

Cooling a Qubit Using n Others

Jake Xuereb^{1,*}, Benjamin Stratton^{2,3,†}, Alberto Rolandi¹, Jinming He⁴, Marcus Huber^{1,5} and Pharnam Bakhshinezhad¹

¹*Vienna Center for Quantum Science and Technology, Atominstitut, TU Wien, 1020 Vienna, Austria*

²*Quantum Engineering Centre for Doctoral Training, H. H. Wills Physics Laboratory and Department of Electrical and Electronic Engineering, University of Bristol, BS8 1FD, United Kingdom*

³*H.H. Wills Physics Laboratory, University of Bristol, Tyndall Avenue, Bristol BS8 1TL, United Kingdom*

⁴*School of Physical Sciences, University of Science and Technology of China, Hefei 230026, China*

⁵*Institute for Quantum Optics and Quantum Information (IQOQI Vienna), Austrian Academy of Sciences, Boltzmannngasse 3, 1090 Vienna, Austria*



(Received 4 July 2025; revised 1 October 2025; accepted 29 October 2025; published 19 December 2025)

In the task of unitarily cooling a quantum system with access to a larger quantum system, known as the machine or reservoir, *how does the structure of the machine impact an agent's ability to cool and the complexity of their cooling protocol?* Focusing on the task of cooling a single qubit given access to n separable, thermal qubits with arbitrary energy structure, we answer these questions by giving two new perspectives on this task. First, we show that a set of inequalities related to the energetic structure of the n -qubit machine determines the protocol which cools the qubit to the coldest reachable state, which parts of the machine contribute to this protocol and give rise to a Carnot-like bound. Second, we show that cooling protocols can be represented as perfect matchings on bipartite graphs enabling the optimization of cost functions, e.g., gate complexity or dissipation. Our results generalize the algorithmic cooling problem, establish new fundamental bounds on quantum cooling and offer a framework for designing novel autonomous thermal machines and cooling algorithms.

DOI: [10.1103/hrph-dbv7](https://doi.org/10.1103/hrph-dbv7)

I. INTRODUCTION

The task of cooling a quantum system using another [1–8] is a foundational point of inquiry in quantum thermodynamics [9–11]. The problem has been studied in a multitude of settings including both closed- and open-system dynamics and has inspired practical computational subroutines such as algorithmic cooling [12–18] and autonomous quantum refrigerators [19–21]. Both have applications to passive error correction and state preparation, with the latter recently experimentally realized to create an autonomous reset mechanism for a superconducting qubit [22].

In this work, we consider the task of cooling a qubit via a unitary interaction with n thermal qubits which will often be referred to as the *machine* throughout this work.

By cooling, we mean increasing the ground-state population of the system. Here, we go beyond prior work [8,13,14] by allowing the n thermal qubits used for cooling to be nonidentical; the qubits can have varying energy gaps, effectively corresponding to different temperatures. This enables an understanding of how the energetic structure of the n -thermal qubits impacts the cooling scenario. Moreover, it brings the scenario closer to practical reality where physical qubits are unlikely to be identical in real hardware, whether intentionally or by design.

Previous works on cooling qubit systems have, in addition to using identical qubits, allowed the machine qubits to be refreshed between successive rounds of unitary cooling [5,6,17], with asymptotic bounds on the cooling found in this case. Alternatively, fixed machines have been considered but have been either asymptotically large [7,25–27], very small [2,19] and/or used identical qubits [8,13–15]. In this work, we remove the assumptions of identical machine qubits and the ability to rethermalize the machine, focusing on finite machines. By removing the ability to rethermalize qubits, we allow the effect of the machine structure on cooling to be isolated and studied. In addition, we ensure that our cooling protocols can be physically implemented as a single quantum circuit and

*Contact author: jake.xuereb@tuwien.ac.at

†Contact author: ben.stratton@bristol.ac.uk

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

Main Results

When unitarily cooling a qubit with energy gap ω at temperature T_S with access to n qubits at temperature T_M and increasing gaps $\Gamma = (\gamma_1, \gamma_2, \dots, \gamma_n) : E_{\text{Max}} = \sum_{j=1}^n \gamma_j$ the optimal-cooling protocol (maximally increasing the ground-state population of S) is specified by: *A ground-state subspace energy level $|0_{Si_M}\rangle$ should be exchanged with an energy level from the excited-state subspace $|1_{Sj_M}\rangle$, for some $i_M, j_M \in \{0, 1\}^{\times n}$ if*

$$\frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < E(i_M),$$

where $E(i_M)$ is the energy corresponding to the level $|0_{Si_M}\rangle$ given by $i_M \cdot \Gamma$. Intuitively, cooling S depends on the energetic structure of M —the individual energies of which can be ordered. The macroscopic quantities forming the left-hand side of the above inequality determine the midpoint in this order. If an energy falls above this midpoint, a corresponding energy may be found below it such that exchanging these energy levels cools S . Using these inequalities we are able to obtain the following in this work:

Sec. III A single condition to determine the minimal set of energy-level exchanges needed for optimal cooling, leading to a set of inequalities that characterize a cooling scenario.

Sec. III A Carnot-like bound on any energy-level exchange that cools the system.

Sec. IV Conditions determining when sets of machine qubits are energetically irrelevant for cooling, characterizing when a machine is reducible (some machine qubits do not contribute) or irreducible (all the machine contributes).

Sec. IV A framework to optimize cooling unitaries relative to a cost using minimum-weight perfect matching (MWPM) on bipartite graphs [23,24].

examined in terms of their gate complexity, given that they contain no nonunitary steps or asymptotic considerations.

However, even in the case of nonidentical machine qubits, this is a *heuristically* solved problem. It is part of the folklore of quantum thermodynamics that *to cool a system S given access to an auxiliary machine M , one need only take their joint system and machine density and reorder its eigenvalues in decreasing order* [6,28,29]. Surprisingly, little attention has been paid to the fine-grained details of these cooling dynamics, namely *what form does this sorting take and how does this depend on the energetic structure of the machine?* This is despite the

potential for vast differences in the complexity of processes which generate these cooling operations. For instance, the energy-level exchanges of a cooling protocol could simplify down to a SWAP between the system and one of the n machine qubits. Alternatively, it could require an operation controlled on all n machine qubits. Recently, it has been shown that in a single-shot state transformation under unitary dynamics, the energetically optimal trajectory can be identified on a polytope of potential population vectors that the system can take [30]. In this work we address this gap by presenting a framework that not only identifies the necessary energy-level permutations for optimal cooling in terms of a set of physically motivated inequalities, but also evaluates the difficulty of implementing them. Specifically, we establish criteria which we term *reducibility conditions* to detect whether a cooling interaction involves all of the machine or merely a part of it. Shockingly, we find that some cooling scenarios only make use of a single qubit from an n -qubit machine. But our criteria can be used as tools to construct *irreducible* machines the parts of which are always useful. We further explore these dynamics by presenting a method for optimizing over the set of optimal cooling unitaries for some given constraint, here considering the example of minimizing gate complexity.

This paper is organized as follows. In Sec. II we present a simple three-qubit example which exposes the main insights of this work. In Sec. III we introduce a set of inequalities that fully characterize a cooling scenario. In Sec. IV we begin by presenting the reducibility conditions for a cooling scenario and move on to the problem optimizing cooling unitaries relative to a cost function which we solve by using a novel graph representation. We follow with closing discussions and conclude, after which technical matter is provided in the appendices.

II. SETTING AND ILLUSTRATIVE EXAMPLE

What does it mean to cool a qubit? The state of a qubit always has a well-defined temperature. This is due to qubit states always being describable by a Gibbs state with respect to a Hamiltonian $H = E_g |g\rangle\langle g| + E_e |e\rangle\langle e|$ at inverse temperature β ,

$$\tau_\beta = \frac{|g\rangle\langle g| + e^{-\beta(E_e - E_g)} |e\rangle\langle e|}{1 + e^{-\beta(E_e - E_g)}}, \quad \beta = \frac{1}{k_B T}, \quad (1)$$

where T is the temperature and k_B is the Boltzmann constant. The goal of cooling can be rephrased as increasing the ground-state population of the qubit, which effectively lowers T when H is fixed. As such, we may, at points, also refer to qubits with larger energy gaps as being “*colder*,” with this alluding to a lower effective temperature with respect to the qubit with the smallest energy gap. Using




FIG. 1. *Summary illustration.* We examine unitary cooling scenarios in which a qubit at temperature β_S with gap ω is coupled to an n -qubit machine at temperature β_M with energy-gap structure $\Gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$. Since the temperature of a two-level system is determined by the ratio of its ground- and excited-state populations, to cool the qubit system an agent must exchange disordered populations on the global population vector from the ground-state subspace $|0_S i_M\rangle$ to the excited-state subspace $|1_S j_M\rangle$. We find that such exchanges only cool if the energy $E(i_M)$ is larger than $1/2(T_M \omega / T_S + E_{\text{Max}})$, exposing a set of inequalities which determine various properties of the cooling scenario. By finding a representation for cooling unitaries as perfect matchings on bipartite graphs we are able to optimize cooling protocols relative to a cost function by using minimum-weight perfect-matching (MWPM) techniques.

this approach, one can apply the results that we present to qubits in general states by *mathematically* assigning such a Hamiltonian to the individual system and machine qubits. Throughout the work, without loss of generality, we take $\hbar = 1$ and fix the computational basis to be $|g\rangle|g\rangle = |0\rangle|0\rangle, |e\rangle|e\rangle = |1\rangle|1\rangle$. We denote the partition function of τ_β as $\mathcal{Z} = \text{tr}\{e^{-\beta H}\}$ and when working with bit strings use superscripts, e.g., $i^n \in \{0, 1\}^{\times n}$ to denote the length of a string and subscripts, e.g., i_2 in $i_1 i_2 i_3$, to denote the position of the index of a bit.

Before introducing the general case we present an expository example. Consider three qubits, all with $\beta = 1$: a system qubit, $\tau_S = 1/(1 + e^\omega)(|0\rangle|0\rangle + e^{-\omega}|1\rangle|1\rangle)$, and two machine qubits, $\tau_{M_1} = 1/(1 + e^{-\gamma_1})(|0\rangle|0\rangle + e^{-\gamma_1}|1\rangle|1\rangle)$ and $\tau_{M_2} = 1/(1 + e^{-\gamma_2})(|0\rangle|0\rangle + e^{-\gamma_2}|1\rangle|1\rangle)$. We assume that $\omega \leq \gamma_1 \leq \gamma_2$, such that M_2 could be colder than M_1 that could be colder than S . Such an ordering is always possible if one aims to cool the hottest qubit S . We now seek to understand *by how much can S be cooled via a joint unitary on S, M_1 , and M_2 for different M_1 and M_2 , and what operations are required to achieve this cooling?*

The state of the system and the machine can be expressed as a diagonal matrix of eight populations,

$$\begin{aligned} \tau_S \otimes \tau_M &= \sum_{i_S, i_{M_1}, i_{M_2} \in \{0, 1\}} \frac{e^{-(i_S \omega + i_{M_1} \gamma_1 + i_{M_2} \gamma_2)}}{\mathcal{Z}_S \mathcal{Z}_M} |i_S i_{M_1} i_{M_2}\rangle \langle i_S i_{M_1} i_{M_2}| \\ &= \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \text{diag} \left(\underbrace{1, e^{-\gamma_2}, e^{-\gamma_1}, e^{-(\gamma_1 + \gamma_2)}}_{|0_S i_M\rangle}, \right. \\ &\quad \left. \underbrace{e^{-\omega}, e^{-(\omega + \gamma_2)}, e^{-(\omega + \gamma_1)}, e^{-(\omega + \gamma_1 + \gamma_2)}}_{|1_S i_M\rangle} \right), \end{aligned} \quad (2)$$

where $\mathcal{Z}_S = (1 + e^{-\omega})$ and $\mathcal{Z}_M = (1 + e^{-\gamma_1})(1 + e^{-\gamma_2})$. It proves useful to adopt a bit-wise product to denote these populations. Specifically, we introduce the machine gap vector $\Gamma = (\gamma_1, \gamma_2)$, and define its product with the 2-bit string $i_M = i_{M_1} i_{M_2}$ that labels a given basis state, as $i_M \cdot \Gamma = i_{M_1} \gamma_1 + i_{M_2} \gamma_2$ giving the energy associated with this level. With this, we are in a position to set the scene. Initially, the ground-state population of S is $p_0 = \sum_{i_M} (e^{-i_M \cdot \Gamma} / \mathcal{Z}_{SM})$ obtained by summing over entries $|0_S i_M\rangle : i_M \in \{0, 1\}^{\times 2}$ in a partial trace and similarly the initial excited-state population of S is $1 - p_0 = \sum_{i_M} (e^{-(\omega + i_M \cdot \Gamma)} / \mathcal{Z}_{SM})$. To optimally cool S , our aim is therefore to maximize the sum over the entries $|0_S i_M\rangle$ under the action of a joint unitary on S, M_1 and M_2 . Any unitary cooling operation $U(\tau_S \otimes \tau_M)U^\dagger$ must preserve these global populations as they are eigenvalues. Thus, to cool S we can only permute the SM basis kets to ensure that the populations contributing to the ground-state population after cooling (p'_0) are the largest. That is, to perform this optimal cooling we seek a permutation that assigns the four largest populations of SM to the basis states $|0_S i_M\rangle$.

To achieve this, we first note that all permutations of the global state can be broken down into a product of exchanges featuring only two energy levels of SM , which we will often term two-level permutations (TLPs). In particular, we are interested in cooling TLPs which exchange $|0_S i_M\rangle \leftrightarrow |1_S j_M\rangle$ if their associated populations satisfy

$$e^{-i_M \cdot \Gamma} < e^{-(\omega + j_M \cdot \Gamma)}, \quad (3)$$

as the new population of the basis state $|0_S i_M\rangle$ will now be larger. Using Eq. (3) one can find a set of TLPs that are necessary to ensure that the four largest populations are associated with the $|0_S i_M\rangle$ basis states.

Due to the nondecreasing gap structure, $\omega \leq \gamma_1 \leq \gamma_2$, the $|0_S i_M\rangle$ populations in nonincreasing order are $1 \geq e^{-\gamma_1} \geq e^{-\gamma_2} \geq e^{-(\gamma_1+\gamma_2)}$. Given that the $|1_S i_M\rangle$ populations are the $|0_S i_M\rangle$ populations multiplied by the constant $e^{-\omega}$, they too are ordered in the same way. The TLP of the basis kets associated with the largest and smallest populations is therefore cooling if $e^{-(\gamma_1+\gamma_2)} < e^{-\omega}$, which holds if $\omega < \gamma_1 + \gamma_2$. This always holds due to the gap structure, and hence the TLP $|0_S 11\rangle \leftrightarrow |1_S 00\rangle$ is always cooling and leads to the largest amount of cooling. Considering the next-largest population in $|1_S i_M\rangle$ against the next-smallest in $|0_S i_M\rangle$, the associated TLP is cooling if $e^{-\gamma_2} < e^{-(\omega+\gamma_1)}$, which holds if $\omega < \gamma_2 - \gamma_1$. This is not always true and depends on the energetic structure of the machine. The remaining potential exchanges $|0_S 10\rangle \leftrightarrow |1_S 01\rangle$ and $|0_S 00\rangle \leftrightarrow |1_S 11\rangle$ can never cool as $|0_S 00\rangle$ and $|0_S 01\rangle$ already hold two of the four largest populations in the state. As we have verified, $|1_S 00\rangle$ holds one of the four largest, and the TLP $|0_S 11\rangle \leftrightarrow |1_S 00\rangle$ always ensures its correct placement. This leaves just one final population, with two possible scenarios: (a) if $\omega \geq \gamma_2 - \gamma_1$, $|0_S 01\rangle$ holds the fourth-largest population already and (b) if $\omega < \gamma_2 - \gamma_1$, $|1_S 10\rangle$ holds the fourth-largest population and the TLP $|0_S 10\rangle \leftrightarrow |1_S 01\rangle$ is needed.

In scenario (a) the only TLP that must be applied is $|0_S 11\rangle \leftrightarrow |1_S 00\rangle$. A unitary $S_{011,100} = |0_S 11\rangle\langle 1_S 00| + |1_S 00\rangle\langle 0_S 11| + \mathbb{1}_{\text{Rest}}$ optimally cools S and can be broken down into a series of `Toffoli` gates as in Fig. 2. The final ground-state population of the system is then $p_0^{(a)} = 1/\mathcal{Z}_S \mathcal{Z}_M (1 + e^{-\gamma_2} + e^{-\gamma_1} + e^{-\omega})$. It can be seen that $p_0^{(a)}$ is larger than the initial ground-state population of M_2 , the coldest qubit, if $p_0^{(a)} > 1/(1 + e^{-\gamma_2})$ which occurs when $\omega > \gamma_2 - \gamma_1$, which is the criterion for scenario (a). In scenario (b) both $|0_S 11\rangle \leftrightarrow |1_S 00\rangle$ and $|0_S 01\rangle \leftrightarrow |1_S 10\rangle$ must be implemented. With both conditions fulfilled, i_{M_1} becomes a free index such that all exchanges $|0_S i_{M_1} 1\rangle \leftrightarrow |1_S i_{M_1} 0\rangle : i_{M_1} \in \{0, 1\}$ must be carried out, which is achieved by SWAP_{S, M_2} . The final achievable ground-state population after cooling is therefore implicitly that of M_2 , i.e., $p_0^{(a)} = 1/(1 + e^{-\gamma_2})$.

This is a curious outcome. In (b), where the energetic gap of S is smaller than the difference between the energy gaps of the machine qubits, M_1 does not contribute to the optimal cooling. The operation for optimal cooling is purely bipartite, with one machine qubit left unused. On the other hand, in (a), where the energetic gap of S is larger than the difference between the energy gaps of the machine qubits, an agent is able to cool further than the temperature of the coldest qubit. However, the dynamics are necessarily tripartite. We can better understand these two cooling scenarios by visualizing them as vertex-weighted graphs. Let each eigenket $|i_S i_{M_1} i_{M_2}\rangle$ correspond to a vertex weighted by $e^{-\beta(i_S \omega + i_M \cdot \Gamma)}$ and let two vertices v_i and v_j be connected by an edge if they correspond to bit strings




FIG. 2. (a) Three qubits at temperature $\beta = 1$ and with gaps $\omega = \log(2)$, $\gamma_1 = \log(3)$ and $\gamma_2 = \log(8)$, with the populations factored by $\frac{4}{9}$ for ease of readability. In orange we see a plane denoting the axis of reflection corresponding to a `SWAP` between the system qubit and the machine qubit 2 (bipartite) which results in the optimal distribution. (b) Three qubits at temperature $\beta = 1$ but with gaps $\omega = \log(2)$ and $\gamma_1 = \log(5)$ and $\gamma_2 = \log(8)$ giving one less disordered pair, with populations factored by $\frac{81}{40}$. Here, the edges in orange correspond to a series of two-level swaps or `Toffolis` in the order $|1_S 00\rangle \leftrightarrow |1_S 10\rangle$ then $|0_S 10\rangle \leftrightarrow |0_S 11\rangle$ and $|0_S 10\rangle \leftrightarrow |1_S 10\rangle$ (tripartite) that give the optimal distribution.

$i_S i_{M_1} i_{M_2}$ and $j_S j_{M_1} j_{M_2}$ that have Hamming distance 1, i.e., differ by one bit. This is the weighted 3-bit hypercube graph Q_3 —the graph of bit strings of length 3 with edges between bit strings differing by 1 bit (Hamming distance 1) [23]—and permutations of the SM eigenkets for cooling can be thought of as operations on Q_3 . An example of each of the two cooling scenarios described is represented on Q_3 in Fig. 2. Even cooling a single qubit with access to two other qubits seems to lead to an intricate set of scenarios, which while complex is seemingly captured by a simple set of relationships between the energy gaps of the system and machine.

III. A SET OF INEQUALITIES

A. Finding the Disordered Populations

Turning our attention to the general problem, imagine a qubit with gap ω at temperature β_S as our system S and n qubits at temperature β_M with an energetic gap structure $\Gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$ such that $\gamma_i \leq \gamma_{i+1} \forall i \in \{1, N\}$, as the machine M in this cooling scenario. The system and machine have the initial joint state

$$\tau_S \otimes \tau_M = \sum_{\substack{i_S \in \{0,1\} \\ i_M \in \{0,1\}^{\times n}}} \frac{e^{-(\beta_S i_S \omega + \beta_M i_M \cdot \Gamma)}}{\mathcal{Z}_S \mathcal{Z}_M} |i_S i_M\rangle\langle i_S i_M|, \quad (4)$$

where $\mathcal{Z}_S = (1 + e^{-\beta_S \omega})$ and $\mathcal{Z}_M = \prod_{j=1}^n (1 + e^{-\beta_M \gamma_j})$. As before, any TLP $|0_S i_M\rangle \leftrightarrow |1_S j_M\rangle$ on $\tau_S \otimes \tau_M$ that an agent performs is cooling on S if their associated populations satisfy $e^{-\beta_M i_M \cdot \Gamma} < e^{-(\beta_S \omega + \beta_M j_M \cdot \Gamma)}$. This implies that any pair of energy eigenkets in the SM Hilbert space that

cool down S when permuted must satisfy

$$\frac{T_M}{T_S}\omega < (i_M - j_M) \cdot \Gamma; \quad (5)$$

i.e., the difference in energy associated with this exchange must be greater than $(T_M/T_S)\omega$. In this case we say that the populations of the eigenkets $|0_S i_M\rangle$ and $|1_S j_M\rangle$ are disordered. To find the TLPs needed for performing optimal cooling, one can then consider the strategy of iteratively swapping the largest populations $|1_S j_M\rangle$ with the smallest populations $|0_S i_M\rangle$, as in the example above. Conveniently, in the general case this can be achieved by always considering the TLPs $|0_S i_M\rangle \leftrightarrow |1_S i_M \oplus 1\rangle$, where $\oplus 1$ is a bit-wise addition mod 2 (for the proof, see Appendix A). Using Eq. (5), a given TLP is therefore necessary for optimal cooling if

$$\frac{T_M}{T_S}\omega < (i_M - (i_M \oplus 1)) \cdot \Gamma, \quad (6)$$

which, given that $(i_M \oplus 1) \cdot \Gamma = \sum_{l=1}^n (i_l \oplus 1)\gamma_l = \sum_{l=1}^n \gamma_l - \sum_{l=1}^n i_l \gamma_l = E_{\text{Max}} - (i_M \cdot \Gamma)$, where $E_{\text{Max}} = \sum_{l=1}^n \gamma_l$, simplifies to

$$\frac{1}{2} \left(\frac{T_M}{T_S}\omega + E_{\text{Max}} \right) < i_M \cdot \Gamma. \quad (7)$$

This is our first main result, an inequality that detects every disordered population that must be exchanged to cool S in terms of whether a change in energy is greater than the product of the ratio of temperatures of system and machine and the energy gap of S . The quantity $\frac{1}{2}((T_M/T_S)\omega + E_{\text{Max}})$ will appear several times in this paper and appears to be a constant that dictates many properties of the cooling scenario such as the order of interaction required to cool and how much a given machine is able to cool.

We can note that if the energy level i_M of the machine has an energy $i_M \cdot \Gamma$ above the threshold $\frac{1}{2}((T_M/T_S)\omega + E_{\text{Max}})$ and another energy level l_M has energy larger than i_M (i.e., $l_M \cdot \Gamma > i_M \cdot \Gamma$), then $|0_S l_M\rangle$ should also be exchanged for cooling if $|0_S i_M\rangle$ is to be exchanged. In fact, considering the energies of the machine we have the largest energy $E_{\text{Max}} = 1^n \cdot \Gamma$, smallest energy $0 = 0^n \cdot \Gamma$ and every energy in between given by $i_M \cdot \Gamma$ for $i_M \in \{0, 1\}^n$. Since $(i_M \oplus 1) \cdot \Gamma = E_{\text{Max}} - i_M \cdot \Gamma$, if $i_M \cdot \Gamma$ is above $E_{\text{Max}}/2$ then $(i_M \oplus 1) \cdot \Gamma$ is below $E_{\text{Max}}/2$. Therefore we find that $E_{\text{Max}}/2$ is the median of the energies. This allows us to reinterpret Eq. (5) as: *if the energy $i_M \cdot \Gamma$ is above the shifted median value $\frac{1}{2}((T_M/T_S)\omega + E_{\text{Max}})$ then there exists at least one energy $j_M \cdot \Gamma$ below the median such that exchanging $|0_S i_M\rangle$ and $|1_S j_M\rangle$ cools S* (for more details and visualizations, see Appendix A).

B. Carnot-Like Bound

Manipulating Eq. (5) we surprisingly find that the change in effective energy gap of the system due to a single energy-level exchange $|0_S i_M\rangle \leftrightarrow |1_S j_M\rangle$ in the system-machine Hilbert space is upper bounded by the Carnot efficiency of the system and machine $\eta_c = 1 - T_M/T_C$. To expose this fundamental bound, divide Eq. (5) by ω , multiply each side by -1 and add 1 to both sides, giving

$$\frac{\Delta E_S}{E_S} = \frac{\omega - (i_M - j_M) \cdot \Gamma}{\omega} < 1 - \frac{T_M}{T_S}, \quad (8)$$

where $\Delta E_S = \omega - (i_M - j_M) \cdot \Gamma$ and $E_S = 0 + \omega$. This bound suggests that the change in effective energy gap of the system induced by a global-energy ket exchange is upper bounded by the Carnot efficiency of SM which is independent of the energetic structure Γ of the n -qubit machine. A conventional Carnot efficiency bound cannot be obtained in this scenario as the reduced state of the machine after the cooling interaction will no longer nor thermal with respect to the machine Hamiltonian. It is interesting to see that despite this the Carnot efficiency still bounds the population changes that a single TLP can induce. This bound is saturated when $(i_M - j_M) \cdot \Gamma = (T_M/T_S)\omega$, which corresponds to energy-level exchange that results in the smallest change in energy that still cools S .

