Exploring Noisy Quantum Thermodynamical Processes via the Depolarizing-Channel Approximation

Jian Li,[1,2] Xiaoyang Wang,[3,4,*] Marcus Huber,[1,5] Nicolai Friis,[1] and Pharnam Bakhshinezhad[1,†]
[1]Technische Universität Wien, Atominstitut & Vienna Center for Quantum Science and Technology (VCQ), Stadionallee 2, 1020 Vienna, Austria
[2]Department of Physics, University of Vienna, Boltzmanngasse 5, 1090 Vienna, Austria
[3]RIKEN Center for Interdisciplinary Theoretical and Mathematical Sciences (iTHEMS), Wako 351-0198, Japan
[4]RIKEN Center for Computational Science (R-CCS), Kobe 650-0047, Japan
[5]Institute for Quantum Optics and Quantum Information—IQOQI Vienna, Austrian Academy of Sciences, Boltzmanngasse 3, 1090 Vienna, Austria

(Received 13 February 2026; revised 16 June 2026; accepted 8 October 2026; published 2 November 2026)

Noise and errors are unavoidable in any realistic quantum process, including processes designed to reduce noise and errors in the first place. In particular, quantum thermodynamical protocols for cooling can be significantly affected, potentially altering both their performance and efficiency. Analytically characterizing the impact of such errors becomes increasingly challenging as the system size grows, particularly in deep quantum circuits where noise can accumulate in complex ways. To address this, we introduce a general framework for approximating the cumulative effect of gate-dependent noise using a global depolarizing channel. We specify the regime in which this approximation provides a reliable description of the noisy dynamics. Applying our framework to the thermodynamical two-sort algorithmic cooling (TSAC) protocol, we analytically derive its asymptotic cooling limit in the presence of noise. Using the cooling limit, the optimal cooling performance is achieved by a finite number of qubits—distinguished from the conventional noiseless TSAC protocol by an infinite number of qubits—and fundamental bounds on the achievable ground-state population are derived. This approach opens new avenues for exploring noisy quantum thermodynamical processes.

DOI: 10.1103/wg2x-z2vm

Introduction—Quantum information processing requires precise control over fragile quantum systems. Tasks such as computation, communication, and sensing often assume ideal noiseless operations such as perfect gates, measurements, and timing. In reality, noise from imperfect control and interactions with the environment inevitably degrade performance. To compensate for noise it is paramount to understand these imperfections, identify error sources, characterize their effect, and develop suitable mitigation strategies. Provided that error probabilities are below certain thresholds, quantum error correction can effectively combat the corresponding errors (see, e.g., Refs. [1–3]). Yet, devices operating in the noisy intermediate-scale quantum (NISQ) [4] regime studied over the past decade have still been far away from these thresholds, and only recently first devices have been reported to operate at or below threshold [5]. In addition, many practical computations are performed without the protection of quantum error correction due to severely limited quantum computational resources. However, in some cases the effects of noise on the results of quantum computations can be predicted and partially ameliorated using quantum error-mitigation schemes [6–10] that leverage physically inspired assumptions on the noise models. Indeed, such mitigation schemes have already been applied to demonstrate the utility of NISQ devices [11–13].

In this Letter, using similar physically inspired assumptions, we study the effect of gate-dependent Markovian noise on a process that is itself a primitive of most (if not all) of the aforementioned tasks: the preparation of sufficiently pure initial states. In the context of quantum thermodynamics this task is routinely referred to as cooling [14–16], as other pure states can be reached from the ground state by suitable unitary transformations. Specifically, we focus on unitary thermodynamical processes implemented via quantum circuits, with finite resources that can be systematically tracked. While realistic noise, which takes into account effects such as dephasing, imprecise timekeeping, and gate imperfections is difficult to treat analytically, the

*Contact author: xiaoyang.wang@riken.jp
†Contact author: pharnam.bakhshinezhad@tuwien.ac.at

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.

depolarizing noise model captures the average effect of more realistic noise models and allows for analytical evaluation [17–19]. These averaging effects have been analyzed using random quantum circuits [20]. However, it is unclear how reasonable this approach is for quantum circuits with fixed gates, which are used in many state-preparation and evolution circuits, as the averaging of gate-dependent Markovian noise as well as the corresponding requirements on the quantum circuits have not yet been systematically studied. Here, we therefore introduce and employ a framework based on what we refer to as the global depolarizing approximation (GDA) to describe how realistic noise is averaged in deep quantum circuits with fixed gates. We show that the depolarizing noise well approximates such realistic noise as long as the circuit depth grows polynomially with the system size.

Using the GDA, we can efficiently predict the noise effect in deep quantum circuits. Specifically, an n-qubit circuit with L quantum gates without parallelization has circuit depth L. Among these gates, $N \leq L$ are occurrences of the noise-affected two-qubit gates with the same error probability p. With additional assumptions on the randomness and the noise strength of the circuit, the GDA predicts the cumulative noise effect via a depolarizing noise strength $\eta = pN(1-q)$, where q is determined by the local noise channel of the two-qubit gates. As an example, we apply the GDA to the two-sort algorithmic cooling (TSAC) protocol [21,22]. We derive an analytical expression for its steady-state cooling limit under noise and validate our model by comparing it with numerical simulations of gate-dependent noise and a mirror-cooling protocol [23]. Our framework effectively predicts final cooling limits, cooling dynamics, and optimal qubit numbers under Markovian noise, and can be combined with noise-tailoring tools, such as randomized compiling, when coherent control errors are present.

