

Simulation of thermal-relaxation noise for quantum error correction

Sean R. Garner^{1,2,*} Nathan M. Myers¹ Meng Wang^{1,3} Samuel Stein¹ Chenxu Liu^{1,†} and Ang Li^{1,2}

¹Physical and Computational Sciences, *Pacific Northwest National Laboratory*, Richland, Washington 99354, USA

²Department of Electrical and Computer Engineering, *University of Washington*, Seattle, Washington 98195, USA

³Department of Electrical and Computer Engineering, *The University of British Columbia*, Vancouver, BC V6T 1Z4, Canada



(Received 15 December 2025; accepted 9 April 2026; published 7 May 2026)

Stabilizer-based simulation of quantum error-correcting codes typically relies on the Pauli-twirling approximation (PTA) to render non-Clifford noise classically tractable, but PTA can distort the behavior of physically relevant channels such as thermal relaxation. Physically accurate noise simulation is needed to train decoders and understand the noise suppression capabilities of quantum error-correcting codes. In this work, we develop an exact stabilizer model without variance overhead for qubit thermal-relaxation noise when $T_2 \leq T_1$. We show that the combined amplitude damping and dephasing channel admits a fully positive Clifford and reset decomposition that enables efficiency in this regime. For $T_2 > T_1$, the resulting decomposition is negative, but the same technique serves to reduce the overhead versus independent channels. We further introduce an approximated error channel with reset for the $T_2 > T_1$ regime that removes the negativity of the decomposition while achieving higher channel fidelity to the true thermal relaxation than PTA and extend our construction to finite-temperature relaxation. We apply the exact combined model to investigate large surface codes and bivariate bicycle codes on superconducting platforms with realistic thermal-relaxation error. The differing logical performances across code states further indicate that noise-model-informed decoders will be essential for accurately capturing thermal-noise structure in future fault-tolerant architectures.

DOI: [10.1103/hgnj-v4j7](https://doi.org/10.1103/hgnj-v4j7)

I. INTRODUCTION

Careful suppression and managing of noise and decoherence are essential for reliable quantum computation, making quantum error correction (QEC) a central focus toward building large-scale fault-tolerant quantum computers [1–3]. Recently, substantial progress has been achieved in both the theoretical design and experimental implementation of QEC across multiple physical platforms, including superconducting circuits [4–7], trapped ions [8,9], and neutral atoms [10–13]. Advances in code constructions, such as surface codes [14–16], color codes [17–19], and quantum low-density parity check (QLDPC) codes [20–22], have enabled experimental demonstrations of multi-round error correction and logical qubit surpassing the breakeven point and showing below-threshold performance. A central challenge that persists, however, is the accurate characterization of logical-level behavior under realistic noise. Thermal noise is especially pervasive throughout all quantum computers, yet has thus far only been captured by approximation or exponential exact

simulation, as the relaxation channel is beyond the constraints of the Gottesman-Knill theorem [23].

Reliable evaluation of code performance, threshold estimation, and decoder benchmarking all depend on understanding how physical noise models propagate through the error-correction process. To better access candidate QEC codes and architectures operating in realistic experimental environments, we need efficient simulation tools that can incorporate more realistic noise models and scale to thousands of physical qubits.

Conventional full-state noise simulation approaches, such as state-vector (SV) and density-matrix (DM) methods, represent arbitrary quantum processes by explicitly evolving the quantum states. The universality allows them to capture any completely positive trace-preserving (CPTP) map acting on the system, but at the cost of exponential scaling with n number of qubits [$O(2^n)$ and $O(4^n)$ for SV and DM, respectively]. This makes them intractable for moderate-sized QEC codes. Previous efforts investigated the QEC code performance under realistic noise using DM simulations; however, they are limited to small code distances and rely on more simplifications and approximations, such as reusing the syndrome qubits in the simulation, which potentially violates the temporal order of syndrome measurements in experimental implementations [24,25]. These challenges make the stabilizer simulation framework necessary for simulating QEC code performance, at the cost of exactness [26].

The stabilizer formalism allows Clifford circuits to be simulated in $O(n^2)$ time and memory by evolving a tableau that tracks stabilizer transformation under Clifford

*Contact author: sean.garner@pnnl.gov

†Contact author: chenxu.liu@pnnl.gov

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

operations [27]. This efficiency makes it widely adopted in most QEC simulation tools, such as Stim [28], and enables exploration of mid- to large-scale QEC code patches in the near fault-tolerant quantum computing (FTQC) regime. However, the efficiency of stabilizer simulation relies on the restriction that the operations are within the Clifford group, with noise channels that are expressed as stochastic Pauli channels. In practice, non-Clifford physical error channels, such as coherent errors [29,30] and thermal errors [25,31], cannot be directly simulated in stabilizer simulators. Instead, the Pauli-twirling approximation (PTA) is commonly employed to replace a general CPTP map by its depolarized Pauli channels [32–34]. Although efficiently simulable, previous studies on small codes have shown that PTA can underestimate performance, producing logical error rates (LERs) that are several times larger than true values [24,35].

A more accurate yet still Clifford-compatible approach is the quasiprobabilistic decomposition (QPD) of a non-Clifford channel into a linear combination of Clifford channels [36]. By sampling from a rescaled non-negative probability distribution and making a weight adjustment, one can exactly reproduce the original channel statistics with only Clifford simulators. However, negativity in the probabilistic distribution results in an exponential sampling cost; i.e., the estimator variance will grow exponentially with the number of non-Clifford error channels [36,37]. To model qubits' thermal-relaxation errors in the QEC circuits, the cost will grow with circuit depth and number of qubits. Multiple works have formalized this cost for QPD and more general noise in terms of channel robustness [38,39]. Consequently, while QPD yields unbiased results, it can introduce exponential sampling cost to offset variance, hindering its implementation for efficiently simulating realistic error models for large-scale QEC code performances. Moreover, QPD often requires explicit stabilizer-state reset operations in some sampling branches of the error channel, which are not directly compatible with Pauli-frame tracking for stochastic Pauli error channels used in state-of-the-art QEC simulation [28]. Recent work on more sophisticated frame representations beyond Pauli frames is promising to reduce the per-sample cost [40]. However, a tool developed to permit more sophisticated frames will still suffer from the exponential variance underlying QPD.

In this paper, we focus on the qubit thermal-relaxation error channel, which is one of the most commonly seen non-Clifford error sources in experiments. We show that combining the amplitude damping (T_1) and dephasing (T_2) channels into a unified relaxation channel significantly reduces the QPD sampling overhead required for exact Clifford-based simulation. Specifically, we find that when $T_2 \leq T_1$, a regime typical for superconducting qubits and several solid-state qubit platforms [6,41], the thermal-relaxation channel admits a fully positive decomposition into Clifford channels. In the regime $T_2 > T_1$, we demonstrate that the resulting decomposition is still less costly to sample than treating amplitude and phase damping separately. We further introduce a reset-based approximation that can preserve complete positivity of the decomposition, while achieving substantially higher channel fidelity than the PTA model. We incorporate these constructions into a message passing interface (MPI)-accelerated and

graphics processing unit (GPU)-capable stabilizer simulator and evaluate the logical performance of both surface codes and bivariate bicycle (BB) codes under realistic thermal-relaxation noise, beyond the scope of previous works that only focus on realistic noise performance on surface code only or other small QEC codes [29,35,42–46]. Our simulations show that PTA can misestimate logical error rates by several factors, either overestimating or underestimating, depending on code distance and logical basis, whereas the composite and reset-based treatments accurately capture the directional bias of relaxation. We also observe that logical error rates depend on the encoded logical state, suggesting that bias-aware, noise-informed decoding strategies are needed for future fault-tolerant architectures.

Our paper is organized as follows. Section II reviews the theoretical foundations of thermal-relaxation error models, including their Kraus representations, the connection to relaxation dynamics induced by a thermal bath, the Pauli-twirling approximation, and the exact quasiprobabilistic decomposition of amplitude-damping processes. We then present our treatment of the thermal-relaxation channel, identify the parameter regimes in which it admits exact Clifford-compatible simulation, and introduce an improved approximation in regimes where this exactness does not hold. Section III applies the resulting error model to evaluate the performance of the surface code and the BB code on superconducting-qubit platforms. We provide further discussion in Sec. IV and conclude in Sec. V.

II. THEORETICAL MODELING OF THERMAL-RELAXATION ERRORS

Thermal relaxation is a fundamental and widely encountered noise mechanism in quantum computing. Because the exact thermal-relaxation channel is inherently non-Pauli and cannot be represented as a stochastic Pauli channel, approximations are required to incorporate it into the current stabilizer-based QEC simulation toolchain.

In this section, we review the main theoretical treatment of thermal-relaxation and introduce the model adopted in this work. Section II A provides a brief review of the thermal-relaxation channel. Section II B reviews the Pauli-twirling approximation and its application to thermal relaxation, highlighting its limitations. Section II C reviews prior study on QPD and its application to the amplitude-damping channel. In Sec. II D, we extend the QPD to general thermal-relaxation error channels, identifying the efficiency gain when combining the transverse and longitudinal relaxation processes and treating them jointly. We then provide an improved approximation of the thermal relaxation in the relevant parameter regime with better accuracy, which is well suited for stabilizer-based QEC simulations.

