Correction of chain losses in trapped ion quantum computers

Nolan J. Coble,[1,2] Min Ye,[1] and Nicolas Delfosse[1]

[1] *IonQ Inc.*

[2] *Department of CS and Joint Center for Quantum Information and Computer Science (QuICS), University of Maryland, College Park, MD 20742*

(Dated: November 21, 2025)

Neutral atom quantum computers and to a lesser extent trapped ions may suffer from atom loss. In this work, we investigate the impact of atom loss in long chains of trapped ions. Even though this is a relatively rare event, ion loss in long chains must be addressed because it destabilizes the entire chain resulting in the loss of all the qubits of the chain. We propose a solution to the chain loss problem based on (1) a quantum error correction code distributed over multiple long chains, (2) beacon qubits within each long chain to detect the loss of a chain, and (3) a decoder adapted to correct a combination of circuit faults and erasures after beacon qubits convert chain losses into erasures. We verify the chain loss correction capability of our scheme through circuit level simulations with a distributed $[[72, 12, 6]]$ BB code with beacon qubits.

I. INTRODUCTION

Trapped ions are one of the most promising platforms for quantum computing thanks to their record fidelity [1] and the ability to move ions, which provides a great flexibility in terms of architecture design as illustrated by the QCCD architecture [2].

To reach large-scale applications, a fault-tolerant architecture capable of correcting faults during a computation is necessary [3]. Faults induced by imperfect qubit operations are corrected using a quantum error correction code such as surface codes in the MUSIQC architecture [4] or quantum LDPC codes adapted to trapped ions in [5, 6].

Another important source of noise is the loss of an ion, *i.e.*, a qubit, which refers to the physical disappearance of the ion. This typically results from heating or collisions with residual background gas molecules in the vacuum chamber [7–9]. Even though the loss rate of trapped ions is typically lower than technologies like neutral atoms [10] (an ion can be trapped for days or even months in some cases [7, 11]), ion losses cannot be ignored because they spread via qubit interactions. If left uncontrolled, the loss of an ion might rapidly propagate to a large fraction of the qubits during a computation. One can deal with ion losses using erasure correction codes [12] combined with loss detection units (originally introduced as a leakage detection in [13], see also [14]), or by replacing measurements by state-selective measurements that distinguish between the three outcomes 0, 1 and loss, together with loss post-selection [15] or a delayed erasure decoder [16]. Fast reloading is also required to replace the lost ions [17]; see also neutral atom reloading [18, 19].

These approaches immediately apply to a trapped ion architecture based on ions confined in their own potential well, and merged during two-qubit gates as in the QCCD architecture [2]. However, the problem of ion loss was not considered in the case of long ion chains where a large number of ions are confined in the same potential well, despite the popularity of long chains in quantum computing experiments. Chains with five ions are demonstrated in [20, 21], chains with 10 to 20 ions in [22–26], a long chain forming a 30-qubit quantum computer is benchmarked in [27], and analog quantum simulations using chains with up to 53 ions are performed in [28, 29].

The problem of ion loss in long chains is not only unexplored, it is also significantly more challenging than in the neutral atom or short chain case because the loss of an ion is likely to induce the loss or complete destabilization of the entire chain, resulting in the simultaneous loss of many qubits at once. Indeed, Table 4.1 of [30], presents experimental data showing that for chains up to six ions, a single ion loss usually results in the entire chain being lost, or only one ion remaining. Even if some ions remain, the loss instantly breaks the delicate force balance between the trap potential and Coulomb repulsion. The remaining ions are forced to rapidly re-equilibrate, resulting in a global shift of their positions and injecting significant motional heat into the system's collective modes. This positional shift has immediate consequences for measurement: since the detection lasers are precisely aimed at the original ion locations, the displaced ions will fall outside the laser focus. Consequently, the measurement system will register no emitted photons, which suggests the ion is in the $|0\rangle$ (dark) state when, in fact, the qubit information is simply lost due to misaligned optics.

In this work, we propose a solution to the ion loss problem in long chains of trapped ions based on three ingredients that we further describe in the next paragraphs: (1) quantum error correction codes distributed over multiple long chains (2) beacon qubits within each long chain to detect the chain loss and (3) a decoder correcting a combination of circuit faults and chain losses.

Distributing the code's data qubits over multiple long chains is necessary because it is impossible to recover for the loss of all the data qubits. This excludes quantum error correction codes supported over a single long chain as in the QEC experiments reported in [24–26] or the protocols designed in [5]. The distributed quantum error correction protocols from [6] seem more robust against chain losses because each chain contains only a fraction of the

data qubits. For example, the [[144,12,12]] bivariate bicycle (BB) code of [31] is distributed over 12 long chains containing 12 data qubits each. However, when distributing a code, the data qubits contained in each chain must be selected carefully to make sure that none of the logical operators is fully supported inside a chain. Otherwise, the loss of this chain, would permanently delete some of the logical information. In this work, we refine the distribution of the BB codes to avoid this issue and to make distributed BB codes robust against chain loss.

Within each ion chain, we use one or a few qubits, that we call *beacon qubits*, to certify that the chain is still present. The basic idea is to use the fact that the outcome 1 during the measurement of a qubit corresponds to the detection of photons emitted by the corresponding ion. If a chain is lost, its ions cannot emit photons anymore and all measurements return the outcome 0. Therefore, to verify that a chain was not lost, we use beacon qubits, kept in state $|1\rangle$, and measured at regular time intervals. As long as the chain is present, the beacon measurement should return the outcome 1. If a beacon returns the outcome 0, it is likely that the chain is lost. This strategy might fail to immediately detect a chain loss in the presence of circuit faults because a fault might flip an outcome 0, hiding a chain loss. To make sure that chain losses are immediately detected, we use multiple beacon qubits per chain. We expect that a few beacons are sufficient to suppress the probability of an undetected chain loss below the logical error rate so that this phenomenon does not significantly affect the code performance. Alternatively, one can use a single beacon qubit per chain and repeat its measurement.

Once a chain loss is detected, it can be replaced by a new chain from a reservoir that is continuously reloaded [17–19]. The qubits of the replacement chain are initialized in the maximally mixed state, which effectively converts the chain loss into an erasure and we adapt the decoder to deal with a combination of circuit faults and erasure using standard techniques [32–36]. The main difference with [16] is that chain losses are immediately detected after each round of circuit operations thanks to our beacon qubits, removing the need for delayed erasure decoding.

Though our protocol differs from previous ideas in that it is distributive, it is compatible with some earlier strategies. For instance, one could replace our beacon qubits with the loss detection units proposed in [13, 14]. We prefer the beacon qubit because it does not consume any two-qubit gate. Our protocol is also compatible with the use of state-selective measurements for ancilla measurements [15, 16]. We prefer the introduction of a beacon qubit that is constantly measured to guarantee an earlier detection of chain losses. However, one could envision combining both approaches.

We perform numerical simulations showing the efficacy of our distributed quantum error correction protocol against circuit faults and chain losses. Our simulations demonstrate that the scheme mitigates these loss events, achieving logical error rates close to the no-chain-loss baseline.

In Section II, we review the distributed quantum computing model of [6] and we extend it to include chain losses. The distributed quantum error correction scheme of [6] is adapted to tolerate chain losses in Section III by introducing beacon qubits and by adjusting the mapping of the data qubits onto chains and the syndrome extraction circuit. Our numerical results are presented in Section IV. We conclude and discuss several future directions in Section V.