The global depolarizing approximation—We consider an n-qubit circuit composed of L sequentially ordered noiseless gates $\hat{\mathcal{U}} = \bigcirc_{l=1}^{L} \hat{\mathcal{U}}_l$, where each $\hat{\mathcal{U}}_l$ is a superoperator acting on density operators ρ as $\hat{\mathcal{U}}_l(\rho) = U_l \rho U_l^\dagger$. To describe the realistic noisy version of this circuit, we assume that the noise channel is a Markovian completely positive trace-preserving (CPTP) map, but non-Markovian and leakage noise are not included due to their subdominant effects [24–28]. A CPTP map can be coherent or incoherent [17,29]. Coherent noise due to inaccurate quantum control preserves state purity, while incoherent noise reduces purity mainly coming from randomness in quantum operations and entanglement of qubits with their environment. By applying additional random Pauli gates, coherent noise can be efficiently converted to incoherent noise using randomized compiling [29,30]. Additionally, we take the following two physically motivated assumptions: (i) two-qubit gates are the dominant source of error, and the circuit contains N occurrences of a single noise-affected two-qubit gate type g; other single-qubit gates are treated as noiseless; (ii) the circuit is sufficiently deep, and the set of noiseless unitaries generated by $\{U_l\}$ is random enough to form approximate unitary 2-designs, i.e., they resemble Haar-random circuit ensembles up to the second-order moments [31].

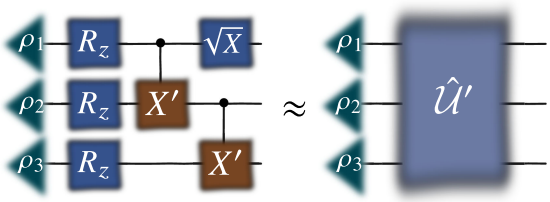


FIG. 1. Global depolarizing approximation. Local noises of two-qubit CNOT gates are twirled by the quantum circuit with fixed gates (left panel), making them effectively act as a global depolarizing channel $\hat{\mathcal{U}}'$ on the observable (right panel).

For each occurrence of the two-qubit gate g, its noisy implementation has the decomposition $\mathcal{M}_g = \mathcal{E}_g \circ \hat{\mathcal{U}}_g$ [32], where $\hat{\mathcal{U}}_g$ is the noiseless superoperator and $\mathcal{E}_g$ is an incoherent noise channel,

$$\mathcal{E}_g(\rho) = (1-p)\rho + p\Lambda(\rho). \tag{1}$$

Here, $0 < p < 1$ is the error probability and $\Lambda(\rho) \neq \rho$ is a fixed local noise channel; its embedding on the qubit pair addressed by a given occurrence of g is implicit. The single-gate-type setting in assumption (i) is relevant to common digital quantum-computing architectures [33–35], since one entangling gate type together with single-qubit gates is sufficient for universality [36]. The second assumption (ii) describes the randomness and twirling effects of the noiseless quantum circuits, which has been found to hold in many dynamic and state-preparation quantum circuits [37–39].

Using the above circuit noise model, we show that the noise in the circuit superoperator $\hat{\mathcal{U}}'$ is twirled by the circuit itself, such that the average effect of the noise acting on an initial state ρ_0 can be approximated by a global depolarizing channel (Fig. 1),

$$\hat{\mathcal{U}}'(\rho_0) \approx (1-\eta)\hat{\mathcal{U}}(\rho_0) + \eta\frac{\mathbb{I}}{d}, \tag{2}$$

where $d = 2^n$ is the Hilbert-space dimension, and η is the effective depolarizing strength given by

$$\eta := pN(1-q). \tag{3}$$

Here, N is the number of noisy occurrences of gate type g in the circuit, and $q := [\mathrm{Tr}(\Lambda) - 1]/(d^2 - 1)$ is a parameter related to the trace of the noise process Λ of the two-qubit gates [40], where instances of Λ implemented on different qubits need not to be distinguished due to assumption (i). Derivations of Eqs. (2) and (3) are given in Supplemental Material (SM) Sec. II A [41].

The GDA in Eq. (2) provides a direct way to predict the effect of noise in deep quantum circuits. For example, it allows us to establish a relation between the noiseless expectation and the noisy one, i.e.,

$$\mathrm{Tr}[O\hat{\mathcal{U}}(\rho_0)] \approx (1-\eta)^{-1}\,\mathrm{Tr}[O\hat{\mathcal{U}}'(\rho_0)] \tag{4}$$

for a traceless Pauli-string observable O (nonidentity), and we can directly calculate η based on the number of noisy two-qubit gates N and the local noise model Λ using Eq. (3). Previous works use Eq. (4) to predict noiseless expectation from the noisy measurement results, conditioned on an additional estimation circuit to effectively evaluate η [17,18]. Here, we do not rely on the additional estimation circuit, but use Eq. (3) to save measurement consumption.

Validity of GDA—The GDA relies on the assumption (ii) that the set of noiseless unitaries generated by $\{U_l\}$ forming approximate unitary 2-designs. On the one hand, the exact 2-design could require deep quantum circuits due to the typically exponential cardinality of the 2-design with respect to the system size [46]. On the other hand, if we focus only on estimating physical quantities such as the fidelity and observable expectations with an additive error, it can be shown that the required circuit depth grows polynomially with the system size.

Using Hoeffding's inequality [47], we prove in SM Sec. II B [41] that, to estimate an observable O within an additive error $\|\langle O\rangle_{\mathrm{meas}} - \langle O\rangle\| \le \varepsilon$ and a confidence level $1-\delta$, the required circuit depth L scales as

$$L \sim n(n-1)\frac{4\|O\|^2}{\varepsilon^2}\ln\left(\frac{2}{\delta}\right), \tag{5}$$

where n is the number of qubits and $\|O\|$ is the spectral norm of the observable O. For fidelity estimation and typical physical observables, where $\|O\| = 1$ and $\|O\|$ scales polynomially with n, respectively, the depth grows polynomially in n. Thus the GDA is scalable to large system sizes, and is validated in quantum algorithms with a polynomially increased circuit depth, as required by many dynamic and state-preparation quantum algorithms [48–50]. Even if the circuit is not deep enough, one can artificially increase the circuit depth without changing the noiseless circuit $\hat{\mathcal{U}}$ using, e.g., unitary gate folding [7,51].