A. Thermal relaxation error channels

Quantum computing systems are inevitably coupled to their environment, which causes qubits to thermalize toward the bath environment. Thermal relaxation can lead to energy loss, causing qubits to relax from the excited state to the ground state and lose the coherence between the computa-

tional states. The relaxation effects are characterized by the qubit energy relaxation time T_1 and the decoherence time T_2 .

To model the effect of thermal noise on the quantum systems, the relaxation dynamics can be described by a master equation and integrated for a time period τ . For a single qubit interacting with a thermal bath, the master equation is

$$\partial_t \rho = \gamma(\langle n_b \rangle + 1) \mathcal{D}[|0\rangle\langle 1|](\rho) + \gamma \langle n_b \rangle \mathcal{D}[|1\rangle\langle 0|](\rho) + \frac{\gamma_\phi}{2} \mathcal{D}[\sigma_z](\rho), \quad (1)$$

where relaxation rate γ is determined by the qubit-bath coupling, $\langle n_b \rangle$ is the thermal occupation of the bath at temperature T_b , given by $\langle n_b \rangle = 1/(e^{\hbar\omega/k_b T_b} - 1)$, and γ_ϕ characterizes the extra dephasing effect on the qubit. The dissipator in Eq. (1) is

$$\mathcal{D}[\hat{A}](\rho) = -\frac{1}{2}(\hat{A}^\dagger \hat{A} \rho + \rho \hat{A}^\dagger \hat{A} - 2\hat{A} \rho \hat{A}^\dagger), \quad (2)$$

for any jump operator $\hat{A}$ and system density operator ρ [47,48].

When the qubit couples to a zero-temperature bath ($\langle n_b \rangle = 0$) for a time τ , the density operator evolves to

$$\rho(\tau) = \begin{pmatrix} 1 - \rho_{11}e^{-\tau/T_1} & \rho_{01}e^{-\tau/T_2} \\ \rho_{10}e^{-\tau/T_2} & \rho_{11}e^{-\tau/T_1} \end{pmatrix}, \quad (3)$$

where $\rho_{i,j} = \langle i|\rho(t=0)|j\rangle$, $T_1 = 1/\gamma$, and $T_2 = 1/(\gamma/2 + \gamma_\phi)$. As the dephasing between the states $|0\rangle$ and $|1\rangle$ comes from two sources, the energy relaxation and pure dephasing effects, the decoherence time satisfies

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}, \quad (4)$$

where $T_\phi = 1/\gamma_\phi$ is the pure dephasing time. From Eq. (4), the qubit dephasing time T_2 is bounded by $2T_1$.

The thermal-relaxation process can be equivalently described by Kraus operators [49]. When there is no pure dephasing, $T_2 = 2T_1$, the process reduces to the amplitude-damping channel as $\mathcal{E}_{\text{ad}}(\rho) = E_0 \rho E_0^\dagger + E_1 \rho E_1^\dagger$, with

$$E_{\text{ad},0} = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p_\gamma} \end{pmatrix}, \quad E_{\text{ad},1} = \begin{pmatrix} 0 & \sqrt{p_\gamma} \\ 0 & 0 \end{pmatrix}, \quad (5)$$

where the channel is characterized by a “relaxation probability” p_γ [49]. Comparing with the density operator evolution in Eq. (3), the relaxation probability is

$$p_\gamma(\tau) = 1 - e^{-\tau/T_1}. \quad (6)$$

In the presence of pure dephasing, i.e., when $T_2 < 2T_1$, the additional dephasing noise can be modeled as an additional stochastic Pauli-Z error with probability p_ϕ , represented in Kraus operators as

$$\mathcal{E}_{\text{pd}}(\rho) = (1 - p_\phi)\rho + p_\phi Z \rho Z. \quad (7)$$

The full thermal-relaxation process is

$$\mathcal{E}_{\text{pd}}[\mathcal{E}_{\text{ad}}(\rho)] \stackrel{\text{def}}{=} (\mathcal{E}_{\text{pd}} \circ \mathcal{E}_{\text{ad}})(\rho). \quad (8)$$

Matching Eq. (3) yields the additional Pauli-Z error probability,

$$p_\phi = \frac{1}{2} \left[1 - e^{-\tau \left(\frac{1}{T_2} - \frac{1}{2T_1} \right)} \right] = \frac{1}{2} (1 - e^{-\tau/T_\phi}). \quad (9)$$

At finite bath temperature, the bath thermal occupation $\langle n_b \rangle \neq 0$, with the value of $\langle n_b \rangle$ determined by the bath temperature T_b . Similarly, we can integrate the full master equation (1) to solve the dynamics of the qubit. Specifically, when $\tau \rightarrow \infty$, the population in the excited state is

$$p_1 = \frac{\langle n_b \rangle}{1 + 2\langle n_b \rangle} = \frac{1}{1 + e^{\hbar\omega_0/k_b T_b}}. \quad (10)$$

The qubit’s total decay rate to the equilibrium state defines the qubit’s T_1 time as $1/T_1 = \gamma(2\langle n_b \rangle + 1)$, while the dephasing rate defines the decoherence time T_2 as $1/T_2 = \gamma(\langle n_b \rangle + 1/2) + \gamma_\phi$.

When there is no pure dephasing ($\gamma_\phi = 0$), the finite-temperature thermal-relaxation process can be described by a generalized amplitude-damping channel as [49]

$$\mathcal{E}_{\text{gad}}(\rho) = \sum_{i=1}^4 E_{\text{gad},i} \rho E_{\text{gad},i}^\dagger, \quad (11)$$

with the Kraus operators

$$E_{\text{gad},0,1} = \sqrt{1-p_1} E_{\text{ad},0,1}, \quad (12)$$

$$E_{\text{gad},2} = \sqrt{p_1} \begin{pmatrix} \sqrt{1-p_\gamma} & 0 \\ 0 & 1 \end{pmatrix}, \quad (13)$$

$$E_{\text{gad},3} = \sqrt{p_1} \begin{pmatrix} 0 & 0 \\ \sqrt{p_\gamma} & 0 \end{pmatrix}, \quad (14)$$

where p_γ defined in $E_{\text{ad},0,1}$ becomes

$$p_\gamma = p'_\gamma = 1 - e^{-\gamma\tau(2\langle n_b \rangle + 1)} = 1 - e^{-\tau/T_1}, \quad (15)$$

p_1 is the population of $|1\rangle$ state when $\tau \rightarrow \infty$ as given in Eq. (10).

When there is additional dephasing ($\gamma_\phi \neq 0$), the thermal-relaxation process can be modeled as $\mathcal{E}_{\text{pd}} \circ \mathcal{E}_{\text{gad}}$, where the Pauli-Z error probability p_ϕ is given in Eq. (9).

B. Pauli twirling and Pauli-twirling approximation to quantum channels

Twirling is a technique used to symmetrize the effect of a noise channel by randomly applying a unitary operation from a chosen symmetry group $\mathcal{G}$ and its inverse. This process effectively averages the noise channel over the group $\mathcal{G}$, which is widely used for both characterizing and mitigating noise in quantum systems [25,50]. Depending on the chosen $\mathcal{G}$, various twirling methods arise, such as unitary twirling [51], Clifford twirling [52], and Pauli twirling [25]. Among these, Pauli twirling is one of the most popular choices due to its relative simplicity in implementation and analysis.

Pauli twirling involves averaging the channel over the n -qubit Pauli group $\mathcal{P}_n$, which transforms the original arbitrary noise channel into stochastic Pauli channels. Approximating the original channel by its Pauli-twirled version is known as the Pauli-twirling approximation. For the original channel $\mathcal{E}$, the channel after Pauli twirling is

$$\mathcal{E}_{\text{pta}}(\rho) = \sum_{\sigma_i \in \mathcal{P}_n} \sigma_i \mathcal{E}(\sigma_i \rho \sigma_i) \sigma_i, \quad (16)$$

where $\mathcal{P}_n$ is the n -qubit Pauli group and σ_i are the Pauli operators. We use the fact that the Pauli operators are unitary and




FIG. 1. The effect of the thermal-relaxation error channel and the channel after PTA. (a) The Bloch sphere after applying the thermal-relaxation channel and the channel after PTA on a single qubit. We notice that the distortion on the Pauli- X and Y directions is nicely captured by PTA channel, while the directional relaxation along the Pauli- Z axis is not. The gray-shaded sphere shows the unit Bloch sphere. (b) The trace distance difference (ΔD) between the two channels. We focus on the states whose Bloch vectors lie along the red great circle in panel (a). Parameters if not specified: $\tau/T_1 = 1.0$, $T_2/T_1 = 1.5$, thermal temperature $T_b = 0$.

Hermitian. Notice that a quantum channel can be equivalently described by its χ matrix as

$$\mathcal{E}(\rho) = \sum_{i,j} \chi_{i,j} \sigma_i \rho \sigma_j. \quad (17)$$

It is shown that the χ matrix of the PTA channel is diagonal,

$$\mathcal{E}_{\text{pt}}(\rho) = \sum_{i,i} \chi_{i,i} \sigma_i \rho \sigma_i, \quad (18)$$

i.e., Pauli twirling removes all the off-diagonal elements in the original χ matrix and retains only the diagonal coefficients.