II. THE $2 \times L$ MODEL WITH CHAIN LOSS

We begin with a summary of the $2 \times L$ module array model proposed in [6].

- **Structure and Rows:** The architecture consists of a $2 \times L$ array of cells labeled $(b, i) \in \mathbb{Z}_2 \times \mathbb{Z}_L$, where each cell contains a module (a register of n qubits). Cells $(0, i), i \in \mathbb{Z}_L$ form the **fixed row**, and Cells $(1, i), i \in \mathbb{Z}_L$ form the **moving row**.
- **Operations:** Two-qubit unitary gates are supported either inside a module or in a pair of **aligned modules** (cells $(0, i)$ and $(1, i)$). Operations acting on different modules can be performed **simultaneously**.
- **Cyclic Shift:** The key mechanism for connectivity is the **cyclic shift** (or s-shift), which moves all modules in the moving row, transporting the module in cell $(1, i)$ to cell $(1, (i + s) \mod L)$.
- **Shift Depth:** Any cyclic shift is assumed to have **depth one**, meaning the shift duration is independent of the shift size s.

The **Sparse Cyclic Layout** (Algorithm 2 in [6]) is a quantum circuit designed to efficiently perform syndrome extraction for bivariate bicycle (BB) codes [31] within this architecture. It produces a **short-depth syndrome extraction circuit** by strategically interleaving cyclic shifts and highly parallel two-qubit gates across different modules. In the next section, we will present a slight variation of Algorithm 2 from [6]. The primary difference is that each module in our layout contains fewer data qubits. This modification ensures that the number of lost data qubits in a single chain loss remains below the code distance, as exceeding the code distance would render the chain loss unrecoverable.

Given that the loss of a single ion destabilizes the entire ion chain, our model treats this as a complete chain-loss event. We integrate this phenomenon into the $2 \times L$ module array model by assuming that each chain is lost independently at each time step with probability p_{loss}. This probabilistic model is detailed further in Appendix A.

Chain loss is particularly detrimental in a distributed setting, as it can escalate to a catastrophic failure of the

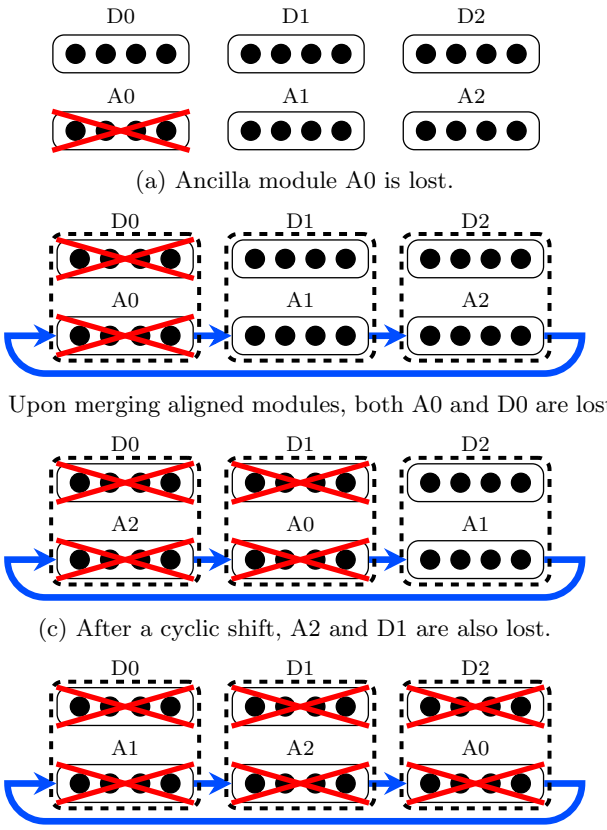

(a) Ancilla module A0 is lost.

(b) Upon merging aligned modules, both A0 and D0 are lost.

(c) After a cyclic shift, A2 and D1 are also lost.

(d) After one more cyclic shift, all modules are lost.

FIG. 1: Catastrophic chain loss: The loss of one module leads to the subsequent loss of all modules after several cyclic shift and merge operations. The top row (D0, D1, D2) contains fixed modules of data qubits, and the bottom row (A0, A1, A2) contains moving modules of ancilla qubits. A dashed rectangle indicates merged modules.

entire system. This is because the loss of a single module propagates to other functioning modules during the cyclic shift and merge operations. As this process repeats, the failure spreads rapidly, ultimately resulting in the loss of the entire array and all encoded data, as illustrated in Fig. 1.

III. CHAIN LOSS CORRECTION USING BEACON QUBITS

To detect chain-loss events, we introduce a scheme based on "beacon qubits". A dedicated beacon qubit is incorporated into each ion chain module, and its sole purpose is to be periodically measured to verify the physical presence and optical alignment of the chain. The protocol operates by initializing the beacon qubit in the $|1\rangle$ ("bright") state, which is defined to scatter photons during measurement. If the measurement detector registers the expected photons, the chain is confirmed to be present. If, however, no photons are detected, the system flags this outcome as a chain-loss event.

A single beacon qubit is susceptible to two failure modes from circuit faults: **(i) False Positives**: A functional chain is incorrectly flagged as lost. This occurs if the beacon qubit, prepared in the $|1\rangle$ state, is measured as $|0\rangle$. This can be caused by a Pauli X error flipping the qubit to $|0\rangle$ or a measurement error where the detector fails to register the emitted photons. **(ii) False Negatives**: A real chain loss is missed. This occurs if the "no photons" ($|0\rangle$) signal from a lost chain is incorrectly measured as $|1\rangle$ (a measurement error).

To address both failure modes, we can employ multiple beacon qubits per ion chain instead of a single one. The decision of whether the chain is lost is then based on a majority vote of the b beacon measurement outcomes. For example, if $b = 3$, a loss is declared only if two or more beacons return $|0\rangle$. This redundancy provides resilience against both failure modes. A single, random fault (e.g., a measurement error flipping one $|1\rangle$ to $|0\rangle$) is no longer sufficient to cause a false positive, as it will be outvoted by the other two $|1\rangle$ outcomes. Similarly, a single fault hiding a real loss (flipping one $|0\rangle$ to $|1\rangle$) will be outvoted by the other beacons that correctly report $|0\rangle$. This method allows for a significant, tunable reduction in both error rates at the cost of increased qubit overhead.

Upon detection of a chain loss, the system immediately replenishes the lost module with fresh qubits. These new qubits are initialized in the maximally-mixed state $I/2$, which signifies a total loss of information as they possess no correlation with the original encoded state. This physical replacement, combined with the knowledge of the loss event's location from the beacon qubit measurements, is equivalent to an erasure error [37, 38] which is typically easier to correct than a Pauli error [39–42].

While this scheme reliably converts chain loss into erasures, successful correction still requires a careful co-design of the layout and syndrome extraction circuit. For instance, if all data qubits supporting a logical operator reside on a single ion chain, the loss of that chain would be uncorrectable. We ensure correctability against a single chain loss by imposing the constraint that the number of data qubits per chain must be less than the code distance. This constraint necessitates modifying the syndrome extraction circuit presented in Algorithm 2 of [6].

