where  $CA_i$  is a transversal  $A_i$  gate conditioned on the  $i$ th ancilla in the GHZ state. Inverting the GHZ generation and then measuring its qubits effectively measures the GHZ state and can identify a single bit flip error [12].

Note that if the transversal operation did not act trivially on the stabilizers of the code, it would generate additional entanglement between the code state and the ancilla. In such a case, the superposition of the ancilla would be destroyed, leaving it in a mixed state over  $|0\rangle^{\otimes n}$  and  $|1\rangle^{\otimes n}$ . A subsequent measurement would then fail to reveal any information about the  $|A\rangle$  state. To avoid this, our initialization sets all stabilizers to the  $+1$  eigenspace, making  $A$ 's automorphism of the stabilizer group trivial.

In the *double-checking* protocol, first introduced in Ref. [16], all the ancilla qubits are initialized in the  $|+\rangle$  state. The controlled transversal operation is applied, which puts the ancilla in the following superposition:

$$\left( \prod_i CA_i \right) |+\rangle^{\otimes n} \otimes |\psi\rangle = \alpha \prod_i (|0\rangle + |1\rangle A_i) |A\rangle + \beta \prod_i (|0\rangle + |1\rangle A_i) |\bar{A}\rangle. \quad (3)$$

Notice that  $\prod_i (|0\rangle + |1\rangle A_i) |A\rangle$  is a  $+1$  eigenstate of the  $X^{\otimes n}$  operator on the ancilla, therefore eigenvalue of  $A$  is encoded in the parity of  $X^{\otimes n}$ . We measure this parity using a series of  $CX$  gates that map by conjugation  $X^{\otimes n}$  to  $X$  operator of one of the ancillas. Measuring this ancilla in the  $X$  basis measures the value of  $A$ . The number of checks is doubled by running the circuit in reverse and measuring all of the ancilla qubits in the  $X$  basis, see Fig. 4.

### D. Code expansion

An important optimization of the cultivation protocol is to gradually increase the code size [11,16]. The protocol described so far aborts following most physical errors, which makes the acceptance rate highly sensitive to the number of active qubits. By activating some of the qubits only in the later stages of the protocol, we can substantially increase the acceptance rate.

Concretely, we first run cultivation on a small  $d = 3$  code, and then use a short circuit to transform the code into a larger  $d = 5$  surface code. As long as the expansion circuit, together with the final projection step, continues to obey Eqs. (1) and (2), the output fidelity retains the same asymptotic scaling. We note that this argument does not work for  $CX$  cultivation, as explained in Appendix A 1.

We will follow the convention of Refs. [16,28], and label each protocol by  $d_1$ , the code distance used at the final

projection step and therefore the distance that determines the asymptotic logical scaling.

Because cultivation relies on postselection, it can be implemented at relatively small code distances, whereas the subsequent fault-tolerant computation must run without postselection and therefore requires a higher distance. Following Ref. [16], we bridge these two regimes by adding a final code-expansion stage that moves the cultivated state onto a larger code suitable for standard error correction, with only limited additional postselection. We label that final distance  $d_2$ . After the last cultivation measurement, each surface-code patch is expanded [5,10,16] to a distance- $d_2 = 11$  code and subjected to ten further rounds of stabilizer measurements [35]. The expansion is performed by initializing each newly added qubit in the eigenbasis of any weight-2 stabilizer or logical operator it extends, so that all stabilizer and logical eigenvalues remain deterministic during the growth process [5]. In our simulations we expanded (un)rotated surface codes to larger (un)rotated surface codes. However, since rotated and unrotated surface codes can both be embedded into sufficiently large codes of either type, the same procedure can in principle be used to expand an unrotated code into a rotated one, which would be more qubit efficient.

At this stage we no longer postselect on every detected syndrome; instead, we follow the soft-decoding strategy of Refs. [16,36,37]. The decoder finds the most probable error resulting in either a logical phase error or no error, and their respective probabilities  $p_Z, p_I$ . These probabilities define the complementary gap  $\text{gap} = \log(p_Z/p_I)$ . By imposing a threshold on this gap, we obtain a family of protocols that trade acceptance rate against output fidelity, and in our results we sweep this threshold to map out the full acceptance–fidelity frontier.

### E. Protocol

We now describe the step-by-step structure of our cultivation protocol, see Fig. 5. The  $H$  and  $H_{XY}$  schemes follow the same sequence of operations, differing only in the specific transversal operator  $A$  and in the details of the injection step. The protocol proceeds as follows:

#### (1) Initialization

A distance-3 surface code is initialized in the  $+1$  eigenstate of the logical  $Z$  operator and in the  $+1$  eigenspace of all stabilizers. This can be achieved by measuring all stabilizers once and applying physical corrections conditioned on the outcomes.

#### (2) State injection with syndrome extraction

A nonfault-tolerant unitary injection circuit is applied to map the encoded state to the target magic state  $|A\rangle$  (see Appendix J for details). We then perform a single round of the standard syndrome-extraction circuit. The run is discarded if a syndrome is observed.

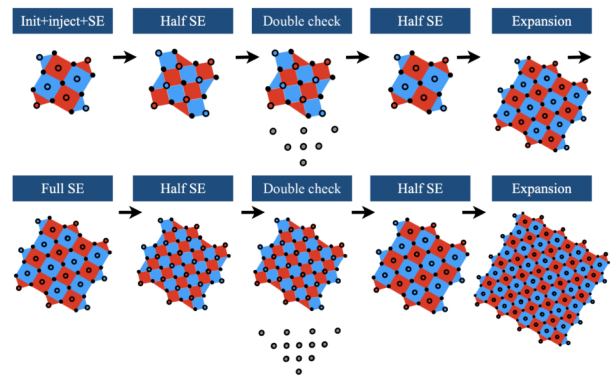


FIG. 5. Full protocol for  $d_1 = 5$  rotated surface code—Step-by-step illustration of the cultivation protocol. Filled black dots denote data qubits. Hollow red (blue) dots denote ancilla qubits used for syndrome extraction of  $X$  ( $Z$ ) stabilizers. Hollow gray dots denote ancilla qubits used for double-checking.  $SE$  refers to the syndrome-extraction circuit, and  $half\ SE$  refers to two of the four  $CX$  gates required for a surface-code  $SE$  round. The red and blue polygons indicate the state of the stabilizer group after each labeled step [34].

#### (3) Double-checking

The logical operator  $A$  is measured using the double-checking circuit described in Sec. II C. The run is aborted if any measurement yields a nontrivial outcome.

#### (4) Second double-checking with code expansion

Only for protocols we label as  $d_1 = 5$ , we expand the code before performing an additional projection round. We initialize the additional qubits required for a  $d = 5$  surface code, half in the  $X$  basis and half in the  $Z$  basis, and apply a short expansion circuit that maps the  $d = 3$  code into a  $d = 5$  code while preserving the  $+1$  eigenvalues of the stabilizers. We run a single syndrome extraction round, and then apply a second *double-checking* circuit.

#### (5) Code expansion and soft-decoded error correction

The accepted state is expanded to a larger surface code. In our simulations we expand to distance  $d_2 = 11$  and perform 10 rounds of syndrome extraction. At this stage we do not postselect on every detected syndrome. Instead, we apply soft decoding and accept the run based on a threshold on the decoder confidence, quantified by the complementary gap.

We present this in full as pseudo-code in Appendix L. For  $CX$  cultivation, we use a slightly different protocol, see Appendix A 1.

## III. SIMULATION

We next turn to numerical simulations. For small code distances, we verified the protocol’s correctness using full state-vector simulations. Specifically, we simulated both  $H$

and  $H_{XY}$  with rotated and unrotated codes with  $d_1 = 3$ , and for  $CX$  cultivation we simulated a representative example with  $d_1 = 2$ .

To estimate the logical performance of our protocols, we introduce noise into the simulations controlled by a single parameter  $p$ . After every single-qubit (two-qubit) gate, we apply a single-qubit (two-qubit) depolarizing channel with probability  $p$ . After each three-qubit operation, we apply a three-qubit depolarizing channel with probability  $3p$ , together with additional two-qubit depolarizing channels between the ancilla control and each data-qubit target, each occurring with probability  $p$ . The full noise model is specified explicitly in Appendix B.

While the computational cost of full state-vector simulation allowed us to simulate the protocol up to distance  $d_1 = 3$  without expansion, it became prohibitive to directly estimate the logical error rate after the final expansion to distance  $d_2$ , or even after the shorter expansion to  $d_1 = 5$ . To access this larger-scale regime, we therefore follow Ref. [16] and employ a closely related circuit that replaces the non-Clifford gates with Clifford gates and projects a different logical operator. This enables efficient sampling via Clifford simulation [16,38–40]. Concretely, in the Clifford simulation [39,40] we replace the non-Clifford  $CA_i$  gates with combinations of  $CX$  designed to measure the logical  $X$  operators. We introduce noise as depolarizing channels as defined in Appendix B. To ensure that our simulation approximates the noise in the system, we apply a uniform depolarizing noise on the support of the

non-Clifford gate, which sometimes includes three-qubit depolarizing errors, see Appendix I. The results of these simulations are shown in Fig. 6.

Similarly to Ref. [16], we observe good agreement in both absolute fidelity and scaling between the full vector simulations and the simplified Clifford simulations, with only a small factor difference before the expansion step. See Appendix D for details. Owing to its smaller qubit footprint, the rotated surface code provides the best overall fidelity, even though its circuit is longer. When idling errors are absent, the shorter circuit of the unrotated surface code performs better, as demonstrated in the next section.

Finally, we compare our protocol with color-code-based magic state distillation [12,16]. Although color-code-based protocols have an initial advantage due to a smaller qubit footprint, our protocol is ultimately superior. The expansion step in our protocol is substantially more efficient than the costly grafting procedure required for color codes. Grafting requires a long idling step with high rates of post selection, necessary to convert from color code to surface code. Our protocol skips this idling step entirely, since we simply expand to a larger surface code. This results in an order-of-magnitude higher distillation rate for representative parameters of  $d_1 = 5, d_2 = 11$ , which are easily comparable with previous work [16]. Our protocol is able to reach  $10^{-9}$  infidelity with the same physical error rate  $p = 10^{-3}$  and  $1/5$  the number of shots.