For a given circuit with N two-qubit gates, we provide two operational criteria on the validity of GDA: (1) evaluate a function $pN\Delta^{\mathrm{out}}_{O,\rho}(N) + 2\|O\|(pN)^2$ and check whether it is below the error tolerance for the observable expectation value. (2) Check whether N falls within a conservative window $N_{\mathrm{req}} \le N < N_{2\mathrm{nd}}$. These criteria and the definition of these notations are summarized in the End Matter and detailed in SM Sec. II C [41].

Application to TSAC—We now demonstrate the utility of the GDA by applying it to analyze the TSAC protocol [21], a quantum algorithmic-cooling method to increase the purity of quantum states. We first review the ideal noiseless protocol, then show how our noise model modifies its core dynamics, and finally derive the analytical cooling limit in this realistic setting.

The TSAC protocol improves the purity and polarization of n_{c} computational qubits by introducing an additional reset qubit. The total system therefore consists of $n = n_{\mathrm{c}} + 1$ qubits. The total Hilbert space is $\mathcal{H} = \mathcal{H}_{\mathrm{C}} \otimes \mathcal{H}_{\mathrm{R}}$, representing the n_{c}-qubit computational space $\mathcal{H}_{\mathrm{C}}$ and the Hilbert space $\mathcal{H}_{\mathrm{R}}$ of the single reset qubit, respectively. The protocol iteratively applies the following two steps. (1) Reset step: the reset qubit is reset to a fresh thermal state $\rho_{\mathrm{R}} \coloneqq z^{-1}\,\mathrm{diag}(e^{\epsilon}, e^{-\epsilon})$, where $z = e^{\epsilon} + e^{-\epsilon}$ and $\epsilon > 0$ is the polarization parameter of the reset qubit. This operation traces out the current state of the reset qubit while preserving the computational subsystem, and then initializes the reset qubit to the fresh thermal state ρ_R. Mathematically, this channel $\mathcal{R}$ acting on the full density matrix ρ is defined as $\mathcal{R}(\rho) \coloneqq \rho_C \otimes \rho_R$, where $\rho_C = \mathrm{Tr}_R(\rho)$ denotes the reduced density matrix of the computational qubits, consistent with the Hilbert space decomposition $\mathcal{H} = \mathcal{H}_C \otimes \mathcal{H}_R$.

(2) Compression step: The compression unitary U_{TS} acts on the computational and reset qubits written in the computational basis $\{|0\rangle, |1\rangle\} \otimes \{|0\rangle, |1\rangle\}^{\otimes n_{\mathrm{c}}}$

$$U_{\mathrm{TS}} = 1 \oplus \sigma_x \oplus \cdots \oplus \sigma_x \oplus 1, \tag{6}$$

where $\oplus$ denotes the direct sum. This constructs a $2^n \times 2^n$ matrix containing $(2^{n_{\mathrm{c}}} - 1)$ Pauli-X blocks on its diagonal, bounded by scalar ones. Because U_{TS} is expressed in this global computational basis, these σ_x terms do not represent local Pauli operations on individual qubits. Rather, each acts on a distinct two-dimensional subspace of the full n-qubit Hilbert space, systematically permuting population levels to increase the polarization of the target qubit [21,22]. The ideal noiseless TSAC protocol iteratively applies $\mathcal{R}$ and U_{TS} to the density matrix of the n qubits.

To study the cooling effect of TSAC, we focus on the density matrix of the n_{c} computational qubits. The dynamics of the diagonal density-matrix elements can be described as a Markov chain. We denote the $d_{\mathrm{c}} = 2^{n_{\mathrm{c}}}$ diagonal elements after t iterations of TSAC as $\vec{v}^t = \{v_1^t, \ldots, v_{d_{\mathrm{c}}}^t\}$. The transition from $\vec{v}^t$ to $\vec{v}^{t+1}$ can be described by a transition matrix T given in [21,52], such that $\vec{v}^{t+1} = T\vec{v}^t$. In the cooling limit $t \to \infty$, the steady state $\vec{v}^{\infty}$ is the eigenvector of T with eigenvalue 1.

The effect of noise on the TSAC protocol can be studied in the GDA framework. The TSAC protocol can be translated to the digital quantum-circuit model by decomposing U_{TS} into CNOT and single-qubit gates [21], and the resulting n-qubit quantum circuit has a layered structure that is sufficiently deep and random to satisfy assumptions (i) and (ii); see SM Secs. II C and II D [41] for details. In

particular, after tracing out the reset qubit, Eq. (2) maps the ideal noiseless TSAC population update $T\vec{v}$ of Ref. [21] to $(1-\eta)T\vec{v}+\eta(1,\ldots,1)^T/d_c$, giving the noisy transition matrix T'_η as

$$T'_\eta = \begin{pmatrix} a_{\eta,+} & a_{\eta,+} & \eta/d_c & \cdots & \eta/d_c & \eta/d_c \\ a_{\eta,-} & \eta/d_c & a_{\eta,+} & \cdots & \eta/d_c & \eta/d_c \\ \eta/d_c & a_{\eta,-} & \eta/d_c & \cdots & \eta/d_c & \eta/d_c \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \eta/d_c & \eta/d_c & \eta/d_c & \cdots & \eta/d_c & a_{\eta,+} \\ \eta/d_c & \eta/d_c & \eta/d_c & \cdots & a_{\eta,-} & a_{\eta,-} \end{pmatrix}, \tag{7}$$

where $a_{\eta,\pm} \coloneqq (\eta/d_c)+(1-\eta)e^{\pm\epsilon}/z$, and $T'_{\eta=0}$ reduces to the noiseless transition matrix T. Similar to the noiseless case, the cooling limit of TSAC in the presence of noise is obtained from the eigenvector of T'_η with eigenvalue 1. For the iterative TSAC protocol comprising applications of $\mathcal{R}$, U_{TS}, and a global depolarizing channel with strength η, this eigenvalue problem has a unique steady-state solution. The resulting probability vector $\vec{v}=\{v_1,\ldots,v_{d_c}\}$ has components

$$v_k = z_1\lambda_1^{k-1}+z_2\lambda_2^{k-1}+\frac{1}{d_c}. \tag{8}$$

The parameters in this solution are fully determined by the physical properties of the system: (i) The eigenvalues λ_1 and λ_2 depend on the noise strength η, the polarization ϵ, and the number of computational qubits n_c. The eigenvalues are related by the noise-independent constraint $\lambda_1\lambda_2=e^{-2\epsilon}$. (ii) The coefficients z_1 and z_2 are normalization factors determined by λ_1 and λ_2, and thus also depend on (η,ϵ,n_c).