C. How Much Can a Machine Cool?

Armed with our bit-wise notation and hierarchy of inequalities we can investigate how much a given machine can cool our system qubit. The figure of merit we are interested in is the optimal change in the ground-state population of S , $\Delta p_0 = p'_0 - p_0$, where p'_0 is the ground-state population after cooling. Examining Δp_0 we begin by noting that the ground-state population after cooling can be expressed as a sum of populations of unchanged energy levels and swapped energy levels,

$$p'_0 = \underbrace{\sum_{i_M \in \{0,1\}^n \setminus \mathbb{S}} \frac{e^{-\beta_M i_M \cdot \Gamma}}{\mathcal{Z}_S \mathcal{Z}_M}}_{\text{unchanged}} + \underbrace{\sum_{j_M \in \mathbb{S}} \frac{e^{-\beta_S \omega - \beta_M (j_M \oplus 1) \cdot \Gamma}}{\mathcal{Z}_S \mathcal{Z}_M}}_{\text{swapped}}, \quad (9)$$

where $\mathbb{S}$ is the set of n -bit strings satisfying Eq. (7). The initial ground-state population can be expressed as $p_0 = 1/\mathcal{Z}_S$ and so when subtracting it from p'_0 acts as $\mathcal{Z}_M$ when taking the common denominator $\mathcal{Z}_S \mathcal{Z}_M$. Then, it can be noted that $\mathcal{Z}_M$ can be expanded to be written as a sum over all n -bit strings $\mathcal{Z}_M = \sum_{i_M \in \{0,1\}^n} e^{-\beta_M i_M \cdot \Gamma}$. From here, it can be seen that the p'_0 contribution in Eq. (9) is nothing more than a sum of two disjoint sets the union of which is $\{0, 1\}^n$. Splitting $\mathcal{Z}_M$ into a contribution from $\mathbb{S}$ and its set complement, the unchanged populations in Eq. (9) cancel




FIG. 3. A comparison of the changes in the ground-state population of S , Δp_0 , induced by the cooling protocol for different machine energy-gap structures as the number of qubits in the machine n increases: (a) linear gaps; (b) exponential gaps; (c) degenerate gaps. We also compare the lower bounds derived to approximate the change in population: in green we see the performance of the bound given in Eq. (11) and in blue we see the bound calculated algorithmically using the partial order presented in Appendix I. We see that while the fixed bound in Eq. (11) performs well for machines with energetic structures that grow quickly it performs poorly for machines with fixed gaps. This is because Eq. (11) captures a fixed set of strings to be exchanged, which in the linear and exponential cases seems to capture most exchanges even as n grows, but in the degenerate case we see that this fixed set of strings contributes less and less to the growing number of strings which need to be exchanged as n grows. In contrast, the blue bound is adaptively constructed via the partial order and so performs better in all scenarios.

to give a sum over $\mathbb{S}$,

$$\Delta p_0 = \sum_{jM \in \mathbb{S}} \frac{e^{-\beta_S \omega - \beta_M (E_{\text{Max}} - jM \cdot \Gamma)} - e^{-\beta_M jM \cdot \Gamma}}{\mathcal{Z}_S \mathcal{Z}_M}, \quad (10)$$

when simplified. This quantity can always be found if the set $\mathbb{S}$ is known, as discussed below for specific examples. If $\mathbb{S}$ is not completely known, but instead a subset of $\mathbb{S}$ is known, a lower bound can be at least be placed on Δp_0 . For example, take the case in which $T_M \omega / T_S < \gamma_1$, and then consider the bit string $b_M = 10^{n-k-1} 1^k$, where $k = n/2$ if n is even and $k = (n-1)/2$ if n is odd. Given the assumption that Γ is ordered in nondecreasing order, this is a bit string such that all bit strings including and above this (in terms of the lexicographic order, cf. Appendix I) are guaranteed to satisfy Eq. (7) for all Γ . By considering performing the swaps associated with just these bit strings, one finds that

$$\Delta p_0 \geq \frac{1}{\mathcal{Z}_S \mathcal{Z}_M^{k,n}} \left(e^{-\beta_S \omega} - e^{-\beta_M E(\tilde{k})} \right) + \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(e^{-\beta_S E(\tilde{k})} - e^{-\beta_S \omega - \beta_M (E_{\text{Max}} - E(\tilde{k}))} \right), \quad (11)$$

where $E(\tilde{k}) = \sum_{i=\tilde{k}}^n \gamma_i$, with $\tilde{k} = n/2 + 1$ if n is even and $\tilde{k} = (n+3)/2$ if n is odd, and $\mathcal{Z}_M^{k,n} = \prod_{i=\tilde{k}}^n (1 + e^{-\beta_M \gamma_i})$ is the partition function of the $n - \tilde{k}$ machine qubits with the largest energy gaps (for more details, see Appendix C). From this example, we can further this line

of thought by considering what is the largest subset of $\mathbb{S}$ that always satisfies Eq. (7). We find that this set can be characterized with the help of the lexicographic partial order of bit strings (cf. Appendix I). Depending on the value of ω it consists of all the bit strings larger than $a_M = 1^{2\ell-n} (01)^{n-\ell}$ when $\sum_{i=1}^{2(\ell-1)-n} \gamma_i \leq T_M \omega / T_S < \sum_{i=1}^{2\ell-n} \gamma_i$ (with $\ell = \lceil (n+1)/2 \rceil, \dots, n$). As opposed to the previous bound with b_M we cannot obtain a formula as explicit as Eq. (11) for the set composed of the elements larger than a_M . However, since the set is fully characterizable as a list of bit strings, we compare both of the bounds in Fig. 3 (green and blue, respectively) against the actual value of Δp_0 calculated numerically.

IV. COOLING UNITARIES: REDUCIBILITY AND MATCHING ON GRAPHS

Having developed a framework to detect which populations must be exchanged in the joint system-machine state to optimally cool S , what remains to be considered is *how*. As we have seen, optimally cooling S given M is simple; *find the machine energy levels $i_M \in \{0, 1\}^n$ satisfying $\frac{1}{2}((T_M/T_S)\omega + E_{\text{Max}}) < i_M \cdot \Gamma$ and then perform a set of TLPs to exchange all $|0_{Si_M}\rangle \leftrightarrow |1_{Si_M} \oplus 1\rangle$* . However, denoting the sets of energy levels to be swapped as $|A\rangle = \{|0_{Si_M}\rangle \in \mathbb{S}\}$ and $|B\rangle = \{|1_{Si_M} \oplus 1\rangle\}$, it can be seen that we have the freedom to pair each $|0_{Si_M}\rangle \in |A\rangle$ with any $|1_{Sj_M}\rangle \in |B\rangle$ and obtain the same increase in ground-state population, Δp_0 . There are therefore at least $|\mathbb{S}|!$ unitaries of the form $U_{\text{cool}} = |A\rangle\langle B| + |B\rangle\langle A| + \mathbb{1}_{\text{Rest}}$ that

can cool S to the same temperature [31]. But, these unitaries, while all valid for cooling, are far from equivalent in implementation or thermodynamic cost. In this section we will begin by examining the order of interaction necessary to carry out U_{cool} for different machines and then move on to represent the problem of choosing a U_{cool} under some cost constraints via a minimum-weight perfect matching (MWPM) on bipartite graphs [23].

A. Reducibility of Machines

Within the setting we have described it is perhaps natural to think that increasing the number of qubits in the machine is always helpful. And indeed it is, as increasing the number of qubits in the machine increases E_{Max} . However, one can ask whether it is only the increased energy that improves the cooling, or also the extra dimensions in the Hilbert space. In other words, we must investigate whether a cooling operation involving S and an n -qubit machine M is ever reducible to an operation involving just some $k < n$ coldest qubits of the machine—in which case, the energetic gap structure Γ extended to n qubits appears dimensionally *wasteful* and could be reduced to a machine with fewer qubits involving only a subset of Γ . To begin this investigation, let us look at the most egregiously wasteful kind of machine, an n -qubit machine that is reducible to a bipartite interaction involving the coldest qubit in the machine and S . Such a machine can be said to be $(n-1)$ -reducible, meaning that its warmest $n-1$ qubits do not participate in the cooling process. Mathematically, this means that U_{cool} acts as the identity on these $n-1$ qubits, leaving them as free indices. This would mean that the set of energy levels to be exchanged can be expressed as $\mathbb{S} = \{|0_S i^{n-1} 1\rangle \mid i^{n-1} \in \{0, 1\}^{\times n-1}\}$ and so every $|0_S i^{n-1} 1\rangle$ satisfies Eq. (7), such that

$$\frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < i^{n-1} 1 \cdot \Gamma \quad \forall i^{n-1} \in \{0, 1\}^{\times n-1}. \quad (12)$$

Most stringently, this must also be valid for $i^{n-1} = 0^{n-1}$ where only the qubit with gap γ_n contributes to cooling. For this string, the above condition becomes

$$\frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < \gamma_n, \quad (13)$$

i.e., *a cooling scenario involving a machine the largest gap of which is larger than half the sum of all the gaps in the machine and the constant $T_M/T_S \omega$ is reducible to a bipartite interaction* (for more details, see Appendix D). This is an unexpected outcome where we see that if the functional form of Γ grows steeply enough with n , then all that matters for cooling is the energetic gap of the largest qubit, and rest of the machine is *useless*. While

surprising, we have already encountered this situation in our introductory example in Sec. II when only one of two qubits in the two-qubit machine was involved in the cooling operation when $\omega < \gamma_2 - \gamma_1$. A quick calculation shows this is precisely the $(n-1)$ -reducibility condition in Eq. (13) for $T_M = T_S$ with a two-qubit machine. While this condition is quite restrictive, it arises naturally in one of the examples we have considered, exponential machines. An n -qubit machine with an exponential gap structure is one the gap structure of which grows exponentially with n such that $\Gamma = (\gamma, \gamma^2, \dots, \gamma^n)$. Applying Eq. (13), we find that a sufficient condition for an exponential machine to be $(n-1)$ -reducible is $\gamma > (T_M/T_S)\omega$ with $\gamma > 1$ (cf. Appendix D 2).

Following this line of reasoning, one can find conditions for a machine to be k -reducible by examining under which conditions all the energy levels $|0_S i^k 1^{n-k}\rangle$ need to be exchanged such that the k bits in i^k are free indices, so that the unitary acts as the identity on the k qubits corresponding to these k bits. This situation is not straightforward, as one must consider $|0_S i^k 1^{n-k}\rangle$ for every possible choice of $i^{n-k} \in \{0, 1\}^{\times n-k}$. Consider Fig. 4, which visualizes the structure of the partially ordered set of inequalities for a three-qubit machine. This poset has three levels corresponding to the machine being 2-reducible, meaning that the optimal-cooling interaction involves only the coldest qubit in the machine, 1-reducible, involving the coldest two qubits in M , and irreducible, involving all the qubits in M . If the machine is 2-reducible, i.e., $|0_S i_1 i_2 1\rangle$ satisfies Eq. (7) for any $i_1, i_2 \in \{0, 1\}$, then every energy level in the hierarchy is also a member of $\mathbb{S}$ and must be exchanged. Ascending one level of the hierarchy we find two scenarios, only one of which leads to 1-reducibility. It must be the case that $|0_S i_1 11\rangle \in \mathbb{S}$ while $|0_S i_1 01\rangle \notin \mathbb{S}$ for $i_1 \in \{0, 1\}$, preventing 2-reducibility. If $|0_S i_1 01\rangle \in \mathbb{S}$, then we also have $|0_S i_1 11\rangle \in \mathbb{S}$, leading again to 2-reducibility. At the final level, two scenarios lead to irreducibility. First, in situation (i) in Fig. 4 we have $|0_S 001\rangle \notin \mathbb{S}$ and $|0_S 101\rangle \in \mathbb{S}$, also implying that $|0_S 011\rangle, |0_S 111\rangle \in \mathbb{S}$. While this gives $|0_S i_1 11\rangle \in \mathbb{S}$ for $i_1 \in \{0, 1\}$, we have $|0_S 101\rangle \in \mathbb{S}$ but $|0_S 001\rangle \notin \mathbb{S}$ and so any exchange involving $|0_S 101\rangle$ is 4-partite. Second, the machine is irreducible if $|0_S 011\rangle \notin \mathbb{S}$ and $|0_S 111\rangle \in \mathbb{S}$, condition (ii) in Fig. 4. The machine must be irreducible in this case as any exchange involving $|0_S 111\rangle \in \mathbb{S}$ cannot involve $|0_S 011\rangle \notin \mathbb{S}$, i.e., i_1 cannot be a free index. Any TLP that cools scenario (ii) acting on $|0_S 111\rangle$ necessarily acts on all qubits in the cooling scenario. This begins to elucidate the complex hierarchical structure of the energetic inequalities we have explored in this paper. In Appendix D we give a visualization for the four-qubit machine hierarchy of inequalities and formal definitions for k -reducibility in general. The general structure of the partial order of these inequalities is explored in Appendix I.




FIG. 4. A diagram of the reducibility inequalities for a three-qubit machine, showing the partial order in the inequalities, where $|0s_{iM}\rangle \implies |0s_{jM}\rangle$ means that if $|0s_{iM}\rangle \in \mathbb{S}$, then so is $|0s_{jM}\rangle \in \mathbb{S}$ —although the reverse implication does not necessarily hold.

B. Irreducible Machines

While general k -reducibility conditions are difficult to obtain, we are able to fully characterize scenarios under which a machine is irreducible. This is arguably the most important reducibility criterion because it teaches us how to design *useful* machines where adding more qubits always leads to more potential for cooling. Such machines exemplify the third law of thermodynamics [7] as an increase in the number of constituents of the machine and thus its dimension leads to further cooling.

An irreducible machine must use every qubit for cooling, even its warmest qubit. This implies that i_1 cannot be a free index in U_{cool} , meaning that there must exist at least one $l^{n-1} \in \{0, 1\}^{\times n-1}$ such that $|0_S 1^{l^{n-1}}\rangle \in \mathbb{S}$ while $|0_S 0^{l^{n-1}}\rangle \notin \mathbb{S}$. If not true, then the joint action of all these permutations renders the first qubit of the machine useless, making the machine reducible. The first condition $|0_S 1^{l^{n-1}}\rangle \in \mathbb{S}$ imposes the upper bound $\frac{1}{2}(\frac{T_M}{T_S}\omega + E_{\text{Max}}) < 1^{l^{n-1}} \cdot \Gamma$, while the second condition $|0_S 0^{l^{n-1}}\rangle \notin \mathbb{S}$ imposes the lower bound $0^{l^{n-1}} \cdot \Gamma \leq \frac{1}{2}(\frac{T_M}{T_S}\omega + E_{\text{Max}})$. Taken together, we come to the condition that *an n -qubit machine with gap structure Γ at temperature T_M acts irreducibly when cooling a qubit at temperature T_S with gap ω if there exists at least one $l^{n-1} \in \{0, 1\}^{\times n-1}$ such that*

$$0 \leq \frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) - l^{n-1} \cdot \Gamma_{2:n} < \gamma_1, \quad (14)$$

where $\Gamma_{k:n} = (\gamma_k, \gamma_{k+1}, \dots, \gamma_n)$ such that $E(l^{n-1}) = l^{n-1}$. $\Gamma_{2:n}$ can be considered an energy of the $n-1$ coldest qubits. Intuitively, one can understand this inequality to mean that there is at least one energy level l^{n-1} of the $n-1$ coldest machine qubits which is only large enough to optimally cool if γ_1 is included—exchanging l^{n-1} without a contribution from the warmest qubit would not cool S . It is worth noting that is clear from Eq. (14) that by investing work and effectively cooling S by increasing ω and keeping Γ constant, one can make any irreducible

machine reducible by trading unitary complexity for thermodynamic work. We investigate this inequality in further depth in Appendix D.

C. Degenerate Machines Are Always Irreducible

A first application of Eq. (14) is to show when a machine of identical qubits, as typically considered in algorithmic and dynamical cooling scenarios [8,13,14,25], is irreducible. For a machine of n qubits at temperature T_M , each with gap γ , note that the energy of any level i_M in the machine can be expressed as $E(i_M) = t\gamma$ where $t = \#(i_M)$ is the number of 1s in i_M . The irreducibility condition in Eq. (14) then takes the form

$$-(n-2t)\gamma \leq \frac{T_M}{T_S}\omega \leq 2\gamma - (n-2t)\gamma, \quad (15)$$

where $E_{\text{Max}} = n\gamma$ and $t \in [0, n - 1]$. Each t gives a condition that is some interval on the real numbers into which $(T_M/T_S)\omega$ must fall. In the degenerate case, these intervals capture the whole number line $[0, n\gamma]$, meaning that whenever $(T_M/T_S)\omega \leq n\gamma$, one can always find a t such that the above condition is satisfied $\square$. Under the reasonable assumptions that $T_M \leq T_S$ and $\omega \leq \gamma$, degenerate machines are therefore always irreducible (for more details, see Appendix D 1 a).

D. Cost-Optimized Cooling Unitaries from Minimum-Weight Perfect Matchings

In the previous section, we have considered what order of interaction is necessary for optimal cooling given the set $\mathbb{S}$. Here, we present a framework to identify which of the $|\mathbb{S}|!$ unitaries of the form $U_{\text{cool}} = |A\rangle\langle B| + |B\rangle\langle A| + \mathbb{1}_{\text{Rest}}$ should be used to implement the cooling protocol given some additional constraint. An agent may wish to optimize different costs of their cooling protocol such as heat dissipated into their machine, work cost invested to implement U_{cool} or the gate complexity of the protocol. While we focus on gate complexity, our methods can be applied to any such cost function. A natural way to characterize implementations of U_{cool} is in terms of a decomposition of the unitary into a product of quantum gates. To do so, one needs to consider a gate set into which a given unitary can be decomposed. We focus on the set $\mathcal{G} = \{X, \text{CNOT}, \text{Toffoli}, C^3X, \dots, C^{n-1}X\}$, where C is used to denote control and CNOT is the controlled not gate, which suffices since the unitaries we are interested in generating are permutations of the energy eigenbasis which, without loss of generality, have been chosen to be in the Z basis. A figure of merit for the complexity of U_{cool} , related to the classical notion of time complexity, is then the *gate complexity*. This the minimum number of gates from $\mathcal{G}$ that are required to implement U_{cool} .

E. Hamming-Weight Lower Bounds Gate Complexity

Proof. Let $S_{j,k}$ denote the TLP unitary that exchanges the n -qubit eigenket $|j\rangle$ for the eigenket $|k\rangle$ expressed as $S_{j,k} = |j\rangle\langle k| + |k\rangle\langle j| + \mathbb{1}_{\text{Rest}}$. For the gate set $\mathcal{G}$, we can show that the minimum number of gates required to implement $S_{j,k}$ using $\mathcal{G}$, denoted $\mathcal{C}_{\mathcal{G}}(S_{j,k})$, is lower bounded by

$$\mathcal{C}_{\mathcal{G}}(S_{j,k}) \geq D_H(j, k), \quad (16)$$

where $D_H(j, k)$ is the Hamming distance between the bit strings j and k . To prove this statement, first note that $D(|i\rangle, U_g |i\rangle) \leq 1 \forall U_g \in \mathcal{G}$ and $|i\rangle$, where $D(|i\rangle, |j\rangle) = D(i, j)$. This is due to each element of $\mathcal{G}$ only changing the Hamming weight of a basis state, that is the number of 1s in its identifying bit string, by at most one. This then leads to a change in distance between the input and output bit string of at most one. Now consider some $S_{j,k}$, where $D(j, k) = l$. This can be decomposed as $S_{j,k} = U_1 U_2 \dots U_l$, where $U_i \in \mathcal{G} \forall i$. Given that $S_{j,k} |j\rangle = |k\rangle$, it must be the case that $l \geq D(j, k)$, as each U_i can only change the distance by at most one and since $l = \mathcal{C}_{\mathcal{G}}(S_{j,k})$ by definition, the statement holds. $\square$

F. Minimum-Weight Perfect Matching

The above result implies that one could find a unitary that both optimally cools and potentially minimizes the gate complexity of U_{cool} by choosing a set of TLPs that satisfy the energetic inequalities but also minimize the Hamming distance between pairs $|0_{SiM}\rangle \in |A\rangle$ and $|1_{SjM}\rangle \in |B\rangle$. We are therefore interested in finding the unitary U_{cool}^* such that

$$U_{\text{cool}}^* = \sum_{(|a\rangle, |b\rangle) \in \mathcal{M}} |a\rangle\langle b| + |b\rangle\langle a| + \mathbb{1}_{\text{Rest}} \quad \text{with} \quad (17)$$

$$\mathcal{M} = \left\{ (|a\rangle, |b\rangle) \in |A\rangle \times |B\rangle \mid \min_{\substack{|a'\rangle \in |A\rangle \\ |b'\rangle \in |B\rangle}} \sum_{\substack{|a'\rangle \in |A\rangle \\ |b'\rangle \in |B\rangle}} d_H(a, b') \right\},$$

where $\mathcal{M}$ is a set of 2-tuples or pairs of bit strings labeling energy levels the Hamming distances of which sum to give the minimum over all possible pairings. Note that this pairing is the opposite of the strategy that we employed to *identify* the smallest set of TLPs necessary for optimal cooling. There, we found the TLPs that maximized the Hamming-weight distance. This emphasizes the point that once we have identified the necessary TLP using that strategy, we are free to implement any of the $|\mathbb{S}|!$ alternative pairings which all achieve optimal cooling but perform differently relative to different costs.

To find U_{cool}^* , consider the complete bipartite graph $G(K, K \oplus 1, E)$, where each node on the left subgraph K represents a bit string in $\mathbb{S}$, and each node on the right subgraph $K \oplus 1$ represents a bit string in $\mathbb{S} \oplus 1$. Each node




FIG. 5. An example of a complete bipartite graph $G(K, K \oplus 1, E)$ for a cooling scenario involving a three-qubit machine as detailed in Appendix B (case 4). Each perfect matching in this bipartite graph corresponds to a choice of U_{cool} that will lead to the same final ground-state population from the system. The MWPM on this graph (orange) corresponds to the unitary U_{cool}^* that minimizes the Hamming weight across pairings as seen using the cost matrix in Eq. (18).

in K is then connected to every other node in $K \oplus 1$ such that an edge represents a potential energy-level exchange or TLP. (for full details on constructing this graph, see Appendix H). A matching M on such a graph is a collection of edges of the graph where each vertex which intersects with the collection intersects precisely with one edge: a matching is said to be perfect if each vertex in a graph intersects with it. That is, a perfect matching connects every vertex in a graph to exactly one other vertex in the graph. We visualize such a complete bipartite graph for $|\mathbb{S}| = 3$ in Fig. 5 for the example of a three-qubit machine, as detailed in Appendix B (case 4). It can be shown that by definition a perfect matching on $G(K, K \oplus 1, E)$ corresponds to a suitable U_{cool} as it pairs each disordered ground-state energy level $|0_{SiM}\rangle \in \mathbb{S}$ to exactly one $|1_{SjM}\rangle$ which need be exchanged for cooling, exhausting all of $\mathbb{S}$ (for the proof, see Appendix H). If we add weights to the edges of a graph, then a *minimum-weight perfect matching* (MWPM) is the perfect matching with pairings that minimize the sum of the edge weights in the perfect matching. By weighting each edge in $G(K, K \oplus 1, E)$ with the Hamming distance between the bit strings labeling the nodes which intersect it, we find that a *minimum-weight perfect matching on the graph corresponds to a U_{cool}^** . In achieving this correspondence we reformulate the task of finding an optimal-cooling unitary subject to some additional constraint as that of finding minimum-weight perfect matchings on complete bipartite graphs. A detailed proof of this correspondence is given in Appendix H.

Finding MWPMs is possible through the *Hungarian algorithm* [24]. The algorithm begins by assigning a cost to every possible pairing and creating a matrix C of these costs. In this case, this matrix can be created by assigning $a_j \in |A\rangle$ to the columns of the matrix and $b_i \in |B\rangle$ to the rows of the matrix and computing the entries of the matrix

as $c_{ij} = D_H(b_i, a_j)$. For the example in Fig. 5, this gives

$$C = \begin{matrix} & |0011\rangle & |0101\rangle & |0111\rangle \\ \begin{matrix} |1100\rangle \\ |1010\rangle \\ |1000\rangle \end{matrix} & \begin{pmatrix} 4 & \textcolor{red}{2} & 3 \\ \textcolor{red}{2} & 4 & 3 \\ 3 & 3 & \textcolor{red}{4} \end{pmatrix} \end{matrix}. \quad (18)$$

From C , a minimum-weight perfect matching corresponds to picking three entries from this matrix using every row and column once such that the sum of these three entries is as low as possible. Here, this gives the unique choice of the cooling unitary that might also minimize the gate complexity as $U_{\text{cool}}^* = |0101\rangle\langle 1100| + |1100\rangle\langle 0101| + |0011\rangle\langle 1010| + |1010\rangle\langle 0011| + |0111\rangle\langle 1000| + |1000\rangle\langle 0111| + \mathbb{1}_{\text{Rest}}$. More generally, the MWPM on $G(K, K \oplus 1, E, d_H)$ will not be unique. While in this section we have focused on optimizing the Hamming distance between energy levels being exchanged, as a proxy for the gate complexity, there is nothing which prevent us from optimizing other constraints by simply choosing another cost function $C(v_i, v_j)$ such as work or dissipated heat with which to weight the edges intersecting nodes $v_i \in K, v_j \in K \oplus 1$. Indeed, adding multiple constraints, e.g., Hamming distance and energetic cost, can lead to a unique choice of cooling unitary, as we discuss in Appendix H. This offers a new way of designing cooling unitaries in quantum thermodynamics which both gives insights into the design of circuits for cooling due to its decomposed structure considering pairs of energy levels to be exchanged at a time, and into the optimization of these circuits relative to a cost function of interest.