### A. Efficient cultivation with long idling time

To enable direct comparison with previous schemes, we adopt the uniform-noise model described in Ref. [16]. This model treats idling as equally harmful as a two-qubit gate, which can be a reasonable approximation for certain superconducting transmon devices [41,42]. However, some quantum devices offer significantly better idling fidelities compared to their 1Q and 2Q gates. For example, Cold-atom devices have exceptionally long  $T_1$  times while not populating the Rydberg state, and their dephasing time  $T_\phi$  can be prolonged with dynamical decoupling, or by shelving in clock states [43,44]. This bias between gates and idling can also be achieved in ions and superconducting cavities [45–47] and possible with spins in solids, like Ref. [48].

We modeled this effect by defining a different noise model, which we call atom-noise model. In this model, we reduce the one-qubit gates error probability to  $p/10$ , and remove any idling noise (Appendix C). The results of simulations with atom-noise model are shown in Fig. 7.

This offers a significant advantage in practice. In a recent estimate for the resources required to break RSA 2048 using superconducting qubits, roughly 20% of the qubits were designated for magic state generation [17]. Using color-code cultivation and targeting  $10^{-7}$  infidelity, the acceptance rate is approximately 5%. With  $d_1 = 3 H_{XY}$

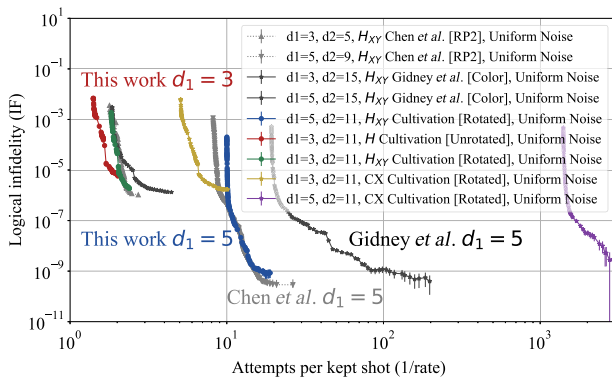


FIG. 6. Logical error rate of cultivation with  $d_2$  expansion—Uniform noise model—Logical infidelity, and acceptance rate of our protocol compared with previous protocols [16,28]. The physical noise parameter is  $p = 10^{-3}$  using uniform noise model. The color code from Ref. [16] is expanded to a variant of the surface code, while the rest are expanded to a standard surface codes, each with their labeled distance  $d_2$ . The different points label various thresholds for the complementary gap post-selection. Error bars represent a single standard deviation. We included the  $d_1 = 5 CX$  case although it has large error bars, giving only an approximate bound on its error rate. Note that Ref. [28] uses different values of  $d_2$ , which we argue in Appendix K has a negligible effect on this figure.

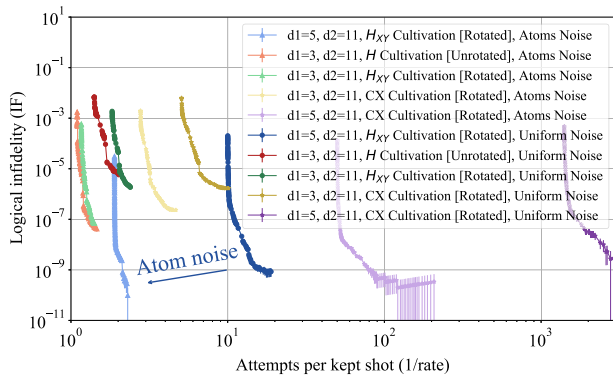


FIG. 7. Logical error rate of cultivation with  $d_2$  expansion—Atom-noise model—Logical infidelity and acceptance rate of surface code  $CX$ ,  $H_{XY}$  and  $H$  cultivation with Uniform and atom-noise model. In the atom-noise model, idling does not introduce any additional error, which approximate the effect of dynamical decoupling or shelving in clock-states. The physical error rate is  $p = 10^{-3}$ , and the full noise model is defined in the Appendix B. The different points label various thresholds for the complementary gap postselection. For practical application that utilize syntillation, achieving  $10^{-7}$  infidelity is possible with  $d_1 = 3$  and approximately 75% rate.

cultivation we have roughly 75% acceptance rate with fewer qubits and roughly half the number of rounds. In total, this results in over an order of magnitude reduction in qubit-rounds cost.

Still, in a resource estimate designed specifically for cold atoms [18], magic state cultivation occupies half the atoms during addition (where most of the Toffoli gates are consumed). Our short circuit and 75% acceptance rate together would reduce the qubit-rounds assumed in Ref. [18] by over a factor of  $\times 20$ . This would directly reduce the qubit count during addition by over 6 million qubit (or roughly 40%).

#### IV. ERASURE QUBITS AND LOSS

Erasure qubits [49–54] are a type of qubits with a specific noise mechanism that can be directly identified. In superconducting qubits, energy relaxation events can be identified by encoding a single qubit in the single excitation subspace of two transmons [50,52] or cavities [55,56], and measuring the zero excitation state. In Rydberg atoms [53,57,58], during a two-qubit gate the population in the Rydberg state decays to many states, most of them outside of the computational subspace. Observing the population in these states identifies this decay. Similarly, loss in Rydberg atoms can be detected through destructive measurements [44,59]. In the context of cultivation, any technique that identifies loss by the end of the protocol effectively functions as erasure detection [59].

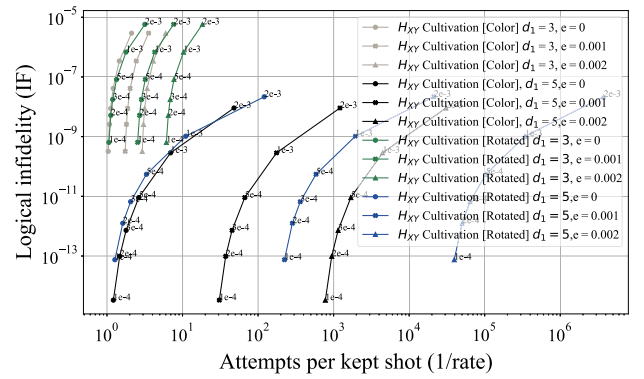


FIG. 8. Logical error rate and Acceptance rate of erasure qubits—Logical error rate and acceptance rate of erasure qubits used for the initialization step of cultivation. The points are labeled by the residual Pauli error rate  $p$ . As mentioned in the main text, the logical error rate does not depend on the erasure rate but only on the residual Pauli error rate. The logical infidelity is estimated by full search on the error graph up to weight 5, and matches sampling for values  $p$ 's where sampling is possible. Importantly, the color-code-expansion step is considerably more expensive than for the surface code, but is not shown in this plot. Still, from this graph we can extract the trade-off when converting between different erasure and Pauli rates. For example, if  $p = 10^{-3}$  can be converted to  $e = 2 \times 10^{-3}$  and  $p = 10^{-4}$ , it would reduce the acceptance rate of  $d_1 = 3$  cultivation by roughly a factor of 5, but improve the fidelity by over  $10^3$ .

The erasure rate does not effect the logical fidelity [32] as long as it can be detected. In this limit, only the acceptance rate is affected by erasure. Since most Pauli errors are identifiable to begin with [32], the effect of erasure  $e$  is approximately to reduce the acceptance rate to the rate of the same circuit with only Pauli noise with rate  $p = e$ . See Appendix F for our analysis and [32], for a detailed investigation of this trade off. We estimated the effect of erasure on cultivation by evaluating its effect on acceptance rate without the final  $d_2$  expansion step in Fig. 8. In Appendix F, we provide an explanation of the estimation process for assessing the impact of erasure on complementary gap decoding. We find that the overhead of adding  $e = 2 \times 10^{-3}$  erasure probability per operation for  $d_1 = 3$  surface code cultivation roughly reduces the acceptance rate by a factor of 5. If in this regime one can achieve a meaningful bias, resulting in  $p = 10^{-4}$ , it can reach infidelity of  $10^{-9}$  with the same rate as  $d_1 = 5, p = 10^{-3}$  cultivation but with significantly fewer qubits and a shorter protocol. Such a level of bias is realistic for certain trapped-atom species [53] and for superconducting qubits [52], although the latter would require some degree of long-range connectivity.

If the bias is small, which is the case for some trapped-atom species [44] where loss represent about 50% of gate fidelity, cultivation is still improved by loss detection. The rate would be the same as we described before, but the final

fidelity would improve by about  $2^{d_1}$ , which as seen in Fig. 9 is between 8 and 50 depending on the specific protocol.

If the bias can reach  $e = 10^{-3}$  and  $p = 10^{-4}$  with only local connectivity, such as for dual-rail superconducting qubits, it would be possible to achieve  $p_L \approx 10^{-15}$  once every 200 attempts with  $d_1 = 5$  color code cultivation of Ref. [16], if we include post selection on the complementary gap. At this scale, the qubit-cycles count of cultivation is  $1.5 \times 10^5$  [60], which is about 20 times less than the  $(15\text{-to-}1)_{13,5,5}^6 \times (15\text{-to-}1)_{29,11,13}^6$  scheme of [7] using nonerasure qubits.

Erasure offers an additional interesting opportunity for cultivation, which is to reach algorithmic interesting infidelities with only a  $d_1 = 2$  unrotated surface code and a single GHZ projection. This protocol requires only nine physical qubits. With this setup and a residual Pauli error rate of  $p = 10^{-4}$ , cultivation can achieve a state infidelity of  $10^{-6}$ , which can then be combined with syntillation [61] to reach algorithmically relevant fidelity. For more information, see Appendix G

## V. DISCUSSION

In this work, we demonstrated how cultivation can be implemented directly on the surface code using nonlocal connection. Although the initialization step of cultivation on the surface code has a slightly lower rate compared with color-code cultivation, this effect is compensated by the simple and high rate expansion step of the surface code to higher distances. This is clearly superior in architectures with all-to-all connectivity such as cold atoms and trapped ions. Even with architectures exhibiting limited long-range connectivity (e.g., superconducting qubits [62]), compiling nonlocal interactions will be quite efficient, given their small number.