We begin by reviewing BB code construction and the sparse cyclic layout from Algorithm 2 of [6], and then present our modification. BB codes are defined by two integers ℓ and m, and six integer pairs $(i_1, j_1), \ldots, (i_6, j_6) \in \mathbb{Z}_\ell \times \mathbb{Z}_m$. The code has $n = 2\ell m$ data qubits and $n/2 = \ell m$ stabilizers of both X and Z types. Data qubits are labeled by $(u, v, w) \in \mathbb{Z}_2 \times \mathbb{Z}_\ell \times \mathbb{Z}_m$. The X and Z ancilla qubits are labeled (X, v, w) and (Z, v, w), respectively, for $(v, w) \in \mathbb{Z}_\ell \times \mathbb{Z}_m$. An X ancilla qubit (X, v, w) measures the weight-6 X stabilizer supported on $(0, v\oplus i_1, w\oplus j_1), (0, v\oplus i_2, w\oplus j_2), (0, v\oplus i_3, w\oplus j_3), (1, v\oplus i_4, w\oplus j_4), (1, v\oplus i_5, w\oplus j_5), (1, v\oplus i_6, w\oplus j_6)$. A Z ancilla qubit (Z, v, w) measures the weight-6 Z stabilizer sup-

ported on $(0, v\ominus i_4, w\ominus j_4), (0, v\ominus i_5, w\ominus j_5), (0, v\ominus i_6, w\ominus j_6), (1, v\ominus i_1, w\ominus j_1), (1, v\ominus i_2, w\ominus j_2), (1, v\ominus i_3, w\ominus j_3)$. Here, $\oplus$ and $\ominus$ denote addition and subtraction modulo ℓ or m, respectively. Note that our description is equivalent to the standard description of BB codes using parity check matrices; see Section IV of [6] for a proof.

In [6], the modules are indexed by $w \in \mathbb{Z}_m$. Each data module consists of the data qubits in the set $M^d_w := \mathbb{Z}_2 \times \mathbb{Z}_\ell \times \{w\}$. Each ancilla module consists of the ancilla qubits in the set $M^a_w := \{X, Z\} \times \mathbb{Z}_\ell \times \{w\}$. Define two sets $\mathcal{A} = \{(i_1, j_1), (i_2, j_2), (i_3, j_3)\}$ and $\mathcal{B} = \{(i_4, j_4), (i_5, j_5), (i_6, j_6)\}$. We use them to introduce shorthand notations for the required sets of CX gates:

- $\mathrm{CX}[(i, j) \in \mathcal{A}]$ refers to the set of CX gates controlled on qubit (X, v, w) targeting qubit $(0, v \oplus i, w \oplus j)$ for all $v \in \mathbb{Z}_\ell$ and $w \in \mathbb{Z}_m$.
- $\mathrm{CX}[(i, j) \in \mathcal{B}]$ refers to the set of CX gates controlled on qubit (X, v, w) targeting qubit $(1, v \oplus i, w \oplus j)$ for all $v \in \mathbb{Z}_\ell$ and $w \in \mathbb{Z}_m$.

Algorithm 1 details the circuit from [6] for measuring the X stabilizers of BB codes.

Algorithm 1: Recap of Algorithm 2 from [6]

```
Prepare all the X ancilla qubits in the state |+⟩.
for each distinct j ∈ {j_1, j_2, ..., j_6} do
    Apply the cyclic shift aligning M^a_0 and M^d_j.
    for i ∈ Z_ℓ such that (i, j) ∈ A do
        Execute all the CX gates in CX[(i, j) ∈ A].
    for i ∈ Z_ℓ such that (i, j) ∈ B do
        Execute all the CX gates in CX[(i, j) ∈ B].
Measure all the ancilla qubits in the X basis.
```

All simulations in this paper use the $[[72, 12, 6]]$ BB code from [31], which is defined by parameters $\ell = m = 6$, $\mathcal{A} = \{(3, 0), (0, 1), (0, 2)\}$, and $\mathcal{B} = \{(0, 3), (1, 0), (2, 0)\}$. For this specific code, the set $\{j_1, \ldots, j_6\}$ contains only 4 distinct elements. Consequently, only 4 cyclic shifts are required to measure the X stabilizers.

However, for the $[[72, 12, 6]]$ BB code, this layout assigns 12 data qubits per module, significantly exceeding the code distance of 6 and rendering the loss of any single module uncorrectable. To address this issue, we modify the layout to assign only 3 qubits to each module. We assume ℓ is even, which holds for the $[[72, 12, 6]]$ code with $\ell = 6$. The data modules in our modified layout are indexed by $(r, s, w) \in \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_m$ and defined as

$$M^d_{(r,0,w)} := \{r\} \times \{0, \ldots, \ell/2 - 1\} \times \{w\},$$
$$M^d_{(r,1,w)} := \{r\} \times \{\ell/2, \ldots, \ell - 1\} \times \{w\}.$$

The ancilla modules in our modified layout are indexed by $(r, s, w) \in \{X, Z\} \times \mathbb{Z}_2 \times \mathbb{Z}_m$ and defined as

$$M^a_{(r,0,w)} := \{r\} \times \{0, \ldots, \ell/2 - 1\} \times \{w\},$$
$$M^a_{(r,1,w)} := \{r\} \times \{\ell/2, \ldots, \ell - 1\} \times \{w\}.$$

With this module assignment, one round of X stabilizer measurements for the $[[72, 12, 6]]$ BB code can be completed using only 6 shifts, as opposed to 4 shifts in the case of 12 qubits per module. The 6 shifts are detailed below.

The first shift aligns $M^a_{(X,s,w)}$ with $M^d_{(0,1-s,w)}$ for all $s \in \mathbb{Z}_2$ and $w \in \mathbb{Z}_6$. This alignment enables the execution of all CX gates in $\mathrm{CX}[(3, 0) \in \mathcal{A}]$.

The second shift aligns $M^a_{(X,s,w)}$ with $M^d_{(0,s,w\oplus 1)}$ for all $s \in \mathbb{Z}_2$ and $w \in \mathbb{Z}_6$. This alignment enables the execution of all CX gates in $\mathrm{CX}[(0, 1) \in \mathcal{A}]$.

The third shift aligns $M^a_{(X,s,w)}$ with $M^d_{(0,s,w\oplus 2)}$ for all $s \in \mathbb{Z}_2$ and $w \in \mathbb{Z}_6$. This alignment enables the execution of all CX gates in $\mathrm{CX}[(0, 2) \in \mathcal{A}]$.

The 4th shift aligns $M^a_{(X,s,w)}$ with $M^d_{(1,s,w\oplus 3)}$ for all $s \in \mathbb{Z}_2$ and $w \in \mathbb{Z}_6$. This alignment enables the execution of all CX gates in $\mathrm{CX}[(0, 3) \in \mathcal{B}]$.

The 5th shift aligns $M^a_{(X,s,w)}$ with $M^d_{(1,s,w)}$ for all $s \in \mathbb{Z}_2$ and $w \in \mathbb{Z}_6$. The 6th shift aligns $M^a_{(X,s,w)}$ with $M^d_{(1,1-s,w)}$ for all $s \in \mathbb{Z}_2$ and $w \in \mathbb{Z}_6$. These two shifts together enable the execution of all CX gates in $\mathrm{CX}[(1, 0) \in \mathcal{B}]$ and $\mathrm{CX}[(2, 0) \in \mathcal{B}]$.

IV. NUMERICAL RESULTS

We perform Monte Carlo simulations using Stim [43] and employ the BP-OSD decoder [44–46] for decoding. Specifically, we use the BP-OSD implementation in [46] with 10,000 min-sum BP iterations followed by order-5 combination-sweep OSD. We assume the beacon qubits provide perfect detection of chain loss. As previously analyzed, the false positives and false negatives of this mechanism can be exponentially suppressed by increasing the number of beacon qubits per chain. Under our model, this addition does not affect the logical error rate, as all qubit measurements are performed in parallel. The only trade-off is the increased qubit overhead.