G. Connections to Other Work and Applications

Algorithmic cooling (AC) [12,13,18] is a collection of techniques to cool a single thermal qubit with access to n identical thermal qubits and the ability to refresh these qubits with a thermal bath. As such, most AC protocols are divided into two parts: (1) *entropy compression*—the process of unitarily ordering the joint-system machine populations in decreasing order; and (2) *relaxation*—where the machine qubits are reset to their initial temperature with access to a thermal bath. These two steps are repeated for many rounds, cooling the target system. In this context, our work gives a general framework for the analysis of the entropy-compression step. Namely, we give three new insights to this part of the procedure. (i) The techniques used for entropy compression such as the partner-pair algorithm [12,13] are heuristic and nonconstructive while our inequalities and minimum-weight perfect-matching methods are constructive. (ii) The heuristic methods used in algorithmic cooling do not cater for optimization relative to costs, e.g., work cost or gate complexity, while the MWPM method provided can be used to optimize

cooling unitaries for any cost function. (iii) Our setting generalizes the scenario from n identical thermal qubits to n thermal qubits with arbitrary energetic structure where the lexicographic order heuristically used in the partner-pair algorithm is no longer enough and the partial order discussed in this work must be considered.

Quantum Error Correction. Several works in this field have pointed out a connection between cooling and error correction [32–37]. Despite experimentally implementing cooling protocols for error correction [33,35] and theoretical works in dissipative quantum error correction investigating how cooling to the ground state of the Hamiltonian given by the negative sum of the stabilizers of the code [32,34,36] naturally restores the system to the code space, a connection with the state of the art of quantum thermodynamics theory is lacking. In Appendix E we explore how the techniques presented in our work can be applied to this scenario by rephrasing the three-qubit error-correcting code as a cooling problem, setting the stage for future work.

Autonomous Thermal Machines. Much of the literature in quantum thermodynamics in the past decade has focused on cooling via the *virtual-qubit subspace swap* [2,6,19,21,38] where the $|0_S 1_M\rangle \leftrightarrow |1_S 0_M\rangle$ exchange is carried out. This exchange carries the most powerful population transfer which in cooling strategies using refreshed machines asymptotically determine how much one can cool [6,38]. In this work we have clarified, using our inequalities, how many other population exchanges $|0 i_M\rangle, |1 i_M \oplus 1\rangle$ with $E(i_M)$ satisfying our main inequality, play a role in a cooling scenario. This opens the door for more efficient cooling protocols using a small finite number of refreshed machines rather than asymptotically many. Embedding the virtual-qubit subspace swap in Lindbladian dynamics where it is the effective steady-state interaction has also inspired autonomous thermal machines used for refrigeration [19], timekeeping [39] and thermometry [40]. Recent works [41,42] even explore using such an exchange toward a computational goal. Since our work identifies cooling operations beyond the virtual-qubit subspace swap, we anticipate that these energy-level permutations can also be embedded in Lindbladian dynamics, leading to more intricate autonomous thermal machines which can perform more complex tasks than those based around the virtual-qubit subspace swap.

Role of Quantum Coherence and Purification. In this research we have chosen to focus on the simplistic setup of machines consisting of thermal qubits in a separable state, allowing in principle only classical correlations to play a role in cooling. An interesting direction to further this work would be to generalize to more complicated machines featuring higher-dimensional systems such as qudits, resonators, many-body systems and, importantly, quantum correlations [7,43–45]. In this direction, in Appendix G, we show how our framework and techniques

can be adapted to the task of amplifying the purity of an arbitrary qubit state with access to n other arbitrary qubit states.

Cooling and Preserving Symmetries. In recent years, the topic of thermalization of quantum systems which preserve non-Abelian symmetries or charges has been a prominent point of inquiry [46,47]. Considering these works provokes one to ponder how preserving symmetries impacts an agent's ability to cool a quantum system within the framework that we have presented. This question has been pursued in Ref. [30], where it has been observed that respecting symmetries restricts that set of unitaries which can cool a quantum system. In Appendix F we investigate this question within our framework for energy and total spin conservation and similarly find that the set of cooling unitaries reduces under these constraints.

V. CONCLUSIONS

In this work, we have taken the folkloric idea that cooling a quantum system given access to another consists of only reordering the eigenvalues of the joint state in descending order, and we have examined the structure of this sorting in detail. We have found that the intricacies of this rearrangement are determined by a simple set of inequalities which depend on the macroscopic properties of the cooling scenario, namely the global temperature of the machine, its total energy E_{Max} and the temperature and energy of the system, elucidating the physics behind the heuristic sorting approach to cooling. These inequalities also allow an agent to cool the system without considering the whole system-machine energy-density matrix by simply checking whether a given energy level of the machine is above the threshold value $\frac{1}{2}(T_M\omega/T_S + E_{\text{Max}})$.

We have also been able to show that these inequalities are sensitive to what number of qubits in the machine are contributing to an optimal-cooling interaction with S . Here, we have found that machines the gaps of which increase exponentially are generally fully reducible to an interaction involving only their coldest qubit $[(n-1)\text{-reducible}]$. Also, we have been able to verify that degenerate machines are generally irreducible, i.e., all their qubits contribute to optimally cooling S . This irreducibility criterion should prove especially helpful when designing quantum thermal machines for cooling quantum systems as it allows one to ensure that the whole machine that they have constructed will be useful at the cooling task at hand. This is especially relevant as autonomous thermal machines are beginning to be experimentally realized [22,48] with the hope of being deployed to cool quantum systems.

Lastly, we have provided a systematic way of decomposing cooling unitaries into a set of two-level permutations. This has led us to provide a novel graph representation for cooling scenarios which is explored further

in the appendices. Importantly, this graph representation has inspired us to consider the optimization of cooling unitaries relative to any cost function of interest as a perfect-matching problem on bipartite graphs. In doing so we have further solidified the ability to represent quantum cooling protocols as quantum circuits [8,49] and developed a new tool for their optimization.

The study of the theoretical limits to cooling a system has been central to every era of thermodynamic inquiry in the past centuries, be that the Carnot cycle for heat engines [50] or the Doppler limit for laser cooling [51]. We are now in the qubit era, forced to imagine if a quantum thermodynamic system could be fully controlled and acted upon using a quantum computer and what the limits and methods of cooling would be in this scenario. Here, we have offered a further step toward contending with the challenge that our era of thermodynamic inquiry faces.

ACKNOWLEDGMENTS

J.X. acknowledges Josef Lauri who taught graph theory at the University of Malta and Xandru Mifsud & Adriana Baldacchino for their hospitality and pointing out that the graphs in earlier versions of this work were hypercube graphs. The authors thank Nayeli Briones Rodriguez for enlightening discussions on the state of the art in algorithmic cooling and Paul Skrzypczyk, Chung-Yun Hsieh & Fabien Clivaz for insightful discussions. The authors acknowledge TU Wien Bibliothek for financial support through its Open Access Funding Programme. B.S. acknowledges support from UK EPSRC (EP/SO23607/1). J.X., M.H. and P.B. acknowledge funding from the European Research Council (Consolidator grant ‘Cocoquest’ 101043705) and financial support from the Austrian Federal Ministry of Education, Science and Research via the Austrian Research Promotion Agency (FFG) through the project FO999914030 (MUSIQ) funded by the European Union – NextGenerationEU. A.R. acknowledges support from the the Swiss National Science Foundation for funding through Postdoc. Mobility (Grant No. P500PT225461).

DATA AVAILABILITY

No data were created or analyzed in this study. The code for generating the presented plots can be found in Ref. [52].

APPENDIX A: HAMMING-WEIGHT CONJUGATE ENERGY EXCHANGES

The central mathematical object of this work is the inequality

$$\frac{T_M}{T_S}\omega < (i_M - j_M) \cdot \Gamma, \quad (\text{A1})$$

which when valid indicates that exchanging the energy levels $|0_S i_M\rangle$ and $|1_S j_M\rangle$ in the joint state of the system at




FIG. 6. Visualization of the energy levels of the machine, showing that (i) Hamming-weight conjugate pairs are equally distant from the extremal energies, i.e., $E(l_M)$ is equally away from 0 as $E(l_M \oplus 1)$ is from E_{Max} and (ii) Hamming-weight conjugate pairs fall on one side or the other of the median energy value $E_{\text{Max}}/2$.

temperature T_M with gap ω and the n -qubit machine at temperature T_M with gaps $\Gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$, cools S . In much of the paper we set $j_M = i_M \oplus 1$, first because it allows us to index the ground- and excited-state subspaces using a single index and second because it is the exchange of $|0_{Si_M}\rangle$ and $|1_{Si_M} \oplus 1\rangle$ which determines whether there is any other $|1_{Sj_M}\rangle$, $j_M \in \{0, 1\}^{\times n}$ which when exchanged with $|0_{Si_M}\rangle$ also cools S . In this appendix we will explain why this is the case, both from the perspective of the energetic structure of the problem and from the mathematical perspective of swapping the largest entry with the smallest entry across the median of a vector.

1. Strategy for Detecting Disordered Populations

a. Informal Argument

Considering Hamming-weight conjugate pairs by setting $|1_{Sj_M}\rangle = |0_{Si_M} \oplus 1\rangle$ gives us a systematic way to go through all energy levels of the system and machine to find the disordered populations which must be swapped. In the inequality in Eq. (A1) which detects which energy levels should be exchanged, the right-hand side only features the machine energy levels i_M and j_M . By making this choice in indexing and recalling that the entries of Γ and $\gamma_j \leq \gamma_{j+1}$ are in nondecreasing order, we are exposing ourselves to a choice in the energetic ordering of these machine energy levels. As visualized in Fig. 6 the bit string $|l_M\rangle$ with energy $E(l_M) = l_M \cdot \Gamma = \sum_{j=1}^n x_j \gamma_j$, where x_j is the j th bit in l_M , is energetically $E(l_M)$ away from 0, while the energy of its Hamming-weight conjugate $E(l_M \oplus 1)$ is also $E(l_M)$ from the maximum energy $E_{\text{Max}} = \sum_{j=1}^n \gamma_j$ as

$$E(l_M \oplus 1) = E_{\text{Max}} - E(l_M). \quad (\text{A2})$$

With this observation we note that if $E(l_M)$ is lower than the median energy value $E_{\text{Max}}/2$, then $E(l_M \oplus 1)$ must be above it; i.e., $E(l_M) < E_{\text{Max}}/2$ implies $E(l_M \oplus 1) > E_{\text{Max}}/2$ since $E(l_M \oplus 1) = E_{\text{Max}} - E(l_M)$.

Perhaps most importantly, if one chooses to exchange $|0_{Si_M}\rangle$ with $|1_{Si_M} \oplus 1\rangle$ proceeding lexicographically through $i_M \in \{0, 1\}^{\times n}$ pairing Hamming-weight conjugate

kets in the inequality, *one always pairs the energetically furthest eigenkets available*. We begin with $|1^n - 0^n\rangle \cdot \Gamma = E_{\text{Max}}$ then $|1^{n-1}0 - 0^{n-1}1\rangle \cdot \Gamma = E_{\text{Max}} - \gamma_n$ and so on, moving in decreasing size of energy difference until $|10^{n-1} - 01^{n-1}\rangle \cdot \Gamma = -E_{\text{Max}} + \gamma_1$ and $|0^n - 1^n\rangle \cdot \Gamma = -E_{\text{Max}}$.

Thus if $|0_{Si_M}\rangle$ does not lead to a valid inequality when paired with $|1_{Si_M} \oplus 1\rangle$ there can be no other energy level $|1_{Sj_M}\rangle$ of a lower energy in the excited-state subspace worth swapping with as $|i_M - j_M\rangle \cdot \Gamma > |i_M - i_M \oplus 1\rangle \cdot \Gamma$ since j_M corresponds to a lower energy than $i_M \oplus 1$ but $\frac{T_M}{T_S}\omega \geq |i_M - i_M \oplus 1\rangle \cdot \Gamma$ by assumption, while machine eigenkets of higher energy would have already been paired up. In this way, pairing Hamming-weight conjugate energy levels exhausts the set of machine levels and detects all disordered populations required to be exchanged for cooling. For a mathematically clearer and complete exposition to this strategy, see below.

Lastly, as we mention in Sec. IV while this pairing is natural and useful to consider for detecting disordered population it might not lead to the pairing for obtaining a gate-complexity optimal-cooling unitary.

b. Heuristic Cooling Algorithm

Given that the initial state of the system and machine is diagonal in the energy eigenbasis, all the populations can be collected into two vectors $\mathbf{q}_0$ and $\mathbf{q}_1$ such that

$$p_0 = \sum_{i=1}^{2^n} q_i, \quad q_i \in \mathbf{q}_0, \quad p_1 = 1 - p_0 = \sum_{i=1}^{2^n} l_i, \quad l_i \in \mathbf{q}_1, \quad (\text{A3})$$

where p_0 and p_1 are the initial ground- and excited-state populations of the system qubit. A heuristic way to perform optimal cooling would be to first order $\mathbf{q}_0$ and $\mathbf{q}_1$ in nondecreasing order, giving $\mathbf{q}_0^\downarrow$ and $\mathbf{q}_1^\downarrow$, respectively. One can then compare the first component of $\mathbf{q}_1^\downarrow$ with the final component of $\mathbf{q}_0^\downarrow$ and perform a permutation between the two levels if the first component of $\mathbf{q}_1^\downarrow$ is greater than




FIG. 7. A diagram of the heuristic cooling algorithm. The vectors of populations are first placed in nondecreasing order. The final component of the ground-state population vector is compared to the first component of the excited-state population vector and a swap between the levels is performed if the former is larger. This process is repeated until a point is reached at which the component from the ground-state population vector is larger than the component from the excited-state population vector. Given that at each stage the largest remaining component from the ground-state population vector is swapped with the smallest remaining component of the excited-state population vector, it represents the optimal (in terms of increasing the system ground state) two-level permutation that can be performed.

the final component of $\mathbf{q}_0^\downarrow$. If valid, this will perform the optimal exchange for increasing the ground-state population of the system qubit in a single two-level permutation. Next, one compares the second component of $\mathbf{q}_1^\downarrow$ with the second-to-last component of $\mathbf{q}_0^\downarrow$ and performs a permutation between the two levels if the second component of $\mathbf{q}_1^\downarrow$ is greater than the second-to-last component of $\mathbf{q}_0^\downarrow$, and so on. Each comparison of components will lead to a two-level permutation that represents the new optimal energy-level exchange for increasing the ground-state population (given that the previous permutations have been performed). If one reaches the k th component of $\mathbf{q}_1^\downarrow$ and it is not greater than the $(2^{n-1} - k)$ th component of $\mathbf{q}_0^\downarrow$, one has successfully moved the populations such that the ground-state population of the system qubit is maximized and, hence, they have performed optimal cooling via unitary dynamics. See Fig. 7 for a pictorial depiction of the heuristic cooling algorithm.

In total, one swaps the k largest elements of $\mathbf{p}_1$ with the k smallest elements of $\mathbf{p}_0$, where the size of k is dependent on the system and machine being used. Note that once the k elements that must be swapped are known, it does not matter for obtaining the largest attainable ground-state population which of the k elements of $\mathbf{q}_0$ are swapped for the k elements of $\mathbf{q}_1$. All that is important is that all k elements are swapped.

c. Cooling without Populations

Now, given some qubit at temperature β_S with energy gap ω , and access to some machine at temperature β_M with energetic gap structure $\Gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$, is it possible to find which two-level swaps must be performed without first calculating all the populations and then ordering them?

To do this, we first note that each element of $\mathbf{q}_0$ and $\mathbf{q}_1$ can be identified with an n -bit string, i_M , as

$$\mathbf{q}_0 = \left\{ \frac{e^{-\beta_M i_M \cdot \Gamma}}{\mathcal{Z}_S \mathcal{Z}_M} : i_M \in \{0, 1\}^{\times n} \right\},$$

$$\mathbf{q}_1 = \left\{ \frac{e^{-\beta_S \omega - \beta_M i_M \cdot \Gamma}}{\mathcal{Z}_S \mathcal{Z}_M} : i_M \in \{0, 1\}^{\times n} \right\}. \quad (\text{A4})$$

Therefore, the k elements of $\mathbf{q}_0$ that must be swapped can be identified with a set of n -bit strings, as can the k elements of $\mathbf{q}_1$ that they must be swapped with. Moreover, from the above definitions it can easily be seen that $\mathbf{q}_1 = e^{-\beta_S \omega} \mathbf{p}_0$, and hence that $\mathbf{q}_1^\downarrow = e^{-\beta_S \omega} \mathbf{p}_0^\downarrow$ —multiplying each element of an ordered list by a positive constant does not change the ordering. Therefore, finding the set of n -bit strings that identify the k largest elements of $\mathbf{q}_1$ is equivalent to finding the set of n -bit strings that identify the k largest elements of $\mathbf{q}_0$. One can therefore consider only $\mathbf{q}_0$ and aim to find the set of n -bit strings that identify the smallest k elements of $\mathbf{q}_0$, denoted by $\mathbb{S}$, and the set of n -bit strings that identify the largest k elements $\mathbf{q}_0$, denoted by $\bar{\mathbb{S}}$. From the above reasoning, the bit strings that identify the k largest elements of $\mathbf{q}_0$, $\bar{\mathbb{S}}$, will also then identify the k largest elements of $\mathbf{q}_1$. We can then consider the following lemma, which states that if the bit string i_M is in the set of n -bit strings that identify one of the k smallest populations of $\mathbf{q}_0$, then its conjugate, $i_M \oplus 1$, is in the set of n -bit strings that identify one of the k largest populations of $\mathbf{q}_1$.

Lemma A1. If $i_M \in \mathbb{S}$, then $i_M \oplus 1 \in \bar{\mathbb{S}}$, where $i_M \in \{0, 1\}^n$.




FIG. 8. Details of the proof strategy used in Lemma A1. One first considers a vector of energies, $\mathbf{E}_0$, ordered in nondecreasing order and generated from the ordered population vectors. First, it can be shown that if i_M identifies the k th-largest energy in $\mathbf{E}_0$, then $(i_M \oplus 1)$ identifies the $(2^n - k)$ th-largest energy in $\mathbf{E}_0$. Then, given that $\mathbf{E}_0$ and $\mathbf{p}_0$ have reverse orderings, if i_M identifies the k th-largest energy in $\mathbf{E}_0$, it identifies the $(2^n - k)$ th-largest population in $\mathbf{p}_0$. Next, as $\mathbf{p}_0$ and $\mathbf{p}_1$ have the same ordering, it can be concluded that if i_M identifies the $(2^n - k)$ th-largest population in $\mathbf{p}_0$, it identifies the $(2^n - k)$ th-largest population in $\mathbf{p}_1$ as well. Hence, the relationships found between the n -bit strings identifying the k largest and smallest energies of $\mathbf{E}_0$ also apply to the n -bit strings identifying the k smallest populations of $\mathbf{p}_0$ and the k largest populations of $\mathbf{p}_1$. The red and blue lines represent components of the respective vectors which are identified by the same n -bit strings. Purple line “A” represents the relationship proved between the sets of n -bit strings identifying the k largest and smallest energies of $\mathbf{E}_0$ that can then be used to infer the relationship between the sets of n -bit strings identifying the k smallest populations of $\mathbf{p}_0$ and k largest populations of $\mathbf{p}_1$, represented by purple line “B”.

Proof. To begin, we first convert the vector of populations to a vector of energies

$$\mathbf{E}_0 = -\frac{1}{\beta_S} \ln(\mathcal{Z}_S \mathcal{Z}_M \mathbf{q}_0) = \{i_M \cdot \Gamma : i_M \in \{0, 1\}^n\}. \quad (\text{A5})$$

The constant $-1/\beta_M \ln(\mathcal{Z}_S \mathcal{Z}_M)$ does not affect the ordering. Then, given that $1/\beta_M$ is a positive constant and $\ln(a) \leq \ln(b)$ if $a \leq b$, $-\mathbf{E}_0$ has the same ordering as $\mathbf{q}_0^\downarrow$; the negative then flips the ordering. Therefore, the set of n -bit strings that identify the k smallest values of $\mathbf{q}_0$ identify the k -largest values of $\mathbf{E}_0$, and vice versa for the k largest values of $\mathbf{q}_0$. Hence, $\mathbb{S}$ now identifies the largest k values of $\mathbf{E}_0$, and $\bar{\mathbb{S}}$ the k largest of $\mathbf{E}_1$ (defined as in Eq. (A5) but using $\mathbf{q}_1$). However, such labeling is irrelevant given that one only cares about the relationship between the n -bit strings in the sets.

Next, we note that

$$i_M \cdot \Gamma + (i_M \oplus 1) \cdot \Gamma = E_{\text{Max}}, \quad \text{where} \quad E_{\text{Max}} = \sum_{i=1}^n \gamma_i, \quad (\text{A6})$$

meaning that the component of $\mathbf{E}_0$ identified by the bit string i_M is paired to the component of $\mathbf{E}_0$ identified by the bit string $(i_M \oplus 1)$. Evidently, the larger $i_M \cdot \Gamma$, the smaller is $(i_M \oplus 1) \cdot \Gamma$ and vice versa. From here, it can already be seen that if a n -bit string leads to one of the k largest values

of $\mathbf{E}_0$, its conjugate will lead to one of the k smallest. More formally, one can define

$$\begin{aligned} E_{\text{sup}} &= \min_{i_M \in \mathbb{S}} i_M \cdot \Gamma = i_{\text{sup}} \cdot \Gamma, \\ E_{\text{inf}} &= \max_{i_M \in \{0, 1\}^n \setminus \mathbb{S}} i_M \cdot \Gamma = i_{\text{inf}} \cdot \Gamma, \end{aligned} \quad (\text{A7})$$

such that E_{sup} is the smallest component of $\mathbf{E}_0$ of the k largest components (acting as a kind of supremum), and E_{inf} is the largest component of $\mathbf{E}_0$ not in the k largest (acting as a kind of infimum), giving $E_{\text{sup}} \geq E_{\text{inf}}$. Therefore,

$$\begin{aligned} E_{\text{Max}} - E_{\text{inf}} &\geq E_{\text{Max}} - E_{\text{sup}} \\ \Rightarrow (i_{\text{inf}} \oplus 1) \cdot \Gamma &\geq (i_{\text{sup}} \oplus 1) \cdot \Gamma. \end{aligned} \quad (\text{A8})$$

All other n -bit strings in $\mathbb{S}$ give components of $\mathbf{E}_0$ larger than E_{sup} and hence their conjugates lead to components smaller than $E_{\text{Max}} - E_{\text{sup}}$. Thus far, we have shown that if i_M is an n -bit string that represents one of the k largest energies of $\mathbf{E}_0$, $(i_M \oplus 1)$ is an n -bit string that represents one of the k smallest energies of $\mathbf{E}_0$.

Now, given that the ordering of $\mathbf{E}_0$ is that of $\mathbf{p}_0$ reversed, if i_M is an n -bit string that represents one of the k largest energies of $\mathbf{E}_0$, it is an n -bit string that represents one of the k smallest populations of $\mathbf{p}_0$. Hence, $(i_M \oplus 1)$ is an n -bit string that represents one of the k largest populations of $\mathbf{p}_0$. To complete the proof, one utilizes the fact that the

ordering of $\mathbf{p}_0$ in nondecreasing order is the same as that of $\mathbf{p}_1$. Hence, the n -bit strings that identify the k largest populations of $\mathbf{p}_0$ also identify the k largest populations of $\mathbf{p}_1$. Finally, given that they are the same set of bit strings, any relationship between the n -bit strings identifying the k largest and smallest energies of E_0 must also hold for those n -bit strings representing the k largest and smallest populations of $\mathbf{p}_0$ and $\mathbf{p}_1$. They then must hold between the n -bit strings identifying the k smallest populations of $\mathbf{p}_0$ and k largest populations of $\mathbf{p}_1$ for the same reason. Using the relationship proved between the n -bit strings representing the k largest and smallest energies of E_0 , it can therefore be concluded that if $i_M \in \mathbb{S}$ then $(i_M \oplus 1) \in \bar{\mathbb{S}}$. In other words, if $E(i_M)$ is larger than the energetic median, $E(i_M \oplus 1)$ is smaller than as visualized in Fig. 6. Interestingly, one can go further and note that if the n -bit string i'_M identifies the k' largest component of $\mathbf{p}_0$, then $(i'_M \oplus 1)$ identifies the k' smallest component (the $2^{n-1} - k'$ largest) of $\mathbf{p}_1$. ■

We now return to the heuristic cooling algorithm. Lemma A1 states that if $\mathbb{S}$ represents the set of n -bit strings that identify the smallest components of $\mathbf{q}_0$, then the set of n -bit strings that identify the largest components of $\mathbf{q}_1$, $\bar{\mathbb{S}}$, are given by $\bar{\mathbb{S}} = \{i_M \oplus 1 : i_M \in \mathbb{S}\}$. Given that

$$\min_{j_M \in \bar{\mathbb{S}}} j_M \cdot \Gamma > \max_{i_M \in \mathbb{S}} i_M \cdot \Gamma, \quad (\text{A9})$$

the smallest component of $\mathbf{q}_1$ to be swapped is still larger than the largest component of $\mathbf{q}_0$ to be swapped (where this ceases to be true, i.e., at what value of k , defines the end of the heuristic cooling algorithm). Which of these k components are paired with each other to be exchanged does not matter—they all lead to positive changes in the ground-state population of the system. Therefore, without loss of generality, one can compare the population of $\mathbf{p}_0$ identified by the n -bit string i_M with the population of $\mathbf{p}_1$ identified by the n -bit string $(i_M \oplus 1)$ to see if the population identified by i_M should be swapped. In fact, given that if the n -bit string $i'_M \oplus 1$ identifies the k' largest component of $\mathbf{p}_1$, then i'_M identifies the k' smallest component of $\mathbf{p}_0$ (the $2^n - k'$ largest component), the strategy of swapping opposite Hamming-weight pairs exactly outputs the same set of two-level swaps as performing the heuristic cooling algorithm, in which at each stage the optimal two-level

swap to increase the ground-state population of the system has been performed.