The acceptance rate could be further improved by providing the decoder with additional information, as in Ref. [28]. In our analysis we adopt the decoder-latency assumption of Ref. [16]: we simulate 10 extra rounds of syndrome extraction after cultivation before declaring the state ready. For superconducting qubits this corresponds to roughly a 10  $\mu\text{s}$  decoding delay. If the decoder responds faster, fewer waiting rounds are needed, consistent with Ref. [28]. This 10-round convention significantly affects the qubits  $\times$  rounds metric. In architectures with longer measurement times, decoder latency can be effectively hidden, so the same waiting rounds would not appear as an added cost in qubits  $\times$  rounds. Although this still slows wall-clock cultivation, it can make cultivation appear cheaper under qubits  $\times$  rounds. We emphasize that, provided  $d_2 \geq 2d_1 + 1$ , both the choice of  $d_2$  and the number of postexpansion rounds have negligible impact on acceptance rate and output fidelity, but a large impact on qubits  $\times$  rounds (see Appendix K). A more direct comparison would incorporate transport and two-qubit gate times, and would choose

$d_2$  based on how the cultivated state is consumed in the target circuit. These effects are not well captured by qubits  $\times$  rounds, so we do not attempt that comparison here and leave it for future work.

Many quantum computing architectures have high idling fidelity. As we have shown, this gives a great benefit for magic state cultivation, both in fidelity and rate. For example, running 2048 bit RSA only requires roughly  $10^{-7}$   $T$  gate infidelity [17,18], which can be achieved with  $d_1 = 3$  cultivation on cold atoms with very minimal loss detection. Since prior work assumed cultivation was successful with such small probability, it either required idling until the magic states are ready [17], or significantly increasing the number of qubits [18]. Our high rates suggest magic states can be generated highly efficiently, which would have very large effect on previous resource estimates.

## ACKNOWLEDGMENTS

We thank Craig Gidney for suggesting the mid-cycle measurement of the  $H_{XY}$  operator, and for making Stim, Sinter and Crumble [40] open source, all of which were instrumental for this project. The authors were made aware of two related publications posted on arXiv concurrently with this work: one from Puri's group, which investigates fold-transversal cultivation on the surface code [63], and another from J. Claes, which demonstrates surface code cultivation using two-qubit gates [64]. We thank the staff from across the AWS Center for Quantum Computing that enabled this project. We also thank Simone Severini, James Hamilton, Nafea Bshara, and Peter DeSantis at AWS, for their involvement and support of the research activities at the AWS Center for Quantum Computing.

## DATA AVAILABILITY

The data that support the findings of this article are openly available [65].

## APPENDIX A: TRANSVERSAL $H_{XY}$ AND $H$

In this Appendix we explain in more detail why the transversal operator  $H_{XY}$  and  $H$ , as defined in the main text, acts as the logical operators on the unrotated surface code and induces an automorphism of its stabilizer group. Since the stabilizer group of the rotated surface code at the middle of the syndrome extraction circuit [34] is the same as the unrotated surface code, the exact same argument applies in the rotated case (Fig. 9).

Recall that, under conjugation by  $H_{XY}$  on a single qubit,

$$X \mapsto Y, \quad Y \mapsto X, \quad Z \mapsto -Z,$$

and that the two-qubit gate  $CZ$  acts by conjugation as

$$CZ(X \otimes I)CZ^\dagger = X \otimes Z, \quad CZ(I \otimes X)CZ^\dagger = Z \otimes X,$$

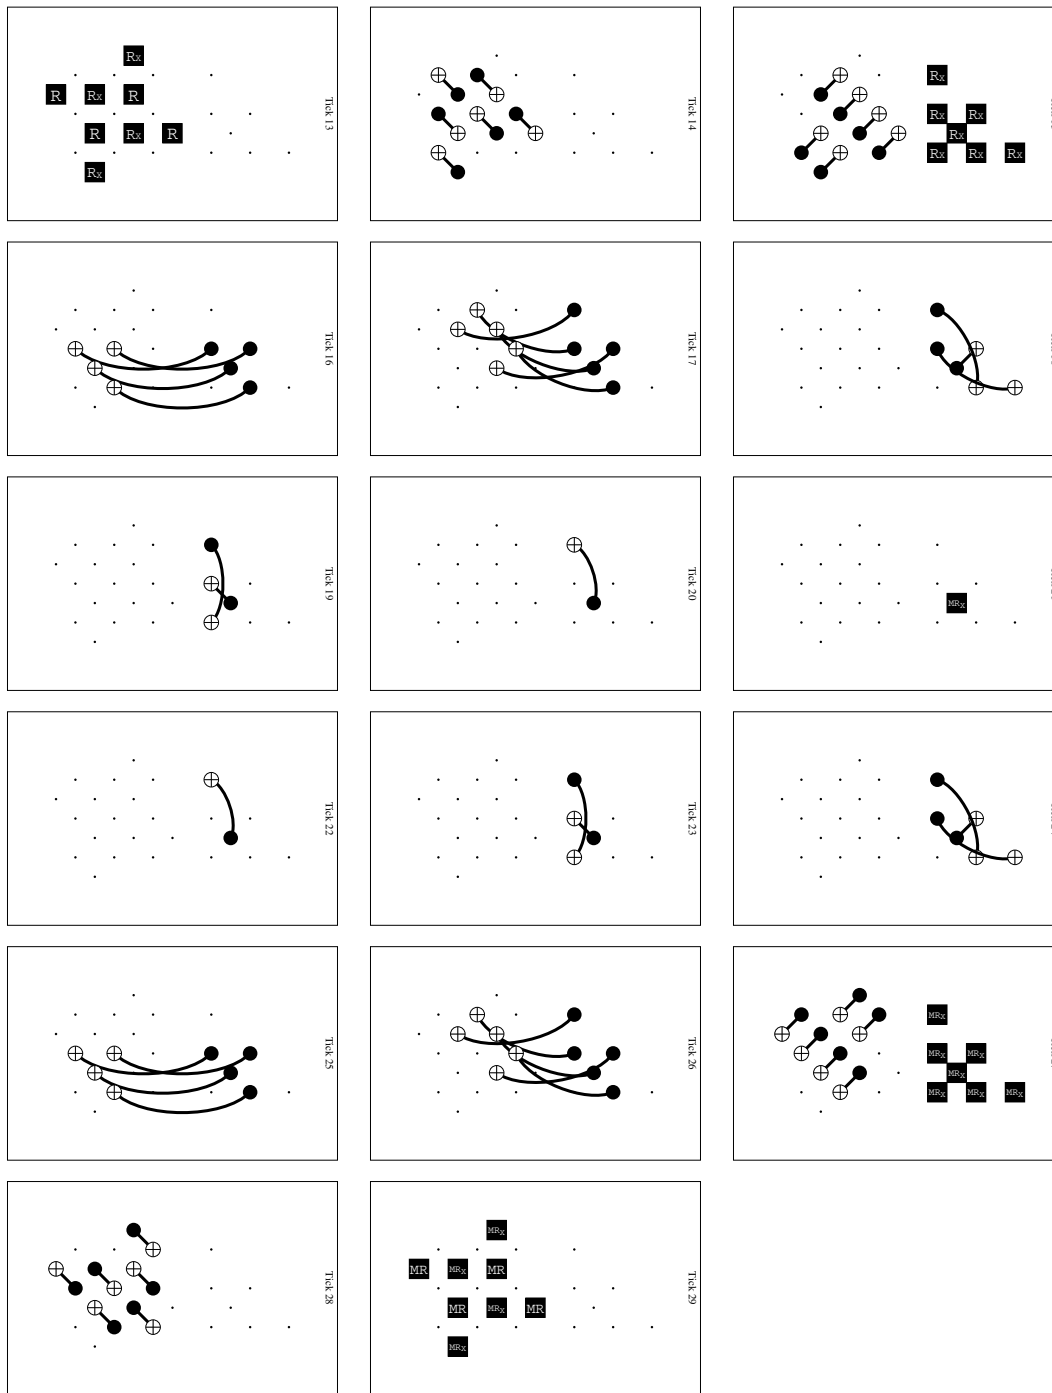


FIG. 9. Circuit measuring the  $X$  operator on a rotated surface code mid SE cycle—A *double-check* of the  $X$  operator on a rotated surface code with  $d_1 = 3$  mid SE cycle. From top to bottom, left to right, each panel represent a time slice of the circuit (labeled with tick index). For non-Clifford simulation, the  $CX$ 's in ticks (16,17,25,26) were replaced with  $CCZ$ ,  $CH_{XY}$  and  $CH_{XY}Z$ . [Click here to view in Crumble.](#)

while commuting with all single-qubit  $Z$  operators. We visualize these transformations in Fig. 10, which is the simplest way to understand the following argument.

First consider the action on the  $Z$ -type stabilizers. Along the diagonal, we apply alternating single-qubit gates  $H_{XY}$

and  $H_{XY}Z$ . Each diagonal  $Z$  operator therefore acquires a factor of  $-1$ . Since every  $Z$ -type stabilizer has even support on the diagonal, the overall phase from the diagonal cancels and each  $Z$  stabilizer is mapped back to itself. The off-diagonal  $CZ$  gates commute with all  $Z$  operators, so

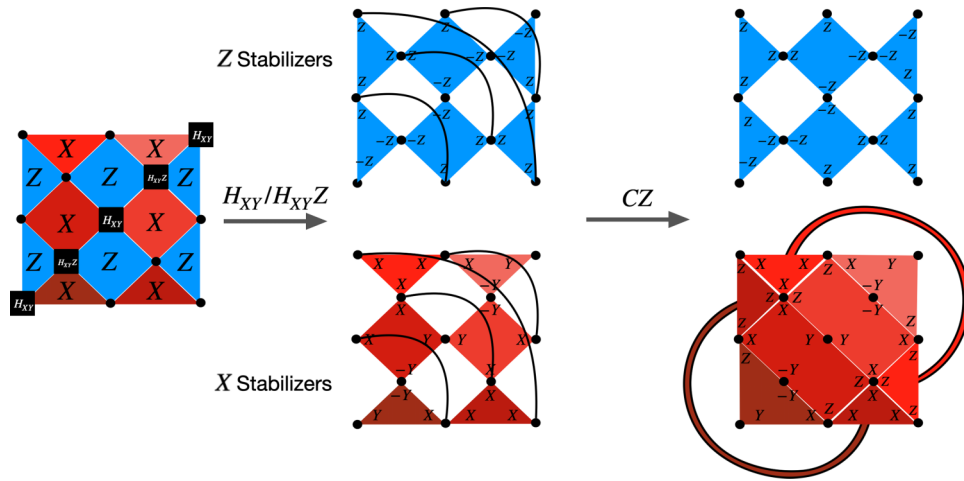


FIG. 10. Action of transversal  $H$  on the surface-code stabilizers. Each panel shows the state of the stabilizer group before the operation depicted in it, and the last panel shows the final form of the stabilizers.

they also preserve every  $Z$ -type stabilizer. In combination, the full transversal operator leaves the  $Z$ -stabilizer group invariant.