The noise model used in our simulation is detailed in Appendix A. Since Stim does not natively support chain or qubit loss events, we developed a sampling-based method to simulate them, as explained in Appendix B. All simulations in this paper use the $[[72, 12, 6]]$ BB code from [31]. In Fig. 2 and Fig. 3, the "logical error rate" on the y-axis refers specifically to the X logical error rate.

In Fig. 2, we plot the logical error rate as a function of the 2-qubit gate error rate p, setting the chain loss probability to $p_{\mathrm{loss}} = p^\alpha$ for $\alpha \in \{1.7, 1.9, 2.1\}$. This simulation assumes an ideal scenario where beacon qubit measurements are instantaneous, i.e., they take zero time and induce no idling noise on other qubits. Meanwhile the measurements of other qubits still take time according to the noise model. Under this assumption, our distributed layout together with the beacon qubit scheme performs extremely well. For low chain loss probabilities $p_{\mathrm{loss}} = p^{1.9}$ and $p_{\mathrm{loss}} = p^{2.1}$, the logical error rate is essentially identical to the no-loss baseline. Even for the

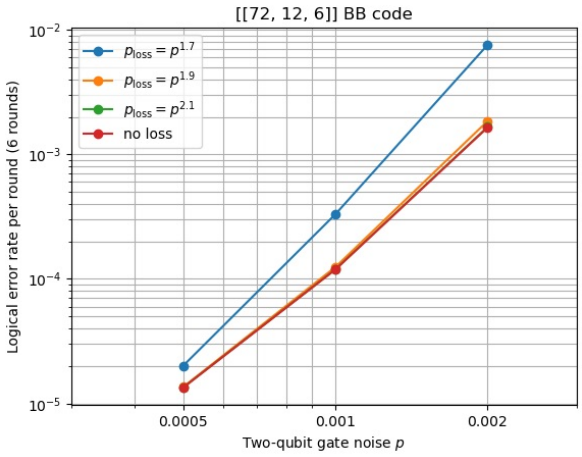

FIG. 2: Logical error rate as a function of p, with $p_{\text{loss}} = p^{\alpha}$ for $\alpha \in \{1.7, 1.9, 2.1\}$.

relatively high loss rate of $p_{\text{loss}} = p^{1.7}$, the logical error rate at $p = 10^{-3}$ remains within a 3x factor of the no-loss case.

In Section C, we present simulation results under a more pessimistic beacon qubit measurement model. In this scenario, beacon measurements are no longer instantaneous and are instead assumed to have the same time duration as all other qubit measurements.

In Appendix D, we investigate the effect of a single chain loss, showing that its impact on the logical error rate is highly dependent on the time step at which it occurs.

V. CONCLUSIONS AND FUTURE DIRECTIONS

We proposed a strategy to correct chain loss in trapped ion architecture using a quantum error correction code distributed over multiple long chains and beacon qubits added to each long chain to detect the chain loss.

Our protocol is based on the distributed quantum error correction scheme proposed in [6]. It would be interesting to adapt other distributed quantum error correction schemes [47–49] to the correction of chain losses.

The problem considered in this paper is analogous to the classical use of erasure correction codes in distributed systems, which must deal with nodes that become temporarily or permanently inaccessible. Reed Solomon codes [50] have been used in this context for decades [51] and are part of Google's file system Colossus [52] and Facebook's file system [53, 54], whereas Microsoft's file system uses local reconstruction codes [55]. Another popular choice is network codes [56], recently generalized to the quantum setting [57]. One could consider importing ideas from classical distributed computing to the quantum setting to deal with chain losses.

Our numerical results show the viability of this approach and demonstrate that this protocol benefits from fast beacon qubit measurements. It would be interesting to explore the design of such fast measurements, exploiting for instance the fact that beacon qubits never interact with other qubits (through two-qubit gates), which makes them less affected by noise, and could reduce the required cooling time which is a major contribution to the operation time of trapped ion quantum computers [58].

To limit the time overhead of our chain loss correction protocol, we may consider measuring the beacon qubit less often and focusing on the most sensitive locations in the circuit, that is the location where a chain loss is the most likely to induce a logical error. Section D makes a first step in that direction by investigating the impact of a chain loss at all possible location of the syndrome circuit.

We focused on the correction of ion loss or chain loss. A related noise source is the leakage error which maps the state of an ion outside of the two levels used to encode a qubit. Leakage errors are less harmful than ion loss for long chain because they do not trigger a global error affecting the whole chain. Therefore, we expect that standard leakage detection and correction techniques can be applied [13, 59–67].

For simplicity, we assume that a lost chain can be immediately replaced by another chain from a reservoir regularly reloaded. We leave the design of the reservoir and the chain routing for future work.

VI. ACKNOWLEDGMENT

The authors thank Jeremy Sage, Aharon Brodutch, Edwin Tham, Felix Tripier, JP Marceaux, Woochang Chung and John Gamble and the whole IonQ team for useful and insightful discussions.

[1] A. Hughes, R. Srinivas, C. Löschnauer, H. Knaack, R. Matt, C. Ballance, M. Malinowski, T. Harty, and R. Sutherland, Trapped-ion two-qubit gates with¿ 99.99% fidelity without ground-state cooling, arXiv preprint arXiv:2510.17286 (2025).

[2] D. Kielpinski, C. Monroe, and D. J. Wineland, Architecture for a large-scale ion-trap quantum computer, Nature **417**, 709 (2002).

[3] P. W. Shor, Fault-tolerant quantum computation, in *Proceedings of 37th conference on foundations of computer science* (IEEE, 1996) pp. 56–65.

[4] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim, Large-scale modular quantum-computer architecture with atomic mem-