2. Graphical Interpretation

The inequality in Eq. (5) also has a geometric interpretation when dividing both sides by the absolute difference of the Hamming distance between the bit strings i_M and j_M , $d_H(i_M, j_M) = |i_M - j_M| = \sum_{l=1}^n |i_l - j_l|$, where i_l and j_l are the l th bits in each string. This gives

$$\frac{T_M \omega}{T_S |i_M - j_M|} < \frac{(i_M - j_M) \cdot \Gamma}{|i_M - j_M|}. \quad (\text{A10})$$

Here, we see that the right-hand side takes on the form of a slope $\Delta y / \Delta x$, where $\Delta x = |i_M - j_M|$ is the Hamming distance between machine energy levels i_M and j_M and $\Delta y = \Delta E = (i_M - j_M) \cdot \Gamma$ is the energy difference between these machine energy levels. The slope captures a sense of how far away the kets that one exchanges have to be to obtain a given energy change.

If this slope is larger than the slope $T_M \omega / T_S |i_M - j_M|$ given by the cooling-scenario constant and the same Δx , then the exchange of the levels $|0_{Si_M}\rangle$ and $|1_{Sj_M}\rangle$ will cool S . Note that pairing inverse Hamming-weight kets and setting $|1_{Mj_M}\rangle = |0_{Si_M} \oplus 1\rangle$ sets Δx to a fixed maximal value of n for any pairing. In this sense we again see the logic of this strategy, that if the slope $((i_M - i_M \oplus 1) \cdot \Gamma) / n$ is not greater than $T_M \omega / T_S n$ there can be no other slope formed by pairing with a j_M of energy lower than $i_M \oplus 1$ which results in a slope that exceeds $T_M \omega / T_S n$.

APPENDIX B: EXAMPLE 2—COOLING USING A THREE-QUBIT MACHINE

To continue to understand the tools we have developed in this paper and the setting at hand, let us consider an exemplary task of cooling a thermal qubit S with gap ω with $\beta_S = 1$ with access to three thermal qubits also with $\beta_M = 1$ and gaps $\Gamma = (\gamma_1, \gamma_2, \gamma_3)$ such that $\omega \leq \gamma_1 \leq \gamma_2 \leq \gamma_3$. We will use our tools to access which swaps should be performed for optimal cooling, which machines are irreducible or reducible, and evaluate a lower bound on the optimal achievable ground-state temperature using a $n = 3$ qubit machine.

The population vectors for the ground and excited states of the system, $\mathbf{p}_0$ and $\mathbf{p}_1$ respectively, are

$$\mathbf{p}_0 = \frac{1}{Z_S Z_M} [1, e^{-\gamma_3}, e^{-\gamma_2}, e^{-(\gamma_2+\gamma_3)}, e^{-\gamma_1}, e^{-(\gamma_1+\gamma_3)}, e^{-(\gamma_1+\gamma_2)}, e^{-(\gamma_1+\gamma_2+\gamma_3)}], \quad (\text{B1})$$

$$\mathbf{p}_1 = \frac{1}{Z_S Z_M} [e^{-\omega}, e^{-(\omega+\gamma_3)}, e^{-(\omega+\gamma_2)}, e^{-(\omega+\gamma_2+\gamma_3)}, e^{-(\omega+\gamma_1)}, e^{-(\omega+\gamma_1+\gamma_3)}, e^{-(\omega+\gamma_1+\gamma_2)}, e^{-(\omega+\gamma_1+\gamma_2+\gamma_3)}] \quad (\text{B2})$$




FIG. 9. The inequality in Eq. (5) can be visualized graphically on a graph the x axis of which represents the bit strings labeling machine energy levels and the Hamming distance between these strings, whereas the y axis represents the energies of these levels and so their energy differences.

in this scenario, where $\mathcal{Z}_S = (1 + e^{-\omega})$ and $\mathcal{Z}_M = (1 + e^{-\gamma_1})(1 + e^{-\gamma_2})(1 + e^{-\gamma_3})$. The heuristic algorithm given in the previous section can be used to determine a set of two-level swaps which will lead to optimal cooling of the system by comparing elements of $\mathbf{p}_0$ and $\mathbf{p}_1$. Alternatively, it has also been shown in the main text and in Appendix A that this set of two-level swaps can be found by considering Eq. (7) for all $i_M \in \{0, 1\}^{\times 3}$. In this case, Eq. (7) takes the form

$$\frac{1}{2}(\omega + \gamma_1 + \gamma_2 + \gamma_3) < i_1\gamma_1 + i_2\gamma_2 + i_3\gamma_3, \quad (\text{B3})$$

with the two-level swap $|0_s i_1 i_2 i_3\rangle \leftrightarrow |1_s(i_1 \oplus 1)(i_2 \oplus 1)(i_3 \oplus 1)\rangle$ cooling the system if Eq. (B3) is satisfied for a given $i_M = i_1 i_2 i_3$. Specifically, the following two-level swaps can be performed to achieve optimal cooling if the related conditions are satisfied:

$$|0_S000\rangle \longleftrightarrow |1_S111\rangle \text{ if } \omega < -\gamma_1 - \gamma_2 - \gamma_3, \quad (\text{B4})$$

$$|0_S 100\rangle \longleftrightarrow |1_S 011\rangle \text{ if } \omega < \gamma_1 - \gamma_2 - \gamma_3, \quad (\text{B5})$$

$$|0_S 010\rangle \longleftrightarrow |1_S 010\rangle \text{ if } \omega < -\gamma_1 + \gamma_2 - \gamma_3, \quad (\text{B6})$$

$$|0_S 110\rangle \longleftrightarrow |1_S 001\rangle \text{ if } \omega < \gamma_1 + \gamma_2 - \gamma_3, \quad (\text{B7})$$

$$|0_S 001\rangle \longleftrightarrow |1_S 110\rangle \text{ if } \omega < -\gamma_1 - \gamma_2 + \gamma_3, \quad (\text{B8})$$

$$|0_S 101\rangle \longleftrightarrow |1_S 010\rangle \text{ if } \omega < \gamma_1 - \gamma_2 + \gamma_3, \quad (\text{B9})$$

$$|0_S 011\rangle \longleftrightarrow |1_S 100\rangle \text{ if } \omega < -\gamma_1 + \gamma_2 + \gamma_3, \quad (\text{B10})$$

$$|0_S 111\rangle \longleftrightarrow |1_S 000\rangle \text{ if } \omega < \gamma_1 + \gamma_2 + \gamma_3. \quad (\text{B11})$$

Due to the assumed gap structure of $\omega \leq \gamma_1 \leq \gamma_2 \leq \gamma_3$, it can be seen that Eqs. (B11), (B10) and (B9) are always satisfied. This therefore means that their associated two-level swaps always cool the system. Note, from the perspective of the heuristic cooling algorithm, that these conditions always being satisfied are equivalent to the fact that the following conditions on the populations from $\mathbf{p}_0$ and $\mathbf{p}_1$ are always satisfied:

$$\omega < \gamma_1 - \gamma_2 + \gamma_3 \iff e^{-(\omega+\gamma_2)} > e^{-(\gamma_1+\gamma_3)}, \quad (\text{B12})$$

$$\omega < -\gamma_1 + \gamma_2 + \gamma_3 \iff e^{-(\omega+\gamma_1)} > e^{-(\gamma_2+\gamma_3)}, \quad (\text{B13})$$

$$\omega < \gamma_1 + \gamma_2 + \gamma_3 \iff e^{-\omega} > e^{-(\gamma_1 + \gamma_2 + \gamma_3)}. \quad (\text{B14})$$

Meanwhile, Eqs. (B6), (B5) and (B4) are never satisfied, meaning that their associated two-level swaps never cool the system. As above, these conditions directly correspond to conditions on the populations from $\mathbf{p}_0$ and $\mathbf{p}_1$ that are never satisfied.

This leaves two possible conditions, namely $\omega < \gamma_1 + \gamma_2 - \gamma_3$ Eq. (B7) and $\omega < -\gamma_1 - \gamma_2 + \gamma_3$ [Eq. (B8)] that one must consider. Given these, there are four possible cases to consider: (1) Eqs. (B7) and (B8) are satisfied, (2) Eq. (B7) is satisfied and Eq. (B8) is not, (3) Eq. (B8) is satisfied and Eq. (B7) is not, and (4) Eqs. (B7) and (B8) are not satisfied.

Case 1. Eqs. (B7) and (B8) are satisfied.

It can easily be seen that this is never possible. This is due to $\omega \geq 0$ and the right-hand sides of Eqs. (B7) and (B8) being the negative of each other.

Case 2. Eq. (B7) is satisfied and Eq. (B8) is not.




FIG. 10. An illustration of the various cooling scenarios for a three-qubit machine by giving a three-dimensional projection of a hypercube. (a) The situation given by case 1 $\gamma_1 + \gamma_2 < \gamma_3$ and $\omega < \gamma_3 - \gamma_1 - \gamma_2$ for an example with $\omega = \log 2$, $\gamma_1 = \log 3$, $\gamma_2 = \log 5$ and $\gamma_3 = \log 32$, where the population vector has been factored by $\frac{40}{90}$. (b) Case 2 for $\gamma_1 + \gamma_2 > \gamma_3$ and $\omega < \gamma_1 + \gamma_2 - \gamma_3$ for $\omega = \log 2$, $\gamma_1 = \log 3$, $\gamma_2 = \log 22$ and $\gamma_3 = \log 32$ factored by $\frac{32}{69}$. (c) The situation for cases 3, 4 and 5 with values satisfying case 5 explicitly, $\omega = \log 2$, $\gamma_1 = \log 2$, $\gamma_2 = \log 16$ and $\gamma_3 = \log 32$, with the population vector factored by $\frac{2048}{5049}$.

In addition to the three always-performed swaps, the additional swap

$$\begin{aligned} |0_S 110\rangle &\longleftrightarrow |1_S 001\rangle \text{ as } \omega < \gamma_1 + \gamma_2 - \gamma_3 \\ &\iff e^{-(\omega+\gamma_3)} > e^{-(\gamma_1+\gamma_2)} \end{aligned} \quad (\text{B15})$$

is needed for optimal cooling in this case. If all swaps are performed, then the final ground-state population in this case is

$$p_0^*(2) = \frac{1 + e^{-\gamma_1} + e^{-\gamma_2} + e^{-\gamma_3}}{\mathcal{Z}_M}, \quad (\text{B16})$$

where

$$\omega < \gamma_1 + \gamma_2 - \gamma_3 \implies p_0^*(2) > \frac{1}{1 + e^{-\gamma_3}}, \quad (\text{B17})$$

meaning that one cools the system further than the coldest qubit.

Case 3. Eq. (B8) is satisfied and Eq. (B7) is not.

In addition to the three always-performed swaps, the additional swap

$$\begin{aligned} |0_S 001\rangle &\longleftrightarrow |1_S 110\rangle \text{ as } \omega < -\gamma_1 - \gamma_2 + \gamma_3 \\ &\iff e^{-(\omega+\gamma_1+\gamma_2)} > e^{-\gamma_3} \end{aligned} \quad (\text{B18})$$

is needed for optimal cooling in this case. If all swaps are performed, then the final ground-state population in this

case is

$$p_0^*(3) = \frac{1}{1 + e^{-\gamma_3}}, \quad (\text{B19})$$

meaning that the optimal cooling one can achieve is the temperature of the coldest qubit. Hence, the optimal unitary for cooling in this case is simply the SWAP between the system and the coldest qubit in the machine.

In the main text and Appendix D the notion of machine reducibility is discussed. Case 3 fits the definition of an $(n-1)$ -reducible machine given that two of the three machine qubits are not involved in the cooling process—the final ground-state population does not depend on γ_1 or γ_2 . In Eq. (13) the following sufficient condition for $(n-1)$ -reducibility is given:

$$\frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < \gamma_n. \quad (\text{B20})$$

In this case, this takes the form

$$\frac{1}{2}(\omega + \gamma_1 + \gamma_2 + \gamma_3) < \gamma_3, \quad (\text{B21})$$

which can equivalently be expressed as

$$\omega < -\gamma_1 - \gamma_2 + \gamma_3. \quad (\text{B22})$$

Hence, Eq. (B8) is exactly the condition for $(n-1)$ -reducibility when $n = 3$. It being satisfied therefore immediately tells us that the optimal cooling that can be

achieved is to the temperature of the coolest machine qubit, and that a full swap between the coldest machine qubit and the system is an optimal-cooling unitary.

Case 4. Eqs. (B7) and (B8) are not satisfied.

This leads to no additional swaps beyond the three always-performed swaps. Performing these gives

$$p_0^*(4) = \frac{1 + e^{-\omega} + e^{-\gamma_1} + e^{-\gamma_2} + e^{-\gamma_3} + e^{-(\omega+\gamma_2)} + e^{-(\omega+\gamma_3)} + e^{-(\gamma_1+\gamma_2)}}{\mathcal{Z}_S \mathcal{Z}_M}. \quad (\text{B23})$$

From here, it can be seen that

$$\omega > -\gamma_1 - \gamma_2 + \gamma_3 \implies p_0^*(3) > \frac{1}{1 + e^{-\gamma_3}}, \quad (\text{B24})$$

where the first condition comes from Eq. (B8) not being satisfied. Hence, one can cool past the temperature of the coldest qubit in this case.

It has been seen in case 3 that the machine was $(n-1)$ -reducible. We now consider whether the machines in cases 2 and 4 are irreducible, meaning that all qubits in the machine contribute to the optimal cooling. In Appendix D, it is shown that a necessary and sufficient condition for a machine to be irreducible is for there to exist an $l_{n-1} \in \{0, 1\}^{\times n-1}$ such that

$$0 \leq \frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) - E(l_{n-1}) < \gamma_1, \quad (\text{B25})$$

where $E(l_{n-1})$ is the energy of the $(n-1)$ coldest qubits. In our case of $n=3$ machine qubits this becomes

$$0 \leq \frac{1}{2} (\omega + \gamma_1 + \gamma_2 + \gamma_3) - E(0_S 0_{i_2} i_3) < \gamma_1. \quad (\text{B26})$$

We first consider case 2, where Eq. (B7) is satisfied and Eq. (B8) is not. For $i_2 = 1$ and $i_3 = 0$ Eq. (B26) becomes

$$-\gamma_1 + \gamma_2 - \gamma_3 < \omega < \gamma_1 + \gamma_2 - \gamma_3. \quad (\text{B27})$$

Given that Eq. (B7) is satisfied the right-hand side of this equation is satisfied. Given the nondecreasing energy structure of the machine, the left-hand side is always satisfied. Hence, the machine in case 2 is irreducible.

We then consider case 4, where Eq. (B7) is satisfied and Eq. (B8) is not. For $i_2 = 0$ and $i_3 = 1$ Eq. (B26) becomes

$$-\gamma_1 - \gamma_2 + \gamma_3 < \omega < \gamma_1 - \gamma_2 + \gamma_3. \quad (\text{B28})$$

Once again, given the nondecreasing energy structure of the machine the right-hand side is always satisfied. Then, given that Eq. (B8) is not satisfied, the left-hand side is satisfied. Therefore, the machine in case 4 is also irreducible. Given that the optimal achievable ground-state populations in cases 2 and 4, $p_0^*(2)$ and $p_0^*(4)$, respectively, depend on all elements in Γ , intuitively one would assume that the machines are irreducible. However, the above calculations confirm this.

Lastly, one can lower bound the optimal achievable change in population of the system, Δp_0^* , using Eq. (11), which, in this case gives

$$\Delta p_0^* \geq \frac{1}{\mathcal{Z}_S(1 + e^{-\gamma_3})} (e^{-\omega} - e^{-\gamma_3}) + \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} (e^{-\gamma_3} - e^{-(\omega+\gamma_1+\gamma_2)}), \quad (\text{B29})$$

where it is noted that the second term is negative if Eq. (B8) is satisfied.

APPENDIX C: HOW COLD CAN WE GO?

Given a thermal qubit at temperature T_S with gap ω and a machine M of n qubits with gaps $\Gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$, a natural question to ask is *how far we can cool S with a single unitary interaction across S and M ?*

Addressing this question within the language of the hierarchy of inequalities we have developed, let us assume that an agent has permuted the energy levels within the joint system-machine Hilbert space such that system cannot be cooled further, i.e., none of the inequalities in Eq. (5) remain valid.

In doing so, the agent obtains a ground-state population for their system qubit after cooling that can be expressed as

$$p_0' = \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(\underbrace{\sum_{i_M \in \{0,1\}^{\times n} \setminus \mathbb{S}} e^{-\beta_M i_M \cdot \Gamma}}_{\text{unchanged}} + \underbrace{\sum_{j_M \in \mathbb{S}} e^{-\beta_S \omega - \beta_M (j_M \oplus 1) \cdot \Gamma}}_{\text{swapped}} \right), \quad (\text{C1})$$

where

$$\mathbb{S} = \left\{ j_M \in \{0, 1\}^{\times n} \mid \frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < E(j_M) \right\}$$

is the set of machine energy levels corresponding to levels in the joint system-machine Hilbert space that will cool the system if swapped $|0_S j_M\rangle \longleftrightarrow |1_S j_M \oplus 1\rangle$, and $\mathcal{Z}_S = (1 + e^{-\beta_S \omega})$ and $\mathcal{Z}_M = \prod_{j=1}^n (1 + e^{-\beta_M \gamma_j})$ are the partition functions of the system and machine, respectively. In Eq. (C1) we have represented the ground-state population after cooling as a sum of two contributions, the first term involving energy levels that did not satisfy Eq. (5) and so

would warm the system if swapped, and the second term involving the populations obtained by swapping the levels that satisfied Eq. (5). Note that the choice of swapping Hamming-weight conjugate levels is nonunique but is useful as it allows us to express changes in energy in terms of one parameter over the set of bit strings.

To investigate the performance of a machine in the task of optimally cooling S we wish to calculate the increase in the ground-state population of S as a result of our protocol. The initial population has been set to be $p_0 = 1/(1 + e^{-\beta_S \omega})$, allowing us to express the difference in populations as

$$\Delta p_0 = p'_0 - p_0 = \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(\sum_{i_M \in \{0,1\}^{\times n} \setminus \mathbb{S}} e^{-\beta_M i_M \cdot \Gamma} + \sum_{j_M \in \mathbb{S}} e^{-\beta_S \omega - \beta_M (j_M \oplus 1) \cdot \Gamma} \right) - \frac{1}{\mathcal{Z}_S} \quad (\text{C2})$$

$$= \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(\sum_{i_M \in \{0,1\}^{\times n} \setminus \mathbb{S}} e^{-\beta_M i_M \cdot \Gamma} + \sum_{j_M \in \mathbb{S}} e^{-\beta_S \omega - \beta_M (j_M \oplus 1) \cdot \Gamma} - \mathcal{Z}_M \right). \quad (\text{C3})$$

But we can also express the partition function of the machine using our bit-wise notation as

$$\mathcal{Z}_M = \prod_{j=1}^n (1 + e^{-\beta_M \gamma_j}) = \sum_{l_M \in \{0,1\}^{\times n}} e^{-\beta_M l_M \cdot \Gamma} \quad (\text{C4})$$

$$= \sum_{l_M \in \{0,1\}^{\times n} \setminus \mathbb{S}} e^{-\beta_M l_M \cdot \Gamma} + \sum_{l_M \in \mathbb{S}} e^{-\beta_M l_M \cdot \Gamma}, \quad (\text{C5})$$

where we have split the partition function into a contribution from S and $\{0, 1\}^{\times n} \setminus \mathbb{S}$ since these set complements form natural disjoint subsets the union of which is $\{0, 1\}^{\times n}$. This reorganization of the partition function allows us to see that the populations of the ground-state subspace that were *unchanged* during cooling cancel out with a contribution of the partition to give

$$\Delta p_0 = \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(\sum_{j_M \in \mathbb{S}} e^{-\beta_S \omega - \beta_M (j_M \oplus 1) \cdot \Gamma} - \sum_{l_M \in \mathbb{S}} e^{-\beta_M l_M \cdot \Gamma} \right). \quad (\text{C6})$$

We now have a difference of two sums over the same set, meaning that l_S is a free index which we can set to $l_S = j_M$, and recalling that $E(j_M \oplus 1) = (j_M \oplus 1) \cdot \Gamma = E_{\text{Max}} -$

$j_M \cdot \Gamma$ we obtain

$$\Delta p_0 = \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(\sum_{j_M \in \mathbb{S}} [e^{-\beta_S \omega - \beta_M (E_{\text{Max}} - j_M \cdot \Gamma)} - e^{-\beta_M j_M \cdot \Gamma}] \right). \quad (\text{C7})$$

1. Change in Population Lower Bounds

The most basic lower bound on this quantity is given by the contribution $j_M = 1^n$ corresponding to the $|0_S 1 \dots 1\rangle$ ket which is swapped in every cooling scenario and is the largest contributor to the change in population with $1^n \cdot \Gamma = \sum_{j=1}^n \gamma_j = E_{\text{Max}}$, giving the lower bound

$$\frac{e^{-\beta_S \omega} - e^{-\beta_M E_{\text{Max}}}}{\mathcal{Z}_S \mathcal{Z}_M} \leq \Delta p_0^*, \quad (\text{C8})$$

where Δp_0^* is the optimal change in the ground-state population of the system qubit. This contribution is due to what is known in the literature as the *virtual-qubit subspace swap* [5,6,38].

As mentioned in the main text, if one considers the case in which $(T_M/T_S)\omega \leq \gamma_1$, then a tighter lower bound can be found by noting that all bit strings above and including $i_M = 10^{n-k-1}1^k$, where $k = n/2$ if n is even and $k = (n-1)/2$ if n is odd, are guaranteed to satisfy Eq. (7) for all Γ . To see this, inputting i_M into Eq. (7) says that the

population of i_M should be swapped with $i_M \oplus 1$ if

$$\frac{T_M}{T_S}\omega < \sum_{j=\tilde{k}}^n \gamma_j - \sum_{j=2}^{\tilde{k}-1} \gamma_j + \gamma_1, \quad \tilde{k} = \begin{cases} \frac{n}{2} + 1, & n \text{ even}, \\ \frac{n}{2} + \frac{3}{2}, & n \text{ odd}. \end{cases} \quad (\text{C9})$$

Now consider the assumption that Γ is ordered in nondecreasing order, so that $\gamma_i \leq \gamma_j$ if $i \leq j \forall i, j$. From this it can be seen that

$$\gamma_1 \leq \sum_{j=\tilde{k}}^n \gamma_j - \sum_{j=2}^{\tilde{k}-1} \gamma_j + \gamma_1. \quad (\text{C10})$$

As we have also assumed that $(T_M/T_S)\omega \leq \gamma_1$, Eq. (C9) must be satisfied. All bit strings above i_M then lead to an Eq. (7) with a right-hand side larger than Eq. (C9), as for all other bit strings above i_M parts of the negative sum in Eq. (C9) are turned positive and hence it can only be larger. Therefore, under the assumption that $(T_M/T_S)\omega \leq \gamma_1$, at least all the bit strings above and including i_M should be swapped.