To derive this form, we rewrite the steady-state condition $T'_\eta\vec{v}=\vec{v}$ as a system of linear equations, which can be formulated as a second-order linear inhomogeneous recurrence relation for the components v_k. The general solution to such a recurrence relation is the sum of a particular solution (a constant related to the uniform background $1/d_c$) and a homogeneous solution, which is a linear combination of terms involving the two roots, λ_1 and λ_2, of the characteristic equation. The coefficients of this combination, z_1 and z_2, are then uniquely determined by the boundary conditions of the system, as detailed in SM Sec. III [41].

This analytical solution reveals the physical mechanism by which noise limits algorithmic cooling. The steady-state distribution is a superposition of a uniform background and two dynamic modes. The leading-order behavior in the noise strength η is $\lambda_1=1+\eta/\tanh\epsilon+\mathcal{O}(\eta^2)>1$ and $\lambda_2=e^{-2\epsilon}(1-\eta/\tanh\epsilon)+\mathcal{O}(\eta^2)<1$.

In the noiseless case ($\eta=0$), λ_1 becomes exactly 1, and its coefficient is $z_1=-1/d_c$. This causes the constant mode to cancel the uniform background $1/d_c$, leaving only the exponentially decaying term driven by $z_2>0$. Accordingly, Eq. (8) reduces to $v_k=z_2e^{-2(k-1)\epsilon}$, recovering the noiseless TSAC steady-state distribution of Ref. [21].

However, for any nonzero noise, the presence of the exponentially growing mode ($\lambda_1>1$) alters the behavior for a large number of computational qubits n_c. To satisfy the physical constraints of probability positivity and normalization, the amplitudes of both modes, $|z_1|$ and $|z_2|$, must be suppressed as the system size $d_c=2^{n_c}$ increases. This forces the entire distribution to flatten towards the uniform high-temperature state, $v_k\approx 1/d_c$.

This behavior is in contrast to the noiseless scenario where increasing n_c always improves cooling performance. Our result shows a trade-off: noise-induced suppression counteracts the benefits of a larger number of qubits. Thus, the analytical solution in Eq. (8) implies the existence of an optimal qubit number for which the best cooling performance is achieved. In the following, we will numerically study this optimal regime.

Optimal cooling performance with physical noise—To verify our analytical results, we numerically study the TSAC protocol with the depolarizing and realistic physical noise models. The depolarizing model is the theoretical noise model underlying the analytical solution in Eq. (8). For the physical noise model, we consider a two-qubit timekeeping-noise channel of strength p associated with each CNOT gate. Following Ref. [53], independent timing fluctuations of successive CNOT control pulses produce gate-local Markovian dephasing in the gate-Hamiltonian eigenbasis. This gate-local CNOT noise is consistent with assumption (i) and is relevant to pulse-controlled platforms. The explicit form of this channel is given in SM Sec. IV [41]. Using Eq. (3), its effective global depolarizing strength η_T on the full n-qubit system is

$$\eta_T = pN\frac{3d^2}{4(d^2-1)}, \tag{9}$$

where N is the number of CNOT gates in the n-qubit quantum circuit and $d=2^n$ is the total dimension. We use η_T to predict the numerical results of algorithmic cooling with timekeeping noise.

We now quantify this trade-off through the optimal qubit number and the corresponding cooling performance. The cooling performance is quantified by the final ground-state population of the target qubit (defined as the first computational qubit). $P_n \coloneqq \lim_{t\to\infty}\mathrm{Tr}(\rho_t|0\rangle\langle 0|)$, where ρ_t is the density matrix of the target qubit after t TSAC iterations, and we define the optimal cooling performance $P_{\max}\coloneqq \max_n P_n$ and the optimal qubit number $n_{\mathrm{opt}}\coloneqq \mathrm{argmax}_n P_n$. The values of n_{opt} and $P_{\max}$ change for different noise strengths. Studying these quantities provides guidance for

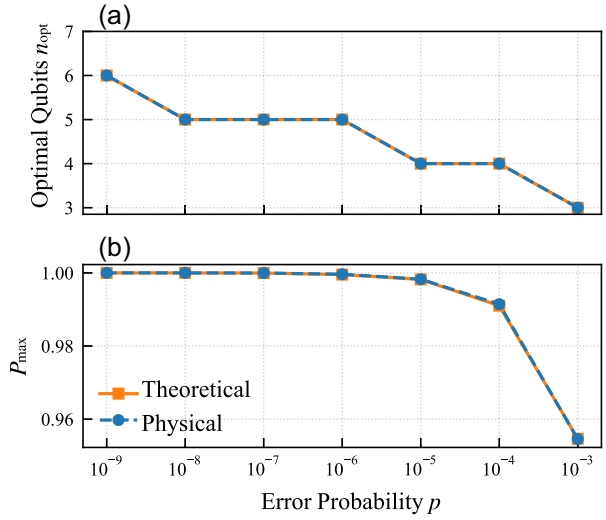


FIG. 2. Summary of optimal performance versus CNOT error probability p (logarithmic scale). (a) Optimal number of total qubits (n_{opt}) required to maximize the final ground-state population. (b) Corresponding maximally achievable final ground-state population (P_{max}). All panels compare results from physical noise simulations (circles, dashed lines) and the theoretical GDA model (squares, solid lines). The initial ground-state population is $P_{initial} = 0.85$.

the practical implementation of the TSAC protocol on noisy quantum hardware with some specific noise strength.