The PTA replaces an arbitrary quantum channel with an incoherent mixture of Pauli error channels. It is widely used in analyzing the effects of complex noise models in quantum algorithms and quantum error correction, because the resulting channel is efficiently simulable using stabilizer simulators for large-scale physical qubits in QEC codes [24,25,35].

Although the PTA provides a way to efficiently investigate the effect of complex noise, the twirled channel may induce significant bias to the true noise channel. Considering the thermal-relaxation error channel in Sec. II A as an example, when the thermal temperature is zero, the channel after the PTA is

$$\mathcal{E}_{\text{th,pta}}(\rho) = \left(1 - \sum_i p_i\right) \rho + \sum_i p_i \sigma_i \rho \sigma_i, \quad (19)$$

where the sum runs over $\{x, y, z\}$ and the error probabilities are

$$p_x = p_y = \frac{p_\gamma}{4}, \quad p_z = \frac{1}{2} - \frac{p_\gamma}{4} - \frac{(1 - 2p_\phi)\sqrt{1 - p_\gamma}}{2}, \quad (20)$$

where p_γ and p_ϕ are defined in Eqs. (6) and (9), respectively [25]. Figure 1 compares the action of the true thermal-relaxation error channel and its PTA counterpart on the Bloch

sphere. The contraction along the X - Y plane is well captured, but the PTA fails to reproduce the directional relaxation along the Z axis.

As the PTA does not accurately capture the full structure of thermal relaxation, Fig. 1(a) illustrates that it can either overestimate or underestimate the true error depending on the input state. Consider two extreme input states, $|0\rangle$ and $|1\rangle$. At zero temperature, thermal relaxation leaves state $|0\rangle$ unchanged, while the PTA introduces a nonzero error probability, leading to overestimating the error. Conversely, as thermal relaxation drives state $|1\rangle$ to state $|0\rangle$, the PTA channel underestimates the error for state $|1\rangle$. To quantify this effect, we compute the trace distance between the noisy states produced by the PTA channel and the noiseless state, and then compare it against the corresponding trace distance under the true thermal-relaxation channel:

$$\Delta D \equiv D(\rho, \mathcal{E}_{\text{th}}^{\text{pta}}[\rho]) - D(\rho, \mathcal{E}_{\text{th}}[\rho]), \quad (21)$$

where D denotes the trace distance operation. We evaluate this trace distance difference for the pure states $|\psi(\theta)\rangle = \cos(\theta/2)|0\rangle + \sin(\theta/2)|1\rangle$, whose Bloch vectors lie on the great circle with $y = 0$ [red curve in Fig. 1(a)]. Figure 1(b) shows the resulting trace distance differences. When $\theta < \pi/3$ (closer to $|0\rangle$), the PTA channel overestimates the error, whereas for $\theta > \pi/3$ (closer to $|1\rangle$), the PTA channel underestimates the error. These observations motivate the need for more accurate representations of error channels than those provided by the PTA.

C. Quasiprobabilistic decomposition to Clifford channels

Reference [36] proposed a method to decompose an arbitrary quantum channel into a composition of Clifford and reset to stabilizer state operations. This approach begins by solving the Pauli-transfer matrix [53,54] of the quantum channel, and then solving for a decomposition into Clifford and reset operations element by element. As a concrete example, the amplitude-damping channel in Eq. (5) can be decomposed as

$$\mathcal{E}_{\text{ad}} = q_+^{(\text{ad})} \mathcal{I} + q_-^{(\text{ad})} \mathcal{Z} + p_\gamma \mathcal{R}_{|0\rangle}, \quad (22)$$

where $\mathcal{I}$ and $\mathcal{Z}$ are identity and Pauli- Z channels, respectively, and $\mathcal{R}_{|0\rangle}$ is the channel corresponding to reset to state $|0\rangle$. The quasiprobabilities are

$$q_{\pm}^{(\text{ad})} = \frac{1 - p_\gamma \pm \sqrt{1 - p_\gamma}}{2}. \quad (23)$$

This decomposition enables Monte Carlo simulation of T_1 relaxation in QEC codes within the stabilizer-based framework.

From Gottesman-Knill theorem [27], it is unlikely that any arbitrary quantum channel can be classically simulated in an efficient manner. The cost of the above method is that probabilistic decomposition does not necessarily yield a completely positive distribution. Equation (22) shows the coefficient of the Pauli- Z channel to be always negative, meaning the decomposition yields a quasiprobabilistic (QP) distribution instead.

Although the resulting QP distribution is not a true probability distribution, it can still be sampled to estimate statistical expectations of any target observable, at the expense of increased sampling variance. For a random variable x , the




FIG. 2. (a) Sampling cost to offset estimator variance $\propto \Gamma^{2n_c}$ if all gates of a code have a uniform time cost, i.e., Γ is the same for every type of gate. Vertical lines denote approximately the current time cost of measurement and two-qubit gates in relation to T_1 coherence on IBM devices. In reality, the time cost will fall in between measurement (M) and two-qubit gates, depending on the relative amount of those gates for each code. Overhead assumes d rounds of syndrome extraction with a distance d code. (b) Approximate quasiprobability sampling cost per-memory experiment. As opposed to panel (a), Γ here is calculated for each gate time, for the number of each gate type in each code, representing the actual simulation. Measurement is approximated to take $T_1/100$, while two-qubit gates take $T_1/1000$. The uncombined quasiprobability variance overhead is independent of T_2 as the probability decomposition for T_2 is always positive.

expectation value of a function $f(x)$ on a QP distribution $q(x)$ is

$$\langle f \rangle_q = \sum_x q(x) f(x). \quad (24)$$

To sample from this QP distribution, a true probability distribution is defined as

$$p(x) = |q(x)|/\Gamma, \quad \Gamma = \sum_x |q(x)|, \quad (25)$$

where Γ measures the negativity of the QP distribution. The expectation value can then be rewritten as

$$\langle f \rangle_q = \sum_x \Gamma \operatorname{sgn}[q(x)] p(x) f(x), \quad (26)$$

which can be sampled using Monte Carlo methods, resulting in an unbiased estimator for $\langle f \rangle$.

However, when sampling the QP distribution using Monte Carlo methods, the variance of the estimator will grow with the negativity of the QP distribution [36,55,56], which requires more Monte Carlo shots to achieve a target precision. For the amplitude-damping channel, the negativity is controlled by the ratio τ/T_1 . When there are multiple quantum channels in a quantum circuit that need to be decomposed independently, the total negativity causes the variance to multiply across each quantum channel,

$$\Gamma_{\text{total}} = \prod_{i=1}^{n_c} \Gamma_i, \quad (27)$$

where Γ_i represents the negativity of each QP distribution and n_c is the total number of decomposed channels. Consequently, the sampling cost must grow exponentially with both the number of qubits and the number of amplitude-damping error channels (the depth of the quantum circuit), to offset the variance introduced by the unbiased estimator. Figure 2

shows the total variance of the uncombined QP distributions when simulating different QEC codes and counting for all the amplitude-damping error channels on each of the operations. In Fig. 2(a), we observe that as the gate time increases (larger error rates in each amplitude-damping channel), the sampling overhead grows rapidly with the error strength. Furthermore, in Fig. 2(b), we show the total variance as a function of the number of total qubits in each code experiment. It is clear that the sampling overhead can make numerical simulation of large-distance codes impractical when trying to determine logical error performance.

D. QPD on the thermal relaxation channel

With the understanding that non-Clifford channels can be decomposed into stabilizer channels using QPD, and the cost to do so, it is natural to seek ways to reduce the negativity by minimizing Γ from the QP distribution or the number of channels n_c .

As thermal relaxation is among the most common error mechanisms in experimental platforms, our goal is to find a low-overhead Clifford-based representation for thermal-relaxation error channels. Observing the decomposition of the amplitude-damping error channel in Eq. (22), the negativity is contributed from the Pauli-Z operation branch. However, thermal relaxation typically includes additional dephasing $\gamma_\phi \neq 0$, which itself can naturally be modeled as a Pauli-Z error channel (see Sec. II A). This suggests that, by creating an affine combination of the amplitude damping and dephasing error channels, the overall negativity of the thermal-relaxation channel can be significantly reduced, yielding an efficient simulation method.