To find a lower bound using Eq. (C7) and this set of bit strings, one can first consider performing all swaps from $|0^{n-k}1^k\rangle \rightarrow |1^n\rangle$, and then removing the contribution arising from the inclusion of $|0^{n-k}1^k\rangle$. Initially, one therefore has

$$\begin{aligned} \Delta p_0^{\text{all}} &= \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(\sum_{i_M \in \{0,1\}^{\times(n-k)}} e^{-\beta_S \omega - \beta_M (E_{\text{Max}} - i_M^{n-k} 1^k \cdot \Gamma)} - e^{-\beta_M i_M^{n-k} 1^k \cdot \Gamma} \right), \\ &= \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(e^{-\beta_S \omega} \sum_{i_M \in \{0,1\}^{\times(n-k)}} e^{-\beta_M (i_M^{n-k} 1^k \oplus 1) \cdot \Gamma} - \sum_{i_M \in \{0,1\}^{\times(n-k)}} e^{-\beta_M i_M^{n-k} 1^k \cdot \Gamma} \right), \end{aligned} \quad (\text{C11})$$

where we have used the fact that $E(i_M) + E(i_M \oplus 1) = E_{\text{Max}}$. We now introduce the notation $\Gamma_{1:l} = (\gamma_1, \gamma_2, \dots, \gamma_l)$ to select parts of Γ . Using this, the above expression becomes

$$\begin{aligned} \Delta p_0^{\text{all}} &= \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(e^{-\beta_S \omega} \sum_{i_M \in \{0,1\}^{\times(n-k)}} e^{-\beta_M (i_M^{n-k} \oplus 1) \cdot \Gamma_{1:n-k}} - e^{-\beta_M \sum_{l=\tilde{k}}^n \gamma_l} \sum_{i_M \in \{0,1\}^{\times(n-k)}} e^{-\beta_M i_M^{n-k} \cdot \Gamma_{1:n-k}} \right), \\ &= \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(e^{-\beta_S \omega} - e^{-\beta_M \sum_{l=\tilde{k}}^n \gamma_l} \right) \sum_{i_M \in \{0,1\}^{\times(n-k)}} e^{-\beta_M (i_M^{n-k} \cdot \Gamma_{1:n-k})}, \end{aligned} \quad (\text{C12})$$

where the constants from the fixed 1s have been factorized out and, the fact that the sets $\{i_M \oplus 1 : i_M \in \{0,1\}^{\times n}\}$ and $\{0,1\}^{\times n}$ are equivalent means that

$$\sum_{i_M \in \{0,1\}^{\times(n-k)}} e^{-\beta_M (i_M^{n-k} \oplus 1) \cdot \Gamma_{1:n-k}} = \sum_{i_M \in \{0,1\}^{\times(n-k)}} e^{-\beta_M (i_M^{n-k} \cdot \Gamma_{1:n-k})}. \quad (\text{C13})$$

From here, it is then noted that this is the partition function over the $(n-k)$ hottest qubits and hence Δp_0^{all} can be written as

$$\begin{aligned} \Delta p_0^{\text{all}} &= \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(e^{-\beta_S \omega} - e^{-\beta_M \sum_{l=\tilde{k}}^n \gamma_l} \right) \prod_{j=1}^{\tilde{k}-1} (1 + e^{-\beta_M \gamma_j}), \\ &= \frac{1}{\mathcal{Z}_S \prod_{j=\tilde{k}}^n (1 + e^{-\beta_M \gamma_j})} \left(e^{-\beta_S \omega} - e^{-\beta_M \sum_{l=\tilde{k}}^n \gamma_l} \right), \\ &= \frac{1}{\mathcal{Z}_S \mathcal{Z}_M^{\tilde{k}:n}} \left(e^{-\beta_S \omega} - e^{-\beta_M \sum_{k:n}^n \gamma_i} \right), \quad \mathcal{Z}_M^{\tilde{k}:n} = \prod_{j=\tilde{k}}^n (1 + e^{-\beta_M \gamma_j}). \end{aligned} \quad (\text{C14})$$




FIG. 11. The set $\mathbb{S}$ of ground-state subspace energy levels worth exchanging with energy levels from the excited-state subspace to cool S under different conditions.

One now needs to take into account the change in population included in Eq. (C14) from swapping $|0^{n-k}1^k\rangle$. Performing this swap would give a change in population of

$$\Delta p_0^{\text{corr}} = \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(e^{-\beta_S \omega - \beta_M (E_{\text{max}} - \sum_k^n \gamma_i)} - e^{-\beta_M \sum_k^n \gamma_i} \right). \quad (\text{C15})$$

A lower bound on Δp_0^* in this case can then be found as

$$\begin{aligned} \Delta p_0^* &\geq \Delta p_0^{\text{all}} - \Delta p_0^{\text{corr}}, \\ &\geq \frac{1}{\mathcal{Z}_S \mathcal{Z}_M^k} \left(e^{-\beta_S \omega} - e^{-\beta_M \sum_k^n \gamma_i} \right) \\ &\quad + \frac{1}{\mathcal{Z}_S \mathcal{Z}_M} \left(e^{-\beta_M \sum_k^n \gamma_i} - e^{-\beta_S \omega - \beta_M (E_{\text{max}} - \sum_k^n \gamma_i)} \right). \end{aligned} \quad (\text{C16})$$

In the main text we have then introduced the notation that $E(k) = \sum_{i=k}^n \gamma_i$ for brevity.

APPENDIX D: REDUCIBILITY

The third law in quantum thermodynamics is often interpreted as a statement that cooling a quantum system to its ground state at zero temperature requires diverging resources. One such resource in our setting is the number of constituent qubits n forming our machine, but as we have observed if the energy of these qubits grows too fast it could render energetically weaker parts of the machine useless. In this appendix we give insight into characterizing this situation by introducing the notion of k -reducibility, the scenario in which the k energetically

weakest qubits of the n -qubit machine do not contribute to the cooling scenario. In particular, here we find that degenerate machines always behave irreducibly for some cooling scenario whereas machines the energy structure of which grows exponentially by a factor larger than 2 are always $(n-1)$ -reducible resulting in a *wasteful* bipartite interaction.

Definition D1. A machine is said to be (i) k -reducible if its first k qubits do not contribute to the cooling interaction or (ii) *precisely* k -reducible if the machine cannot be further reduced, i.e., it is not $(k+1)$ -reducible:

- (i) $|0_S 1^k l^{n-k}\rangle \in \mathbb{S} \implies |0_S i^k l^{n-k}\rangle \in \mathbb{S} \forall i^k \in \{0, 1\}^{\times k}$,
- (ii) $\exists |0_S 1^{k+1} l^{n-k-1}\rangle \in \mathbb{S} : |0_S i^{k+1} l^{n-k-1}\rangle \notin \mathbb{S} \forall i^{k+1} \in \{0, 1\}^{\times k+1}$.

Energetically, this definition implies that a k -reducible machine is one that results in a cooling scenario in which if $|0_S 1^k l^{n-k}\rangle$ corresponds to a population that should be exchanged so must every other energetically weaker machine energy level $|0_S i^k l^{n-k}\rangle$. In particular $|0_S 0^k l^{n-k}\rangle \in \mathbb{S}$, giving

$$(i) \quad \frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < \sum_{j=k+1}^n \gamma_j, \quad (\text{D1})$$

which is a necessary condition on the energetic structure of the machine for M to be k -reducible. On the other hand from (ii) we find that a precisely k -reducible machine

would have an energetic structure satisfying

$$(ii) \quad i^{k+1} l^{n-k-1} \cdot \Gamma \leq \frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < 1^{k+1} l^{n-k-1} \cdot \Gamma, \quad (D2)$$

for some $i^{k+1} \in \{0, 1\}^{\times k+1}$ and $l^{n-k-1} \in \{0, 1\}^{\times n-k-1}$.

1. Irreducible Machines

With this definition in mind an irreducible n -qubit machine, one the energetic structure of which grows with increasing n ensuring that all its n qubits contribute to the cooling scenario, cannot be 1-reducible. This will be ensured if condition (i) is never satisfied for $k = 1$; i.e., an irreducible machine induces a swappable set $\mathbb{S}$ which features at least one $|0_S 1^{n-1}\rangle \in \mathbb{S}$ and $|0_S 0^{n-1}\rangle \notin \mathbb{S}$ —or equivalently, if condition (ii) is always satisfied for $k = 0$. Energetically, (ii) being satisfied for $k = 0$ implies that there is always at least one energy level $|0_S 1^{n-1}\rangle$ that must be exchanged for cooling while its bit conjugate $|0_S 0^{n-1}\rangle$ does not cool when exchanged, giving the inequality

$$0 l^{n-1} \cdot \Gamma \leq \frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < 1 l^{n-1} \cdot \Gamma. \quad (D3)$$

This inequality can be restated in two ways which give different insight into when irreducibility occurs. A machine is irreducible if there is at least one energy level $l^{n-1} \in \{0, 1\}^{\times n-1}$ of the $n - 1$ coldest qubits of the machine which satisfies

$$0 \leq \frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) - E(l^{n-1}) < \gamma_1, \quad (D4)$$

where γ_1 is the energy of the warmest qubit in the machine and

$$\gamma_2 \leq E(l^{n-1}) \leq \sum_{i=2}^n \gamma_i. \quad (D5)$$

Alternatively by noting that $E_{\text{Max}} - 2E(l^{n-1}) = E(l^{n-1} \oplus 1) - E(l^{n-1})$ one can restate this inequality as a set of intervals given by energy changes induced by swapping Hamming-weight conjugate pairs,

$$\begin{aligned} & E(l^{n-1}) - E(l^{n-1} \oplus 1) \\ & \leq \frac{T_M}{T_S} \omega < E(l^{n-1}) - E(l^{n-1} \oplus 1) + 2\gamma_1. \end{aligned} \quad (D6)$$

This inequality can be visualized by taking the number line, marking the negative energy differences of Hamming-weight conjugate energies $-\widetilde{\Delta E}(l^{n-1}) = E(l^{n-1}) - E(l^{n-1} \oplus 1)$ and considering intervals of

$[-\widetilde{\Delta E}, -\widetilde{\Delta E} + 2\gamma_1)$ and asking if $\frac{T_M}{T_S} \omega$ falls in these regions. If it does, we have an irreducible machine. The distance between neighboring differences—say $\widetilde{\Delta E}(l^{n-1})$ and $\widetilde{\Delta E}(l^{n-1} + 1)$, where $l^{n-1} + 1$ is an abuse of notation to denote the next string above l^{n-1} in lexicographic order—could be much larger than $[-\widetilde{\Delta E}, -\widetilde{\Delta E} + 2\gamma_1)$ giving the possibility of reducibility.

a. When Are Degenerate Machines Reducible?

Using the criteria we have derived we are now in a position to prove that degenerate machines, such as those used in algorithmic and dynamical cooling [8, 14] are irreducible and so increase their cooling potential as the number of constituents forming the machine increases. First, the simple condition in Eq. (D4) takes the form

$$0 \leq \frac{1}{2} \left(\frac{T_M}{T_S} \omega + n\gamma - 2t\gamma \right) < \gamma, \quad (D7)$$

where $E_{\text{Max}} = n\gamma$, $E(l^{n-1}) = t\gamma$ and $t \in [0, n - 1]$. Simplifying, we arrive at

$$-(n - 2t)\gamma \leq \frac{T_M}{T_S} \omega \leq 2\gamma - (n - 2t)\gamma. \quad (D8)$$

As $0 < \omega$, it must be the case that the left-hand side is always positive, which is true if $n/2 < t$. Hence, one gets a set of conditions for each integer $n/2 \leq t \leq n - 1$, where each t gives a condition that is some interval on the real-numbers. In the degenerate case, these intervals capture the whole number line $[0, n\gamma]$, meaning that whenever $(T_M/T_S)\omega \leq n\gamma$, one can find a t such that the above condition is satisfied. More specifically, under the credible assumptions that $T_M \leq T_S$ and $\omega \leq \gamma$, degenerate machines are therefore always irreducible.

As a second approach we consider energy changes between Hamming-weight conjugate pairs and see that $\widetilde{\Delta E}(l^{n-1}) = E(l^{n-1} \oplus 1) - E(l^{n-1}) = (n - t)\gamma - t\gamma$, where t is the number of 1s in l^{n-1} . Comparing with a neighboring energy difference $\widetilde{\Delta E}(l^{n-1} + 1) = (n - t - 1)\gamma - (t + 1)\gamma$ we find that the distance between two points on the number line

$$\begin{aligned} & |\widetilde{\Delta E}(l^{n-1} + 1) - \widetilde{\Delta E}(l^{n-1})| \\ & = |(n - 2t - 2)\gamma - (n - 2\gamma)| = 2\gamma \end{aligned} \quad (D9)$$

for all l^{n-1} indexed by t . Thus in the case of the degenerate machine, we have the remarkable situation that the distance between neighboring differences is always 2γ ; i.e., $|\widetilde{\Delta E}(l^{n-1}), -\widetilde{\Delta E}(l^{n-1} + 1)| = 2\gamma$ for any $l^{n-1} \in \{0, 1\}^{\times n-1}$ which grants degenerate machines the property of being always irreducible. This scenario is visualized in Fig. 12.




FIG. 12. Visualization of the inequality Eq. (D6) which turns irreducibility into a question of whether or not $\frac{T_M}{T_S}\omega$ falls within an interval $[-\widetilde{\Delta E}, -\widetilde{\Delta E} + 2\gamma_1)$. Remarkably, for degenerate machines $\widetilde{\Delta E}(l^n + 1) - \widetilde{\Delta E}(l^{n-1}) = 2\gamma$ so the blue region is equal to the distance between differences, implying that a degenerate machine will always behave irreducibly for some cooling scenario $\frac{T_M}{T_S}\omega$, as there are no gaps between blue regions.

2. $(n-1)$ -Reducible Machines

The most wasteful kind of n -qubit machine constructed to cool a quantum system would be one where the optimal cooling interaction is merely bipartite, involving only the energetically strongest (coldest) qubit in the machine and rendering the other $n-1$ qubits in the machine useless. In such a machine, these $n-1$ qubits play no role in the cooling interaction and are in this sense only useful in the construction of the machine, to ramp up to the energy of the last qubit γ_n . But if one had access to a singular qubit with energy equal to that of the coldest qubit in an $(n-1)$ -reducible machine, one could cool equivalently. Let us investigate under what energetic conditions this can happen. First, condition (i) for $(n-1)$ -reducibility states that $|0_S 1^n\rangle$ being an energy level the exchange of which is cooling implies that every energetically weaker $|0_S i^{n-1} 1\rangle$ is also worth swapping $\forall i^{n-1} \in \{0, 1\}^{n-1}$. In particular even the weakest level $|0_S 0^{n-1} 1\rangle \in \mathbb{S}$ giving the first condition for $(n-1)$ -reducibility,

$$\frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < \gamma_n. \quad (\text{D10})$$

A precisely $(n-1)$ -reducible machine must still be able to cool, which is to say it is not n -reducible and there is at least one energy γ_j for which

$$0 \leq \frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) - \gamma_j < \gamma_1, : j \in [2, n]. \quad (\text{D11})$$

a. Exponential Machines

We call an n -qubit machine with gap vector $\Gamma = (\gamma, \gamma^2, \gamma^3, \dots, \gamma^n)$ where the energetic gap of each qubit is exponentially larger than that of its predecessor, an exponential machine. Exponential machines are susceptible to being reducible if the growth of their energy structure is so rapid that a single qubit is able to bring the system to its ground state on its own. To avoid this scenario and figure out if exponential machines can ever be irreducible let us investigate the k -reducibility of exponential machines.

To begin with, recall that finite progressions of the geometric series sum to give $\sum_{i=0}^n ar^i = a(1 - r^{n+1}/1 - r)$ for $r \neq 1$. We can use this to express the sum of gaps in the exponential machine as $E_{\text{Max}} = (1 - \gamma^{n+1}/1 - \gamma) - 1 = \gamma(1 - \gamma^n/1 - \gamma)$. Now checking the $(n-1)$ reducibility we find

$$\frac{1}{2} \left(\frac{T_M}{T_S} \omega + \gamma \left(\frac{1 - \gamma^n}{1 - \gamma} \right) \right) < \gamma^n \quad (\text{D12})$$

$$\frac{T_M}{T_S} \omega < \frac{\gamma(2\gamma^n - 2\gamma^n - 1)}{1 - \gamma} \quad (\text{D13})$$

which, since $1 < (2\gamma^{n-1} - 2\gamma^n - 1)/(1 - \gamma)$ for all $1 < \gamma$, means that Eq. (D13) holds if $(T_M/T_S)\omega < \gamma$ and $1 < \gamma$. Hence, this is a sufficient condition for an exponential machine to be $(n-1)$ -reducible.

APPENDIX E: TOWARD CONNECTIONS BETWEEN COOLING AND ERROR-CORRECTING CODES

Several works in the community have pointed out the connection between error correction and cooling quantum systems. Despite this, explicit connections framing error-correcting protocols as quantum cooling protocols, particularly within the framework of thermal machines, are still to be made. In this section we explore this connection using the tools we have developed in this work by examining the most basic quantum error-correcting code, the three-code repetition code, rephrasing it as a cooling problem.

1. Repetition Code

This code encodes a single-qubit pure state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ into the three-qubit subspace spanned by $|000\rangle, |111\rangle$ as $|\psi_L\rangle = \alpha|000\rangle + \beta|111\rangle$. Here, $|000\rangle, |111\rangle$ are known as the code words of the code and the space they span, the logical subspace or code space $\mathcal{C}$. The code can detect and correct for any single-qubit σ_X or bit-flip error as any $\sigma_X^{(i)}$ error applied to one of the qubits rotates the encoded state to one of three orthogonal




FIG. 13. The three-qubit repetition code that corrects for single-qubit σ_x errors. (a) In the computational basis we have the code words $|000\rangle, |111\rangle$ giving the encoding $|\psi_L\rangle = \alpha|000\rangle + \beta|111\rangle$. (b) In the Greenberger-Horne-Zeilinger- (GHZ) pair basis we have the code words $|\mu_{000}^+\rangle = \frac{|000\rangle + |111\rangle}{\sqrt{2}}, |\mu_{000}^-\rangle = \frac{|000\rangle - |111\rangle}{\sqrt{2}}$ giving the encoding $|\psi_L\rangle = c_+|\mu_{000}^+\rangle + c_-|\mu_{000}^-\rangle$, where $c_{\pm} = \sqrt{2}(\alpha \pm \beta)$.

error subspaces as depicted in Fig. 13(a), $\mathcal{E}_3$ spanned by $|001\rangle, |110\rangle$ if $\sigma_x^{(3)}$ occurs, $\mathcal{E}_2$ given by $|010\rangle, |101\rangle$ if $\sigma_x^{(2)}$ occurs and $\mathcal{E}_1$ spanned by $|100\rangle, |011\rangle$ if $\sigma_x^{(1)}$ occurs. This code has stabilizers $S_1 = \sigma_z^{(1)}\sigma_z^{(2)}$ and $S_2 = \sigma_z^{(2)}\sigma_z^{(3)}$ as both S_1 and S_2 leave the code words invariant but together their expectation values can jointly determine which error has occurred via the syndrome Table I.

2. Change of Basis

In trying to reformulate this protocol into a thermodynamic one, two problems arise. The first, a general state $|\psi\rangle$ encoded for protection against error can have coherences in the code space. This means that the setting we have considered for cooling considering only permutations on the diagonal of a joint system-machine state may not *a priori* suffice. Second, the code words $|000\rangle, |111\rangle$ are very disparate in energy with respect to a local qubit Hamiltonian, i.e., the code space is spanned by the energetically lowest and highest states preventing a straightforward reformulation of the problem into a thermodynamic one. A change of basis may address both issues. Consider the rotation

$$|000\rangle \rightarrow \frac{|000\rangle + |111\rangle}{\sqrt{2}}, \quad |111\rangle \rightarrow \frac{|000\rangle - |111\rangle}{\sqrt{2}}$$

on the code words. In the computational basis, this addresses both issues as it block diagonalizes the coherences of the encoded state and renders the code words energy degenerate with respect to a local qubit Hamiltonian. For the general unitary, split the computational basis into two disjoint sets of Hamming-weight conjugate pairs, in this case $x \in S = \{000, 001, 010, 100\}$ and $\bar{x}_i = x_i \oplus 1 \in \bar{S}$ such that $\{0, 1\}^{\times 3} = S \cup \bar{S}$ and let us define

$|\mu_x^{\pm}\rangle = (|x\rangle + |\bar{x}\rangle)/\sqrt{2}$. Then the desired change of basis is obtained via the unitary

$$V = \sum_{x \in S} |\mu_x^+\rangle\langle x| + |x\rangle\langle \mu_x^+| + |\mu_x^-\rangle\langle \bar{x}| + |\bar{x}\rangle\langle \mu_x^-|, \quad (\text{E1})$$

which can be simply obtained by applying a Hadamard gate to the first qubit, followed by CNOT₁₂ and CNOT₁₃ for any three-qubit string. This change of basis rotates the code, giving the encoding $|\psi_L\rangle = c_+|\mu_{000}^+\rangle + c_-|\mu_{000}^-\rangle$, where $c_{\pm} = \sqrt{2}(\alpha \pm \beta)$, while preserving the stabilizers as one can see in the syndrome Table I.

3. Error Correction as a Cooling Problem

In this basis, an error or correction merely permutes block-diagonal subspaces bringing us to a setting closer to that considered in the cooling protocols presented in this work. To fully reformulate this problem as a thermodynamic one, we must now introduce a Hamiltonian which energetically discriminates between each subspace. Consider the Hamiltonian

$$H_{\text{Code}} = \sum_{x \in S} E_x |\mu_x^{\pm}\rangle\langle \mu_x^{\pm}|, \quad (\text{E2})$$

where $E_x < E_{x'}$ if $x < x'$ in lexicographic order. This Hamiltonian energetically sees the states spanning a given subspace as energy degenerate but can energetically discriminate between the code space and individual error subspaces. In this way, the error-free logical state $|\psi_L\rangle = c_+|\mu_{000}^+\rangle + c_-|\mu_{000}^-\rangle$ is a zero-temperature thermal state of this Hamiltonian. The application of an amplitude-damping channel, i.e., probabilistic bit-flip errors, results in a thermal distribution over the Hamiltonian H_{Code} , giving a state $\rho_{\beta} = e^{-\beta H_{\text{Code}}} / \mathcal{Z}$. With access to an auxiliary

TABLE I. Repetition code and rotated repetition code syndrome tables.

Subspace	Span	$\langle \sigma_z^{(1)} \sigma_z^{(2)} \rangle$	$\langle \sigma_z^{(2)} \sigma_z^{(3)} \rangle$
$\mathcal{C}$	$ 000\rangle, 111\rangle$	1	1
$\mathcal{E}_1$	$ 100\rangle, 011\rangle$	-1	1
$\mathcal{E}_2$	$ 010\rangle, 101\rangle$	-1	-1
$\mathcal{E}_3$	$ 001\rangle, 110\rangle$	1	-1

Subspace	Span	$\langle \sigma_z^{(1)} \sigma_z^{(2)} \rangle$	$\langle \sigma_z^{(2)} \sigma_z^{(3)} \rangle$
$\mathcal{C}$	$\frac{ 000\rangle \pm 111\rangle}{\sqrt{2}}$	1	1
$\mathcal{E}_1$	$\frac{ 100\rangle \pm 011\rangle}{\sqrt{2}}$	-1	1
$\mathcal{E}_2$	$\frac{ 010\rangle \pm 101\rangle}{\sqrt{2}}$	-1	-1
$\mathcal{E}_3$	$\frac{ 001\rangle \pm 110\rangle}{\sqrt{2}}$	1	-1

quantum system such as a thermal machine this three-qubit system can be unitarily cooled using the methods developed in this work, giving

$$\rho_{\beta'} = \text{tr}_M \left\{ U_{\text{Cool}}(\rho_\beta \otimes \rho_M) U_{\text{Cool}}^\dagger \right\}, \quad (\text{E3})$$

where $\beta' > \beta$, bringing one closer to the logically encoded state.

APPENDIX F: COOLING WITH SYMMETRIES

In the main text, we have assumed access to any unitary when cooling the system (although ultimately we have argued that only permutations in the energy eigenbasis are necessary). Here, we consider restricting the allowed dynamics due to the existence of some symmetry, and assess its effect on our ability to optimally cool the system.

Given some observable Q , we will now consider unitaries U such that $[U, Q] = 0$. Under such dynamics Q is said to be symmetric as the commutator implies that $Q = UQU^\dagger$. Physically, this means that under evolution via U the probabilities of getting a given measurement outcome of Q will remain a constant. As a consequence, the expectation value of Q will also remain constant. Hence, this symmetry in the dynamics is said to have led to the conservation of Q , with this being an example of the Noether theorem.

1. Energy Conservation

First, we will consider $H_T = H_S + H_M$, where

$$H_T = \sum_{j_S \in \{0,1\}} \sum_{i_M \in \{0,1\}^n} (j_S \cdot \omega + i_M \cdot \Gamma) |j_S i_M\rangle \langle j_S i_M|, \quad (\text{F1})$$

such that H_S and H_M are the Hamiltonians of the system and machine, respectively, so that H_T is the total Hamiltonian. The symmetry in the dynamics U , i.e., $[U, H_T] = 0$, then leads to the total energy of the system and machine being conserved. A unitary U commutes with H_T if it acts an independent unitary within each degenerate subspace of

H_T ,

$$U = \bigoplus_c U_c, \quad U_c U_c^\dagger = \mathbb{I}, \quad (\text{F2})$$

such that

$$U_c = \sum_{j_S \in \{0,1\}} \sum_{i_M \in \{0,1\}^n} u_{i_S, i_M} |j_S i_M\rangle \langle j_S i_M|$$

$$: j_S \cdot \omega + i_M \cdot \Gamma = E_c, \quad (\text{F3})$$

where E_c is the energy of the c th degenerate subspace and u_{i_S, i_M} is some matrix coefficient of U_c in the eigenbasis of H_T . Now applying the same logic as in the main text to each degenerate subspace, it can be seen that optimal cooling can be achieved by applying permutations of the energy eigenbasis within each degenerate subspace of H_T . As before, a permutation can only change the ground-state population of the system qubit if it is of the form

$$|0_S i_M\rangle \leftrightarrow |1_S k_M\rangle : i_M, k_M \in \{0, 1\}^{\times n}. \quad (\text{F4})$$

Hence, to achieve any cooling there must exist at least one $i_M, k_M \in \{0, 1\}^{\times n}$ such that $i_M \cdot \Gamma = \omega + k_M \cdot \Gamma$, as this means that there will exist a degenerate subspace of H_T that can facilitate such a TLP. As before, this TLP will be cooling if

$$\frac{T_M}{T_S} \omega < (i_M - k_M) \cdot \Gamma, \quad (\text{F5})$$

although, due to the above equality, this will be satisfied if $T_S > T_M$. This result is expected, as the thermal state of H_T is

$$\tau_{SM} = \frac{1}{\mathcal{Z}_{SM}} e^{-\beta H_T} = \frac{e^{-\beta H_S}}{\mathcal{Z}_S} \otimes \frac{e^{-\beta H_M}}{\mathcal{Z}_M}, \quad (\text{F6})$$

such that the system and machine are locally thermal, each at the some temperature β . As energy-conserving dynamics act invariantly on the thermal state, i.e., $U \tau_{SM} U^\dagger = \tau_{SM}$ if $[U, H_T] = 0$, no cooling of the system can be achieved if $T_M = T_S$ (as no populations can be moved). However, if the system and machine are at different temperatures, they are not initially in a thermal state of H_T . Hence, the

initial state is not invariant under energy-conserving unitaries and there exists the possibility to cool. Therefore, under energy-conserving dynamics, the system can only be cooled using a machine if (a) H_T has degeneracies facilitating TLP of the form given in Eq. (F4), and (b) the initial temperature of the machine is less than the initial temperature of the system, $T_S < T_M$.