Next we use the fold duality of the unrotated surface code. For each  $X$ -type stabilizer  $S_X$  there is a corresponding  $Z$ -type stabilizer  $S_Z$  obtained by reflecting across the diagonal. Conjugation by the transversal  $H_{XY}$  maps each  $X$ -type generator to the product  $S_X S_Z$ , as illustrated in Fig. 10. Since the  $Z$ -stabilizer group is preserved, this mapping takes the  $X$ -stabilizers to elements of the stabilizer group and thus defines an automorphism of the full stabilizer group.

Finally, we consider the logical operators. The logical  $Z_L$  string intersects the diagonal in a single qubit, so it picks up a minus sign from the action of  $H_{XY}$  on that qubit, and is otherwise unchanged. Hence  $Z_L \mapsto -Z_L$ . The logical  $X_L$  string crosses  $Z_L$  at a single qubit, and on that qubit  $X$  is mapped to  $Y$  by  $H_{XY}$ . Combined with mapping  $X$  terms to an  $X \otimes Z$  product with their fold-dual qubit, this implies that the logical operator transforms as  $X_L \mapsto iX_L Z_L = Y_L$ , as required from a logical  $H_{XY}$ .

For the  $H$  operator, under conjugation, we have

$$X \mapsto Z, \quad Z \mapsto X.$$

To achieve this on the surface code, we use the  $H$  operator on all the physical qubits, and then reflect the code across its diagonal using SWAP gates. As can be seen in Fig. 11, the effect of the  $H$  gates is to exchange  $X$  and  $Z$  stabilizers, while the SWAP recovers the original structure of the surface code. Note that for every pair of dual stabilizer  $S_X$  and  $S_Z$ , the result of the transformation is  $S_X \mapsto S_Z$  and  $S_Z \mapsto S_X$ , which is an automorphism on the stabilizer group. The  $X$  and  $Z$  logical operators map to each other.

### 1. Logical scaling of $CX$ cultivation

The logical error rate of  $CX$  measurement has different properties compared with  $H/H_{XY}$  measurement. In the  $H/H_{XY}$  measurement schemes, the magic state is injected into the code, making any following measurement deterministic (absent any errors).

Generally, injection improves the rate by reducing the probability of logical error by an additional order of  $p$ . It is similarly possible to inject a  $|CX\rangle$  state at the beginning of the cultivation procedure, but some errors at that stage would not be identifiable later. For example, a failed injection that resulted in the state  $|0, +\rangle$  would not be detected by any subsequent  $CX$  measurement, since  $|0, +\rangle$  is a  $+1$  eigenstate of  $CX$  but is not the  $|CX\rangle$  state.

Instead, for  $CX$  cultivation, the two surface codes are initialized in the  $|+\rangle, |0\rangle$  state by initializing each physical qubit in the  $|+\rangle$  and  $|0\rangle$  states, respectively, and measuring all stabilizers (see Appendix J). The  $CX$  measurement is repeated  $d_1$  times, and the state is discarded following any detection event. The initial measurement projects the state to the  $|CX\rangle$  state with probability  $3/4$ , and any subsequent measurement is deterministic (again, absent any




FIG. 11. Action of transversal  $H$  on the surface-code stabilizers. Action of the transversal  $H_{XY}$ , split between  $Z$  (top) and  $X$  (bottom) stabilizers. Each panel shows the state of the stabilizer group before the operation depicted in it, and the last panel shows the final form of the stabilizers. Double lines connect stabilizers with spatially separated support.

errors). In this way, the  $CX$  scheme has a fault distance  $d_1$ , at the expense of an additional  $CX$  measurement and nondeterministic initialization.

For the same reason, the codes cannot grow during the  $CX$  cultivation stage. The  $CX$  measurement projects into an eigenspace that contains multiple states, so some weight 3 errors are not identifiable in the  $d = 5$  state. For example, consider a logical  $Z$  error on the control qubit during the  $d_1 = 3$  stage. This error has weight 3, but it will not be identifiable by a subsequent  $d_1 = 5$   $CX$  measurement, as it commutes with  $CX$ . To achieve fault distance of 5, the surface codes should be initialized and measured at  $d_1 = 5$  for the entire cultivation protocol.

In Clifford simulation of  $CX$  cultivation, our circuit measures the  $XX$  operator on two surface codes instead of the  $CX$  operator. Just like the  $CX$  projection, the  $XX$  measurement projects a pair of codes into a subspace with an additional degree of freedom, namely the  $ZZ$  operator, and not to a specific state like  $H$  and  $H_{XY}$  cultivation schemes.

Initialization in the Clifford simulation is achieved by initializing the two surface code in the state  $|+, 0\rangle$  state. We then apply a transversal  $CX$  gate between them, resulting in the state:

$$|00\rangle + |11\rangle,$$

which is a  $+1$  eigenstate of both of their logical  $X_1X_2$  and  $Z_1Z_2$  operators. We then apply a logical  $Z_1$  operator with probability  $1/4$ , which simulates the probabilistic nature of the initial  $CX$  projection.

Figure 12 shows our phase-kickback implementation for the Clifford simulation, where the  $CX$  cultivation with  $CCX$  is replaced with  $CXX$ .

Last, we turn to our definition of our complementary gap in the  $CX$  Clifford simulation. At the end of that simulation, both  $XX$  and  $ZZ$  are measured perfectly. The log-likelihood of the matching for all four possible values of  $XX$  and  $ZZ$  is evaluated, and the most likely option is chosen. The complementary gap is defined as the log-difference between the two most likely options. The  $ZZ$  operator is not measured during the protocol, but its initialization is fault-tolerant. This simulates the additional degrees of freedom of the  $+1$  subspace of  $CX$  that are not probed by the projections, but are initialized fault-tolerantly.

## APPENDIX B: NOISE MODEL

We define the noise model similarly to uniform noise model as defined in Ref. [16], adding an additional term describing the noise coming from three-qubit operation. We introduce a three-qubit depolarizing channel that applies one of the 63 (nontrivial) Pauli channels to the three-qubits with equal probability  $p/63$ . The noise channels are defined in Fig. 13 and the uniform noise model, used throughout the paper, is defined in Fig. 14. This noise

model is motivated by our decomposition of  $CSWAP$  into  $CCZ$ ,  $CX$  and single qubit rotations, see Fig. 18.

For Fig. 16, we used a simpler noise model, where single and two-qubit depolarizing errors follow every single and two-qubit gate with the same probability  $p$ , with no idling or measurement errors.

## APPENDIX C: ALTERNATIVE NOISE MODEL FOR RYDBERG ATOMS

Since our simulation assumes all-to-all connectivity, it would be most efficient to implement on a device with long-range gates such as a Rydberg Atoms based quantum computer [66,67]. Qubits based on Rydberg atoms can have coherence time much longer than their gate time. The dephasing error rate is dominated by decay during the excitation to the Rydberg state, which only happens during two-qubit gates. To label this effect, we changed the noise model of Fig. 14 by removing any idling error, and reducing the probability of single qubit errors to  $p/10$ . Since our scheme rarely uses one-qubit gates to begin with, the dominant effect is removing the idling errors.

In Fig. 7, we can see improved logical fidelities, as high-weight errors are significantly less probable. As was demonstrated in Refs. [29,30], there is a native implementation of a  $CCX$  gate Rydberg atom. It is therefore reasonable to expect future experiments would reduce the overhead of implementing  $CCX$  compared with the standard combination of multiple single and two-qubit gates.

## APPENDIX D: FULL VECTOR SAMPLING

We used a full vector sampling in order to benchmark our protocol spending roughly  $10^6$  CPU hours resulting in roughly  $10^7$  shots for each point. We only benchmarked the  $H$  cultivation protocol using the unrotated surface code, and  $H_{XY}$  protocol using the rotated surface code. The results are shown in Fig. 15, in good agreement with the Clifford approximation, up to roughly a factor of 2.

We estimate the error bars in Fig. 15 by considering variance of our Bernoulli variable, which is  $\text{Var}(X) = (1-p)p \approx p$ . For  $N$  accepted shots with  $n_{\text{error}}$  errors, we set the error bars on our estimated  $p$  to be

$$\sigma^2 = \frac{n_{\text{errors}}}{N^2}$$

## APPENDIX E: PREPARATION AND MEASUREMENT OF THE GHZ ANCILLA STATE

The GHZ ancilla state is generated by first initializing all of its qubits in the  $|0\rangle$  state beside a single *middle* qubit in the  $|+\rangle$  state, and expanding the state in stages by repeated application of  $CX$  as seen in Fig. 3 and in Ref. [12].