1. Zero-temperature thermal relaxation

We first consider the QPD of the thermal-relaxation error channels at temperature zero, i.e., $\mathcal{E}_{\text{th}} = \mathcal{E}_{\text{pd}} \circ \mathcal{E}_{\text{ad}}$. The QPD

of the thermal channel can be derived from Eqs. (22) and (7) as

$$\mathcal{E}_{\text{th},0} = q_+^{(\text{th},0)} \mathcal{I} + q_-^{(\text{th},0)} \mathcal{Z} + p_\gamma \mathcal{R}_{|0\rangle}, \quad (28)$$

where the quasiprobabilities are

$$\begin{aligned} q_\pm^{(\text{th},0)} &= \frac{1 - p_\gamma \pm (1 - 2p_\phi)\sqrt{1 - p_\gamma}}{2} \\ &= [e^{-\gamma\tau} \pm e^{-(\gamma/2 + \gamma_\phi)\tau}] / 2 \\ &= [e^{-\tau/T_1} \pm e^{-\tau/T_2}] / 2. \end{aligned} \quad (29)$$

From Eq. (29), the decomposed identity channel from the thermal channel always has positive probabilities, while the Pauli-Z channel can have positive or negative quasiprobability, depending on the relative strength of the pure dephasing and energy relaxation process. Specifically, when $T_2 \leq T_1$, $q_- \geq 0$, implying that the combined thermal-relaxation channel admits a fully positive decomposition in this regime, removing all QPD sampling overhead. Combining the amplitude damping with the dephasing, therefore, enables an exponential advantage over decomposing them separately.

Importantly, this parameter regime is not only of theoretical interests, however, it is typical of many current quantum platforms, including superconducting and semiconductor spin qubits. In these systems, the environmental dephasing typically limits the coherence of the qubits, such that the T_2 time is comparable to T_1 [57,58]. Thus, many practically relevant quantum devices naturally fall into this “sweet spot” where the decomposition results in a fully positive distribution, costing no extra sampling overhead. Specifically, in the limiting case $T_1 = T_2$, the thermal-relaxation channel further simplifies to a pure stochastic reset process.

When $T_1 < T_2$, the decomposition is no longer completely positive, but the affine combination can still be simulated exactly by sampling over the quasiprobabilities. In this case, combining phase decoherence with the amplitude damping helps to lower the overhead, as seen in Fig. 3. Once T_2 coherence is at the relaxation limit, i.e., there is effectively no pure phase decoherence, the composite method converges to the standalone relaxation stabilizer decomposition found in Ref. [36]. This is consistent with the findings in Ref. [55], where combining Pauli-X error channels with coherent errors can reduce the sampling overhead.

Beyond exactly sampling from the QP distribution, in the $T_2 \gtrsim T_1$ regime, one may approximate the thermal-relaxation channels by the reset error channel, trading accuracy for efficiency by maintaining the complete positivity of the distribution. Specifically, we consider the quantum channel at absolute zero

$$\mathcal{E}_{\text{reset}} = q_+^{(\text{th},0)} \mathcal{I} + (1 - q_+^{(\text{th},0)}) \mathcal{R}_{|0\rangle}, \quad (30)$$

which matches the exact QPD thermal-relaxation channel in Eq. (28) when $T_2 = T_1$. When $T_2 > T_1$, this reset channel neglects the negative term $q_-^{(\text{th},0)}$, and the reset probability is rescaled.

To better quantify the accuracy of this channel, we compute the process fidelity to the exact thermal-relaxation error channel by

$$F_{\mathcal{E}} = F((\mathcal{E}_{\text{th},0} \otimes \mathcal{I})[\rho_{\text{me}}], (\mathcal{E} \otimes \mathcal{I})[\rho_{\text{me}}]), \quad (31)$$




FIG. 3. Sampling overhead for the quasiprobability relaxation channel on a single qubit, decomposed using Eq. (28). The overhead is measured by the negativity of the quasiprobabilistic distribution, which is captured by Γ . When $T_2/T_1 \leq 1$, the QP distribution is completely positive, while $T_2/T_1 = 2$, there is no additional dephasing, the overhead is the same as the QP distribution of an amplitude-damping channel.

where ρ_{me} is a maximally entangled state defined on two copies of the Hilbert space that the quantum channels $\mathcal{E}$ and $\mathcal{E}_{\text{th},0}$ act on [59]. In our case, the maximally entangled state is a Greenberger–Horne–Zeilinger (GHZ) state $|00\rangle + |11\rangle$. We compare it with the Pauli-twirled thermal-relaxation channel. Figure 4 shows the channel fidelity gain of $\mathcal{E}_{\text{reset}}$ with respect to the Pauli-twirled channel $\mathcal{E}_{\text{th},\text{pta}}$ in Eq. (19). We focus on the regime where no exact positive QPD exists, $2 \gtrsim T_2/T_1 > 1$. In the experimentally relevant regimes, the reset channel $\mathcal{E}_{\text{reset}}$ shows constantly better channel fidelity to the exact thermal-relaxation channel. We notice that the difference between the two error channels decreases as τ/T_1 decreases.




FIG. 4. We compare the channel fidelity of the reset error and the Pauli-twirled thermal channel to the exact thermal channel at zero temperature. We plot the difference of the channel fidelity $\Delta F = F_{\text{reset}} - F_{\text{pta}}$. In this regime, the reset channel can always give a better channel fidelity to the exact thermal-relaxation error channel compared to the Pauli-twirled error channel.

This is because as we reduce the ratio τ/T_1 , the error probability decreases, which makes both the PTA and the reset channel closer to the identity channel, thus reducing the distance between the two channels. However, the reset channel $\mathcal{E}_{\text{reset}}$ always has better channel fidelity to the exact thermal-relaxation channel. As the thermal-relaxation error becomes increasingly pronounced, the advantage of the reset channel becomes more significant. We attribute this to the fact that the reset error channel can efficiently capture the directional bias of the thermal-relaxation error channel in the low-dephasing regime while the PTA channel fails (see Fig. 1).

2. Finite-temperature thermal relaxation

We now extend the decomposition to the thermal-relaxation channel at nonzero temperature. We start from the general amplitude-damping channels, whose Kraus operators are given in Eqs. (12)–(14). As the Kraus operators $E_{\text{gad},0,1}$ only differ from the amplitude-damping Kraus operators by only a constant scaling factor, this partial channel can be decomposed similarly as in Eq. (22). The branches that describe the thermal excitation $E_{\text{gad},2,3}$ can be viewed as an amplitude-damping channel conjugated by Pauli- X , i.e.,

$$\mathcal{E}_{\text{gad}} = (1 - p_1)\mathcal{E}_{\text{ad}}[p_\gamma] + p_1\mathcal{X} \circ \mathcal{E}_{\text{ad}}[p'_\gamma] \circ \mathcal{X}, \quad (32)$$

where $\mathcal{X}$ is the quantum channel for the Pauli- X , and p_γ and p'_γ are given in Eq. (15). Thus, the QPD of the general amplitude-damping channel becomes

$$\mathcal{E}_{\text{gad}} = q_+^{(\text{gad})}\mathcal{I} + q_-^{(\text{gad})}\mathcal{Z} + (1 - p_1)p_\gamma\mathcal{R}_{|0\rangle} + p_1p_\gamma\mathcal{R}_{|1\rangle}, \quad (33)$$

where the quasiprobabilities are

$$q_\pm^{(\text{gad})} = (1 - p_1)q_\pm^{(\text{ad})}[p_\gamma] + p_1q_\pm^{(\text{ad})}[p'_\gamma] = q_\pm^{(\text{ad})}. \quad (34)$$

Including the dephasing effect into consideration, the full finite-temperature thermal-relaxation channel becomes

$$\mathcal{E}_{\text{th}} = q_+^{(\text{th})}\mathcal{I} + q_-^{(\text{th})}\mathcal{Z} + (1 - p_1)p_\gamma\mathcal{R}_{|0\rangle} + p_1p_\gamma\mathcal{R}_{|1\rangle}, \quad (35)$$

where the quasiprobabilities are

$$q_\pm^{(\text{th})} = q_\pm^{(\text{th},0)}[p_\gamma], \quad (36)$$

$q_\pm^{(\text{th},0)}$ is given in Eq. (29), and the thermal relaxation and excitation probability p_γ are defined in Eq. (15). We notice that compared to the zero-temperature thermal-relaxation channel, the negativity of the distribution remains the same as the zero-temperature thermal channel with the same p_γ , indicating that higher temperatures do not increase or decrease the sampling cost. In the high-temperature limit where $k_bT_b \gg \hbar\omega$, $\langle n_b \rangle \gg 1$, making the probability of the excited state $p_1 \sim 1/2$. In this regime, the thermal relaxation tends to thermalize toward the maximally mixed state $(|0\rangle\langle 0| + |1\rangle\langle 1|)/2$ and the directional feature is lost. The thermal-relaxation channel can be exactly represented by the PTA. However, for experimentally relevant thermal bath temperature $k_bT_b < \hbar\omega$, the reset error channel can still be a good approximation to the thermal-relaxation error channel,

$$\mathcal{E}_{\text{reset}} = q_+^{(\text{th})}\mathcal{I} + (1 - q_+^{(\text{th})})(1 - p_1)\mathcal{R}_{|0\rangle} + (1 - q_+^{(\text{th})})p_1\mathcal{R}_{|1\rangle}. \quad (37)$$




FIG. 5. The channel fidelity gain of the reset-based Clifford channel to the PTA channel to approximate a finite-temperature thermal-relaxation channel. We set $T_2 = 1.5T_1$. We sweep the excited state population up to $p_1 = 0.1$. To give a more intuitive understanding of experimental relevance, we compute the corresponding thermal temperature assuming the qubit frequency is 5 GHz, which is labeled on the right.