The change in ground-state temperature that can be achieved under energy-conserving unitaries will then depend on both the dimension of the degenerate subspaces that can facilitate permutations of the form of Eq. (F4)—as this dictates the number of TLPs that can be performed within the given degenerate subspace—and the temperature difference between the system and machine. To see this, it is noted that the change in the ground-state population from performing the TLP $|0_S i_M\rangle \leftrightarrow |1_S k_M\rangle$, within the E_c th degenerate subspace, is given by

$$\begin{aligned} \Delta p_c &= \frac{e^{-\beta_S \omega - \beta_M k_M \cdot \Gamma} - e^{-\beta_M i_M \cdot \Gamma}}{\mathcal{Z}_S \mathcal{Z}_M} \\ &= \frac{e^{-\beta_M E_c} (e^{\omega(\beta_M - \beta_S)} - 1)}{\mathcal{Z}_S \mathcal{Z}_M}, \end{aligned} \quad (\text{F7})$$

where we have used the fact that $i_M \cdot \Gamma = \omega + k_M \cdot \Gamma = E_c$. Hence, it can be seen that the larger the disparity between T_S and T_M , the larger $\beta_M - \beta_S$, and the larger is the possible change in ground-state energy.

To solidify the notions introduced here, we will now consider energy-conserving unitaries in the motivating example of a two-qubit machine considered in the main text. Given the assumption of the gap vector being ordered in nonincreasing order, $\omega \leq \gamma_1 \leq \gamma_2$, there are two possible ways in which H_T can have degeneracies that facilitate TLPs of the form of Eq. (F4): (1) $\omega + \gamma_1 = \gamma_2$, and (2) $\omega = \gamma_1 = \gamma_2 = \gamma$.

In case (1), there is a single degenerate subspace: $\text{Span}\{|0_S 01\rangle, |1_S 10\rangle\}$. Hence, the TLP $|0_S 01\rangle \leftrightarrow |1_S 10\rangle$ is all that can be performed, leading to an optimal change in ground-state population of the system of

$$\Delta p^*(1) = \frac{e^{-\beta_M \gamma_2} (e^{\omega(\beta_M - \beta_S)} - 1)}{\mathcal{Z}_S \mathcal{Z}_M}. \quad (\text{F8})$$

In case (2), there are two degenerate subspaces:

$$\begin{aligned} &\text{Span}\{|0_S 01\rangle, |0_S 10\rangle, |1_S 00\rangle\}, \\ &\text{Span}\{|0_S 11\rangle, |1_S 01\rangle, |1_S 10\rangle\}. \end{aligned} \quad (\text{F9})$$

In the first subspace, either the TLP $|0_S 01\rangle \leftrightarrow |1_S 00\rangle$ or $|0_S 10\rangle \leftrightarrow |1_S 00\rangle$ must be performed. An equivalent choice of TLPs then exists for the other degenerate subspace. Taken together, they lead to an optimal change in

ground-state population of

$$\Delta p^*(2) = \frac{(e^{-\beta_M \omega} + e^{-2\beta_M \omega}) (e^{\omega(\beta_M - \beta_S)} - 1)}{\mathcal{Z}_S \mathcal{Z}_M}. \quad (\text{F10})$$

2. Additional Conserved Quantities

In addition to energy, one can also consider other observables that must be conserved under the cooling dynamics. For example, one can consider symmetry with respect to

$$\begin{aligned} J_Z &= \sigma_Z^S \otimes \mathbb{I}_M + \mathbb{I}_S \otimes \sum_{i=1}^n \sigma_Z^i, \\ \sigma_Z^i &= \mathbb{I}_1 \otimes \mathbb{I}_2 \dots \mathbb{I}_{i-1} \otimes \sigma_Z \otimes \mathbb{I}_{i+2} \dots \mathbb{I}_n, \end{aligned} \quad (\text{F11})$$

which is the total spin observable with constants set to 1. As before, the unitaries that conserve total spin are those that commute with J_Z , $[U, J_Z] = 0$, and hence they act invariantly within the degenerate subspaces of J_Z . If both H and J_Z are conserved, one therefore needs to consider unitaries U such that $[U, H] = [U, J_Z] = 0$. Such a condition can be very restrictive, allowing arbitrary unitaries to act only within the joint degenerate eigenspaces of H and J_Z , although if one assumes that the energy eigenvectors and spin eigenvectors align, such that $[H, J_Z] = 0$, the problem becomes more tractable. For example, in the case of the two-qubit motivating example, the total spin operator in the joint eigenbasis under this assumption is

$$J_Z = \text{diag}(3, 1, 1, -1, 1, -1, -1, -3). \quad (\text{F12})$$

Hence, it has the same degenerate subspaces as case (2) detailed above. Therefore, in case (2) both H_T and J_Z can be conserved while performing cooling. However, in case (1) there are no overlapping degenerate subspaces between J_Z and H_T , meaning that it is not possible to perform any cooling while conserving both quantities.

An interesting future extension in this direction would be to consider cooling with multiple conserved quantities, as has been briefly explored in Ref. [30]. While having two conserved quantities in the above example of case (1) has proved restrictive enough to remove any ability to cool, in larger systems it could still be possible to perform cooling despite added restriction. In addition, it would be interesting to characterize in what situations symmetric dynamics become so restrictive that cooling cannot occur. Moreover, it would be timely and interesting to also consider situations in which the conserved quantities do not commute, e.g., $[H_T, J_Z] \neq 0$. Considering noncommuting conserved quantities in the process of thermalization has already proved interesting [46]. It would now be pertinent to consider their role in opposite task: cooling.

APPENDIX G: COOLING A QUBIT WITH n INTERACTING OTHER QUBITS—GENERAL PURIFICATION

Throughout this work we have focused on the task of cooling a qubit, which has been shown to be equivalent to increasing the probability of being in the ground state of the system Hamiltonian. However, given that we allow the application of arbitrary unitaries (when not considering symmetries), once we have increased the probability of being in the ground state we could arbitrarily change the basis of the state. As cooling is only applicable when considering the energy eigenbasis, and increasing the purity of a qubit just means increasing its largest eigenvalue, it would be more appropriate to view the protocol detailed here as a *purity-enhancement* protocol. If our restriction to unitaries that are permutations of the energy eigenbasis was a physical restriction, rather than just a result of them being sufficient for cooling, our protocol would be strictly cooling.

In the same vein, our protocol can easily be adapted to perform purity enhancement of an arbitrary system qubit given access to a machine of arbitrary qubits, i.e., it can be applied more generally than just to thermal states. This can be achieved given that all qubit states can be considered to be a thermal state with respect to some Hamiltonian. Specifically, given some state ρ , with spectral decomposition

$$\rho = \lambda_0 |\lambda_0\rangle\langle\lambda_0| + (1 - \lambda_0) |\lambda_1\rangle\langle\lambda_1|, \quad \lambda_0 \geq 2^{-1}, \quad (\text{G1})$$

an effective Hamiltonian can be defined as

$$H^{\text{eff}} = \gamma |\lambda_1\rangle\langle\lambda_1|, \quad \gamma = -\frac{1}{\beta} \ln \left(\frac{1}{\lambda_0} - 1 \right), \quad (\text{G2})$$

where β is some effective temperature that can be chosen to tune the gap. Hence, given some arbitrary set of states

$$\rho_S \otimes \rho_M = \rho_S \otimes \rho_{M_1} \otimes \rho_{M_2} \otimes \rho_{M_3} \otimes \dots \otimes \rho_{M_n}, \quad (\text{G3})$$

one can define an effective Hamiltonian for the system and each machine qubit such that

$$\rho_S = \frac{e^{-\beta_S H_S^{\text{eff}}}}{Z_S} \quad \text{and} \quad \rho_{M_i} = \frac{e^{-\beta_M H_i^{\text{eff}}}}{Z_i} \quad \forall i, \quad (\text{G4})$$

with total effective Hamiltonian

$$H_T^{\text{eff}} = H_S^{\text{eff}} + \sum_{i=1}^n H_i^{\text{eff}}. \quad (\text{G5})$$

From this, an effective system gap, ω^{eff} , and an effective gap vector, Γ^{eff} , can be defined. The gaps in the machine can be ordered such that the gap vector is in nondecreasing order. Then, if the chosen system qubit has $\lambda_0 \neq 2^{-1}$,

such that the state is not maximally mixed, a β_S can always be chosen for any ρ such that $\omega^{\text{eff}} \leq \gamma_1^{\text{eff}} : \gamma_1^{\text{eff}} \in \Gamma^{\text{eff}}$. The framework is therefore the same as presented above, meaning that the same protocol can be performed, although TLPs of the joint eigenbasis of the system and machine are now considered, rather than TLPs of the energy eigenbasis, e.g., TLPs of the form

$$|\lambda_0^S \lambda_{i_1}^{M_1} \lambda_{i_2}^{M_2} \dots \lambda_{i_n}^{M_n}\rangle \leftrightarrow |\lambda_1^S \lambda_{j_1}^{M_1} \lambda_{j_2}^{M_2} \dots \lambda_{j_n}^{M_n}\rangle, \quad (\text{G6})$$

where $\{|\lambda_0^S\rangle\langle\lambda_0^S|, |\lambda_1^S\rangle\langle\lambda_1^S|\}$ and $\{|\lambda_0^{M_k}\rangle\langle\lambda_0^{M_k}|, |\lambda_1^{M_k}\rangle\langle\lambda_1^{M_k}|\}$ are the eigenbasis of the system (the chosen qubit being purified) and the k th machine qubit, respectively. Therefore, $i_M = i_1 i_2 \dots i_n$ and $j = j_1 j_2 \dots j_n$ are bit strings labeling the eigenstates of H_T which are tensor products of the local eigenstates. These bit strings are then used in determining which inequalities are satisfied, i.e., which TLPs should be performed to purify (cool) the system. The purifying unitary generated by our protocol will then be in this joint eigenbasis, but can of course be decomposed into any basis once it has been found. Finally, we emphasize that the local eigenbasis of each qubit can be different, with the eigenbasis of different qubits not needing to be orthogonal.

In this alternative setup, all of the above results still apply, with the change in ground-state population now being reinterpreted as a change in purity. Hence, given access to any number of qubits, our protocol can be used to find the unitary that maximally increases the purity of some chosen target qubit.

APPENDIX H: CONSTRUCTING COOLING UNITARIES USING HYPERCUBE GRAPHS

1. Hypercube Graphs

The graph $G = (V, E)$ with the set of vertices V corresponding to the set of n -bit strings $\{0, 1\}^{\times n}$ and edges E connecting vertices v_j and v_k if the bit strings j and k differ only at one bit—i.e., they are Hamming distance 1 from each other—is known as Q_n the n -bit hypercube graph [23,53]. These graphs have 2^n vertices each having n edges giving a total of $n2^{n-1}$ edges on the graph. These graphs are well studied in the field of combinatorics due to their simple structure and ubiquity in classical coding theory. Note that any Hamiltonian path on Q_n corresponds to a Gray code between two bit strings.

2. Representing Cooling Problems on Hypercube Graphs

The state of system and machine

$$\tau_S \otimes \tau_M = \sum_{\substack{i_S \in \{0,1\} \\ i_M \in \{0,1\}^{\times n}}} \frac{e^{-(\beta_S i_S \omega + \beta_M i_M \cdot \Gamma)}}{Z_{SM}} |i_S i_M\rangle\langle i_S i_M| \quad (\text{H1})$$

exists in a Hilbert space spanned by eigenkets which are indexed by the set of $n + 1$ bit strings $i_{S i_M} \in \{0, 1\}^{\times n}$. As such, the state $\tau_S \otimes \tau_M$ can be represented on an $n + 1$ hypercube graph since each eigenket $|i_{S i_M}\rangle$ corresponds to a vertex $v_j \in V = \{0, 1\}^{\times n+1}$. To represent the state we add weights $w(v_j)$ to each vertex to give the vertex set W , where $w(v_j) = (e^{i_S \omega + i_M \Gamma}) / \mathcal{Z}_S \mathcal{Z}_M$ for $v_j = i_{S i_M} \in V$ mapping each bit string to its corresponding population in to give $\tau_S \otimes \tau_M$, where $Q_{n+1}^W = (V, E, W)$ is the weighted hypercube graph of $n + 1$ bit strings.

The disordered ground-state subspace energy levels $|0_S k\rangle$ which need be exchanged for cooling correspond to the subgraph $K \subseteq Q_{n+1}$ with vertices v_{0k} with n bit strings $k \in \mathbb{S}$ satisfying the condition

$$\frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right) < k \cdot \Gamma. \quad (\text{H2})$$

This in turn defines the subgraph $K \oplus 1 \subseteq Q_{n+1}$ with vertices $v_{1k \oplus 1}$ of disordered excited-state energy levels $|1_S k \oplus 1\rangle$ satisfying the property

$$(k \oplus 1) \cdot \Gamma \leq \frac{1}{2} \left(\frac{T_M}{T_S} \omega + E_{\text{Max}} \right), \quad (\text{H3})$$

ensuring that whenever the vertices in Q_{n+1} are permuted such that a $v_l \in K$ and $v_m \in K \oplus 1$ are exchanged, S is cooled. Note that since the bit-flip operation is an automorphism of Q_{n+1} , $K \oplus 1$ is a reflection of K through an axis in Q_{n+1} , i.e., $K \simeq K \oplus 1$ the subgraphs of disordered populations are isomorphic to each other.

3. Cooling Unitaries as Minimum-Weight Perfect Matchings

Any *optimal* cooling unitary U_{cool} that cools S with access to M to a state with the largest obtainable ground-state population

$$\Delta p^* = \max_{U_{SM} \in \mathcal{SU}(2^{n+1})} \langle 0_S | \text{tr}_M \left\{ U_{SM} (\tau_S \otimes \tau_M) U_{SM}^\dagger \right\} | 0_S \rangle \quad (\text{H4})$$

under unitary interaction must necessarily take the form

$$U_{\text{cool}} = |A\rangle\langle B| + |B\rangle\langle A| + \mathbb{1}_{\text{Rest}}, \quad (\text{H5})$$

which exchanges the set of disordered ground-state subspace energy levels $|A\rangle = \{|0_S i_M\rangle \mid i_M \in \mathbb{S}\}$ with the disordered excited subspace energy levels $|B\rangle = \{|1_S i_M \oplus 1\rangle \mid i_M \in \mathbb{S}\}$.

As described in Sec. IV a central problem for constructing cooling unitaries is that there are $|\mathbb{S}|!$ such unitaries that achieve Δp^* with potentially different resource costs as quantified by a chosen cost function, e.g., gate complexity, dissipation or energetic cost. In this section we have

also given a proof that the *the Hamming-weight difference between strings j and k lower bounds the gate complexity of a circuit $S_{j,k}$ that permutes their corresponding eigenkets* for the gate set $\mathcal{G} = \{X, CX, CCX, \dots, C^n X\}$. Minimizing the sum of Hamming-weight differences $D_H(a, b)$ between all energy levels to be exchanged for cooling would give a cooling unitary U_{cool}^* which not only obtains the optimal ground-state population Δp^* but does so with minimal gate complexity,

$$U_{\text{cool}}^* = \sum_{\substack{|a\rangle \in |A\rangle \\ |b\rangle \in |B\rangle}} |a\rangle\langle b| + |b\rangle\langle a| + \mathbb{1}_{\text{Rest}} : \min_{\substack{|a\rangle \in |A\rangle \\ |b\rangle \in |B\rangle}} \sum_{\substack{|a\rangle \in |A\rangle \\ |b\rangle \in |B\rangle}} d_H(a, b). \quad (\text{H6})$$

But, how does one pair all the $|a\rangle \in |A\rangle$ and $|b\rangle \in |B\rangle$ to achieve this?

a. Perfect Matching

In graph theory [23] a *matching* $M \subseteq E$ is a subset of the edge set of a graph where each vertex is incident to only one edge in M ; i.e., no two edges in M share a common vertex. A matching is said to be *perfect* if every vertex of the graph is incident once to a single edge in the matching. In other words, a perfect matching M is a collection of edges pairing each vertex to exactly one other vertex in the graph.

If a graph $G = (V, E, W_E)$ also has an edge weight set W_E given by some cost function $w(E)$, then a *minimum-weight perfect matching* is a perfect matching M^* on the graph G where the sum of the weights of the edges in the matching is the minimum out of all possible perfect matchings on G , i.e.,

$$M^* = \arg \min_{\substack{M \subseteq E \\ M \text{ perfect}}} \sum_{e \in M} w(e). \quad (\text{H7})$$

Lemma H1. There is a one-to-one correspondence between edges of the bipartite graph $G(K, K \oplus 1, E)$ and energy-level-exchanging unitaries $S_{j,k}$, where $K = \{k \in \{0, 1\}^{\times n} \mid 0k \in \mathbb{S}\}$ is the subgraph related to disordered ground-state subspace energy levels and $K \oplus 1 = \{k' \in \{0, 1\}^{\times n} \mid k' = 1k \oplus 1 \forall k \in K\}$ disordered excited-state subspace energy levels.

Proof. Let $\mathfrak{S}$ denote the set of unitary-basis ket permutations exchanging exactly two levels,

$$\mathfrak{S} = \{S_{j,k} \in \mathcal{SU}(2^n) \mid S_{j,k} = |j\rangle\langle k| + |k\rangle\langle j| + \mathbb{1}_{\text{Rest}}, j < k\}, \quad (\text{H8})$$

and consider the function $f : \mathfrak{S} \rightarrow E$ which maps the unitary $S(j, k) \in \mathfrak{S}$ to the edge $\{v_j, v_k\} \in E \subseteq \{0, 1\}^{\times n} \times \{0, 1\}^{\times n}$. To show that this function is bijective we must

show that it is injective and surjective. For injectivity consider two edges $\{v_j, v_k\} = \{v_l, v_m\}$: since $v_j, v_k, v_l, v_m \in K \times K \oplus 1$, then $j, k, l, m \in \{0, 1\}^{\times n}$ and so the assumption implies that these n bit strings are equal, $j = l$ and $k = m$. The n -bit strings index the eigenkets of the machine so $|j\rangle = |l\rangle$, $|k\rangle = |m\rangle$ and we find $S_{j,k} = |j\rangle\langle k| + |k\rangle\langle j| + \mathbb{1}_{\text{Rest}} = |l\rangle\langle m| + |m\rangle\langle l| + \mathbb{1}_{\text{Rest}} = S_{l,m}$ giving injectivity, i.e., $f(x) = f(y) \implies x = y$. For surjectivity we wish to show that $\forall e \in E, \exists S_{l,m} \in \mathfrak{S} : f(S_{l,m}) = e$. To see this consider an arbitrary $e = \{v_j, v_k\}$ corresponding to $j, k \in \{0, 1\}^{\times n}$. Then clearly one can construct $S_{j,k} = |j\rangle\langle k| + |k\rangle\langle j| + \mathbb{1}_{\text{Rest}}$ such that $f(S_{j,k}) = \{v_j, v_k\} = e$ as desired. Therefore f is a bijection. ■

As a direction consequence of this lemma we have the main statement of this section, that the problem of finding U_{cool}^* is equivalent to a minimum-weight perfect-matching problem.

Proposition H1. The problem of finding an optimal-cooling unitary U_{cool}^* which minimizes the gate complexity relative to $\mathcal{G} = \{X, CX, C^2X, \dots, C^nX\}$ to cool S with gap ω at temperature T_S given n qubits with gaps $\Gamma = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$ at temperature T_M is equivalent to finding a minimum-weight perfect matching on the bipartite graph $G(K, K \oplus 1, E, d_H)$.

Proof. First, U_{cool}^* can be expressed as a product $U_{\text{cool}}^* = S_{a_1, b_1} S_{a_2, b_2} \dots S_{a_n, b_n}$ of $S_{j,k}$ since each $|a\rangle \in |A\rangle$ is paired to exactly one $|b\rangle \in |B\rangle$ so none of the pairs intersect. Then, by Lemma H1 each $S_{j,k}$ can be represented as an edge $e = \{v_j, v_k\}$ in the bipartite graph $G(K, K \oplus 1, E)$, giving a collection of edges M . From the definition of U_{cool}^* we can infer that M is a matching on $G(K, K \oplus 1, E)$ since each $|a\rangle \in |A\rangle$ corresponds to $v_a \in K$ and each $|b\rangle \in |B\rangle$ corresponds to $v_b \in K \oplus 1$, exhausting K and $K \oplus 1$ such that no vertex is incident on more than one edge. Lastly, considering $G(K, K \oplus 1, E, d_H)$ where the Hamming distance $D_H(j, k)$ between two strings j, k labeling vertices v_j and v_k is added as a weight $w(e) = D_H(v_j, v_k)$ and the

definition of U_{cool}^* where the pairing of eigenkets being exchanged satisfies

$$\min_{\substack{|a\rangle \in |A\rangle \\ |b\rangle \in |B\rangle}} \sum_{\substack{|a\rangle \in |A\rangle \\ |b\rangle \in |B\rangle}} D_H(a, b), \quad (\text{H9})$$

which implies by construction that M is minimum-weight perfect matching on G since it is a matching satisfying $\min_{M' \subseteq E} \sum_{e \in M'} w(e)$ since $w(e) = D_H(v_j, v_k)$ and U_{cool}^* minimizes the Hamming weight across pairs of exchanged energy levels. Therefore any U_{cool}^* corresponds to a minimum-weight perfect matching. ■

4. Hungarian Algorithm for Perfect Matching and Complexity of Cooling

Having shown that the problem of finding a unitary to cool S given M that minimizes the gate complexity given the gate set $\mathcal{G}$ can be recast as a minimum-weight perfect-matching problem we now describe how MWPMs can be found algorithmically and discuss complexity-theoretic implications on quantum cooling in light of this connection.

Minimum-weight perfect matching is a well-studied and seminal problem in the field of combinatorial optimization with applications in network theory. For a more in-depth introduction see Ref. [24]. For complete bipartite graphs, such as $G(K, K \oplus 1, E, d_H)$, the bipartite graph connecting the subgraphs of disordered populations $K, K \oplus 1 \subset Q_{n+1}$ (as visualized in Fig. 14), minimum-weight perfect matchings can be found using the Hungarian algorithm [24].

The matrix formulation of this algorithm for a complete weighted bipartite graph $B(V_1, V_2, E, C)$ with $2m$ vertices is as follows. Consider the $m \times m$ cost matrix C with entries $c_{ij} = c(v_i^{(1)}, v_j^{(2)}) \forall v_i^{(1)} \in V_1, v_j^{(2)} \in V_2$, where $c(\cdot)$ is the cost function determining the weight of an edge $\{v_i^{(1)}, v_j^{(2)}\} \in E$. A perfect matching corresponds to picking m numbers from this matrix, picking once from each




FIG. 14. (a) The hypercube-graph visualization Q_4 for one of the scenarios given in Appendix B. (b) A visualization of the complete bipartite graph $G(K, K \oplus 1, E)$ in this scenario with a minimum-weight perfect matching in orange.

row and column of the matrix. Without loss of generality one can fix the columns of C and permute its m rows using the permutation matrix $P_\pi : \pi \in S_m$. This defines a new matrix $\tilde{C} = P_\pi C$ the diagonal of which corresponds to the cost of a given perfect matching and so the trace of $\tilde{C} = P_\pi C$, $\text{tr}\{\tilde{C}\} = \sum_{i=1}^m \tilde{c}_{ii} = \sum_{e \in M} c(e)$ corresponds to the weight of this perfect matching M . Minimizing this trace over all permutations

$$\min_{\pi \in S_m} \text{tr}\{P_\pi C\} = \text{tr}\{P_{\pi^*} C\}, \quad (\text{H10})$$

one finds the minimum-weight perfect matching M^* given by the set of edges $E^* = \{\{v_i^{(1)}, v_{\pi^*(i)}^{(2)} : i \in (1, 2, \dots, m)\}\}$. The time complexity of this algorithm is $\mathcal{O}(m^3)$ [24].

a. Example

Let us make the connection between cooling unitaries and minimum-weight perfect matching more tangible by returning to examples from Appendix B. In Fig. 14(b) we observed case 2 from Appendix B, i.e., $\omega < \gamma_1 + \gamma_2 - \gamma_3$ for $T_M = T_S$ which we visualize in Fig. 14 for $\omega = \log 2$, $\gamma_1 = \log 3$, $\gamma_2 = \log 22$ and $\gamma_3 = \log 32$ factored by $\frac{32}{69}$. Here, we have the disordered energy levels $|A\rangle = \{|0_S 011\rangle, |0_S 101\rangle, |0_S 110\rangle, |0_S 111\rangle\}$ and $|B\rangle = \{|1_S 100\rangle, |1_S 010\rangle, |1_S 001\rangle, |1_S 000\rangle\}$. Thinking through the Hungarian algorithm for our example we have the cost matrix C with entries given by the Hamming distance between the bit strings labeling energy levels:

$$C = \begin{matrix} & \begin{matrix} |0011\rangle & |0101\rangle & |0110\rangle & |0111\rangle \end{matrix} \\ \begin{matrix} |1100\rangle \\ |1010\rangle \\ |1001\rangle \\ |1000\rangle \end{matrix} & \begin{pmatrix} 4 & \textcolor{red}{2} & 2 & \textcolor{blue}{3} \\ \textcolor{red}{2} & 4 & \textcolor{blue}{2} & 3 \\ 2 & \textcolor{blue}{2} & 4 & \textcolor{red}{3} \\ \textcolor{blue}{3} & 3 & \textcolor{red}{3} & 4 \end{pmatrix} \end{matrix}. \quad (\text{H11})$$

A minimum-weight perfect matching now corresponds to picking four entries from this matrix using every row and column once such that the sum of these four entries is as low as possible. A valid MWPM (highlighted in red) gives the exchanges $|0_S 011\rangle \leftrightarrow |1_S 010\rangle = 2$, $|0_S 101\rangle \leftrightarrow |1_S 100\rangle = 2$, $|0_S 111\rangle \leftrightarrow |1_S 001\rangle = 3$ and $|0_S 110\rangle \leftrightarrow |1_S 000\rangle = 3$ and $2 + 2 + 3 + 3 = 10$ minimizing the sum of interest. This corresponds to the unitary

$$U_{\text{cool}}^* = (S_{0011,1010} \oplus S_{0101,1100}) \oplus (S_{0111,1001} \oplus S_{0110,1000}) \\ = \text{SWAP}_{S, M_3} \oplus S_{011,100_{SM_1 M_2}}, \quad (\text{H12})$$

where we see that a complexity reduction has been obtained since the first pair of two-level swaps features two free indices leading to a bipartite SWAP and the second

pair features one free index corresponding to an irreducible 4-partite operation.