As mentioned in the main text, the measurement is just the expansion step done in reverse. This is equivalent to the measurement in Ref. [12], but requires a few additional

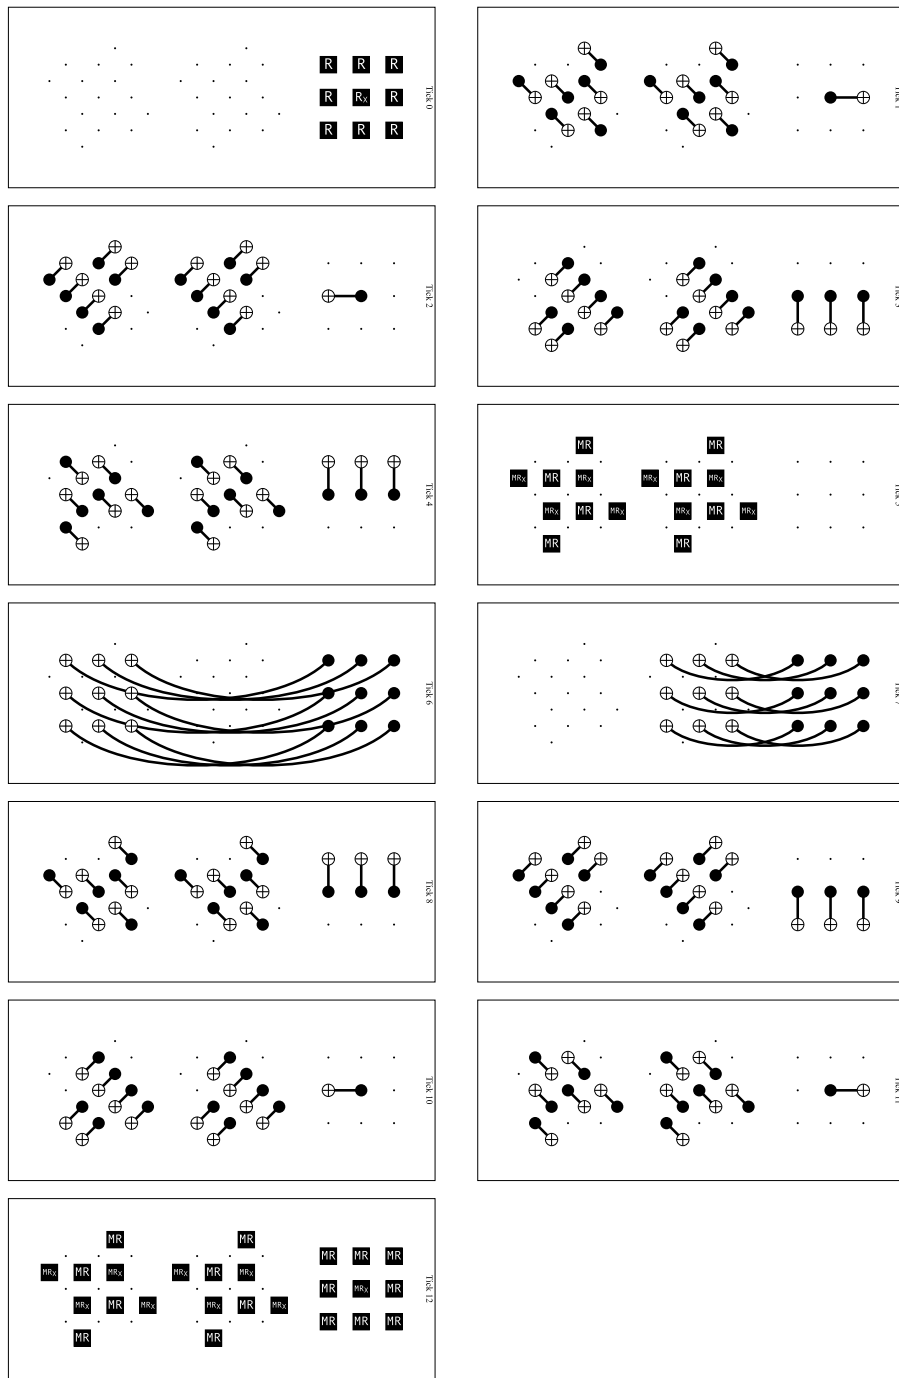


FIG. 12. Circuit measuring the  $XX$  operator—A GHZ measurement of the  $XX$  operator on a pair of surface codes. From top to bottom, left to right, each panel represent a time slice of the circuit (labeled with tick index). During the collapse and expansion stages of the GHZ state, the two surface codes are performing their respective stabilizer measurements. To measure  $CX$ , Ticks 6–7 should show a transversal  $CCX$  gates instead of  $CXX$ . [Click here to view in Crumble](#).

time steps. In Ref. [12], the stabilizers of the repetition code are measured in separate flag qubits, followed by measurement of all the qubits composing the GHZ state in  $X$  basis. Here, the application of the expansion step means that the stabilizer information can be measured by

measuring the non-*middle* qubit in the  $Z$  basis. The logical measurement of the GHZ state can be recovered by measuring the *middle* qubit in the  $X$  basis.

Because of this equivalence, our GHZ measurement behaves as a one-flag circuit as defined in Ref. [12].

Noise channel	Probability distribution of effects
MERR( $p$ )	$1 - p \rightarrow$ (report previous measurement correctly) $p \rightarrow$ (report previous measurement incorrectly; flip its result)
XERR( $p$ )	$1 - p \rightarrow I$ $p \rightarrow X$
ZERR( $p$ )	$1 - p \rightarrow I$ $p \rightarrow Z$
DEP1( $p$ )	$1 - p \rightarrow I$ $p/3 \rightarrow X$ $p/3 \rightarrow Y$ $p/3 \rightarrow Z$
DEP2( $p$ )	$1 - p \rightarrow I$ $p/15 \rightarrow \{A \otimes B \mid A, B \in \{I, X, Y, Z\}, A \otimes B \neq I \otimes I\}$
DEP3( $p$ )	$1 - p \rightarrow I$ $p/63 \rightarrow \{A \otimes B \otimes C \mid A, B, C \in \{I, X, Y, Z\}, A \otimes B \otimes C \neq I \otimes I \otimes I\}$

FIG. 13. Definition of the various noise channels in the Clifford simulation—The table was adapted from Ref. [16] with the addition of DEP3.

### APPENDIX F: CULTIVATION WITH ERASURE QUBITS

The benefit of using erasure qubits for cultivation is estimated analytically by assuming perfect detection and estimating the resulting acceptance rate. Any missed detection can be regarded as an additional Pauli error simply depolarizing the erased gate without a signal, which we assume for measurement error  $q$ , erasure rate  $e$  and Pauli rate  $p$  obeys  $eq \ll p$ . To simulate erasure, we define the Pauli noise level by a parameter  $p$ , as defined in Appendix B. At every noise location, we introduce additional erasure error with probability  $e$ . We assume the state is postselected out following every erasure error.

Figure 8 shows the acceptance rate for cultivation on both the surface code for various values of  $p, e$ , in a system with erasure bias. To compare the benefit of erasure qubits, we assume there is some additional cost associated with erasure detection [54]. A nonerasure qubit circuit with Pauli error rate  $p_0$ , is compared to an erasure circuit with erasure  $e$  and Pauli rate  $p$  with more overall noise ( $p + e > p_0$ ) and bias for erasure ( $p \ll p_0$ ). For a nonerasure qubit with Pauli rate  $p_0 = 10^{-3}$ , we can compare to  $e \in \{1 \times 10^{-3}, 2 \times 10^{-3}\}$  erasures and  $p = 10^{-4}$  residual Pauli noise. For  $e = 2p_0$ , we say that the “overhead” due to erasure conversion is two, and we see in Fig. 8 that the effect on  $d_1 = 3, H_{XY}$  surface code cultivation is to reduce the rate by a factor of 5. At this cost, the logical fidelity improves by  $\sim (p/p_0)^{d_1} \approx 10^3$ . The acceptance

Ideal gate	Noisy gate
(single qubit unitary, including idle) $U_1$	DEP1( $p$ ) · $U_1$
$CX$	DEP2( $p$ ) · $CX$
$CH_{XY}$	DEP2( $p$ ) · $CH_{XY}$
$C\bar{H}_{XY}$	DEP2( $p$ ) · $C\bar{H}_{XY}$
$CCX$	DEP3( $3p$ ) · DEP2 <sub>12</sub> ( $p$ ) · DEP2 <sub>13</sub> ( $p$ ) · $CCX$
$CSWAP \cdot CH^{\otimes 2}$	DEP3( $3p$ ) · DEP2 <sub>12</sub> ( $p$ ) · DEP2 <sub>13</sub> ( $p$ ) · $CSWAP \circ CH^{\otimes 2}$
$CCZ$	DEP3( $3p$ ) · DEP2 <sub>12</sub> ( $p$ ) · DEP2 <sub>13</sub> ( $p$ ) · $CCZ$
(reset) $R_X$	ZERR( $p$ ) · $R_X$
$R_Z$	XERR( $p$ ) · $R_Z$
$M_X$	DEP1( $p$ ) · MERR( $p$ ) · $M_X$
$M_Z$	DEP1( $p$ ) · MERR( $p$ ) · $M_Z$

FIG. 14. Uniform noise model—Each perfect operation is replaced by a noisy version, with the noise channels defined in Fig. 13. The non-Clifford operations are replaced by  $CX$ ’s which effectively measures the logical  $X$  or  $XX$  operators. The noise model does not change when replacing non-Clifford gates with Clifford gates. Subscripts in label the subset of qubits depolarized, with labeled as 1-Control, 2,3-Targets. Table adapted from Ref. [16] with some additions.

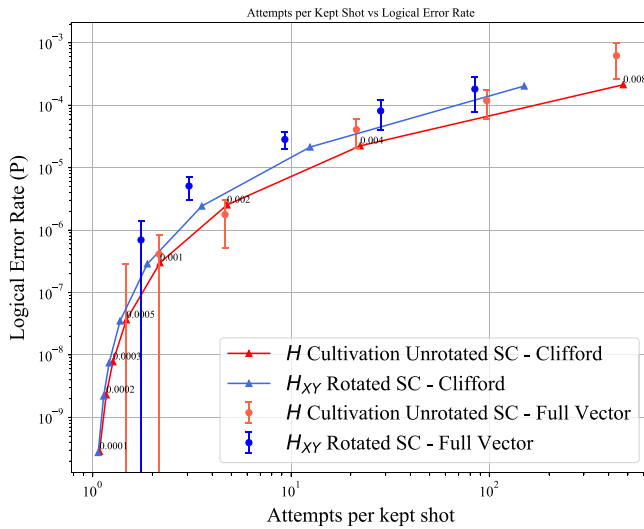


FIG. 15. Logical error rate of cultivation—comparing Clifford and Full Vector simulation—Logical error rate and acceptance rate of unrotated surface code  $H$  cultivation with Uniform noise model using both a full vector simulation (black) and an approximation using a Clifford simulator. The simulation does not include the expansion step and post selects on all stabilizer measurements before a final (perfect) logical measurement in the Magic state basis.

rate decays exponentially in both qubit count and erasure rate. Therefore, doubling the erasure rate squares the required shots.