In Fig. 5, we plot the channel-fidelity improvement of the reset-based model relative to the PTA channel with finite temperature. We focus on the regime $p_1 \leq 0.1$, corresponding to $k_bT_b/\hbar\omega \sim 0.45$. For a 5 GHz superconducting qubit, this corresponds to a bath temperature of approximately $T_b \sim 109$ mK. We take $T_2/T_1 = 1.5$ as an example. Across this regime, the reset model consistently achieves higher channel fidelity to the finite-temperature thermal-relaxation channel than the PTA channel. As the thermal temperature increases, however, the advantage of the reset model gradually diminishes, consistent with the analysis above. We provide further analysis by directly simulating a small Steane code in Appendix B, connecting the improved fidelity of the reset-only model to an improved approximation of logical error rate under the actual noise channel.

III. NUMERICAL ANALYSIS OF QUANTUM MEMORIES UNDER THERMAL ERROR

Superconducting qubits are both well established and a promising avenue for scaled quantum computers [60]. Practical developments in multiple QEC directions on superconducting qubits have been made, including planar codes and quantum low-density parity check codes [6,7,61]. The rotated surface code [14] and bicycle-bivariate QLDPC [21] are two particularly promising code variants toward fault tolerance. Rotated surface code offers an efficient planar layout with well-understood operators, while BB code offers a very high encoding rate with minimal layers of connectivity to support the sparsity necessary for QLDPC. Understanding the code performance under realistic noises, especially thermal relaxation, can greatly help to optimize the code design and support the experimental realization.

Prior studies have investigated small to medium surface code performance under thermal relaxation with PTA [34,35]. For distance-3, comparisons with the exact thermal relaxation using density-matrix simulation were also performed [24,35,42]. PTA in these simulations was found to

TABLE I. Median coherence times (T_1 and T_2) of IBM Heron QPUs [62]. Superconducting qubits are one example where T_1 tends to exceed T_2 , although coherence times can vary significantly on current devices.

Device	T_1 (μ s)	T_2 (μ s)	T_1/T_2
Boston	275.27	338.82	0.81
Fez	142.41	98.43	1.45
Kingston	261.36	131.93	1.98
Marrakesh	185.77	104.16	1.78
Pittsburgh	300	324.17	0.93
Torino	183.67	131.48	1.40
Average	242.75	188.165	1.29

have a 2–10 $\times$ overestimation of error depending on the code and error parameters. Even distance-3 codes are difficult to fit within the density-matrix simulation. For previous studies, syndrome qubits were reused to “truncate” the code, which does not preserve the time order that physical surface codes are measured in. Surface codes larger than distance-3 have too many data qubits for exponential methods with or without truncation to be practical.

On the side of QLDPC codes, the noise performance of the BB code is thus far under investigated. BB codes with the best encoding rate are out of reach for density-matrix simulation. To our knowledge, the BB code logical error rate simulation has only been performed using uniform depolarizing errors to fit existing stabilizer simulators. BB codes beyond [18, 4, 4] are targeted for fault tolerance and high encoding rates, but are far too large for simulating thermal noise with exponential methods. By using the decomposition model discussed in Sec. IID to simulate T_1 and T_2 errors with stabilizer sampling, we are able to explore targeted QLDPC codes and preserve syndrome qubit layouts with both accuracy and efficiency.

It happens that transmon qubit technology tends to fall within $T_1 \approx T_2$, and thus can be simulated under composite stabilizer decomposition with good efficiency and accuracy. Table I provides coherence times at the time of writing for the current generation Heron processors, which vary with calibrations and individual operations. Current transmon-based hardware has $T_2 \sim T_1$ with fluctuations. To perform very precise qubit-by-qubit noise simulation, the combined thermal noise channel can be applied to most qubits where $T_1 \geq T_2$. Outlier qubits can then be well approximated using the reset channel in Sec. IID or simulated exactly by quasiprobability sampling the reduced-cost combined channel. The reduced-cost channel results from some pure phase decoherence reducing the negativity in Eq. (35). In the reduced-cost case, the total simulation overhead is the product of the respective reduced negativity (reduced Γ) for each qubit where $T_1 < T_2$. In our numerical investigation of the surface and BB codes, we assume $T_1 = T_2$ for simplicity. So long as $T_1 \geq T_2$, the combined decomposition method has both exactness and good efficiency.

In the rest of this section, we adopt the thermal-relaxation error model developed in Sec. IID and integrate it into our highly efficient Clifford simulator. Section III A introduces the simulator and experimental settings. The surface code and BB

code logical performances are reported in Secs. III B and III C, respectively.

A. Clifford simulator

To perform these simulations, we use a stabilizer simulator [63] with MPI and GPU support for many-shot simulation and large-code speedup. Like other state-of-the-art simulators, including Stim [28], our simulator is derived from the stabilizer and destabilizer tableaux introduced in CHP: CNOT-Hadamard-phase simulator [26]. Thermal relaxation and decoherence errors are modeled as probabilistic gates injected into the circuit during significant timesteps and idles. Every error gate performs its own Monte Carlo roll via random number generation to choose between potential Clifford or reset operations based on the stabilizer decomposition probability distribution. For $T_1 \geq T_2$, the thermal relaxation can be exactly decomposed into a probability distribution, inducing no sampling overhead. When $T_1 < T_2$, the negativity in the QP distribution can be handled in two ways: (1) the QP distribution can be sampled exactly with a smaller exponential number of shots by renormalizing the QP distribution to the corresponding probability distribution shown as in Eq. (26), or (2) the thermal-relaxation channel can be approximated by the reset channel introduced in Eq. (30), and sample it using Monte Carlo methods when the sampling cost is a concern.

Our framework provides a few standout differences for QEC simulations versus Stim [28]: (1) Stim is optimized for distributing single-thread CPU performance; however, large stabilizer tableaux can be greatly sped up with GPU parallelization, and (2) the Pauli frame tracking that underlies Stim’s fast sampling in noisy simulations cannot admit probabilistic reset operations, limiting it to Pauli-twirled approximations only. For the numerical simulations here, we distribute shots with MPI across CPU threads to probe lower error rates. Our CPU simulator provides better performance versus GPU in the tens to low hundreds of qubit range that single logical-qubit memories will use [63].

B. Rotated surface code

Rotated surface code fits the design constraints of local connectivity for planar superconducting hardware architecture while saturating the Bravyi-Poulin-Terhal bound in two dimensions [64]. It uses fewer resources than the planar surface code, and supports universal logical operations via state injection and lattice surgery [65]. To understand how a rotated surface code would perform as memory under exact thermal relaxation, we apply the thermal-relaxation error channels before measurement operations. This is because the measurements are typically much slower on superconducting devices than one and two-qubit gates, and it is the most heavily affected by the thermal-relaxation errors. We specifically focus on thermal channels under PTA and the exact channel with the error model discussed in Sec. II A. This application of noise focuses on the incoherent effects of relaxation and decoherence, complementing recent work on better modeling for coherent gate noise [55].

In Fig. 6, we investigate rotated surface code memories, with particular interest in larger distance codes to understand




FIG. 6. Rotated surface code $|0\rangle_L$ [panel (a)] and $|+\rangle_L$ [panel (b)] memory experiments (rounds = distance) decoded using minimum weight perfect matching, 10^7 shots per data point. We only track the thermal-relaxation errors during the syndrome qubit measurement, whose time (τ) is scaled by the device T_1 . We calculate the LER (per experiment) of the surface code with the exact thermal-relaxation error channel (solid lines) and the channel under PTA (dashed lines). Error bars show a 95% confidence interval for the number of logical errors given 10 million shots.

how performance scales under our exact model. We simulate the logical performance of surface code with distance $d = 3, 5, 7$, and 11 . In each memory experiment, we run d rounds of syndrome checks and use the minimum-weight-perfect-matching (MWPM) decoder to decode the syndromes and solve for the logical error rates. The LERs under thermal relaxation when the code block is in logical $|0\rangle_L$ [Fig. 6(a)] and logical $|+\rangle_L$ state [Fig. 6(b)] are computed. Since measurement dominates the idling time in superconducting systems, thermal errors are injected on all qubits while they await measurement extraction on syndrome qubits. We observe that PTA channels often predict the code performance reasonably with a constant offset, consistent with the previous study [24]. However, with access to low error rate and large-distance regimes, we find scaling is not perfect. When the logical state is $|0\rangle_L$, the PTA channel underestimates the LER, while in the state $|+\rangle_L$, the PTA overestimates.

By undergoing a nonuniform channel like thermal relaxation, codes will perform differently based on the logical state stored. On hardware, classical Pauli frame tracking and other runtime techniques will be used to minimize unnecessary physical Pauli operations and make the quantum computation as effective as possible. The rotated surface code experiments shown in Fig. 6 act differently with the same damping models, as the data qubits initialized in Z and X can accumulate errors for many rounds of syndrome measurement without being reset. In low-distance Z-memory experiments, where the overall error probability is low, random bit flips on the data qubit are more easily decoded since our MWPM decoder cannot account for the state precession toward $|0\rangle$ in its priors. With

the data qubits initialized in a $|+\rangle$ state for X memory, decoder performance is more consistent with error rate and distance. In this case, noise bias from damping is only one component of undetected logical errors in data qubits. Probabilistic phase flips that arise from pure phase decoherence can be accounted for in MWPM priors such that the LER from PTA and stabilizer decomposition channels match.