ory and photonic interconnects, Physical Review A **89**, 022317 (2014).
[5] M. Ye and N. Delfosse, Quantum error correction for long chains of trapped ions, arXiv:2503.22071 (2025).
[6] E. Tham, M. Ye, I. Khait, J. Gamble, and N. Delfosse, Distributed fault-tolerant quantum memories over a $2\times L$ array of qubit modules, arXiv:2508.01879 (2025).
[7] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, Trapped-ion quantum computing: Progress and challenges, Applied physics reviews **6** (2019).
[8] N. P. De Leon, K. M. Itoh, D. Kim, K. K. Mehta, T. E. Northup, H. Paik, B. Palmer, N. Samarth, S. Sangtawesin, and D. W. Steuerman, Materials challenges and opportunities for quantum computing hardware, Science **372**, eabb2823 (2021).
[9] T. Strohm, K. Wintersperger, F. Dommert, D. Basilewitsch, G. Reuber, A. Hoursanov, T. Ehmer, D. Vodola, and S. Luber, Ion-based quantum computing hardware: Performance and end-user perspective, arXiv:2405.11450 (2024).
[10] 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).
[11] G. Gabrielse, X. Fei, L. Orozco, R. Tjoelker, J. Haas, H. Kalinowsky, T. Trainor, and W. Kells, Thousandfold improvement in the measured antiproton mass, Physical review letters **65**, 1317 (1990).
[12] J. Vala, K. B. Whaley, and D. S. Weiss, Quantum error correction of a qubit loss in an addressable atomic system, Physical Review A—Atomic, Molecular, and Optical Physics **72**, 052318 (2005).
[13] J. Preskill, Fault-tolerant quantum computation, Introduction to quantum computation and information **213** (1998).
[14] H. Perrin, S. Jandura, and G. Pupillo, Quantum error correction resilient against atom loss, Quantum **9**, 1884 (2025).
[15] B. W. Reichardt, A. Paetznick, D. Aasen, I. Basov, J. M. Bello-Rivas, P. Bonderson, R. Chao, W. van Dam, M. B. Hastings, R. V. Mishmash, *et al.*, Fault-tolerant quantum computation with a neutral atom processor, arXiv preprint arXiv:2411.11822 (2024).
[16] G. Baranes, M. Cain, J. Ataides, D. Bluvstein, J. Sinclair, V. Vuletic, H. Zhou, and M. D. Lukin, Leveraging atom loss errors in fault tolerant quantum algorithms, arXiv preprint arXiv:2502.20558 (2025).
[17] C. D. Bruzewicz, R. McConnell, J. Chiaverini, and J. M. Sage, Scalable loading of a two-dimensional trapped-ion array, Nature communications **7**, 13005 (2016).
[18] Y. Li, Y. Bao, M. Peper, C. Li, and J. D. Thompson, Fast, continuous and coherent atom replacement in a neutral atom qubit array, arXiv:2506.15633 (2025).
[19] N.-C. Chiu, E. C. Trapp, J. Guo, M. H. Abobeih, L. M. Stewart, S. Hollerith, P. L. Stroganov, M. Kalinowski, A. A. Geim, S. J. Evered, *et al.*, Continuous operation of a coherent 3,000-qubit system, Nature , 1 (2025).
[20] T. Monz, D. Nigg, E. A. Martinez, M. F. Brandl, P. Schindler, R. Rines, S. X. Wang, I. L. Chuang, and R. Blatt, Realization of a scalable shor algorithm, Science **351**, 1068 (2016).
[21] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Demonstration of a small programmable quantum computer with atomic qubits, Nature **536**, 63 (2016).
[22] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, 14-qubit entanglement: Creation and coherence, Physical Review Letters **106**, 130506 (2011).
[23] Y. Nam, J.-S. Chen, N. C. Pisenti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf, *et al.*, Ground-state energy estimation of the water molecule on a trapped-ion quantum computer, npj Quantum Information **6**, 33 (2020).
[24] L. Egan, D. M. Debroy, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Newman, M. Li, K. R. Brown, M. Cetina, *et al.*, Fault-tolerant control of an error-corrected qubit, Nature **598**, 281 (2021).
[25] L. N. Egan, *Scaling quantum computers with long chains of trapped ions*, Ph.D. thesis, University of Maryland, College Park (2021).
[26] L. Postler, S. Heuβen, I. Pogorelov, M. Rispler, T. Feldker, M. Meth, C. D. Marciniak, R. Stricker, M. Ringbauer, R. Blatt, *et al.*, Demonstration of fault-tolerant universal quantum gate operations, Nature **605**, 675 (2022).
[27] J.-S. Chen, E. Nielsen, M. Ebert, V. Inlek, K. Wright, V. Chaplin, A. Maksymov, E. Páez, A. Poudel, P. Maunz, *et al.*, Benchmarking a trapped-ion quantum computer with 30 qubits, Quantum **8**, 1516 (2024).
[28] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, Nature **551**, 601 (2017).
[29] F. Kranzl, M. K. Joshi, C. Maier, T. Brydges, J. Franke, R. Blatt, and C. F. Roos, Controlling long ion strings for quantum simulation and precision measurements, Physical Review A **105**, 052426 (2022).
[30] G. D. Vittorini, *Stability of ion chains in a cryogenic surface-electrode ion trap*, Ph.D. thesis, Georgia Institute of Technology (2013).
[31] 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).
[32] N. Delfosse and N. H. Nickerson, Almost-linear time decoding algorithm for topological codes, Quantum **5**, 595 (2021).
[33] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, Erasure conversion for fault-tolerant quantum computing in alkaline earth rydberg atom arrays, Nature communications **13**, 4657 (2022).
[34] A. Kubica, A. Haim, Y. Vaknin, H. Levine, F. Brandão, and A. Retzker, Erasure qubits: Overcoming the t 1 limit in superconducting circuits, Physical Review X **13**, 041022 (2023).
[35] M. Kang, W. C. Campbell, and K. R. Brown, Quantum error correction with metastable states of trapped ions using erasure conversion, PRX Quantum **4**, 020358 (2023).
[36] S. Gu, A. Retzker, and A. Kubica, Fault-tolerant quantum architectures based on erasure qubits, Physical Review Research **7**, 013249 (2025).
[37] M. Grassl, T. Beth, and T. Pellizzari, Codes for the quantum erasure channel, Physical Review A **56**, 33 (1997).

[38] C. H. Bennett, D. P. DiVincenzo, and J. A. Smolin, Capacities of quantum erasure channels, Physical Review Letters **78**, 3217 (1997).

[39] N. Delfosse and G. Zémor, Linear-time maximum likelihood decoding of surface codes over the quantum erasure channel, Physical Review Research **2**, 033042 (2020).

[40] K. Chang, S. Singh, J. Claes, K. Sahay, J. Teoh, and S. Puri, Surface code with imperfect erasure checks, arXiv preprint arXiv:2408.00842 (2024).

[41] N. Connolly, V. Londe, A. Leverrier, and N. Delfosse, Fast erasure decoder for hypergraph product codes, Quantum **8**, 1450 (2024).

[42] M. Gökduman, H. Yao, and H. D. Pfister, Erasure decoding for quantum ldpc codes via belief propagation with guided decimation, in *2024 60th Annual Allerton Conference on Communication, Control, and Computing* (IEEE, 2024) pp. 1–8.

[43] C. Gidney, Stim: a fast stabilizer circuit simulator, Quantum **5**, 497 (2021).

[44] P. Panteleev and G. Kalachev, Degenerate quantum LDPC codes with good finite length performance, Quantum **5**, 585 (2021).

[45] J. Roffe, D. R. White, S. Burton, and E. Campbell, Decoding across the quantum low-density parity-check code landscape, Physical Review Research **2** (2020).

[46] J. Roffe, LDPC: Python tools for low density parity check codes (2022).

[47] J. Sinclair, J. Ramette, N. Breuckmann, and V. Vuletic, Fault-tolerant connection of error-corrected qubits with noisy links, in *APS Division of Atomic, Molecular and Optical Physics Meeting Abstracts*, Vol. 2023 (2023) pp. X04–005.

[48] S. de Bone, P. Möller, C. E. Bradley, T. H. Taminiau, and D. Elkouss, Thresholds for the distributed surface code in the presence of memory decoherence, AVS Quantum Science **6** (2024).

[49] E. Sutcliffe, B. Jonnadula, C. L. Gall, A. E. Moylett, and C. M. Westoby, Distributed quantum error correction based on hyperbolic floquet codes, arXiv:2501.14029 (2025).

[50] I. S. Reed and G. Solomon, Polynomial codes over certain finite fields, Journal of the society for industrial and applied mathematics **8**, 300 (1960).

[51] J. S. Plank, A tutorial on reed–solomon coding for fault-tolerance in raid-like systems, Software: Practice and Experience **27**, 995 (1997).

[52] E. Ma, Colossus: Successor to the google file system (gfs), `https://www.systutorials.com/colossus-successor-to-google-file-system-gfs/` (2012), accessed: 2025-11-15.