We found the expansion stage very difficult to probe numerically. Introducing erasure into circuits for decoding requires generating a new matching graph for each circuit, which is costly. It is possible to probe higher noise rates which increase the logical error rates, and would then require fewer shots to estimate, but that too required too many shots due to the increased postselection rate.

A coarse but informative estimate of the postselection overhead associated with the expansion step can be obtained by comparing the relative overhead between the initialization and final gap-estimation stages for the case  $p = e$ . This approach assumes that erasure errors have, on average, a comparable effect on gap decoding, while allowing for a much simpler numerical evaluation. When combined with the initialization acceptance rates presented in Fig. 8, this approximation provides a lower bound on the overall cultivation rate under erasure-biased noise. It is reasonable to assume erasure errors result in higher overall logical fidelity at the expense of less postselection on the complementary gap. Erasure errors convey more information that can be utilized by the decoder (which is why the erasure threshold is much higher than the Pauli threshold for circuit model noise [50]).

In order to demonstrate this effect, we used a simpler model that does not include postselection on the stabilizers. A standard surface-code memory experiment is




FIG. 16. Logical error rate and rejection probability of surface code memory experiment—Logical error rate of  $d$  rounds memory experiments. Different points label different threshold for the complementary gap. We compare two different regimes, one with only Pauli noise and one with mix of erasure and Pauli noise obeying  $e = 2p$ . When erasure dominates the noise model, it can achieve the same logical fidelity and significantly higher rate with a factor of 2 overhead.

postselected using only the complementary gap following  $d$  rounds, with various values of erasure probability  $e$  and Pauli noise  $p$  in the same manner described earlier. The noise model is defined in Appendix B. As can be seen in Fig. 16, even when  $e = 2p$ , the gap estimation can identify a correct decoding much more easily.

## APPENDIX G: CULTIVATION WITH ERASURE QUBITS AND $d_1 = 2$

With erasure qubits, the residual Pauli error can reach low error rates. This suggests a simple proof of concept experiment for small devices with erasure qubits. Using a  $d_1 = 2$  unrotated surface code that contains five data qubits and four ancilla, the logical  $H$  operator can be measured using a single three-qubit gate (with the ancilla's doubling as GHZ ancilla). Because of the small distance of the code, only a single GHZ measurement is required. We simulated this protocol without erasure noise for various values of the Pauli noise using both of our noise models. The simulation includes the expansion to  $d_2 = 11$ . If we want to estimate how such experiment will behave with  $e = 10^{-3}$  and  $p = 10^{-4}$ , the logical fidelity will approximately be the same as our simulation with  $p = 10^{-4}$ , and the rate can be deduced from the  $p = 10^{-3}$  simulation, see Fig. 17. With these assumptions, we expect to reach  $10^{-6}$  infidelity with rate exceeding 70% of shots. Without idling noise, we can further reach infidelities of  $10^{-7}$  with over 85% acceptance rate.

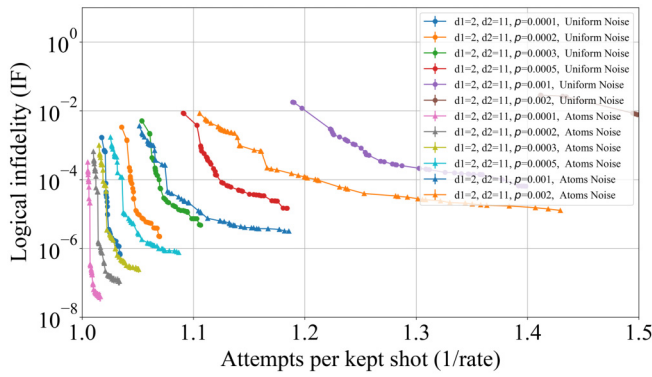


FIG. 17. Logical error rate of  $H$  cultivation with expansion for  $d_1 = 2$ —Logical error rate and acceptance rate of unrotated surface code  $H$  cultivation with Uniform and Atom-noise model for  $d_1 = 2$ , following an expansion to  $d_2 = 11$  unrotated surface code.

### APPENDIX H: $CX$ CULTIVATION WITH CODES THAT ENCODE MULTIPLE QUBITS

Since  $CX$  is transversal for all CSS codes, it can be used on pairs of codes that encode more than a single logical qubit. The eigenvalues of the transversal  $CX$  operator would still be  $\pm 1$ , but measuring it would not project every single logical qubit into the  $|CX\rangle$  state. Instead, a  $+1$  measurement would project the qubits into an entangled state.

For a pair of  $[[n, k, d]]$  quantum codes, the projection operator  $P_{\pm}^i$  projects two logical qubits, one from each code, to the  $\pm 1$  eigenspace of their combined logical  $CX$  operator. The index  $i$  labels the index of *both* qubits, each at their respective codes. If all of the stabilizers of both codes are initialized in the  $+1$  eigenstate, then a  $+1$  measurement of the transversal  $CX$  operator would effectively apply the following projection to the state:

$$\sum_{a \subseteq N} \prod_{i \notin a} P_+^i \prod_{j \in a} P_-^j, \quad (\text{H1})$$

where  $N = \{a \mid a \subseteq \{1 \dots k\}, |a| \text{ is even}\}$ , i.e., all sets of qubits with even size. This is different from projecting all pairs of qubits into the  $|CX\rangle$  state, which is achieved by the  $\prod_{i \leq k} P_+^i$  operator.

More explicitly, define the state  $|\overline{CX}\rangle = |1, -\rangle$  state, which is the single eigenstate with eigenvalue  $-1$  of the  $CX$  operator. Projecting two  $k = 2$  codes into the  $+1$  eigenstate of the transversal  $CX$  operator would project the state into the following state:

$$|CX\rangle \otimes |CX\rangle + |\overline{CX}\rangle \otimes |\overline{CX}\rangle. \quad (\text{H2})$$

Since this state is not composed of two copies of the  $|CX\rangle$  state, the method described in Ref. [9] cannot directly translate it into the  $CCX$  state. It is possible that future




FIG. 18. Circuit Decomposition of 3-qubit Controlled SWAP in terms of Clifford and CCZ gates—Circuit decomposition of Controlled SWAP gate using native gates for cold atoms—a single CCZ and two  $CX$  gates. We verified that this decomposition is in fact optimal using the code provided in Ref. [68].

schemes will allow it to be used to generate multiple useful magic states, that can effectively increase the rate of the cultivation procedure.

### APPENDIX I: THREE-QUBIT GATE COMPILATION

Our cultivation schemes make use of three types of three-qubit gates. As noted above, certain trapped-atom platforms natively support  $CCZ$  gates, which simplify the implementation of  $CCX$ . For the  $CSWAP$  gate, we employed the compilation from Ref. [68] (see Fig. 18), yielding an optimal decomposition that requires only two additional  $CX$  gates.

In our noise model, each three-qubit gate is followed by a  $\text{DEPOLARIZE}_3(3p)$  channel with probability  $3p$ , where  $p$  denotes the two-qubit gate error rate for the rest of the circuit. Since three-qubit gates are used sparingly in the protocol, the overall performance is not highly sensitive to this assumption. Figure 19 shows the effect of varying the noise strength of the  $\text{DEPOLARIZE}_3(p_{3Q})$  term while fixing  $p = 10^{-3}$ . Even when  $p_{3Q}$  is increased by more than an order of magnitude, the rate and fidelity are affected by less than a factor of two.

### APPENDIX J: INITIALIZATION AND INJECTION

Here, we provide here additional details on the initialization steps of our cultivation protocols. Each protocol begins by resetting all data qubits to either the  $|0\rangle$  or  $|+\rangle$  state. Specifically, for  $H$  cultivation we initialize in  $|0\rangle$ ; for  $H_{XY}$  we initialize in  $|+\rangle$ ; and for  $CX$  cultivation we set the control patch to  $|+\rangle$  and the target patch to  $|0\rangle$ .

As is standard for CSS codes, this initialization fixes half of the stabilizers to the  $+1$  eigenstate—namely, those aligned with the basis of the initialized data qubits. The remaining stabilizers are left in random eigenstates. To enforce that all stabilizers are in the  $+1$  eigenstate, we perform a syndrome-extraction step to identify the randomly assigned stabilizers and then apply conditional single-qubit gates to flip any  $-1$  outcomes to  $+1$ .

Following these corrections, we proceed with state injection. For  $H_{XY}$  cultivation, we collapse the weight-3 logical  $Z$  operator onto a single qubit using a pair of

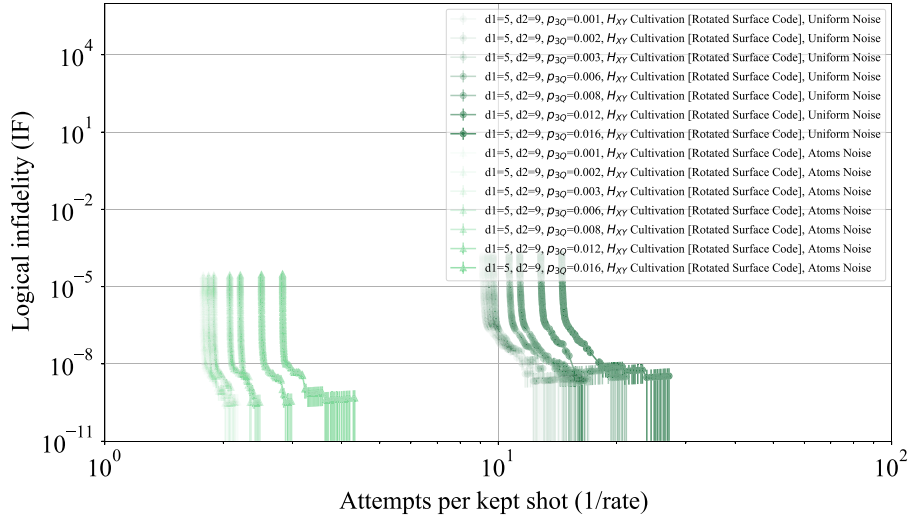


FIG. 19. Logical fidelity and acceptance rate as a function of  $p_{3Q}$ . Logical error rate and acceptance rate of rotated surface code  $H_{XY}$  cultivation with Uniform and Atom-noise model, with varying strength for  $p_{3Q}$  and constant  $p = 10^{-3}$ . In the main text, all of our simulations used  $p_{3Q} = 3p = 3 \times 10^{-3}$ , as described in B.