In Fig. 7, we plot the average excitation of the surface code block when the code block is in logical states $|0\rangle_L$, $|1\rangle_L$, and $|+\rangle_L$. We noticed that the excited state population of the logical $|0\rangle_L$ and $|+\rangle_L$ states differs only slightly, indicating that their LER differences are driven primarily by the MWPM decoder's response to the noise bias, rather than by the small differences in excited state population. A distance d rotated surface code requires a minimum of d qubits to support a logical operator, meaning a logical $|1\rangle_L$ state memory can be constructed such that the typical excited state population is very close to the logical $|0\rangle_L$ state. In testing LER curves for both error models in the $|1\rangle_L$ state are nearly identical to the curves in the $|0\rangle_L$ state (see Appendix A). The difference in scaling between memory states in Fig. 6 suggests that noise bias-aware decoding will be necessary in codes that can be solved with generalizations of MWPM.

C. Bivariate bicycle code

Bivariate bicycle code is a high encoding rate QLDPC scheme with minimal overlapping connection layers, helpful for superconducting architectures. BB code is not yet known to support universal quantum operations on its own;




FIG. 7. Excited state population of the rotated surface code block when the code is in logical $|0\rangle_L$, $|1\rangle_L$, or $|+\rangle_L$ state. We include both the data and syndrome qubits after the logical state preparation circuit for each logical state. By not postselecting any syndrome check outcomes, all the possible realizations of the same logical state are effectively being averaged. The generation circuits are shown in Appendix C. Each point uses 10 million samples.

however, as a memory, it is capable of suppressing logical errors very effectively with small improvements in physical error rate, creating an efficient fault-tolerant universal architecture via lattice surgery with other codes [20,66]. Previous studies developing and introducing these codes have primarily used uniform depolarizing errors [21] for simplicity and simulation performance.

Sparsely connected codes like BB code cannot be decoded with the same minimum weight matching graph method commonly employed for surface code and other locally connected codes. Instead, the state-of-the-art decoding algorithm for BB codes is belief propagation (BP), which propagates informed messages along the edges of the parity check Tanner graph [67]. BP decoding is often followed by ordered statistics decoding (OSD) refinement, which refines the BP output by flipping the least reliable bits and checking candidate error patterns at a combinatorial time cost [68]. In our study, we use BP-OSD with up to 1000 BP iterations and 10 OSD iterations as a common baseline decoder to compare the composite thermal-relaxation error model against the PTA.

In Fig. 8, we investigate the impact of the thermal-relaxation error on the high-encoding-rate BB code memories. Similar to the surface code experiments, we perform d rounds of syndrome extraction for distance d codes. Thermal-relaxation errors are injected before each measurement round to all the qubits. The resulting syndromes are processed using the BP-OSD decoder [69,70] to numerically compute the LER. The PTA thermal-relaxation channel shown in Fig. 8 serves as a baseline for real thermalization performance, while previous studies use uniform depolarizing errors without a direct connection to noise. The decomposed channel shows an offset from PTA, with approximately $1.5\text{--}2\times$ lower error rate relative to PTA, consistent with the trends observed in surface code experiments.

In BB codes, different logical qubits can share subsets of data qubits, determined by the generating polynomials




FIG. 8. The LER (per experiment) of bicycle bivariate code memory experiments (rounds = distance) decoded with BP-OSD decoder. During each round of syndrome extraction, measurement resets are preceded and followed by thermal-relaxation channels, with measurement duration τ . LERs obtained using the exact thermal relaxation are shown as solid lines, while those using the PTA channels are plotted as dashed lines. Error bars show a 95% confidence interval given 5 million samples.

of the code. This code structure disperses errors across the physical qubits, reducing the likelihood that a small number of localized faults produce a logical failure, but at the cost of increased connectivity and a more demanding decoding procedure that benefits from numerical refinement.

For the purpose of error modeling, we find that this disbursement aggregates to a nearly constant offset between the LER of the PTA channel and the exact thermal-relaxation error channel, similar to what was found in earlier low-distance surface code studies. OSD refinement can further smooth out logical error behavior by correcting unreliable bits, e.g., a qubit that has decayed and no longer has the expected behavior in future parity checks. Additionally, when comparing the LER of the BB code with that of the surface code under realistic thermal-relaxation noise, we observe that the BB code can achieve comparable or even lower logical error rates while encoding substantially more logical qubits, highlighting the advantages provided by its high encoding rate.

IV. DISCUSSION

This work demonstrates that combining amplitude damping and dephasing into a unified thermal-relaxation channel can substantially reduce or even eliminate the QP sampling overhead required for exact simulation of non-Clifford noises. This finding opens several interesting avenues for further

research in both theoretical understanding and large-scale simulation of realistic quantum error-correcting codes.

A natural extension of this work is to incorporate nonuniform, qubit-wise thermal noise models into the stabilizer simulation framework for better characterizing the hardware performance in QEC code implementations. Real devices exhibit spatial and temporal varying coherence properties. An efficient stabilizer simulator that can incorporate hardware-specific, qubit-wise T_1 and T_2 values from device calibration data and apply distinct thermal-relaxation errors to each qubit would enable substantially more accurate predictions of the QEC code performance, especially the logical error rates.

Our findings also connect to the broader questions of the robustness and magic of the quantum channels. Recent work by Howard *et al.*, building on the robustness of the quantum channels [71,72], showed that compositing coherent error channels and Pauli error channels can reduce the channel robustness and the overhead for simulation, which our work is consistent with Ref. [55]. This opens up a question: What other useful channels, previously thought of as too sophisticated for stabilizer simulation, can this technique be applied to? While we focus on relaxation and decoherence for this work, there are other physical processes that continuously affect qubit operation and can potentially be efficiently simulated by combining with other operations. This requires a systematic framework for identifying and optimizing the combinations to expand the capability of classical simulators. Furthermore, it is possible that commonly reused non-Clifford gate blocks, which contain a lot of magic can be combined in such a way that makes a minimally negative non-Clifford channel. If these channels can be combined with little negativity, there could be shortcuts in algorithm design, or at least a rethink of the most critical advantages that quantum computing brings us.

On the applications side, the approximate thermal-relaxation models developed here, such as the reset-based approximation, may also prove valuable for decoder design, training, and calibration. Recent development of machine-learning-based decoders shows great advantages in decoding speed and scalability to larger systems [73,74]. Using PTA or reset channels as pretraining models could help adapt decoders to the biased noise channels from the thermal relaxation and to hardware realities. Such noise modeling may improve logical performance in thermal-noise regimes.

From a theoretical perspective, the connection between channel-level metrics (e.g., channel fidelity, state distance) and logical-level properties (e.g., error threshold, bias sensitivity, or error suppression rate) remains incompletely understood. Establishing quantitative relationships between microscopic channel structure and macroscopic fault-tolerant behavior would provide a principled foundation for noise-tailored code design and decoder optimization.

Recent ongoing efforts in rare event simulation may also be helpful in combination with more accurate error modeling to investigate fault-tolerant designs. Higher-distance codes presented in Figs. 6 and 8 require tens of millions or more shots to encounter a logical error, even given the relaxation and measurement times of only current superconducting qubits, which are still improving. Many-shot fast stabilizer simulation and distribution via Sinter and Stim [28] have played a

significant role in enabling works up to this point. However, for large fault-tolerant operations including lattice surgery between logical qubits, works in recruiting parallelization for single-shot speedup [63] and circuit-level rare event simulation [75] are promising for the viability of future quantum error-correction work.

V. CONCLUSION

In this work, we developed a stabilizer-compatible method for efficiently simulating the thermal-relaxation error channels by combining amplitude damping and dephasing into a unified channel. We showed that in the experimentally relevant regimes $T_2 \leq T_1$, this composite error channel admits a completely positive Clifford decomposition with no quasiprobability sampling overhead. This decomposition enabled efficient and scalable simulation of realistic thermal-relaxation processes within the stabilizer frameworks for QEC study and beyond.

Using this approach, we integrate the thermal-relaxation error channels into our GPU-accelerated stabilizer simulator and perform large-scale simulations of QEC codes to understand their logical performance under realistic noise conditions. We focused on superconducting qubit platforms where the qubits' T_1 and T_2 times fall within the sweet spot of the Clifford decomposition. We investigated the logical performance of the two most promising QEC codes in superconducting platforms, i.e., surface codes and the BB code, and compared their noisy performance with the channels under PTA. Across memory experiments for both code examples, we observe that PTA can misestimate logical error rates by 2–10 $\times$ in either direction, depending on code distance and code type. We also observed that the rate of error suppression with exact thermal relaxation could depend on the logical states of the code block, which we attribute to the decoder given the closeness in excited state population.