To determine n_{opt} and P_{max}, we simulate the TSAC protocol for various noise strengths and total number of qubits. The dependence of the optimal performance metrics n_{opt} and P_{max} versus error probability are summarized in Fig. 2. We assume that the initial state ρ_i of the target qubit has the ground-state population $P_{initial} = \mathrm{Tr}(\rho_i|0\rangle\langle 0|) = 0.85$. Figure 2(a) displays the optimal qubit number n_{opt} as a function of the CNOT error probability p. The value of p serves as the fundamental noise parameter for both curves. The "physical" curve (circles) shows the result of a direct gate-level simulation using the timekeeping-noise model (see SM Sec. IV [41]) with strength p. The "theoretical" curve (squares) shows the analytical prediction from Eq. (8), where we use the GDA strength $\eta = \eta_T$ given by Eq. (9). Figure 2(b) shows the corresponding maximum achievable ground-state population P_{max} with the corresponding n_{opt} from Fig. 2(a). Results comparing the physical noise simulation and the GDA model show agreement, validating our GDA framework's prediction on algorithmic cooling. Additionally, we see that both P_{max} and n_{opt} decrease monotonically as the error probability increases, which is consistent with the intuition that noise causes the cooling performance to deviate from the noiseless case. Specifically, currently available superconducting chips can achieve $p \sim 10^{-3}$ [54,55], with the optimal qubit number $n_{opt} = 3$. Therefore, unlike in the noiseless case, under the current noise level, further increasing the total number n of qubits does not improve the cooling performance.

We also extend the GDA framework to dynamic cooling (DC), where a single-shot global unitary redistributes entropy across the system. This approach qualitatively captures the cooling performance of the mirror protocol, with predicted effective temperatures aligning closely with physical noise simulations (see End Matter), as illustrated in Fig. 3. This agreement validates the model's reliability for analyzing diverse cooling architectures beyond iterative schemes [23,56].

Conclusion and outlook—We introduced the GDA, a scalable theoretical framework for analyzing the cumulative effects of Markovian noise in deep quantum circuits. The GDA approximates complex local noise channels with a single, system-wide depolarizing channel, requiring a circuit depth that only scales polynomially with the system size to emulate a unitary 2-design. Applying the GDA to the TSAC protocol, we derived a novel analytical expression for the asymptotic cooling limit under realistic noise [Eq. (8)]. This analysis identified a fundamental noise-induced trade-off: The exponential gain in the ideal cooling power with increasing qubit number n is limited by the exponential increase in noise accumulation in deep quantum circuits. This competition dictates an optimal qubit number of the cooling protocol and sets a sharp upper bound on the achievable ground-state population (Fig. 2). The accuracy of the GDA in predicting this optimal boundary was also validated against detailed physical noise simulations for both TSAC and DC protocols.

A key direction for future research is the broad application of the GDA framework to other noisy input-output circuit-based protocols, which include most quantum thermodynamical processes. This approach offers a powerful avenue for the analytical investigation of fundamental performance bounds [14,57–59] for thermodynamic tasks under realistic noise. Crucially, the GDA offers a critical tool for quantifying the thermodynamic cost of imperfect control. Physical constraints inherent in finite-resource implementations directly lead to operational errors using the depolarizing strength η featuring in the GDA. By analytically linking the thermodynamic resources (e.g., energy, time, or control precision) [16,53,60] required to achieve a target error tolerance to η, this framework facilitates the rigorous study of fundamental resource-error-performance trade-offs—a necessary step toward optimizing finite-resource quantum heat engines.

Acknowledgments—The authors thank Nayeli Rodríguez-Briones, Jake Xuereb, and Philip Taranto for insightful discussions. We acknowledge funding from the European Research Council (Consolidator grant "Cocoquest" No. 101043705) as well as from the Austrian Federal Ministry of Education, Science and Research via the Austrian Research Promotion Agency (FFG) through the

flagship Project No. FO999897481 (HPQC), the project FO999914030 (MUSIQ), the project FO999921407 (HDcode), and the project FO999921415 (Vanessa-QC) funded by the European Union—NextGenerationEU. This research was also funded in part by the Austrian Science Fund (FWF) [10.55776/P36478]. X. W. is supported by the RIKEN TRIP initiative (RIKEN Quantum) and the UTokyo Quantum Initiative.

Data availability—Simulation code supporting this study is publicly available in Ref. [61].