$CX$  gates. We then rotate that qubit by  $\pi/8$  around the  $Z$  axis and apply the  $CX$  gates in reverse, preparing the surface-code patch in the  $|T\rangle$  state.

For  $H$  cultivation, we instead perform two successive rotations: a  $\pi/8$  rotation about the logical  $X$  axis, followed by a  $\pi/4$  rotation about the logical  $Z$  axis. These operations initialize the surface-code patch in the  $|T_H\rangle$  state.

We then perform an additional syndrome extraction round, and post select on any faulty syndrome.

### APPENDIX K: INSENSITIVITY TO $d_2$ AND POSTEXPANSION ROUNDS

After the final cultivation measurement, we expand each patch from the postselected distance- $d_1$  code to a larger distance- $d_2$  code and then perform additional rounds of stabilizer measurements with soft decoding. In the main text we follow the convention of Ref. [16] and include 10 extra syndrome-extraction (SE) rounds after expansion as a simple model of decoder latency.

The role of the postexpansion rounds is to produce a full-distance syndrome history for the enlarged patch and to enable reliable soft-decoding of the final state via the complementary gap. For  $d_2 \geq 2d_1 + 1$  the expanded code is large enough to support the subsequent fault-tolerant computation without changing the asymptotic logical scaling set by  $d_1$ .

Empirically, we find that the acceptance rate and the output fidelity are essentially insensitive to the precise value of  $d_2$  and to the number of postexpansion rounds. The intuitive reason is that the dominant postselection events originate from faults occurring on the original distance- $d_1$  region.

Figure 20 illustrates this effect for  $d_1 = 3$   $H_{XY}$  cultivation on a rotated surface code. We vary the expansion distance  $d_2 \in \{7, 9, 11\}$  and also vary the number of post-expansion syndrome-extraction rounds,  $r \in \{4, \dots, d_2\}$ . To

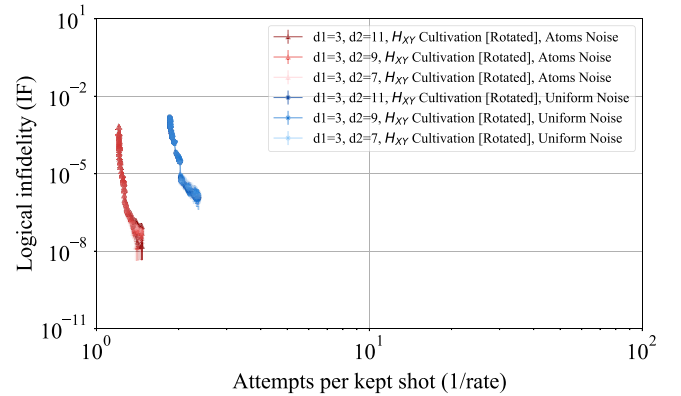


FIG. 20. Sensitivity of the postexpansion stage to the choice of  $d_2$  and to the number of postexpansion rounds  $r$  for  $d_1 = 3$   $H_{XY}$  cultivation on a rotated surface code. We plot the logical infidelity (IF) versus the number of attempts per kept shot (1/Rate) while varying the expansion distance  $d_2 \in \{7, 9, 11\}$ , for both the atom and uniform noise models. For each  $d_2$ , we sweep the postselection threshold on the decoder's complementary gap and repeat the sweep for multiple values of the number of postexpansion syndrome-extraction rounds  $r$ . All curves for different  $r \in \{4, \dots, d_2\}$  are overlaid in the same panel using low opacity for readability. The resulting trade-off curves are all within the variance of the sampling noise, across  $d_2$  and  $r$  with no observable trend. This indicates that once  $d_2 \geq 2d_1 + 1$  and sufficient postexpansion syndrome information is collected, the acceptance rate and output fidelity are largely insensitive to the specific postexpansion parameters.

present the full sweep compactly, we overlay all curves for different  $r$  on the same axes with low opacity. The resulting trade-off curves between logical infidelity and attempts per kept shot are nearly unchanged across both  $d_2$  and  $r$ , for both the uniform and atom-noise models. This supports treating acceptance rate and output fidelity as the most portable metrics across architectures, while qubits  $\times$  rounds remains strongly dependent on  $d_2$  and  $r$ .

## APPENDIX L: FULL PROTOCOL AS PSEUDOCODE

Below we provide pseudocode for the  $H$  cultivation protocol on the unrotated surface code. In the rotated surface code, each double-check is performed by inserting it between the two halves of a syndrome extraction  $CX$  cycle. We follow this by postselecting on the absence of any  $-1$  syndrome outcomes.

ALGORITHM 1. Pseudocode for cultivation on an unrotated surface code.

---

**Require:** Projection operator  $P \in \{H, H_{XY}\}$ , code distance parameter  $d_1 \in \{3, 5\}$ , complementary gap threshold  $\tau$ , final distance  $d_2$

- 1: Initialize a  $d = 3$  unrotated surface code
- 2: **if**  $P = H$  **then**
- 3:   Initialize data qubits in  $|0\rangle$
- 4: **else if**  $P = H_{XY}$  **then**
- 5:   Initialize data qubits in  $|+\rangle$
- 6: **end if**
- 7: Run a single round of syndrome extraction
- 8: Post-select on  $Z$  stabilizers (for  $H$ ) or  $X$  stabilizers (for  $H_{XY}$ ) in the  $-1$  state
- 9: **if**  $P = H$  **then**
- 10:   Apply  $Z$  corrections to set all  $X$  stabilizers to  $+1$
- 11:   Rotate around the logical  $X$  axis by  $\pi/8$
- 12:   Rotate around the logical  $Z$  axis by  $\pi/4$
- 13: **else if**  $P = H_{XY}$  **then**
- 14:   Apply  $X$  corrections to set all  $Z$  stabilizers to  $+1$
- 15:   Rotate around the logical  $Z$  axis by  $\pi/8$
- 16: **end if**
- 17: Run a single round of syndrome extraction
- 18: Post-select on any stabilizers in the  $-1$  state
- 19: *Double-check* the logical operator
- 20: **if**  $d_1 = 5$  **then**
- 21:   Expand the code to  $d = 5$  with all stabilizers in the  $+1$  eigenstate using a unitary circuit
- 22:   Run a single round of syndrome extraction
- 23:   Post-select on any stabilizers in the  $-1$  state
- 24:   *Double-check* the logical operator
- 25: **end if**
- 26: Expand the code to  $d_2$
- 27: Run 10 rounds of syndrome extraction
- 28: Decode and post-select using the complementary gap threshold  $\tau$

---

- [1] D. Aharonov and M. Ben-Or, Fault-tolerant quantum computation with constant error rate, *SIAM J. Comput.* **38**, 1207 (2008).
- [2] E. Knill, R. Laflamme, and W. H. Zurek, Resilient quantum computation, *Science* **279**, 342 (1998).
- [3] B. Eastin and E. Knill, Restrictions on transversal encoded quantum gate sets, *Phys. Rev. Lett.* **102**, 110502 (2009).
- [4] L. Lao and B. Criger, Magic state injection on the rotated surface code, in *Proceedings of the 19th ACM International Conference on Computing Frontiers, CF '22* (Association for Computing Machinery, New York, NY, USA, 2022), p. 113.
- [5] C. Gidney, Cleaner magic states with hook injection, [arXiv:2302.12292](https://arxiv.org/abs/2302.12292).
- [6] S. Bravyi and J. Haah, Magic-state distillation with low overhead, *Phys. Rev. A—At., Mol. Opt. Phys.* **86**, 052329 (2012).
- [7] D. Litinski, Magic state distillation: Not as costly as you think, *Quantum* **3**, 205 (2019).
- [8] 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).
- [9] E. Dennis, Toward fault-tolerant quantum computation without concatenation, *Phys. Rev. A* **63**, 052314 (2001).
- [10] H. Goto, Step-by-step magic state encoding for efficient fault-tolerant quantum computation, *Sci. Rep.* **4**, 7501 (2014).
- [11] H. Goto, Minimizing resource overheads for fault-tolerant preparation of encoded states of the Steane code, *Sci. Rep.* **6**, 19578 (2016).
- [12] C. Chamberland and K. Noh, Very low overhead fault-tolerant magic state preparation using redundant ancilla encoding and flag qubits, *npj Quantum Inf.* **6**, 91 (2020).
- [13] T. Itogawa, Y. Takada, Y. Hirano, and K. Fujii, Efficient magic state distillation by zero-level distillation, *PRX Quantum* **6**, 020356 (2025).
- [14] Y. Hirano, T. Itogawa, and K. Fujii, in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Montreal, QC, Canada, 2024), Vol. 1.
- [15] R. S. Gupta, N. Sundaresan, T. Alexander, C. J. Wood, S. T. Merkel, M. B. Healy, M. Hillenbrand, T. Jochym-O'Connor, J. R. Wootton, T. J. Yoder *et al.*, Encoding a magic state with beyond break-even fidelity, *Nature* **625**, 259 (2024).
- [16] C. Gidney, N. Shutty, and C. Jones, Magic state cultivation: Growing  $T$  states as cheap as CNOT gates, [arXiv:2409.17595](https://arxiv.org/abs/2409.17595).
- [17] C. Gidney, How to factor 2048 bit RSA integers with less than a million noisy qubits, [arXiv:2505.15917](https://arxiv.org/abs/2505.15917).
- [18] H. Zhou, C. Duckering, C. Zhao, D. Bluvstein, M. Cain, A. Kubica, S.-T. Wang, and M. D. Lukin, *Resource Analysis of Low-Overhead Transversal Architectures for Reconfigurable Atom Arrays*, in *Proceedings of the 52nd Annual International Symposium on Computer Architecture, ISCA '25* (Association for Computing Machinery, New York, NY, USA, 2025), p. 1432.
- [19] H. Bombin and M. A. Martin-Delgado, Topological quantum distillation, *Phys. Rev. Lett.* **97**, 180501 (2006).
- [20] A. Kubica, B. Yoshida, and F. Pastawski, Unfolding the color code, *New J. Phys.* **17**, 083026 (2015).