In the regime where $T_2 > T_1$, we further introduced a reset-based approximation to the thermal-relaxation error channel. In this regime, we have shown that this reset approximation can consistently outperform the PTA while maintaining the complete positivity of the decomposition and compatibility with the stabilizer simulation framework. This model provides a practical route for realistic noise simulations without the cost of the full sampling overhead of QPD.

These results establish a practical route for incorporating physically realistic thermal-relaxation noise into fast stabilizer-based simulation frameworks. Approximations that preserve Clifford compatibility while more accurately capturing the intrinsic bias and directional structure of thermal-relaxation noise can play an important role in designing and assessing quantum error-correction codes for fault-tolerant quantum computing.

ACKNOWLEDGMENTS

We thank Shifan Xu, Ming Wang, and Shuwen Kan for fruitful discussions on BB code and surface code decoding. The research on noise modeling was supported by Pacific Northwest National Laboratory's Quantum Algorithms and Architecture for Domain Science (QuAADS) Laboratory




FIG. 9. The logical error rate of the surface code memory experiments when the code is in logical $|0\rangle_L$ (solid colors) and logical $|1\rangle_L$ states (hollow markers).




FIG. 10. Decoded Steane code memory comparing the exact Kraus channel to the reset approximation introduced in Sect. II D and Pauli-twirling approximation discussed in Sec. II B. Direct simulation of a small example allows us to connect the better channel fidelity shown in Fig. 4 to a closer approximation of code performance under the true noise channel. Each state vector point uses 2000 samples, while the stabilizer-compatible methods are sampled with 10 000 shots each. The bottleneck of this analysis is state vector sampling.




FIG. 11. Logical state preparation circuit for state $|0\rangle_L$.

Directed Research and Development (LDRD) Initiative. The Clifford simulation work was supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Science Center (QSC). N.M.M. was supported by the “Embedding QC into Many-body Frameworks for Strongly Correlated Molecular and Materials Systems” project, which is funded by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences (BES), the Division of Chemical Sciences, Geosciences, and Biosciences under FWP 72689. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract No. DE-AC05-00OR22725. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. The Pacific Northwest National Laboratory is operated by Battelle for the U.S. Department of Energy under Contract No. DE-AC05-76RL01830.

DATA AVAILABILITY

The data that support the findings of this article are openly available [76].

APPENDIX A: LER OF SURFACE CODE LOGICAL $|0\rangle_L$ AND $|1\rangle_L$ STATES

In Fig. 9, we show the comparison of the LERs of the surface code memory experiments when the code is in state $|1\rangle_L$ and $|0\rangle_L$. We observe that the performance of the logical $|1\rangle_L$ state is nearly identical to the logical $|0\rangle_L$ state under both noise methods.

APPENDIX B: CONNECTING IMPROVED CHANNEL FIDELITY TO CLOSER LER APPROXIMATIONS

In Sec. II D, we introduce a reset-only decomposition approximation that has better channel fidelity to true relaxation versus Pauli-twirling approximation. This approximation is relevant when $T_2 < T_1 \leq 2T_2$, i.e., where there are not enough decoherence errors to make the combined distribution entirely positive. Although this regime is not efficiently simulable as discussed in Sec. II C, for very small codes, exact methods are possible to determine baseline performance.

As a concrete example to connect the channel fidelity improvement to closer LER approximations, we simulate a three

round Steane code memory decoded with Chromobius [77], using three different methods of sampling. First, we simulate the exact Kraus channel in Eq. (5) using state vector simulation. Then, we simulate the reset approximation in Eq. (30) and PTA in Eq. (20). We then calculate the logical error rate of the final logical observable after each error modeling method. Error injection in Fig. 10 is performed the same way as in Figs. 6 and 8, that is, applying the relaxation channel around measurement blocks. However, relaxation in Fig. 10 is simulated at the limit $T_1 = 2T_2$. At this worst-case scenario, the combined QPD would provide no advantage as there is no pure dephasing to reduce the negativity in Eq. (35).

APPENDIX C: CODE BLOCK OCCUPATION CIRCUIT

In Fig. 11, we provide the circuit we use to compute the average occupation of the surface code block when the logical state is $|0\rangle_L$. We initially prepare all the physical qubits in state $|0\rangle$ and then perform one round of syndrome checks. At the end of the circuit, Pauli-Z measurements are performed on all the qubits, including the data and syndrome qubits, to compute the average excitation probability. We do not perform postselection on the outcomes of the syndrome qubits, and hence it is an average of all the possible physical realizations of the logical state $|0\rangle_L$. The other two logical states, $|+\rangle_L$ and $|1\rangle_L$, are prepared similarly.

-
- [1] J. Preskill, Fault-tolerant quantum computation, in *Introduction to Quantum Computation and Information* (World Scientific, Singapore, 1998), pp. 213–269.
 - [2] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
 - [3] A. Katabarwa, K. Gratsea, A. Caesura, and P. D. Johnson, Early fault-tolerant quantum computing, *PRX Quantum* **5**, 020101 (2024).
 - [4] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, The future of quantum computing with superconducting qubits, *J. Appl. Phys.* **132**, 160902 (2022).
 - [5] K. S. Chou *et al.*, A superconducting dual-rail cavity qubit with erasure-detected logical measurements, *Nat. Phys.* **20**, 1454 (2024).
 - [6] R. Acharya *et al.*, Quantum error correction below the surface code threshold, *Nature (London)* **638**, 920 (2025).
 - [7] N. Lacroix *et al.*, Scaling and logic in the colour code on a superconducting quantum processor, *Nature (London)* **645**, 614 (2025).
 - [8] A. Paetzniak *et al.*, Demonstration of logical qubits and repeated error correction with better-than-physical error rates, *arXiv:2404.02280*.
 - [9] B. W. Reichardt, D. Aasen, R. Chao, A. Chernoguzov, W. v. Dam, J. P. Gaebler, D. Gresh, D. Lucchetti, M. Mills, S. A. Moses, B. Neyenhuis, A. Paetzniak, A. Paz, P. E. Siegfried, M. P. d. Silva, K. M. Svore, Z. Wang, and M. Zanner, Demonstration of quantum computation and error correction with a tesseract code, *arXiv:2409.04628*.
 - [10] K. Sahay, J. Jin, J. Claes, J. D. Thompson, and S. Puri, High-threshold codes for neutral-atom qubits with biased erasure errors, *Phys. Rev. X* **13**, 041013 (2023).
 - [11] Q. Xu, J. P. Bonilla Ataides, C. A. Pattison, N. Raveendran, D. Bluvstein, J. Wurtz, B. Vasić, M. D. Lukin, L. Jiang, and H. Zhou, Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays, *Nat. Phys.* **20**, 1084 (2024).
 - [12] D. Bluvstein *et al.*, A fault-tolerant neutral-atom architecture for universal quantum computation, *Nature (London)* **649**, 39 (2026).
 - [13] 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, S. H. Li, X. Lyu, L. M. Peters, D. Bluvstein, T. T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin, Continuous operation of a coherent 3,000-qubit system, *Nature (London)* **646**, 1075 (2025).
 - [14] 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).
 - [15] T. J. Yoder and I. H. Kim, The surface code with a twist, *Quantum* **1**, 2 (2017).
 - [16] C. Gidney, M. Newman, P. Brooks, and C. Jones, Yoked surface codes, *Nat. Commun.* **16**, 4498 (2025).
 - [17] H. Bombin and M. A. Martin-Delgado, Exact topological quantum order in $D = 3$ and beyond: Branyons and brane-net condensates, *Phys. Rev. B* **75**, 075103 (2007).
 - [18] A. G. Fowler, 2D color code quantum computation, *Phys. Rev. A* **83**, 042310 (2011).
 - [19] A. Kubica, B. Yoshida, and F. Pastawski, Unfolding the color code, *New J. Phys.* **17**, 083026 (2015).
 - [20] 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 (London)* **627**, 778 (2024).