[53] S. Muralidhar, W. Lloyd, S. Roy, C. Hill, E. Lin, W. Liu, S. Pan, S. Shankar, V. Sivakumar, L. Tang, *et al.*, f4: Facebook's warm {BLOB} storage system, in *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)* (2014) pp. 383–398.

[54] K. V. Rashmi, N. B. Shah, D. Gu, H. Kuang, D. Borthakur, and K. Ramchandran, A" hitchhiker's" guide to fast and efficient data reconstruction in erasure-coded data centers, in *Proceedings of the 2014 ACM conference on SIGCOMM* (2014) pp. 331–342.

[55] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin, Erasure coding in windows azure storage, in *2012 USENIX Annual Technical Conference (USENIX ATC 12)* (2012) pp. 15–26.

[56] A. G. Dimakis, K. Ramchandran, Y. Wu, and C. Suh, A survey on network codes for distributed storage, Proceedings of the IEEE **99**, 476 (2011).

[57] N. Delfosse and G. Zémor, Correction of circuit faults in a stacked quantum memory using rank-metric codes, arXiv preprint arXiv:2411.09173 (2024).

[58] S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn, *et al.*, A race-track trapped-ion quantum processor, Physical Review X **13**, 041052 (2023).

[59] P. Aliferis and B. M. Terhal, Fault-tolerant quantum computation for local leakage faults, arXiv preprint quant-ph/0511065 (2005).

[60] M. Suchara, A. W. Cross, and J. M. Gambetta, Leakage suppression in the toric code, in *2015 IEEE International Symposium on Information Theory (ISIT)* (IEEE, 2015) pp. 1119–1123.

[61] J. M. Auger, S. Bergamini, and D. E. Browne, Blueprint for fault-tolerant quantum computation with rydberg atoms, Physical Review A **96**, 052320 (2017).

[62] R. Stricker, D. Vodola, A. Erhard, L. Postler, M. Meth, M. Ringbauer, P. Schindler, T. Monz, M. Müller, and R. Blatt, Experimental deterministic correction of qubit loss, Nature **585**, 207 (2020).

[63] I. Cong, H. Levine, A. Keesling, D. Bluvstein, S.-T. Wang, and M. D. Lukin, Hardware-efficient, fault-tolerant quantum computation with rydberg atoms, Physical Review X **12**, 021049 (2022).

[64] K. C. Miao, M. McEwen, J. Atalaya, D. Kafri, L. P. Pryadko, A. Bengtsson, A. Opremcak, K. J. Satzinger, Z. Chen, P. V. Klimov, *et al.*, Overcoming leakage in quantum error correction, Nature Physics **19**, 1780 (2023).

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

[66] C.-C. Yu, Z.-H. Chen, Y.-H. Deng, M.-C. Chen, C.-Y. Lu, and J.-W. Pan, Processing and decoding rydberg decay error with mbqc, arXiv preprint arXiv:2411.04664 (2024).

[67] C.-C. Yu, Y.-H. Deng, M.-C. Chen, C.-Y. Lu, and J.-W. Pan, Locating rydberg decay error in swap-lru, arXiv preprint arXiv:2503.01649 (2025).

[68] A Pauli error with probability 0.75 (i.e., applying X, Y, Z, or I with 1/4 probability each) sends any qubit to the maximally-mixed state, regardless of its original state.

Appendix A: Noise model in the simulation

We used the chain model of [5] as our noise model to simulate qubit operations inside modules. The model assumes that two-qubit gates are sequential inside a module, but gates acting on distinct modules can be performed simultaneously.

- **Two-qubit gates**: error rate p.
- **Single-qubit operations** (including unitary gate, reset, and measurement): error rate $p/10$.
- **Idle qubits**: error rate $p/100$.

- **Measurement duration (τ_m)**: We assume unmeasured qubits undergo $\tau_m = 30$ rounds of idle noise during a measurement operation.

- **Cyclic Shift duration (τ_s)**: A cyclic shift is followed by depolarizing noise on all qubits with a rate of $\tau_s p/100$, where $\tau_s = 20$. This means all qubits suffer from τ_s rounds of idle noise during a shift.

Our simulation model assumes all operations (unitary gates, resets, measurements, and cyclic shifts) have depth one in the syndrome extraction circuit, executing in a single time step. In practice, the chain loss probability for a given time step may depend on the specific operations executed, as different operations usually have different durations. However, we neglect such dependence and simply model chain loss as an independent event occurring at any time step, to any module, with a uniform probability $p_{\text{loss}} = p^\alpha$. Here, p is the two-qubit gate error rate and α is a model parameter. To reflect that chain loss is a rare event relative to gate errors, we set $\alpha > 1$. In our simulations, we specifically investigate $\alpha \in \{1.5, 1.7, 1.9, 2.1\}$, which ensures the chain loss probability p_{loss} is significantly smaller than p.

In the syndrome extraction circuit, data and ancilla modules alternate between merging with each other and standing alone. We assume that if the loss occurs while two modules are merged, the entire merged module (all 3 data and 3 ancilla qubits) is erased to the maximally-mixed state. On the other hand, if the loss occurs during an unmerged step, only the 3 qubits in the lost module are erased. We further assume that the loss probability is independent of module size, meaning a merged data/ancilla module is just as likely to be lost as an unmerged one.

Appendix B: Method for simulating chain loss in Stim

One difficulty with Stim is that it does not natively support chain or qubit loss events. We circumvent this limitation by first simulating conditional logical error probabilities and then using the law of total probability to compute the overall logical error rate. Specifically, by fixing a loss event $\text{Loss}(i, j)$, denoting the loss of the i-th chain at the j-th time step, we can efficiently simulate the conditional probability $\mathbb{P}(\text{logical error}|\text{Loss}(i, j))$. Indeed, simulating $\mathbb{P}(\text{logical error}|\text{Loss}(i, j))$ only requires two modifications to the original syndrome extraction circuit (which does not include chain loss or beacon qubits): (1) Add idling noise to unmeasured qubits during the measurements of beacon qubits; (2) At time step j, apply an independent single-qubit Pauli error with probability 0.75 to every qubit in the i-th chain [68]. Based on this modified circuit, Stim compiles a new detector error model, which the BP-OSD decoder uses as its Tanner graph. This approach generalizes, allowing us to simulate conditional error probabilities for two or more loss events, such as $\mathbb{P}(\text{logical error}|\text{Loss}(i_1, j_1) \text{ and } \text{Loss}(i_2, j_2))$.