- [21] C. Gidney and C. Jones, New circuits and an open source decoder for the color code, [arXiv:2312.08813](#).
- [22] A. Y. Kitaev, Fault-tolerant quantum computation by anyons *Ann. Phys. (N. Y.)* **303**, 2 (2003).
- [23] J. Niu, L. Zhang, Y. Liu, J. Qiu, W. Huang, J. Huang, H. Jia, J. Liu, Z. Tao, W. Wei *et al.*, Low-loss interconnects for modular superconducting quantum processors, *Nat. Electron.* **6**, 235 (2023).
- [24] T. J. Yoder, E. Schoute, P. Rall, E. Pritchett, J. M. Gambetta, A. W. Cross, M. Carroll, and M. E. Beverland, Tour de gross: A modular quantum computer based on bivariate bicycle codes, [arXiv:2506.03094](#).
- [25] F. Marxer, A. Vepsäläinen, S. W. Jolin, J. Tuorila, A. Landra, C. Ockeloen-Korppi, W. Liu, O. Ahonen, A. Auer, L. Belzane *et al.*, Long-distance transmon coupler with CZ-gate fidelity above 99.8%, *PRX Quantum* **4**, 010314 (2023).
- [26] K. Wang, Z. Lu, C. Zhang, G. Liu, J. Chen, Y. Wang, Y. Wu, S. Xu, X. Zhu, F. Jin *et al.*, Demonstration of low-overhead quantum error correction codes, *Nat. Phys.* **22**, 308 (2026).
- [27] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, *Nature* **627**, 778 (2024).
- [28] Z.-H. Chen, M.-C. Chen, C.-Y. Lu, and J.-W. Pan, Efficient magic state cultivation on  $\mathbb{R}\mathbb{P}^2$ , *PRX Quantum* **7**, 010315 (2026).
- [29] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler *et al.*, Parallel implementation of high-fidelity multiqubit gates with neutral atoms, *Phys. Rev. Lett.* **123**, 170503 (2019).
- [30] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara *et al.*, High-fidelity parallel entangling gates on a neutral-atom quantum computer, *Nature* **622**, 268 (2023).
- [31] Y. Lu, S. Zhang, K. Zhang, W. Chen, Y. Shen, J. Zhang, J.-N. Zhang, and K. Kim, Global entangling gates on arbitrary ion qubits *Nature* **572**, 363 (2019).
- [32] S. Jacoby, Y. Vaknin, A. Retzker, and A. L. Grimsmo, Magic state injection with erasure qubits *PRX Quantum* **6**, 040323 (2025).
- [33] P. W. Shor, in *Proceedings of 37th Conference on Foundations of Computer Science* (IEEE, NW Washington, DC United States, 1996), p. 56.
- [34] Z.-H. Chen, M.-C. Chen, C.-Y. Lu, and J.-W. Pan, Transversal logical Clifford gates on the rotated surface code with reconfigurable neutral atom arrays, *Phys. Rev. Lett.* (2026).
- [35] In Appendix K we show that our protocol is insensitive to these exact numbers, which should only be understood as a convention.
- [36] H. Bombín, M. Pant, S. Roberts, and K. I. Seetharam, Fault-tolerant postselection for low-overhead magic state preparation, *PRX Quantum* **5**, 010302 (2024).
- [37] C. Gidney, M. Newman, P. Brooks, and C. Jones, Yoked surface codes, [arXiv:2312.04522](#).
- [38] D. Gottesman, in *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, edited by S. P. Corney, R. Delbourgo, and P. D. Jarvis (International, Cambridge, MA, 1999), p. 32.
- [39] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A—At., Mol. Opt. Phys.* **70**, 052328 (2004).
- [40] C. Gidney, Stim: A fast stabilizer circuit simulator, *Quantum* **5**, 497 (2021).
- [41] A. Hashim, R. K. Naik, A. Morvan, J.-L. Ville, B. Mitchell, J. M. Kreikebaum, M. Davis, E. Smith, C. Iancu, K. P. O’Brien *et al.*, Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor, *Phys. Rev. X* **11**, 041039 (2021).
- [42] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Z. Chen, K. Satzinger, R. Barends *et al.*, Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms, *Phys. Rev. Lett.* **125**, 120504 (2020).
- [43] H. J. Manetsch, G. Nomura, E. Bataille, K. H. Leung, X. Lv, and M. Endres, A tweezer array with 6100 highly coherent atomic qubits, *Nature* **647**, 60 (2025).
- [44] D. Bluvstein, A. A. Geim, S. H. Li, S. J. Evered, J. Ataiades, 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).
- [45] C. Ryan-Anderson *et al.*, Implementing fault-tolerant entangling gates on the five-qubit code and the color code, [arXiv:2208.01863](#).
- [46] C. Ryan-Anderson, J. G. Bohnet, K. Lee, D. Gresh, A. Hankin, J. P. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. C. Brown, T. M. Gatterman, S. K. Halit, K. Gilmore, J. A. Gerber, B. Neyenhuis, D. Hayes, and R. P. Stutz, Realization of real-time fault-tolerant quantum error correction, *Phys. Rev. X* **11**, 041058 (2021).
- [47] T. LeBlond, P. Groszkowski, J. G. Lietz, C. M. Seck, and R. S. Bennink, Logical error rates for the surface code under a mixed coherent and stochastic circuit-level noise model inspired by trapped ions, *Phys. Rev. Res.* **7**, 043184 (2025).
- [48] P. Steinacker, N. D. Stuyck, W. H. Lim, T. Tanttu, M. Feng, A. Nickl, S. Serrano, M. Candido, J. D. Cifuentes, F. E. Hudson *et al.*, Industry-compatible silicon spin-qubit unit cells exceeding 99% fidelity, *Nature* **646**, 81 (2025).
- [49] S. Bartolucci, P. Birchall, H. Bombin, H. Cable, C. Dawson, M. Gimeno-Segovia, E. Johnston, K. Kieling, N. Nickerson, M. Pant *et al.*, Fusion-based quantum computation, *Nat. Commun.* **14**, 912 (2023).
- [50] A. Kubica, A. Haim, Y. Vaknin, H. Levine, F. Brandão, and A. Retzker, Erasure qubits: Overcoming the T1 limit in superconducting circuits, *Phys. Rev. X* **13**, 041022 (2023).
- [51] A. Retzker and A. Kubica, US patent US20240104416A1 (2024).
- [52] H. Levine, A. Haim, J. S. Hung, N. Alidoust, M. Kalae, L. DeLorenzo, E. A. Wollack, P. Arrangoiz-Arriola, A. Khala-jhedayati, R. Sanil *et al.*, Demonstrating a long-coherence dual-rail erasure qubit using tunable transmons, *Phys. Rev. X* **14**, 011051 (2024).
- [53] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays, *Nat. Commun.* **13**, 4657 (2022).
- [54] S. Gu, Y. Vaknin, A. Retzker, and A. Kubica, Optimizing quantum error-correction protocols with erasure qubits, *PRX Quantum* **6**, 040354 (2025).

- [55] K. S. Chou, T. Shemma, H. McCarrick, T.-C. Chien, J. D. Teoh, P. Winkel, A. Anderson, J. Chen, J. C. Curtis, S. J. de Graaf *et al.*, A superconducting dual-rail cavity qubit with erasure-detected logical measurements, *Nat. Phys.* **20**, 1454 (2024).
- [56] A. Koottandavida, I. Tsioutsios, A. Kargioti, C. R. Smith, V. R. Joshi, W. Dai, J. D. Teoh, J. C. Curtis, L. Frunzio, R. J. Schoelkopf *et al.*, Erasure detection of a dual-rail qubit encoded in a double-post superconducting cavity, *Phys. Rev. Lett.* **132**, 180601 (2024).
- [57] S. Ma, G. Liu, P. Peng, B. Zhang, S. Jandura, J. Claes, A. P. Burgers, G. Pupillo, S. Puri, and J. D. Thompson, High-fidelity gates and mid-circuit erasure conversion in an atomic qubit, *Nature* **622**, 279 (2023).
- [58] P. Scholl, A. L. Shaw, R. B.-S. Tsai, R. Finkelstein, J. Choi, and M. Endres, Erasure conversion in a high-fidelity Rydberg quantum simulator, *Nature* **622**, 273 (2023).
- [59] G. Baranes, M. Cain, J. Ataiades, D. Bluvstein, J. Sinclair, V. Vuletic, H. Zhou, and M. D. Lukin, Leveraging qubit loss detection in fault-tolerant quantum algorithms, *Phys. Rev. X* **16**, 011002 (2026).
- [60] This is the qubit-cycles count of  $H$  cultivation with  $p_0 = 10^{-3}$  as reported in Ref. [16]. In the erasure case, we assume the erasure overhead of Pauli and erasure are the same, and the additional  $p = 10^{-4}$  Pauli errors have minimal additional overhead [32]. Since gap-estimation of erasure scales better than Pauli noise, this result should be understood as the bounds on qubit-cycles count.
- [61] C. Jones, Low-overhead constructions for the fault-tolerant Toffoli gate, *Phys. Rev. A* **87**, 022328 (2013).
- [62] Z. He, A. Cowtan, D. J. Williamson, and T. J. Yoder, QLDPC architectures for efficient Pauli-based computation, [arXiv:2503.10390](https://arxiv.org/abs/2503.10390).
- [63] K. Sahay, P.-K. Tsai, K. Chang, Q. Su, T. B. Smith, S. Singh, and S. Puri, Fold-transversal surface code cultivation, [arXiv:2509.05212](https://arxiv.org/abs/2509.05212).
- [64] J. Claes, Cultivating  $T$  states on the surface code with only two-qubit gates, [arXiv:2509.05232](https://arxiv.org/abs/2509.05232).
- [65] Y. Vaknin, Github, <https://github.com/tomirendo/surface-code-magic-state-cultivation> (2025).
- [66] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with Rydberg atoms, *Rev. Mod. Phys.* **82**, 2313 (2010).
- [67] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter *et al.*, Logical quantum processor based on reconfigurable atom arrays, *Nature* **626**, 58 (2024).
- [68] É. Gouzien and N. Sangouard, Provably optimal exact gate synthesis from a discrete gate, [arXiv:2503.15452](https://arxiv.org/abs/2503.15452).