An issue is that there can be multiple minimum-weight perfect matchings obtained through the minimization over C , each of which correspond to cooling unitaries with equal capability to potentially minimize gate complexity—but that do not have equal circuit complexity, as some might lead to unitaries with a lower order of interaction. This is the case comparing with the other valid solution to this example highlighted in blue, given by the exchange $|0_S 111\rangle \leftrightarrow |1_S 100\rangle$, $|0_S 110\rangle \leftrightarrow |1_S 010\rangle$, $|0_S 101\rangle \leftrightarrow |1_S 001\rangle$, $|0_S 011\rangle \leftrightarrow |1_S 000\rangle$, giving the unitary

$$U'_{\text{cool}} = S_{0110,1010} \oplus S_{0101,1001} \oplus S_{0111,1100} \oplus S_{0011,1000} \\ = S_{010,100_{SM_1 M_3}} \oplus S_{0111,1100} \oplus S_{0011,1000}, \quad (\text{H13})$$

which is more complex than U_{cool}^* .

b. Optimizing Energetic Cost

Adding an energetic constraint to reduce the change in energy in the machine to the cost matrix by considering the cost function $c(i, j) = D_H(i, j)(|i - j| \cdot \Gamma)$, one can optimize the *work cost* of the cooling unitary which also turns out to make the *optimal* MWPM unique in this example for $\omega = 1$, $\Gamma = (1, 2, 4)$ unique, giving

$$C = \begin{matrix} & \begin{matrix} |0011\rangle & |0101\rangle & |0110\rangle & |0111\rangle \end{matrix} \\ \begin{matrix} |1100\rangle \\ |1010\rangle \\ |1001\rangle \\ |1000\rangle \end{matrix} & \begin{pmatrix} 16 & \textcolor{red}{6} & 2 & \textcolor{blue}{15} \\ \textcolor{red}{6} & 8 & \textcolor{blue}{0} & 12 \\ 2 & \textcolor{blue}{0} & -8 & \textcolor{red}{6} \\ \textcolor{blue}{15} & 12 & \textcolor{red}{6} & 24 \end{pmatrix} \end{matrix}, \quad (\text{H14})$$

where we see that the cost of U_{cool}^* is 24 which is lower than the cost of U'_{cool} which is 30. Of course the reduction in the number of viable MWPMs one gets by choosing this additional constraint depends on the structure of Γ . It is left to future work to determine the fascinating interplay which we allude to here between unitary complexity, energetic cost and minimum-weight perfect matchings for cooling unitaries.

c. Complexity of Finding an Optimal-Cooling Unitary

The Hungarian method for obtaining a minimum-weight perfect matching on a complete bipartite graph of $2m$ nodes requires $\mathcal{O}(m^3)$ operations [24]. In our setting the complete bipartite graph formed by the disordered energy levels $G(K, K \oplus 1, E, d_H)$ has $2|\mathbb{S}|$ vertices and so finding U_{cool}^* has an algorithmic complexity of $\mathcal{O}(|\mathbb{S}|^3)$ and since $0 \leq |\mathbb{S}| \leq 2^n$ depending on the structure of Γ , where n is




FIG. 15. A visual representation of the elements of the automorphism group of the cube graph Q_3 in terms of quantum gates. A σ_x operation applied to one of these three qubits generates the subgroup $\mathbb{Z}_2^3$ by reflecting two opposite faces of the cube together through a vertical or horizontal line. We use X_i to denote the action of σ_x on the i th qubit. The SWAP gate generates the permutation subgroup S_3 the action of which reflects two opposite faces of the cube through a diagonal plane. Controlled operations act on one face only and so do not preserve the symmetry of $\text{Aut}(Q_3)$. We denote the action of controlled operations with dotted lines and anticcontrolled operations (which act conditioned on 0) using solid lines.

the number of machine qubits, the complexity of this task can be anywhere from polynomial to exponential.

Note that before obtaining U_{cool}^* via minimum-weight perfect matching an agent need necessarily acquire knowledge of $\mathbb{S}$ by examining at most 2^n inequalities. Thus the complexity of cooling is more accurately of order $\mathcal{O}(|\mathbb{S}|) + \mathcal{O}(|\mathbb{S}|^3)$ where the first contribution stems from determining the disordered populations by verifying the inequalities described in previous sections.

5. Automorphisms of the Hypercube Graph

A graph automorphism is an operation on the graph that preserves the vertex-edge connectivity structure in the graph. The automorphism group of a graph $\text{Aut}(G)$ is the set of all such operations. While U_{cool} is generally not an element of the automorphism group, studying the structure of this group can give us insights into which gates are needed for which cooling scenarios.

The automorphism group of the hypercube graph $\text{Aut}(Q_n)$ is the hyperoctohedral group [23] which can be expressed as $\mathbb{Z}_2^n \rtimes S_n$, where $\mathbb{Z}_2^n$ is the group of operations that add a bit string $x \in \{0, 1\}^n$ modulo 2 to the bit strings corresponding to the vertices of the graph effectively given by the possible bit flips and S_n is the group of permutations of n objects.

a. Two-Qubit Machine Example: $\text{Aut}(Q_3)$

In the given example, we have $z_x \in \mathbb{Z}_2^3$ such that $z_x(v_j) = j_0 j_1 j_2 \oplus x_0 x_1 x_2$ for all bit strings j and since there are 2^3 such bit strings x we can add then $|\mathbb{Z}_2^3| = 2^3$. And $S_3 = \{(1)(2)(3), (12)(3), (13)(2), (23)(1), (123), (132)\}$ in cycle notation such that $\pi_{123} |abc\rangle = |cab\rangle$. Therefore $\text{Aut}(Q_3)$ has $48 = 2^3 \times 3!$ elements given by the product of these group which in this case can easily be visualized as the possible rotations on the cube (including inversions). In this way $\mathbb{Z}_2^n \rtimes S_n$ in general has order $2^n \times n!$. Of course this does not give one the full set of $2^n!$ permutations of the diagonal elements of the density matrix but this because we are restricting ourselves to Hamming-weight-1 increasing operations in a given step and as such it is only in very many steps that we can approach the full set of permutation unitaries.

b. Can You Cool a Qubit with Access to Two Qubits Optimally without a Toffoli?

Now that we know that we have these group-theoretic tools at our disposal, can we use them to show that $G_0 = \{X, CX\}$ are not enough to cool a qubit to a temperature lower than that of a machine qubit? In the case in which $|0_S 01\rangle \leftrightarrow |1_S 10\rangle$ and $|0_S 11\rangle \leftrightarrow |1_S 00\rangle$ are both cooling energy exchanges we know that G_0 is enough since $U_{\text{cool}} = \text{SWAP}_{S, M_2}$ which can be generated by CX . But in the case in which $U_{\text{cool}} = |0_S 11\rangle\langle 1_S 00| + |1_S 00\rangle\langle 0_S 11| + \mathbb{1}_{\text{Rest}} = S_{011, 100}$ we can show that this cannot be generated by G_0 as no $S_{j,k}$ can be generated by G_0 .

Proposition H2. It is not possible to cool a qubit S with gap ω and temperature T_S given access to two qubits with temperature T_M and gaps γ_1, γ_2 without a Toffoli gate if

$$\gamma_2 \leq \frac{T_M}{T_S} \omega + \gamma_1.$$

Proof. First, if $\gamma_2 \leq (T_M/T_S)\omega + \gamma_1$, then the only energy-level exchange which cools is $|0_S 11\rangle \leftrightarrow |1_S 00\rangle$ that is $U_{\text{cool}} = S_{011, 100}$. But no nontrivial $S_{j,k}$ can be produced by X and CX . To see this, note that the application of X and SWAP give proper subgroups of $\text{Aut}(Q_3)$ and so give a coset group structure which partitions $\text{Aut}(Q_3)$. We can represent the action of the subgroups the semiproduct of which gives the automorphism group on the cube using the quantum gates X and SWAP as

$$\mathbb{Z}_2^3 = \{\mathbb{I}, X_1, X_2, X_3, X_1 X_2, X_1 X_3, X_2 X_3, X_1 X_2 X_3\},$$

$$S_3 = \{\mathbb{I}, \text{SWAP}_{1,2}, \text{SWAP}_{2,3}, \text{SWAP}_{1,3}, \text{SWAP}_{1,2} \text{SWAP}_{2,3}, \text{SWAP}_{1,3} \text{SWAP}_{2,3}\}.$$

As such, xS_3 for every $x \in \mathbb{Z}_2^3$ partitions the action of the group into eight orbits each of size 6 and similarly we have

the partitioning $s\mathbb{Z}_2^3$ for every $s \in S_3$ which gives six orbits each of size 8. A basis ket permutation $S_{j,k}$ flips a single edge of the cube breaking the edge connectivity of the graph and so is not an automorphism of Q_3 and does not respect the symmetry of $\text{Aut}(Q_3)$. As such, it cannot be carried out using products of X and SWAP . CX and CSWAP are similarly not automorphisms of Q_3 as they act on a single face of the cube and break edge connectivity so perhaps they could generate an $S_{j,k}$. But note that the actions CX and CSWAP act within the orbits of the partitions xS_3 and $s\mathbb{Z}_2^3$, respectively. Focusing on CX , this operation acts on a single face of the cube, flipping two edges at once. Let us say that one desires one of the edge flips and not the other to give one $S_{j,k}$. Then, one could carry out $S_{j,k}$ if the action of the unwanted flip could be reversed under the action of another X or CX , but this is not possible since both X and CX act on partitions xS_3 and so reversing the action of an unwanted flip of a given CX_i is not possible using X_j or CX_j from another orbit since they form a partition. Therefore, it is not possible to generate $S_{j,k}$ using $\{X, CX, \text{SWAP}\}$, as SWAP can be generated by CX . As a result, in the scenario in which $U_{\text{cool}} = S_{011,100}$ one cannot optimally cool using $\{X, CX, \text{SWAP}\}$. The necessity of CCX is then apparent by noting any Gray code from 011 to 100, e.g., 011, 111, 110, 100, each step of which can be done with X and CCX . In summary, we have shown that X, CX are not sufficient to generate $S_{011,100}$ and that X, CCX are sufficient.

An alternative way to see this is that to optimally cool the system in this case an odd permutation is needed, as only a single pair of elements is permuted. As X and CX apply even permutations, and combinations of even permutations are also even permutations, no combination of these gates can implement the necessary odd permutation to optimally cool. However, when restricting to three qubits CCX is able to implement this odd permutation. Hence, given access to CCX one can optimally cool in this case. ■

Examining how the general automorphism group of the n -bit hypercube-graph partitions under $G = \{X, CX, \text{CCX}, \dots, C^{n-1}X\}$ and deriving general no-go theorems in different k -reducibility scenarios could be an interesting direction for future research.

APPENDIX I: THE SWAPPABLE SET

The problem of cooling a qubit with access to n other qubits reveals a mathematical structure on the set of n bit strings in terms of two partial orders. The goal of this section is to introduce these partial orders and in doing so to characterize the *swappable set* $\mathbb{S} := \{j_M \in \{0, 1\}^{\times n} \mid \frac{1}{2}((T_M/T_S)\omega + E_{\text{Max}}) < E(j_M)\}$, where $E(j_M) := j_M \cdot \Gamma$, which contains the bit strings denoting

energy levels with populations that need to be swapped in order to maximally cool the system.

One can combine these partial orders to reveal a forced structure on the energetics of the system and machine. This mathematical structure can be used to obtain some results that stand at the root of problem of how to identify which states need to be swapped and which should not. Fascinatingly, these results hold no matter the energy structure of the system and machine, as they arise directly from the Hilbert-space structure of multiple qubits. We can therefore identify two parts of the general problem of identifying which are the elements of the swappable set: (1) the mathematical structure imposed by the Hilbert-space structure of qubits, and (2) the specification of the energetic structure of the system and machine.

1. Lexicographic Partial Order

By definition of the swappable set, we can note that if for two bit-strings j_M and k_M we have $E(j_M) \geq E(k_M)$, then $k_M \in \mathbb{S} \implies j_M \in \mathbb{S}$. That is, if k_M corresponds to an energy level that should be exchanged, then j_M should also be exchanged to cool $\mathbb{S}$ —in this $\mathbb{S}$ is ordered.

With knowledge of the refined energy structure of Γ , energy imposes a total order on the bit strings. That is, the map $E : \{0, 1\}^{\times n} \rightarrow \mathbb{R}$ defines a total order on the bit strings of the swappable set. However, we can weaken the total order defined by E to find a complete partial order that is independent of the detailed structure of Γ and only its global properties, i.e., that it is formed of nondecreasing energetic gaps. In what follows, we will describe and examine these partial orders on $\mathbb{S}$. To begin, we will denote the partial-order relation between bit strings (in the text) by the symbol $\succeq$ and between kets (in figures) by $\rightarrow$.

Definition II. Hamming-weight partial order. $j_M, k_M \in \{0, 1\}^{\times n}$ are $j_M \succeq_{\#} k_M$ if and only if $j_i \geq k_i$ for all $i = 1, \dots, n$.


This partial order encodes the action of flipping 0s of a bit string into 1s: if we can generate j_M by changing 0s of k_M into 1s, then we have $j_M \succeq_{\#} k_M$ and $|k_M\rangle \rightarrow_{\#} |j_M\rangle$. The three defining properties of a weak partial order are: (1) *reflexivity* $a \succeq a$, (2) *antisymmetry* if $a \succeq b$ and $b \succeq a$ then $a = b$, and (3) *transitivity* if $a \succeq b$ and $b \succeq c$ then $a \succeq c$. These three properties are satisfied naturally by $\succeq_{\#}$. Furthermore the condition, if $j_M \succeq_{\#} k_M$ and $\exists i$ such that $j_i > k_i$ renders this partial order a strict or strong partial order: $j_M \succ_{\#} k_M$. Naturally, this relation carries over to the energy of the corresponding ket,

$$E(j_M) = \sum_{i=1}^n j_i \gamma_i \geq \sum_{i=1}^n k_i \gamma_i = E(k_M),$$

where we have used the positivity of the gaps. We can therefore conclude that the total order induced by $E(\cdot)$

follows from the Hamming-weight partial order with additional ordering constraints.

Graphically, for $n = 3$ we can map this order on the kets to obtain the following directed graph:



This graph is simply a directed version of Q_3 . At this point we can see that this graph seems to suggest to group the bit strings by column, i.e., $\{0, 1\}^{\times n} = \bigcup_{k=0}^n \mathbb{B}_k : \bigcap_{k=0}^n \mathbb{B}_k = \emptyset$, where $\mathbb{B}_k$ contains all the bit strings with fixed Hamming weight k —i.e., $j_M \in \{0, 1\}^{\times n}$ is an element of $\mathbb{B}_k$ if and only if its number of ones $\#(j_M)$ is equal to k . The Hamming-weight partial order cannot compare two elements with the same Hamming weight. We can therefore use the fact that we can always reorder the machine gaps so that they are increasing $\gamma_1 \leq \gamma_2 \leq \dots \leq \gamma_n$ to obtain a partial order within each $\mathbb{B}_k$.

For $j_M \in \mathbb{B}_\ell$ let us label the indices where j_M has a 1 as $j_M(i)$, denoting that the i th bit of j_M is a 1.

Definition I2. Positional partial order. $j_M, k_M \in \mathbb{B}_\ell$ are $j_M \succeq_p k_M$ if and only if $j_M(i) \geq k_M(i)$ for all $i \in [1, \ell]$.


This partial order encodes the action of swapping a 10 for a 01 in the bit string. Therefore if we can construct a j_M from a k_M via this action, then we have $j_M \succeq_p k_M$ and $|k_M\rangle \rightarrow_p |j_M\rangle$. Similarly to $\succeq_\#$, $\succeq_p$ satisfies the defining properties of a weak partial order, can be made into the strong partial order $\succ_p$ by requiring $\exists i$ such that $j_M(i) > k_M(i)$, and the order carries to the energies of the kets

$$E(j_M) = \sum_{i=1}^{\ell} \gamma_{j_M(i)} \geq \sum_{i=1}^{\ell} \gamma_{k_M(i)} = E(k_M),$$

where we have used that we can rewrite the energy as $E(k_M) = \sum_{j=1}^n k_j \gamma_j = \sum_{i=1}^{\ell} \gamma_{k_M(i)}$ since $k_j = 1$ for $j \in \{k_M(i)\}_{i=1}^{\ell}$ and $k_j = 0$ if $j \notin \{k_M(i)\}_{i=1}^{\ell}$. This allows us to conclude that the total order induced by E follows from the positional partial order with additional conditions. Conceptually, this partial order compares which out of two bit strings i_M, j_M has more 1s on the right and so features higher energetic gaps contributing to their energies $E(i_M)$ and $E(j_M)$.

If we map this order onto a directed graph of kets we obtain a directed graph the structure of which is analogous to that of an antisymmetric ℓ -tensor over an n -dimensional

space. For $n = 5$ and $\ell = 2$ we obtain




This graph can be obtained by assigning a *direction* or degree of freedom for each 1 and moving along the designated direction when the corresponding 1 is shifted, resulting in a graph that expands in ℓ directions—e.g., for $\ell = 3$ vertically for the third 1, horizontally (x) for the second, and horizontally (y) for the first.


Since the Hamming-weight partial order only compares elements across two different $\mathbb{B}_\ell$ and the positional partial order only compares within a single $\mathbb{B}_\ell$, it is very natural to try to combine the two by finding a single partial order that does not depend on the *fine details of the energy gaps* and extends both of these.

Definition I3. Lexicographic partial order. $j_M, k_M \in \{0, 1\}^{\times n}$ with $J = \#(j_M)$ and $K = \#(k_M)$. Then $j_M \succeq k_M$ if and only if $J \geq K$ and $j_M(i + J - K) \geq k_M(i)$ for all $i = 1, \dots, K$.

Reflexivity and antisymmetry follow from $\succeq_p$ being a partial order. For transitivity we only need to note that if $j_M \succeq k_M$ and $k_M \succeq l_M$, then we can shift $i + J - K$ by $K - L$ to obtain $j_M(i + J - L) \geq k_M(i + K - L) \geq l_M(i)$ for all $i = 1, \dots, L = \#(l_M)$.

Similarly to the previous two orders, we can make the lexicographic order strict: $j_M \succ k_M$ if and only if $j_M \succeq k_M$ and either $J > K$ or $j_M \succ_p k_M$. This order is truly an extension of both the positional order and the Hamming-weight order: we can see this because we can define $l_M \in \mathbb{B}_K$ such that $l_M(i) = j_M(i + J - K)$ for $i = 0, \dots, K$; then by construction we have $l_M \succeq_p k_M$ and $j_M \succeq_\# l_M$. This intermediary element also allows us to conclude that the total order induced by E is an extension of the lexicographic order with further ordering constraints. By mapping it onto a graph for $n = 3$ and $n = 4$ we obtain





We underline here that the lexicographic order is a partial order and not a total order as noncomparable elements still remain, e.g., $|0001\rangle$ and $|1100\rangle$, unlike in an ordering of these bit strings with complete knowledge of the energetic structure which induces a total order. We can therefore apply this lexicographic order to the original observation to obtain the following lemma.

Lemma II. Let $j_M, k_M \in \{0, 1\}^{\times n}$ such that $j_M \geq k_M$ and $k_M \in \mathbb{S}$. Then $j_M \in \mathbb{S}$.

Proof. Since the total order induced by $E(\cdot)$ is an extension of $\geq$, we also have $E(j_M) \geq E(k_M)$. Therefore, by the definition of $\mathbb{S}$, we have $j_M \in \mathbb{S}$. ■

From this lemma we can conclude that the characterization of $\mathbb{S}$ without knowledge of each individual $\gamma_j \in \Gamma$ comes down to finding a small collection of bit strings that are incomparable to each other, from which we can generate all the other elements of $\mathbb{S}$.

2. Always- and Never-Swappable Sets

The structure induced by the partial order allows us to note some interesting aspects about the characterization of the swappable set that do not depend on the energetic structure of the machine. In particular we seek to analyze if there are some elements of $\{0, 1\}^{\times n}$ that are always or never part of $\mathbb{S}$; i.e., ground-state subspace energy levels which when exchanged for excited subspace levels always cool or heat S regardless of the structure of M . In order to do so, it is worth noting that the definition of $\mathbb{S}$ is dependent on the system gap ω : as we make ω larger then we can reduce the size of $\mathbb{S}$ to the point that it becomes the empty set for $\omega \geq E_{\text{Max}}$. Therefore, we will consider for a moment the set $\mathbb{S}_0 := \{j_M \in \{0, 1\}^{\times n} | \frac{1}{2}E_{\text{Max}} < E(j_M)\}$. Note that, by definition, $\mathbb{S} \subseteq \mathbb{S}_0$.

To characterize $\mathbb{S}_0$ we are trying to find which energy levels are always larger or smaller than the median $\frac{1}{2}E_{\text{Max}}$. Since $|\{0, 1\}^{\times n}| = 2^n$, which is even, the median $\frac{1}{2}E_{\text{Max}}$ is not defined by a specific energy level $E(j_M)$, but from the average of two energy levels, $E(j_M)$ and $E(j_M \oplus 1)$. Therefore, almost everywhere for $\Gamma \in (\mathbb{R}^+)^n$, there is no j_M such that $E(j_M) = \frac{1}{2}E_{\text{Max}}$. In turn this implies that

$|\mathbb{S}_0| = 2^{n-1}$ by definition of the median, for almost every energy structure. It is important to note that counterexamples can easily be found to this statement: it is sufficient to construct an energy structure such that for some j_M we have $E(j_M) = E(j_M \oplus 1)$; however, it is easy to see that these examples are not robust to arbitrarily small noise on the energy gaps. Furthermore, it is also worth noting that $\mathbb{S}_0$ corresponds to the swaps that *need* to be performed in order to fully cool the system. However, swaps on the states corresponding to the bit strings such that $E(j_M) = E(j_M \oplus 1) = \frac{1}{2}E_{\text{Max}}$ neither cool nor heat the system. Therefore, when $\omega = 0$, one can always bring the cardinality of the set of states that are being swapped to 2^{n-1} .

Two main observations are at play for the characterization of which states are guaranteed to be in (or out of) $\mathbb{S}_0$: the first is Lemma II, and the second is the existence of *incompatible bit strings*. We say that two elements of $\{0, 1\}^{\times n}$ are incompatible when the belonging of one in $\mathbb{S}_0$ is equivalent to the exclusion of the other in $\mathbb{S}_0$. A natural example of such two elements can always be constructed with any $j_M \in \{0, 1\}^{\times n}$ by picking its conjugate $\tilde{j}_M := j_M \oplus 1$: by definition of $E(\cdot)$ we have $E(\tilde{j}_M) = E_{\text{Max}} - E(j_M)$, which allows us to conclude that $j_M \in \mathbb{S}_0 \Rightarrow \tilde{j}_M \notin \mathbb{S}_0$. If we further assume that $E(j_M) \neq \frac{1}{2}E_{\text{Max}}$, then we obtain the other direction: $j_M \notin \mathbb{S}_0 \Rightarrow \tilde{j}_M \in \mathbb{S}_0$. Therefore there are three scenarios: (i) $j_M \in \mathbb{S}_0$ and $\tilde{j}_M \notin \mathbb{S}_0$, (ii) $j_M \notin \mathbb{S}_0$ and $\tilde{j}_M \in \mathbb{S}_0$, and (iii) $j_M \notin \mathbb{S}_0$ and $\tilde{j}_M \notin \mathbb{S}_0$, with (iii) happening if and only if $E(j_M) = \frac{1}{2}E_{\text{Max}}$. Then it is worth noting that if for a bit string k_M there exists j_M such that $j_M \geq k_M$ and $\tilde{j}_M \geq k_M$ then $k_M \notin \mathbb{S}_0$ regardless of the energy structure. And, almost always, we will also have $\tilde{k}_M \in \mathbb{S}_0$. This is particularly interesting for cases in which $\tilde{k}_M \geq k_M$, where we can pick $j_M = k_M$ and reach the same conclusion. It turns out this is the defining and only characteristic of the bit strings that are either never or (almost) always in $\mathbb{S}_0$ and that it generalizes to the entire set $\mathbb{S}$.

With these observations, we outline with the following theorem what are the bit strings that correspond to the states that always need to be swapped and those that never need to be swapped.