[1] D. A. Lidar and T. A. Brun, *Quantum Error Correction* (Cambridge University Press, Cambridge, 2013).
[2] B. M. Terhal, Quantum error correction for quantum memories, Rev. Mod. Phys. **87**, 307 (2015).
[3] J. Roffe, Quantum error correction: An introductory guide, Contemp. Phys. **60**, 226 (2019).
[4] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).
[5] Google Quantum AI and Collaborators, Quantum error correction below the surface code threshold, Nature (London) **638**, 920 (2025).
[6] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, Nature (London) **549**, 242 (2017).
[7] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, Phys. Rev. Lett. **119**, 180509 (2017).
[8] S. Endo, S. C. Benjamin, and Y. Li, Practical quantum error mitigation for near-future applications, Phys. Rev. X **8**, 031027 (2018).
[9] L. Funcke, T. Hartung, K. Jansen, S. Kühn, P. Stornati, and X. Wang, Measurement error mitigation in quantum computers through classical bit-flip correction, Phys. Rev. A **105**, 062404 (2022).
[10] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, Quantum error mitigation, Rev. Mod. Phys. **95**, 045005 (2023).
[11] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, Nature (London) **618**, 500 (2023).
[12] S. Cao *et al.*, Generation of genuine entanglement up to 51 superconducting qubits, Nature (London) **619**, 738 (2023).
[13] S. Guo *et al.*, Experimental quantum computational chemistry with optimized unitary coupled cluster ansatz, Nat. Phys. **20**, 1240 (2024).
[14] D. Reeb and M. M. Wolf, An improved Landauer principle with finite-size corrections, New J. Phys. **16**, 103011 (2014).
[15] F. Clivaz, R. Silva, G. Haack, J. B. Brask, N. Brunner, and M. Huber, Unifying paradigms of quantum refrigeration: A universal and attainable bound on cooling, Phys. Rev. Lett. **123**, 170605 (2019).
[16] P. Taranto, F. Bakhshinezhad, A. Bluhm, R. Silva, N. Friis, Maximilian P. E. Lock, G. Vitagliano, F. C. Binder, T. Debarba, E. Schwarzhans, F. Clivaz, and M. Huber, Landauer versus Nernst: What is the true cost of cooling a quantum system?, PRX Quantum **4**, 010332 (2023).
[17] M. Urbanek, B. Nachman, V. R. Pascuzzi, A. He, C. W. Bauer, and W. A. deJong, Mitigating depolarizing noise on quantum computers with noise-estimation circuits, Phys. Rev. Lett. **127**, 270502 (2021).
[18] J. Vovrosh, K. E. Khosla, S. Greenaway, C. Self, M. S. Kim, and J. Knolle, Simple mitigation of global depolarizing errors in quantum simulations, Phys. Rev. E **104**, 035309 (2021).
[19] Y. Takahashi, Y. Takeuchi, and S. Tani, Classically simulating quantum circuits with local depolarizing noise, in *45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 170, edited by J. Esparza and D. Král' (Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2020), pp. 83:1–83:13.
[20] A. M. Dalzell, N. Hunter-Jones, and F. G. S. L. Brandão, Random quantum circuits transform local noise into global white noise, Commun. Math. Phys. **405**, 78 (2024).
[21] S. Raeisi, M. Kieferová, and M. Mosca, Novel technique for robust optimal algorithmic cooling, Phys. Rev. Lett. **122**, 220501 (2019).
[22] P. Taranto, F. Bakhshinezhad, P. Schüttelkopf, F. Clivaz, and M. Huber, Exponential improvement for quantum cooling through finite-memory effects, Phys. Rev. Appl. **14**, 054005 (2020).
[23] L. Bassman Oftelie, A. De Pasquale, and M. Campisi, Dynamic cooling on contemporary quantum computers, PRX Quantum **5**, 030309 (2024).
[24] V. Vadimov, J. Tuorila, T. Orell, J. Stockburger, T. Ala-Nissila, J. Ankerhold, and M. Möttönen, Validity of born-Markov master equations for single- and two-qubit systems, Phys. Rev. B **103**, 214308 (2021).
[25] Z. Chen *et al.*, Measuring and suppressing quantum state leakage in a superconducting qubit, Phys. Rev. Lett. **116**, 020501 (2016).
[26] Z. Li, P. Liu, P. Zhao, Z. Mi, H. Xu, X. Liang, T. Su, W. Sun, G. Xue, J.-N. Zhang, W. Liu, Y. Jin, and H. Yu, Error per single-qubit gate below 10^{-4} in a superconducting qubit, npj Quantum Inf. **9**, 111 (2023).
[27] V. Tripathi, M. Khezri, and A. N. Korotkov, Operation and intrinsic error budget of a two-qubit cross-resonance gate, Phys. Rev. A **100**, 012301 (2019).
[28] F. Bao *et al.*, Fluxonium: An alternative qubit platform for high-fidelity operations, Phys. Rev. Lett. **129**, 010502 (2022).
[29] J. J. Wallman and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, Phys. Rev. A **94**, 052325 (2016).
[30] Z. Cai, X. Xu, and S. C. Benjamin, Mitigating coherent noise using Pauli conjugation, npj Quantum Inf. **6**, 17 (2020).
[31] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, Phys. Rev. A **80**, 012304 (2009).
[32] E. Magesan, J. M. Gambetta, and J. Emerson, Scalable and robust randomized benchmarking of quantum processes, Phys. Rev. Lett. **106**, 180504 (2011).