- [21] 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](#) [quant-ph].
- [22] D. Ruiz, J. Guillaud, A. Leverrier, M. Mirrahimi, and C. Vuillot, LDPC-cat codes for low-overhead quantum computing in 2D, *Nat. Commun.* **16**, 1040 (2025).
- [23] D. Gottesman, The Heisenberg representation of quantum computers, [arXiv:quant-ph/9807006](#).
- [24] A. Katabarwa and M. R. Geller, Logical error rate in the Pauli twirling approximation, *Sci. Rep.* **5**, 14670 (2015).
- [25] M. R. Geller and Z. Zhou, Efficient error models for fault-tolerant architectures and the Pauli twirling approximation, *Phys. Rev. A* **88**, 012314 (2013).
- [26] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
- [27] D. Gottesman, Stabilizer codes and quantum error correction, [arXiv:quant-ph/9705052](#).
- [28] C. Gidney, Stim: A fast stabilizer circuit simulator *Quantum* **5**, 497 (2021).
- [29] S. Bravyi, M. Englbrecht, R. König, and N. Peard, Correcting coherent errors with surface codes, *npj Quantum Inf.* **4**, 55 (2018).
- [30] E. Huang, A. C. Doherty, and S. Flammia, Performance of quantum error correction with coherent errors, *Phys. Rev. A* **99**, 022313 (2019).
- [31] K. Georgopoulos, C. Emary, and P. Zuliani, Modelling and simulating the noisy behaviour of near-term quantum computers, *Phys. Rev. A* **104**, 062432 (2021).
- [32] D. DiVincenzo, D. Leung, and B. Terhal, Quantum data hiding, *IEEE Trans. Inf. Theory* **48**, 580 (2002).
- [33] W. Dür, M. Hein, J. I. Cirac, and H.-J. Briegel, Standard forms of noisy quantum operations via depolarization, *Phys. Rev. A* **72**, 052326 (2005).
- [34] J. Ghosh, A. G. Fowler, and M. R. Geller, Surface code with decoherence: An analysis of three superconducting architectures, *Phys. Rev. A* **86**, 062318 (2012).
- [35] Y. Tomita and K. M. Svore, Low-distance surface codes under realistic quantum noise, *Phys. Rev. A* **90**, 062320 (2014).
- [36] R. S. Bennink, E. M. Ferragut, T. S. Humble, J. A. Laska, J. J. Nutaro, M. G. Pleszkoch, and R. C. Pooser, Unbiased simulation of near-Clifford quantum circuits, *Phys. Rev. A* **95**, 062337 (2017).
- [37] M. Katsuda, K. Mitarai, and K. Fujii, Simulation and performance analysis of quantum error correction with a rotated surface code under a realistic noise model, *Phys. Rev. Res.* **6**, 013024 (2024).
- [38] K. Mitarai and K. Fujii, Overhead for simulating a non-local channel with local channels by quasiprobability sampling, *Quantum* **5**, 388 (2021).
- [39] Y. Shao, Z. Chen, Z. Wei, and Z. Liu, Diagnosing quantum circuits: Noise robustness, trainability, and expressibility, [arXiv:2509.11307](#) [quant-ph].
- [40] J. Denzler, J. Carrasco, J. Eisert, and T. Guaita, Simulation of noisy quantum circuits using frame representations, [arXiv:2601.05131](#) [quant-ph].
- [41] M. AbuGhanem, IBM quantum computers: Evolution, performance, and future directions, *J. Supercomput.* **81**, 687 (2025).
- [42] D. Puzzioli, C. Granade, H. Haas, B. Criger, E. Magesan, and D. G. Cory, Tractable simulation of error correction with honest approximations to realistic fault models, *Phys. Rev. A* **89**, 022306 (2014).
- [43] M. Gutiérrez and K. R. Brown, Comparison of a quantum error-correction threshold for exact and approximate errors, *Phys. Rev. A* **91**, 022335 (2015).
- [44] Á. Márton and J. K. Asbóth, Coherent errors and readout errors in the surface code, *Quantum* **7**, 1116 (2023).
- [45] A. S. Darmawan and D. Poulin, Tensor-network simulations of the surface code under realistic noise, *Phys. Rev. Lett.* **119**, 040502 (2017).
- [46] Z. Schwartzman-Nowik, L. Shirizly, and H. Landa, Modeling error correction with Lindblad dynamics and approximate channels, *Phys. Rev. A* **111**, 022613 (2025).
- [47] M. O. Scully and M. S. Zubairy, *Quantum Optics* (Cambridge University Press, Cambridge, 1997).
- [48] C. W. Gardiner and P. Zoller, *Quantum Noise: A Handbook of Markovian and Non-Markovian Quantum Stochastic Methods with Applications to Quantum Optics*, 3rd ed. (Springer, Berlin, New York, 2004).
- [49] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th anniversary ed. (Cambridge University Press, Cambridge, 2011).
- [50] J. J. Wallman and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, *Phys. Rev. A* **94**, 052325 (2016).
- [51] D. J. Stonner, J. Kysela, G. Weir, J. Novotny, G. Alber, and I. Jex, Deterministic twirling with low resources, *Phys. Lett. A* **384**, 126179 (2020).
- [52] K. Tsubouchi, Y. Mitsuhashi, K. Sharma, and N. Yoshioka, Symmetric Clifford twirling for cost-optimal quantum error mitigation in early FTQC regime, *npj Quantum Inf.* **11**, 104 (2025).
- [53] J. M. Chow, J. M. Gambetta, A. D. Córcoles, S. T. Merkel, J. A. Smolin, C. Rigetti, S. Poletto, G. A. Keefe, M. B. Rothwell, J. R. Rozen, M. B. Ketchen, and M. Steffen, Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits, *Phys. Rev. Lett.* **109**, 060501 (2012).
- [54] L. Hantzko, L. Binkowski, and S. Gupta, Pauli transfer matrices, *Phys. Scr.* **100**, 075125 (2025).
- [55] S. Hakkaku, K. Mitarai, and K. Fujii, Sampling-based quasiprobability simulation for fault-tolerant quantum error correction on the surface codes under coherent noise, *Phys. Rev. Res.* **3**, 043130 (2021).
- [56] C. Piveteau, D. Sutter, and S. Woerner, Quasiprobability decompositions with reduced sampling overhead, *npj Quantum Inf.* **8**, 12 (2022).
- [57] S. Shankar, A. M. Tyryshkin, J. He, and S. A. Lyon, Spin relaxation and coherence times for electrons at the Si/SiO₂ interface, *Phys. Rev. B* **82**, 195323 (2010).
- [58] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Corcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, Superconducting qubit in waveguide cavity with coherence time approaching 0.1 ms, *Phys. Rev. B* **86**, 100506(R) (2012).
- [59] K. Mayer and E. Knill, Quantum process fidelity bounds from sets of input states, *Phys. Rev. A* **98**, 052326 (2018).

- [60] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, *Phys. Rev. A* **76**, 042319 (2007).
- [61] K. Wang *et al.*, Demonstration of low-overhead quantum error correction codes, *Nat. Phys.* **22**, 308 (2026).
- [62] IBMQ, Compute resources | IBM quantum platform (2025), accessed: Nov. 18, 2025.
- [63] S. Garner, C. Liu, M. Wang, S. Stein, and A. Li, STABSim: A parallelized Clifford simulator with features beyond direct simulation, [arXiv:2507.03092](https://arxiv.org/abs/2507.03092).
- [64] S. Bravyi, D. Poulin, and B. Terhal, Tradeoffs for reliable quantum information storage in 2D systems, *Phys. Rev. Lett.* **104**, 050503 (2010).
- [65] D. Horsman, A. G. Fowler, S. Devitt, and R. V. Meter, Surface code quantum computing by lattice surgery, *New J. Phys.* **14**, 123011 (2012).
- [66] S. Stein, S. Xu, A. W. Cross, T. J. Yoder, A. Javadi-Abhari, C. Liu, K. Liu, Z. Zhou, C. Guinn, Y. Ding, Y. Ding, and A. Li, HetEC: Architectures for heterogeneous quantum error correction codes, in *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)* (Association for Computing Machinery, New York, NY, USA, 2025), Vol. 2, pp. 515–528.
- [67] D. Poulin and Y. Chung, On the iterative decoding of sparse quantum codes, *Quantum Inf. Comput.* **8**, 987 (2008).
- [68] C.-F. Kung, K.-Y. Kuo, and C.-Y. Lai, On belief propagation decoding of quantum codes with quaternary reliability statistics, in *2023 12th International Symposium on Topics in Coding (ISTC)*, Brest, France (IEEE, 2023), pp. 1–5.
- [69] J. Roffe, D. R. White, S. Burton, and E. Campbell, Decoding across the quantum low-density parity-check code landscape, *Phys. Rev. Res.* **2**, 043423 (2020).
- [70] J. Roffe, LDPC: Python tools for low density parity check codes (2022), <https://pypi.org/project/ldpc/>.
- [71] M. Howard and E. Campbell, Application of a resource theory for magic states to fault-tolerant quantum computing, *Phys. Rev. Lett.* **118**, 090501 (2017).
- [72] J. R. Seddon and E. T. Campbell, Quantifying magic for multi-qubit operations, *Proc. R. Soc. A* **475**, 20190251 (2019).
- [73] J. P. B. Ataides, A. Gu, S. F. Yelin, and M. D. Lukin, Neural decoders for universal quantum algorithms, [arXiv:2509.11370](https://arxiv.org/abs/2509.11370).
- [74] J. Bausch, A. W. Senior, F. J. H. Heras, T. Edlich, A. Davies, M. Newman, C. Jones, K. Satzinger, M. Y. Niu, S. Blackwell, G. Holland, D. Kafri, J. Atalaya, C. Gidney, D. Hassabis, S. Boixo, H. Neven, and P. Kohli, Learning high-accuracy error decoding for quantum processors, *Nature (London)* **635**, 834 (2024).
- [75] C. Mayer, A. Ganti, U. Onunkwo, T. Metodi, B. Anker, and J. Skryzalin, Rare event simulation of quantum error-correcting circuits, [arXiv:2509.13678](https://arxiv.org/abs/2509.13678).
- [76] S. Garner, N. M. Myers, and C. Liu, Composite decomposition repository, 2026, <https://github.com/seangarn32/Composite-Decomposition>.
- [77] C. Gidney and C. Jones, New circuits and an open source decoder for the color code, [arXiv:2312.08813](https://arxiv.org/abs/2312.08813).