Theorem I2 (Always- and never-swappable levels). We define the following sets:

- (a) the ℓ -minimal swappable set $\mathbb{S}_{\min}^\ell = \{j_M \in \{0, 1\}^{\times n} | j_M \geq j_M \oplus 1 \text{ and } \#(j_M) > n/2 + \ell\}$, with $\ell = -1, 0, 1, \dots, \lfloor (n-1)/2 \rfloor$
- (b) the never-swappable set $\tilde{\mathbb{S}}_n := \{j_M \in \{0, 1\}^{\times n} | j_M \oplus 1 \geq j_M\}$

Then the following statements hold true:

- (i) $\tilde{\mathbb{S}}_n$ is the largest set such that $\tilde{\mathbb{S}}_n \cap \mathbb{S} = \emptyset$ for all $0 < \gamma_1 \leq \gamma_2 \leq \dots \leq \gamma_n$ and $0 \leq (T_M/T_S)\omega$;

- (ii) $\mathbb{S}_{\min}^\ell$ is the largest set such that $\mathbb{S}_{\min}^\ell \subseteq \mathbb{S}$ for all $0 < \gamma_1 \leq \gamma_2 \leq \dots \leq \gamma_n$ and $0 \leq (T_M/T_S)\omega < \sum_{i=1}^L \gamma_i$ for $\ell \geq 0$;
- (iii) $\mathbb{S}_{\min}^{-1}$ is the largest set such that $\mathbb{S}_{\min}^{-1} \subseteq \mathbb{S}$ for all $0 < \gamma_1 \leq \dots \leq \gamma_n$ and $(T_M/T_S)\omega < \min(\gamma_1 + \gamma_2, \gamma_n - \gamma_1)$ for n even;

where $L = 2(\ell + 1)$ for n even and $L = 2\ell + 1$ for n odd.

Proof of (ii) and (iii). We can rephrase statement (ii) as $j_M \in \mathbb{S}$ for all energy structures with $(T_M/T_S)\omega < \sum_{i=1}^L \gamma_i$ if and only if $j_M \in \mathbb{S}_{\min}^\ell$ ($\ell \geq 0$). ($\Leftarrow$ of ii) Assuming that $j_M \in \mathbb{S}_{\min}^\ell$, we have $j_M \geq \tilde{j}_M$ and $J = \#(j_M) > n/2 + \ell$. By definition of the lexicographic partial order and of conjugate bit strings, we necessarily have

$$\begin{aligned} E(j_M) &= \sum_{i=1}^J \gamma_{j_M(i)} \geq E(\tilde{j}_M) + \sum_{i=1}^{2J-n} \gamma_{j_M(i)} \\ &\geq E(\tilde{j}_M) + \sum_{i=1}^{2J-n} \gamma_i. \end{aligned} \quad (11)$$

By using $E(\tilde{j}_M) = E_{\text{Max}} - E(j_M)$, we find $E(j_M) \geq \frac{1}{2}(E_{\text{Max}} + \sum_{i=1}^{2J-n} \gamma_i)$. For n even we have $J \geq n/2 + 1 + \ell$, and therefore $2J - n \geq 2(\ell + 1)$; while for n odd we have $J \geq (n + 1)/2 + \ell$, and therefore $2J - n \geq 2\ell + 1$. Therefore $j_M \in \mathbb{S}$ for $(T_M/T_S)\omega < \sum_{i=1}^L \gamma_i$.

Similarly to (ii), we can rephrase statement (iii) $j_M \in \mathbb{S}$ for all energy structures with $(T_M/T_S)\omega \leq \min(\gamma_1 + \gamma_2, \gamma_n - \gamma_1)$ if and only if $j_M \in \mathbb{S}_{\min}^{-1}$ when n is even. ($\Leftarrow$ of iii) It is sufficient to extend the argument for ($\Leftarrow$ of ii) with $\ell = 0$ to the elements of $\mathbb{S}_{\min}^{-1}$ that do not belong in $\mathbb{S}_{\min}^0$. Assuming that $j_M \in \mathbb{S}_{\min}^{-1}$ and $j_M \notin \mathbb{S}_{\min}^0$ we have $\#(j_M) = \#(\tilde{j}_M) = n/2$. Furthermore we also have that $j_M(i) > \tilde{j}_M(i)$ for all $i = 1, \dots, n/2 = \#(j_M) = \#(\tilde{j}_M)$, therefore $j_M(n/2) = n$ and $\tilde{j}_M(1) = 1$, which implies $E(j_M) \geq E(\tilde{j}_M) + \gamma_n - \gamma_1$. Similarly to the argumentation in ($\Leftarrow$ of ii), we get $E(j_M) \geq \frac{1}{2}(E_{\text{Max}} + \gamma_n - \gamma_1)$ and therefore $j_M \in \mathbb{S}$ for $(T_M/T_S)\omega < \gamma_n - \gamma_1$. This concludes the proof of ($\Leftarrow$ of iii).

($\Rightarrow$ of ii and iii) Conversely, if we now assume that $j_M \notin \mathbb{S}_{\min}^\ell$, then either $j_M \in \mathbb{S}_{\min}^{\ell-1}$ (for $\ell \geq 0$) or $j_M \not\geq \tilde{j}_M$. In both cases it is sufficient to show that there exists an energy structure (that satisfies the requirements of the corresponding part of the theorem) such that $j_M \notin \mathbb{S}$. We start with the part that is only required for the proof of (ii), assuming that $j_M \in \mathbb{S}_{\min}^{\ell-1}$ and $j_M \notin \mathbb{S}_{\min}^\ell$ (with $\ell \geq 0$), then $\#(j_M) = \lfloor n/2 \rfloor + \ell$. For the energy structure $\gamma_1 = \gamma_2 = \dots = \gamma_n$ this means that $E(j_M) = \frac{1}{2}(E_{\text{Max}} + 2\ell\gamma_1)$ for n even and $E(j_M) = \frac{1}{2}(E_{\text{Max}} + (2\ell - 1)\gamma_1)$ for n odd. By picking $(T_M/T_S)\omega = 2\ell\gamma_1$ we have $(T_M/T_S)\omega < (2\ell + 1)\gamma_1 = \sum_{i=1}^{2\ell+1} \gamma_i$ and $E(j_M) \leq \frac{1}{2}(E_{\text{Max}} + (T_M/T_S)\omega)$, therefore $j_M \notin \mathbb{S}$.

Assuming that $j_M \not\geq \tilde{j}_M$ we have that either (a) $J < n/2$ or (b) $\exists \ell \in \{1, \dots, n - J\}$ such that $\tilde{j}_M(\ell) > j_M(\ell + 2J - n)$. For (a) we pick $\gamma_1 = \gamma_2 = \dots = \gamma_n$. Then it is immediate that $E(\tilde{j}_M) > E(j_M)$, and therefore $\frac{1}{2}E_{\text{Max}} > E(j_M)$ and $j_M \notin \mathbb{S}$. For (b) we pick $\gamma_1 = \dots = \gamma_{j_M(\ell)-1}$ and $\gamma_{j_M(\ell)} = \dots = \gamma_n > \gamma_1$. This leads to $E(\tilde{j}_M) = (\ell - 1)\gamma_1 + (n - J - \ell + 1)\gamma_n$. While for j_M we have

$$\begin{aligned} E(j_M) &= \sum_{i=1}^J \gamma_{j_M(i)} = (\ell + 2J - n)\gamma_1 + \sum_{i=\ell+2J-n+1}^J \gamma_{j_M(i)} \\ &\leq (\ell + 2J - n)\gamma_1 + (n - J - \ell)\gamma_n. \end{aligned} \quad (12)$$

Therefore we can pick γ_n large enough such that $E(\tilde{j}_M) > E(j_M)$. Therefore $j_M \notin \mathbb{S}$ for all ω . This concludes the proof of (ii) and (iii).

Proof of (i). The statement can be rephrased as $j_M \notin \mathbb{S}$ for all energy structures if and only if $j_M \in \tilde{\mathbb{S}}_n$. By definition of the lexicographic order, $j_M \in \tilde{\mathbb{S}}_n$ satisfies $\#(j_M) \leq \#(\tilde{j}_M)$. Therefore it is immediate that $\tilde{\mathbb{S}}_n = \{j_M \oplus 1 | j_M \in \mathbb{S}_{\min}^{-1}\}$. We obtain ($\Leftarrow$) immediately from Lemma 11 and incompatibility of j_M and $j_M \oplus 1$.

($\Rightarrow$) Conversely, we now assume $j_M \not\geq \tilde{j}_M$. By definition of the lexicographic order, either (a) $J = \#(j_M) > n/2$ or (b) there exists an $\ell \in \{1, \dots, J\}$ such that $\tilde{j}_M(\ell + n - 2J) < j_M(\ell)$. In both cases we construct examples of energy structures where $j_M \in \mathbb{S}$: for (a) $J > n/2$ we can pick $\gamma_1 = \dots = \gamma_n$, from which it is immediate that $E(j_M) > E(\tilde{j}_M)$ and therefore $E(j_M) > \frac{1}{2}E_{\text{Max}}$. For (b) we can pick $\gamma_1 = \dots = \gamma_{j_M(\ell)-1}$ and $\gamma_{j_M(\ell)} = \dots = \gamma_n > \gamma_1$. Similarly to the proof of (ii) and (iii) we obtain $E(j_M) = \sum_{i=1}^J \gamma_{j_M(i)} = (J - \ell + 1)\gamma_n + (\ell - 1)\gamma_1$ and $E(\tilde{j}_M) = \sum_{i=1}^{n-J} \gamma_{\tilde{j}_M(i)} \leq (J - \ell)\gamma_n + (\ell + n - 2J)\gamma_1$. Therefore we can choose γ_n large enough such that $E(j_M) > E(\tilde{j}_M)$. Therefore we conclude that if $j_M \notin \tilde{\mathbb{S}}_n$, then there are energy structures such that $E(j_M) > \frac{1}{2}E_{\text{Max}}$ and therefore for which $j_M \in \mathbb{S}$. This concludes the proof of (i) and the theorem. ■

Let us briefly unpack the statements made by this theorem by considering first the never-swappable set, as it is more intuitive. Starting from observations made above that use the lexicographic partial order, we are able to find a set of states that, regardless of the energy structure of the machine, will not need to be swapped. Furthermore, the set $\tilde{\mathbb{S}}_n$ is maximal, in the sense that for any bit string $j_M \notin \tilde{\mathbb{S}}_n$ then we can find an energy structure where one needs to act on the population of the corresponding state $|0_{Sj_M}\rangle$ to cool the system.

Ideally, we would want to make a similarly simple analysis on the set of states that always need to be acted upon in order to cool the system. However, there is a complication to be taken into account: no state is *always* worth swapping as one can simply pick a value of ω that is large enough so that it is not worth swapping. Therefore one has

to work around this problem by keeping track of how cold the initial system is already. We can start by considering the case in which $\omega = 0$, for which the set of states that always need to be swapped, regardless of the energy structure of the machine, is given by the $\mathbb{S}_{\min}^{-1}$. Incidentally, and intuitively, the elements of this set are the conjugate elements of the never-swappable set. Similarly to the never swappable set, $\mathbb{S}_{\min}^{-1}$ is maximal in the sense that for any bit string $j_M \notin \mathbb{S}_{\min}^{-1}$ we can find a corresponding energy structure (that satisfies $\omega = 0$) such that the state $|0_S j_M\rangle$ does not need to be swapped in order to cool the system. As one increases the value of ω there will be some states that stop being always in the swappable set. Therefore one needs to “adjust” the mathematical statement accordingly. In the theorem we do so for every value of ω until the machine becomes completely useless—which happens when the system starts so cold that the machine cannot be used to cool it further.

a. Constructing the ℓ -Minimal Swappable Set

We now move on to consider which are the elements of $\mathbb{S}_{\min}^\ell$. For the practicality of this section, we will decompose this set into its Hamming-weight sectors: $A_k := \{j_M \in \{0, 1\}^{\times n} \mid j_M \succeq j_M \oplus 1 \text{ and } \#(j_M) = k\}$, so that

$$\mathbb{S}_{\min}^\ell = \bigcup_{k > \frac{n}{2} + \ell} A_k.$$

The defining property of the lexicographic partial order is that a bit string j_M is larger than a bit string k_M when for each 1 in k_M we can identify a unique 1 in j_M that is “further to the right.” Therefore the condition $j_M \succeq j_M \oplus 1$ means that for every 0 in j_M there is a 1 uniquely associated with this 0 further to the right in the bit string. Very naturally, we can immediately see that A_ℓ is empty for $\ell < n/2$ since there are more zeros than ones in those bit strings. However the problem becomes less trivial for $\ell \geq n/2$ (which we will assume for the rest of this section), as there are bit strings such as 1110 which cannot be part of the set A_3 for $n = 4$.

In practice one can simply brute-force check every element, and therefore one would obtain the list of elements of A_k for a given n . However one can do better than this—and even obtain analytical expressions that can be used for bounds—by using the property that $j_M \succeq k_M \Leftrightarrow k_M \oplus 1 \succeq j_M \oplus 1$ (this follows immediately from the fact that $\oplus 1$ swaps the zeros and ones). In particular, let us consider $j_M, k_M \in \mathbb{B}_\ell$ such that $j_M \succeq k_M$. Then it is immediate from the above property (and transitivity) that $k_M \in A_\ell \Rightarrow j_M \in A_\ell$.

Let us consider the decreasing operation of the positional partial order: within the bit string we change a 01 for a 10. Since for any element of A_k one can find a unique 1 on

the right of each 0, then we can apply the decreasing operation to the bit string until we obtain $a_M^k = 1^{2k-n}(01)^{n-k}$. We can note that $a_M^k \in A_k$ and that if we apply any decreasing operation to a_M^k then the resulting bit string is not an element of A_k . Therefore a_M^k is the smallest element of A_k by the positional partial order. Since A_k has a smallest element, by using the property mentioned above, we can construct all of A_k by applying iteratively the increasing operation $10 \rightarrow 01$ to a_M^k . Furthermore, it is also worth noting that $0^{n-k}1^k$ is an element of A_k by construction and is the largest element of $\mathbb{B}_k$, and therefore is also the largest element of A_ℓ .

Finally, we can also note that if we change any 0 of a_M^k for a 1, then we obtain an element that is larger or equal to a_M^{k+1} . The same remains true for any other element of A_k , since they are larger than a_M^k . But we also have $a_M^{k+1} \succeq a_M^k$, which allows us to extend our generative method from one A_k to all of them. Therefore we can simplify the definition of the ℓ -minimal swappable set to be all the elements that are larger than $a_M^{\lfloor \frac{n}{2} + \ell + 1 \rfloor}$ by the lexicographic partial order

$$\mathbb{S}_{\min}^\ell = \{j_M \in \{0, 1\}^{\times n} \mid j_M \succeq a_M^{\lfloor \frac{n}{2} + \ell + 1 \rfloor}\}. \quad (\text{I3})$$

And one can generate this set by considering all the elements obtained by starting from $a_M^{\lfloor n/2 + \ell + 1 \rfloor}$ and either changing a 0 into a 1 or changing the sequence 10 to 01 in the bit string.

We can therefore rephrase statements (ii) and (iii) of the theorem as

For $\sum_{i=1}^{2(\ell-1)-n} \gamma_i \leq (T_M/T_S)\omega < \sum_{i=1}^{2\ell-n} \gamma_i$ (with $\ell = \lceil (n+1)/2 \rceil, \dots, n$) the largest subset of $\mathbb{S}$ that is agnostic of Γ is given by $\{j_M \in \{0, 1\}^{\times n} \mid j_M \succeq 1^{2\ell-n}(01)^{n-\ell}\}$. Furthermore, when n is even and $(T_M/T_S)\omega < \min(\gamma_1 + \gamma_2, \gamma_n - \gamma_1)$ we can expand this set to $\{j_M \in \{0, 1\}^{\times n} \mid j_M \succeq (01)^{n/2}\}$.

-
- [1] D. Reeb, and M. M. Wolf, An improved Landauer principle with finite-size corrections, *New J. Phys.* **16**, 103011 (2014).
 - [2] R. Silva, G. Manzano, P. Skrzypczyk, and N. Brunner, Performance of autonomous quantum thermal machines: Hilbert space dimension as a thermodynamical resource, *Phys. Rev. E* **94**, 032120 (2016).
 - [3] H. Wilming, and R. Gallego, Third law of thermodynamics as a single inequality, *Phys. Rev. X* **7**, 041033 (2017).
 - [4] L. Masanes, and J. Oppenheim, A general derivation and quantification of the third law of thermodynamics, *Nat. Commun.* **8**, 14538 (2017).
 - [5] F. Clivaz, R. Silva, G. Haack, J. B. Brask, N. Brunner, and M. Huber, Unifying paradigms of quantum refrigeration: Fundamental limits of cooling and associated work costs, *Phys. Rev. E* **100**, 042130 (2019).
 - [6] 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).
- [7] P. Taranto, F. Bakhshinezhad, A. Bluhm, R. Silva, N. Friis, M. 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).
- [8] L. Bassman Oftelie, A. De Pasquale, and M. Campisi, Dynamic cooling on contemporary quantum computers, *PRX Quantum* **5**, 030309 (2024).
- [9] J. Goold, M. Huber, A. Riera, L. del Rio, and P. Skrzypczyk, The role of quantum information in thermodynamics—a topical review, *J. Phys. A: Math. Theor.* **49**, 143001 (2016).
- [10] F. Binder, L. A. Correa, C. Gogolin, J. Anders, and G. Adesso, eds. *Thermodynamics in the Quantum Regime: Fundamental Aspects and New Directions*, Fundamental Theories of Physics Vol. 195 (Springer, Cham, Switzerland, 2018).
- [11] P. Strasberg, *Quantum Stochastic Thermodynamics: Foundations and Selected Applications*, Oxford Graduate Texts (Oxford University Press, Oxford, United Kingdom, 2022).
- [12] D. K. Park, N. A. Rodríguez-Briones, G. Feng, R. Rahimi, J. Baugh, and R. Laflamme, in *Electron Spin Resonance (ESR) Based Quantum Computing*, edited by T. Takui, L. Berliner, and G. Hanson (Springer New York, New York, 2016), p. 227.
- [13] L. J. Schulman, and U. V. Vazirani, in *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, STOC '99 (Association for Computing Machinery, New York, 1999), p. 322.
- [14] L. J. Schulman, T. Mor, and Y. Weinstein, Physical limits of heat-bath algorithmic cooling, *Phys. Rev. Lett.* **94**, 120501 (2005).
- [15] S. Raeisi, and M. Mosca, Asymptotic bound for heat-bath algorithmic cooling, *Phys. Rev. Lett.* **114**, 100404 (2015).
- [16] N. A. Rodríguez-Briones, E. Martín-Martínez, A. Kempf, and R. Laflamme, Correlation-enhanced algorithmic cooling, *Phys. Rev. Lett.* **119**, 050502 (2017).
- [17] Á. M. Alhambra, M. Lostaglio, and C. Perry, Heat-bath algorithmic cooling with optimal thermalization strategies, *Quantum* **3**, 188 (2019).
- [18] 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).
- [19] N. Linden, S. Popescu, and P. Skrzypczyk, How small can thermal machines be? The smallest possible refrigerator, *Phys. Rev. Lett.* **105**, 130401 (2010).
- [20] A. Levy, and R. Kosloff, Quantum absorption refrigerator, *Phys. Rev. Lett.* **108**, 070604 (2012).
- [21] M. T. Mitchison, Quantum thermal absorption machines: Refrigerators, engines and clocks, *Contemp. Phys.* **60**, 164 (2019).
- [22] M. A. Aamir, P. J. Suria, J. A. M. Guzmán, Claudia Castillo-Moreno, Jeffrey M. Epstein, Nicole Yunger Halpern, and Simone Gasparinetti, Thermally driven quantum refrigerator autonomously resets a superconducting qubit, *Nat. Phys.* **21**, 318 (2025).
- [23] C. Godsil, and G. F. Royle, *Algebraic Graph Theory*, Graduate Texts in Mathematics No. 207 (Springer, 2001).
- [24] C. H. Papadimitriou, and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Prentice-Hall, Inc., USA, 1982).
- [25] A. E. Allahverdyan, K. V. Hovhannisyan, D. Janzing, and G. Mahler, Thermodynamic limits of dynamic cooling, *Phys. Rev. E* **84**, 041109 (2011).
- [26] 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).
- [27] P. Taranto, P. Lipka-Bartosik, N. A. Rodríguez-Briones, M. Perarnau-Llobet, N. Friis, M. Huber, and P. Bakhshinezhad, Efficiently cooling quantum systems with finite resources: Insights from thermodynamic geometry, *Phys. Rev. Lett.* **134**, 070401 (2025).
- [28] A. E. Allahverdyan, R. Balian, and T. M. Nieuwenhuizen, Maximal work extraction from finite quantum systems, *Europhys. Lett.* **67**, 565 (2004).
- [29] M. Horodecki, and J. Oppenheim, Fundamental limitations for quantum and nanoscale thermodynamics, *Nat. Commun.* **4**, 2059 (2013).
- [30] R. Silva, P. Bakhshinezhad, and F. Clivaz, Optimal unitary trajectories under commuting target and cost observables; applications to cooling, *ArXiv:2412.07291*.
- [31] This is because we have two lists of size $|\mathcal{S}|$. Starting with one element in $|A\rangle$ we have $|\mathcal{S}|$ choices to pair it with in $|B\rangle$, whereas the next element has $|\mathcal{S}| - 1$ possibilities and so on, giving $|\mathcal{S}| \times |\mathcal{S}| - 1 \times |\mathcal{S}| - 2 \times \dots \times 1 = |\mathcal{S}|!$ possible pairings in general. Other unitaries may be considered that cool optimally but also act on levels outside $|A\rangle$ without any effect on the cooling.
- [32] K. Temme, and M. J. Kastoryano, How fast do stabilizer Hamiltonians thermalize? *ArXiv:1505.07811*.
- [33] F. Reiter, A. S. Sørensen, P. Zoller, and C. A. Muschik, Dissipative quantum error correction and application to quantum sensing with trapped ions, *Nat. Commun.* **8**, 1822 (2017).
- [34] Z. Weinstein, G. Ortiz, and Z. Nussinov, Universality classes of stabilizer code Hamiltonians, *Phys. Rev. Lett.* **123**, 230503 (2019).
- [35] B. de Neeve, T. L. Nguyen, T. Behrle, and J. P. Home, Error correction of a logical grid state qubit by dissipative pumping, *Nat. Phys.* **18**, 296 (2022).
- [36] T. Hillmann, and F. Quijandria, Quantum error correction with dissipatively stabilized squeezed-cat qubits, *Phys. Rev. A* **107**, 032423 (2023).
- [37] O. Shtanko, Y.-J. Liu, S. Lieu, A. V. Gorshkov, and V. V. Albert, Bounds on autonomous quantum error correction, *Quantum* **9**, 1804 (2025).
- [38] N. Brunner, N. Linden, S. Popescu, and P. Skrzypczyk, Virtual qubits, virtual temperatures, and the foundations of thermodynamics, *Phys. Rev. E* **85**, 051117 (2012).
- [39] P. Erker, M. T. Mitchison, R. Silva, M. P. Woods, N. Brunner, and M. Huber, Autonomous quantum clocks: Does thermodynamics limit our ability to measure time? *Phys. Rev. X* **7**, 031022 (2017).

- [40] P. P. Hofer, J. B. Brask, M. Perarnau-Llobet, and N. Brunner, Quantum thermal machine as a thermometer, *Phys. Rev. Lett.* **119**, 090603 (2017).
- [41] P. Lipka-Bartosik, M. Perarnau-Llobet, and N. Brunner, Thermodynamic computing via autonomous quantum thermal machines, *Sci. Adv.* **10**, eadm8792 (2024).
- [42] J. Xuereb, A temperature change can solve the Deutsch-Josza problem: An exploration of thermodynamic query complexity, [ArXiv:2505.15887](https://arxiv.org/abs/2505.15887).
- [43] M. Perarnau-Llobet, K. V. Hovhannisyan, M. Huber, P. Skrzypczyk, N. Brunner, and A. Acín, Extractable work from correlations, *Phys. Rev. X* **5**, 041011 (2015).
- [44] T. Yada, N. Yoshioka, and T. Sagawa, Measuring multipartite quantum correlations by thermodynamic work extraction, [ArXiv:2407.04058](https://arxiv.org/abs/2407.04058).
- [45] 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).
- [46] S. Majidy, W. F. Braasch, A. Lasek, T. Upadhyaya, A. Kalev, and N. Yunger Halpern, Noncommuting conserved charges in quantum thermodynamics and beyond, *Nat. Rev. Phys.* **5**, 689 (2023).
- [47] S. Majidy, Noncommuting charges can remove non-stationary quantum many-body dynamics, *Nat. Commun.* **15**, 8246 (2024).
- [48] P. P. Hofer, M. Perarnau-Llobet, J. B. Brask, R. Silva, M. Huber, and N. Brunner, Autonomous quantum refrigerator in a circuit QED architecture based on a Josephson junction, *Phys. Rev. B* **94**, 235420 (2016).
- [49] L. G. DiFranco, QuL: Programming library for computational cooling of qubits, *ACM Trans. Quantum Comput.* **6**, 1 (2025).
- [50] N. L. S. Carnot, *Réflexions sur la Puissance Motrice du feu et sur les Machines Propres à développer Cette Puissance* (Bachelier, Paris, 1824).
- [51] D. J. Wineland, and W. M. Itano, Laser cooling of atoms, *Phys. Rev. A* **20**, 1521 (1979).
- [52] J. Xuereb, <https://github.com/jqed-xuereb/Cooling-a-Qubit-w-N-Others> (2025).
- [53] F. Harary, J. P. Hayes, and H.-J. Wu, A survey of the theory of hypercube graphs, *Comput. Math. Appl.* **15**, 277 (1988).