[33] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, Nature (London) **574**, 505 (2019).
[34] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis, Demonstration of the trapped-ion quantum CCD computer architecture, Nature (London) **592**, 209 (2021).
[35] P. Jurcevic *et al.*, Demonstration of quantum volume 64 on a superconducting quantum computing system, Quantum Sci. Technol. **6**, 025020 (2021).
[36] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th Anniversary Edition (Cambridge University Press, Cambridge, United Kingdom, 2010).
[37] X. Wang, X. Feng, L. Funcke, T. Hartung, K. Jansen, S. Kühn, G. Polykratis, and P. Stornati, Model-independent error mitigation in parametric quantum circuits and depolarizing projection of quantum noise, Proc. Sci. **396**, 603 (2022).
[38] K. Seki, Y. Kikuchi, T. Hayata, and S. Yunoki, Simulating Floquet scrambling circuits on trapped-ion quantum computers, Phys. Rev. Res. **7**, 023032 (2025).
[39] T. Yada, R. Suzuki, Y. Mitsuhashi, and N. Yoshioka, Non-Haar random circuits form unitary designs as fast as Haar random circuits, Phys. Rev. Lett. **136**, 030401 (2026).
[40] J. Emerson, R. Alicki, and K. Życzkowski, Scalable noise estimation with random unitary operators, J. Opt. B **7**, S347 (2005).
[41] See Supplemental Material at http://link.aps.org/supplemental/10.1103/wg2x-z2vm for details of randomized compiling, derivations and validity checks of the global depolarizing approximation, the analytical cooling limit, noise-model analysis, and numerical implementation, which includes Refs. [42–45].
[42] S. T. Flammia and J. J. Wallman, Efficient estimation of Pauli channels, ACM Trans. Quantum Comput. **1**, 3 (2020).
[43] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta, Quantum computing with Qiskit arXiv:2405.08810.
[44] K. Shende, Arvind, and K. Dorai, State-independent robust heat-bath algorithmic cooling of nuclear spins, Phys. Rev. Appl. **21**, 024017 (2024).
[45] G. Difranco and L. B. Oftelie, QuL: Programming library for computational cooling of qubits, ACM Trans. Quantum Comput. **6**, 1 (2025).
[46] S. Bravyi and D. Maslov, Hadamard-free circuits expose the structure of the Clifford group, IEEE Trans. Inf. Theory **67**, 4546 (2021).
[47] W. Hoeffding, Probability Inequalities for sums of bounded random variables, J. Am. Stat. Assoc. **58**, 13 (1963).
[48] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, Rev. Mod. Phys. **86**, 153 (2014).
[49] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, Nature (London) **607**, 667 (2022).
[50] T. Albash and D. A. Lidar, Adiabatic quantum computation, Rev. Mod. Phys. **90**, 015002 (2018).
[51] L. Hour, M. Go, and Y. Han, Improving zero-noise extrapolation for quantum-gate error mitigation using a noise-aware folding method, IEEE Access **14**, 37362 (2026).
[52] Z. Farahmand, R. Aghaei Saem, and S. Raeisi, Quantum noise can enhance algorithmic cooling, Phys. Rev. A **105**, 022418 (2022).
[53] J. Xuereb, P. Erker, F. Meier, M. T. Mitchison, and M. Huber, Impact of imperfect timekeeping on quantum control, Phys. Rev. Lett. **131**, 160204 (2023).
[54] W.-J. Lin, H. Cho, Y. Chen, M. G. Vavilov, C. Wang, and V. E. Manucharyan, 24 days-stable CNOT gate on fluxonium qubits with over 99.9% fidelity, PRX Quantum **6**, 010349 (2025).
[55] E. A. Sete, V. Tripathi, J. A. Valery, D. Lidar, and J. Y. Mutus, Error budget of a parametric resonance entangling gate with a tunable coupler, Phys. Rev. Appl. **22**, 014059 (2024).
[56] N. A. Rodríguez-Briones, Novel heat-bath algorithmic cooling methods, Ph.D. thesis, University of Waterloo, 2020.
[57] P. Lipka-Bartosik, G. F. Diotallevi, and P. Bakhshinezhad, Fundamental limits on anomalous energy flows in correlated quantum systems, Phys. Rev. Lett. **132**, 140402 (2024).
[58] J. Xuereb, B. Stratton, A. Rolandi, J. He, M. Huber, and P. Bakhshinezhad, Cooling a qubit using n others, PRX Quantum **6**, 040368 (2025).
[59] J. Neuser, J. Xuereb, P. Bakhshinezhad, and M. Huber, Quantum vs classical erasure: Equal bounds but unequal costs, arXiv:2607.27341.
[60] R. Silva, P. Bakhshinezhad, and F. Clivaz, Optimal unitary trajectories under commuting target and cost observables; applications to cooling, arXiv:2412.07291.
[61] J. Li, NOISY_HBAC: Simulation code for exploring noisy quantum thermodynamical processes via the depolarizing-channel approximation, GitHub (2026), source code, revision 35a7740, https://github.com/Jian887109/NOISY_HBAC/tree/35a7740d8d64b982c8dd8d2e6cbe230547d44b7a.
[62] L. J. Schulman and U. V. Vazirani, Molecular scale heat engines and scalable quantum computation, in *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, STOC '99 (Association for Computing Machinery, New York, NY, USA, 1999), pp. 322–329.
[63] J. Lin, N. A. Rodríguez-Briones, E. Martín-Martínez, and R. Laflamme, Thermodynamic analysis of algorithmic cooling protocols: Efficiency metrics and improved designs, Phys. Rev. A **110**, 022215 (2024).

End Matter

Application to dynamic cooling—To further analyze the performance of the GDA framework, we consider a different class of algorithms known as single-shot or dynamic cooling (DC) [23,56,62,63]. Unlike the iterative TSAC protocol using one auxiliary reset qubit, DC protocols apply a global unitary transformation only

once across the target qubit and a set of auxiliary qubits to transfer entropy from the target to the auxiliaries.

Specifically, the DC process begins with an initial state given by the tensor product of n identical thermal states, ρ_{th}, that is $\rho_{\mathrm{in}} = \rho_{\mathrm{th}}^{\otimes n}$. A global unitary operator, U_{DC}, is then applied to the entire system. The final, cooled state of the target qubit, ρ_{T}, is obtained by tracing out the $n-1$ auxiliary qubits (labeled here as subsystems 2,3, …, n):

$$\rho_{\mathrm{T}} = \mathrm{Tr}_{2,3,\ldots,n}(U_{\mathrm{DC}}\,\rho_{\mathrm{in}}\,U_{\mathrm{DC}}^{\dagger}). \tag{A1}$$

Several specific unitaries U_{DC} have been proposed [23]. For our analysis, we use the computationally efficient mirror protocol [23].

Numerical verification for DC—We verify the GDA model against a physical noise simulation for the DC mirror protocol. To better resolve performance differences in the high-purity regime and to maintain consistency with related work [23], we use the final effective temperature of the target qubit as the performance metric. This effective temperature, $\mathcal{T}_{\mathrm{eff}}$, is derived from the final diagonal elements $(p_0,\ p_1)$ of the target qubit's density matrix ρ_{T} in the energy basis (the eigenbasis of $H = \frac{1}{2}\hbar\omega\sigma_z$ [23]) via the Boltzmann relation $p_1/p_0 = e^{-\hbar\omega/(k_{\mathrm{B}}\mathcal{T}_{\mathrm{eff}})}$, where k_{B} is the Boltzmann constant.