Our simulation runs 6 rounds of syndrome extraction for the $[[72, 12, 6]]$ BB code. In this circuit, there are $N = 13,824$ possible loss events $\text{Loss}(i, j)$, where i is the chain index and j is the time step. The probability of exactly k chain loss events occurring during the entire syndrome extraction circuit is given by $p_k = \binom{N}{k} p_{\text{loss}}^k (1-p_{\text{loss}})^{N-k}$. Under a pessimistic assumption $\mathbb{P}(\text{logical error}|k \text{ chain losses}) = 1$ for all $k \geq 3$, the overall logical error rate can be approximated by the upper bound

$$\begin{aligned}\mathbb{P}(\text{logical error}) \lesssim\ & p_0 \mathbb{P}(\text{logical error}|0 \text{ chain losses}) \\ & + p_1 \mathbb{P}(\text{logical error}|1 \text{ chain loss}) \\ & + p_2 \mathbb{P}(\text{logical error}|2 \text{ chain losses}) \\ & + (1 - p_0 - p_1 - p_2).\end{aligned} \tag{B1}$$

The $\mathbb{P}(\text{logical error}|0 \text{ chain losses})$ term is obtained by simulating the original syndrome extraction circuit without chain loss or beacon qubits. Simulating the 1-loss term, $p_1 \mathbb{P}(\text{logical error}|1 \text{ chain loss}) = p_{\text{loss}}(1 - p_{\text{loss}})^{N-1} \sum_{(i,j)} \mathbb{P}(\text{logical error}|\text{Loss}(i, j))$, would require generating a separate circuit for all $N = 13,824$ possible loss events $\text{Loss}(i, j)$. Due to resource constraints, we approximate this by randomly sampling 600 pairs of (i, j) and simulating their respective conditional probabilities, $\mathbb{P}(\text{logical error}|\text{Loss}(i, j))$. We then estimate the full sum $\sum_{(i,j)} \mathbb{P}(\text{logical error}|\text{Loss}(i, j))$ as N times the average conditional error probability of these 600 samples. A similar sampling approach is used for the third term, which is estimated from the average of 600 random double-loss samples of the form $\mathbb{P}(\text{logical error}|\text{Loss}(i_1, j_1) \text{ and } \text{Loss}(i_2, j_2))$.

Appendix C: Simulation results under a pessimistic beacon qubit measurement model

In Fig. 3, we plot the logical error rate as a function of the 2-qubit gate error rate p with $p_{\text{loss}} = p^\alpha$ for fixed α. We simulate two settings: (1) The solid curve ("instantaneous") represents beacon qubit measurements that take 0 time. This models an ideal scenario of free and perfect loss detection, while other measurements still take time according to the noise model. (2) The dashed curve ("normal") represents standard beacon qubit measurements, which are assumed to take the same amount of time as other measurements. In both settings, all beacon qubits are measured after every time step. This allows for the immediate detection of a chain loss, which is then transformed into an erasure by replacing the lost qubits with fresh ones in the maximally-mixed state.

In Fig. 3, the solid curves for $\alpha = 1.9$ and 2.1 are nearly identical, indicating that for $\alpha \geq 1.9$, chain loss events constitute a negligible part of the overall logical error rate, as further decreasing the chain loss probability yields very little improvement. In contrast, decreasing α

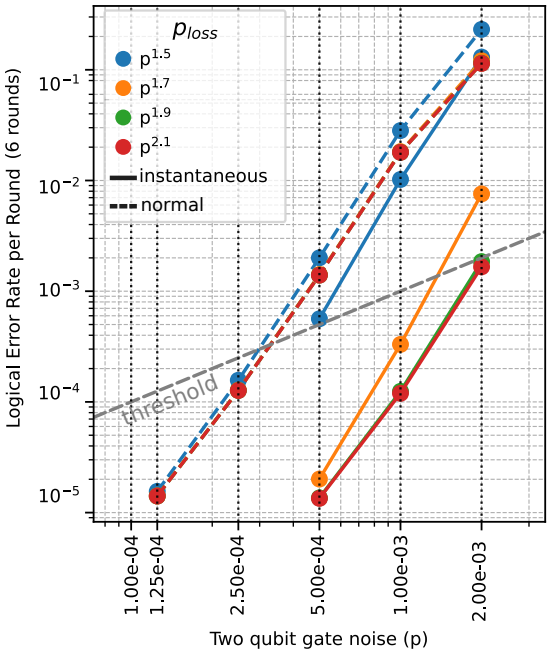

FIG. 3: Logical error rate as a function of p, with $p_{\text{loss}} = p^{\alpha}$ for $\alpha \in \{1.5, 1.7, 1.9, 2.1\}$. We plot dashed curves only for $\alpha \in \{1.5, 2.1\}$, as the 4 dashed curves are closely clustered.

from 1.9 to 1.5 causes a degradation of nearly two orders of magnitude, showing that chain loss becomes the dominant source of logical errors in this regime. This is also reflected in the threshold (the break-even point where the logical error rate per round equals p): for $\alpha \geq 1.9$, the threshold is roughly $2{\cdot}10^{-3}$, while for $\alpha = 1.5$, it degrades to $5 \cdot 10^{-4}$. The gap between the dashed and solid curves highlights the benefit of fast beacon qubit measurements. For $p_{\text{loss}} = p^{1.5}$, the two curves differ only by a factor of about 3, indicating that chain loss detection based on beacon qubits is nearly as effective as free and perfect loss detection for large p_{loss}. However, for $p_{\text{loss}} = p^{2.1}$, the curves differ by a factor of more than 100. This large gap shows that fast beacon qubit measurements become significantly more valuable when p_{loss} is small.

Appendix D: Effect of single chain loss

We investigate the effect of single chain loss within the 6-round, 480-timestep syndrome extraction circuit for the $[[72, 12, 6]]$ BB code. The system comprises 36 active modules: 24 data modules (indexed 0-23) and 12 ancilla modules (indexed 24-35). Note that while 24 ancilla modules exist in total (12 for X and 12 for Z), only 12 are active at any given time, as X and Z stabilizer measurements never overlap in the circuit. The number of independent chain modules (potential loss targets) varies throughout the circuit. At time steps for operations like ancilla preparation or measurement, all 36 modules are separate. During two-qubit gate operations, the 12 ancilla modules merge with 12 data modules, resulting in 12 merged modules and 12 un-merged data modules. Summing the number of available modules at each of the 480 time steps, we find the total number of possible single chain loss events $\text{Loss}(i, j)$ is $N = 13,824$. All plots in this appendix (Fig. 4 and Fig. 5) assume "instantaneous" beacon qubit measurements (i.e., they take 0 time).

We first examine the effect of losing only data module 0 at different time steps in the circuit, with the results shown in Fig. 4. As previously mentioned, this data module alternates between being merged with an ancilla module and standing alone. Our simulation assumes that if the loss occurs while the module is merged, the entire merged module (all 3 data and 3 ancilla qubits) is erased to the maximally-mixed state. Conversely, if the loss occurs during an unmerged step, only the 3 data qubits are erased.

In Fig. 4, the X logical error rate spikes if the data module is lost during Z stabilizer measurements, while the Z logical error rate spikes if the loss occurs during X stabilizer measurements. This phenomenon is independent of the measurement order. As Fig. 4 illustrates, the top plot (X-then-Z order) and the bottom plot (Z-then-X order) exhibit the same spiking behavior in their respective logical error rates. We also note that the X logical error rate is consistently slightly higher than the Z logical error rate across all four simulated circuits (2 values of p and both X/Z measurement orders).

In Fig. 5, we simulate the effect of a single, random chain module loss occurring at a random time step. We assume the loss probability is independent of module size, meaning a merged data/ancilla module is just as likely to be lost as an unmerged one. We observe that the logical error rate decreases exponentially as the 2-qubit gate error rate decreases.