For these simulations, the qubits are initialized thermal states $\rho_{\mathrm{th}} = e^{-\beta H}/\mathcal{Z}$, where $\mathcal{Z} = \mathrm{Tr}(e^{-\beta H})$ and $\beta = 1/(k_{\mathrm{B}}\mathcal{T}_{\mathrm{initial}})$ [23]. In the numerical simulation, we use

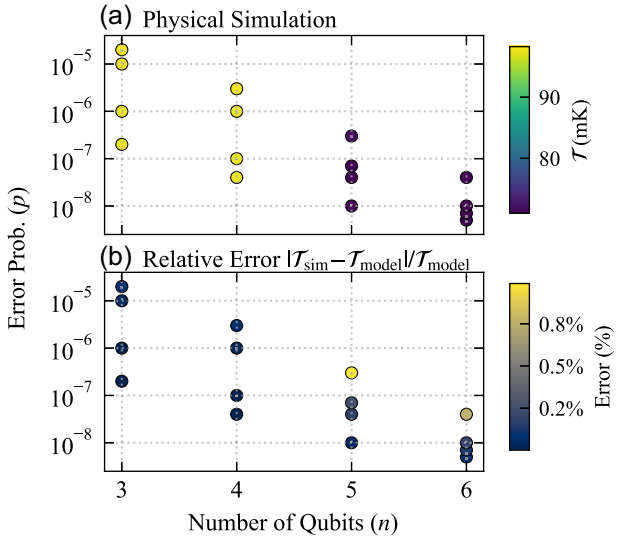


FIG. 3. Verification of the GDA model for the DC mirror protocol across a range of qubit numbers n and CNOT error probabilities p. The panels show (a) the final effective temperature $\mathcal{T}_{\mathrm{sim}}$ obtained from a physical noise simulation; (b) the relative error $|\mathcal{T}_{\mathrm{sim}} - \mathcal{T}_{\mathrm{model}}|/\mathcal{T}_{\mathrm{model}}$, where $\mathcal{T}_{\mathrm{model}}$ is the result from GDA. Temperatures in (a) are shown in mK; (b) is shown as a percentage. The plots demonstrate agreement of the model with the simulation results.

$\omega/2\pi = 10$ GHz and an initial temperature $\mathcal{T}_{\mathrm{initial}} = 163$ mK ($P_{\mathrm{initial}} \approx 0.95$), which is consistent with typical parameters for superconducting qubit quantum computers [23].

Figure 3 shows this final effective temperature (in mK) as a function of the qubit number n and the CNOT error probability p. Figure 3(a) shows the results obtained from the physical noise simulation ($\mathcal{T}_{\mathrm{sim}}$). The final effective temperature predicted by GDA is denoted as $\mathcal{T}_{\mathrm{model}}$. As shown in Fig. 3(b), the relative error of the final temperature $|\mathcal{T}_{\mathrm{sim}} - \mathcal{T}_{\mathrm{model}}|/\mathcal{T}_{\mathrm{model}}$ is at the level of 1%. This result validates our GDA framework's applicability to the DC mirror protocol.

Operational noisy-gate-number window—For a fixed compiled circuit, input state ρ, and observable O, let N be the number of noisy two-qubit gates. We use $\Delta^{\mathrm{out}}_{O,\rho}(N)$ to quantify the observable-level residual of the first-order noise contribution. It compares the expectation value of O obtained from the noisy circuit truncated to include at most one Pauli channel, with the value obtained from the depolarizing channel given in the GDA. Thus, $\Delta^{\mathrm{out}}_{O,\rho}(N)$ measures the finite-depth mixing error for the chosen pair (ρ, O). Its explicit definition is given in SM Sec. II C [41].

For $pN \le 1$, the observable-level finite-depth bound gives

$$\left|\mathrm{Tr}\left[O(\hat{\mathcal{U}}'(\rho) - \hat{\mathcal{U}}^{(1)}_{\mathrm{GDA}}(\rho))\right]\right| \le pN\Delta^{\mathrm{out}}_{O,\rho}(N) + 2\|O\|(pN)^2, \tag{C1}$$

where the first-order GDA output can be written as

$$\hat{\mathcal{U}}^{(1)}_{\mathrm{GDA}}(\rho) = (1-\eta)\hat{\mathcal{U}}(\rho) + \eta\frac{\mathbb{I}}{d}, \quad \eta = pN(1-q). \tag{C2}$$

On the one hand, the first term in the right-hand-side of Eq. (C1) quantifies insufficient finite-depth mixing, which is larger for smaller N. The second term, on the other hand,

TABLE I. Operational noisy-gate-number windows for selected TSAC circuits; "error" is the absolute target-qubit population difference after one TSAC round between the physical simulation and GDA. The last two columns list the operational noisy-gate-count window $[N_{\mathrm{req}}, N_{\mathrm{2nd}})$ obtained with the confidence level $\delta = 0.1$ and the compiled noisy gate count N of the TSAC circuit.

p	n	τ_O	error	$[N_{\mathrm{req}}, N_{\mathrm{2nd}})$	N
10^{-3}	3	0.01	1.5×10^{-3}	[8, 71)	20
10^{-3}	4	0.02	3.3×10^{-3}	[25, 100)	56
2×10^{-4}	5	0.01	4.5×10^{-3}	[20, 354)	148
2×10^{-4}	6	0.02	8.3×10^{-3}	[90, 500)	364

bounds higher-order contributions, which are larger for larger N. Therefore, to ensure that the GDA is accurate up to an observable tolerance τ_O

$$pN\Delta_{O,\rho}^{\mathrm{out}}(N) + 2\|O\|(pN)^2 \leq \tau_O, \tag{C3}$$

we can obtain an operational depth window such that both the first and the second term are small enough.

For the TSAC cooling observable $O = |0\rangle\langle 0|_t \otimes \mathbb{I}$, Table I reports the operational depth windows for different noise strengths p, numbers of qubits n, and observable tolerances τ_O. Using the Hoeffding-based estimate in Eq. (5), together with an additional effective-sampling assumption, we derive a conservative lower gate count N_{req} required for sufficient mixing, and an upper gate count $N_{2\mathrm{nd}}$ set by the higher-order contribution. We see that $N_{\mathrm{req}} < N_{2\mathrm{nd}}$ for all cases shown in the list, confirming the existence of the operational depth window. More details of the table are given in SM Sec. II C [41].