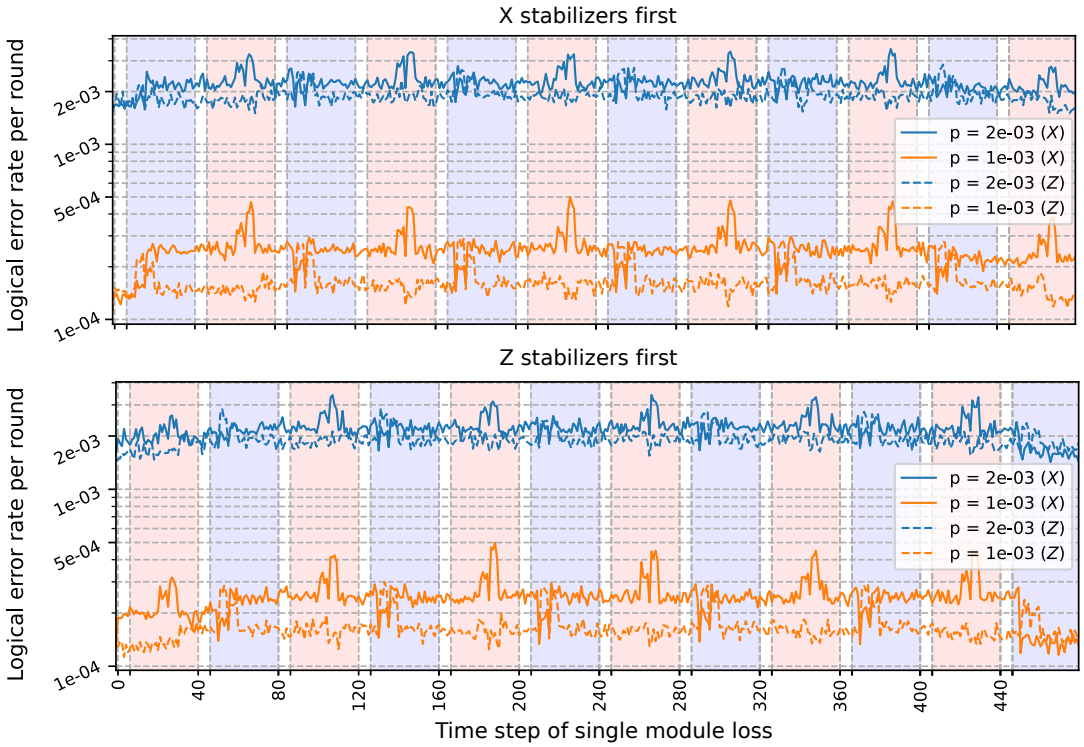

FIG. 4: Effect of losing data module 0 at different time steps. The top plot shows results for the X-then-Z stabilizer measurement order, while the bottom plot uses the reverse Z-then-X order. Shaded areas (blue for X, red for Z) indicate stabilizer measurement steps, during which some modules are merged to execute 2-qubit gates. White areas represent ancilla reset/measurement steps, during which all modules are separate. Solid curves represent the X logical error rate, and dashed curves represent the Z logical error rate. p in the legend denotes the physical 2-qubit gate error rate.

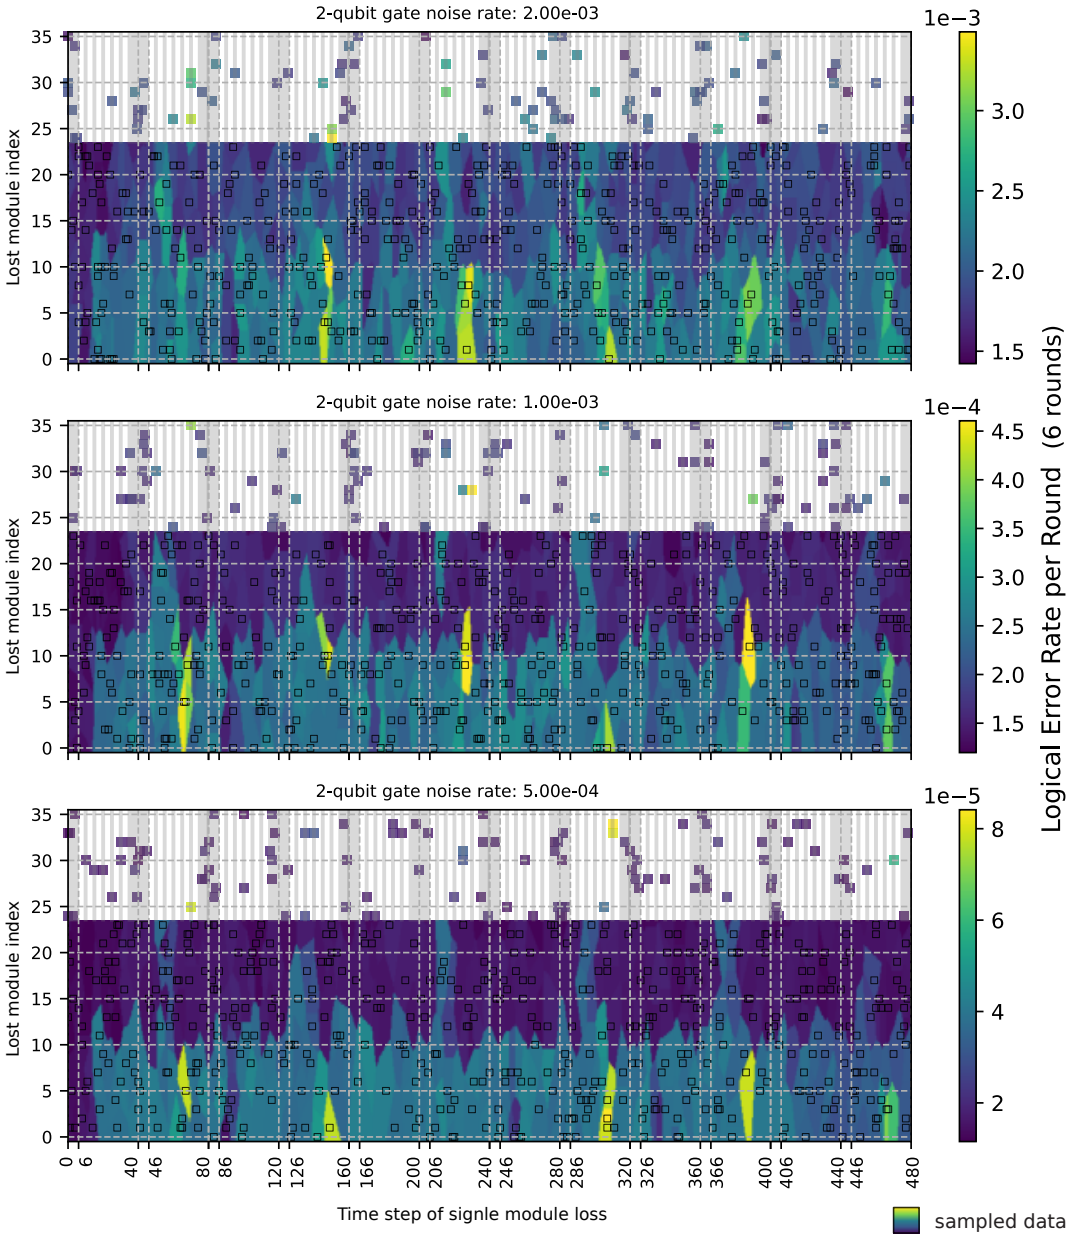

FIG. 5: Effect of a single random chain loss occurring at a random time step. Each of the three plots corresponds to a different 2-qubit gate error rate. The colored squares show the logical error rate for 600 randomly sampled chain loss events, defined by the time step (x-axis) and the lost module index (y-axis). During merged steps (unshaded regions), an ancilla module cannot be lost independently; it is lost jointly with the data module it is merged with. For this reason, no independent loss events are sampled for ancilla modules (indices 24-35) in the unshaded regions. To aid interpretation, the plot uses nearest-neighbor interpolation for unsampled data module (indices 0-23) loss events. Therefore, for the data modules, the color map represents true simulated data only at the locations of the colored squares.